

Teradata® Viewpoint

User Guide - 23.04

Release 23.04




April 2023

Copyright and Trademarks

Copyright © 2008 - 2023 by Teradata. All Rights Reserved.

All copyrights and trademarks used in Teradata documentation are the property of their respective owners. For more information, see [Trademark Information](#).

Product Safety

Safety type	Description
 NOTICE	Indicates a situation which, if not avoided, could result in damage to property, such as to equipment or data, but not related to personal injury.
 CAUTION	Indicates a hazardous situation which, if not avoided, could result in minor or moderate personal injury.
 WARNING	Indicates a hazardous situation which, if not avoided, could result in death or serious personal injury.

Third-Party Materials

Non-Teradata (i.e., third-party) sites, documents or communications ("Third-party Materials") may be accessed or accessible (e.g., linked or posted) in or in connection with a Teradata site, document or communication. Such Third-party Materials are provided for your convenience only and do not imply any endorsement of any third party by Teradata or any endorsement of Teradata by such third party. Teradata is not responsible for the accuracy of any content contained within such Third-party Materials, which are provided on an "AS IS" basis by Teradata. Such third party is solely and directly responsible for its sites, documents and communications and any harm they may cause you or others.

Warranty Disclaimer

Except as may be provided in a separate written agreement with Teradata or required by applicable laws, all designs, specifications, statements, information, recommendations and content (collectively, "content") available from the Teradata Documentation website or contained in Teradata information products is presented "as is" and without any express or implied warranties, including, but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or noninfringement, which are hereby disclaimed. In no event shall Teradata corporation, its suppliers or partners be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of content, even if advised of the possibility of such damage.

The Content available from the Teradata Documentation website or contained in Teradata information products may contain references or cross-references to features, functions, products, or services that are not announced or available in your country. Such references do not imply that Teradata Corporation intends to announce such features, functions, products, or services in your country. Please consult your local Teradata Corporation representative for those features, functions, products, or services available in your country.

The Content available from the Teradata Documentation website or contained in Teradata information products may be changed or updated by Teradata at any time without notice. Teradata may also make changes in the products or services described in the Content at any time without notice.

The Content is subject to change without notice. Users are solely responsible for their application of the Content. The Content does not constitute the technical or other professional advice of Teradata, its suppliers or partners. Users should consult their own technical advisors before implementing any Content. Results may vary depending on factors not tested by Teradata.

Machine-Assisted Translation

Certain materials on this website have been translated using machine-assisted translation software/tools. Machine-assisted translations of any materials into languages other than English are intended solely as a convenience to the non-English-reading users and are not legally binding. Anybody relying on such information does so at his or her own risk. No automated translation is perfect nor is it intended to replace human translators. Teradata does not make any promises, assurances, or guarantees as to the accuracy of the machine-assisted translations provided. Teradata accepts no responsibility and shall not be liable for any damage or issues that may result from using such translations. Users are reminded to use the English contents.

Feedback

To maintain the quality of our products and services, e-mail your comments on the accuracy, clarity, organization, and value of this document to: docs@teradata.com.

Any comments or materials (collectively referred to as "Feedback") sent to Teradata Corporation will be deemed nonconfidential. Without any payment or other obligation of any kind and without any restriction of any kind, Teradata and its affiliates are hereby free to (1) reproduce, distribute, provide access to, publish, transmit, publicly display, publicly perform, and create derivative works of, the Feedback, (2) use any ideas, concepts, know-how, and techniques contained in such Feedback for any purpose whatsoever, including developing, manufacturing, and marketing products and services incorporating the Feedback, and (3) authorize others to do any or all of the above.

Confidential Information

Confidential Information means any and all confidential knowledge, data or information of Teradata, including, but not limited to, copyrights, patent rights, trade secret rights, trademark rights and all other intellectual property rights of any sort.

The Content available from the Teradata Documentation website or contained in Teradata information products may include Confidential Information and as such, the use of such Content is subject to the non-use and confidentiality obligations and protections of a non-disclosure agreement or other such agreements to protect Confidential Information that you have executed with Teradata.

Contents

Chapter 1: Teradata Viewpoint Basics	7
Welcome to Teradata Viewpoint User Guide	7
Teradata Viewpoint Portal and Portlets	8
Profile	24
Chapter 2: Viewpoint RESTful APIs	30
Viewpoint RESTful APIs Overview	30
Chapter 3: Administration Portlets	31
System Administration	31
Alert Setup	32
Backup	77
Certificates	79
General	86
LDAP Servers	87
Monitored Systems	92
Portlet Library	152
Query Group Setup	153
Roles Manager	157
Shared Pages	177
User Manager	182
Chapter 4: Alert Viewer	189
Alert Viewer Overview	189
Chapter 5: Application Queries	194
Application Queries Overview	194
Chapter 6: Canary Response Times	206
Canary Response Times Overview	206
Chapter 7: Completed Queries	210
Completed Queries Overview	210
Chapter 8: Elastic Performance	230
Elastic Performance Overview	230
Chapter 9: Hadoop Services	235
Hadoop Services Overview	235

Chapter 10: Lock Viewer	244
Lock Viewer Overview	244
Chapter 11: MAPS Manager	256
MAPS Manager Overview	256
Chapter 12: Metric Heatmap	287
Metric Heatmap Overview	287
Chapter 13: Metrics Analysis	305
Metrics Analysis Overview	305
Chapter 14: My Queries	317
My Queries Overview	317
Chapter 15: Node Monitor	344
Node Monitor Overview	344
Chapter 16: Node Resources	361
Node Resources Overview	361
Chapter 17: Performance Data Collection	368
Performance Data Collection Overview	368
Chapter 18: Productivity	380
Productivity Overview	380
Chapter 19: Query Groups	384
Query Groups Overview	384
Chapter 20: Query Log	413
Query Log Overview	413
Chapter 21: Query Monitor	429
Query Monitor Overview	429
Chapter 22: Query Spotlight	474
Query Spotlight Overview	474
Chapter 23: Remote Console	487
Remote Console Overview	487
Chapter 24: Space Usage	493
Space Usage Overview	493

Chapter 25: SQL Scratchpad	505
SQL Scratchpad Overview	505
Chapter 26: Stats Manager	515
Stats Manager Overview	515
Chapter 27: System Health	558
System Health Overview	558
Chapter 28: Today's Statistics	567
Today's Statistics Overview	567
Chapter 29: TVS Monitor	572
TVS Monitor Overview	572
Chapter 30: Viewpoint Calendar	578
Viewpoint Calendar Overview	578
Chapter 31: Viewpoint Dashboard	584
Viewpoint Dashboard Overview	584
Chapter 32: Viewpoint Mobile	641
Viewpoint Mobile Overview	641
Chapter 33: Viewpoint Monitoring	677
Viewpoint Monitoring Overview	677
Chapter 34: Workload Management	684
Workload Management Scenarios	684
Workload Designer Overview	693
Workload Health Overview	796
Workload Monitor Overview	805
Appendix A: Metric Sources	843
Appendix B: Additional Information	848

Teradata Viewpoint Basics

Welcome to Teradata Viewpoint User Guide

Using *Teradata® Viewpoint User Guide*

The Teradata Viewpoint portal is a framework where web-based applications, known as portlets, are available. This guide covers the basics of the Teradata Viewpoint portal/mobile and explains the portlets.

Application developers and users, database administrators, and system administrators can customize the portlets to manage and monitor Teradata systems using supported web browsers.

Why Would I Use this Content?

This content explains how to work with the Teradata Viewpoint portal and portlets. You can get acquainted with:

- Logging on and logging out from the Viewpoint portal
- Accessing Viewpoint help
- Using notifications
- Setting up the portals for portlets
- Managing the portlets
- Using Viewpoint Mobile

How Do I Use this Content?

If you are new to Teradata Viewpoint, start with [Teradata Viewpoint Portal and Portlets](#) and get familiar with the product. Then proceed to the different portlets according to your use.

If you are a Viewpoint administrator, refer to [System Administration](#) to know more about administration portlets in the Viewpoint portal. These portlets allow you to provide access to Teradata Viewpoint resources and information.

If you are a database administrator, refer to [Metric Sources Overview](#) to understand how the Viewpoint metrics are derived.

How Do I Get Started?

This guide explains about Portlets in Viewpoint. To become familiar with the Viewpoint portal, read:

- [Getting Started](#): Learn how to log on and log off from the Viewpoint portal, access the help, use notifications, and use common keyboard features to browse the portal and portlets.
- [Setting Up the Portal for Portlets](#): Learn how to set up pages and portlets in the Viewpoint portal.
- [Working with Portlets in the Portal](#): Learn how to manage data in the portlet using various features and controls provided in the portal.

Teradata Viewpoint Portal and Portlets

The Teradata® Viewpoint portal is a framework where web-based applications, known as portlets, are displayed. IT professionals and business users can customize their portlets to manage and monitor their Teradata systems using a web browser.

Portlets enable users across an enterprise to customize tasks and display options to their specific business needs. You can view current data, run queries, and make timely business decisions, reducing the database administrator workload by allowing you to manage your work independently. Portlets are added to a portal page using the **Add Content** screen. The Teradata Viewpoint Administrator configures access to portlets based on your role.

Supported Web Browsers for Viewpoint Portal

Browser	Minimum Version
Microsoft Edge	97.0.1072.69
Mozilla Firefox	96.0.2
Google Chrome	97.0.4692.99
Safari	12.1.1

Supported Web Browsers for Viewpoint Mobile

Browser	Minimum Version
Google Chrome for Android	70.0.3538.102
Safari for iOS	9.3.5

Getting Started

Learn how to log on and log off from Teradata Viewpoint, access the help, use notifications, and use common keyboard features to browse the portal and portlets.

Logging on to the Viewpoint Portal

Log on to the Viewpoint portal.




1. In a browser, enter the address for your Viewpoint portal.
2. Log on to the Viewpoint portal with the username and password to perform system level tasks.

Note:

Your account will be locked after five unsuccessful attempts. After lockout, wait for five minutes before retrying to login. You can configure the failure count and duration of lockout using `failureCount` and `lockOutTime` parameters in `/opt/teradata/viewpoint/portal/conf/server.xml`. Any changes to these parameters requires restart of Viewpoint service.

Accessing Viewpoint Help


The **Teradata Viewpoint Help** provides information about how to use Teradata Viewpoint to manage and monitor your Teradata systems. You can search the help to get a list of relevant topics that match your criteria.

1. Do any of the following:
 - Click , located in the upper right of the portal frame, to open a browser window containing **Teradata Viewpoint Help** for all portlets.
 - Click , located to the right of a portlet name on the **Add Content** screen, to open a browser window containing **Teradata Viewpoint Help** for a specific portlet.
 - Click , located in the upper right of the portlet frame, to open a browser window containing **Teradata Viewpoint Help** for that specific portlet.

Using Notifications

Teradata Viewpoint uses notifications to communicate the status of Viewpoint operations. Notifications are displayed in the Teradata Viewpoint portal for 72 hours from the last status change.

Notifications appear if they are enabled for your role.

1. Click  to read the notifications.
2. [Optional] Click **Dismiss All** to hide the notifications.
The icon appears when there is at least one notification and indicates the number of active notifications.

Keyboard Accessibility

You can use common keyboard features to browse the portal and portlet interfaces, enabling Teradata Viewpoint to be accessible to all users. For example, press the **Tab** key to move focus from one item to the next and press **Esc** to leave a dialog box.

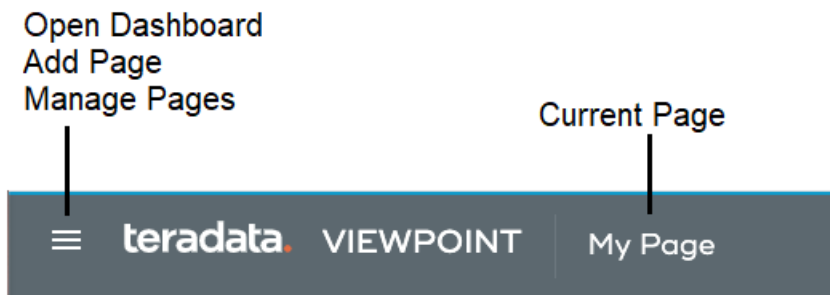
Logging Out

1. From the Teradata Viewpoint portal menu , click **Log Out** to exit the portal.

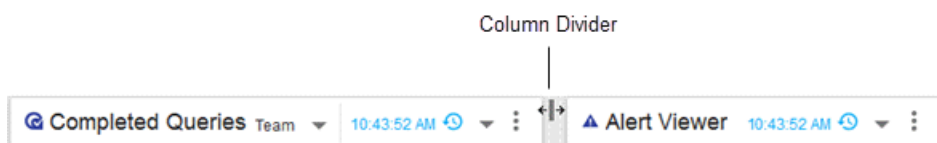
All page content is preserved for your next Teradata Viewpoint session. Any open **Teradata Viewpoint Help** window remains open.

Setting Up the Portal for Portlets

Teradata Viewpoint provides portlets to manage and monitor systems. A *portal page* is a virtual work space where you decide which portlets to display and how to arrange them on the page. Examples of ways to organize your work include defining a page for each system being monitored, or for each type of query or user. The currently selected page is called the *active page*.



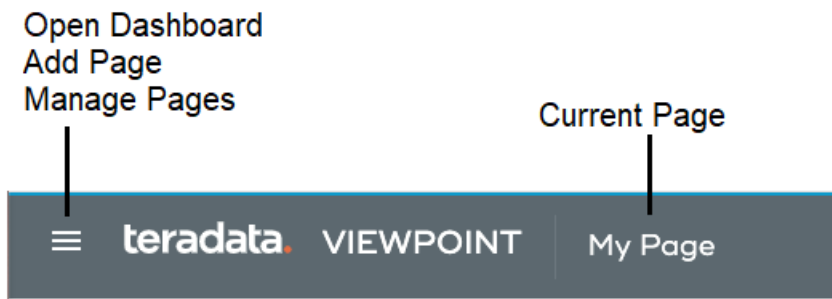
A portal page has two columns that can be resized. After adding portlets to the portal, you can drag and drop portlets into either column. You can resize the width of the portal columns on each page by dragging the column divider until the minimum size of a column is reached. Portlets that can be horizontally resized will expand to fill the width of the column. Portlets that cannot be horizontally resized are not affected by resizing the portal columns. Resizing the browser window adjusts the columns proportionally unless a column is at its minimum size. Not all Viewpoint portlets and products support resizing or two-column layout.



A shared page contains portlets that are applicable to your role. The Teradata Viewpoint Administrator controls the type of shared page you have access to and can discuss limitations you may have with read-only and mandatory pages. Your role may be prevented from adding portlets, deleting portlets, moving portlets on a page, as well as other functions you may have been accustomed to doing with regular portal pages.

Adding and Naming a Page

Organize your portlets by adding pages to the Teradata Viewpoint portal. When a page is added, it becomes the active page. When you add many pages, a page list appears where you can access all pages, including those that may not be visible.



1. In the portal, select > .
2. Select **Blank**, enter a name for the page up to 30 characters, then select **Create**.
A new page appears.
3. [Optional] Select **Add Content** to add portlets.

Renaming a Page

You can change the name of the active page at any time during a Teradata Viewpoint session.

1. In the portal, select then select the page.
2. Select to the right of the name, then select **Rename**.
3. Enter the new name up to 30 characters, then select **Save**.

Arranging Pages

New portal pages are added to the top of the list. Arranging pages does not change page content.


1. In the portal, select .
2. In the list of pages, select and drag a page up or down in the list.


Deleting a Page

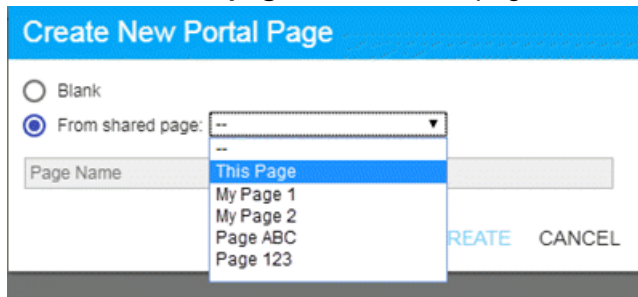
Delete pages you no longer need. You cannot delete a mandatory shared page.

1. In the portal, select then select the page to delete.
2. Select to the right of the name, then select **Delete**.
A confirmation message appears.
3. Click **OK**.
The portal page and its contents are deleted.

Adding a Shared Page

A shared page is a pre-defined page for your role. Mandatory shared pages automatically appear when you log on and cannot be added from the add page list. An  appears next to read only pages.

1. In the portal, select  > +.
2. Select **From shared page**, then select a page from the



list.


The new page becomes the active page.

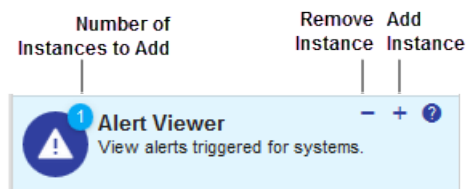
3. [Optional] Select **Add Content** to add portlets if the shared page is not read-only.

Adding Portlets

Add one or more portlets to the active page at any time during a session. Add multiple instances of a portlet as needed, but only one instance of the **Viewpoint Calendar** and **Workload Designer** portlet is permitted.

1. In the portal, click **Add Content**.
2. Click the portlet name to add an instance.


You can search for portlet names or keywords in the portlet descriptions using the search box .



3. Click **Add**.
The selected portlets are added to the current page.
4. [Optional] Drag portlets to reposition the portlets on the portal page.

Removing a Portlet

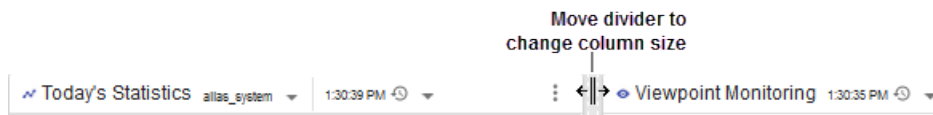
Remove a portlet from the active portal page at any time.

1. Select  in the portlet frame, then select **Remove**.
A confirmation message appears.
2. Select **OK**.
The portlet is removed from the active portal page.

Resizing Portal Columns

You can resize the width of most portlets. Portlets that are two columns wide, such as the **Workload Designer** portlet, cannot be resized and always appear at the top of the portal page. To resize the width of a portlet, follow these instructions.

1. Hover over the column divider.
The cursor changes to a bidirectional arrow.



2. Drag the column divider to the left or to the right to resize the column width.
The portlets resize after the column divider is released.

Working with Portlets in the Portal

The Viewpoint portlets allow you to adjust size of the portlets. You can rewind, pause, and filter data. You can also create a shared portlet.

Portlet Controls

The Teradata Viewpoint Administrator controls which portlets and what features you are authorized to use. If a portlet or feature is not available, check with the Teradata Viewpoint administrator for permission.

Each portlet instance has its own settings and controls, as well as the features common to all portlets. Not all portlets have all controls.

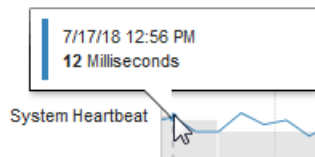
You can access the following portlet features and controls.

Back

Returns to the previous view.

Information Balloons

Displays detailed information when you hover over graphs and other objects in the portlet.



▼ Pause, Resume

Pauses or resumes refreshing data for a portlet.

If rewind is on, you can select to **Unlink from Rewind** or **Link to Rewind**.

If rewind is off, you can select **Refresh** to bring the most recent data collection from the Viewpoint server to the portlet.

Portlet Frame




Contains portlet controls.

Portlet Name

Displays the portlet name.

Rewind

Provides the status of the rewind function for a portlet:

-  Rewind is available.
-  Rewind is on.
-  Rewind is unlinked.

Restore Height

Resets the portlet to its default height. The button appears in the bottom right of the portlet frame after you change the default height of the portlet.

Selection Menus

Contains menus for selecting the data to be displayed in the portlet.

Table Actions

Provides controls for managing information displayed in the table.

Time Stamp

Displays the last time data was retrieved from the Teradata server. The time is based on the time zone set in the **Profile** portlet.

Toolbar

Contains buttons for quick access to information.

You can access the following portlet controls from .

Settings

Accesses portlet settings. Settings are used to specify what information is displayed, time intervals for reporting, and other features that help you customize the portlet functions.

Maximize

Resizes the portlet to span two columns in the portal page.

Minimize

Resizes the portlet to one column wide.

Share

Captures a customized version of a portlet for use by other users. The Teradata Viewpoint Administrator must make the customized portlet available for sharing.

Help

Opens a browser window containing Teradata Viewpoint Help for a specific portlet.

Remove

Removes the portlet from the active portal page.

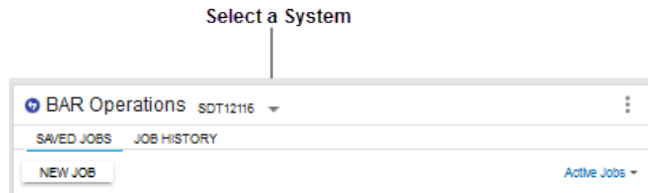
The screenshot shows the 'Completed Queries' portlet interface. Annotations point to the following controls:

- Select a System**: Points to the dropdown menu showing '_PD'.
- Rewind, Pause, or Refresh Portlet**: Points to the circular refresh icon.
- Access Settings**: Points to the gear icon in the top right corner.
- Change a Filter**: Points to the filter bar showing '809' and 'All'.
- Configure Columns**: Points to the 'DURATION' dropdown menu.
- Restore Height**: Points to the vertical double-headed arrow on the right side of the table.

START TIME	USER	LINK NAME	STATUS	DURATION
7/16/18 4:21:15 PM	TESTUSER	link4625to4935	Successful	0:00:22
7/16/18 4:25:03 PM	TESTUSER	l4625-4935-o1	Successful	0:00:17
7/16/18 4:25:23 PM	TESTUSER	link4625to4935	Successful	0:00:17
7/16/18 4:30:01 PM	TESTUSER	l4625-4935-o2	Successful	0:01:27
7/16/18 4:32:29 PM	TESTUSER	sd04625sd04935	Successful	0:01:24
7/16/18 4:50:02 PM	TESTUSER	link4625to4935	Successful	0:00:30
7/16/18 4:50:05 PM	TESTUSER	sd04625-sd04935	Successful	0:02:33
7/16/18 4:50:07 PM	TESTUSER	sd04625sd04935	Successful	0:06:26
7/16/18 4:50:07 PM	TESTUSER	link4625to4935	Successful	0:07:04
7/16/18 4:50:11 PM	TESTUSER	sd04625sd04935	Successful	0:00:20
7/16/18 5:00:03 PM	TESTUSER	l4625-4935-o0	Successful	0:00:33
7/16/18 5:03:02 PM	TESTUSER	sd04625-sd04935	Successful	0:00:34
7/16/18 5:10:03 PM	TESTUSER	l4625-4935-o2	Successful	0:00:30
7/16/18 5:12:07 PM	TESTUSER	sd04625sd04935	Successful	0:00:26
7/16/18 5:15:03 PM	TESTUSER	l4625-4935-o0	Successful	0:00:27
7/16/18 5:17:19 PM	TESTUSER	sd04625-sd04935	Successful	0:00:36
7/16/18 5:20:04 PM	TESTUSER	l4625-4935-o1	Successful	0:00:26
7/16/18 5:21:27 PM	TESTUSER	link4625to4935	Successful	0:00:26
7/16/18 5:25:03 PM	TESTUSER	l4625-4935-o1	Successful	0:00:18
7/16/18 5:25:24 PM	TESTUSER	link4625to4935	Successful	0:00:17
7/16/18 5:30:02 PM	TESTUSER	l4625-4935-o2	Successful	0:01:30
7/16/18 5:32:33 PM	TESTUSER	sd04625sd04935	Successful	0:01:26
7/16/18 5:50:02 PM	TESTUSER	link4625to4935	Successful	0:00:29
7/16/18 5:50:06 PM	TESTUSER	sd04625-sd04935	Successful	0:02:38

Page 1 of 2 (809 rows total)

Selecting a System to Monitor



1. Select ▼ in the portlet frame and select a system.

Summary Table Controls

Some portlets feature data in summary tables. If available, you can filter the displayed content with the toolbar, change columns, or export data.

Configuring Columns to Display

Use the **Configure Columns** dialog box to set thresholds and select, lock, and order columns. You can resize columns in the table. In some portlets, such as **Alert Viewer**, the column selection, order, and lock settings that you choose are applied to the minimized and maximized views separately.

1. Select ▼ in the table header then select **Configure Columns**.
2. In the **Configure Columns** dialog box, select the check boxes of columns to display. Hover over the column name to see a description of the data the column displays.
3. [Optional] To lock the column position, select 🔒 next to the column name. The columns at the top of the list can be locked in the table to remain on the left when scrolling horizontally.
4. [Optional] Select **Set Threshold**, type a threshold value, then select **OK**. Threshold values can be set only in certain portlets, and only for certain columns. In the **Query Monitor** portlet, threshold values can be set only for the Analytics Database. Qualifying data is highlighted in the table.
5. [Optional] Select ≡ and drag the row to reorder the column.
6. Click **OK**.
7. [Optional] In the table, drag the column heading border ↔ in either direction to resize the column.

Filters and Sorting

Filters allow you to display only rows that match your filter criteria. You can narrow the search further by filtering on multiple columns. Depending on the portlet, you can:

- Use greater than or less than symbols when filtering columns with numeric values. Exact matches may not produce the expected results due to rounding of numbers containing decimals.
- Use wildcard characters or symbols in the filter to include number or word variations in the filter match.
- Use the filter bar to display only specific states, severity levels, or time periods.

Sorting allows you to change the order of rows in a table based upon criteria in a column and applies across all pages of the table. Sort on a column by clicking the column header. A second click reverses the sort order. Sort on two columns consecutively by sorting the first column and then using **Ctrl+Click** once in the column header of the second column. Using **Ctrl+Click** one time sets the secondary sort order. Primary sort order is indicated by a single arrow, and secondary sort order is indicated by a double arrow. The secondary sort order is removed when the primary or secondary sort column has been modified. After modification, you see a single arrow to indicate the secondary sort order has been removed.

In portlets that can be maximized, the filtering, sorting, and page number settings that you choose for the default or minimized view are not used when you switch to the maximized view.

Filter Symbols

Use wildcard characters or symbols in a filter to expand or limit the search criteria. Some symbols are used for string filtering and others are used for numeric filtering. String filters are not case-sensitive and an asterisk wildcard is applied to the end of the filter string.




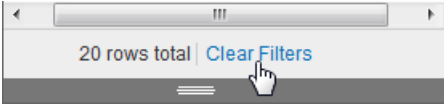
Wildcard or Symbol	Description	Example	Results	Supported Portlets
?	Use this wildcard character to match alpha or numeric characters in the position it occupies. Type this wildcard character at the beginning, middle, or end of your search. This wildcard character can be used more than once in the same search. This wildcard character can be used in conjunction with any other symbol.	cat? ?cat cat?l cat???	cats, catalog scat, Scatter catalog catalog	All summary tables
		p???er	packer, parser, proper	
*	Use this wildcard character to match zero, one, or multiple alpha or numeric characters in the position it occupies. Type this symbol at the beginning, middle, or end of your search. This wildcard character can be used more than once in the same search. This wildcard character can be used in conjunction with any other symbol.	*cat cat*l cat*	cat, cats, catalog, s catter, wild cat catalog cat, cats, catalog	All summary tables
		*ews*er	newscaster, newspaper	
		*%	what is 100%?	
=	Use this symbol at the beginning of your search to match alpha or numeric	=CAT	CAT	All summary tables

Wildcard or Symbol	Description	Example	Results	Supported Portlets
	characters literally. The search results are case-sensitive. This symbol can be used in conjunction with * and ?.	=cat? =Cat* =Cat_ =Cat\?	cats Cat, Cats, Catalog Cat_ Cat?	
\	Use this symbol in front of a wildcard character so the wildcard is interpreted as a regular character and not as a wildcard. This symbol can be used in conjunction with *, ?, and \.	*\? =cat\ cat\?	what is 100%? cat* cat?, Cat?	All summary tables
>	Use this symbol to match any number that is greater than the specified value. This symbol is only used to perform mathematical comparisons. For columns containing percentages, avoid using % in the filter. Type only > and the number.	>60	61, 62, 70, 500, and so on	All summary tables except in the Query Group Setup portlet
>=	Use these symbols in conjunction to match values greater than or equal to the specified number. For columns containing percentages, avoid using % in the filter. Type only >= and the number.	>=60	60, 61, 62, 70, 500, and so on	All summary tables except in the Query Group Setup portlet
<	Use this symbol to match any number that is less than the specified value. This symbol is only used to perform mathematical comparisons. For columns containing percentages, avoid using % in the filter. Type only < and the number.	<60	59, 58, 50, 8, and so on	All summary tables except in the Query Group Setup portlet
<=	Use these symbols in conjunction to match values less than or equal to the specified number. For columns containing percentages, avoid using % in the filter. Type only <= and the number.	<=60	60, 59, 58, 50, 8, and so on	All summary tables except in the Query Group Setup portlet
!	Use this symbol at the beginning of your search to match alpha or numeric characters that do not contain the alpha or numeric characters. This symbol can be used in conjunction with any other symbol or wildcard character.	!cat33	cat32, cat34, and so on	All summary tables except in the Space Usage and Query Group Setup portlets

Clearing Filters

You can clear the column filters from the table.

1. Do one of the following:

Option	Description
Clear individual column filters	Select  on the filter box. 
Clear all column filters	Select  in the table header and select Clear Filters .
Clear all column filters	Select Clear Filters at the bottom of the table. 

Exporting Table Data

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains all available data up to one million rows.

If filters are used, only filtered data is exported for most portlets. If a portlet displays more than one million rows of data, then use the filters to display the data you want before exporting it.

The format for the time, date, and some numeric values differs in the view and exported .csv file.

1. Click  in the table header and select **Export**.
2. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.

Database Login Settings

Several portlets require you to log in to the database to take actions such as aborting a query or session, changing a workload, changing a priority, or releasing a query from the delay queue.

Database login information is used with your session and is retained until your current Viewpoint session ends.

The login dialog box fields are listed as following.

Field Name	Description
Username	Database user ID.
Password	Database password.
Account String	Characteristics for profile member sessions in the database. An account string identifies which account is charged for the space used by the user and session and may include a priority-level Performance Group prefix code which establishes the session priority.
Authentication Mechanism	Method of identifying an account to the database. Options might include TD2 or LDAP. If an authentication mechanism is not selected, the login uses the default authentication mechanism for the database.
Character Set	Character set the system uses to communicate with the database. If not defined, Teradata Viewpoint uses Teradata system settings.

Rewind Feature

The rewind feature lets you view data from the past and compare it to data for a different date and time. You can rewind the data for some or all portlets on a portal page to a previous point in time, such as when a job failed. Rewinding portlet data helps identify and resolve issues.

You can rewind data as far back as data is available. The rewind feature is not available for portlets that have portlet-specific methods for reviewing data over time. The rewind feature works uniquely in the **Query Spotlight** portlet. See the Query Spotlight section for further details.

You can enter a specific date and time or select data in increments such as 1 day. All available portlets display data that corresponds to the selected time.



If portlet data can be rewound, ⌚ appears in the portlet frame.

If the rewind toolbar does not appear on the portal page, you are not authorized to use the feature. If ⌚ does not appear in the portlet frame, the rewind feature is not available for that portlet.

Rewinding a Portal Page

Some portlets let you rewind data. Use rewind to see changes in data over time among portlets displayed on the portal page.

1. In the upper right corner, select **Rewind On** to turn on the rewind feature.
The rewind toolbar appears along the top of the portal page.



2. Select a date, time, or time increment.
3. View the older data.

Pausing a Portlet During Rewind

You can pause an individual portlet on a portal page while other portlets on the page display rewind data. Pausing an instance of a portlet freezes the display of data at the moment the portlet is placed in pause mode, allowing you to compare data for multiple rewind time periods.

1. Select ▼ in the portlet frame.
2. Select **Pause**.
Data displayed in the portlet freezes and does not change until you resume the portlet refresh.



Resuming a Portlet During Rewind

You can resume using the rewind feature for an individual portlet that is in pause mode on a portal page.


1. Select ▼ in the portlet frame.
2. Select **Resume**.
Data displayed in the portlet is updated to the selected rewind time and does not change until the selected date or time for the rewind feature is changed.

Stopping Rewind Participation

You can prevent an individual portlet from participating in rewind when the portlet frame displays .

1. Next to , select ▼.
You can stop rewind participation when the portlet is in pause mode or actively responding to the rewind date and time controls.
2. Select **Unlink from Rewind**.
The icon changes to  and the portlet no longer responds to rewind date or time controls and is updated to display the most recent data available from the Teradata server.

Restarting Rewind Participation

You can have an individual portlet not participating in rewind resume participation when the portlet frame displays .

1. Select ▼ in the portlet frame.
2. Select **Link to Rewind**.
The portlet returns to the mode it was in when it was unlinked from participation in rewind. For example, if the portlet was in pause mode when unlinked from rewind, it is in pause mode when participation in rewind resumes.

Turning Off the Rewind Feature

You can turn off the rewind feature and resume display of current data in all portlets on the portal page.

1. In the upper left corner of the page, select **Off** to turn off the rewind feature.
All portlets on the portal page that were participating in rewind activities end participation and display the most recent data available.

Pause Feature

Any rewindable portlet can be paused from refreshing data. Pausing an instance of a portlet freezes the display of data at the moment the portlet is placed in pause mode, allowing you to compare data for different time periods. The rewind feature does not need to be active to pause the portlet. When a portlet is in pause mode, the time stamp flashes, indicating that the data in the portlet is no longer being refreshed.

Pausing Portlet Refresh

You can pause an individual portlet on a portal page.

1. Select ▼ in the portlet frame.
2. Select **Pause**.
Data displayed in the portlet freezes and does not change until you resume the portlet refresh.

Resuming Portlet Refresh

You can resume a portlet refresh when the portlet is in pause mode.

1. Select ▼ in the portlet frame.
2. Do one of the following:
 - Select **Refresh** to bring the most recent data collection from the Viewpoint server to the portlet.
 - Select **Resume** to resume the portlet refresh.

Shared Portlets

Use the shared portlets feature to make a customized or preconfigured version of a Teradata Viewpoint portlet available to other users. When a shared portlet is created, all the settings you selected when you customized the portlet are used.

Any user can customize a shareable portlet and submit it for sharing. Depending on your Teradata Viewpoint system configuration, the portlet is not available to authorized users until the Teradata Viewpoint Administrator approves the request. Only the Teradata Viewpoint Administrator can delete the shared portlet entirely. If deleted, existing instances of a shared portlet are removed from the user's page.

The shared portlets feature allows you to:

- Share customized views of a Teradata Viewpoint portlet with other users.
- Create different configurations of a portlet for different roles.
- Assign a preconfigured portlet to a role as a default instance of the portlet.
- Assign a name and description to the preconfigured portlet.

The portlet name and description appears under the **SHARED PORTLETS** category on the **Add Content** screen. This category only appears when at least one shared portlet is enabled for the user who is logged on.



Creating a Shared Portlet

Use the shared portlets feature to customize a Teradata Viewpoint portlet and make it available to other users.

The Teradata Viewpoint Administrator can use **Portlet Library** to edit the shared portlet description after it has been created and approved.

Teradata Viewpoint Administration portlets cannot be shared and the following portlets cannot be shared:


- **Application Queries**
- **Data Mover**
- **Elastic Performance**
- **Hadoop Services**
- **Performance Data Collection**
- **Query Log**
- **SQL Scratchpad**
- **Stats Manager**
- **Unity**
- **Viewpoint Calendar**
- **Workload Designer**

1. Add a portlet instance that you want to share to your portal page.
2. Click  in the portlet frame and select **Settings**.
3. Customize the portlet using the available settings.
4. Click **OK** to save the settings.
5. Click  and select **Share**.
6. Enter a name up to 25 characters.
7. Enter a description.
8. Click **Submit**.
9. [Optional] To make the shared portlet available to the user in the **Add Content** screen, do one of the following:
 - In **Portlet Library**, make sure the portlet has been enabled.


- In the **General** portlet, make sure the **Enable newly created shared portlets** check box is selected.

Editing a Shared Portlet

You can change the name and description of a shared portlet.

1. From the **Portlets** view, select the **Shared Portlets** tab.
2. Select .
3. [Optional] Enter a new name up to 25 characters.
You can use spaces, alphanumeric characters, and underscores (_).
4. [Optional] Enter a description.
5. Click **OK**.
Your changes appear in the **Shared Portlets** tab.

Deleting a Shared Portlet

1. From the **Portlets** view, select the **Shared Portlets** tab.
2. Select  that corresponds to the shared portlet you want to delete.
A confirmation message appears.
3. Select **Delete**.
The portlet name no longer appears on the **Add Content** screen in the portal, and the portlet is no longer available to Teradata Viewpoint users. Existing instances of a shared portlet are removed from the user's page.

Profile

Use the **Profile** portlet to edit your personal user account settings and manage the association between your Teradata Viewpoint account and your database profile information:

Personal Info

Change your name, password, email address, and locale, and select the time zone.

Note:

LDAP users cannot change some of these settings.

Teradata Accounts

Configure settings for each of your database user accounts. Link your Teradata Viewpoint account to the Teradata systems you are authorized to access.

Access the **Profile** portlet from the Teradata Viewpoint portal.

Personal Info Tab

Use the **Personal Info** tab in the **Profile** portlet to view account activity and manage your Teradata Viewpoint user account settings:

- Change your first and last name.

Note:

This option is not available for LDAP users.

- Change your email address.

Note:

This option is not available for LDAP users.

- Change your locale.
- Select a time zone.
- Change password.
- View account activity, including Previous Login, Previous IP, Current Login, and Current IP.

Note:

The previous logon and previous IP appear in the footer of the portal page.

Changing a Name

You can change the first and last name for a Teradata Viewpoint user account.

Note:

This option is not available for LDAP users.

1. From the **Profile** portlet, select the **Personal Info** tab.
2. Change the name in the **First Name** and **Last Name** fields.
3. Click **Apply**.
A confirmation message appears.

Setting a Contact Email Address

You can set your **Email** address using the **Personal Info** tab. This email address is used for identification and communication in the Teradata Viewpoint system, including notifications and reminders.

The email address cannot contain uppercase letters. Uppercase characters are automatically converted to lowercase.

Note:

This option is not available for LDAP users.

1. From the **Profile** portlet, select the **Personal Info** tab.
2. Enter your full address in the **Email** field.
3. Click **Apply**.
A confirmation message appears.

Managing Passwords

You can manage Teradata Viewpoint account passwords. A password must be at least 8 and no more than 25 characters. It must include 1 uppercase, 1 lowercase, 1 digit, and 1 special character.

Users cannot change their password if their account is externally authenticated. The fields required to change the password are not available when external authentication is used. To change the password for an externally authenticated profile, follow your company policy for password management.

Note:

This option is not available for LDAP users.

1. From the **Profile** portlet, select the **Personal Info** tab.
2. Select **Change Password**.
3. Enter your **Current Password**.
4. Enter a **New Password**.
5. Confirm your new password by retyping it in **Re-enter Password**.
6. Select **Change Password**.

Setting the Locale

You can choose a different locale in the **Profile** portlet to override the locale set by the Teradata Viewpoint Administrator.

The locale is used to display content in the Viewpoint portal. Dates, times, and numbers are displayed in Viewpoint portlets in the format used by the selected locale. For example, the **English (US)** locale displays content in English and dates in a *mm/dd/yy* format. If the language pack for the selected locale is installed, all content appears in the selected language. If the language pack for the selected locale is not installed, content appears in English and dates, times, and numbers appear in the format for the selected locale.

1. From the **Profile** portlet, select the **Personal Info** tab.

2. Select the locale from the list.

Option	Date, Time, and Number Display Format
English (US)	mm/dd/yy 12h 1234.57
English (UK)	dd/mm/yy 24h 1234.57
Chinese (China)	yy-mm-dd 24h 1234.57
Chinese (Taiwan)	yy/mm/dd 24h 1234.57
French	dd/mm/yy 24h 1234,57
German	dd.mm.yy 24h 1234,57
Japanese	yy/mm/dd 24h 1234.57
Korean	yy. mm. dd 12h 1234.57
Spanish (Argentina)	dd/mm/yy 24h 1234,57
Spanish (Mexico)	dd/mm/yy 24h 1234.57
Spanish (Spain)	dd/mm/yy 24h 1234,57

The Teradata Viewpoint application server locale is the default.

3. Click **Apply**.

Setting the Time Zone

You can adjust your local Teradata Viewpoint time zone setting to compensate for different time zones. Select the local time zone to display all time and date information in the local date and time, even if the Teradata system you are monitoring is located in a different time zone.

The time zone you enter compensates for the time difference between the monitored system location and the Teradata Viewpoint reporting location.

1. From the **Profile** portlet, select the **Personal Info** tab.
2. Select a time zone from the list.
Time zones shown are relative to Greenwich Mean Time (GMT) and include a reference location. The Teradata Viewpoint application server time zone is the default. Use the scroll bar to view more time zones.
3. Click **Apply**.
A confirmation message appears.

Teradata Accounts Tab

Use the **Teradata Accounts** tab in the **Profile** portlet to manage Teradata Viewpoint account login and profile information for each Teradata system you have access to. Teradata account information is used by the **My Queries** portlet. For each Teradata system, you can add and delete Teradata Viewpoint accounts.

System accounts added using the **Teradata Accounts** tab must already exist on the monitored system. The system validates the **USERNAME**, **PASSWORD**, **ACCOUNT STRING**, and **AUTHENTICATION** fields prior to confirmation.

	SYSTEM	USERNAME	ACCOUNT STRING	AUTHENTICATION
▼	ACCT	ADMIN		DEFAULT
▼	PROD	BOB		DEFAULT
▼	SKATE	DIANE		DEFAULT
▼	TEST	JOE		DEFAULT

Adding Database Login Information to a Viewpoint User Account

This task describes how to associate your Teradata Viewpoint account with Teradata systems.

1. From the **Profile** portlet, select the **Teradata Accounts** tab.
2. Select **Add Account**.
3. Select a system from the list.
4. Enter the following information in the appropriate fields:
 - Username
 - Password
 - [Optional] Account string
5. [Optional] Select an **Authentication Mechanism** from the list for use on the Teradata system.
6. Select **Add**.

Deleting a Teradata Viewpoint User Account from a Teradata System

This task describes how to delete a Teradata Viewpoint user account from a Teradata system.

1. From the **Profile** portlet, select the **Teradata Accounts** tab.
2. Select ▼ next to a row then select **Delete**.

A confirmation message appears.

Viewpoint RESTful APIs

Viewpoint RESTful APIs Overview

Teradata Viewpoint provides a number of RESTful APIs that enable you to perform various operations programmatically without using the web portal. The Viewpoint Web Services can be enabled and disabled through the Viewpoint portal using the **Roles Manager** portlet.

APIs are available as a part of Viewpoint software deployment and do not require any additional installation.

For more information, see *Teradata® Viewpoint and Teradata® Data Lab API Reference Guide*.

Administration Portlets

System Administration

Administration Portlets

The Teradata Viewpoint administration portlets allow the Teradata Viewpoint Administrator to provide access to Teradata Viewpoint resources and information.

You can access these portlets from the Teradata Viewpoint portal page if your role has permission.

Alert Setup

Configure the alert delivery settings and actions.

Backup

Configure the backup of Teradata Viewpoint server data.

Certificates

Manage trusted certificate authorities and HTTPS certificates.

General

Configure Teradata Viewpoint settings.

LDAP Servers

Configure the LDAP servers for Teradata Viewpoint to authenticate users and assign user roles.

Monitored Systems

Add the systems and configure the data collectors that provide data to portlets. You also can add and configure a managed system available to display in the **Viewpoint Monitoring** portlet.

Portlet Library

View the installed portlets and specify which portlets can be enabled.

Query Group Setup

Manage the sets of queries available to users in the **Query Groups** and **Application Queries** portlets. Can also be used to define the criteria that associate a query with a particular application in the **Query Log** portlet.

Roles Manager

Manage roles and specify the level of access users are given.

Shared Pages

Manage shared pages and how they are viewed by users.

User Manager

Manage user accounts and assign users to roles.

Security and Permissions

The Teradata Viewpoint administration portlets allow the Teradata Viewpoint Administrator to control portlet access at different levels by setting permissions for Teradata Viewpoint users.

Global

Enable or disable access to portlets globally using **Portlet Library**. This setting takes precedence over all other portlet permission levels.

Role


Assign permissions to classes of users called roles using the **Roles Manager** portlet. Each tab controls access to features.

User

Assign available roles to users and set role precedence using **User Manager** portlet. Roles assigned higher priority take precedence over lower-priority roles.

Accessing Administration Portlets

You must have permission to view these portlets.

1. From the Teradata Viewpoint portal page, select  in the top right corner of the page.
2. Do one of the following to access a portlet:
 - Select a portlet from the **Select Portlet** list.
 - Select any portlet on the page.
3. Do one of the following to change the page view:
 - Select **Show All** to return to the **Viewpoint Administration** list.
 - Select **Close** to close the Viewpoint Administration view and return to the Teradata Viewpoint portal page.

Alert Setup

The **Alert Setup** portlet enables you to configure delivery settings and actions that are triggered when the alert service generates an alert. After setting alert actions in the **Alert Setup** portlet, you can use the following portlets to define threshold levels and triggers to generate alerts.

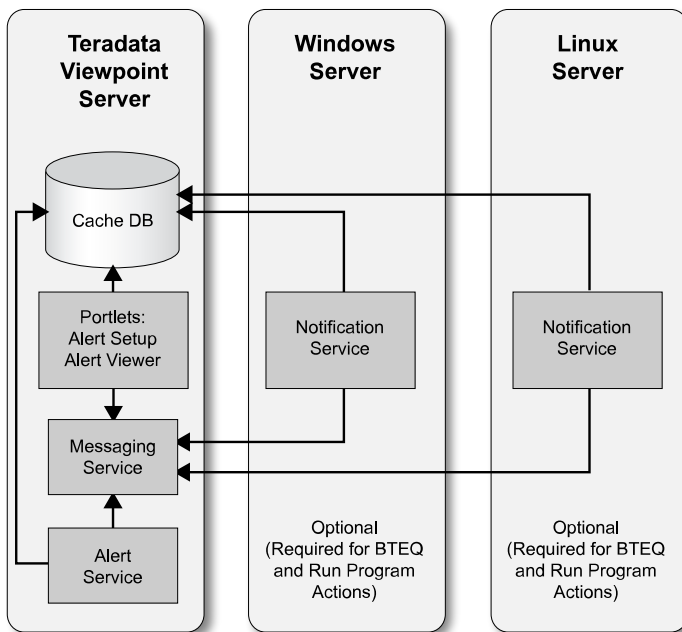
Portlet	Interaction
Ecosystem Configuration	Configure rules to trigger alerts. See <i>Teradata® Ecosystem Manager User Guide</i> for more information.
Monitored Systems	Add alerts to a system and set rules to trigger alerts.
Performance Data Collection	Add job and database space alert configuration to trigger alerts.
Server Management	Select actions for Server Management summary alerts. See <i>Teradata® Server Management Web Services User Guide</i> for more information.
Workload Designer	Configure rulesets to trigger an alert when an exception to the workload occurs.

For example, you can do the following:

- Use the **Ecosystem Configuration** portlet to trigger a BTEQ script to be run when a Heartbeat event occurs for a defined application server.
- Use the **Monitored Systems** portlet to configure an SNMP trap to be sent if the current perm space used on any database in a Teradata system exceeds 80%.
- Use the **Performance Data Collection** portlet to configure an email alert to be sent when a job fails to complete successfully or when database space reaches a specified threshold.
- Use the **Server Management** portlet to configure a trap to be sent to an SNMP management server when a database alerts with a Warning severity.
- Use the **Workload Designer** portlet to configure an email alert to be sent when the CPU usage for an Analytics Database request exceeds 1 million CPU seconds.

Overview of Teradata Alerts Integration

Teradata Alerts consists of several components that are included in the Teradata Viewpoint server. Teradata Alerts can also use the Windows and Linux Notification Services to run BTEQ scripts and executable programs.



Cache DB

A local repository instance used by the Teradata Viewpoint server. The Cache DB stores delivery settings and action set definitions used by the alert service.

Portlets

The **Alert Setup** portlet is used to configure alerts.

The **Alert Viewer** portlet displays logged alerts.

Messaging Service

The Java Message Service provider (Apache ActiveMQ) that enables message exchange between the software components.

Alert Service

The Java daemon that processes alerts.

Notification Service

The service that can execute BTEQ scripts and executable programs to actively respond to generated alerts. You must activate this service on a Windows or Linux server you provide.

Alert Configuration Workflow

The alert configuration workflow has the following phases:

1. Configure delivery settings and alert action sets.
2. Define the triggers, conditions, and thresholds that cause alert actions to be generated.
3. View generated alerts in the **Alert Viewer** portlet.

The alert configuration workflow has several access points:

Alert Setup

Define the hours during which alert actions run, and define notifications in **Alert Exceptions** to notify you if an alert action set fails to execute properly.

Alert Triggers

After alert action sets are configured, they are available to other portlets in the Teradata Viewpoint portal. Define triggers, conditions, and threshold levels and select corresponding alert action sets.

Database

Integrate the database with Teradata Alerts directly by inserting rows in the AlertRequest and MonitorRequest tables to trigger alert actions. The Alert Request data collector forwards the row contents to the alert service for processing.

Alert Viewer

The **Alert Viewer** portlet displays alerts generated from many sources. The alert includes the date, time, alert name, type, source, and system for each alert generated if it was configured to be logged during alert setup. Additional details, including the alert property and metric that triggered the alert, can also be displayed.

Delivery Settings

You can configure delivery settings for the following alert actions:



- Write to the alert log
- Send an email notification
- Send an SNMP notification
- Run a SQL query
- Run a BTEQ script
- Run a Windows or Linux executable program



Adding a Login Account

The Teradata Viewpoint Administrator or Teradata System Administrator must define a login configuration before you can create an action set that runs a BTEQ script or a SQL action.

A login configuration must be enabled before BTEQ scripts or SQL queries in an action set may be executed.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Authentication** list, select **Teradata Login**.
3. From the **Teradata Login** list, do one of the following:

- To add a login configuration, select .
 - To copy a login configuration, select  next to the login configuration you want to copy.
 - To edit a login configuration, select the login configuration name.
4. Enter the **TDPID** of the database you want to log in to.
The BTEQ scripts and SQL queries you provide as alert actions run against the TDPID you specify.
 5. Select the **Enable** check box to enable the login configuration.
 6. Under **Login**, enter a **Username** and associated **Password** to be used for BTEQ and SQL actions.
Account String is optional.
 7. [Optional] For BTEQ, enter **BTEQ Authentication Options**.
 - **Logmech Name** is the name of the login mechanism.
 - **Logmech Data** is the login parameters or credentials associated with the selected login mechanism.
 8. For SQL, under **JDBC Authentication Options**, select the **Logmech Name** from the list.
Login mechanism data is automatically generated from the user name and password, if necessary.
 9. Select the **Session Character Set** from the list.
 10. [Optional] To test the Teradata login:
 - a. Select the login configuration you want to test from the **Test Teradata Logon** list. If the host name is known, it is displayed.
 - b. Select **Test** to verify that the alert service can execute the login.

If the operation is successful,  appears. If the operation fails,  appears.

 - If a JDBC login fails, verify that the alert service is running properly on the Viewpoint server.
 - If a BTEQ login fails, verify the Teradata Notification Service is running properly on the Windows or Linux server where it was installed. In addition, examine the `cam.properties` file (Linux: `/etc/opt/teradata/cam/cam.properties`, Windows: `C:\etc\opt\teradata\cam\cam.properties`) to make sure that the `cam.activemq.host` property is set to the hostname of the Viewpoint server.
 11. Click **Apply**.

Defining Alert Retention

You can set limits that define how long to retain alert log data and how much to retain.

If you set a retention period, alert log data is deleted in weekly increments after the specified period expires. If you set a file-size restriction, alert log data is deleted in weekly increments after the specified file size is exceeded, and the oldest data is deleted first.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Delivery Types** list, select **Alert Viewer**.
3. Under **Delete Alert Viewer Data**, do any of the following:
 - To set the retention period, select **After**.

Type an integer number in the box and select a time period from the list.

- To set the file size restriction, select **Over**.

Type an integer number in the box and select a file-size unit from the list.

4. Click **Apply**.

Defining Email Configuration

You can configure the alert service to send alert notifications by email. Once you configure the email notification using these steps, you must also create an action set.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Delivery Types** list, select **Email**.
3. In the **SMTP Host** box, enter the SMTP host address of the outgoing (SMTP) email server.
4. Select the **Enable** check box.
5. [Optional] Select the **Use SSL** check box to use Secure Sockets Layer as the transmission protocol. If you use SSL, you must also install the SMTP Host certificate using the **Certificates** administration portlet.
6. Under **Port**, select one of the following:

Option	Action
Use default port	Use the default communications port on the SMTP host for outgoing email.
Enter port number	Type a number in the box to use as the custom port number on the SMTP host for outgoing email.



7. In **Server Timeout**, type the number of seconds to wait for a connection with the SMTP host before the system times out.
The default is 30 seconds.
8. Under **Reply To**, type the **Email Address** and **Display Name** to appear in alert email messages.
9. Under **Login**, select one of the following:

Option	Action
Anonymous login	Log in to the SMTP host anonymously.
Enter credentials	Type a Username and Password to log in to the SMTP host.

10. Leave the **Advanced (Optional Java mail)** box blank.

Note:

The alert service uses this box for diagnostic purposes. Type information in this box *only* under the guidance of a Teradata Technical Support Specialist.

11. To verify your email delivery settings, type a valid email address for **Test Recipient** and select **Test**.
If the operation is successful,  appears. If the operation fails,  appears.
If the operation is successful, the alert service sends an email message to the test recipient address.
12. Click **Apply**.

Postrequisite:



After the email configuration is completed, create an action set.

Related Information:

[Adding and Editing Action Sets](#)

Adding SNMP Configurations



You can configure the alert service to deliver alert notifications to third-party management applications using SNMP. After you define the SNMP configuration using these steps, you must also create an action set.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Delivery Types** list, select **SNMP**.
3. From the **SNMP** list, do one of the following:
 - To add an SNMP configuration, select  next to **SNMP**.
 - To copy an SNMP configuration, select  next to the SNMP configuration you want to copy.
 - To edit an SNMP configuration, click the SNMP configuration name.
4. Type a name in the **Configuration Name** box.
5. Select the **Enable** check box.
6. [Optional] Select **Default for Alert Request Collector** to make this the default configuration for SNMP requests that originate from the AlertRequest and MonitorRequest tables in the dbcmngr database.



Note:

When you make this configuration the default, the alert service overwrites the previous default configuration for SNMP requests.

7. Type the IP address or host name of the destination for SNMP alert notifications in the **Destination** box.
If using TCP, update **Destination** to `tcp:hostname:port`. The port number should be taken from the snmpTrapdAddr field in the snmptrapd conf.
8. Type an SNMP community name in the **Community** box.
The default value is public.

9. [Optional] You can add  or remove  **Destination** and **Community** pairs.
10. Under **Include in SNMP Trap Details**, configure the information to display in the SNMP trap details. Select at least one of the check boxes. If *both* are selected, the SNMP trap details include a description (the conditions that triggered an alert) followed by the message.
 - Select **Alert message** to include the message in the SNMP Trap details. For Monitored Systems alerts, this message is composed when defining the alert rules.
 - Select **Alert Conditions** to display the conditions that triggered the alert in the SNMP trap details.
11. [Optional] To test the SNMP trap:
 - a. Select a Management Information Base (MIB) from the **Test SNMP Trap** list. The following table provides information on the MIB location and OID for the alert service and Ecosystem Manager.

Provider	MIB Location	OID
Alert Service	/opt/teradata/cam/alert/config/teradataCamMib.txt on the Teradata Viewpoint server	iso.org.dod.internet.private.enterprises.teradata.td-products.tdcam
Ecosystem Manager	/opt/teradata/emserver/conf/Teradata-TSM-MIB1.txt	iso.org.dod.internet.private.enterprises.teradata.td-products.tmsm.dualActive

- b. Select **Test** to verify the SNMP configuration on the specified destination host or hosts. If the operation is successful,  appears. If the operation fails,  appears. Verify the settings are correct and try again.

If the operation is successful, the alert service sends an SNMP trap (notification) to each specified destination. Verify the SNMP trap is received at each destination.

12. Click **Apply**.

Postrequisite:

After the SNMP configuration is completed, create an action set.

Related Information:

[Adding and Editing Action Sets](#)

Adding SQL Query Configurations

You can configure the alert service to run SQL query alert actions for Teradata systems. After you configure the alert service using these steps, you must also create an action set.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Delivery Types** list, select **SQL Queries**.

3. From the **SQL Queries** list, do one of the following:
 - Select ☐ next to **SQL Queries** to add a SQL query.
 - Select ☐ next to the SQL Query configuration you want to copy.
 - Click the SQL Query name to edit its configuration.
4. Type a name for the SQL configuration in the **Name** box.
5. Select the **Enable** check box.
6. [Optional] Type the name of the default database for the SQL query in the **Default Database** box.
7. Type the SQL query for the alert action in the **SQL** box.
8. [Optional] Under **Timeout**, select an action to perform if a SQL query is still running after a specified time period.
 - Select the **Notify After** check box to be notified, and indicate the time period. You configure notifications by selecting **Alert Exceptions** from the **Setup Options** list.
 - Select the **Terminate After** check box to terminate the SQL query, and indicate the time period.

A notification is always sent when a SQL query is terminated.
9. Click **Apply**.

Postrequisite:

After the SQL query configuration is done, you must create an action set.

Related Information:

[Configuring Alert Exceptions](#)
[Adding and Editing Action Sets](#)

Configuring BTEQ Script Settings

The alert service can run BTEQ scripts as alert actions for Teradata systems.

When the Teradata Notification Service is installed and the service is running, the **Alert Setup** portlet automatically displays a list of the BTEQ scripts that have been installed on the Windows or Linux server.

After you configure the BTEQ script settings, you must also create an action set.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Notification Service** list, select the operating system on which the BTEQ script is installed.
3. Select the name of the BTEQ script that you want to configure.
4. [Optional] Under **Timeout**, select an action to perform if a BTEQ script is still running after a specified time period.
 - Select the **Notify After** check box to be notified, and indicate the time period. You configure notifications by selecting **Alert Exceptions** from the **Setup Options** list.
 - Select the **Terminate** check box to terminate the BTEQ script and indicate the time period after which the script is terminated.

- Click **Apply**.

Postrequisite:

After the BTEQ configuration is done, you must create an action set.

Related Information:

[Adding and Editing Action Sets](#)

[Configuring Alert Exceptions](#)

BTEQ Script Processing

When you configure Teradata Alerts to run a BTEQ script as an alert action, the process creates a logon session with the Teradata system, runs the script to completion, and ends the session. Teradata Alerts processes BTEQ scripts sequentially as shown in the following table.

BTEQ Command	Description	Details	Executed
.LOGON.	Creates a logon session	<p>Enters the TDPID, Account ID, Username and Password that you used when creating the BTEQ configuration</p> <p>Note: Because you define login credentials in the Alert Setup portlet, make sure that BTEQ scripts do not contain the .LOGON. command.</p>	Automatically
.LOGMECH.	Provides extra information for a logon session	Enters the optional Logmech Name	Must define Logmech Name when creating the BTEQ configuration
.LOGDATA.	Provides extra information for a logon session	Enters the optional Logmech Name and the corresponding Logmech Data	Must define Logmech Name and the corresponding Logmech Data when creating the BTEQ configuration
.RUN FILE.	Executes the BTEQ script	<p>Executes the commands in the BTEQ script in <i>InstallFolder\Teradata\Client\16.20\Teradata Notification Service\sql</i> that is available from the Script list when you created a BTEQ action set</p>	Automatically

BTEQ Command	Description	Details	Executed
.QUIT.	Logs out of the session		Automatically

Configuring User Program Settings

You can configure the alert service to run a program as an alert action.

When the Teradata Notification Service is installed and the service is running, the **Alert Setup** portlet automatically displays a list of the programs that have been installed on the Windows or Linux server.

After you configure the user program using these steps, you must also create an action set.

1. From the **Setup Options** list, select **Delivery Settings**.
2. From the **Notification Service** list, select the operating system on which the program is installed.
3. Select the name of the user program you want to configure.
4. [Optional] Under **Timeout**, select an action to perform if the user program is still running after a specified time period.
 - Select the **Notify After** check box to be notified, and indicate the time period. You configure notifications by selecting **Alert Exceptions** from the **Setup Options** list.
 - Select the **Terminate** check box to terminate the user program and indicate the time period after which the program is terminated.
5. Click **Apply**.

Postrequisite:

After the user program configuration is done, you must create an action set.



Related Information:

[Adding and Editing Action Sets](#)

[Configuring Alert Exceptions](#)

Disabling Alert Configurations

You can disable an alert delivery configuration so that triggered alerts using that configuration are not executed.

When you disable email as a delivery type,  displays next to **Email** under **Delivery Types**. When you disable a BTEQ script or user program,  displays next to the configuration name.

Even though a delivery configuration is disabled, you can still configure action sets that use it. When you do so, the disabled components are highlighted in red on the action set configuration screen.

If the **Include in Alert Viewer** option is selected when configuring an action set, a triggered alert is logged in the alert log and can be viewed in the **Alert Viewer** portlet, regardless of whether any other delivery configuration is enabled.

1. From the **Setup Options** list, select **Delivery Settings**.
2. Do one of the following, depending on the type of configuration you want to disable:

Option	Description
Email	<ol style="list-style-type: none"> a. From the Delivery Types list, select Email. b. Clear the Enable check box next to the SMTP Host box.
SNMP, SQL Query	<ol style="list-style-type: none"> a. From the Delivery Types list, select SNMP or SQL Queries. b. Under SNMP or SQL Queries, select the configuration name. c. Clear the Enable check box, located to the right of the Name box.
BTEQ Script, Program	<ol style="list-style-type: none"> a. From the Notification Service list, select Windows or Linux, depending on the operating system the Teradata Notification Service is running on. b. Under BTEQ Scripts or Programs, click the configuration name. c. Clear the Enable check box, located to the right of the Name box.

3. Click **Apply**.


Deleting Alert Configurations

You can delete an email, SNMP, or SQL query delivery configuration.

You cannot delete a BTEQ script or a user program using the **Alert Setup** portlet. However, if you remove scripts or programs from the Windows or Linux server running the Teradata Notification Service, they are not listed and cannot be configured in the **Alert Setup** portlet.

The following scenarios result in an error:

- You attempt to delete an alert delivery configuration that is part of a defined action set.
 - You attempt to delete the last Teradata login configuration, but BTEQ or SQL action sets that use the Get from Alert option still remain.
1. From the **Setup Options** list, select **Delivery Settings**.
 2. Do one of the following, depending on the type of configuration you want to delete:

Delivery Type	Delete Action
Email	<ol style="list-style-type: none"> a. From the Delivery Types list, select Email. b. Select Delete. A confirmation message appears. c. Select OK.
SNMP, SQL query, Teradata Login	<ol style="list-style-type: none"> a. From the Delivery Types list, select SNMP or SQL Queries. b. Select , located to the right of the name.

Delivery Type	Delete Action
	<p>A confirmation message appears.</p> <p>c. Select OK.</p>

Alert Presets

You can configure alert presets for the following:

- Core hours of operation
- Action sets of multiple alert actions that run in a single operation
- Groups of multiple alert action sets that run in a single operation

Core Hours

The alert service operates 24 hours a day, 7 days a week.

You can schedule alert actions to run during a combination of core, evening, or weekend hours as described in the following table.

Hours	Description
Core Hours	<p>The days and times during the week that you define:</p> <ul style="list-style-type: none"> • A period to include one or more consecutive days of the week. • Time as a period that includes one or more consecutive hours or as all day (24 hours).
Evening Hours	The hours opposite the core hours for core days.
Weekend Hours	The ending time of the ending core day to the beginning time of the beginning core day.

Time Interval

If the core hours are Monday through Friday, 9:00 AM to 5:00 PM, then evening hours are Monday through Friday, 5:00 PM to 9:00 AM, and weekend hours are from Friday at 5:00 PM through Monday at 9:00 AM.

All Day

If the core hours are Monday through Friday, 24 hours, then there are no evening hours, and weekend hours are Saturday through Sunday, 24 hours.

Setting Core Hours

This task describes how to set core hours of operation for the alert service. When you make selections in this view, the **Evening Hours** and **Weekend Hours** are updated in the box to the right.

1. From the **Setup Options** list, select **Alert Presets**.
2. From the **Preset Options** list, select **Core Hours**.
3. Under **Core Hours Details**, select the beginning and ending core **Days** from the two lists.
4. For **Time**, do one of the following:
 - Select **24 hours** to specify the entire day on core days.
 - Select beginning and ending times from the two lists.

The hours listed under **Evening Hours** and **Weekend Hours** adjust automatically according to the hours you selected for the core hours.

5. For **Time Zone**, select the time zone to use for the core hours.
6. Click **Apply**.



Adding and Editing Action Sets

Alert actions can be combined into action sets. An action set enables you to run multiple alert actions in a single operation.

Note:

You must define alert delivery settings by selecting **Delivery Settings** from the **Setup Options** list, and then selecting the appropriate option under **Delivery Types** or **Notification Service** before you can create action sets that use those alert actions.

Any disabled components are highlighted in red on the actions set configuration screen. You can configure action sets with disabled components, but disabled delivery types will not be executed. To enable a disabled component, select **Delivery Settings** from the **Setup Options** list, and then under the **Delivery Types** or **Notification Service** list select the **Enable** check box for the disabled components.

1. From the **Setup Options** list, select **Alert Presets**.
2. From the **Preset Options** list, select **Action Sets**.
3. From the **Action Sets** list, do one of the following:
 - Select  to add an action set.
 - Select  in the row of the action set you want to copy.
 - Select the name of the action set you want to edit.
4. In the **Action Set Name** box, type a name.
5. Under **Times**, select the applicable check boxes to run the action set at the hours defined in **Core Hours** under the **Preset Options** list.
6. Under **Actions**, select the check box for any of the following:

Action	Description
Include in Alert Viewer	Logs the alert details so you can view alerts for this action set in the Alert Viewer portlet.

Action	Description
Email recipients	<p>Designates the recipient of an email alert as a Viewpoint user, Viewpoint role, or email address.</p> <ul style="list-style-type: none"> • Select Bcc to hide the recipients. • Select To to make all recipients visible in the email notification. • Use semicolons to separate multiple recipients and use any combination of user, role, or email addresses. <p>The Email recipients action is available if you configured the email settings in Delivery Settings > Email.</p>
SNMP	<p>Lists the SNMP configurations that are defined under Delivery Types.</p> <ul style="list-style-type: none"> • Select an SNMP configuration from the list. <p>The SNMP action is available if you configured the SNMP settings in Delivery Settings > SNMP.</p>
SQL	<p>Lists the SQL queries defined under Delivery Types.</p> <p>a. From the Query list, select a query and from the TDPID list, do one of the following:</p> <ul style="list-style-type: none"> • Select the TDPID of the database on which to run the SQL query, regardless of which system generated the alert. • Select Get from alert to run the SQL query on any database that generated the alert. <p>For Get from alert option, the TDPID must be defined in the Alert Setup and Monitored Systems portlets. If the system is not defined, the alert service writes the failed login attempt to /var/opt/teradata/cam/alert/logs/alertservice.log on the Viewpoint server.</p> <p>The SQL action is available if you configured a SQL query and a Teradata login.</p>
BTEQ	<p>Lists the BTEQ scripts that are located on the Windows or Linux system that is hosting the Teradata Notification Service.</p> <p>a. From the Location list, select a location from where the script is run.</p> <p>b. From the Script list, select BTEQ and from the TDPID list, do one of the following:</p> <ul style="list-style-type: none"> • Select the TDPID of the database on which to run the BTEQ script, regardless of which system generated the alert. • Select Get from alert to run the BTEQ script on any database that generated the alert. <p>For Get from alert option, the TDPID must be defined in the Alert Setup and Monitored Systems portlets. If the system is not defined, the alert service writes the failed login attempt to /var/opt/teradata/cam/alert/logs/alertservice.log on the Viewpoint server.</p> <p>The BTEQ action is available if these conditions are met:</p> <ul style="list-style-type: none"> • The BTEQ delivery settings are configured in the Alert Setup portlet in Delivery Settings > Teradata Login. • The Teradata Notification Service must be installed and running on the Windows or Linux server. • The BTEQ script is located on the Windows server at: <i>InstallFolder\Teradata\Client\16.20\Teradata Notification Service\sql</i>. It is located on the Linux server at: /opt/teradata/client/cam/tdnotification/sql.
Run a program	<p>Lists the programs and batch files located on the Windows and Linux servers that run the Teradata Notification Service.</p> <p>a. From the Location list, select the location where the program is located.</p> <p>b. From the Program list, select the program or batch file to run.</p>

Action	Description
	<p>c. In the Arguments box, type any command line arguments you want to pass as part of the alert.</p> <p>For example, if the batch file on the Windows server is <code>currentSpace.bat</code> and the batch file takes arguments -db for the name of the database and -sp for the percentage of space used, you can type -db TEST1 -sp 90 in the Arguments box. This specifies that the batch file runs with TEST1 as the database name, and 90 as the percentage of space used.</p> <p>Run a program is available if the following conditions are met:</p> <ul style="list-style-type: none"> • The Teradata Notification Service is installed and running on the Windows or Linux server. • The program or batch file is located on the Windows server at: <code>InstallFolder\Teradata\Client\16.20\Teradata Notification Service\usrcmd</code>. It is located on the Linux server at: <code>/opt/teradata/client/cam/tdnotification/usrcmd</code>.

7. Click **Apply**.

Adding and Editing Groups

Alert actions can be combined into action sets and action sets can be combined into groups in the alert service. A group enables you to run multiple action sets in a single operation. You must define action sets before you can include them in a group.

1. From the **Setup Options** list, select **Alert Presets**.
2. From the **Preset Options** list, select **Groups**.
3. From the **Groups** list, do one of the following:
 - Select ☐ next to **Groups** to add a group.
 - Select ☐ next to the group you want to copy.
 - Select a group name to edit the group.
4. Type a name in the **Group Name** box.
5. Under **Include the Following Action Sets**, select the action set to include.
6. [Optional] Add ☐ or remove ☐ action sets.
7. Click **Apply**.

Example: Different Alert Actions for Weekday and Weekend Hours

To designate a set of different alert actions to be performed for an event depending on the time period, define time-specific action sets and assign them to a group.

Suppose you want the alert to provide SNMP notification and send an email message under the following conditions:

- An SNMP trap when a space-usage threshold is exceeded on a weekday

- An SNMP trap and an email message sent to the Teradata DBA when the threshold is exceeded on a weekend

To do this, define two action sets:

- One SNMP notification action for core, evening, and weekend hours
- One email notification action for weekend hours only

Create a group and assign both action sets to the group. In the Teradata Viewpoint administrative portlet where you want to define the alert, assign the group as the alert action.

Example: Managing Email Alert Recipients

Groups and action sets can be used to manage alert recipients, similar to email distribution lists.


Define action sets that send email notifications to different sets of people in your organization:

- Define one action set to send an email notification to the Teradata System Administrator
- Define another action set to send an email notification to end users in the finance department

Create a group and assign both action sets to the group. In the Teradata Viewpoint administrative portlet where you want to define the alert, assign the group as the alert action. When an alert is issued, both sets of users receive an email notification.

Deleting Action Sets or Groups

If you delete a group in the alert service, the group configuration is lost, but the action sets combined in the group remain. You cannot delete an action set if it is used in a group or an alert rule. You cannot delete a group if it is used in an alert rule.

1. From the **Setup Options** list, select **Alert Presets**.
2. From the **Preset Options** list, select **Action Sets** or **Groups**.
3. Select  next to the action set or group configuration you want to delete.
A confirmation message appears.
4. Click **OK**.

Alert Exceptions

An alert action may fail. For example, a BTEQ script may not run because it was deleted from the server. Or an email alert may not be delivered because the Viewpoint user specified in the email action was deleted. You can [configure notifications](#) for such alert exceptions.

Configuring Alert Exceptions

You can designate that an email notification be sent or an alert logged if an alert action fails.

1. From the **Setup Options** list, select **Alert Exceptions**.

2. From the **Configuration** list, select **General** to configure how to be notified about exceptions encountered while processing alert action sets.
3. Under **Repeat**, type the number of minutes in the box.
This setting prevents an alert from being sent more than once in the time specified.

Zero is the default value and is a valid entry, indicating there is no alert inactivity period.
4. For the alert action types for which you want to generate alert exceptions:
 - Under **Severity**, select the level associated with the alert type.
 - [Optional] Under **Send Email**, select the check box to configure an email message to be sent.
 - [Optional] Under **Alert Viewer**, select the check box if you want the alert to be displayed in the **Alert Viewer** portlet.
5. [Optional] Under **Email**, enter one or more email addresses to which an email message is sent if an alert exception occurs.
Separate multiple addresses with semicolons.

You can enter email addresses even if you have not selected any **Send Email** check boxes in the previous step. However, emails are only generated when **Send Email** is enabled.

Database User Alerts



You can set up an alert action so that the database user is notified by email if there is a session or database space alert.

Notification Type	Description
Domain suffix	Define an email domain suffix to append to the database user name. For example, if you define <i>teradata.com</i> as the suffix, the database user <i>user1</i> is emailed at <i>user1@teradata.com</i>
User mapping	<p>Define a database table or view from which Viewpoint selects and return two columns, <i>username</i> and <i>email address</i>. This approach requires that the Database Administrator creates a table or view before the database user alert is configured. The Database Administrator adds rows for specific users to receive emails. Both set and multiset are supported with Latin and Unicode as supported character sets. NOT NULL is not supported. The following example shows a sample table with a username and email.</p> <pre>CREATE MULTISET TABLE sysdba.user_mail ,FALLBACK , NO BEFORE JOURNAL, NO AFTER JOURNAL, CHECKSUM = DEFAULT, DEFAULT MERGEBLOCKRATIO (Username VARCHAR(30) CHARACTER SET UNICODE NOT CASESPECIFIC, EmailAddress VARCHAR(30) CHARACTER SET UNICODE NOT CASESPECIFIC) PRIMARY INDEX (Username);</pre> <p>The User mapping option requires the credentials to execute the query. The results of the query are cached and updated by Viewpoint once a day.</p>

After database user alerts are configured in the **Alerts Setup** portlet, they are available in the **Monitored Systems** portlet for session and database space alerts.

Configuring Database User Alerts

1. From the **Setup Options** list, select **Database User Alerts**.
2. To disable database user alerts, select **Disable database user alerts**.
3. To use database user alerts, select one of the following options:

Option	Description
Domain suffix	Enter the suffix you want to append to the database user name to construct an email address. For example: domain.com
User mapping	<ol style="list-style-type: none"> a. Under System Details, select the System, and specify the credentials, authentication, and character set to use for the system. b. Under Table or View Name, enter a name for the user mapping table or view. c. Select Test to verify that the connection is successful, that the mapping table or view exists, and that a select statement against the mapping object returns two columns containing strings. If the operation is successful,  appears. If the operation fails,  appears. You can hover over the icon to see information about what caused the failure.

Alert Properties

The following table lists the available properties for all alert types, regardless of the alert source. The associated property names enable you to customize alert actions in the **Alert Setup** portlet.

Property	Name
Alert Type	<i>alertType</i>
Alert Name	<i>alertName</i>
Source	<i>source</i>
Timestamp	<i>timestamp</i>
Severity	<i>severity</i>
System	<i>systemName</i>
Additional property available for Run a Program arguments and in SQL queries:	
Message	<i>message</i> Note: When using the Message property as a Run a Program argument, enclose the property token in quotes: "{\$message}".

Note:

If an alert contains a source-specific property having the same name as a common property (for example, *severity*) Teradata Alerts resolves it to the source-specific property by default. You can access the common property by preceding the property name with `cam:`. For example: `cam:severity`

Program Argument to Display All Properties

The Run a Program arguments box supports a special `${***}` token representing all properties associated with the alert. This includes the common properties for all alert types, regardless of the alert source, as well as any application domain-specific properties, such as the database and Teradata Server Management alerts.

Each property is assigned an individual command line argument of the form `name=value`. For example, suppose you define program arguments (on the Action Set screen) like this: `myAlertTest ${***}`. If your program is subsequently executed for a Viewpoint database space alert, it receives arguments similar to this:

```
argv[1]: myAlertTest
argv[2]: cam:alertId=123
argv[3]: cam:alertName=dbspace80
argv[4]: cam:alertType=databaseSpace
argv[5]: cam:source=Viewpoint
argv[6]: cam:timestamp=2013-09-05T12:00:00.000-04:00
argv[7]: cam:severity=MEDIUM
argv[8]: cam:systemName=devhost
argv[9]: cam:tdpid=devhost
argv[10]: cam:description=(${maxSpaceUsedPct} > 80.0)
argv[11]: cam:message=Teradata alert for hr_db
argv[12]: databaseName=hr_db
argv[13]: spaceUsedPct=74.512
argv[14]: maxSpaceUsedPct=80.678
argv[15]: currentSpool=0.0
argv[16]: peakSpool=596.0
```

SNMP Overview

Teradata Alerts provides the ability to send an SNMP Trap Notification as an alert action. The system takes the incoming alert that was generated in response to an alert rule, maps values from the arriving alert into the SNMP trap fields, and then forwards an SNMP trap to the configured destination using the `snmptrap` Linux command.

This section is an overview of the MIB fields and describes how they are mapped from the alert that is being processed, based on how the alert was generated.

Teradata Alerts SNMP MIB File

The Teradata Alerts SNMP MIB file is shipped with the camalerts package, in the following location: /opt/teradata/cam/alert/config/teradataCamMib.txt. The following table shows the Teradata Alerts SNMP MIB file syntax.

Field	Type	Values
alertId	Integer	Unique integer generated by the alert service when the alert is generated.
alertName	String	The alert name as defined by the user.
source	String	The Teradata application that requested the alert or generated the event that triggered the alert. These include: <ul style="list-style-type: none"> • Teradata Viewpoint • Server Management • Teradata Alerts Ecosystem Manager alerts use the Ecosystem Manager MIB, which is not discussed here.
subcomponent	String	Unused.
timestamp	String	A timestamp, formatted as: 2013-05-04T10:06:21.227-08:00.
severity	String	Low, Medium, High, or Critical.
systemName	String	The system that generated the alert. This varies by the source of the alert.
details	String	The alert message or conditions that indicate the reason for the alert. If both the alert message and the alert conditions are selected, details contains both the conditions, then the message, separated by a comma.

SNMP Configurations

The SNMP configurations define where the SNMP trap notifications should be sent. You can configure these in the **Alert Setup** portlet by selecting **Delivery Settings** from the **Setup Options** list, and then **SNMP** from the **Delivery Types** list. Use the **Include in SNMP Trap Details** option to include the alert message or description.

Selecting the **Default for Alert Request Collector** check box identifies the default configuration for the alert request collector. This configuration is used for SNMP alert requests using the DBCMNGR AlertRequest alert mechanism. Although multiple SNMP configurations can be defined, only one can be identified as this default.

Teradata Viewpoint SNMP Alerts

Teradata Viewpoint alerts are defined in the **Monitored Systems** portlet by selecting **Alerts** under **Setup**. While the exact thresholds differ based on the alert type being configured (for example, **System**, **Node**, **Vproc**), the basic components of the alert rule are the same.

The following table illustrates the mapping.

Teradata Alerts SNMP Fields	Value
alertName	AlertName
source	Viewpoint
systemName	The system being configured
severity	The severity
details	The message and/or description: <ul style="list-style-type: none"> • message = Message • conditions = Alert rules displayed as a string. For example, CPU Usage > 70.

Server Management SNMP Alerts

Server Management service-required summary alerts are defined in the **Server Management** portlet, under the **Alerts** configuration.

The following table illustrates the mapping.

Teradata Alerts SNMP Fields	Value
alertName	AlertName
source	Server Management
systemName	The site ID being configured
severity	Severity is automatically mapped based on alert severity: <ul style="list-style-type: none"> • Low = Unknown, OK, information • Medium = Warning • High = Degraded • Critical = Critical, fatal
details	The message or description: <ul style="list-style-type: none"> • message = Message • conditions = Alert rules displayed as a string. For example, Synopsis contains "restart".

Alert Request Data Collector Alerts

The Alert Request data collector monitors the DBCMNGR.AlertRequest table in the database. When rows are inserted into the table, the alert requests are forwarded to the Teradata Alerts alert service.

The following table illustrates the mapping.

Teradata Alerts SNMP Fields	Value
alertId	0 (unused)
alertName	AlertRequest.jobName
source	Viewpoint
systemName	Teradata system for which the Alert Request data collector is configured
severity	Low
details	<p>The message and/or description, as configured in the default SNMP delivery settings:</p> <ul style="list-style-type: none"> • message = AlertRequest.message • conditions = AlertRequest.description <p>Note: For AlertRequests, the description column of the dbcmngr table is used for the Alert Conditions option of the SNMP delivery settings.</p>

Workload Designer SNMP Alerts

Workload Designer can generate alerts for either an event or an exception.

- Events can be:
 - Planned or scheduled, such as a weekend time period begins
 - Unplanned, such as a node going down
- Exceptions are generated when a query exceeds some user-defined threshold.

TDWM generates the alerts using the DBCMNGR.AlertRequest table, using the same mapping as the Alert Request data collector. The contents of the jobName, message, and description columns are specified by the TDWM software.

The following table illustrates the mapping.

Teradata Alerts SNMP Fields	Value	TDWM-Specific Contents
alertId	0 (unused)	

Teradata Alerts SNMP Fields	Value	TDWM-Specific Contents
alertName	AlertRequest.jobName	TDWM: HostID= <i>hostId</i> , SesNum= <i>sessionNumber</i> , ReqNum= <i>requestNumber</i>
source	Viewpoint	
systemName	Teradata system for which the Alert Request data collector is configured	
severity	Low	
details	The message or description, as configured in the default SNMP delivery settings: <ul style="list-style-type: none"> • message = AlertRequest.message • conditions = AlertRequest.description 	Message for exception: TDWM: WD <i>workloadName</i> encountered exception ' <i>exceptionName</i> ' which resulted in ' <i>action</i> ' Message for event: TDWM: Expression ' <i>eventName</i> ' was triggered Description for exception: TDWM: WDIId: <i>workloadId</i> , ExceededCriteria type= <i>type</i> Description for exception: <i>EventName</i>

Alert Configuration Scenarios

This section presents step-by-step scenarios for how you can configure various types of alerts.

The process for setting up alerts is specific to the products and versions that are installed in your environment and the permissions for your role. The details in these scenarios may not apply to your particular system.

How Can I Customize Email Alerts?

This example describes how to customize the body of an email message used in an alert. You can define a threshold for a Teradata system metric when a specific event occurs. You can then specify that the event triggers an alert. The alert service has access to information on several properties of a monitored system, depending on the data collectors that you enable. If you configure the alert to send text, such as an email message, you can select which of the available properties to display in the message.

The following is an overview of the process:

1. Configure the email delivery settings in the **Alert Setup** portlet.
2. Create an action set in the **Alert Setup** portlet.
3. Configure the Teradata system in the **Monitored Systems** portlet.

4. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
5. Define alert rules in the **Monitored Systems** portlet.
6. Review the alert messages that arrive in the email inbox.

Alert Property Syntax

The following guidelines apply to alert properties:

- Enclose the alert property with curly brackets ({ }) and precede the expression with the dollar sign (\$), as in `${alert_property}`.

For example, type `${databaseName}` to return the name of the database that triggers an alert.





- Make sure you enter alert properties with the correct case, because they are case-sensitive.
- If an alert property is misspelled, uses incorrect case, or does not exist in the list of properties, the alert service displays the string literally.

For example, if you use `${databaseNME}` instead of `${databaseName}` when adding the alert property, the alert service returns `${databaseNME}` instead of displaying the name of the database that triggers the alert.

Example: Customizing Email Alerts

Suppose you want to send an email message to `dba@example.com` with medium severity if the current perm space that is used in any database on Teradata system TEST1 exceeds 80%. On the system TEST1, finance is one of the databases you want to monitor.

1. In the **Alert Setup** portlet, configure delivery settings.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **Email**.
 - c. In the **SMTP Host** box, type `smtp.example.com` as the SMTP host address of the outgoing (SMTP) email server.
 - d. Select the **Enable email** check box.
 - e. From the **Port** list, select **Use default port** to use the default port 25.
 - f. In the **Server Timeout** box, type a default of 30 seconds.
 - g. Keep the default **Anonymous login**.
 - h. In the **Reply-To** box, type `doNotReply@example.com`.
 - i. Leave the **Advanced Options** box blank.
 - j. In the **Test Recipient** box, type `dba@example.com`.
 - k. Select **Test** to verify the SMTP server delivers the email.
 - l. Click **Apply**.
2. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Action Sets**.

- c. Select  next to **Action Sets**.
 - d. In the **Action Set Name** box, type CustomEmailAlert.
 - e. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - f. From the **Actions** list, select **Email recipients**.
 - g. Select the **Email recipients** check box, click **Bcc** or **To** , and in the **Email** box, type dba@example.com.
 - h. Click **Apply**.
3. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure TEST1.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In the **System Nickname** box, type TEST1.
 - c. Select the **Enable system** check box to activate the TEST1 system for monitoring.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.
 4. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list , select **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select **Database Space**.
 - d. Select the **Enable Database Space Collector** check box and keep the default settings.
 - e. Click **Apply**.
 5. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Database Space**.
 - d. Select  next to **Alerts**.
 - e. In the **Alert Name** box, type PermExceeds80.
 - f. Select the **Enable alert** check box.
 - g. In the **Severity** list, select **Medium**.
 - h. From the **Threshold** list, select **Current Perm**.
 - i. In the box to the right of the greater-than symbol (>), type 80.
 - j. Select , select the **All databases except the selected databases below** option and include all databases, and select **OK**.
 - k. From the **Action** list, select **CustomEmailAlert**.
 - l. Leave the value unchanged in the **Do not run twice in ... minutes** box.
 - m. In the **Message** box, enter: The database \${databaseName} is more than 80% full.
 - n. Click **Apply**.

Note:

You can include alert properties in the **Message** box by enclosing the property name in curly brackets ({}), and preceding the expression with a dollar sign (\$) as shown previously for the *databaseName*. See [Teradata Alert Metrics and Properties](#) for a description of the alert properties that are available for Teradata Viewpoint alerts.

6. In the email inbox for the account dba@example.com, check for messages from doNotReply@example.com.

Example Email Message

The email message contains the following formatted information when it arrives in the inbox.

```
From: doNotReply@example.com
Sent: Tuesday, August 7, 2012 2:20 PM
To: dba@example.com
Subject: [Alert] permExceeds80 (System: TEST1, Database: finance)
```

The database finance is more than 80% full.

Event Timestamp: 2012-08-07T14:20:27.016-07:00

```
Database Name=finance
Space Used Percentage=85.0
Max Space Used Percentage=85.28
```

Description: (Space Used Percentage > 80.0)

How Can I Configure and Send an SNMP Trap as an Alert Action?




The following is an overview of the process:


1. Configure the SNMP delivery type in the **Alert Setup** portlet.
2. Create an action set in the **Alert Setup** portlet.
3. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
4. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
5. Define alert rules in the **Monitored Systems** portlet.
6. Review alerts in the **Alert Viewer** portlet.

Example: Customizing and Sending an SNMP Trap as an Alert Action

Suppose you want to send a trap to an SNMP management server when the current perm space that is used by any database on Teradata system TEST1 exceeds 80%. You want the trap to include the name of the database on which the threshold was exceeded as well as the actual percent of perm space in use for that database.

For this example, the SNMP management package is running on a system named prodMgr1. To configure and send an SNMP trap as an alert action, perform the following steps:

1. In the **Alert Setup** portlet, create an SNMP configuration.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **SNMP**.
 - c. Next to **SNMP**, select .
 - d. In the **Configuration Name** box, type snmpConfig.
 - e. Select the **Enable SNMP** check box.
 - f. In the **Destination** box, type prodMgr1 as the hostname for the trap.
 - g. Select the **Alert message** check box.
 - h. Select the **Alert Conditions** check box.
 - i. Click **Apply**.
2. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
 - c. Verify the core hours displayed correspond to the weekday hours.
 - d. From the **Preset Options** list, select **Action Sets**.
 - e. Select  next to **Action Sets**.
 - f. In the **Action Set Name** box, type sendTrap.
 - g. Under **Times**, select all the check boxes: **Core**, **Evening**, and **Weekend**.
 - h. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - i. Under **Actions**, select the **SNMP** check box and from the **Config** list select snmpConfig.
 - j. Click **Apply**.
3. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure TEST1.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In the **System Nickname** box, type TEST1.
 - c. Select the **Enable system** check box to activate the TEST1 system for monitoring.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.

4. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select **Database Space**.
 - d. Select the **Enable Database Space Collector** check box and keep the default settings.
 - e. Click **Apply**.
5. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Database Space**.
 - d. Select ☐ next to **Alerts**.
 - e. In the **Alert Name** box, type PermExceeds80.
 - f. Select the **Enable alert** check box.
 - g. From the **Threshold** list, select **Current Perm**.
 - h. In the box to the right of the greater-than symbol (>), type 80.
 - i. Select , select the **All databases except the selected databases below** option and include all databases, then select **OK**.
 - j. From the **Action** list, select **sendTrap**.
 - k. In the **Message** box, enter: Database \${databaseName} current perm use is \${spaceUsedPct}%.
 - l. Click **Apply**.

Note:

You can include alert properties in the **Message** box by enclosing the property name in curly brackets ({}) and preceding the expression with a dollar sign (\$) as shown previously for *databaseName* and *spaceUsedPct*. See [Teradata Alert Metrics and Properties](#) for a description of the alert properties that are available for Teradata Viewpoint alerts.

When the perm space usage for a database exceeds 80%, a trap is sent to the configured trap destination, in this example, *prodMgr1*. Because you configured the SNMP delivery settings with both the **Alert conditions** and **Alert message** options enabled, the trap **Details** box contains the alert rule conditions of the event that generated the trap followed by the message you provided in the PermExceeds80 alert rule: (Space Used Percentage > 80. 0), Database finance current perm use is 81.9%.

6. In the **Alert Viewer** portlet, review triggered actions.

How Can I Enable a BTEQ Script as an Alert Action?



This example shows how to enable a BTEQ script when creating action sets in the **Alert Setup** portlet. The following is an overview of the process:

1. Copy BTEQ scripts to the server that is running the Teradata Notification Service.
2. Verify that the Teradata Notification Service is running.
3. Configure the delivery type for the Teradata system in the **Alert Setup** portlet.
4. Create an action set in the **Alert Setup** portlet.
5. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
6. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
7. Define alert rules in the **Monitored Systems** portlet.
8. Review alerts in the **Alert Viewer** portlet.

Example: Enabling a BTEQ Script to Run as an Alert Action

In this example, run a BTEQ script as an alert action for Teradata system TEST1, weeknights only, and track alerts in the **Alert Viewer** portlet. The BTEQ script, `cpuUsageReport.txt` contains the following macro that returns a problem analysis of daily workloads: `EXECUTE ResPmaTotal;`

1. Copy the BTEQ script `cpuUsageReport.txt` to the Windows server at: `InstallFolder\Teradata\Client\16.20\Teradata Notification Service\sql`.
2. Verify the Teradata Notification Service is running on the Windows server.
 - a. At the Windows command prompt, type `services.msc`.
 - b. Check that the status for the Teradata Notification Service is started in the list of services.
 - c. If the service is stopped, right-click the service name and select **Start**.
3. In the **Alert Setup** portlet, configure delivery settings.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Authentication** list, select **Teradata Login**.
 - c. Select ☐ next to **Teradata Login**.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. Select the **Enable login** check box.
 - f. Enter a login name and password.
 - g. [Optional] From the **Test Teradata Logon** list, select **BTEQ (Windows)**.
 - h. [Optional] Click **Test** to verify that the login settings are correct.
 - i. Click **Apply**.
4. In the **Alert Setup** portlet, create an action set.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
Verify the **Evening Hours** displayed in the box correspond to the weeknight hours for this alert action.
 - c. From the **Preset Options** list, select **Action Sets**.
 - d. Select ☐ next to **Action Sets**.
 - e. In **Action Set Name**, type `cpuUsageReport`.

- f. From the **Times** check boxes, select the **Evening** check box, and clear both **Core** and **Weekend** check boxes.
 - g. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - h. From the **Actions** list, select the **BTEQ** check box.
 - i. From the **Location** list, select your Windows server.
 - j. From the **Script** list, select **cpuUsageReport.txt**.
If the script does not appear in the list, verify that the script `cpuUsageReport.txt` exists at:
`InstallFolder\Teradata\Client\16.20\Teradata Notification Service\sql`.
 - k. From the **TDPID** list, select **TEST1**.
 - l. Click **Apply**.
5. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure TEST1.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In **System Nickname**, type TEST1.
 - c. Select the **Enable system** check box to activate the TEST1 system for monitoring.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.
 6. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, click **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, click **Resource Usage**, select the **Enable Stats Manager Collector** check box, and click **Apply**.
 - d. From the **Data Collectors** list, click **System Stats**, click the **Enable Stats Manager Collector** check box, and click **Apply**.
 7. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, click **System**.
 - d. Click  next to **Alerts**.
 - e. In **Alert Name**, type `cpuUsageReportAlert`.
 - f. Select the **Enabled alert** check box.
 - g. In the **Severity** list, select **High**.
 - h. From the **Alert Rules** list, select **Any**.
 - i. From the lists, select **Node CPU Usage** and **is greater than**.
 - j. In the box to the right of the greater-than symbol (>), type 80.
 - k. From the **Alert Action** list, select **cpuUsageReport**.
 - l. Click **Apply**.

8. In the **Alert Viewer** portlet, review triggered actions.
 - a. Select **High** to filter the alerts in the portlet by high severity.
 - b. In the **SYSTEM NAME** filter box, type TEST1 to view alerts for the TEST1 system.

How Can I Submit an SQL Statement as an Alert Action?

The following is an overview of the process:

1. Configure login credentials for the Teradata system in the **Alert Setup** portlet.
2. Configure an SQL query delivery type in the **Alert Setup** portlet.
3. Create an action set in the **Alert Setup** portlet.
4. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
5. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
6. Define alert rules in the **Monitored Systems** portlet.
7. Enable the alert to be generated.
8. Review the SQL output.
9. Review alerts in the **Alert Viewer** portlet.

Example: Configuring an SQL Statement as an Alert Action

Suppose you want to run the following SQL insert statement as an alert action for Teradata system TEST1.

This example applies only to core hours.

```
INSERT INTO admin.sessionHist (TheTime, UserName, SessionCount)
VALUES (CURRENT_TIMESTAMP, ${userName}, ${sessionsForUser})
```

This insert is to be performed when the session count for any user on the TEST1 system exceeds ten.


Notice the alert properties `${userName}` and `${sessionsForUser}` in the insert statement before.


When you configure an alert action to run SQL statements using the **SQL Queries** delivery type, the text can include alert properties that replaces with runtime values. See [Teradata Alert Metrics and Properties](#) for a description of the alert properties available for Teradata Viewpoint alert types.

For this example, assume the target table (admin.sessionHist) has already been created on TEST1. For example:



```
CREATE TABLE admin.sessionHist
(TheTime TIMESTAMP NOT NULL,
UserName VARCHAR(256) NOT NULL,
SessionCount INTEGER NOT NULL)
```


To configure an SQL statement as an alert action for a Teradata system, perform these steps.

1. In the **Alert Setup** portlet, configure the login credentials for TEST1.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Authentication** list, select **Teradata Login**.
 - c. Select  next to **Teradata Login**.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. Select the **Enable login** check box.
 - f. Enter a login name and password.
 - g. From the **Session Character Set** list, select **UTF8**.
 - h. [Optional] From the **Test Teradata Logon** list, select **JDBC**.
 - i. [Optional] Select **Test** to verify the login settings are correct.
 - j. Click **Apply**.

2. In the **Alert Setup** portlet, configure the SQL query for TEST1.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **SQL Queries**.
 - c. Select  next to **SQL Queries**
 - d. In **NAME**, type insertSessionHistory.
 - e. Select the **Enable** check box.
 - f. In **SQL**, type:

```
INSERT INTO admin.sessionHist (TheTime, UserName, SessionCount)
VALUES (CURRENT_TIMESTAMP, ${userName}, ${sessionsForUser})
```

- g. Click **Apply**.
3. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Action Sets**.
 - c. Select  next to **Action Sets**.
 - d. In **Action Set Name**, type runSessionCount.
 - e. From the **Times** list, select the **Core** check box, and clear both the **Evening** and **Weekend** check boxes.
 - f. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - g. From the **Actions** list, select the **SQL** check box.
 - h. From the **Query** list, select **insertSessionHistory**.
 - i. From the **TDPID** list, select **TEST1**.
 - j. Click **Apply**.
4. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure TEST1.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In **System Nickname**, type TEST1.

- c. Select the **Enable system** check box to activate the TEST1 system for monitoring.
- d. In the **TDPID** box, type the TDPID of the TEST1 system.
- e. Enter a login name and password.
- f. Click **Apply**.
5. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select **Sessions**.
 - d. Select the **Enable Sessions Collector** check box and keep the default settings.
 - e. Click **Apply**.
6. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Session**.
 - d. Click  next to **Alerts**.
 - e. In **Alert Name**, type userSessionCountAlert.
 - f. Select the **Enabled alert** check box.
 - g. From the **Severity** list, select **High**.
 - h. From the **Match** list, select **All**.
 - i. From the lists, select **Sessions Per User** and **is greater than**.
 - j. In the empty box after the lists, type the threshold number 10.
 - k. From the **Action** list, select **runSessionCount**.
 - l. Click **Apply**.
7. If no single user on TEST1 has more than ten sessions, then create the conditions for the alert to be raised by opening 11 sessions for a Teradata user.
For example, using BTEQ:

```
.set sessions 11
.logon TEST1/myUser,myPassword
```

8. On the TEST1 system, review the **sessionHist** table. Use a query tool such as BTEQ or SQL Scratchpad to query the table.
For example:

```
select * from admin.sessionHist order by TheTime
```

Note:

You may need to wait a few minutes for the alert to be triggered, depending on the Teradata session monitoring rate on TEST1.

9. In the **Alert Viewer** portlet, review triggered actions.

How Can I Run a Program as an Alert Action?



The following is an overview of the process:

1. Verify the Teradata Notification Service is running.
2. Copy programs or batch files to the server that is running the Teradata Notification Service.
3. Create an action set in the **Alert Setup** portlet.
4. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
5. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
6. Define alert rules in the **Monitored Systems** portlet.
7. Review alerts in the **Alert Viewer** portlet.

Example: Running a Program as an Alert Action

Suppose you want to use the batch file `echoUsage.bat` as an alert action for the Teradata system `PROD1`, weekends only, and track alerts in the **Alert Viewer** portlet. The batch file `echoUsage.bat` opens the program that monitors the performance of your system. The program output includes information about the space usage on `PROD1` if the current perm space that is used in any database on the `PROD1` system exceeds 80%.

1. Verify the Teradata Notification Service is running on the Windows server.
 - a. At the Windows command prompt, type `services.msc`.
 - b. Make sure the status for the Teradata Notification Service is *started* in the list of services.
 - c. If the service is stopped, right-click the service name and select **Start**.
2. Copy the batch file `echoUsage.bat` to the Windows server at `InstallFolder\Teradata\Client\16.20\Teradata Notification Service\usrcmd`.
3. In the **Alert Setup** portlet, create an action set.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
 - c. Verify the **Weekend Hours** displayed in the box correspond to the weekend hours that you want for this alert action.
 - d. From the **Preset Options** list, select **Action Sets**.
 - e. Select ☐ next to **Action Sets**.
 - f. In **Action Set Name**, type `echoUsageReport`.
 - g. From the **Times** check boxes, select the **Weekend** check box, and clear the **Core** and **Evening** check boxes.
 - h. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected. This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - i. From the **Actions** list, select **Run a program**.

- j. From the **Location** list, select your Windows server.
 - k. From the **Program** list, select **echoUsage.bat**.
If the program or script does not appear in the list, verify it exists at *InstallFolder\Teradata*
Client\16.20\Teradata Notification Service\usrcmd\echoUsage.bat.
 - l. Leave the **Arguments** box blank.
 - m. Click **Apply**.
4. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure PROD1.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In **System Nickname**, type Prod1.
 - c. Select the **Enable system** check box to activate the PROD1 system for monitoring.
 - d. In the **TDPID** box, type the **TDPID** of the PROD1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.
 5. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select **Database Space**.
 - d. Select the **Enable Database Space Collector** check box and keep the default settings.
 - e. Click **Apply**.
 6. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Database Space**.
 - d. Select  next to **ALERTS**.
 - e. Enter the **Alert Name** echoUsageReportAlert.
 - f. Select the **Enable alert** check box.
 - g. In the **Severity** list, select **Medium**.
 - h. From the **Threshold** list, select **Current Perm**.
 - i. In the box to the right of the greater-than symbol (>), type 80.
 - j. In the **Action** list, select **echoUsageReport**.
 - k. Leave the value unchanged in the **Do not run twice in ... minutes** box.
 - l. [Optional] Enter a **Message**.
 - m. Click **Apply**.
 7. In the **Alert Viewer** portlet, review triggered actions.
 - a. [Optional] Select **Medium** to filter the alerts in the portlet by normal severity.
 - b. [Optional] In the **SYSTEM NAME** filter box, type PROD1 to view alerts for the PROD1 system.

Note:

The alert or alert detail information the **Alert Viewer** portlet displays does not include the output of the program or batch file.

How Can I Add Alert Properties as Arguments to the Run a Program Alert Action?

This example explores how to add alert properties as arguments to the **Run a Program** action. You can define a threshold for a database metric when a specific event occurs. You can then specify that the event triggers an alert. The alert service has access to information on several properties of a monitored system, depending on the data collectors that you enable. If you configure the alert to run a program and the program accepts command-line parameters, you can add custom parameters that the program can access.

The following is an overview of the process:


1. Verify the Teradata Notification Service is running.
2. Copy programs or batch files to the server that is running the Teradata Notification Service.
3. Create an action set that adds alert properties as arguments under the **Run a program** action in the **Alert Setup** portlet.
4. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
5. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
6. Define alert rules in the **Monitored Systems** portlet.
7. Review alerts in the **Alert Viewer** portlet.

Example: Adding Alert Properties as Arguments to Run a Program

Suppose you want to use the script `echoUsage.sh` as a medium severity alert action for the Teradata system `PROD1`, weekends only, and track alerts in the **Alert Viewer** portlet. `echoUsage.sh` opens the program that monitors the performance of your system. The script reads in any arguments that appear on the command-line and runs the program using these arguments as run-time parameters. The program output includes information about the space usage on `PROD1` if the current perm space that is used in any database on the `PROD1` system exceeds 80%.

1. Verify the Teradata Notification Service is running on the Linux server:
 - a. Use the following Linux command to validate that the service is running correctly:
`/etc/init.d/tdnotification status`
 - b. If the result is not "running," use: `/etc/init.d/tdnotification start`
2. Copy the script `echoUsage.sh` to the Linux server at : `/opt/teradata/client/cam/tdnotification/usrcmd/echoUsage.sh`.
 - a. Make sure the script `echoUsage.sh` is executable by the `tdnotification` service user account.

3. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
 - c. Verify the **Weekend Hours** displayed correspond to the weekend hours that you want for this alert action.
 - d. From the **Preset Options** list, select **Action Sets**.
 - e. Select ☐ next to **Action Sets**.
 - f. In **Action Set Name**, type `echoUsageSpaceUsedPct`.
 - g. From the **Times** check boxes, select **Weekend**, and clear the **Core** and **Evening** check boxes.
 - h. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - i. From the **Actions** list, select **Run a program**.
 - j. From the **Location** list, select your Linux server.
 - k. From the **Program** list, select **echoUsage.sh**.
If the script does not appear in the list, verify that it exists at `/opt/teradata/client/cam/tdnotification/usrcmd/echoUsage.sh`.
 - l. In the **Arguments** box, type `-db ${databaseName} -sp ${spaceUsedPct}`.
 - m. Click **Apply**.
4. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure PROD1.
 - a. Next to **Systems**, select ☐, then select **Add Teradata System**.
 - b. In **System Nickname**, type `PROD1`
 - c. Select the **Enable system** check box to activate the PROD1 system for monitoring.
 - d. In the **TDPID** box, type the **TDPID** of the PROD1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.
5. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select **Database Space**.
 - d. Select the **Enable Database Space Collector** check box and keep the default settings.
 - e. Click **Apply**.
6. In the **Monitored Systems** portlet, define alert rules.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Database Space**.
 - d. Select ☐ next to **Alerts**.
 - e. In **Alert Name**, type `echoPermSpaceLow`.
 - f. Select the **Enable alert** check box.

- g. From the **Severity** list, select **Medium**.
 - h. From the **Alert Rules** list, select **Current Perm**.
 - i. In the box to the right of the greater-than symbol (>), type 80.
 - j. Select  then select the **All databases except the selected databases below** option and include all databases.
 - k. From the **Action** list, select **echoUsageSpaceUsedPct**.
 - l. Leave the value unchanged in the **Do not run twice in** minutes box.
 - m. Click **Apply**.
7. In the **Alert Viewer** portlet, review triggered actions.
- a. [Optional] Click **Medium** to filter the alerts in the portlet by normal severity.
 - b. [Optional] In the **SYSTEM NAME** filter box, type PROD1 to view alerts for the PROD1 system.

Note:

The alert or alert detail information that the **Alert Viewer** portlet displays does not include the output of the program or batch file.

How Can I Check the Logon Timeout with Canary Query?

This example explores how to create an alert action that is based on an expired timeout period when connecting to a Teradata system.

The following is an overview of the process:

1. Configure email delivery options in the **Alert Setup** portlet.
2. Create an action set in the **Alert Setup** portlet.
3. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
4. Define a canary query in the **Monitored Systems** portlet to run on the Teradata system.
5. Define alert rules in the **Monitored Systems** portlet.
6. Review alerts in the **Alert Viewer** portlet.
7. Review alert messages that arrive in the email inbox.

Example: Checking the Canary Query Logon Timeout

Suppose you want to monitor when locks that occur on data-dictionary tables prevent you from logging on to Teradata system PROD1. You create an alert action that is based on an expired logon-timeout period using a canary query. Each time the query runs, it attempts to log in to PROD1 and then executes a simple SQL statement. If the logon fails, it times out and emails an alert to dba@example.com and logs the alert in the **Alert Viewer** portlet.

1. In the **Alert Setup** portlet, configure delivery settings.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **Email**.

- c. In the **SMTP Host** box, type `smtp.example.com` as the SMTP host address of the outgoing (SMTP) email server.
 - d. Select the **Enable email** check box.
 - e. From the **Port** list, select **Use default port** to use the default port 25.
 - f. Enter a default **Server Timeout** of 30 seconds.
 - g. In the **Reply-To** box, type `doNotReply@example.com`.
 - h. Keep the default **Anonymous login**.
 - i. Leave the **Advanced** box blank.
 - j. In the **Test Recipient** box, type `dba@example.com`.
 - k. Select **Test** to verify the SMTP server delivers the email to `dba@example.com`.
 - l. Click **Apply**.
2. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Action Sets**.
 - c. Select ☐ next to **Action Sets**.
 - d. In **Action Set Name**, type `Email dba`.
 - e. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - f. From the **Actions** list, select **Email recipients**.
 - g. Select the **Email recipients** check box, click **Bcc** or **To** , and in the **Email** box, type `dba@example.com`.
 - h. Click **Apply**.
3. In the **Monitored Systems** portlet for Teradata Viewpoint monitoring, configure PROD1.
 - a. Next to **Systems**, select ☐ , then select **Add Teradata System**.
 - b. In **SYSTEM NICKNAME**, type `PROD1`.
 - c. Select the **Enable system** check box to activate the PROD1 system for monitoring.
 - d. In the **TDPID** box, type the **TDPID** of the PROD1 system.
 - e. Enter a login name and password.
 - f. Click **Apply**.
4. In the **Monitored Systems** portlet, define a canary query to run on the Teradata system.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Canary Queries**.
 - c. Select ☐ next to **Canary Queries**.
 - d. In **Name**, type `LogonTimeout`.
 - e. Select the **Enable query** check box.
 - f. In the **SQL** box, enter `SELECT COUNT(*) FROM DBC.DBCInfo;`
 - g. In **Default Database**, type `DBC`.
 - h. From the **Login** list, select the login credentials you selected for system tasks.

- i. Select **Log in each time the query executes**.
 - j. In the **Time out the login attempt after box**, type 20.
This is a timeout period for the logon.
 - k. Leave the other default values unchanged.
 - l. Click **Apply**.
5. In the **Monitored Systems** portlet, define a canary query alert rule.
 - a. From the **Systems** list, select **PROD1**.
 - b. From the **Setup** list, select **Alerts**.
 - c. From the **Alert Types** list, select **Canary Queries**.
 - d. Select ☐ next to **Alerts**.
 - e. In **Alert Name**, type Canary Query Logon Timeout.
 - f. Select the **Enable alert** check box.
 - g. In the **Severity** list, select **Medium**.
 - h. In the **Match** list, select **All**.
 - i. From the metric lists, select **LogonTimeout**, and **Login Timeout Occurred**.
 - j. For **Only trigger if alert rule(s) are met for**, leave 0 as the default number of minutes.
 - k. In the **Action** list, select **Email dba**.
 - l. For the **Do not run twice in box**, leave 0 as the default number of minutes.
 - m. In the **Message** box, type Logon Timeout occurred on PROD1.
 - n. Click **Apply**.
6. In the **Alert Viewer** portlet, review triggered actions.
 - a. [Optional] Select **Medium** to filter the alerts in the portlet by normal severity.
 - b. [Optional] In the **SYSTEM NAME** filter box, type PROD1 to view alerts for the PROD1 system.
7. In the email inbox for the account dba@example.com, check for messages from doNotReply@example.com.

Sample Email Message

The email message contains the following formatted information when it arrives in the inbox.

```

From: doNotReply@example.com
Sent: Tuesday, August 7, 2012 2:23 PM
To: dba@example.com
Subject: [Alert] PROD1 - Canary Query Logon Timeout (Source: Viewpoint, Type:
Canary Query)

Logon timeout occurred on PROD1.

Event Timestamp: 2012-08-07T14:23:01.331-07:00

Canary Id=90
  
```

```
Query Name=LogonTimeout
Response Time=
Login Timeout Occurred=true
```

```
Description: ((Canary Id = 90 and Login Timeout Occurred is true))
```

How Can I Define an Exception in the Workload Designer Portlet to Send an Alert?

This example makes the following assumptions:


- The example is implemented first in a test environment. Because only one ruleset can be active in Workload Designer at a given time, the ruleset you create with this procedure replaces the currently active ruleset. Applying an untested ruleset to a production system can degrade performance.
- The example only generates alerts for queries that are classified under the **WD-Default** workload for the **ALWAYS** planned environment.

The following is an overview of the process:





1. Configure the Teradata system for Teradata Viewpoint to monitor in the **Monitored Systems** portlet.
2. Configure data collectors to monitor the configured Teradata system in the **Monitored Systems** portlet.
3. Configure email delivery options in the **Alert Setup** portlet.
4. Create an action set in the **Alert Setup** portlet.
5. Create a ruleset in the **Workload Designer** portlet.
6. Activate the ruleset in the **Workload Designer** portlet.
7. Review alert messages that arrive in the email inbox.

Example: Defining an Exception to Send an Alert

Suppose you want to send an email alert to dba@example.com during core hours when a period event begins and ends.


1. In the **Monitored Systems** portlet, configure TEST1 as a Teradata Active System Management (TASM) system.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In **System Nickname**, type TEST1.
 - c. Select the **Enable system** check box.
 - d. In the **TDPID** box, type the TDPID of the TEST1 system.
 - e. [Optional] Enter the site ID assigned by Teradata Customer Services.
 - f. Enter a login name and password.
 - g. Click **Apply**.
2. In the **Monitored Systems** portlet, configure the data collectors.

- a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select the following collectors, enable them, and select **Apply** for each:
 - TASM Config
 - TASM Distribution
 - TASM Exception
 - TASM State
 - TASM Summary
3. In the **Alert Setup** portlet, configure delivery settings.
- a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **Email**.
 - c. In the **SMTP Host** box, type `smtp.example.com` as the SMTP host address of the outgoing (SMTP) email server.
 - d. Select the **Enable email** check box.
 - e. From the **Port** list, select **Use default port** to use the default port 25.
 - f. In **Server Timeout**, type 30 seconds.
 - g. Keep the default **Anonymous login**.
 - h. In the **Reply-To** box, type `doNotReply@example.com`.
 - i. Leave the **Advanced** box blank.
 - j. In the **Test Recipient** box, type `dba@example.com`.
 - k. Select **Test** to verify the SMTP server delivers the email to `dba@example.com`.
 - l. Click **Apply**.
4. In the **Alert Setup** portlet, create an **Action Set**.
- a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
 - c. Verify the core hours are correctly set for your environment.
 - d. From the **Preset Options** list, select **Action Sets**.
 - e. Select ☐ next to **Action Sets**.
 - f. Type an **Action Set Name**, such as `CoreHours`.
 - g. From the **Times** list, select the **Core** check box, and clear both the **Evening** and **Weekend** check boxes.
 - h. Under **Actions**, make sure the **Include in Alert Viewer** check box is selected.
This option is selected by default so you can view alerts for this action set in the **Alert Viewer** portlet.
 - i. From the **Actions** list, select the **Email recipients** check box.
 - j. Select the **Email recipients** check box, click **Bcc** or **To** , and in the **Email** box, type `dba@example.com`.
 - k. Click **Apply**.





5. In the **Workload Designer** portlet, create a period event that triggers the alert email.
 - a. From the **For System** list, select **TEST1**.
 - b. Select .
 - c. From the **General** tab, in the **Name** box, type **TEST1_alert_email**.
 - d. Click **Save**.
 - e. From the ruleset toolbar, click **States**.
 - f. Hover over **Planned Events**, and click .
 - g. Select  next to **Available Events**.
 - h. Select **Period Event**.
 - i. In **Event Name**, type **CoreEnvironment**.
 - j. Select **Create New Corresponding Planned Environment** to automatically create a planned environment that this event triggers.
 - k. Select **Day of Week**, and select all weekdays.
 - l. Select **Start Time**, with a 9 am **Start Time** and 5 pm **End Time**.
 - m. In the **Send Alert** row under **Configure Notifications**, click **Start Action** and **End Action**, and select **CoreHours** from the list.
 - n. Click **OK**.
 - o. Select **Close**.
 - p. Click **Save**.
6. In the **Workload Designer** portlet, activate the ruleset.
 - a. Under **Working**, select the **TEST1_alert_email** .
 - b. Select **Make Active**.
 - c. In the **Confirm Activation Request** dialog box, select **Activate**.
7. In the email inbox for the account **dba@example.com**, check for messages from **doNotReply@example.com**.

Example: Defining an Exception to Send an Alert to Database User on Abort/Abort Selects

Suppose you want to send an email alert to the dbauser having the e-mail id **dba@example.com**, during core hours when a period event begins and ends.

1. In the **Monitored Systems** portlet, configure **TEST1** as a Teradata Active System Management (TASM) system.
 - a. Next to **Systems**, select , then select **Add Teradata System**.
 - b. In **System Nickname**, type **TEST1**.
 - c. Select the **Enable system** check box.
 - d. In the **TDPID** box, type the TDPID of the **TEST1** system.
 - e. [Optional] Enter the site ID assigned by Teradata Customer Services.
 - f. Enter a login name and password.

- g. Click **Apply**.
2. In the **Monitored Systems** portlet, configure the data collectors.
 - a. From the **Systems** list, select **TEST1**.
 - b. From the **Setup** list, select **Data Collectors**.
 - c. From the **Data Collectors** list, select the following collectors, enable them, and select **Apply** for each:
 - TASM Config
 - TASM Distribution
 - TASM Exception
 - TASM State
 - TASM Summary
 - Alert Request
3. In the **Alert Setup** portlet, configure delivery settings.
 - a. From the **Setup Options** list, select **Delivery Settings**.
 - b. From the **Delivery Types** list, select **Email**.
 - c. In the **SMTP Host** box, type `smtp.example.com` as the SMTP host address of the outgoing (SMTP) email server.
 - d. Select the **Enable email** check box.
 - e. From the **Port** list, select **Use default port** to use the default port 25.
 - f. In **Server Timeout**, type 30 seconds.
 - g. Keep the default **Anonymous login**.
 - h. In the **Reply-To** box, type `doNotReply@example.com`.
 - i. Leave the **Advanced** box blank.
 - j. In the **Test Recipient** box, type `dba@example.com`.
 - k. Select **Test** to verify the SMTP server delivers the email to `dba@example.com`.
 - l. Click **Apply**.
4. In the **Alert Setup** portlet, create an **Action Set**.
 - a. From the **Setup Options** list, select **Alert Presets**.
 - b. From the **Preset Options** list, select **Core Hours**.
 - c. Verify the core hours are correctly set for your environment.
 - d. Under **Database User Alerts**, select **User Mapping**.
 - e. Select the system.
 - f. Enter the username and password.
 - g. Enter the table or view name as configured in the selected Teradata Viewpoint system. Refer [Database User Alerts](#) for details.
 - h. Verify the test.
 - i. Click **Apply**.
5. In the **Workload Designer** portlet, create a period event that triggers the alert email.
 - a. From the **For System** list, select **TEST1**.


- b. Select .
 - c. From the **General** tab, in the **Name** box, type TEST1_alert_email.
 - d. Click **Save**.
 - e. From the ruleset toolbar, click **States**.
 - f. Hover over **Planned Events**, and click .
 - g. Select  next to **Available Events**.
 - h. Select **Period Event**.
 - i. In **Event Name**, type CoreEnvironment.
 - j. Select **Create New Corresponding Planned Environment** to automatically create a planned environment that this event triggers.
 - k. Select **Day of Week**, and select all weekdays.
 - l. Select **Start Time**, with a 9 am **Start Time** and 5 pm **End Time**.
 - m. In the **Exceptions**, after configuring the exception criteria, select only **Abort/Abort Select** from **Actions**.
 - n. Select **Notify Database User from Notifications**.
 - o. Click **OK**.
 - p. Select **Close**.
 - q. Click **Save**.
6. In the **Workload Designer** portlet, activate the ruleset.
 - a. Under **Working**, select the **TEST1_alert_email** .
 - b. Select **Make Active**.
 - c. In the **Confirm Activation Request** dialog box, select **Activate**.
 7. In the email inbox for the account dba@example.com, check for messages from doNotReply@example.com.
The alert collection upon TDWM abort happens from Alert Request Collector rate.

Backup

The **Backup** portlet allows you to enable and disable automatic Teradata Viewpoint system backups and to select the daily backup time as well. You can select whether to store the backup files on the local system or at a network location and set the number of days the backups are stored. The Teradata Alerts (*cam*) database is backed up at the same time.

The **Backup** portlet displays the backup process icon with either a successful completion or an error for the last attempt and lists the last successful backup date and time.

Backup Errors

The **Backup** dialog box displays  if an error occurs during the backup process. To investigate the error, you can view the backup log file: /opt/teradata/viewpoint/dcs/logs/backup.log.




Scheduling Viewpoint Backups


From the **Backup** view, a status displays. If the last backup was successful, the date, time, and size (MBs) displays.

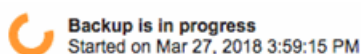
1. From the **Backup** view, select one of the following:

Option	Description
Enable backup	Create a daily backup. [Recommended]
Backup Now	Create an immediate backup in the defined location. Use only when needed.

2. [Optional] From the **Daily Backup Time** list, select a start time for a daily backup.
3. From the **Keep Backups For** list, select the number of days to retain the backup.
4. Under **Location**, select one of the following options for backup files:

Option	Description
Local	Backs up Teradata Viewpoint databases to the local server.
Network	<p>Backs up Teradata Viewpoint databases to a networked file server to prevent data loss during catastrophic failure.</p> <ol style="list-style-type: none"> a. If a postgres user has not been created on the NFS server with write access to the specified backup path, contact the Teradata System Administrator. b. Enter a Host to specify the host system name. The host is the IP address or DNS name of the networked file server. c. Enter a Path to specify the absolute path name where the backups are stored on the host system. For example, to run the command on Linux: <code>mount <host>:<path>/path/to/local/mount/point</code> d. [Optional] Click Test to verify that the login settings are correct. If the operation is successful,  appears. If the operation fails,  appears. Verify the settings are correct and try again. <p>Note: Select  to view the result of the test.</p>



5. Click **Apply**.
6. [Optional] To update the status of the backup, select  next to **Status**.
The throbber indicates the status of the backup in progress.



Configuring Alert for Failed Backup Processes


You can enable notification for the failed backup processes.

1. Select the **Enable Alert** check-box from the **Configure Alert** section of the **Backup** portlet.
2. Select any of the following from the **Severity** drop-down:
 - **Critical**
 - **High**
 - **Medium**
 - **Low**
3. Select an action from the **Action** drop-down.
The action sets are defined in the **Alert Setup** portlet.
4. Enter a message for the alert (Optional).
5. Click **Apply**.


Backup
 Configure the backup of Teradata Viewpoint server data.
 

Details

Status


Last Configuration Only Backup Successful
 Feb 6, 2022 12:02:44 AM (9.85MB)

BACKUP NOW

Setup

☒ Enable backup
☒ Full backup ☐ Configuration only backup
 Daily Backup Time: 12:00 AM ▾
 Keep Backups For: 7 ▾ days

Backup Location

☒ Local
☐ Network

Configure Alert

☒ Enable Alert
 Alert Name: Backup Job Alert
 Severity: Critical ▾
 Action: BackupTestAlert ▾
 Message: This is a backup test alert

APPLY RESET

Certificates

Use the **Certificates** portlet to manage certificates on the Viewpoint server. Certificates are used to help secure communications to and from a Viewpoint server. You can install a certificate to secure communication from a browser to the Viewpoint server and from the Viewpoint server to external servers.

You can do any of the following:

- Install one or more trusted root certificates that are issued by a trusted certificate authority (CA).
- Create a certificate signing request (CSR) and install it after it is signed by a CA.
- Create and install a self-signed certificate (SSC).
- Install a PKCS #12 file if you want to install a certificate without having to first create a CSR.

The installed certificate applies to only this Viewpoint server.

Certificate Authorities

A certificate authority (CA) issues certificates. To initiate secure communications with an external server (for example, a corporate LDAP server), the Viewpoint server must trust the CA that issued the certificate of the external server.

The root certificates of certain standard certificate authorities are pre-installed on the Viewpoint server and certificates issued by them are automatically accepted. However, if a non-standard CA issued the certificate, the root certificate of this CA must be installed on the Viewpoint server for the certificate to be accepted.

Installing Trusted Root Certificates

You can install root certificates that are issued by a trusted CA from either of the following:

- A file with a PEM or DER encoded certificate
 - A trusted SSL-enabled service
1. From the **Setup** list, select **Certificate Authority**.
 2. Click **Install Certificate**.
 3. Type an alias for the certificate, up to 30 characters.
 4. Select an option and do one of the following:

Option	Description
A file with a PEM or DER encoded certificate	Select Choose File and browse to select a certificate on your system.
A trusted SSL-enabled service	<ol style="list-style-type: none"> a. Type a host name. b. Type a port number.

5. Select **Install**.

The alias, expiration date, and authority of the new certificate appears under **Trusted Certificate Authorities** in the **Certificate Authority** view.

Deleting Trusted Root Certificates

1. From the **Setup** list, select **Certificate Authority**.
2. Select ☐ next to the certificate you want to delete then select **Delete**.
A confirmation message appears.

HTTPS Connections

HTTPS connections are secure connections between a browser and the Viewpoint server. To enable HTTPS connections, the Viewpoint server needs a certificate so that browsers can verify that passwords or other sensitive information is not intercepted or sent to an untrusted source. Starting with Viewpoint 16.20.00.03, a self-signed HTTPS certificate is created if a certificate does not exist.

After a certificate is installed, HTTPS connections immediately start using the new certificate and no restart of the Viewpoint server is required.

Installing a Certificate from a Certificate Signing Request

A certificate signing request (CSR) is used to obtain a signed certificate for the server. After the request is created, the certificate signing request must be submitted to a CA to be signed. After the request is signed, the resulting signed certificate can be installed.

Note:

For Viewpoint 16.20.00.03 and later, a self-signed certificate is visible after you click **HTTPs Config**.

1. From the **Setup** list, select **HTTPS Config**.
2. Under **Certificate Signing Request**, do the following under **Step 1**:
 - a. Select **Create File**.
 - b. Enter the following information:


Field	Description
Common Name	Domain name
Subject Alternative Name	[Optional] IP address or DNS name. Select <input type="checkbox"/> to see the required format.
Organizational Unit	Business unit
Organization	Company name
City or Locality	City or locality of the organization
State or Province	State or province of the organization

Field	Description
Country	Country code where the organization is located
Email	[Optional] Email address of the requester

- c. Click **Create**.
- d. If a certificate was previously installed, select **Replace**.

A csr file is created and the file name appears as a link under **Step 2**.

- 3. Under **Step 2**, do one of the following:

Option	Description
Copy	<ul style="list-style-type: none"> a. Select the link. b. View the contents of the CSR in the box or copy it to the Clipboard for further use. c. Select Close.
View CSR contents before downloading	<ul style="list-style-type: none"> a. Select the link. b. Select Download. <p>The file is saved to your download area or to a location that you specify, depending on the browser settings.</p>
Download without viewing CSR contents	<ul style="list-style-type: none"> a. Select . <p>The file is saved to your browser's download area.</p>

- 4. Outside of Teradata Viewpoint, do the following:
 - a. Send the CSR to a CA for signing.
 - b. After the request is signed by the CA, get the certificate bundle to your local system and extract its contents, namely ServerCertificate.crt, Intermediate.crt and Root.crt.
 - c. Generate a PKCS#7 formatted file through the command "openssl crl2pkcs7 -nocrl -certfile /tmp/ServerCertificate.crt -certfile /tmp/Intermediate.crt -certfile /tmp/Root.crt -out /tmp/final.p7b" using ServerCertificate.crt, Intermediate.crt and Root.crt obtained from the CA
- 5. Under **Step 3**, do the following:
 - a. Select **Install**.
 - b. In the **Certificate** box, select **Choose File** and browse to the saved final.p7b file obtained from step 4.c
- 6. Select **Install**.

The name of the new certificate appears in the **Certificate** box.

HTTPS connections immediately start using the new certificate and no restart of the Viewpoint server is required.

Installing a PKCS #12 File

You can install a PKCS #12-formatted private key and public key certificate chain. The benefit of PKCS #12 is that you can install a certificate without having to first create a CSR. Depending on the particular PKCS #12 file you are installing, a password may or may not be required.


1. From the **Setup** list, select **HTTPS Config**.
2. Under **Install PKCS #12 File**, select **Install**.
3. From the **Certificate** box, select **Choose File**.
4. If required, enter a password.
5. Select **Install**.
 - a. If a certificate was previously installed, select **Replace**.

The name of the new certificate appears in the **Certificate** box.

Creating and Installing a Self-Signed Certificate

A self-signed certificate (SSC) can provide security for an internal site or server as a temporary measure while a fully authenticated certificate is obtained from a CA. Installing a certificate replaces the existing certificate.

1. From the **Setup** list, select **HTTPS Config**.
2. Under **Self-Signed Certificate**, do the following:
 - a. Select **Create and Install**.
 - b. Enter the following information:

Field	Description
Common Name	Domain name
Subject Alternative Name	[Optional] IP address or DNS name. Select  to see the required format.
Organizational Unit	Business unit
Organization	Company name
City or Locality	City or locality of the organization
State or Province	State or province of the organization
Country	Country code where the organization is located
Email	[Optional] Email address of the requester
Expiration	Select an expiration period

- c. Select **Create and Install**.


- d. If a certificate was previously installed, select **Replace**.

The name of the new certificate appears in the **Certificate** box.

HTTPS connections immediately start using the new certificate and no restart of the Viewpoint server is required.

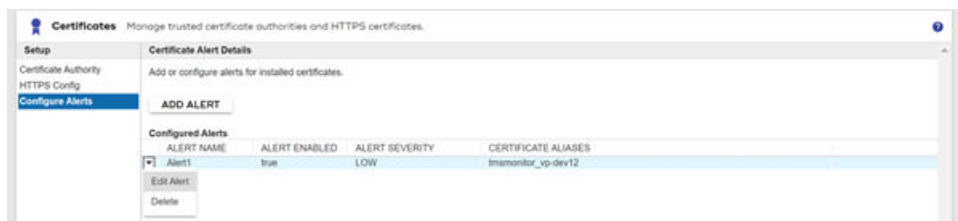
Deleting an Installed Certificate

If you delete an installed certificate, HTTPS is disabled and the less secure HTTP becomes the default protocol.

1. From the **Setup** list, select **HTTPS Config**.
 2. From the **Certificate** box, select .
- A confirmation message appears.

Certificate Expiry Alert Configuration

You can configure alerts to get prior notifications about the upcoming expiry date of the certificates in the Teradata Viewpoint key store. The **Configure Alerts** maintains a list of certificate expiry alerts. You can add new alerts, and edit or delete existing alerts. The **Configure Alerts** shows the alert name, status, severity, and the certificate name for which you have configured the alert.



By default, the Viewpoint system checks the enabled alerts everyday at 12 AM. If the set criteria of any of the alert satisfies, a notification triggers.

Note:

In the `/opt/teradata/viewpoint/dcs/config/dcs.properties`, the **certalert.enabled** is to enable or disable the system check for expiry alert, and the **certalert.schedule** is to modify the system check timing and frequency. Only a server administrator can modify the values of these parameters.

Creating an Expiry Alert

To create an expiry alert to **Configure Alerts**:

1. From the **Setup** list, select **Configure Alerts**.
 2. Under **Certificate Alert Details**, click **ADD ALERT** button.
- Add Alert** window opens.

3. In the **Add Alert** enter the details as following:

- Enter an alert name.
 - Set the alert severity to low, medium or high.
 - Select a certificate for which you are creating the alert.
 - Set a time duration before the certificate expiry date, from which the notification will start triggering.
 - Select an action for the alert, that occurs when the alert activates.
 - Add any message for the alert, if required.
4. Click **SAVE**.
The new alert appears in the **Certificate Alert Details** page.

Editing an Expiry Alert

To edit an existing expiry alert from the **Configure Alerts**:

1. Click the drop-down menu beside the alert.
2. Select **Edit Alert**.
3. Update the alert settings.
4. Click **SAVE**.

The updated alert appears in the **Certificate Alert Details** page.

Deleting an Expiry Alert

To delete an existing expiry alert from the **Configure Alerts**:

1. Click the drop-down menu beside the alert.
2. Select **Delete**.

General

The **General** portlet allows a Teradata Viewpoint Administrator to set certain default settings for a Viewpoint server.

Locale

The locale is used to display content in the Viewpoint portal. Dates, times, and numbers are displayed in Viewpoint portlets in the format used by the selected locale. For example, the **English (US)** locale displays content in English and dates in a *mm/dd/yy* format. If the language pack for the selected locale is installed, all content appears in the selected language. If the language pack for the selected locale is not installed, content appears in English and dates, times, and numbers appear in the format for the selected locale.

Users can choose a different locale in the **Profile** portlet to override the default locale selection.

Session Timeout

Length of time a user account can remain idle before being logged out of the system. The default setting is 1 hour and 30 minutes.

Allow multiple sessions per login

Allows a user to open two or more sessions using the same user name.

Enable newly created shared portlets

Shared portlets are automatically enabled upon creation.

Require access via HTTPS

Make sure Teradata Viewpoint is accessed over a secure connection.

Allow autofill of username and passwords

Allows the username and password fields to have their values automatically completed by the browser. This setting applies where users are prompted for login credentials to Teradata Viewpoint, as well as to the systems being monitored by Viewpoint.

Display a message on the login screen

Create and display a message at the bottom of the Teradata Viewpoint portal login page.

Administrator Alerts

Select an alert action for Viewpoint-generated alerts. **Administrator Default** is selected by default and sends an email to all users in the Administrator role. Instead of using the **Administrator Default** setting, you can select a custom action set that was created in the **Alert Setup** portlet.

Setting General Options

You can set the default locale, set the session timeout, set multiple sessions per login, enable newly created shared portlets, set to require users to use a secure connection when accessing Viewpoint, set to automatically fill login credentials, create and display a message when users log in, and set administrator alerts.

1. For **Locale**, select a locale from the list.
The default locale in Teradata Viewpoint is **English (US)**. Changing the default locale affects all Teradata Viewpoint users who have not explicitly set a locale in their user profile.
2. For **Session Timeout**, select **hours** and **minutes**.
3. Select the **Allow multiple sessions per login** check box to enable a user to open two or more sessions using the same user name.
4. Select the **Enable newly created shared portlets** check box if you want shared portlets to be automatically enabled when they are created.
5. Select the **Require access via HTTPS** check box to require users to use a secure connection when accessing Viewpoint.
Starting with Viewpoint 16.20.00.03, a self-signed certificate that enables HTTPS access is generated if a certificate is not found. Additionally, new installations have **Require access via HTTPS** enabled by default. The **Require access via HTTPS** setting is unchanged for existing installations. This check box is enabled only when accessing Viewpoint using a secure connection.
6. Select or clear the **Allow autofill of username and passwords** check box to control the automatic completion of login credentials by the browser:
 - Select the check box to allow the login credentials to have their values automatically completed by the browser.
 - Clear the check box to prohibit the login credentials from having their values automatically completed by the browser.
7. Select or clear the **Display a message on the login screen** check box to display or hide a customized message at the bottom of the Teradata Viewpoint portal login page.
 - Select the check box and type a message in the box that displays on the login page.
 - Clear the check box to hide the message on the login page.

To test that the message displays at the bottom of the login page of the Viewpoint portal, you can log out and log back in to Viewpoint.
8. Under **Administrator Alerts**, select an alert action from the **Action** list.
Administrator Default is selected by default and sends an email to all users in the Administrator role.
Instead of using the **Administrator Default** setting, you can select a custom action set that was created in the **Alert Setup** portlet.
9. Click **Apply**.

LDAP Servers

The **LDAP Servers** portlet allows you to configure LDAP (Lightweight Directory Access Protocol) servers to authenticate users and assign user roles in Teradata Viewpoint. The **LDAP Servers** portlet displays the names of currently configured LDAP servers or is empty if an LDAP server is not configured.






Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), the names of currently configured LDAP servers are not listed, and the **Details** pane displays the **Auto-Provisioning** and **Role Mapping** sections only.

Adding an LDAP Server

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you cannot add LDAP servers.

1. Select  next to **Servers**.
2. Enter a server nickname for the LDAP server, up to 8 characters.
3. [Optional] Select the **Enable server** check box to activate LDAP upon successfully adding the LDAP server to Teradata Viewpoint.
4. Enter the URL address of at least one LDAP server.
5. [Optional] Use     to delete, add, and change the order of multiple URLs.
The order of the URLs determines the order in which LDAP servers are accessed for authentication.
6. From **Name Matching**, choose a method from the following options:
 - Select **DN Pattern Bind** and specify one or more patterns that represent a distinguished name (DN) in the LDAP directory.

Example: CN={0},OU=User Accounts,DC=td,DC=acme,DC=com, where {0} is replaced by the username the user enters when logging in to Teradata Viewpoint.
 - Select **User Search**.
7. If you selected **User Search**, configure the search options:
 - a. [Optional] To search using a specific account, enter a valid service account DN and password.
If these fields are empty, the Viewpoint portal connects to the LDAP server anonymously.
 - b. Enter a search pattern.
Example: CN={0}, where {0} is replaced by the username the user enters when logging in to Teradata Viewpoint.
 - c. Enter a search base to specify a base directory for the search pattern.
Example: OU=User Accounts,DC=td,DC=acme,DC=com
 - d. [Optional] For **Search Extent**, select the **Recursive scan** check box to include subdirectories of the base directory in the search.



If the check box is not selected, the search is limited to the specified base directory.

8. Under **Key User Information**, enter the LDAP attribute names for the user attributes:

Attribute	Example User Attribute Name
LDAP First Name Attribute	givenName
LDAP Last Name Attribute	surname
LDAP Email Attribute	mail

9. [Optional] Under **Settings Test** you can test whether your settings are correct.

- Enter a username and password.
- Select **Test**.

If the operation is successful,  appears. If the operation fails,  appears.

If the test is successful, you receive the full DN and key user information for the user, and if the test fails, you receive the error message returned from the LDAP server.

10. Select **Apply**.

The **LDAP Servers** view remains open if you want to use the **Advanced Options**.

Auto-Provisioning and Role Mapping

In Teradata Viewpoint, users are assigned roles to organize and control their access to portlets, metrics, and features. You can use the **Roles Manager** portlet to assign permissions to the roles. When configuring an authentication server for Teradata Viewpoint, you can automatically assign Viewpoint roles through auto-provisioning or role mapping.

Auto-provisioning assigns the same roles to all users the first time they access Teradata Viewpoint. You cannot use auto-provisioning to add or remove roles on subsequent logins.

Role mapping assigns roles to each user based on the user's authentication server group membership or LDAP attributes. Use role mapping to do the following:








- Assign roles for a new user
- Add roles that were assigned to the user since the last user login
- Remove roles that were removed since the last user login

Only roles already defined using the **Roles Manager** portlet can be mapped. The Teradata Viewpoint Administrator can change a role at any time using the **Roles Manager** portlet.







Configuring Auto-Provisioning (LDAP)

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), see [Configuring Auto-Provisioning \(JWT\)](#).

1. Select  next to **Advanced Options**.
2. Select the **Turn on auto-provisioning** check box to automatically add a Viewpoint account for an LDAP user the first time that user successfully logs in to Teradata Viewpoint.
3. Select a role or roles from the **Automatically assign these roles** list.
4. [Optional] Use     to delete, add, and change the order of roles.
The order of roles affects priority, which in turn can affect the user's settings in portlets.
5. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears.


Configuring Auto-Provisioning (JWT)







1. Select a role or roles from the **Automatically assign these roles** list.
2. [Optional] Use     to delete, add, and change the order of roles.
The order of roles affects priority, which in turn can affect the user's settings in portlets.
3. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears.

Configuring Role Mapping (LDAP)

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), see [Configuring Role Mapping \(JWT\)](#).

1. Select  next to **Advanced Options**.
2. Configure **Role Mapping** by selecting **Add Role Mapping** at the bottom of the dialog box.
3. Select one of the following role mapping methods from the **Type** list:
 - **Attribute**
 - **Group**
4. If you selected the **Attribute** method of role mapping, do the following:
 - a. Enter an **Attribute Name**.
 - b. Enter an **LDAP Value**.
 - c. Select a **Viewpoint Role** from the list.
5. If you selected the **Group** method of role mapping, do the following:
 - a. For **Group Search Base**, enter a value.
 - b. For **Group Attribute Name**, enter a value.
 - c. [Optional] For **Group Search Extent**, select the **Search subtree** check box to include subtree directories during the binding process.
 - d. Enter a value for **LDAP Value**.
 - e. Select a **Viewpoint Role** from the list.

6. [Optional] Use     to change the order of roles, to add roles, and to delete them.
The order of roles affects their priority, which in turn can affect the user's settings in portlets.
7. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears.

Example: Assigning a Teradata Viewpoint Role Based on an LDAP Record Attribute

In the following example, **LDAP Value** represents the distinguished name (DN) of a group. When **Type** is **Attribute**, the **Attribute Name** field represents the user attribute where the group DN is stored.

This example assigns users to a Teradata Viewpoint role based on an attribute of the user's record in LDAP. Assume you have a user with the following record in LDAP:

```
dn: cn=User1,ou=Users,dc=teradata,dc=com
memberOf: cn=USA,ou=Groups,dc=teradata,dc=com
```

To map User1 to the Teradata Viewpoint role US Users, create the following role mapping:

1. Set **Type** to **Attribute**.
2. Type the **Attribute Name** `memberOf`.
3. Type the **LDAP Value** `cn=USA,ou=Groups,dc=teradata,dc=com`.
4. Set **Viewpoint Role** to **US Users**.

Example: Assigning a Teradata Viewpoint Role Based on LDAP Group Membership

In the following example, **LDAP Value** represents the distinguished name (DN) of a group. When **Type** is **Group**, the **Group Attribute Name** field is used for mapping and represents the group attribute where the user DN is stored.







This example assigns users to a Teradata Viewpoint role based on membership in an LDAP group with the following structure:

```
dn: cn=DBAs,ou=Groups,dc=teradata,dc=com
uniqueMember: cn=User2,ou=Users,dc=teradata,dc=com
```

To map User2 to the Teradata Viewpoint role Administrator, create the following role mapping:

1. Type `ou=Groups,dc=teradata,dc=com` in the **Group Search Base** field.
2. Type `uniqueMember` in the **Group Attribute Name** field.
3. Set **Type** to **Group**.
4. Type the **LDAP Value** `cn=DBAs,ou=Groups,dc=teradata,dc=com`.
5. Set **Viewpoint Role** to **Administrator**.

Configuring Role Mapping (JWT)

1. Do the following map a user's group to a Viewpoint role:
 - a. Enter a value for **Group Value**.
 - b. Select a **Viewpoint Role** from the list.
2. [Optional] Use     to change the order of roles, to add roles, and to delete them.
The order of roles affects their priority, which in turn can affect the user's settings in portlets.
3. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears.

Monitored Systems

The **Monitored Systems** portlet allows the Teradata Viewpoint Administrator to add, configure, enable, and disable systems, as well as view the amount of disk space used and set a threshold for the total disk usage alert.

General

Configure the system nickname, TDPID, site ID, host ID or hostname, login credentials, authentication method, and access to the Teradata, Aster, or Hadoop systems, and managed systems. Configure the system nickname, host ID, and login credentials for Teradata QueryGrid™. Test the connection and add or delete logins.

Data Collectors

Enable, disable, and configure data collectors to capture and retain portlet, disk usage, and resource data.

System Health

Enable metrics for the **System Health** and **Productivity** portlets. Configure degraded and critical thresholds for each metric. Create customized status and tooltips.

Canary Queries

Configure canary queries used to test Teradata system response times. The **System Heartbeat** canary query cannot be removed.

Alerts

Add, delete, copy, and configure alerts, or migrate existing Teradata Manager alerts.

Monitor Rates

Set Analytics Database internal sample rates for sessions, node logging, and vproc logging.

Log Table Cleanup

Configure log table cleanup to remove unnecessary information and reallocate space in the Analytics Database.

Cleanup Schedule

Create a schedule for cleaning up Teradata system log tables.

You can view the amount of disk space used and set a threshold for a disk usage alert.

Total Disk usage

View disk usage under /data directory on the Teradata Viewpoint server.

Viewpoint DB Disk usage

View disk usage of the Viewpoint database under /data/viewpointdb for all systems over time, and the distribution of data on the Teradata Viewpoint server.

12 hour change on Viewpoint DB

View the percentage of Viewpoint database disk space under /data/viewpointdb used or has become available over the past twelve hours. Positive indicates the amount of space used has increased. Negative indicates the amount of space used decreased.

Edit Alert


Set a threshold that sends an administrator alert when total disk usage in the DCS exceeds the threshold.

General



General system settings help you define the logins used to perform tasks on the system, such as data collection, canary queries, clean up, and TASM ruleset management. For certain systems you can define character sets and select authentication options.

Adding an Analytics Database

Add an Analytics Database to the **Monitored System** portlet to begin collecting data. Your system may be set up for you. Follow these steps if your Analytics Database is not set up.



1. Next to **Systems**, select , then select **Add Teradata System**.
2. Under **General System Details**, enter a system nickname, up to 30 characters.
3. [Optional] Select the **Enable system** check box.
After it is enabled, the system starts collecting data.
4. Enter the TDPID of the Analytics Database if not shown.
5. [Optional] Enter the site ID assigned by Teradata Customer Services.
After enabling a system, the **Site ID** becomes read-only if Viewpoint detects a value for the system.

A site ID is required when using Elastic Performance on Demand.

6. Enter a login name and password.
Specify an account string and authentication, add more logons, test the logon settings, and grant access to an Analytics Database. Add any logins used for Viewpoint data collection to the system-level bypass. To define system-level bypass settings, use the **Workload Designer** portlet.
7. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
8. [Optional] Under **Authentication Options**, specify which authentication mechanisms are available in the portlets.
9. [Optional] Under **Character Set**, select default character sets and enter a **JDBC Flag** value:
 - Select a character set from the **Session** list.
 - Select a character set from the **Monitor** list.
 - Enter a **JDBC Flag** value.
10. Under **Time Zone**, select the time zone the system uses to log DBQL and ResUsage data.
11. [Optional] Under **Collectors**, select **Enable Data Collectors (Configure in Data Collectors)** to begin collecting data.

The data collectors can be enabled and configured individually in **Data Collectors**.






Portlet	Data Collectors
Elastic Performance	Elastic Limit and Elastic Usage
TVS Monitor	Virtual Storage
Stats Manager	Stats Manager
Query Log	Query Log data collectors for historical DBQL data

12. [Optional] Under **Enhanced TASM Functions**, select **Enable this option if your Teradata has license entitlement to TASM**.
The workload management features are available with your TASM license on SLES11 or higher.
13. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Adding a Teradata Machine Learning Engine


Adding a Machine Learning Engine includes registering a certificate, testing login settings, and adding data collectors. Your Machine Learning Engine may be already set up for you. Follow these instructions if your Machine Learning Engine is not set up.

1. Before adding a Machine Learning Engine system to Viewpoint, download and register the SSL CA certificate from `https://master-node-IP:30009/` where *master-node-IP* is the IP address of the master node.

2. Next to **Systems**, select , then select **Add Teradata Machine Learning Engine**.
3. Under **General System Details**, enter a system nickname, up to 30 characters.
4. [Optional] Select the **Enable system** check box.
After it is enabled, the system starts collecting data.
5. Enter the **Host ID** of the system.
The **Host ID** is the IP address or host name of the system.
6. [Optional] Enter the site ID assigned by Teradata Customer Services.
After enabling a system, the **Site ID** becomes read-only if Viewpoint detects a value for the system.
7. Enter a login name and password.
The logon is case-sensitive.
8. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
9. [Optional] Under **Collectors**, select **Enable all data collectors (Configure in Data Collectors)** to collect data on this system.
The data collectors can be enabled and configured individually in **Data Collectors**.
10. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Adding a Hadoop System





The Hadoop system must be added and the logins created before you can configure the data collectors to monitor the system.

1. [After installing the Hadoop CA certificate](#), restart Viewpoint.
2. Select  next to **Systems** then select **Add Hadoop System**.
3. Under **General System Details**, enter a system nickname, up to 30 characters.
4. [Optional] Select the **Enable system** check box.
After it is enabled, the system starts collecting data.
5. Enter the host ID of the Hadoop system.
The host ID is the IP address or host name of the node running the Ambari service or Cloudera Manager on the Hadoop system. The port also needs to be specified if Ambari or Cloudera Manager is not running on port 80. For example, to indicate that Ambari or Cloudera Manager is running on node `hadoop1` and listening on port 8081, enter `hadoop1:8081` in this field.
6. [Optional] Select the **Use HTTPS** check box.
If you select this check box, a secure connection between the Viewpoint server and the Hadoop system is established after you install a trusted root certificate on the Viewpoint server. For more information, see [Certificates](#).
7. [Optional] Enter the site ID assigned by Teradata Customer Services.
After enabling a system, the **Site ID** becomes read-only if Viewpoint detects a value for the system.
8. Select one of the following platforms for this server:

Note:

This setting cannot be changed after you add the system.


Option	Description
Teradata Hadoop appliance	Select this option if this is a Teradata Hadoop appliance.
Licensed for Hadoop commodity system	Select this option if this system is running on commodity hardware and you are licensed to monitor the system with Teradata Viewpoint.

9. Enter a login name and password.
10. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
11. [Optional] Under **Collectors**, select the check box to activate Hadoop data collectors in order to collect data on this system.
The data collectors can be enabled and configured individually in **Data Collectors**.
12. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.
13. Restart the Viewpoint GUI.
14. [Optional] Depending on the version of Hadoop you are using, do one of the following to change port default settings:





HDP Version	Description
<ul style="list-style-type: none"> • CDH • HDP 2.1 and later 	Under Ports , enter a different port number to change the default setting for the HBase Master UI, Namenode UI, or ResourceManager UI ports.
<ul style="list-style-type: none"> • HDP 1.3 	Under Ports , enter a different port number to change the default setting for the Namenode UI, Jobtracker UI, or HBase Master UI ports.

Adding a Managed System

Teradata Viewpoint monitors and collects system-specific information. After you add a managed system in Teradata Viewpoint, you can test the configuration settings and enable it.






1. Select  next to **Systems**.
2. Select **Add Managed System** and then the system type.
 - **Viewpoint/Other:** Monitors managed systems for system-level metrics and collects information for Teradata Viewpoint servers.
 - **Unity Environment:** Enables the Unity multisystem environment for use with the high availability feature of the **Data Labs** portlet. Only one Unity environment is supported, and the name of

the managed server defaults to "Unity Environment." Before adding a Unity environment to **Monitored Systems**, Unity must be installed and configured, and all Teradata systems managed by the Unity Environment must be added to **Monitored Systems**.




3. Under **General System Details**, enter a system nickname, up to 30 characters.
4. [Optional] Select the **Enable system** check box.
If you cannot enable the managed system, contact the Teradata Viewpoint Administrator.
5. [Viewpoint and TMSMonitor 16.50.02.00 or later] To enable TMSMonitor to use HTTPS, select **Use HTTPS**.
[Viewpoint and TMSMonitor 16.50.04.00 or later] For TMSMonitor, HTTPS is enabled by default and HTTP connection is no longer supported.
6. Enter the host, login, and server information, depending on the type of managed system.
For a Viewpoint system, the hostname is the IP address or host name of the system. The default values only need to be changed if you change the login and password settings during the installation of TMS Monitor.
7. [Optional] Select **Test** to verify the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
8. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Adding QueryGrid

QueryGrid may already be set up for you. Follow these steps if your QueryGrid is not set up.

1. Register the cluster's SSL CA certificate.
2. Next to **Systems**, click , then select **Add QueryGrid**.
3. Under **General System Details**, enter a system nickname, up to 30 characters.
4. [Optional] Select the **Enable system** check box.
After it is enabled, the system starts collecting data.
5. Enter the IP address as the **host ID** of QueryGrid.
6. Enter a login name and password.
7. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
8. [Optional] Under **Collectors**, select **Enable all data collectors (Configure in Data Collectors)**.
The data collectors can be enabled and configured individually in **Data Collectors**.
9. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.






Enabling a System

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **General**.
3. Under **General System Details**, do any of the following:
 - Select the **Enable system** check box to enable a system.
After it is enabled, the system starts collecting data.
 - Clear the **Enable system** check box to disable a system.
When disabled,  appears next to the system name in the **Systems** list. The system cannot be accessed in other Teradata Viewpoint portlets, but the configured settings are maintained.
4. Select **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Enabling a Managed System

After a managed system is added, you can enable it to collect metrics for use in the **Viewpoint Monitoring System** portlet.

1. From the **Systems** list, select the name of the managed system you want to update.
2. Do one of the following:
 - Select the **Enable system** check box to enable a server.
 - Clear the **Enable system** check box to disable a server.

When disabled,  appears next to the server name in the **Systems** list.
If you cannot enable or disable the managed system, contact the Teradata Viewpoint Administrator.
3. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
4. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Deleting a Database System

Deleting a system removes all information configured for that system, such as data collectors, alerts, and canary queries. This is a permanent procedure that differs from disabling a database system.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **General**.
3. Under **General System Details**, select **Delete System**.

A confirmation message appears.

4. Select **Delete**.

Deleting a Managed System




This is a permanent procedure that differs from [disabling the managed system](#).

1. From the **Systems** list, select the name of the managed system you want to delete.
2. Under **General System Details**, select **Delete System**.
A confirmation message appears.
3. Select **Delete**.

Adding a Login Account

An unlimited number of login credentials can be configured for use in the **Monitored Systems** portlet.

The tdwm login is required to load and activate Workload Designer rulesets. The tdwm login and authentication cannot be removed or renamed, but you can type a password and an account string, and test the tdwm login settings.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **General**.
3. Under **Login**, select .
4. Type a **Login** name and password.
You also have the option to specify an account string, select an authentication mechanism, add more logins, test that the login settings are correct, and grant access to a Teradata system.
5. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Granting Privileges

Granting privileges provides Viewpoint data collectors with access to Teradata system resources.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **General**.
3. Under **Login**, select **Grant** for the user to be granted privileges.
4. Select a **Version**, if no Analytics Database version is found.
Listed privileges may vary depending on version selected.
5. [Optional] Copy the SQL.
You can execute the SQL using an external tool.
6. Click **Continue**.
7. Enter login information to authenticate the Teradata user.

The user must be authorized to grant privileges.

8. Select **Run**.
9. Review results and select **OK**.

Deleting a Login Account

At least one login account must be configured for a system. You cannot remove the last login account or the TDWM login.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **General**.
3. Under **Login**, select ☐ on the login row to be removed.
4. Select **Delete**.

The row disappears. The login is removed permanently.

Adding Portlet Authentication Options

Some Viewpoint portlets require authentication to access Analytics Database, for example, for changing the priority of a query. You can enable authentication options for Viewpoint portlets in addition to the default Analytics Database authentication mechanism.

1. From the **Systems** list, select the name of the system you want to update.
2. Under **Authentication Options**, select the check box to specify which authentication mechanisms to make available in the portlets that require them.
3. Click **Apply**.

Character Set Definitions

When you configure the general setup of a Teradata system, the **Character Set** settings allow you to define the *charset* that the system uses to communicate with Teradata Viewpoint. The character set definitions are optional. If you do not define them, Teradata Viewpoint uses the Teradata system settings.

Character Set Fields	Character Set Definition Suggestions
Session	The UTF8 session character set is strongly recommended for the Teradata Viewpoint session to provide end-to-end fidelity of character data by avoiding conversion between character sets.
Monitor	The UNICODE session character set is supported by Teradata Database 14.10 and later.
JDBC Flag	The JDBC Flag setting is a legacy-support feature to assist the transition away from the unsupported use of Analytics Database to store non-Latin characters in a Latin column, and the unsupported access of non-ASCII in an ASCII session character set. The JDBC Flag option specifies the CLIENT_CHARSET connection parameter, which can be used to override the Teradata JDBC Driver fixed-mapping of the Teradata session character set

Character Set Fields	Character Set Definition Suggestions
	<p>to the Java character set. The valid values are the actual Java character set names. Here are two examples of valid JDBC Flag settings:</p> <ul style="list-style-type: none"> • Windows-936 from Chinese (PRC) • Windows-950 for Chinese (Taiwan)



The default Teradata system settings are as following.

Teradata Database	Language Support Mode	Data Dictionary Charset	Teradata Viewpoint Session Charset	Teradata Viewpoint Monitor Charset	Limitations
14.10 and later	Standard	UNICODE	UTF8	UTF8	
14.10 and later	Japanese	UNICODE	UTF8	UTF8	

Viewing System Status

You can investigate the details of a system that is disabled or functioning with errors. Each level provides additional icons that represent system status.

1. From the **Systems** list, select the system whose name appears with one of the following icons:

Icon	Description
 Error	<p>Appears next to a system name or data collector name when one of these errors has been detected:</p> <ul style="list-style-type: none"> • Processing exception • Incorrect system privileges • Configuration errors, such as an incorrect password • System is offline
 Disabled	Appears next to a system name or data collector name if it is currently disabled.

Data Collectors

Use the Teradata Viewpoint Data Collection Service (DCS) to monitor systems. After a system has been configured in Teradata Viewpoint, data collectors can be configured to monitor the system. Data collectors gather information from different sources and make the data available to Teradata Viewpoint portlets. Each data collector has a sample rate, or frequency, used to collect data from the system and an initial retention policy used to keep the collected data for a time period or up to a certain size.

Most data collectors have similar behaviors as described in the data collector catalog, but a few are unique:

- Elastic Usage can automatically upload usage data to Teradata ServiceConnect.

- Query Log collects DBQL data in order to display database usage by users and applications, as well as suspect queries.
- Sessions collects SQL and Explain data when thresholds are met in the session.
- Sessions defines values to use when a Lower Priority action executes for an alert in the session.
- System Config provides critical system data to portlets and other collectors, so it cannot be disabled.

Configuring Data Collectors

After a system has been configured in Teradata Viewpoint, data collectors can be configured to monitor the database.

The login values used by data collectors are established in the **Monitored Systems** portlet by selecting a system, and then selecting **General** from the **Setup** list, under the **Login** section.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Data Collectors**.
3. Select the name of the data collector you want to configure.
4. [Optional] Select the **Enable data collector name Collector** check box to enable the data collector to begin processing data.
5. If **Login** appears, select the login you want to use to run the system task.
6. Under **Sample Rate**, select one of the following options:

Option	Description
Recommended sample rate	Uses the default sample rate for a data collector.
Custom sample rate	Adjusts the recommended sample rate to your system requirements.

7. Under **Delete Data**, select one of the following options:

Option	Description
Date	<ol style="list-style-type: none"> a. Select the After check box to set the calendar retention period. b. Enter a number and select a calendar value.
Size	<ol style="list-style-type: none"> a. Select the Over check box to set the file size retention period. b. Enter a number and select a file size restriction. <p>Note: The size of the data actually retained by the collector is always larger than the specified number because data clean up only drops table partitions that exceed this number. A 1024 byte kilobyte is used, so 1 GB is actually 1,073,741,824 bytes.</p>

Teradata Viewpoint uses a retention policy to keep collected data for a time period or up to a certain size. For more information, see [Data Collector Sample Rates](#).

8. If the collector details include additional options, apply the appropriate settings.

- Click **Apply**.

Configuring Elastic Usage Collector to Change the Reporting Interval

Complete the following task to change the Elastic Performance On Demand (EPOD) reporting interval from one-hour to 10-minute intervals. Teradata recommends setting the node sample rate on the Teradata system to less than or equal to 600 seconds to use this option.

Note:

Enabling this option changes the elastic performance on demand usage tracking calculation and billing model. Legacy EPP customers should contact the Teradata Account Team before enabling this option.

- From the **Systems** list, select the name of the system you want to update.
- From the **Setup** list, select **Data Collectors**.
- Select the **Elastic Usage** data collector.
- Check **Enable 10-minute interval reporting**.
- Click **Apply**.

Configuring Elastic Usage with Automatic Uploading


The Elastic Usage data collector can be enabled to collect usage data for the **Elastic Performance** portlet if you are using Elastic Performance on Demand. The collector can be configured to automatically upload usage data for billing purposes using Teradata ServiceConnect.

- From the **Systems** list, select the name of the system you want to update.
- From the **Setup** list, select **Data Collectors**.
- Select the **Elastic Usage** data collector.
- Under **Teradata ServiceConnect**, select the automatic uploading check box and enter the connection information:

Field	Description
Host	Host name or IP address of the SWS server or CMIC where Teradata ServiceConnect is hosted.
Port	Port number on which the SSH server is listening.
Username and Password	Login credentials that have the necessary permissions to upload usage data using the SWS server.

The public key is used by Teradata Corporation to verify the integrity of your usage data.

- [Optional] Select **Test Upload** to verify the Teradata ServiceConnect settings are correct.

If the connection is successful, it performs an on-demand of data upload from the start of the month.
If the operation fails,  appears with the error message.

6. Click **Apply**.

Configuring Query Log Collection

The Query Log data collector can be enabled to collect DBQL data from the PDCRDATA.DBQLLogTbl_Hst table in Analytics Database, which allows users to see database usage by users and applications as well as suspect queries in the **Query Log** and **Application Queries** portlets.

Up to 10 million rows of DBQL data is collected per day. When the number of DBQL rows exceeds 10 million, you can specify the name of a database view to load a subset of the query log data to make sure only the most important data is collected.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Data Collectors**.
3. Select the **Query Log** data collector.
4. [Optional] Select the **Enable Query Log Collector** check box to enable the data collector to begin processing data.
5. If **Login** appears, select the login you want to use to run the system task.
6. Under **Collection Time**, select a time zone to define the 24-hour period (midnight to midnight) for summarizing data.
7. Select one of the following to specify when the data collector begins collecting data:

Option	Description
After PDCR daily job	Select to collect data within an hour of when the daily PDCR job runs.
Daily at	Select and then specify a time to start the daily data collection relative to the specified time zone.

8. Under **Delete Summary Data**, select one of the following options to set the calendar or file size retention period.

Summary data refers to the data aggregated up to a system, user, or an application per day. This data is displayed in the **Trends** tab in the **Query Log** and **Application Queries** portlets.

Option	Description
Date	<ol style="list-style-type: none"> a. Select the After check box to set the calendar retention period. b. Enter a number and select a calendar value.
Size	<ol style="list-style-type: none"> a. Select the Over check box to set the file size retention period.

Option	Description
	<p>b. Enter a number and select a file size restriction.</p> <p>Note: The size of the data actually retained by the collector is always larger than the specified number because data clean up only drops table partitions that exceed this number. A 1024 byte kilobyte is used, so 1 GB is 1,073,741,824 bytes.</p>

9. Under **Delete Detail Data**, type a number up to 999 days to control how many days of detailed data is kept.

Detail data refers to the individual queries logged to the DBQL.

This setting controls the number of days you can view detailed query information in the **Query Log** and **Application Queries** portlets.

10. Under **Suspect Query Definition**, set any of the following options to define the criteria for what classifies a suspect query.

This setting controls what queries are flagged as being suspect in the **Query Log** and **Application Queries** portlets.

Option	Description
CPU Skew	Type a percentage up to 100 to flag as suspect those queries having a CPU skew over this percentage.
I/O Skew	Type a percentage up to 100 to flag as suspect those queries having an I/O skew over this percentage.
Product Join Indicator	Type a number up to 100 to flag as suspect those queries having a product join indicator over this number.
Unnecessary I/O	Type a number up to 100 to flag as suspect those queries having unnecessary I/O over this number.

- a. For **Tolerance**, type a number to set a threshold that is used to classify which queries are identified as suspect.

The number must be between 0 and 99999.

Increasing the **Tolerance** decreases the number of suspect queries by increasing the threshold for queries that consume smaller amounts of resources. For more information, see [Suspect Query Formulas](#).

11. Under **Row Limiting View**, enter a database view, in the format `databaseName.viewName`, that can be used to filter out less important data when the number of rows exceeds 10 million.
12. Click **Apply**.

Suspect Query Formulas

Formulas and Default Settings

Queries are flagged as suspect according to the suspect query formulas listed here. The tolerance value controls how many queries are flagged as suspect.

When the tolerance is zero, any queries that consume 1 AMP CPU second or greater and exceed the thresholds are flagged as suspect. It is normal for queries that consume smaller amounts of resources to exceed these thresholds and, increasingly so, on larger systems. The tolerance value prevents these smaller consuming queries from being flagged as suspect if they are not far from the threshold in proportion to the resources they consume. The larger the Teradata system, the larger the tolerance values need to be to make sure queries that are performing normally are not identified as suspect.

Metric	Formula	Suspect Query Formula	Default Threshold	Default Tolerance
CPUSkew	$100 * (1 - (\text{AMPCpuTime} / (\text{MaxAMPCPUTime} * \text{NumOfActiveAMPs})))$	$\text{CPUSkew} > 100 - ((100 - \text{Threshold}) * \text{AMPCPUTime} / (\text{Tolerance} + \text{AMPCPUTime}))$	40	10
IOSkew	$100 * (1 - (\text{TotalIOCount} / (\text{MaxAmplIO} * \text{NumOfActiveAMPs})))$	$\text{IOSkew} > 100 - ((100 - \text{Threshold}) * \text{TotalIOCount} / (\text{Tolerance} * 20,000 + \text{TotalIOCount}))$	40	10
PJI	$(\text{AMPCPUTime} * 1000) / \text{TotalIOCount}$	$\text{PJI} > (10 * \text{Tolerance}) / (\text{AMPCPUTime} - \text{Tolerance} + 1) + \text{Threshold}$ AND $\text{AMPCpuTime} \geq \text{Tolerance}$	3	10
UII	$\text{TotalIOCount} / (\text{AMPCPUTime} * 1000)$	$\text{UII} > (200,000 * \text{Tolerance}) / (\text{TotalIOCount} - (\text{Tolerance} * 1000) + 1) + \text{Threshold}$ AND $\text{TotalIOCount} \geq \text{Tolerance} * 1000$	20	10

Approaches to Obtaining an Optimal Tolerance Value

There are two different approaches to obtaining the optimal tolerance value.

With the first approach, you can use a visualization tool to scatter plot queries by CPU Skew or PJI versus AMPCPUTime, or IO Skew or UII versus TotalIOCount for a single day's worth of DBQL data. Then graph the corresponding suspect query formula on top of the scatter plot. Queries before the suspect query line are flagged as suspect. Adjust the tolerance in the suspect query formula to shape the line so the visual outliers are before the line, but the majority of queries are after the line.

With the second approach, you can run a SQL query that counts the number of rows that exceed the value of the suspect queries formulas for a days worth of DBQL data and adjust the tolerance values so that an actionable number of queries are returned.

Configuring Sessions for Lower Priority Actions

In a Teradata system, the Sessions data collector can define a value to use when a Lower Priority action executes for a session alert.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Data Collectors**.
3. Select the **Sessions** data collector.
4. Under **New priority for Lower Priority action**, enter one of the following values:
 - **\$L** is for low (default)
 - **\$M** is for medium
 - **\$H** is for high
 - **\$R** is for rush

When the Lower Priority action executes for an alert, the session that met the alert rules has its priority changed to the value you enter.

5. Click **Apply**.

Sessions with SQL and Explain Tabs

The Sessions data collector collects all details for a session and the currently running request. SQL and Explain text is collected when the **Enable SQL and Explain** check box is selected and when SQL and Explain thresholds are met. Query band information is collected when the **Enable SQL and Explain** check box is selected and when either SQL and Explain thresholds are met or for utility sessions.

The SQL and Explain thresholds can be set high or low and corresponding SQL and Explain text appears in the SQL and Explain tabs in query details views. The SQL associated with a query can help you target and optimize the query. The Explain text shows a query's progress through the Explain plan so you can optimize the query and change the Explain plan.

The Teradata Viewpoint system is configured to collect longer-running queries when SQL and Explain collection is enabled. Set the threshold lower to include shorter running queries, which may increase performance overhead.

The explain plan progress for all queries is updated at the collection rate of the Sessions collector. If you want to fetch the latest explain plan progress for a particular query more frequently than that rate, you can do so in the **Query Monitor** or **Query Groups** portlets. You can enable this feature per role by granting the **Get Explain from DB** permission for one or both of these portlets in **Roles Manager**.

Requirements for SQL and Explain Collection

The following are the requirements for SQL and Explain collection:

- Request is in the DBC/SQL partition.
- Request must be active when collected.

- Request must exceed the threshold of 10 elapsed seconds or 5 CPU seconds (active thresholds are per request, not per session).

In addition, note the following for the utilities:

- Utility partition data sessions do not run SQL.
- Utility Control sessions in the DBC/SQL partition normally do not exceed the threshold.

Configuring Sessions with SQL and Explain Tabs

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Data Collectors**.
3. Select the **Sessions** data collector.
4. Select the **Enable SQL and Explain** check box.
5. Under **SQL and Explain Thresholds**, set one or more thresholds:
 - **Elapsed seconds**
 - **Total CPU seconds**

When a threshold is exceeded, SQL and Explain text is collected for the session and the SQL and Explain tabs appear in query details views.

Note:

Setting SQL and Explain thresholds too low might impact your system performance.

6. Click **Apply**.

Configuring Stats Manager

Prerequisite:

You must make sure the database version information has been collected before you configure the Stats Manager data collector. The Stats Manager data collector will not display in the list of collectors until the database version information is collected.





The Stats Manager data collector can be enabled to collect statistics data for the **Stats Manager** portlet.

The **Stats Manager** portlet and data collector are available for systems running Teradata Database 14.10 and later.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Data Collectors**.
3. Select the **Stats Manager** data collector.
4. Enter a login name and password.

The login name must be unique from the other logins already configured for the system to allow overhead of statistics collection and analysis to be measured separately from other Teradata Viewpoint monitoring overhead.

You also have the option to specify an account string and authentication mechanism.

5. [Optional] Click **Test** to verify that the login settings are correct.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the login credentials are valid and the host can be reached.
6. [Optional] Select **Grant** to provide Viewpoint data collectors with access to Teradata system resources.
7. Click **Apply**.
If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Data Collector Catalog

Data collectors gather information from different sources and make the data available to Teradata Viewpoint portlets.

Database Data Collectors

Database Data Collector	Description	Portlets Using Data
Account Info	<p>Queries the DBC.AccountInfo table in the database and collects a list of available account strings.</p> <p>The Queries portlets use this data to list the available account strings for the Change Priority function.</p> <p>The Workload Designer portlet uses this data to provide a list of available account strings when creating classification criteria.</p>	<ul style="list-style-type: none"> • My Queries • Query Groups • Query Monitor • Workload Designer
Alert Request	<ul style="list-style-type: none"> • Monitors the dbc.mngr.AlertRequest table on the Teradata system for incoming alert requests. If the table row contains valid data, the contents are forwarded to the Alert Service to process the alert action. • Monitors the dbc.mngr.MonitorRequest table on the Teradata system. Any database utility or user program can request Teradata Viewpoint to monitor its progress by inserting rows into the dbc.mngr.MonitorRequest table. <p>Each row includes fields that indicate the date and time by which the next row is inserted. If a new row is not inserted before the specified date and time, the collector forwards the contents to the Alert Service to process the alert action.</p> <ul style="list-style-type: none"> • Monitors the dbc.sw_event_log table on the database. 	<ul style="list-style-type: none"> • Alert Viewer
AWT Info	<p>Collects AMP worker task (AWT)-related information using the AMP Load (ampload) utility. The collector executes the</p>	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis

Database Data Collector	Description	Portlets Using Data
	<p>MONITOR AWT RESOURCE request of PM/API instead of AMP Load (ampload). If this collector is not enabled, the portlets do not display AWT information.</p> <p>The AWT Info collector is also required for alerts that depend on AWT information.</p>	<ul style="list-style-type: none"> • Node Resources • System Health • Today's Statistics
Database Space	<p>Queries the DBC.DataBaseSpace and DBC.DBase tables to collect database space usage metrics.</p> <p>The Space Usage and Metric Heatmap portlets use the data collected by the Database Space collector. If this collector is not enabled, these portlets do not display current database space data.</p>	<ul style="list-style-type: none"> • Metric Heatmap • Space Usage
Dictionary	<p>Collects and caches information from the Teradata Data Dictionary.</p> <p>If this collector is not enabled, the Query Group Setup and Workload Designer portlets do not display the most current usernames.</p> <p>The Dictionary collector is required to generate Database Space alerts.</p>	<ul style="list-style-type: none"> • Data Labs • Lock Viewer • Monitored Systems • Query Group Setup • Query Monitor • Query Spotlight • Roles Manager • Space Usage • Workload Designer • Workload Monitor
Disk Space	<p>Obtains disk-space and database space usage data by querying the DBC.DiskSpace view of the DBC database.</p> <p>If this collector is not enabled, these portlets do not display the most current disk-space data.</p>	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis • System Health
	<p>Collects the perm space data and skew percentage for the DBC.TransientJournal table.</p>	<ul style="list-style-type: none"> • Alert Viewer
Elastic Limit	<p>Collects the current system CPU limit set in the Schmon utility or Workload Designer portlet.</p> <p>Only enable this collector for Teradata systems leveraging Elastic Performance On Demand.</p>	<ul style="list-style-type: none"> • Elastic Performance
Elastic Usage	<p>Collects the overall node CPU usage from the DBC.ResUsageSPMA table.</p> <p>Only enable this collector for Teradata systems leveraging Elastic Performance On Demand.</p>	<ul style="list-style-type: none"> • Elastic Performance
Lock Info	<p>Uses the data written to the DBC.DBQLXMLLockTbl table to capture lock information. For the Lock Info collector to run against Teradata Database 14.10 and later, DBQL query logging must be enabled with the WITH LOCK option.</p> <p>The Lock Viewer portlet displays updated lock data only for the period of time when the collector is enabled.</p>	<ul style="list-style-type: none"> • Lock Viewer

Database Data Collector	Description	Portlets Using Data
	For more information, see Lock Info Collector .	
Query Count	Database query count and query log data are collected by querying the DBC.DBQLogTbl and DBC.DBQLSummaryTbl tables. Query logging in the database must be enabled for a query to be counted. For more information, see Query Count Collector .	<ul style="list-style-type: none"> • Metric Heatmap • Productivity • Today's Statistics • Workload Designer
Query Log	Collects DBQL data from the PDCRDATA.DBQLogTbl_Hst table. Up to 10 million rows of DBQL data is collected per day. When the number of DBQL rows exceeds 10 million, you can specify the name of a database view in the Monitored Systems portlet to load a subset of the query log data to make sure only the most important data is collected.	<ul style="list-style-type: none"> • Application Queries • Query Log
Resource Usage	Collects node resource usage data from the DBC.ResUsageSPMA table and DBC.ResUsageIPMA table. Resource Usage collects vproc usage data from the DBC.ResUsageSVPR table. The recommended collection is the rate that this information is logged to these tables in the database. You can set this rate in the Monitored Systems portlet by selecting Monitor Rates from the Setup list.	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis
Sessions	Collects snapshots of database session data for each login from the Teradata Monitor partition by using the Teradata Performance Monitor API. For more information, see Sessions Collector .	<ul style="list-style-type: none"> • Metric Heatmap • My Queries • Query Groups • Query Monitor • Query Spotlight • SQL Scratchpad • Stats Manager • System Health • Today's Statistics • Workload Monitor
Stats Manager	Collects and retains statistic data used by Teradata system tables from the DBC.Dbase, DBC.TVM, DBC.Indexes, DBC.StatsTbl, TDStats.StatsTbl, TDStats.ObjectList, and TDStats.ExcludedTbl tables.	<ul style="list-style-type: none"> • Stats Manager
System Config	Collects and retains Teradata system configuration data used by portlets and collectors. The collector retrieves the database version from DBC.DBCInfo. This information is critical to provide proper functionality of several portlets and other data collectors. To gather physical system configuration data, the collector also uses the MONITOR PHYSICAL CONFIG request of the PM	<ul style="list-style-type: none"> • Alert Viewer • Canary Response Times • Lock Viewer • Metric Heatmap • Metrics Analysis • My Queries

Database Data Collector	Description	Portlets Using Data
	<p>/API and executes STATUS PDE of the Vproc Manager console partition utility.</p> <p>This collector cannot be disabled. The sample rate and initial retention policy for this collector cannot be modified.</p> <p>This collector does not call Vproc Manager console partition utility.</p>	<ul style="list-style-type: none"> • Node Resources • Productivity • Query Monitor • Query Spotlight • Space Usage • System Health • Today's Statistics • TVS Monitor • Viewpoint Monitoring • Workload Designer • Workload Health • Workload Monitor
System Stats	<p>Collects system, node, and vproc statistics data by querying the Teradata Monitor partition using the PM/API. Physical and virtual resource statistics are collected from Teradata system memory at a configurable sample rate.</p> <p>Collecting system statistics data more frequently than the resource sample rate in database results in duplicate data being collected and wasted resource usage on the database. You can set this rate in the Monitored Systems portlet by selecting Monitor Rates from the Setup list.</p> <p>If this collector is not enabled, these portlets do not display system statistics data.</p> <p>The Workload Designer portlet uses node and vproc counts from this collector to assist with the creation of rulesets.</p> <p>The System Stats collector is required for system, node, vproc, and System Health alerting.</p>	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis • Node Resources • System Health • Today's Statistics • Workload Designer • Workload Health • Workload Monitor
Table Space	<p>Collects and retains data about the space used by Teradata system tables from the DBC.Dbase, DBC.TVM, DBC.Indexes, and DBC.DataBaseSpace tables.</p> <p>The Space Usage and Metric Heatmap portlets use the data collected by the Table Space collector. If this collector is not enabled, these portlets do not display the most current table space data.</p>	<ul style="list-style-type: none"> • Metric Heatmap • Space Usage • System Health
TASM Config	<p>Obtains configuration information about the TDWM version and active ruleset in the database from the TDWM database in Teradata.</p> <p>The TASM Config collector needs to be enabled to use Workload Designer for a new Teradata system. If the collector is not enabled, the TDWM version information is not collected, and the Teradata system does not appear in the list of available Teradata systems in Workload Designer.</p>	<ul style="list-style-type: none"> • My Queries • Query Groups • Query Monitor • Workload Designer • Workload Health • Workload Monitor

Database Data Collector	Description	Portlets Using Data
TASM Distribution	Collects information about the average percentage of CPU consumption used by each workload for each node in the Teradata system. The TASM Distribution data collector executes <code>schmon -M -p -P</code> in the Teradata Console partition to collect information for the CPU Consumption graph in the Distribution view. The collector executes the MONITOR WD request of PM/API instead of <code>schmon</code> .	<ul style="list-style-type: none"> • Workload Monitor
TASM Exception	Obtains TASM Exception data from the DBC. TDWMEExceptionLog table.	<ul style="list-style-type: none"> • Workload Monitor
TASM State	Collects and retains system regulation data. You must enable both the TASM Config and TASM State collectors to be able to associate the active workload name with the session information displayed in these portlets.	<ul style="list-style-type: none"> • Workload Health • Workload Monitor
TASM Summary	Obtains data for each active workload in the database. This collector sends queries to the Teradata Monitor partition using PM/API.	<ul style="list-style-type: none"> • Workload Health • Workload Monitor
Virtual Storage	Collects data on cylinder temperature and storage grade from space types managed by the Teradata Virtual Allocation Manager (TVAM) utility. This collector uses the GetTvsUdfViewpointSummary user-defined function, created by the database DIP scripts, to collect this data.	<ul style="list-style-type: none"> • TVS Monitor

Teradata Aster Data Collectors

Teradata Aster Data Collector	Description	Portlets Using Data
Cluster Status	Collects cluster status and node status information for each node.	<ul style="list-style-type: none"> • Node Monitor • System Health
Component Stats	Collects statistics for each node of the Teradata Aster system, such as CPU, disk I/O, net I/O, and memory use.	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis • Node Monitor • System Health
Nodes	Collects detailed information for each node, including the node storage and node virtual worker information.	<ul style="list-style-type: none"> • Node Monitor • Space Usage • System Health
Processes	Collects data about sessions logged into the system and processes. For active processes, detailed information about the statement and phases of the process is also collected.	<ul style="list-style-type: none"> • Completed Queries • Query Monitor • System Health

Teradata Aster Data Collector	Description	Portlets Using Data
Rulesets	Collects the active ruleset information for the Teradata Aster system.	<ul style="list-style-type: none"> • Workload Designer
Space Usage	Collects database, table, and node storage information.	<ul style="list-style-type: none"> • Node Monitor • Space Usage • System Health
System Config	Collects the site ID and Aster Database version.	<ul style="list-style-type: none"> • Monitored Systems • Workload Designer
Workloads	Collects workload policy and workload service class data.	<ul style="list-style-type: none"> • Completed Queries • Query Monitor

Teradata Machine Learning Engine Collectors

Teradata Machine Learning Engine Collector	Description	Portlets Using Data
Cluster Status	Collects cluster status and information about Teradata Machine Learning Engine queen and workers.	<ul style="list-style-type: none"> • Node Monitor • System Health
Component Stats	Collects statistics on Teradata Machine Learning Engine queen and workers such as CPU, disk I/O, net I/O, and memory use.	<ul style="list-style-type: none"> • Metric Heatmap • Metrics Analysis • Node Monitor • System Health
Pods	Collects detailed information, including storage and virtual worker information.	<ul style="list-style-type: none"> • Node Monitor • System Health
Processes	Collects data about sessions logged into the system and processes. For active processes, detailed information about the statement and phases of the process is also collected.	<ul style="list-style-type: none"> • Completed Queries • Query Monitor • System Health
System Config	Collects the site ID and system version.	<ul style="list-style-type: none"> • Monitored Systems

Teradata Hadoop Data Collectors

Teradata Hadoop Data Collector	Description	Portlets Using Data
HBase	Collects the number of region servers that are running and configured and the amount of time the master server has been running.	<ul style="list-style-type: none"> • Hadoop Services

Teradata Hadoop Data Collector	Description	Portlets Using Data
HDFS	Collects the number of files and directories in HDFS, the space capacity not currently used by HDFS, datanode process information, and the amount of time a name node is running.	<ul style="list-style-type: none"> • Hadoop Services • Metric Heatmap • Metrics Analysis • Node Monitor • Space Usage • System Health
MapReduce	(HDP 1.3) Collects map and reduce jobs statistics, jobtracker information, the number of map slots occupied and reserved, the number of reduce slots occupied and reserved, and the number of tasktrackers running and configured.	<ul style="list-style-type: none"> • Hadoop Services • Metric Heatmap • Metrics Analysis • Node Monitor • System Health
Services	Collects Hadoop services information.	<ul style="list-style-type: none"> • Hadoop Services • Node Monitor • System Health
System Config	Collects the list of nodes that make up the Teradata Hadoop system and the site ID from the system. This data is used by portlets and other collectors, so this collector cannot be disabled or modified.	<ul style="list-style-type: none"> • Hadoop Services • Node Monitor
System Stats	Collects system statistics data from Hadoop Services.	<ul style="list-style-type: none"> • Hadoop Services • Metric Heatmap • Metrics Analysis • Node Monitor
YARN	(CDH and HDP 2.1 and later) Collects YARN statistics, ResourceManager information, the amount of cluster memory allocated and reserved, the number of containers allocated and reserved, and the number of applications running and pending.	<ul style="list-style-type: none"> • Hadoop Services • Metric Heatmap • Metrics Analysis • Node Monitor • System Health

Teradata QueryGrid Data Collectors

Teradata QueryGrid Data Collector	Description	Portlets Using Data
Issues	(Teradata QueryGrid™ 16.10 and later) Collects Teradata QueryGrid™ issues.	<ul style="list-style-type: none"> • QueryGrid
Queries	Collects information about queries run on a Teradata QueryGrid system.	<ul style="list-style-type: none"> • QueryGrid

Teradata QueryGrid Data Collector	Description	Portlets Using Data
System Config	(Teradata QueryGrid™ 16.10 and later) Collects and retains the latest snapshot of Teradata QueryGrid™ instances. This data is used by portlets and other collectors, so this collector cannot be disabled or modified.	<ul style="list-style-type: none"> • QueryGrid

Lock Info Collector

This topic provides additional information about the Lock Info data collector that may be helpful.

The following parameters are passed to the Lock Logger database utility.

Parameter	Description
Run continuously	No
Number of sessions	1
Character set	* (default character set)
Table name to write lock log entries	lockLoggerTableyyyyMMddHHmmss
Utility to create table	Yes
Time constraint at or later than	One hour before previous collection time
Time constraint at or prior to	* (current time)
Object constraint	* (all objects)

When the Lock Info collector is not executing, there should be only one lock table in the database of the database user being used for data collection. The table name starts with "lockLoggerTable" and ends with a unique timestamp. If there is more than one lock table, older tables can be safely dropped as they are not being used by the DCS.

Query Count Collector

This topic provides additional information about the Query Count data collector that may be helpful.

When enabling query logging, it is important to manage the size of the DBQL tables. Teradata recommends clearing the DBQL tables nightly. The queries performed by Teradata Viewpoint against DBQL require an all-row scan. If the size of the DBQL tables is not managed, queries against them can cause unnecessary use of database resources. Teradata Professional Services has a Data Collection and Capacity Planning offering that includes the movement and cleanup of DBQL data on a nightly basis. If DBQL is not cleaned up nightly, using the Query Count collector is not recommended.

The **Productivity** portlet uses query count data to show total query counts and query counts by application hourly. The **Today's Statistics** portlet uses the data to show query counts and query log data for the last hour of collected data, grouped by duration. The **Metric Heatmap** portlet shows query count and query log data collected over the hour. The **Workload Designer** portlet uses this data to populate several select lists when creating classification criteria.

The default collection rate for querying DBQL is one hour. To get the portlets to display the most current query counts, you can increase the collection frequency, but at the cost of using more database resources to query DBQL more frequently.

Sessions Collector

This topic provides additional information about the Sessions data collector that may be helpful.

Session-level statistics are collected in memory by the database at a configurable sample rate. The default option in Teradata Viewpoint collects session data at the same rate as sampled by the database. Collecting session data more frequently than the session sample rate in the database results in duplicate data collection and wasted CPU. You can set this rate in the **Monitored Systems** portlet by selecting **Monitor Rates** from the **Setup** list.

Note:

If you want to see a query with an interval shorter than 60 seconds on Query Monitor, change the interval value in the ruleset for the system.

The data is also used for the Active Sessions metric in the **System Health**, **Metrics Analysis**, and **Metric Heatmap** portlets.

If this collector is not enabled, the Queries portlets always display zero sessions. If this collector is enabled and then disabled, the Queries portlets display stale session data. The Sessions collector and the System Stats collector must be enabled for the Active Sessions metric to display data in the **System Health** and **Metric Heatmap** portlets.

The Sessions collector retains utility data for the **Query Monitor** utility view. The Sessions collector is required for Session alerting.

SQL and Explain data is collected by querying the Teradata Monitor partition using the PM/API. It gathers SQL, Explain and query band information from currently running sessions for requests that meet the SQL and Explain thresholds. SQL, Explain and query band information displays in the details view in the following portlets only if SQL and Explain data is collected for that time period:

- **My Queries**
- **Query Groups**
- **Query Monitor**
- **Query Spotlight**
- **Workload Monitor**

Configuring TDWMStatistics Timeout

To configure the TDWM statistics timeout, perform the following steps:

1. Open `/opt/teradata/viewpoint/dcs/config/dcs.properties`
2. Uncomment the property `tdwmstatistics.collector.sqlquerytimeout`
3. Set the timeout (in seconds) as per your preference
4. Save the file and restart the dcs service: `service dcs restart`

Data Collector Sample Rates

Teradata Viewpoint uses a sample rate to collect data from the system and uses a retention policy to keep the collected data for a time period or up to a certain size. The sample rate and retention policy vary for each data collector and most can be adjusted as needed.

Not applicable under the **Initial Retention Policy** column indicates that data collector does not store data.

Teradata Data Collectors

Data Collector	Recommended Sample Rate	Initial Retention Policy
Account Info	12 hours	1 week
Alert Request	5 minutes	Not applicable
AWT Info	60 seconds	1 year
Database Space	1 day	10 GB
Dictionary	1 hour	Not applicable
Disk Space	1 hour	1 year
Elastic Limit	60 seconds	1 year
Elastic Usage	1 hour	1 year
Lock Info	1 hour	10 GB
Query Count	DBQLFlushRate value from the DBS Control utility display command, otherwise 1 hour	10 GB
Query Log	Within an hour after the PDCR daily job completes	2 years for summary data 14 days for detail data
Resource Usage	Node/Vproc logging rate as set in Monitor Rates , otherwise 1 hour	10 GB
Sessions	Session sampling rate as set in Monitor Rates , otherwise 30 seconds	10 GB
Stats Manager	1 hour	3 months

Data Collector	Recommended Sample Rate	Initial Retention Policy
System Config	12 hours	1 week
System Stats	Node/vproc collection as set in Monitor Rates , otherwise 30 seconds	1 year
Table Space	1 day	10 GB
TASM Config	10 minutes	None (both disabled)
TASM Distribution	60 seconds	1 year
TASM Exception	Sample rate returned by PM/API TDWM Summary, otherwise 60 seconds	1 year
TASM State	60 seconds	1 year
TASM Summary	Sample rate returned by PM/API TDWM Summary, otherwise 60 seconds	1 year
Virtual Storage	12 hours	1 year

Aster Data Collectors

Data Collector	Recommended Sample Rate	Initial Retention Policy
Cluster Status	10 seconds	1 year
Component Stats	30 seconds	1 year
Nodes	5 minutes	1 year
Processes	10 seconds	10 GB
Rulesets (Database 15.0 and later)	10 minutes	Not applicable
Space Usage	1 day	10 GB
System Config	12 hours	10 GB
Workloads	1 hour	1 year

Teradata Machine Learning Engine Data Collectors

MLE Data Collector	Recommended Sample Rate	Initial Retention Policy
Cluster Status	10 seconds	1 year
Component Stats	60 seconds	1 year
Pods	300 seconds	1 year
Processes	60 seconds	10 GB

MLE Data Collector	Recommended Sample Rate	Initial Retention Policy
System Config	12 hours	10 GB

Teradata Hadoop Data Collectors

The range of allowed values for the sample rate of the Hadoop data collectors are 15-3600 seconds.

Data Collector	Recommended Sample Rate	Initial Retention Policy
HBase	60 seconds	1 year
HDFS	60 seconds	10 GB
MapReduce (HDP 1.3)	60 seconds	10 GB
Services	60 seconds	10 GB
System Config	5 minutes	Not applicable
System Stats	60 seconds	10 GB
YARN (CDH and HDP 2.1 and later)	60 seconds	10 GB



Teradata QueryGrid Data Collectors

Data Collector	Recommended Sample Rate	Initial Retention Policy
Issues	30 seconds	1 year or 10 GB
Queries	Query summary frequency as set in Managers > Settings in the QueryGrid portlet, otherwise 30 seconds	1 year or 10 GB
System Config	12 hours	Not applicable

Viewing Data Collector Status



You can investigate the details of a data collector that is disabled or functioning with errors.

- From the **Systems** list, select the system that has one of the following icons:

Icon	Description
 Error	Appears next to a system name or data collector name when one of these errors has been detected: <ul style="list-style-type: none"> Processing exception Incorrect Teradata system privileges Configuration errors, such as an incorrect password Teradata system is offline
 Disabled	Appears next to a system name or data collector name if it is currently disabled.

- From the **Setup** list, select **Data Collectors**.

The **Data Collectors** list refreshes the condition of each collector, identifying collectors that are disabled or have errors.

- Do one of the following:
 - Select  next to the collector name of the log you want to view.
 - Select  next to the name of the collector you want to enable.

System Health

You can customize system status and tooltips and configure metrics and thresholds. The thresholds are settings for the data collected by canary queries and the Disk Space, Sessions, and System Stats data collectors.

- For Teradata systems, the system status, tooltips, metrics, and thresholds appear in the **System Health** and **Productivity** portlets.
- For Teradata Aster and Hadoop systems, the system status, tooltips, metrics, and thresholds appear only in the **System Health** portlet.

Configuring System Health Values

You can define threshold values for your system status calculation.

- From the **Systems** list, select the name of the system you want to update.
- From the **Setup** list, select **System Health**.
- Under **Thresholds**, select one of the following options for each metric:

Option	Description
Enabled	Makes the metric visible in the System Health portlet. Uses the threshold values in the system status calculation. User , System , and Wait I/O is factored into the CPU Utilization calculation when each of these threshold values is enabled.
Disabled	Omits the metric in the System Health portlet. Does not use threshold values in the system status calculation.
View Only	Makes the metric visible in the System Health portlet. Does not use threshold values in the system status calculation.

- Adjust the degraded or critical thresholds for each metric.
Only critical thresholds can be adjusted for **Services Bad (Components)** for CDH.
- If CPU skew or I/O skew metrics are **Enabled**, do one of the following:
 - Select the CPU skew or I/O skew qualification check box and specify a threshold.

If the system CPU is under the defined threshold, the skew metric is excluded from the system health calculation. The metric displays in the **System Health** portlet as **View Only**.

- Clear the CPU skew or I/O skew qualification check box to include the skew metric in the system health calculation, regardless of system CPU.

6. Click **Apply**.

System Health Metrics

You can enable the metrics and thresholds that appear in the **System Health** and **Productivity** portlets. Available metrics and the portlets in which the data appears differs depending on the type of system you choose.

Teradata System Metrics

Metric	Description
Active Sessions	Number of sessions with active queries
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP
AMP Worker Tasks	Average number of AMP worker tasks in use on each AMP
Component Down	Number of components, such as BYNETs or AMPs, that are down
(CPU) System	Average CPU time spent executing kernel system calls or servicing I/O and timer hardware interrupts
(CPU) User	Average CPU time spent executing code on behalf of operating system user processes
CPU Utilization	Average node CPU use. CPU is calculated as the sum of the user CPU, system CPU, and wait I/O usage percentages, depending on which of these are enabled by the Teradata Viewpoint Administrator in the Monitored Systems portlet
(CPU) Wait I/O	Average percentage of CPU time spent waiting for I/O
DBC Disk Space	Available DBC disk space in use
Enabled CPU Usage Note: In Viewpoint Mobile view this metric is currently not available.	CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled
Max Disk by AMP	Available disk space currently in use
Max Spool by AMP	Available spool space in use
Memory Used	Represents an increase or decrease in the node-level memory. A null value or negative number reports as 0.

Metric	Description
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of disk use on the busiest node to the average node
System Heartbeat	Canary query showing the response time of the most recent system heartbeat in milliseconds
Total Disk Space	Percentage of total disk space currently in use
Additional Canaries	Additional canary queries that are defined for the system and used as part of the monitoring of system health

Aster System Metrics

Metric	Description
Active Sessions	Number of users and applications currently connected to database
Component Down	Number of nodes that are not available
Component Passive	Number of nodes not processing queries but can be made ready to process queries when needed
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Memory	Average node memory use
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Queen Disk Space	Percentage of used disk space on the queen node
Replication Factor	Number of copies of the user data
Total Space	Percentage of used space to overall storage capacity

Teradata Machine Learning Engine Metrics

Metric	Description
Active Sessions	Number of users and applications currently accessing the engine
CPU	Average worker CPU and queen use
CPU Skew	Comparison of CPU use on the busiest worker/queen to the average worker/queen

Teradata Hadoop System Metrics

CDH

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Services Bad	Number of services in a critical state
Services Concerning	Number of services in a degraded state
Total Space	Percentage of used space to overall storage capacity

HDP 2.1 and later

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances

Metric	Description
Components Down	Number of services not started
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Total Space	Percentage of used space to overall storage capacity

HDP 1.3

Metric	Description
Blocks Corrupt	Number of blocks whose replicas are all corrupt
Blocks Missing	Number of blocks with no replicas anywhere in the cluster
Components Down	Number of service components not running
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.
Jobs Failed	Number of jobs that failed
Jobs Running	Number of jobs currently executing in the system
Job Tracker CPU	CPU use for the node running the jobtracker service
Map Tasks Running	Number of map tasks executing in the system
Map Tasks Waiting	Number of map tasks waiting to execute
Max Disk by Node	Amount of used disk space on the node with the most disk space in use
Name Node CPU	Node CPU use for the node running the namenode service
Name Node Heap	Percentage of heap space used in the namenode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node

Metric	Description
Reduce Tasks Running	Number of reduce tasks executing in the system
Reduce Tasks Waiting	Number of reduce tasks waiting to execute
RPC Latency JT	Average wait time in queue for jobtracker service calls
RPC Latency NN	Average wait time in queue for namenode service calls
Total Space	Percentage of used space to overall storage capacity

Configuring Custom Status and Tooltips

You can change the default definitions of status and tooltips by enabling and defining custom status and tooltips for your system.

For Teradata systems, these changes appear in the **System Health** and **Productivity** portlets. For Aster and Hadoop systems, these changes appear only in the **System Health** portlet.


1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **System Health**.
3. Under **Status and Tooltips**, select **Enable custom status**.
4. Change the status and tooltip definitions for your system.
5. [Optional] Select ☐ at the end of the tooltip to restore the default definition.
6. Click **Apply**.

Copying System Health Settings between Systems

You can copy the custom status, tooltips, and threshold settings to other compatible systems. You cannot copy canary query threshold settings to other systems.

Note:

All copied settings overwrite the target system settings.

1. From the **Systems** list, select the source system name.
2. From the **Setup** list, select  next to **System Health**.
3. [Optional] Select **Copy status and tooltips** to copy your custom status.
4. [Optional] Select **Copy thresholds** to copy your threshold settings, except canary queries.
5. Select target systems.
6. Select **Copy**.

Canary Queries

After a Teradata system is configured in Teradata Viewpoint, the System Heartbeat canary query checks the responsiveness of the system. You can also create user-defined canary queries that gather and retain data to display in the **Canary Response Times**, **Productivity**, and **System Health** portlets. Canary queries are not data collectors.

Adding or Copying a Canary Query

1. From the **Systems** list, select the name of the Teradata system you want to update.
2. From the **Setup** list, select **Canary Queries**.
3. Do one of the following:

Option	Description
Add	Select <input type="checkbox"/> to add a canary query.
Copy	Select <input type="checkbox"/> in the row of the canary query you want to copy.

4. Enter a name for the canary query.
5. [Optional] Select the **Enable query** check box to enable the canary query.
6. Under **SQL**, enter or modify the SQL for the query.
7. Under **Default Database**, specify the default database.
8. If **Login** appears, select the login you want to use to run the system task.
 - a. [Optional] Select the **Log in each time the query executes** check box to force a new log in connection to the Teradata system each time that the query runs.
 - b. [Optional] Enter a timeout value. The value indicates the number of seconds before abandoning the attempt to log in to the Teradata system.
Enter a value only if you selected the **Log in each time the query executes** check box. You cannot add an alert action for a log in timeout for a **System Heartbeat** canary query.
9. For **Collect Data Every**, specify the data collection frequency in minutes.
The recommended sample rate is 2 minutes. The initial retention policy is 1 year.
10. Under **Execute**, select one of the following options:



Option	Description
24/7	Select 24/7 to execute the query 24 hours a day, 7 days a week.
Custom	<ol style="list-style-type: none"> a. Select Custom to execute the query during a time frame. b. Specify the start time and end time. c. Specify the start day and end day.

11. Under **Delete Data**, select one of the following options:

Option	Description
Date	a. Select the After check box to set the calendar retention period. b. Enter a number and select a calendar value.
Size	a. Select the Over check box to set the file size retention period. b. Enter a number and select a file size restriction. Note: The size of the data actually retained by the collector is always larger than the specified number because data clean up only drops table partitions that exceed this number. A 1024 byte kilobyte is used, so 1 GB is actually 1,073,741,824 bytes.

Teradata Viewpoint uses a retention policy to keep collected data for a time period or up to a certain size. For more information, see [Data Collector Sample Rates](#).

12. Click **Apply**.

If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Deleting a Canary Query

System Heartbeat is a canary query that cannot be deleted. It is used to check whether the Teradata system is responsive.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Canary Queries**.
3. Select ☐ in the row of the canary query you want to delete.
A confirmation message appears.
4. Click **OK**.
The canary query row disappears.

Alerts

Alerts monitor the performance of a system and automatically take action when events occur. You can add, copy, and configure alerts, as well as migrate Teradata Manager alerts.

You can activate alert actions that send a notification, or take some other type of action, when a metric exceeds a threshold.

After you add alert action sets in the **Alert Setup** portlet, they appear in the **Monitored Systems** portlet.

The types of alert actions you can choose are:

Send an Email

Choose a defined email address and text message. Event information, such as date, time, event name, threshold, and actual value, is automatically added to the body of the email message.

Run a Program

Choose a program to execute (.exe).

Run a BTEQ Script

Choose a BTEQ script.

Run a SQL Query

Choose a SQL query.

Notify SNMP System

Choose an SNMP trap.

You can also choose one of the following alert actions for database session alert types:

Abort Session

Abort the session for which an alert was detected. This action is only available for sessions.

Lower Priority

Set the priority of the session for which an alert was detected to the one specified in the **New priority for lower priority action** box, located under **Data Collectors** setup for **Sessions**.

Notify session user

Sends an email to notify the database session user of the alert.

Notify session proxy user

Sends an email to notify the database proxy session user of the alert.

Alert Types

You can set alerts that take place when performance or database space events occur on one or more systems.

Teradata System Alert Types

Alert Type	Description	Data Collector Required
Canary Queries	Set an alert to occur when the response time for a canary query exceeds the threshold.	Canary Queries

Alert Type	Description	Data Collector Required
Data Collectors	Set an alert to occur when any errors are encountered by the DCS for the specified data collectors.	
Database Space	Set an alert to occur when the space used on any resource, such as an AMP, exceeds the specified percentage.	Database Space
Events	Set an alert to occur when a database restarts.	Alert Request
Node	Set an alert to occur when the performance on a database reaches the specified threshold at the node level.	System Stats
Session	Set an alert to occur when the performance on a database reaches the specified threshold at the session level.	Session
System	Set an alert to occur when the performance on a database reaches the specified threshold.	System Stats
System Health	Set an alert to occur when the system health of a database reaches the specified level.	
Table Space	Set an alert to occur when one or both of the following thresholds are exceeded for the DBC.TransientJournal table: <ul style="list-style-type: none"> Space used exceeds the specified size Space usage skew percent is greater than the specified percentage 	Table Space
Vproc	Set an alert to occur when the performance on a database reaches the specified threshold at the vproc level.	System Stats

Aster System Alert Types

Alert Type	Description	Data Collector Required
Data Collectors	Set an alert to occur when any errors are encountered by the DCS for the specified data collectors.	
Node	Set an alert to occur when the performance on Aster reaches the specified threshold at the node level.	Nodes
Process	Set an alert to occur when the performance on Aster reaches the specified threshold at the process level.	Processes
System	Set an alert to occur when the performance on Aster reaches the specified threshold.	Component Stats
System Health	Set an alert to occur when the system health of Aster reaches the specified level.	

Teradata Machine Learning Engine Alert Types

Alert Type	Description	Data Collector Required
Data Collectors	Set an alert to occur when any errors are encountered by the DCS for the specified data collectors.	
Pod	Set an alert to occur when performance reaches the specified threshold at the queen or worker level.	Pods
Process	Set an alert to occur when the performance on Teradata Machine Learning Engine reaches the specified threshold at the process level.	Processes
System	Set an alert to occur when the performance on Teradata Machine Learning Engine reaches the specified threshold.	Component Stats
System Health	Set an alert to occur when the system health of Teradata Machine Learning Engine reaches the specified level.	

Hadoop System Alert Types

Alert Type	Description	Data Collector Required
Data Collectors	Set an alert to occur when any errors are encountered by the DCS for the specified data collectors.	
HDFS	Set an alert to occur when HDFS metrics on Hadoop reach the specified threshold.	HDFS
MapReduce (HDP 1.3)	Set an alert to occur when MapReduce metrics on Hadoop reach the specified threshold.	MapReduce
System	Set an alert to occur when the performance on Hadoop reaches the specified threshold.	System Stats
System Health	Set an alert to occur when the system health of Hadoop reaches the specified level.	
YARN (HDP 2.1 and later)	Set an alert to occur when the YARN metrics on Hadoop reach the specified threshold.	YARN

QueryGrid Alert Types

Alert Type	Description	Data Collector Required
Data Collectors	Set an alert to occur when any errors are encountered by the DCS for the specified data collectors.	
Manager	Set an alert to occur for issues pertaining to QueryGrid manager components.	Issues

Alert Type	Description	Data Collector Required
System	Set an alert to occur for issues pertaining QueryGrid system components.	Issues

Teradata Alert Metrics and Properties

These metrics are available for Teradata alerts. The associated property names allow you to customize alert actions in the **Alert Setup** portlet or customize the message in the **Monitored Systems** portlet.

Metrics for Node and Vproc Alert Types

Metric	Description	Property Name
AMP CPU Skew (system level)	Comparison of CPU use on the busiest AMP to the average AMP on the system	<i>ampCpuSkew</i>
AMP CPU Usage (system level)	Average percent of CPU usage of all online AMPs in the configuration on the system	<i>aveAmpCpuUsage</i>
AMP Disk Usage (system level)	Average percent of disk usage of all online AMPs in the configuration on the system	<i>aveAmpDiskUsage</i>
Available AWT	Total available AMP worker tasks	<i>availableAWTCount</i>
CIC Usage	Percent of Channel Interface Controller usage for this resource	<i>cicUse</i>
CPU Usage	Percent of CPU usage that is not idle (includes system CPU and user CPU)	<i>cpuUse</i>
Disk I/O	Number of disk I/Os	<i>diskIO</i>
Disk Out Req Avg	Average number of outstanding disk requests for this resource	<i>diskOutReqAvg</i>
Host I/O	Number of host I/Os	<i>hostIO</i>
Message Count	Messages waiting on the vproc	<i>msgCount</i>
Net A Usage	Percent of CPU usage that is not Total BYNET utilization (BYNET receiver usage)	<i>netAUse</i>
Net A Usage (system level)	Total BYNET utilization (average of the online BYNETs) on the system	<i>systemNetAUse</i>
Node CPU Skew (system level)	Comparison of CPU use on the busiest node to the average node on the system	<i>nodeCpuSkew</i>
Node CPU Usage (system level)	Average percent of CPU usage of all online nodes in the configuration on the system	<i>aveNodeCpuUsage</i>
Node Disk Usage (system level)	Average percent of disk usage of all online nodes in the configuration on the system	<i>aveNodeDiskUsage</i>

Metric	Description	Property Name
PE CPU Usage (system level)	Average percent of CPU usage of all online PEs in the configuration on the system	<i>avePeCpuUsage</i>
Status	Status of the node or vproc, where U means the resource is up and D means the resource is down Note: With the is equal to comparison operator, The existing vproc status alerts starting with U/u and D/d migrate to U and D respectively.	<i>status</i>
Swaps	Total number of swap reads and swap writes	<i>swaps</i>
Total Sessions	Number of sessions logged in to the PE	<i>sessLogCount</i>
CPU I/O Wait	Percentage of highest CPU I/O wait on node	<i>percntKernel</i>

The following metrics are not available in the dialog box, but can be entered in the message included with the alert by using the property name with the syntax *\${alertProperty}*. These properties can also be used to customize alert actions in the **Alert Setup** portlet.

Metric	Description	Property Name
ID	ID of the process	<i>proclD</i>
Node ID	ID of the node	<i>nodeId</i>
Vproc Number	Number of the vproc	<i>vprocNo</i>

Metrics for Database Space Alert Types

Metric	Description	Property Name
Current Perm %	Percentage of the total permanent disk space the database is currently using (Current Perm divided by Max Perm)	<i>spaceUsedPct</i>
Current Perm Max %	Current Perm Max * Number of Vprocs / Max Perm	<i>maxSpaceUsedPct</i>
Current Spool Space %	Percentage of total temporary spool space the database is currently using (Current Spool divided by Max Spool)	<i>currentSpoolPer</i>
Effective Perm Space %	Larger of two values: Percentage of perm space used, accounting for the configured skew limit, on the highest utilized AMP or percentage of Perm space used in the database as a whole.	<i>effectivePermSpace</i>
Peak Spool Space %	Highest percentage of total temporary spool space the database has used (Peak Spool divided by Max Spool)	<i>peakSpoolPer</i>

The following metric is not available in the dialog box, but can be entered in the message included with the alert by using the property name with the syntax `${alertProperty}`. These properties can also be used to customize alert actions in the **Alert Setup** portlet.

Metric	Description	Property Name
Database Name	Name of the database	<i>databaseName</i>

Metrics for Session Alert Types

Metric	Description	Property Name
Account	Account from which a query was submitted	<i>userAccount</i>
Active Time	Amount of time the query has been running	<i>activeTime</i>
AMP CPU	Current elapsed CPU time, in seconds, used on all AMPs by the associated session for executing requests	<i>ampCPUsec</i>
AMP I/O	Current number of logical Reads and Writes issued across all AMPs by the associated session	<i>ampIO</i>
Blocked Time	How long the query has been blocked	<i>blockedTime</i>
Connection Time	How long the session has been connected	<i>timeLoggedOn</i>
CPU delta	Total CPU usage time consumed, in seconds, since the last sample	<i>deltaCPU</i>
CPU Skew	CPU skew during the last sample	<i>cpuSkew</i>
Defer Time	Amount of time a session has been deferred by Arrival Rate Meter rules	<i>deferTime</i>
Delay Time	Amount of time a session has been delayed by utility, system throttle, workload throttle, or workload group throttle rules	<i>delayTime</i>
Delta I/O	I/O count since the last sample	<i>deltaIO</i>
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session	<i>hostId</i>
Hot AMP CPU	CPU time of the highest CPU utilized AMP during the collection interval	<i>hotAmp1CPU</i>
Impact CPU	The total CPU seconds consumed across all AMP's that are active for a session	<i>impactCPU</i>
Idle Time	How long the query has been idle	<i>idleTime</i>
Logon Source	Logon source containing the application ID and client ID	<i>logonSource</i>
Partition	Partition to which the session is connected (SQL, CONSOLE, MONITOR)	<i>partName</i>

Metric	Description	Property Name
Request AMP CPU	Current elapsed CPU time, in seconds, used on all AMPs by the current request	<i>requestAmpCPU</i>
Request AMP I/O	Current number of logical Reads and Writes issued across all AMPs by the current request	<i>requestAmpIO</i>
Responding Time	How long a completed query is taking to send (respond) spooled data back to the user	<i>respondingTime</i>
Sessions Per User	Number of sessions logged in under this username	<i>sessionsForUser</i>
Spool Space	Amount of spool space used across all AMPs by this request	<i>requestAmpSpool</i>
Utility Session	Session from one of the Teradata utilities, indicated by True and False	<i>utility</i>

The following metrics are not available in the dialog box, but can be entered in the message included with the alert by using the property name with the syntax *#{alertProperty}*. These properties can also be used to customize alert actions in the **Alert Setup** portlet.

Metric	Description	Property Name
Blocking Session Number 1	Session ID of the first blocking query	<i>blk1SessNo</i>
Blocking Session Number 2	Session ID of the second blocking query	<i>blk2SessNo</i>
Blocking Session Number 3	Session ID of the third blocking query	<i>blk3SessNo</i>
Blocking User Name 1	Name of the user running the first blocking query	<i>blk1Username</i>
Blocking User Name 2	Name of the user running the second blocking query	<i>blk2Username</i>
Blocking User Name 3	Name of the user running the third blocking query	<i>blk3Username</i>
Session Number	Number of the session	<i>sessionNo</i>
User	User of the session	<i>userName</i>

Metrics for System Alert Types

Metric	Description	Property Name
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP	<i>ampCpuSkew</i>
AMP CPU Usage	Average percent of CPU usage of all online AMPs in the configuration	<i>aveAmpCpuUsage</i>
Blocked Sessions	Total number of blocked sessions	<i>blockedSessions</i>
Net A Usage	Total BYNET utilization (average of the online BYNETs) on the system	<i>netAUse</i>

Metric	Description	Property Name
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	<i>nodeCpuSkew</i>
Node CPU Usage	Average percent of CPU usage of all online nodes in the configuration	<i>aveNodeCpuUsage</i>
Node Disk Usage	Average percent of disk usage of all online nodes in the configuration	<i>aveNodeDiskUsage</i>
PE CPU Usage	Average percent of CPU usage of all online PEs in the configuration	<i>avePeCpuUsage</i>
Total Sessions	Total number of logged in sessions	<i>totalSessions</i>

Metrics for System Health Alert Types

Metric	Description	Property Name
Health	Name of the system health state	<i>health</i>

Metrics for Table Space Alert Types

Metric	Description	Property Name
Current Perm	Amount of permanent disk space the DBC.TransientJournal table is currently using	<i>totalSize</i>
Skew	Percentage of uneven distribution of disk space used for the DBC. TransientJournal table across all AMPs	<i>skewPct</i>

Metrics for Canary Queries Alert Types

Metric	Description	Property Name
Canary Query ID	ID of the canary query	<i>canaryId</i>
Login Timeout Occurred	Timeout for the login, indicated by True or False	<i>loginTimeoutOccurred</i>
Query Name	Name of the query	<i>queryName</i>
Response Time	Number of milliseconds it took for the canary query to return a response	<i>responseTime</i>

Aster Alert Metrics and Properties

These metrics are available for Aster alerts. The associated property names allow you to customize alert actions in the **Alert Setup** portlet or customize the message in the **Monitored Systems** portlet.

Metrics for Node Alert Types

Metric	Description	Property Name
% Disk Full	Percent of used disk space on the node. Includes user data, replica data, and system data	<i>percentFull</i>
CPU	Amount of available processing power used by the node	<i>cpuUse</i>
Disk In	Rate of disk traffic into the node in bytes per second	<i>diskWrite</i>
Disk Out	Rate of disk traffic out of the node in bytes per second	<i>diskRead</i>
Memory	Memory used on the node	<i>memUse</i>
Network In	Rate of network traffic into the node in bytes per second	<i>netIn</i>
Network Out	Rate of network traffic out of the node in bytes per second	<i>netOut</i>
State	State of the node	<i>status</i>
Type	Values for node types as follows: <ul style="list-style-type: none"> • Worker nodes hold the data and process the queries • Queen nodes manage the process, including the performance of the other nodes and delegation of queries • Loader nodes transfer data in and out of virtual workers 	<i>type</i>

Metrics for Process Alert Types

Metric	Description	Property Name
Database	Name of the database on which the process ran	<i>dbName</i>
Duration	How long the process ran	<i>processExecutionTime</i>
User	Name of the user who submitted the process	<i>userName</i>

Metrics for System Alert Types

Metric	Description	Property Name
Active Processes	Number of processes with active queries	<i>activeSessions</i>
CPU	Average node CPU use	<i>cpu</i>
Memory	Available amount of memory used by node	<i>memory</i>
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	<i>nodeCpuSkew</i>
Node I/O Skew	Comparison of I/O use on the busiest node to the average node	<i>nodeIoSkew</i>
Replication Factor	Number of copies of user data	<i>replicationFactor</i>

Metrics for System Health Alert Types

Metric	Description	Property Name
Health	Name of the system health state	<i>health</i>

Teradata Machine Learning Engine Alert Metrics and Properties

These metrics are available for Teradata Machine Learning Engine alerts. The associated property names allow you to customize alert actions in the **Alert Setup** portlet or customize the message in the **Monitored Systems** portlet.

Metrics for Pod Alert Types

Metric	Description	Property Name
CPU	Amount of available processing power used	<i>cpuUse</i>
Network In	Rate of network traffic into the worker or queen in bytes per second	<i>netIn</i>
Network Out	Rate of network traffic out of the worker or queen in bytes per second	<i>netOut</i>
Type	Values as follows: <ul style="list-style-type: none"> Workers hold the data and process the queries Queens manage the process, including the performance of the workers and delegation of queries 	<i>type</i>

Metrics for Process Alert Types

Metric	Description	Property Name
Duration	How long the process ran	<i>processExecutionTime</i>
User	Name of the user who submitted the process	<i>userName</i>

Metrics for System Alert Types

Metric	Description	Property Name
Active Processes	Number of processes with active queries	<i>activeSessions</i>
CPU	Average CPU use	<i>cpu</i>
Pod CPU Skew	Comparison of CPU use on the busiest queen or worker to that of the average one	<i>podCpuSkew</i>

Metrics for System Health Alert Types

Metric	Description	Property Name
Health	Name of the system health state	<i>health</i>

Hadoop Alert Metrics and Properties

These metrics are available for Hadoop alerts. The associated property names allow you to customize alert actions in the **Alert Setup** portlet or customize the message in the **Monitored Systems** portlet.

Note:

Unless otherwise indicated, metrics are common to both HDP 2.1 and later and HDP 1.3.

Metrics for HDFS Alert Types

Metric	Description	Property Name
Blocks Corrupt	Number of blocks whose replicas are all corrupt	<i>corruptBlocks</i>
Blocks Excess	Number of blocks that exceed their target replication for the file they belong to	<i>excessBlocks</i>
Blocks Missing	Number of blocks with no replicas anywhere in the cluster	<i>missingBlocks</i>
Blocks Pending Deletion	Number of blocks waiting for deletion	<i>pendingDeletionBlocks</i>
Blocks Pending Replication	Number of blocks waiting to be replicated	<i>pendingReplicationBlocks</i>
Blocks Scheduled for Replication	Number of blocks scheduled for replication	<i>scheduledReplicationBlocks</i>
Blocks Under Replicated	Number of blocks that do not meet their target replication for the file they belong to	<i>underReplicatedBlocks</i>
Disk Capacity Used	Number of bytes of disk space currently used by HDFS	<i>capacityUsed</i>
Disk Usage	Percentage of available disk space used by HDFS	<i>hdfsDiskUsage</i>
Files + Directories	Total number of files and directories in HDFS	<i>filesTotal</i>
Total Load	Number of connections to HDFS	<i>totalLoad</i>

Metrics for YARN Alert Types (HDP 2.1 and later)

Metric	Description	Property Name
Applications Running	Number of YARN applications currently executing	<i>appsRunning</i>
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances	<i>clusterMemUsed</i>
Containers Allocated	Number of YARN containers currently allocated	<i>allocatedContainers</i>

Metrics for MapReduce Alert Types (HDP 1.3)

Metric	Description	Property Name
Jobs Running	Number of jobs currently executing on the system	<i>jobsRunning</i>
Map Tasks Running	Number of map tasks running	<i>runningMaps</i>
Map Tasks Waiting	Number of map tasks queued to run	<i>waitingMaps</i>
Reduce Tasks Running	Number of reduce tasks running	<i>runningReduces</i>
Reduce Tasks Waiting	Number of reduce tasks queued to run	<i>waitingReduces</i>

Metrics for System Alert Types

Metric	Description	Property Name
CPU Idle	Percentage of CPU time not processing any commands and the system not having an outstanding disk I/O request	<i>cpuIdle</i>
CPU Nice	Percentage of CPU time spent executing at the user level with nice priority	<i>cpuNice</i>
CPU Skew	Comparison of CPU use on the busiest node to the average node	<i>cpuSkew</i>
CPU System	Percentage of CPU time spent running kernel code	<i>cpuSystem</i>
CPU Usage	The sum of the CPU user and CPU system usage percentages.	<i>cpuUse</i>
CPU User	Percentage of CPU time spent running non-kernel code	<i>cpuUser</i>
CPU Wait I/O	Percentage of CPU time spent waiting for I/O	<i>cpuWio</i>
Disk Skew	Comparison of disk space on the fullest node to the average node	<i>diskSkew</i>
Disk Use	Percentage of disk space being used on a system	<i>diskUse</i>
Load average last 15 minutes	Average number of jobs in the job queue over the last 15 minutes	<i>loadFifteen</i>

Metric	Description	Property Name
Load average last 5 minutes	Average number of jobs in the job queue over the last 5 minutes	<i>loadFive</i>
Load average last minute	Average number of jobs in the job queue over the last minute	<i>loadOne</i>
Memory Usage	Average memory use of the system during a sample period	<i>memUse</i>
Network In	Rate of incoming network traffic in bytes per second	<i>bytesIn</i>
Network Out	Rate of outgoing network traffic in bytes per second	<i>bytesOut</i>

Metrics for System Health Alert Types

Metric	Description	Property Name
Health	Name of the system health state	<i>health</i>

Migrating Alerts

After a Teradata system has been configured in Teradata Viewpoint, all of the alerts except event alerts can be migrated from Teradata Manager to Teradata Viewpoint. Configured rates cannot be migrated. You must configure data collection rates in Teradata Viewpoint.



Alerts can only be migrated when there are no alerts in Teradata Viewpoint for a Teradata system; therefore, migrate existing Teradata Manager alerts before adding and copying alerts.





1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Alerts**.
3. From the **Alert Types** list, select an alert type.
4. Under **Alerts**, select **Migrate Alerts**.
5. Enter the log in information for an account that has permission to read data from the migrating database.
6. Select an **Authentication Mechanism** for the migrating database.
7. [Optional] For **SNMP Config**, select the SNMP configuration that you assigned to any migrated action sets containing an SNMP action.
This is required when the alert actions being migrated contain any SNMP actions. The SNMP configurations are defined in the **Alert Setup** portlet.
8. Click **Apply**.
A message appears in the dialog box with the number of alerts that were migrated.







Adding and Copying Alerts









After you configure a system in Teradata Viewpoint, you can add an alert. Available alert types differ depending on the type of system you choose.

Because you can migrate alerts only when no alerts currently exist in Teradata Viewpoint for a Teradata system, migrate existing Teradata Manager alerts before adding and copying alerts.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Alerts**.
3. From the **Alert Types** list, select an alert type.
4. From the **Alerts** list, do one of the following:
 - Select  to add an alert.
 - Select  in the row of the alert you want to copy.
 - Select the name of the alert you want to edit.
5. Enter a name for the alert.
6. [Optional] Select the **Enable alert** check box to enable the alert.
7. Under **Severity**, select a severity from the list.
8. Under **Alert Rules**, do the following:

Alert Type	Trigger
<ul style="list-style-type: none"> • Canary Queries • HDFS • MapReduce (HDP 1.3) • Node • Process • Query • System • Vproc • YARN (HDP 2.1 and later) 	<ol style="list-style-type: none"> Under Thresholds, select a scope to apply to the defined rules. <ul style="list-style-type: none"> • All requires all criteria match to trigger the alert. • Any requires any one criterion match to trigger the alert. Select a metric from the list and enter a threshold value. You can add  or remove  rules. [Optional] Enter the number of minutes in the Only trigger if alert rule(s) are met for box. The alert triggers if the rule is true for the number of minutes that you entered.
<ul style="list-style-type: none"> • Data Collectors 	<p>Select one or more data collectors from the Available box and select  to move the data collectors to the Selected box.</p> <p>The alert triggers if any errors are encountered by the DCS for the selected data collectors.</p>
<ul style="list-style-type: none"> • Database Space 	<ol style="list-style-type: none"> Under Threshold, select a metric and enter a threshold value in the box that appears. The alert triggers if the metric exceeds the value that you enter. [Optional] Select . Select one of the following to include or exclude databases when triggering the alert: <ul style="list-style-type: none"> • All databases except the selected databases below • Only the selected databases below

Alert Type	Trigger
	<p>d. If you want to filter the list of databases in the Match string box, do any of the following:</p> <ul style="list-style-type: none"> Type a value. Use the * wildcard character to match zero or more characters in the position it occupies. Type this symbol at the beginning, middle, or end of your search. This wildcard character can be used more than once in the same search. Type a wildcard statement. <p>e. If you entered values in the Match string box, do one of the following:</p> <ul style="list-style-type: none"> Select  to move the values to the Selected databases box. Select , located to the right of the box, to clear the values. <p>f. Select one or more databases from the Available databases box and select  to move the databases to the Selected databases box.</p> <p>Use Ctrl or Shift to select multiple databases.</p> <p>The alert triggers if the metric conditions are encountered for the included or excluded databases or wildcard statements.</p>
<ul style="list-style-type: none"> Events 	<p>Under Alert Rules, the Database Restart metric is selected.</p> <p>The alert triggers if the database is restarted.</p>
<ul style="list-style-type: none"> Manager 	<p>a. Under Alert Rules, select one of the following:</p> <ul style="list-style-type: none"> Critical and warning to trigger an alert when an issue has a severity level of either critical or warning. Critical only to trigger an alert when an issue has a severity level of critical. <p>b. [Optional] Enter the number of minutes in the Only trigger if alert rule(s) are met for box.</p> <p>The alert triggers if the rule is true for the number of minutes you entered.</p>
<ul style="list-style-type: none"> Session 	<p>a. Under Thresholds, use the Match list to select a scope to apply to the defined rules:</p> <ul style="list-style-type: none"> All requires all criteria match to trigger the alert. Any requires any one criterion match to trigger the alert. <p>b. Select a metric from the list and enter a threshold value.</p> <p>If you set up an alert rule whose only condition is Sessions per User, only <i>userName</i> and <i>sessionsForUser</i> metrics are available when customizing Teradata Alerts alert actions.</p> <p>c. [Optional] Add  or remove  rules.</p> <p>d. [Optional] Select  next to Users.</p> <p>e. Select one of the following to include or exclude users when triggering the alert:</p> <ul style="list-style-type: none"> All users except the selected users below Only the selected users below <p>f. If you want to filter the list of users in the Match string box, do any of the following:</p> <ul style="list-style-type: none"> Type a value. Use the * wildcard character to match zero or more characters in the position it occupies. Type this symbol at the beginning, middle, or end of

Alert Type	Trigger
	<p>your search. This wildcard character can be used more than once in the same search.</p> <ul style="list-style-type: none"> Type a wildcard statement. <p>g. If you entered values in the Match string box, do one of the following.</p> <ul style="list-style-type: none"> Select  to move the values to the Selected users box. Select , located to the right of the box, to clear the values. <p>h. Select one or more users from the Available users box and select  to move the users to the Selected users box.</p> <p>Use Ctrl or Shift to select multiple users.</p> <p>i. [Optional] Under Users, select Ignore alert action if blocking and blocker user are the same.</p> <p>Note: This checkbox is only visible when the Blocked Time alert is selected as the Thresholds option.</p> <p>j. [Optional] Select  next to Account Strings</p> <p>k. Select one of the following to include or exclude account strings when triggering the alert:</p> <ul style="list-style-type: none"> All account strings except the selected account strings below Only the selected account strings below <p>l. If you want to filter the list of account strings in the Match string box, do any of the following:</p> <ul style="list-style-type: none"> Type a value. Use the * wildcard character to match zero or more characters in the position it occupies. Type this symbol at the beginning, middle, or end of your search. This wildcard character can be used more than once in the same search. Type a wildcard statement. <p>m. If you entered values in the Match string box, do one of the following:</p> <ul style="list-style-type: none"> Select  to move the values to the Selected Account Strings box. Select , located to the right of the box, to clear the values. <p>n. Select one or more account strings from the Available Account Strings box and select  to move the account strings to the Selected Account Strings box.</p> <p>Use Ctrl or Shift to select multiple account strings.</p> <p>The alert triggers if the metric conditions are encountered for the included or excluded account strings or wildcard statements.</p>
<ul style="list-style-type: none"> System (QueryGrid only) 	<p>a. Under Alert Rules, select one of the following:</p> <ul style="list-style-type: none"> Alert on all systems in QueryGrid. Alert on selected systems in QueryGrid <p>b. If you selected Alert on selected systems in QueryGrid, select one or more systems from the Available box and select  to move the systems to the Selected box.</p> <p>c. Under Trigger for all issues of severity, select one of the following:</p> <ul style="list-style-type: none"> Critical and warning to trigger an alert when an issue has a severity level of either critical or warning.

Alert Type	Trigger
	<ul style="list-style-type: none"> • Critical only to trigger an alert when an issue has a severity level of critical. <p>d. [Optional] Select Alert only for active configuration if you want alerts issued on active versions (not on pending or previous versions) of the configuration for fabrics, connectors, links, communication policies, and networks.</p> <p>e. [Optional] Enter the number of minutes in the Only trigger if alert rule(s) are met for box. The alert triggers if the rule is true for the number of minutes that you entered.</p>
<ul style="list-style-type: none"> • System Health 	<p>a. Select a health level from the Trigger if system health is this level or worse list.</p> <p>b. [Optional] Enter the number of minutes in the Only trigger if alert rule(s) are met for box. The alert triggers if the rule is true for the number of minutes that you entered.</p>
<ul style="list-style-type: none"> • Table Space 	<p>[Optional] Under Table, do any of the following for the DBC. TransientJournal table:</p> <ul style="list-style-type: none"> • Select the Current Perm check box and enter the amount of total permanent disk space the table is currently using. • Select the Skew check box and enter the percentage of uneven distribution of disk space used for the table across all AMPs. <p>The alert triggers if the metric exceeds the value that you entered.</p>

9. Under **Alert Action**, do the following:

a. [Optional] For session alerts only, select one of the following check boxes:

- **Abort Session**

This aborts the blocked database session after performing the selected alert action.

- **Lower Priority**

This modifies the priority level for the database session while performing the selected alert action. The priority level is configured in the Session data collector under **New priority for Lower Priority action**.

- **Notify Session User**

This sends an email to notify the database session user, if database user alerts have been configured in the **Alert Setup** portlet.

- **Notify Session Proxy User**

This sends an email to notify the database session proxy user, if database user alerts have been configured in the **Alert Setup** portlet.

- If only the selected threshold is set to **Blocked Time**, then the following check-boxes appear:

- **Abort Blocked Session**

This aborts the blocked database session after performing the selected alert action.

- **Abort Blocking Session**

This aborts the blocking database session after performing the selected alert action.



- b. [Optional] For database space alerts only, select the **Notify database user** check box to send an email to notify the database user, if database user alerts have been configured in the **Alert Setup** portlet.

This option uses the database name alert property to map to an email address.

- c. Select an alert action from the **Action** list.
- d. Type a limit for **Do not run twice in ... minutes**.
The alert action does not run twice during the minutes that you entered. The alert action does not run more frequently than the data collector sample rate.
- e. [Optional] For session alerts only, select **Group email by session user** to send a single email that contains the IDs of each session matching the rule criteria.
If you do not select this check box, a separate email will be generated for every session that matches the criteria.

- 10. [Optional] Under **Message**, enter a message that appears when the alert action executes.

- 11. Click **Apply**.


If the operation is successful,  appears. If the operation fails,  appears. If you receive an error, verify that the settings are correct and try again.

Copying Alerts between Systems

You can copy all alerts, except canary query alerts, from one system to another compatible system.

Note:

Alerts with duplicate names are overwritten in the target system.

1. From the **Systems** list, select the system from which you want to copy alerts.
2. From the **Setup** list, select  next to **Alerts**.
3. Do one of the following:
 - Select a target system.
 - Select multiple target systems with **Shift** or **Ctrl**.
4. Select **Copy**.

Deleting Alerts

An alert can be deleted any time after it has been added to a system.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Alerts**.

3. From the **Alert Types** list, select an alert type.
4. From the **Alerts** list, select the alert name.
5. Select ☐ in the row of the alert you want to delete.
A confirmation message appears.
6. Click **OK**.
The alert row disappears.

Monitor Rates

After a Teradata system is configured in Viewpoint, you can set internal sample rates for database sessions, node/vproc collection, and node/vproc logging. These rates determine how often PM/API data are refreshed. Note the following when setting the monitoring rate:

- The rates are based on the numbers set in the database.
- The sample rates for node/vproc collection and node/vproc logging must be an integer divisor of 3600.
- The logging rate is independent of the sample rate.

Setting Monitor Rates

You can set internal sample rates for database sessions, system statistics, and resource usage. The sample rate you set becomes the recommended sample rate for each collector. Enabling the monitoring and logging process allows each collector to begin collecting data.

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Monitor Rates**.
3. Under **Session Sampling**, select the **Session sample rate** check box and enter the session sample rate in the box for the **Sessions** collector.
4. Under **Node Vproc Collection**, select the **Resource collection rate** check box and select the collection rate for the System Stats collector.
5. Under **Node Vproc Logging**, select the **ResUsage logging rate** check box and select the logging rate for the **Resource Usage** collector.
6. Click **Apply**.

Log Table Cleanup

You can configure log table cleanup to remove unnecessary information and reallocate space in the database. After you enable the cleanup process for a table and configure the retention period, you must set the time for the cleanup schedule. The cleanup process permanently deletes data and does not archive it.

Configuring Log Table Cleanup

1. From the **Systems** list, select the name of the system you want to update.

2. From the **Setup** list, select **Log Table Cleanup**.
3. From the **Tables** list, select a table.
4. [Optional] Under **Details**, do one of the following:
 - Select the **Enable table cleanup** check box to activate cleanup and include the table in the cleanup schedule.
 - Clear the **Enable table cleanup** check box to disable cleanup and exclude the table from the cleanup schedule.
5. For **Login**, select the user you want to run the cleanup task.
6. For **Cleanup After**, set the retention period using a number and calendar value.
7. Click **Apply**.

Cleanup Tables

The following cleanup processes delete the associated system log tables.

Cleanup Process Name	System Log Table Name
Access Log	<ul style="list-style-type: none"> • DBC.AccLogTbl
DBQ Log	<ul style="list-style-type: none"> • DBC.DBQLExplainTbl • DBC.DBQLObjTbl • DBC.DBQLogTbl • DBC.DBQLParamTbl (Teradata Database 15.0 and later) • DBC.DBQLStepTbl • DBC.DBQLSqlTbl • DBC.DBQLSummaryTbl • DBC.DBQLUtilityTbl (Teradata Database 15.0 and later) • DBC.DBQLXMLTbl
Event Log	<ul style="list-style-type: none"> • DBC.EventLog
Resource Usage	<ul style="list-style-type: none"> • DBC.ResUsageSpma • DBC.ResUsageSvpr • DBC.ResUsageShst • DBC.ResUsageIpma • DBC.ResUsageIvpr • DBC.ResUsageScpu • DBC.ResUsageSldv • DBC.ResUsageSawt • DBC.ResUsageSps • DBC.ResUsageSpdsk • DBC.ResUsageSvdsdsk
SW Event Log	<ul style="list-style-type: none"> • DBC.SW_Event_Log
TDStats	<ul style="list-style-type: none"> • TDStats.AutomateHistoryTbl (Teradata Database 14.10 and later)

Cleanup Process Name	System Log Table Name
	<ul style="list-style-type: none"> • TDStats.AnalyzeHistoryTbl (Teradata Database 14.10 and later) • TDStats.CommandsHistoryTbl (Teradata Database 14.10 and later)
TDWM	<ul style="list-style-type: none"> • DBC.TDWMSummaryLog • DBC.TDWMEventLog • DBC.TDWMEExceptionLog

Cleanup Schedule

After the log table cleanup has been configured for the database, you can set a schedule. The schedule sets the time to delete all log tables that have been enabled for the cleanup process. The cleanup process permanently deletes data and does not archive it.

Scheduling Log Table Cleanup

1. From the **Systems** list, select the name of the system you want to update.
2. From the **Setup** list, select **Cleanup Schedule**.
3. [Optional] Under **Cleanup Schedule Details**, do one of the following:
 - Select the **Enable cleanup scheduler** check box to activate the clean up schedule.
 - Clear the **Enable cleanup scheduler** check box to disable the clean up schedule.
4. For **Login**, select the user you want to collect the current time and database version.
5. If you are collecting data with the **Performance Data Collection** portlet, you can select **Clean up after all Performance Data Collection daily jobs have completed successfully** to schedule cleanup to occur after the jobs have completed, ensuring that this data is archived before the cleanup occurs.
This option appears only if you have the **Performance Data Collection** portlet activated.
6. For **Cleanup Time**, enter a time at which the cleanup task is run on the Viewpoint server by the DCS. **Current Teradata Time** indicates the current setting or **unknown** indicates a connection to the database cannot be made.
7. Click **Apply**.

Viewpoint DB Disk Usage

View disk usage of the Viewpoint database under /data/viewpointdb for all systems over time, and the distribution of data on the Teradata Viewpoint. The sparkline and graph provide different aspects of the data.

Graphics	Description
Sparkline	Shows 3 months of data. Information balloons show detailed data about a point in time.

Graphics	Description
Graph	Shows how the space used by the Teradata Viewpoint server is divided among systems and their data collectors. Hover over or click the system boxes in the graph to see detailed disk space and data collector information. If a system contains a region that is not associated with a collector, the information balloon displays Unknown and the amount of data space.

Viewing Viewpoint DB Disk Usage

Your database size is always larger than the number specified in the collector for retention by size. Data cleanup only deletes whole tables that exceed this number. A 1024 byte kilobyte is used, so 1 GB is actually 1,073,741,824 bytes.

1. Select the **Viewpoint DB Disk usage** link in the **Monitored Systems** portlet heading.
2. [Optional] Hover over the sparkline and graph to activate the information balloons and see detailed information.
3. [Optional] Select a system box in the graph to view the disk usage for each data collector on that system.
4. [Optional] Select the collector name in the graph to see the corresponding data collector configuration settings.

Setting a Threshold for the Total Disk Usage Alert

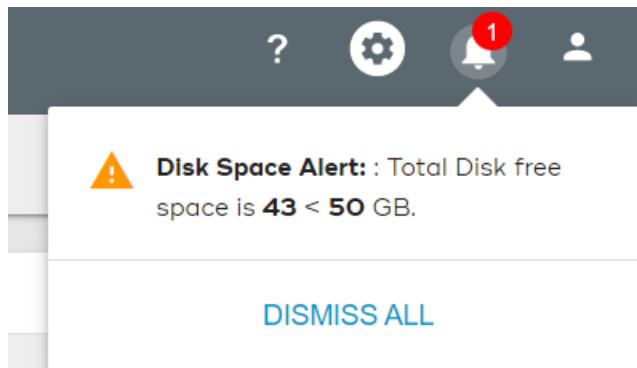
You can set a threshold that sends an administrator alert when total disk usage in the DCS exceeds the threshold.

By default, the administrator alert is set to **Administrator Default** in the **General** portlet and sends an email to all users in the Administrator role. The recipients continue to receive alert emails for every hour until disk space usage no longer exceeds the specified threshold.

Instead of using the **Administrator Default** setting, you can select a custom action set that was created in the **Alert Setup** portlet.

When you edit the alert, the average daily space usage for the last 30 days displays at the bottom and is intended to help guide you in setting the threshold.

1. Select the **Viewpoint DB Disk usage** link in the **Monitored Systems** portlet heading.
2. Select **Edit Alert**.
3. Type a number of, 1 or more GB.
The number must be a whole number and cannot be a negative number.
4. Click **OK**.
If the available disk space in /data mount is lesser that the set thresold, a notification appears as following:



Reclaiming Disk Space from Deleted Table Data

You can reclaim storage space on the Teradata Viewpoint server from deleted table data on demand or on a schedule. When reclaim is running, Viewpoint locks tables on the Postgres database.

1. Select the **Viewpoint DB Disk usage** link in the **Monitored Systems** portlet heading.
2. Select **Advanced Settings**, then select one of the following:

Option	Description
Reclaim disk space on demand	a. Select Reclaim Now .
Schedule cleanup	a. Select the Schedule cleanup check box, then select the schedule options. b. Select Save .

3. [Optional] Select **Refresh Status** to update status of the reclaim process.

Configuring Cleanup of Space Usage Collector Tables

1. Open `/opt/teradata/viewpoint/dcs/config/dcs.properties`.
2. Set `spacecollector.cleanUp.enabled=true` to clean up the disk usage tables. The default setting is false.

```
#spacecollector.cleanUp.enabled=true
#The interval time for each cleanup job run. Default value is 1 week.
#spacecollector.cleanUpRate=604800000
#Parameter for data retention in terms of cache database size. Allowed
size units (KB/MB/GB)
#spacecollector.retentionSize=10GB
#Parameter for data retention in terms of duration. Allowed time units
```

```
(DAY/DAYS/WEEK/WEEKS/MONTH/MONTHS/YEAR/YEARS)
#spacecollector.retentionTime=10MONTHS
```

3. Save the file and restart the dcs service:
`service dcs restart`

Portlet Library

Use the **Portlet Library** to enable or disable portlets globally. Even if a portlet is enabled for a role, it must be enabled in **Portlet Library** for a user in the role to have access to it.

The **Portlets** tab displays installed portlets, grouped by category, and provides the following information:

- Portlet name
- Version number
- Publisher name
- Bundle name
- Installation date
- Portlet description

The **Shared Portlets** tab displays information about shared portlets. A shared portlet is a user-defined version of a portlet. The **Parent Portlet** column identifies the original portlet before it was customized as a shared portlet. Portlet names and descriptions can be edited. Portlets can be deleted. Shared portlet permissions can be edited using **Roles Manager**.

The **Web Services** tab displays information about installed web services and provides the following information:

- Version number
- Publisher name
- Bundle name
- Installation date
- Web service description

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), the **User Management** web service will not be listed.

Enabling or Disabling Portlets

1. In the **Portlets** view, do one of the following:
 - Enable the portlet by selecting the corresponding check box.
 - Disable the portlet by clearing the corresponding check box.


Portlets in the **Admin** category cannot be disabled. They are included in the view to give you information, such as the version number.

2. Click **Apply**.


If you disable a portlet, it is immediately removed from the portlet selection list. Any portlet that was on a user portal page is no longer available for users to access.

Editing a Shared Portlet

You can change the name and description of a shared portlet.

1. From the **Portlets** view, select the **Shared Portlets** tab.
2. Select .
3. [Optional] Enter a new name up to 25 characters.
You can use spaces, alphanumeric characters, and underscores (_).
4. [Optional] Enter a description.
5. Click **OK**.
Your changes appear in the **Shared Portlets** tab.

Deleting a Shared Portlet

1. From the **Portlets** view, select the **Shared Portlets** tab.
2. Select  that corresponds to the shared portlet you want to delete.
A confirmation message appears.
3. Select **Delete**.
The portlet name no longer appears on the **Add Content** screen in the portal, and the portlet is no longer available to Teradata Viewpoint users. Existing instances of a shared portlet are removed from the user's page.

Granting Web Service Access and Permissions

Note:

Disabled web services are not available for any role, even if permissions have been enabled for roles in **Roles Manager**.

1. From the **Portlet Library** view, select **Web Services**.
2. Select the check box of a web service to enable it.
3. Click **Apply**.

Query Group Setup

The **Query Group Setup** portlet allows the Teradata Viewpoint Administrator to manage the sets of queries available to users in the **Query Groups** and **Application Queries** portlets and define new applications that can appear in the **Query Log** portlet.

You can define query groups based on a set of criteria, assign users who can access the group, and delete query groups. In addition, query groups are automatically generated when an *app* or *applicationName* query band is found in DBQL by the Query Log data collector. The criteria for an automatically generated query group cannot be changed, but the access rights can.


Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure displayed data is current in the portlet.

Query Group Setup <small>Manage the sets of queries available in Query Groups and Application Queries.</small>				
Systems	Groups +			
Dev1		GROUP NAME	DESCRIPTION	PERMISSIONS
Prod1	▼	DCSIntTestSession		0 users and 0 roles
Dev2	▼	DBCAccountStgiring	Session with DBC	0 users and 1 roles
Dev3	▼	VP-AlertRequest		1 users and 0 roles
				ORIGIN
				Auto-generated
				User-defined
				User-defined

Creating a Query Group

Create a query group in the **Query Group Setup** portlet to allow users to view their most relevant queries in the **Query Groups**. You create a query group to filter the queries that are displayed and limit the users that have access to the group.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Access** tab is current.

1. Select a system from the **Systems** list.
A system cannot be changed after you create a query group.
2. Select the **Add Group**  button.
3. Enter a **Query group name**, up to 30 characters.
4. [Optional] Enter a **Description**, up to 256 characters.
5. For the **This group contains queries with attributes that match at least one value in** field on the **Criteria** tab, do one of the following:
 - Select **Every criteria below** to match at least one value in each defined sub-tab. Criteria on each sub-tab are joined by AND.
 - Select **Any criteria below** to match at least one value in one or more of the defined sub-tabs. Criteria on each sub-tab are joined by OR.

Sub-Tabs

Query Group Setup Manage the sets of queries available in Query Groups and Application Queries.

Group Details

Edit the query group below. To save this query group, specify a group name and at least one criteria and value.

Query group name: *
 Origin:
 Description:

Criteria **Access**

System Name: Dev1 This group contains queries with attributes that match at least one value in:

Account String Criteria
 Query Band Criteria
 Username Criteria
Workload Name Criteria

Select from available workloads or add a wildcard string.

Match string: (wildcard = *)


Available:



- H-WD
- L-WD
- M-WD
- T-WD
- WD-Default
- WD1
- WD2

Selected:

6. Select a criterion from the **Add Criteria** list.
7. Select values for whichever criteria you chose:

Criteria	Select values by doing any of the following:
Account String Username Workload Name	<ul style="list-style-type: none"> • In the Match string box, enter a value to reduce the number of available values. • Select a value from the Available box then select to move it to the Selected box. • In the Match string box, enter a value to create custom criteria, then select to move it to the Selected box. The value can contain ? to match exactly one character in the position it occupies, * to match zero or more characters, or = at the beginning of the string to match alpha and numeric characters literally.
Query Band	<ul style="list-style-type: none"> • Select a Query band name from the list and a value from the Available Values box, then select to move them to the Selected box.

Criteria	Select values by doing any of the following:
	<ul style="list-style-type: none"> Enter a Custom Query band name and value pair, then select  to move them to the Selected box.

- In the **Access** tab, designate the users who can view the group by doing any of the following:
 - In the **User List Filter** box, enter a value to reduce the list of available users. The value can contain ? to match exactly one character in the position it occupies, * to match zero or more characters, or = at the beginning of the string to match alpha and numeric characters literally.
 - In the **Available Users** box, select one or more Viewpoint users and click  to move them to the **Selected Users** box.
 - In the **Available Roles** box, select one or more Viewpoint roles and click  to move them to the **Selected Roles** box.
- [Optional] Select **Reset** to clear the query group name, description, and the selections on the **Criteria** and **Access** tabs.
- Click **Apply**.

Viewing Query Groups


Use the **Query Group Setup** portlet to view the query groups that have been created for enabled systems. If a system is disabled in the **Monitored Systems** portlet, it does not appear in the **Query Group Setup** portlet. The defined criteria is retained.

- Select a system to view query groups related to that system.
- [Optional] Select a row in the query group list to see the details for that group.

Editing a Query Group

Edit a query group when criteria or users require updating or if a user changes roles.

If the query group is auto-generated, you can only specify which users can view the query group.

- Select the system associated with the query group.
- Select a row in the list for the query group you want to edit.
- [Optional] Enter a new name or change the description.
- If the query group is user-defined, change the criteria by doing one of the following in the **Criteria** tab:
 - Select **Every criteria below** to match at least one value in each defined sub-tab. Criteria on each sub-tab are joined by AND.
 - Select **Any criteria below** to match at least one value in one or more of the defined sub-tabs. Criteria on each sub-tab are joined by OR.
- For each criteria you select, do any of the following:
 - Edit the criteria.
 - Select  to delete the criteria and associated values.
 - Select unnecessary or obsolete criteria from the **Selected** box then select **Remove**.

6. In the **Access** tab, edit the list of Viewpoint users and roles that can view the query group.
Use the **User List Filter** to reduce the list of available users.
7. [Optional] Select **Reset** to revert the query group name and description, and any changes made in the **Criteria** or **Access** tabs to the previously applied values.
8. Click **Apply**.

Deleting a Query Group

Delete a query group when it is no longer useful.

1. Select a system.
2. Select ☐ on the row of the query group you want to delete then select **Delete**.

Roles Manager

The **Roles Manager** portlet allows the Teradata Viewpoint Administrator to assign permissions efficiently by creating classes of users called roles.

The Teradata Viewpoint Administrator can perform the following tasks:

- Add and configure new roles
- Edit the configuration and settings of existing and default roles
- Copy roles, saving time in creating new roles
- Enable or disable portlets for a role
- Delete roles that are no longer needed
- Configure Secure Zones for a database system
- Grant metric and portlet permissions for a role
- Configure default portlet settings
- Grant a role access to portlets, **Viewpoint Dashboard**, and **Viewpoint Mobile**
- Grant web services access

Note:

- If a user with role VP_User_Manager accesses the Roles Manager portlet, the user can only view or modify their role and roles created by them.
 - If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), the **User Management** web service is not listed.
-

Roles

Teradata Viewpoint includes preconfigured roles that cannot be removed from Teradata Viewpoint. Change the following roles according to your needs:

Administrator

This role has all permissions and can be assigned to any account. It is recommended that this role be used only by the Teradata Viewpoint Administrator.

User

This role is assigned to every Teradata Viewpoint user. It is recommended that this role be set with no permissions.

TD_WebServices

This role has all permissions necessary to access enabled web services.

VP_User_Manager

This role has all permissions (excluding Bar Portlet and Bar Operations) and can be assigned to any account. It is recommended that this role be used only by the Teradata Viewpoint Administrator who must not have access to Bar Portlet and Bar Operations.

Teradata recommends that you copy **Administrator** and **User** to configure new roles with partial permissions that are appropriate to all users in that role. Each role you create controls access to specific systems, portlets, metrics, settings, and permissions in portlets.

Example Roles

The following examples show how roles can be defined to manage portlet usage and control permissions within Teradata Viewpoint.

AppDBA

The AppDBA role is defined for users that are application DBAs. The rewind feature and some portlets are enabled for this role so users can review graphs in the **Productivity** portlet.

AcctDept

The AcctDept role is defined for all users in the Accounting Department. Some portlets are enabled for this role. These users are granted permissions so they can edit all calendar events in the **Viewpoint Calendar** portlet and access the **My Queries** portlet to monitor their queries.

OpDBA

The OpDBA role is defined for users that are operations DBAs. The rewind feature and all portlets are enabled for this role so users can monitor performance trends. Users in the OpDBA role access the **Workload Monitor** portlet to track database request arrivals and completions.

Roles Manager View

The **Roles Manager** view lets the Teradata Viewpoint Administrator manage roles, including assigning users to roles, and granting permissions to roles. Customize a new role using the following tabs:

General

Manage a role, rewind, **Viewpoint Dashboard**, **Viewpoint Mobile**, and notifications for a role. The notifications appear in the Teradata Viewpoint portal.

Systems

Manage the systems for a role, select the metrics to display for each system, and configure Secure Zones for a database system.

Note:

Systems must also be enabled in **Monitored Systems**.

Portlets

Manage portlets for a role, select permissions, and configure default settings per portlet for each role. Provide permissions to users in this role to modify their own settings and share customized versions of the portlet with other users. Configure which portlet features the role can access.

Note:

Portlets must also be enabled in **Portlet Library**.

Web Services

Enable or disable access to web services for a role and select permissions.

Note:

Web services must also be enabled in **Portlet Library**.

Users

Manage the list of users assigned to the selected role.

Notifications

Teradata Viewpoint uses notifications to communicate the status of Viewpoint operations. Notifications are displayed in the Teradata Viewpoint portal for 72 hours from the last status change.

You can enable notifications to provide status of cluster setup, restore, and upgrade data migration.

- Cluster setup notifications may display when the Teradata System Administrator sets up or recovers a Teradata Viewpoint cluster after a promotion.
- Restore notifications may display when a Viewpoint server has been restored from a backup.
- Data migration notifications may display after a major upgrade when data is migrated to a new format. Migration begins with the newest data.

You can enable notifications for any role. The User role has restore and upgrade data migration notifications enabled by default. The Administrator role has restore and upgrade data migration as well as cluster setup notifications enabled by default. These notifications are not enabled by default for any other roles.

The cluster setup notification displays information about the following cluster setup status:

- Percentage of cluster setup completed
- If the setup failed or completed

The restore or data migration notifications displays information about the following restore or migration status:

- Percentage of restore or migration completed
- Earliest available date for restored or migrated data
- If the restore or migration failed or completed

Adding a Role

You can add a role and configure the settings for users assigned to the role.

1. From the **Roles Manager** view, click **Add Role**.
2. Enter a name for the role, up to 30 characters.
You can use spaces, alphanumeric characters, and underscores (_).
3. [Optional] Select **Enable role** to activate the role for use in Teradata Viewpoint.
4. Enter a description, up to 255 characters.
5. [Optional] Under **Portal Permissions**, select **Enable rewind** to grant permission for this role to use the rewind feature, when available.
6. [Optional] Specify notifications that appear in the Teradata Viewpoint portal for users assigned to the role.
For more information, see [Notifications](#).
7. Click **Apply**.
8. [Optional] Specify additional role settings on each of the tabs.

Tabs	Description
Systems	Enable or disable the systems available to the role, select the metrics to display for each system, and configure Secure Zones. For more information, see Secure Zones .
Portlets	<p>Enable or disable portlets for a role, select permissions, and configure default settings for each portlet.</p> <p>Control whether users in this role can modify their own settings and share customized versions of the portlet with other users, and configure which portlet features the role can access.</p> <p>Note: Portlets must also be enabled in Portlet Library.</p>

Tabs	Description
Web Services	Enable or disable access to each web service for this role and select permissions.
Users	Specify users that are assigned to the role.

Note:

If a user with role 'VP_User_Manager' adds another role, the user cannot assign Bar Operations, Bar Setup, and Shared Bar Portlets to the new role.

Editing a Role

You can edit a role, but you cannot change the name of an existing role. To change the name, you must either add a new role or copy the role, and then change the role name.

1. From the **Roles Manager** view, select a role from the list.
The **General** tab appears.
2. [Optional] Select **Enable role** to activate the role for use in Teradata Viewpoint.
3. [Optional] Enter a description, up to 255 characters.
4. [Optional] Under **Portal Permissions**, select **Enable rewind** to grant permission for this role to use the rewind feature, when available.
5. [Optional] Specify notifications that appear in the Teradata Viewpoint portal for users assigned to the role.
For more information, see [Notifications](#).
6. Select **Apply**

Note:

A user with the role "VP_User_Manager" can only view the user role and the roles created by the user. If a user with the role "VP_User_Manager" edits a role, the user cannot assign Bar Operations, Bar Setup, and Shared Bar Portlets to the role.

7. [Optional] Specify additional role settings on each of the tabs.


Tabs	Description
Systems	Enable or disable the systems available to the role, select the metrics to display for each system, and configure Secure Zones. For more information, see Secure Zones .
Portlets	Enable or disable portlets for a role, select permissions, and configure default settings for each portlet. Control whether users in this role can modify their own settings and share customized versions of the portlet with other users, and configure which portlet features the role can access.

Tabs	Description
	Note: Portlets must also be enabled in Portlet Library .
Web Services	Enable or disable access to each web service for this role and select permissions.
Users	Specify users that are assigned to the role.

Copying a Role

When a role is copied, all settings on all tabs are inherited from the original role, except for the role name.

When the Administrator role is copied, all information on all tabs is copied. However, the copied role does not have administrative privileges and does not have access to administrative portlets from the portal.

- To copy a role, do any of the following:
 - From the **Roles Manager** view, select  next to the role you want to copy and select **Copy Role**.
 - From the **General** tab, select **Copy Role**.
- [Optional] Enter a name for the role, up to 25 characters.
You can use spaces, alphanumeric characters, and underscores (_).
The name from the original copied role is prefaced with *Copied from* and remains unless you change it.
- [Optional] Select **Enable role** to activate the role for use in Teradata Viewpoint.
- [Optional] Enter a description, up to 255 characters.
The description from the original copied role remains unless you change it.
- [Optional] Under **Portal Permissions**, select **Enable rewind** to grant permission for this role to use the rewind feature, when available.
- [Optional] Specify notifications that appear in the Teradata Viewpoint portal for users assigned to the role.
For more information, see [Notifications](#).
- Select **Apply**.
- [Optional] Specify additional role settings on each of the tabs.

Tabs	Description
Systems	Enable or disable the systems available to the role, select the metrics to display for each system, and configure Secure Zones. For more information, see Secure Zones .
Portlets	<p>Enable or disable portlets for a role, select permissions, and configure default settings for each portlet.</p> <p>Control whether users in this role can modify their own settings and share customized versions of the portlet with other users, and configure which portlet features the role can access.</p> <p>Note: Portlets must also be enabled in Portlet Library.</p>


Tabs	Description
Web Services	Enable or disable access to each web service for this role and select permissions.
Users	Specify users that are assigned to the role.

Note:

If a user with the role “VP_User_Manager” copies a role, the user cannot assign Bar Operations, Bar Setup, and Shared Bar Portlets to the role.

Deleting a Role

You can delete a role. A role cannot be restored once deleted. This action affects all users assigned to the role.

- To delete a role, do any of the following:
 - From the **Roles Manager** view, select  next to the role you want to copy and select **Delete**.
 - From the **General** tab, click **Delete**.

A confirmation message appears.

- Select **Delete**.
The role is deleted.


Note:

A user with the role “VP_User_Manager” can only delete roles created by the user.

Assigning Users to a Role


Search for users and assign them to a role.

If auto-provisioning is in use, a default role is assigned to new users the first time they log in. The default role is typically *User*.

- From the **Roles Manager** view, select a role from the list.
- Select the **Users** tab.
- [Optional] Search for a user:
 - Select a filter from the list.
 - Enter the search criteria in the **contains** field.
 - Select **Find**.
The **Available Users** list displays only users meeting the search criteria.
- Select one or more users from the left pane.
Press **Ctrl** while selecting additional names to select multiple users.
- Select .

6. Click **Apply**.

Removing Users from a Role

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Users** tab.
3. Select one or more users from the **Users in Role** list.
You can select multiple users by holding the **Ctrl** key and clicking additional names.
4. Select .
5. Click **Apply**.

Secure Zones

Secure Zones, available in Teradata Database 15.10 and later, allows you to create one or more separate and exclusive database hierarchies, called *zones*, within a single Teradata system. A database user with the required database privileges can create and administer Secure Zones on a Teradata system, outside of Viewpoint. For more information, see *Teradata Vantage™ - Analytics Database Security Administration*.

In Viewpoint, user access to monitoring data from Secure Zones is configured using Viewpoint roles. The Teradata Viewpoint Administrator can grant access to a Secure Zone for a role so data from that Secure Zone is available to users in the role. A role can be given access to view data from multiple zones on one or more systems.

Monitoring data from Secure Zones is restricted in the following portlets:

- Application Queries
- Lock Viewer
- Query Groups
- Query Log
- Query Monitor
- Query Spotlight
- Space Usage

For the **Viewpoint Dashboard** and **Viewpoint Mobile**, the **Queries** view inherits the Secure Zones that are set for the **Query Monitor** portlet.

N/A displays if the system does not support Secure Zones or if the system supports Secure Zones, but no zones have been configured.

If there are zones and the role has been granted access to specific zones, and the Dictionary collector has been enabled, the **Secure Zones** button displays x of y, representing the number of zones enabled out of the total possible.

Configuring Secure Zones

1. From the **Roles Manager** view, select a role from the list.

2. Select the **Systems** tab.
3. Select the **Enable System** check box to enable a system.
After it is enabled, the system starts collecting data. Systems must also be enabled in **Monitored Systems**.
Metrics are enabled or disabled per system and not by Secure Zone.
4. Select the **x of y** button under the system for which you want to enable the Secure Zone.
5. Select one or more secure zones from the **Available zones** box then select **>** to move them to the **Allow access to these zones** box.
6. Click **OK**.
The **Secure Zones** button displays x of y, representing the number of zones enabled out of the total possible.
7. Click **Apply**.

Viewpoint Dashboard Permissions

The Teradata Viewpoint Administrator can set permissions for a role to access the **Viewpoint Dashboard**.

All permissions for the **Viewpoint Dashboard** inherit portlet permissions so the role must be granted portlet access as described as following, before the users in the role can begin using the specified views in the **Viewpoint Dashboard**.

Portlets	Views
Alert Viewer	Alerts
Hadoop Services	Hadoop Services
Metrics Analysis	Trends
Query Monitor	Queries
System Health	System Health
Workload Monitor	Workloads

In addition, the Teradata Viewpoint Administrator may need to grant access to portlet views. For more information, see [Portlet Settings and Permissions](#).

Granting Access to Viewpoint Dashboard for a Role

1. From the **Roles Manager** view, select a role from the list.
The **General** tab appears.
2. Under **Viewpoint Dashboard**, select **Enable Dashboard** to grant permission for this role to use the **Viewpoint Dashboard**.
3. Click **Apply**.
4. If this same role has not already been granted portlet access, click the **Portlets** tab and enable access to the following portlets:
 - **Alert Viewer**

- **Hadoop Services**
- **Metrics Analysis**
- **Query Monitor**
- **System Health**
- **Workload Monitor**

Viewpoint Mobile Permissions

The Teradata Viewpoint Administrator can set permissions for a role to access **Viewpoint Mobile**.

All permissions for **Viewpoint Mobile** inherit portlet permissions so the role must be granted portlet access as described as following, before the users in the role can begin using the specified views in **Viewpoint Mobile**.

Portlets	Views
Alert Viewer	Alerts
Metrics Analysis	Trends
Query Monitor	Queries
System Health	System Health
SQL Scratchpad	SQL Mobile

In addition, the Teradata Viewpoint Administrator may need to grant access to portlet views. For more information, see [Portlet Settings and Permissions](#).

Granting Access to Viewpoint Mobile for a Role

1. From the **Roles Manager** view, select a role from the list.
The **General** tab appears.
2. Under **Viewpoint Mobile**, select **Enable Mobile** to grant permission for this role to use **Viewpoint Mobile**.
3. Click **Apply**.
4. If this same role has not already been granted portlet access, click the **Portlets** tab and enable access to the following portlets:
 - **Alert Viewer**
 - **Metrics Analysis**
 - **Query Monitor**
 - **System Health**
 - **SQL Scratchpad**

Portlet Settings and Permissions

The Teradata Viewpoint Administrator can set permissions for a role to access portlets, settings, and tabs in the **Settings** view. Higher-level permissions override lower level permissions.

Portlet Access

Portlet access is the highest permission level for portlet access and gives this role access to the entire portlet. If granting a role portlet access in **Roles Manager**, you must also enable the portlet in **Portlet Library**.

Settings Access

Settings access is the highest permission level for settings and gives this role access to change all settings in this portlet.

Tab Access

Tab access is the lowest permission level for settings and denies this role access to tabs in the **Settings** view. For portlets that do not contain tabs, you can deny the role from modifying default portlet settings shown in the **Settings** view in the portlet.


Granting Access to a Portlet for a Role

To grant access to a portlet for a role, you must enable the portlet using **Portlet Library**.

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Portlets** tab.
3. Select the check box to the left of the portlet name.
4. Click **Apply**.

The role can access the portlet from the **Add Content** screen after you enable the portlet using **Portlet Library**.


Granting Access to the Settings View for a Role

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Portlets** tab.
3. Select  for the particular portlet in which you want the role to access the **Settings** view.
4. Select the **Enable portlet settings** check box.
The **Enable portlet settings** check box is only available for portlets that have a **Settings** view.
5. Click **OK**.
6. From the **Portlets** tab, make sure the portlet is enabled.
7. Click **Apply**.

The role can access the **Settings** view for the portlet.

Denying Access to a Settings View Tab for a Role


You can set defaults in the **Settings** view that users cannot to see or change. This allows you to hide tabs in the **Settings** view from the user. For portlets that do not contain tabs, this denies a role from modifying default portlet settings shown in the **Settings** view in the portlet.

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Portlets** tab.
3. Select  for the particular portlet in which you want to deny the role tab access in the **Settings** view.
4. Select the **Do not allow users in this role to view or change these values** check box.
5. Click **OK**.
6. Click **Apply**.

Configuring Default Portlet Settings

You can configure default portlet settings for a role. The settings you configure become the defaults in the **Settings** view for each portlet.

You can disable user settings if you do not want the **Settings** view to be accessible for a portlet. By selecting the **Do not allow users in this role to view or change these values** check box, you can prevent users from modifying the default portlet settings you configure.

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Portlets** tab.
3. Select  for the portlet you want to configure.
4. Modify the settings.
5. Click **OK**.

Granting Metric Permissions for a Role

You can grant permissions for metrics for a role, enabling users in a role to see metrics on a system.


The metric permissions granted to a role apply to only the following portlets when you enable a system:

- **Metric Heatmap** (Teradata, Aster, Hadoop)
- **Metrics Analysis** (Teradata, Aster, Hadoop)
- **Today's Statistics** (Teradata)

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Systems** tab.
3. Select the **Enable System** check box to enable the system where the metrics reside.
4. Select the check box for each metric on each system you want users in the selected role to see.
5. Click **Apply**.

Granting Portlet Permissions for a Role

You can enable specific features of a portlet by system and role.

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Portlets** tab.
3. Select the check box to the left of the portlet name to enable the portlet.
4. Select  for the portlet you want to set permissions.
5. Select the permissions you want to grant to the selected role for each available system by doing any of the following:
 - Select the **Enable portlet settings** check box to allow users to access and configure the **Settings** view in the portlet.

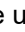
Note:

This check box is only available for portlets that have a **Settings** view.

- Select the **Share portlet** check box to allow users to share customized versions of the portlet with other users.
 - Select the system type from the list.
 - Select the **Select all** check box to select or clear all permissions for that particular system.
 - Click the permission name to select or clear that particular permission for all systems.
 - Select a check box to select or clear a single permission for a single system.
6. Click **OK**.

Portlet Permissions

The Teradata Viewpoint Administrator can grant access to portlet views using the **Portlets** tab.

Permission Name	Description	Dependency	Portlets Displaying Setting
Admin for All Workflows	Enables a non-admin role to edit and manage all workflows		MAPS Manager
Enable portlet settings	Enables role to access the Settings view in this portlet.		The Enable portlet settings check box is only available for portlets that have a Settings view.
Share portlet	Enables role to share customized versions of the portlet with other users. After this setting is enabled, the user can click  in the portlet and select Share .		The Share portlet check box is only available for portlets that can be shared.

Permission Name	Description	Dependency	Portlets Displaying Setting
Abort Host	Enables role access to the Abort Host console utility		Remote Console
Abort Job	Enables role to abort jobs that are running	• View System	Stats Manager
Abort Process	Enables role access to the <input type="checkbox"/> list in the summary view to abort non-idle processes that are running on an Aster system	• View Summary	Query Monitor
Abort Query	Enables role access to the <input type="checkbox"/> list in the summary view to abort queries	• View Summary	My Queries, Query Groups, Query Monitor, Workload Monitor
Add Space	Enables role to reallocate permanent disk space from one database to another		Space Usage
Allocation Grid	Enables role access to the Cylinder Allocation by Temperature and Grade view		TVS Monitor
Approve Recommendations	Enables role to approve or reject recommendations	• View System	Stats Manager
Automate	Enables role to automate and deautomate statistics	• View System	Stats Manager
BAR Admin	Enables role to administrate the BAR Operations portlet		Bar Operations
Change Priority	Enables role access to the <input type="checkbox"/> list in the summary view allowing role to change the priority of a query or session		My Queries, Query Groups, Query Monitor, Workload Monitor
Change Workload	Enables role access to the <input type="checkbox"/> list in the summary view to change the workload of a query or session		My Queries, Query Groups, Query Monitor, Workload Monitor
Check Table	Enables role access to the Check Table console utility		Remote Console
Clear Alert	Enables role to clear an alert for a component, process, or resource		Unity
Clear Operation	Enables role to clear Unity operation details		Unity
Configure	Enables role access to the Configure console utility		Remote Console

Permission Name	Description	Dependency	Portlets Displaying Setting
Create Event	Enables role to create calendar events		Viewpoint Calendar
Current Distribution	Enables role access to the Cylinder Distribution by Grade view		TVS Monitor
DBS Control	Enables role access to the DBS Control console utility		Remote Console
Edit Any Event	Enables role to edit any events, including events created by other users		Viewpoint Calendar
Edit Collect Settings	Enables role to set collect settings and user importance	• View System	Stats Manager
Edit Configurations	Enables role to edit the QueryGrid configuration, run a diagnostic check, run a bandwidth test and delete active issues		QueryGrid
Edit Job	Enables role to create, edit, and delete jobs	• View System	Stats Manager, Performance Data Collection
Edit Own Event	Enables role to only edit their own calendar events		Viewpoint Calendar
Edit Rulesets	Enables role to edit rulesets for Teradata and Aster systems		Workload Designer
Enable Clear Alert	Enables role to clear an alert for a component, process, or resource		Ecosystem Explorer
Enable Partial Size	Enables role access to run SQL queries on the system to retrieve and calculate the precise size of a partial table, a table with a WHERE clause. Affects partial table size only		Data Mover
Enable State Change	Enables role to change a state		Ecosystem Explorer
Enable Table Validation	Enables role to validate tables in multiple systems		Ecosystem Explorer
Ferret	Enables role access to the Ferret console utility		Remote Console
Gateway Global	Enables role access to the Gateway Global console utility		Remote Console
Get Explain from DB	Enables a single query refresh option. Allows the user to view the latest Explain steps retrieved from the database		My Queries, Query Groups, Query Monitor

Permission Name	Description	Dependency	Portlets Displaying Setting
Historical Distribution	Enables role access to the Historical Cylinder Distribution by Grade view		TVS Monitor
Lock Display	Enables role access to the Lock Display console utility		Remote Console
Manage Lab Groups	Enables role to create, edit, and delete lab groups		Data Labs Note: Select the View Reports permission when this permission is selected.
Manage Maps	Enables role to view map details, create and delete a sparse map, and manage or revoke grants		MAPS Manager
Operator Console	Enables role access to the Operator Console console utility		Remote Console
Priority Scheduler	Enables role access to the Priority Scheduler console utility		Remote Console
Query Configuration	Enables role access to the Query Configuration console utility		Remote Console
Query Session	Enables role access to the Query Session console utility		Remote Console
Recovery Manager	Enables role access to the Recovery Manager console utility		Remote Console
Release Query	Enables role access to the <input type="checkbox"/> list in the summary view to release a query that is in the delay queue and waiting to run		My Queries, Query Groups, Query Monitor, Workload Monitor
Routing Actions	Enables role to deploy and import session routing rules		Unity Setup
Run Job	Enables role to select Run now for analyze and collect jobs	• View System	Stats Manager, Performance Data Collection
Session Action	Enables role to move, fail over, kill, or abort sessions		Unity
Session Drilldown	Enables role to view details of each session from the Account String, Session, User, or Utility reports for Teradata and Aster systems	• View Summary • Account	Query Monitor

Permission Name	Description	Dependency	Portlets Displaying Setting
		<ul style="list-style-type: none"> Account Drilldown All My Criteria User User Drilldown Utility Partition Drilldown 	
Show Locks	Enables role access to the Show Locks console utility		Remote Console
Space by Database	Enables role to select By Database in the Select Report dialog box for Teradata and Aster systems		Space Usage
Space by Node	Enables role to select By Node in the Select Report dialog box for Aster and Hadoop systems		Space Usage
Space by Table	Enables role access to the details view with a list of tables in each database for Teradata and Aster systems		Space Usage
Space by Vproc	Enables role to select By Vproc in the Select Report dialog box for Teradata systems		Space Usage
State Change	Enables role to change the state in the Table Health view of the Unity portlet		Unity
Teradata DWM Dump	Enables role access to the Teradata DWM Dump console utility		Remote Console
Unlock Rulesets	Enables role to unlock any ruleset for Teradata and Aster systems	<ul style="list-style-type: none"> View Rulesets 	Workload Designer
Update Space	Enables role access to the Update Space console utility		Remote Console
View Block	Enables role to see that a query is blocked. The Blocked By tab is visible, but contains no details	<ul style="list-style-type: none"> View Summary 	SQL Scratchpad
View Block Details	Enables role access to the content of the Blocked By tab	<ul style="list-style-type: none"> View Summary View Block 	SQL Scratchpad

Permission Name	Description	Dependency	Portlets Displaying Setting
View Blocked By	Enables role access to the Blocked By tab	<ul style="list-style-type: none"> View Summary Session Drilldown 	My Queries, Query Groups, Query Monitor, Query Spotlight, Workload Monitor
View Blocking	Enables role access to the Blocking tab when connected to Teradata Database 15.0 and later	<ul style="list-style-type: none"> View Summary Session Drilldown 	Query Monitor
View Canary Query	Enables role to view canary queries in the summary view		Canary Response Times, Productivity
View DBQL Metrics	Enables role to view the Queries Per Hour metrics		Productivity
View Delay	Enables role access to the Delay tab	<ul style="list-style-type: none"> View Summary Session Drilldown 	My Queries, Query Groups, Query Monitor, Workload Monitor
View Defer	Enables role access to the Defer tab	<ul style="list-style-type: none"> View Summary Session Drilldown 	My Queries, Query Groups, Query Monitor, Workload Monitor
View Detail	Enables role access to the details view	<ul style="list-style-type: none"> View Summary 	Alert Viewer, Completed Queries, My Queries, Node Monitor, Query Groups, Query Spotlight, System Health, Workload Health, Workload Monitor
View Explain	Enables role access to the Explain tab containing the confidence level indicator	<ul style="list-style-type: none"> View Summary Session Drilldown 	Completed Queries, My Queries, Query Groups, Query Monitor, Query Spotlight, Workload Monitor
View Explain Simple	Enables role access to the Explain tab	<ul style="list-style-type: none"> View Summary 	SQL Scratchpad
View Maps	Enables role access to the Maps tab		MAPS Manager
View Query Band	Enables role access to the Query Band tab	<ul style="list-style-type: none"> View Summary 	My Queries, Query Groups, Query Monitor,

Permission Name	Description	Dependency	Portlets Displaying Setting
		<ul style="list-style-type: none"> Session Drilldown 	Query Spotlight, Workload Monitor
View Query Detail	Enables role access to the details view for a Teradata QueryGrid query		My Queries, Query Groups, Query Monitor, Completed Queries
View Reports	Enables role access to the Reports tab		Data Labs Note: Select this permission when the Manage Lab Groups permission is selected.
View Rulesets	Enables role to view rulesets for Teradata and Aster systems		Workload Designer
View Skew	Enables role to view the Skew tab	<ul style="list-style-type: none"> View Summary Session Drilldown 	Query Monitor
		<ul style="list-style-type: none"> View Summary View Detail 	My Queries, Query Groups, Workload Monitor
View SQL	Enables role access to the SQL tab	<ul style="list-style-type: none"> View Summary Session Drilldown 	Application Queries, Completed Queries, My Queries, Query Groups, Query Log, Query Monitor, Query Spotlight, Workload Monitor
View Statistics	Enables role access to see statistics in the Today's Statistics view		Today's Statistics
View Summary	Enables role access to the summary view		Application Queries, Completed Queries, Canary Response Times, Data Labs, Data Mover Elastic Performance, Hadoop Services, Lock Viewer, My Queries, Node Monitor, Node Resources, Productivity, Query Groups, Query


Permission Name	Description	Dependency	Portlets Displaying Setting
			Log, Query Monitor, Query Spotlight, Remote Console, Space Usage, SQL Scratchpad, System Health, TVS Monitor, Viewpoint Monitoring, Workload Health, Workload Monitor
View System	Enables role to select the system from the system selection list		MAPS Manager, Stats Manager, Performance Data Collection
View System Monitoring	Enables role access to the System Monitoring tab		MAPS Manager
View Workflows	Enables role access to the Object Workflows tab		MAPS Manager
Vproc Manager	Enables role access to the Vproc Manager console utility		Remote Console

Granting Web Service Access and Permissions

You can enable access and select permissions for web services by system and role.

Note:

Web services must also be enabled in **Portlet Library**.

1. From the **Roles Manager** view, select a role from the list.
2. Select the **Web Services** tab.
3. Select the check box to the left of the web service name to enable the web service.
4. Select  for the web service you want to set permissions.
5. Select the permissions you want to grant to the selected role for each available system by doing any of the following:

Note:

The available choices are dependent on your role.

- Select the system type from the list.
- Select the **Select all** check box to select or clear all permissions for that particular system.
- Click the permission name to select or clear that particular permission for all systems.

- Select a check box to select or clear a single permission for a single system.
6. Click **OK**.
 7. Click **Apply**.

Shared Pages

The **Shared Pages** portlet displays a list of the shared pages by role and allows you to create, copy, delete, and configure shared pages.

A shared page is a pre-defined page created by the Teradata Viewpoint Administrator or a user with role **VP_User_Manager** that contains the portlets, systems, and settings a role uses most frequently. A shared page can be assigned to only one role and inherits the permissions of the role. Multiple shared pages can be assigned to a role.

Shared pages can be editable, read-only, or mandatory.

- Pages are editable by default with users having full control over the content, layout, and settings of their instance of the shared page. The shared page in the **Shared Pages** portlet is not affected.
- Read-only pages can be added or removed from the portal by users but the content and layout cannot be changed.
- Mandatory pages are similar to read-only pages except they are present in the user portal at all times, therefore, users cannot add or delete the page from the portal.

To allow users to administer shared pages, create a role using **Roles Manager** and then enable the **Shared Pages** portlet for this role.

Following is an example of the summary view.

PAGE	ROLE	ENABLED	FIRST LOGIN	PORTLETS	DESCRIPTION
Finance Dashboard	Business_User	✓	✓	Metric Heatmap	Business use
Network Dashboard	Power_User	✓		Workload Health, Quer	
Sales Dashboard	End_User	✓		My Queries, SQL Scrati	Sales users a

Shared Page Controls

Summary View


You can access these controls from the summary view.

Control	Description
Copy Page	Creates an exact copy of the shared page and opens the Copy shared page dialog box where you can name and configure settings for the page.
Delete	Deletes the page from the Shared Pages portlet. Removes the page from the add page list for users in this role, and removes the mandatory page from the portal the next time the user logs in or refreshes the page.

Control	Description
Add Shared Page	Opens a dialog box where you can create a new shared page.

Shared Page Editor

You can access the **Shared page editor** by clicking a row in the summary view.

Control	Description
Add Content	Opens the Add Content screen where you can add portlets to a shared page.
Page Properties	Opens a dialog box where you can edit the page name, set page controls, and add an optional page description. You cannot change the assigned role after the page has been saved.
Modes	<p>Edit and preview Use this mode to edit the page. This mode shows how the page appears to the user in the target role.</p> <p> Configure settings Use this mode to configure portlet settings that are not accessible to users in the target role. After using this mode, return to the Edit and preview mode.</p>
Save	Commits changes that have not already been saved. When a page is enabled and saved, users in the associated role are able to immediately access the shared page.
Cancel	Returns to the summary view without saving changes to content or layout.

Shared Page Properties

You can access the shared page properties controls by selecting **Page Properties** while in the **Shared page editor**.

Controls	Description
Page Name	Allows you to add a new or edit an existing name for a shared page.
Enable page	Adds the non-mandatory page to the add page list.
Show this page the first time a user logs in	Displays the page when a new Viewpoint user logs in for the first time as a member of the role.
Read only	Prevents the user from adding portlets, or making changes to the content, layout or page name of the shared page. The shared page can be added from the add page list and deleted from the portal.
Mandatory	Same as read-only, except the shared page is present in the user portal at all times. The shared page cannot be added from the add page list or deleted from the portal.

Controls	Description
Description	Allows you to add a new or edit an existing description for a shared page.
Last Modified	Displays the date and time the shared page was last modified.
Date Created	Displays the date and time the shared page was created.

Read-Only and Mandatory Shared Pages

You can prevent users from modifying shared pages by making a shared page read-only or mandatory.

Read-only prevents a user from adding portlets, or making changes to the content, layout or name of the shared page. The user, however, can add the page using the add page list and delete the page from the portal. Read-only pages can be added to the portal only once.

Mandatory pages are always present for all users. The restrictions are the same as read-only pages, except the user cannot add a mandatory shared page using the add page list or delete the page from the portal.

When you update a read-only or mandatory shared page, it automatically updates the next time the user logs on or refreshes the page.

When a user is removed from a role, a read-only shared page remains in the portal until the user deletes it. However, a mandatory shared page is removed from the portal the next time the user logs in or the page refreshes.

Creating a Shared Page

Create a shared page when you want a role to use specific portlets, systems, and settings, or to give users a pre-defined page as a starting point.

To allow users to administer shared pages, create a role using **Roles Manager** and then enable the **Shared Pages** portlet for this role.

1. From the **Shared Pages** portlet, click **Add Shared Page**.
2. Enter the new page name up to 30 characters.
Page names must be unique.
3. Select a role to assign to the page from the list. A role cannot be changed after you create a page. Only roles that you have permission to see are displayed in the list.

Note:

If a user with the role VP_User_Manager creates the page, only the user's role and the roles created by the user is visible to the user.


4. Select the **Enable page** check box.

Enabling a non-mandatory page makes the page available for the user to add. Enabling a mandatory page adds the page to the user's portal after it is saved.

5. [Optional] Select the **Show this page the first time a user logs in** check box to display the page the first time a new Viewpoint user logs in.
6. [Optional] Select the **Read only** check box to prevent users in this role from adding portlets, or making changes to the content, layout or page name of the shared page.
7. [Optional] Select the **Mandatory** check box to make the shared page present in the user portal at all times.
8. [Optional] Add a description up to 500 characters.
The description appears in the summary view in the **Shared Pages** portlet.
9. Select **Create**.
The **Shared page editor** opens. If the page is mandatory, it automatically appears in the user's portal as soon as you click **Create**.
10. Select **Add Content** to select the portlets you want to add to the shared page, and select **Add**.
Selecting **Cancel** does not cancel the page creation.

Note:

If a user with role VP_User_Manager creates the page, Bar Operations portlet, and shared bar portlets is not visible.



11. Edit the shared page by doing any of the following:
The **Edit and preview** button is selected by default. This preview mode provides a way for you to see how the page appears to the user in the target role.
 - Arrange the portlets on the page in the order they must appear to the user.
 - Select **Add Content** to add additional portlets, then select **Add**.
12. [Optional] Configure settings not accessible to users in the target role.
For example, if a role is not given access to change what columns are displayed, you can configure the settings for each portlet to display specific columns that you determine are needed by the role.
 - a. Select .
 - b. Configure the settings for each portlet.
 - c. Select **Edit and preview** to see how the page will appear to users in the target role.
13. Select **Save**.
Changes are committed and the new shared page appears in the summary view in the **Shared Pages** portlet.

Content, layout, and portlet setting changes are immediately available to users in the role. Read-only and mandatory pages automatically update the next time the user logs in or refreshes the page. Unrestricted pages must be added by the user from the add page list.

Copying a Shared Page

Copying an existing shared page is useful when the majority of the page characteristics are applicable to another role.

You can change the role only before you save the copied page. If you change the role, review the shared page thoroughly to make sure the permissions of the new role are displaying the portlet content as you expected.

1. From the **Shared Pages** portlet, select  located on the row of the page you want to copy then select **Copy Page**.
2. Enter the new page name up to 30 characters.
Page names must be unique.
3. Specify the page properties.
4. Click **Create**.
The page is accessible to users as soon as you select **Create**.
5. [Optional] Edit the shared page.
6. [Optional] Configure settings not accessible to users in the target role.
For example, if a role is not given access to change what columns are displayed, you can configure the settings for each portlet to display specific columns that you determine are needed by the role.
 - a. Select .
 - b. Configure the settings for each portlet.
 - c. Select **Edit and preview** to see how the page will appear to users in the target role.
7. Select **Save**.
Changes are committed and the new shared page appears in the summary view in the **Shared Pages** portlet.

Content, layout, and portlet setting changes are immediately available to users in the role. Read-only and mandatory pages automatically update the next time the user logs in or refreshes the page. Unrestricted pages must be added by the user from the add page list.

Editing a Shared Page

Edit a shared page to add portlets or change portlets and portlet settings. You cannot edit the assigned role.


1. From the **Shared Pages** portlet, select the row of the page you want to edit.
The **Shared page editor** opens. The page name displays at the top with the role name for the shared page shown in parentheses.
2. [Optional] Edit the shared page by doing any of the following:
 - Select **Add Content** to add portlets, then select **Add**.

Note:

If a user with the role VP_User_Manager edits the page, Bar Operations portlet and shared bar portlets is not visible.


- Select **Page Properties** to modify the settings, then select **Save**.
 - Rearrange the portlets on the page.
 - Remove portlets.
3. [Optional] Configure settings not accessible to users in the target role.

For example, if a role is not given access to change what columns are displayed, you can configure the settings for each portlet to display specific columns that you determine are needed by the role.

- a. Select .
 - b. Configure the settings for each portlet, including, but not limited to, selecting systems.
 - c. Select **Edit and preview** to see how the page will appear to users in the target role.
4. Select **Save**.
Changes are committed and the updated shared page appears in the summary view in the **Shared Pages** portlet.

Content, layout, and portlet setting changes are immediately available to users in the role. Read-only and mandatory pages automatically update the next time the user logs in or refreshes the page. Unrestricted pages must be added by the user from the add page list in order for the user to see the updated page.

Deleting a Shared Page

1. From the **Shared Pages** portlet, select  located on the row of the page you want to delete then select **Delete**.
2. Select **Delete** to confirm that you want to delete this shared page.
The shared page no longer appears in the **Shared Pages** portlet.

A non-mandatory page is removed from the add page list for users in this role. A read-only shared page becomes a normal page. A mandatory shared page is removed from the portal the next time the user logs in or refreshes the page.

User Manager

The **User Manager** portlet allows the Teradata Viewpoint Administrator to view and manage Teradata Viewpoint user accounts so that you can add usernames, modify user accounts, assign roles, reset passwords, and export data. Each portal username must be unique. You must have Teradata Viewpoint Administrator permission to access the **User Manager** portlet.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you can only edit or delete a Viewpoint local user, and assign a user role. No other operations are allowed.

Users View

The **Users** view has a list of all Teradata Viewpoint system users with last name, first name, portal username, email address, the last logon time, the creation time, the modification time, the user who performed the modification, and assigned roles.

The last logon information remains blank, if a newly created user has never logged into the system.

When you create an user, the creation time appears in the **Users** view. When you make any change to the first name, last name, email address, password, or role of the user, the **Modification Time** and **Modified by** fields show the change details.

In this view, you can view, select, and search for individual user details, add or delete a Teradata Viewpoint user account, and export data. See [Summary Table Controls](#).

Adding a User

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you cannot add a user.

Each portal username must be unique to avoid being logged off when someone logs on with the same username, such as Administrator.

If your Teradata Viewpoint portal is configured to use auto-provisioning, a user account is created automatically the first time a user logs on. Auto-provisioning validates portal usernames with the LDAP directory server and assigns the default role and email domain.

1. From the **Users** view, select **Add User**.
2. Enter the portal username for the user account.
3. Choose an authentication method:
 - Select the **Authenticate externally** check box to authenticate the selected user from an external source, such as LDAP.
 Externally authenticated passwords are stored in the external source, such as LDAP, and validated from that source.
 - Clear the **Authenticate externally** check box to authenticate the selected user from the Teradata Viewpoint internal database.
4. Enter the first name, last name, and email address for the user account.
5. If the authentication method is not external, enter a password for the user account and enter it again to validate.
6. Enter your admin password to authorize the changes.
7. Select **Apply**.

Deleting Users

Note:

If JWT integration with Vantage is enabled (Viewpoint versions 16.50.05.00 and later), you can only delete a Viewpoint local user account.

1. From the **Users** view, do one of the following:

Option	Description
To delete one user	<ol style="list-style-type: none"> Browse the list of users or enter text to find the user. Select <input type="checkbox"/> on the row you want to delete then select Delete.
To delete multiple users	<ol style="list-style-type: none"> Select the Table Actions list then select Delete User. Select one or more check boxes next to the users you want to delete then select Next.

A confirmation message appears.

- Enter your admin password to authorize the changes.
- Select **Delete**.

The user account is deleted. The user no longer appears in the **Users** view.

User Details View

In the user details view you can edit and delete Teradata Viewpoint user accounts, assign roles to users, set role priority, and change password of the user if it is not externally authenticated. It also allows you to view the user details audit and roles audit information.

Assigning User Roles

A role must be defined in the **Roles Manager** portlet before it can be assigned to a user.

You can assign roles to an existing Teradata Viewpoint user account. If your Teradata Viewpoint portal is configured to use auto-provisioning, a user account is created automatically the first time a user logs on to Teradata Viewpoint. By default, auto-provisioned accounts are authenticated externally and assigned a default role. The LDAP authentication source and default role are set in the **LDAP Servers** portlet.

When assigning multiple roles to the same user, you can set role priority.

- From the **Users** view, browse the list of users or use the filters to find users.
- Select a user name.
- Select the **Roles** tab.

Available roles are listed in the **Available Roles** pane. Roles assigned to the user are shown in the **Selected Roles** pane.

- Select a role from the **Available Roles** pane or select multiple roles by pressing **Shift** or **Ctrl**.
- Select ☐.

The selected roles appear in the **Selected Roles** pane.

- [Optional] Change the role priority to be applied to the default portlet settings for a role when a user is assigned multiple roles.

For more information, see [Setting the Priority of Default Portlet Settings for a Role](#).

- Select **Apply**.

Note:



If a user with the role VP_User_Manager updates roles, it only assigns/revokes its own role or roles create by it. The other roles are disabled.

Setting the Priority of Default Portlet Settings for a Role

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you cannot set the priority of default portlet settings for a role.

You can set the role priority to be applied to the default portlet settings when a user is assigned multiple roles. The highest ordered role dictates what displays in the portlet. For example, if a user has two roles, User and Administrator, and Administrator has the highest priority, then the user is given the default portlet settings based on the Administrator role. You can set the default portlet settings using the **Portlets** tab in **Roles Manager**.

1. From the **Users** view, browse the list of users or use the filters to find users.
2. Select a user name.
3. Select the **Roles** tab.
4. Select a role in the **Selected Roles** pane.
5. Do any of the following:
 - Select  to increase the priority of the role, which moves the role higher in the list.
 - Select  to decrease the priority of the role, which moves the role lower in the list.
6. Select **Apply**.

Editing a User

Note:

If JWT integration with Vantage is enabled (Viewpoint versions 16.50.05.00 and later), you can only edit a Viewpoint local user account.

You can edit a Teradata Viewpoint user account, except the portal username. The username is defined when the user account is created using the **User Manager** portlet or it is stored in LDAP when the user is externally authenticated.

1. From the **Users** view, browse the list of users or use the filters to find users.
2. Select a user name.
3. Choose an authentication method:
 - Select the **Authenticate externally** check box to authenticate the selected user from an external source, such as LDAP.

Externally authenticated passwords are stored in the external source, such as LDAP, and validated from that source.

- Clear the **Authenticate externally** check box to authenticate the selected user from the Teradata Viewpoint internal database.
4. [Optional] Change the first name, last name, or email address for the user account.
 5. Select **Apply**.
 6. [Optional] Change the password if it is not authenticated externally.
For more information, see [Changing a User Password](#).

Changing a User Password

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you cannot change a user password (**Change Password** is not visible).

You can change a Teradata Viewpoint user password if it is not authenticated externally. When a password is authenticated externally, change the password in the external source.

1. From the **Users** view, browse the list of users or use the filters to find users.
2. Select a user name.
3. From the **General** tab, select **Change Password**.
4. Enter your admin password to authorize the changes.
5. Enter the new password.
6. Re-enter the password exactly as it was entered previously.
7. Select **Change Password**.

Viewing User Details Audit

The **USER DETAILS AUDIT** tab shows the modifications done to the user. The view has the user first name, last name, email id, password modification status, modification timestamp, and details of the user who modified.

For the externally authenticated users the **USER DETAILS AUDIT** tab remains disable.

1. From the **Users** view, browse the list of users or use the filters to find users.
2. Select a user name.
3. Select the **USER DETAILS AUDIT** tab.



LOG TIME	FIRST NAME	LAST NAME	EMAIL	PASSWORD MODIFIED	MODIFIED BY
9/17/22 8:45:10 AM	User1	User	user1@teradata.com	true	admin
9/17/22 8:45:44 AM	User1	User	user1@teradata.com	false	admin
9/17/22 8:45:39 AM	User1	User	user@teradata.com	false	admin

4. [Optional] You can configure columns or export data of this tab.

Viewing Roles Audit

The **ROLES AUDIT** tab shows the modifications done to the roles assigned to an user. The view has the role name, the performed action (ADDED/REMOVED), modification timestamp, and details of the user who modified.

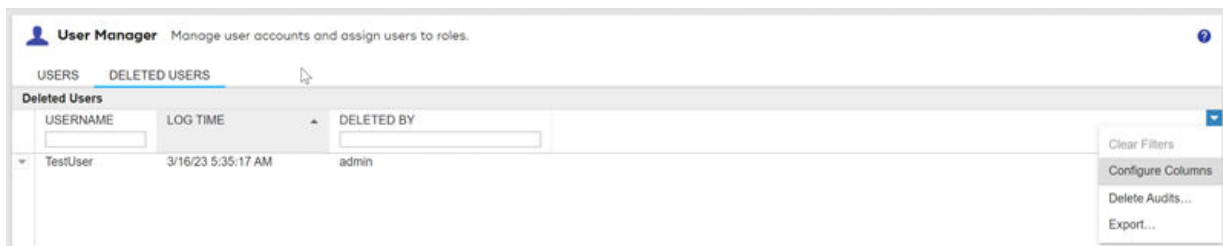
1. From the **Users** view, browse the list of users or use the filters to find users.
2. Select a user name.
3. Select the **ROLES AUDIT** tab.



4. [Optional] You can configure columns or export data of this tab.

Viewing Deleted Users Audit

1. Open the **User Manager** portlet.
2. Select **Deleted Users** tab.



3. [Optional] You can configure columns or export data of this tab.

Removing Deleted Users Audit

1. From the **Deleted Users** view, do one of the following:

Option	Description
To remove deleted audit for one user	<ol style="list-style-type: none"> Browse the list of deleted users audit or enter text to find the user. Select <input type="checkbox"/> on the row you want to delete, then select Delete.
To remove multiple deleted users audit	<ol style="list-style-type: none"> Select the Table Actions list then select Delete Audits. Select one or more checkboxes next to the deleted users audit you want to delete, then select Next.

A confirmation message appears.

2. Enter your admin password to authorize the changes.
3. Select **Delete**.

Alert Viewer

Alert Viewer Overview

The **Alert Viewer** portlet allows users to view triggered alerts and are specific to your system. The alert information in the summary view is refreshed every 30 seconds. Every alert has a timestamp, displaying the date and time at which the alert was issued.

Alerts

An alert is an event that the Teradata System Administrator defines as significant. The Teradata System Administrator assigns alert severity levels to rank alerts, and can also include a message. The severity levels are critical, high, medium, or low. The alerts displayed in the **Alert Viewer** portlet are specific to your system.

Alert Example

The Teradata System Administrator can define that a database exceeding a certain amount of storage usage triggers an alert. After the usage level is exceeded, an alert appears in the **Alert Viewer** portlet to inform the portlet user that the threshold was exceeded.

Alert Viewer View

The **Alert Viewer** view displays information about alerts, including their type, source, severity, and any messages. The message text displays the alert property token, in the format `${alertProperty}`. When you drill down to view alert details, the value of the alert property is displayed in the message.

You can filter the alerts by severity, time period, type, or name. You can also combine the filters to narrow the results further.

Severity Filter Bar

Displays a count of the alerts for each severity. Click any severity in the bar to change the displayed data in the summary table to show only the alerts of that severity.

▼ **Last 1 hour** Sets the time period for the alerts in the **Severity Filter Bar**.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information about alerts in columns. The current view is configured in the **Columns** dialog box. The view is refreshed every 30 seconds. Click a row in the table to see details.

Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

Hide Alerts enables you to select and hide alerts so that they are no longer displayed in the summary list. Hidden alerts can be viewed and then displayed again in the summary list.

For more information on **Clear Filters**, **Configure Columns**, and **Export** table actions, see [Summary Table Controls](#).

The screenshot shows the Alert Viewer interface. At the top, there's a header with 'Alert Viewer' and a time '12:02:53 PM'. Below the header is a 'Severity Filter Bar' with buttons for 'All' (190), 'Critical' (18), 'High' (20), 'Medium' (11), and 'Low' (141). To the right of the filter bar is a 'Time Period' dropdown set to 'Last 1 hour'. Below the filter bar is a 'Filters' section with input fields for 'SEV...', 'TIMESTAMP', 'ALERT TYPE', and 'ALERT NAME'. To the left of the table is a 'Hide Alert' column with checkboxes. The table itself has columns for 'SEV...', 'TIMESTAMP', 'ALERT TYPE', 'ALERT NAME', and 'S'. The table shows 7 rows of alert data. At the bottom of the table, it says '190 rows total'.

SEV...	TIMESTAMP	ALERT TYPE	ALERT NAME	S
	1/3/14 11:03:26 AM	Canary Query	skmCQ	V
	1/3/14 11:03:26 AM	Canary Query	Test	V
	1/3/14 11:05:26 AM	Canary Query	skmCQ	V
	1/3/14 11:05:26 AM	Canary Query	Test	V
	1/3/14 11:07:26 AM	Canary Query	skmCQ	V
	1/3/14 11:17:26 AM	Canary Query	skmCQ	V
	1/3/14 11:19:26 AM	Canary Query	skmCQ	V

Alert Metrics

Metric	Description
Alert Name	Name of the alert
Alert Type	Type of alert
Message	Message providing more information about the alert
Severity	Severity level assigned to the alert
Severity Icon	Colored icon representing the severity level
Source	Source (Viewpoint, Teradata Alerts, Server Management, Performance Data Collection, or Teradata Ecosystem Manager)
System Name	Name of the configured database system or site ID
Timestamp	Date and time when the alert was issued

Hiding Alerts Basics

By default, the **Alert Viewer** portlet displays all alerts that have been generated. You can hide some alerts so they no longer appear in the summary list. For example, you might want to hide all alerts generated for database space conditions or alerts associated with conditions that were subsequently resolved.

You can also display hidden alerts, and then unhide them. When you display hidden alerts, the alert rows appear with a line through them, and their associated severity indicators appear as unfilled squares. If you click the alert to view its details, the severity indicator next to the alert name also indicates the hidden state of the alert.

If you hide alerts in one of the following, the alerts are hidden in the other two:

- **Alert Viewer** portlet
- **Viewpoint Dashboard**
- **Viewpoint Mobile**

The **Settings** view allows you to display previously hidden alerts.

Hiding Alerts

You can select and hide alerts from the **Alert Viewer** portlet.

- To hide only one alert, do the following:

Option	Description
From the Alert Viewer summary view	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> in the row of the alert you want to hide. 2. Select Hide Alert. <p>The alert no longer appears in the view.</p>
From the Alert Viewer summary view	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> in the table header to open the Table Actions list. 2. Select Hide Alerts. 3. Select the check box in the row of the alert you want to hide. 4. Click Save. <p>The alert no longer appears in the view.</p>
From the details view for an alert	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> next to the alert name. 2. Select Hide Alert. <p>The colored severity indicator square appears unfilled to indicate that the alert is hidden. The alert no longer appears in the summary view.</p>

- To hide one or more alerts, do the following:
 1. From the **Alert Viewer** summary view, select ☐ in the table header to open the **Table Actions** list.
 2. Select **Hide Alerts**.
 3. Select the check boxes in the rows of the alerts you want to hide.

4. Click **Save**.
The alerts no longer appear in the summary view.

Unhiding Alerts

Prerequisite:

You must display alerts that were previously hidden before you can unhide them.

You can select and unhide alerts from the **Alert Viewer** portlet. When you unhide an alert, it is displayed again in the summary view and the severity indicator associated with the alert is displayed as a filled square.

- To unhide only one alert, do the following:

Option	Description
From the Alert Viewer summary view	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> in the row of the alert you want to unhide. 2. Select Unhide Alert.
From the Alert Viewer summary view	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> in the table header to open the Table Actions list. 2. Select Hide Alerts. 3. Clear the check box in the row of the alert you want to unhide. 4. Click Save.
From the details view for an alert	<ol style="list-style-type: none"> 1. Select <input type="checkbox"/> next to the alert name. 2. Select Unhide Alert.

The alert displays in its normal active state again in the summary view and the severity indicator associated with it is displayed as a filled square.

- To unhide one or more alerts, do the following:
 1. From the **Alert Viewer** summary view, select ☐ in the table header to open the **Table Actions** list.
 2. Select **Hide Alerts**.
 3. Clear the check boxes in the rows of the alerts you want to unhide.
 4. Click **Save**.

The alert displays in its normal active state again in the summary view and the severity indicator associated with it is displayed as a filled square.

Related Information:

[Displaying Hidden Alerts](#)

Viewing Alert Details


1. In the alert summary view, click anywhere in a row to display the alert details view for that alert. The alert details for an alert displays details about the selected alert and displays the following:
 - **General:** Alert type, severity, source, timestamp, resulting action, and alert criteria.
 - **Properties :** Alert properties and their associated values. For metric descriptions, see [Alerts](#).
 - **Messages:** Any messages issued.
2. Select **< > Previous/Next** to view the details for each alert without returning to the summary view.

Settings View

The **Settings** view allows you to designate that hidden alerts be displayed.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Displaying Hidden Alerts

You can designate that hidden alerts be displayed. Hidden alerts are displayed with a line through them, and their associated severity indicators appear as unfilled squares.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Display hidden alerts** check box.
3. Click **OK**.

Application Queries

Application Queries Overview

The **Application Queries** portlet enables you to view key information about the queries submitted by applications, based on the historical DBQL data in Analytics Database. You can review application queries designated as suspect based on thresholds set by a Viewpoint Administrator. You can view key metrics for any date for which query log data has been collected and retained by Viewpoint for the applications you have access to, and drill down to see details about specific queries. You can also view the trends of key performance indicators, aggregated daily or weekly and going back over a time frame from the last six months to two years, specific to an application and application version.

You can only see data in the **Application Queries** portlet if the optional query logging (DBQL) feature has been enabled in Analytics Database. The **Application Queries** portlet displays all queries that have been logged for the applications you have access to.

The **Application Queries** portlet does not participate in the rewind feature. The portlet provides a date selector to select the date of the query log information you want to view.

Application Queries View

The **Application Queries** view displays summary information for each application that submitted queries on the selected date.

Date Selector

Shows the date of the information displayed and allows you to select a different date.

Summary Table

Displays summary information about the applications that submitted logged queries on the selected date.

Filters

Filters allow you to change displayed data by showing only rows that match your filter criteria. You can also sort on the column headers to find the details you want.

▾ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data up to one million rows. If more than one million rows of data are available, then use the filters to display the data you want before exporting it.

For more information, see [Summary Table Controls](#).

Summary Table Filters

Table Actions Date Selector

APPLICATION	VERSION	SYSTEM	QUERY COU...	SUSPECT ...	SYSTEM UTIL...	TOTAL AMP...	TOTAL PARS...	IMPACT CPU	TOTAL I/O ...
PDCR	15.10.00.00	Dev1	2.82K	23	0.02	1.08K	191.5	1.72K	5.99M
PDCR	15.10.00.00	Dev2	26	1	0	2.104	2.1	4.896	29.3K
PDCR	15.10.00.00	Dev3	1.99K	15	0.008	278.4	197.6	737.4	3.35M
StatsManager	15.10.00.00	Dev1	23	0	0	0	0.016	0	0
StatsManager	15.10.00.00	Dev2	162	0	0	3.728	1.548	11.95	96.9K
StatsManager	15.10.00.00	Dev3	24	0	0	0	0.02	0	0
Viewpoint	15.10.00.00	Dev1	42.3K	456	0.048	2.52K	488.5	5.2K	25.6M
Viewpoint	15.10.00.00	Dev2	89.3K	624	0.089	4.43K	1.12K	9.9K	39.1M
Viewpoint	15.10.00.00	Dev3	51.1K	192	0.041	1.95K	629.3	4.03K	19.3M

10 rows total

Changing the Date

Use the date selector to designate the date of the information you want to view.

The date selection is retained until you change it or until your current Viewpoint session ends.

1. Click the date selector box.
2. Select a date from the calendar.

Application Queries View Metrics

Metrics available for display are listed as following.

Metric	Description
Application	Application that submitted queries on the selected date
Burn Rate	Percentage of the overall system capacity that was effectively consumed by all queries during their execution
Gating Efficiency	Percentage indicating the overall query duration time spent executing versus being delayed. A value of 100% indicates all time was spent executing.
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.
Max Spool	Total amount of temporary spool space used
Query Counts	Total number of logged queries
Suspect Queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector
System	System of the application that submitted queries on the selected date

Metric	Description
System Utilization	Percentage of the system's CPU that was utilized by the application for the particular day by logged queries
Time Active	Amount of time in seconds that queries spent executing. Excludes any time queries were delayed before being executed.
Time Delay	Amount of time in seconds that queries were delayed before being executed
Time Expansion	Ratio of how long the query took to execute versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded. For example, a value of 10x to indicates an expansion of 10 times the ideal execution time.
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries
Total I/O	Total number of I/O requests
Total Parser CPU	Total amount of CPU consumed by the parsing engine in processing queries
Version	Version of the application that submitted queries on the selected date

Viewing Queries by Application

You can view summary information about all of the queries submitted by an application on the selected date.

1. Click the row of an application listed in the summary table.

Application Queries Details View

When you click an application listed in the **Application Queries** view, the system and application are listed at the top, and you can select the version of the application from the **Version** list.

Information for the queries submitted by the application is displayed on the **Overview**, **Queries**, and **Trends** tabs.

Overview Tab

The **Overview** tab includes a **Summary** section that lists key metric values and a **Counts** section that displays bar charts representing query counts for a selected metric.

Summary

The **Summary** section displays values for the following metrics.

Metric	Description	Type
Burn Rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent

Metric	Description	Type
Gating Efficiency	Percentage indicating the percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates all time was spent executing.	Percent
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
Logged Queries	Number of queries for which query logging was enabled for the selected date	Number
Max Spool	Total amount of temporary spool space used	Number
Suspect Queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector	Number
System Utilization	Percentage of the system's CPU utilized by the user or application	Percent
Time Active	Number of seconds the query spent executing. Excludes any time the query was delayed before being executed.	Number
Time Delay	Number of seconds the query was delayed before being executed	Number
Time Expansion	Ratio of how long the query took to execute versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	Number
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries	Number
Total I/O Count	Total number of I/O requests the query used	Number
Total Parser CPU	Total amount of CPU consumed by the parsing engine in processing the query	Number

Counts

The **Counts** section displays bar charts representing total counts for the following metrics.

Metric	Description
Account String	Number of queries submitted by the account string represented by the bar
App ID	Number of queries associated with the displayed App ID. The App ID is the ID of the application used to access Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.
Cache Flag	Number of queries that used the cache flag type represented by the bar. The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated.

Metric	Description
	<ul style="list-style-type: none"> A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed.
Client Address	Number of queries submitted by the client IP address represented by the bar
ClientID	Number of queries submitted by the client ID represented by the bar
Error Code	Number of queries that have the error code number represented on the bar
Number of Active AMPs	Number of queries that used the number of AMPs represented by the bar
Statement Type	Number of queries that used the SQL statement type represented by the bar
Username	Number of queries submitted by the Teradata user represented by the bar
Workload end	Number of queries that finished execution in the workload represented by the bar
Workload start	Number of queries that started execution in the workload represented by the bar
Workload start -> end	Number of queries that started and finished execution in the same workload represented by the bar

Queries Tab

The **Queries** tab enables you to view two subtabs:

- **Logged Queries:** All queries that were logged for the selected application and application version.
- **Suspect Queries:** All suspect queries that were logged for the selected application and application version.

Logged Queries Tab

The **Logged Queries** tab provides key metrics for queries logged on the selected date.

Suspect Queries Tab

The **Suspect Queries** tab displays information for all logged queries that are designated as suspect. Suspect queries are those whose values surpass thresholds defined for the Query Log data collector in the **Monitored Systems** portlet. Thresholds can be set for the following metrics:

- CPU Skew
- I/O Skew
- Product Join Indicator
- Unnecessary I/O

Any queries that result in more than one AMP CPU second consumed and that exceed the defined thresholds are listed in the **Suspect Queries** tab, and the value that exceeds the threshold is displayed in red.

Logged Queries and Suspect Queries Tabs Metrics

Metrics available for display are listed as following.

Metric	Description	Type
Account	Account of the user who submitted the query	String
Active AMPs	Number of AMPs involved in executing the query	Number
AMP CPU Time	Total number of CPU seconds consumed by AMPs in processing the query	Number
App ID	The ID of the application used to access the Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.	String
Burn Rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent
Cache Flag	The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated. • A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed. 	String
Client Address	IP address of the client	String
Client ID	ID of the client that submitted the query	String
CPU Skew	CPU skew of the entire query	Percent
Error Code	Error code associated with the query	String
Final Workload	Name of the workload under which the query finished	String
Flagged	Denotes whether the query has been flagged for further analysis	String
Gating Efficiency	The percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates that all time was spent executing.	Percent
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
I/O Skew	I/O skew of the entire query	Percent

Metric	Description	Type
Max AMP CPU Time	The maximum CPU consumed by a single AMP while executing the query	Number
Max AMP I/O	The maximum IOs submitted by a single AMP while executing the query	Number
Parser CPU Time	Total amount of CPU seconds consumed by the parsing engine in processing the query	Number
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proc ID	ID of the parsing engine that run the query	String
Query ID	System-wide unique ID that identifies the query	String
Row Count	Number of result rows returned by the query	Number
Spool Usage	Amount of spool consumed by the query, using a 1024 byte kilobyte as the base	Number
Start Time	The time the query was submitted	Number
Starting Workload	Name of the workload under which the query started execution	String
Statement Type	SQL statement type the query used	String
Step Count	Number of steps in the query plan	Number
Time Active	Number of seconds the query spent running. Excludes any time the query was delayed before being executed.	Number
Time Delay	Number of seconds the query was delayed before running.	Number
Total Duration	Total elapsed time it took for the query to run once it was submitted	Number
Time Expansion	Ratio of how long the query took to run versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	Number
Total I/O Count	Total number of I/O requests the query used	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring in proportion to the CPU consumed	Number
Username	Name of the user who submitted the query	String

Trends Tab

The trend chart displays a visual representation of the magnitude of the selected metric over the selected time period. The horizontal axis represents the date, and the vertical axis represents the value of the metric selected for display. The green trend line plots the metric values that occurred during the selected time frame, until the day before the current date. As you move your mouse pointer over any point on the trend line, a balloon displays the actual magnitude of the selected metric on that date.

You can choose a daily or weekly aggregation period. Selecting a daily aggregation period plots the aggregate value each day. Selecting a weekly period shows the aggregate value calculated for the entire week, ending on the selected date. For example, the burn rate aggregated by week is calculated as the total number of CPU seconds consumed during the week divided by the system capacity of the system while the queries were active.

You can view a time frame of the last six months, the last year, or the last two years. You can also choose to display vertical lines indicating the date of Teradata system version, application version, and ruleset changes using the check boxes on the upper right.

By default, one chart is displayed. You can add  and remove  charts.

Trend Chart Metrics

Trends for the following metrics can be displayed.

Metric	Description
Burn rate	Percentage of the overall system capacity that was effectively consumed by all queries during their execution
Gating efficiency	Percentage indicating the overall query duration time spent executing versus being delayed. A value of 100% indicates all time was spent executing.
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the sum of the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.
Max spool	Maximum amount of temporary spool space used by any one query, using a 1024 byte kilobyte as the base
Query counts	Total number of logged queries
Suspect queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector
System utilization	Percentage of the system's CPU seconds that was utilized by logged queries
Time active	Amount of time in seconds that was recorded as utilized by logged queries. Excludes any time queries were delayed before execution.
Time delay	Amount of time in seconds that queries were delayed before being executed

Metric	Description
Time expansion	Ratio of how long queries took to execute versus the ideal execution time. The larger the number, the more the execution time of the query was expanded. For example, a value of 10x indicates an expansion of 10 times the ideal execution time.
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries
Total I/O	Total number of I/O requests
Total parser CPU	Total amount of CPU consumed by the parsing engine in processing queries

Query Details View

You can view query details by drilling down on any of the queries listed in Logged Queries or Suspect Queries tables.

Query details are displayed on an **Overview** and a **SQL** tab. If applicable, **Query Band** and **Errors** tabs are also displayed.

Viewing Application Query Details

You can display detailed information about a query, including its statement type, row count and active AMPs, its duration and performance, and workload and session information.

1. From the **Applications Queries** view, select the row of an application.
2. Select the **Queries** tab.
3. Select a query listed in the **Logged Queries** or **Suspect Queries** tab.

Overview Tab

The **Overview** tab on the query details view displays five sections containing information about the selected query. The sections and the metrics each displays are shown as following.

Query Information

Metric	Description	Type
Query ID	System-wide unique ID that identifies the query	String
Proc ID	ID of the parsing engine that runs the query	String
Start time	The time the query was submitted	Number
Statement type	SQL statement type the query used	String
Row count	Number of result rows returned by the query	Number
Active AMPs	Number of AMPs involved in executing the query	Number

Metric	Description	Type
Step count	Number of steps in the query plan	Number
Cache flag	The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated. • A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed. 	String

Query Duration

Metric	Description	Type
Time active	Number of seconds the query spent running. Excludes any time the query was delayed before being executed.	Number
Time delay	Number of seconds the query was delayed before running	Number
Total duration	Total elapsed time it took for the query to run once it was submitted	Number

Query Performance

Metric	Description	Type
AMP CPU time	Total number of CPU seconds consumed by AMPs in processing the query	Number
Parser CPU time	Total number of CPU seconds consumed by the parsing engine in processing the query	Number
Total I/O count	Total number of I/O requests the query used	Number
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the sum of the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
CPU skew	CPU skew for the entire query	Percent
I/O skew	I/O skew for the entire query	Percent
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time	Number
Spool usage	Amount of spool consumed by the query	Number
Gating efficiency	Percentage indicating the percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates all time was spent executing.	Percent

Metric	Description	Type
Burn rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent
Time expansion	Ratio of how long the query took to run versus how long it could have ideally run. The larger the number, the more the run time of the query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	Number

Workload Information

Metric	Description	Type
Starting workload	Name of the workload under which the query started execution	String
Final workload	Name of the workload under which the query finished	String

Session Information

Metric	Description	Type
Username	Name of the user who submitted the query	String
Application	Name of the application that submitted the query	String
Version	Version of the application that submitted the query	String
Account	Account of the user who submitted the query	String
App ID	The ID of the application used to access the Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.	String
Client ID	ID of the client that submitted the query	String
Client address	IP address of the client	String

SQL Tab

The **SQL** tab displays the SQL for the selected query. The SQL information is presented from the PDCRDATA.DBQLSqlTbl_Hst table, if available. If not available, the information from the PDCRDATA.DBQLLogTbl_Hst table is displayed.

The information is read-only.

By default, the SQL is formatted. Clear the **Format SQL** check box to display unformatted SQL.

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Errors Tab

The **Errors** tab has a text field that displays any errors for the selected query. The information is read-only. The tab is available only when the query resulted in an error.

Flagging a Query for Analysis

You can choose to flag a query for further analysis. This tags the query so that you can load it for analysis in Index Analyzer.

1. From the query details view, select the **Flag for analysis** check box at the top right.

Canary Response Times

Canary Response Times Overview

The **Canary Response Times** portlet allows you to analyze trends in single-system performance over a 60-minute or a 120-minute period using *canary queries*. Canary queries measure the time in milliseconds it takes for a user-defined query to run to completion. The same query is run repeatedly at a preselected interval to compare system workload over time. The longer the query takes to complete, the heavier the system workload at that time.

Canary response times can help you plan maintenance activities or schedule your workloads appropriately. For example, you can use this portlet to identify heavy workload periods on a Teradata system so that you can avoid assigning tasks when system performance is marginal. You can also use canary response times to compare current performance with average performance measured across one or more weeks of historical data. Additionally, you can monitor and compare the performance of multiple Teradata systems by adding a **Canary Response Times** portlet to the portal page for each monitored system.

The summary view displays a graphic overview of selected canary query metrics monitored over the previous 60-minute or 120-minute period. Samples are taken once every minute, and a data point is added to the graph for each selected canary query metric.

The Teradata Viewpoint Administrator defines default metrics for the **Canary Response Times** portlet in the **Monitored Systems** portlet. If enabled for your profile, you can set metrics that override the defaults.

The **Canary Response Times** portlet also gives you a means for comparing current and historical data by displaying past performance as an average of metric data points collected over time. The duration over which the averages are calculated can be changed using the **Past Averages** tab.

Canary Response Times View

The **Canary Response Times** portlet monitors the workload of a single Teradata system over time. You can monitor the performance of multiple systems by adding a **Canary Response Times** portlet instance for each Teradata system you want to monitor. The **Export** button allows you to create a .csv file containing selected data.

The summary view shows current and historical performance data using sparklines. The *sparkline* is a horizontal graph, showing the time (in milliseconds) a query takes to complete. Canary response times, such as **System Heartbeat**, are represented by a data point that is added to the sparkline once every 60 seconds. Hover over a sparkline to see an information balloon containing detailed information about the data point.

The sparkline types are:

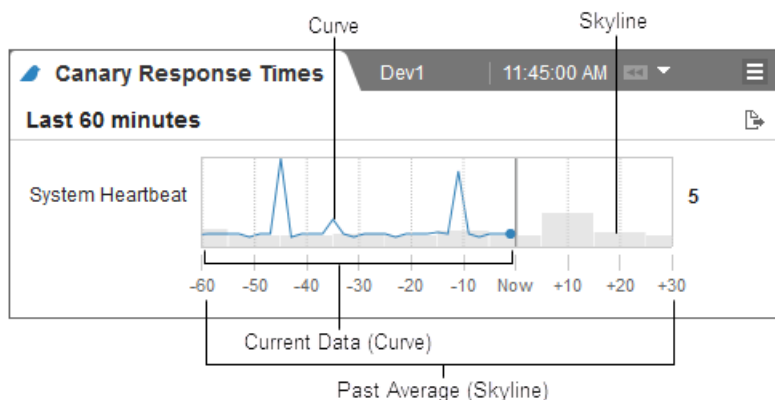
Current Data


Shows the time (in milliseconds) that a query takes to complete. Each data point shown on the sparkline is the average of all response times recorded during a 120-second interval, by default. The result is a curve sparkline. A large dot at the **NOW** point on the sparkline indicates the value of the last data point captured. A number to the right of the sparklines also indicates the **NOW** data-point value.

The Teradata Viewpoint Administrator sets the default data collection rate.

Past Average

Presents the averages for a user-selected number of weeks. The sparkline is shown as a skyline chart with flat, solid data points. Use this sparkline to compare the current performance from the previous 60-minute or 120-minute period with the system performance during the same 60-minute or 120-minute period 1 or more weeks in the past. Use the past-average data shown from **NOW** to 30 minutes or 60 minutes in the future to estimate system workload and enhance decision-making regarding workload assignments and resource allocation.



A 90-minute period is shown in the minimized view and a 180-minute period is shown in the maximized view. You can change these views from .

Minimized

The skyline shows 90 minutes of the average-workload data values from a user-selected number of weeks in the past. The curve sparkline shows the actual data for the past 60 minutes using 1-minute data points.

Maximized

The skyline shows 180 minutes of the average-workload data values from a user-selected number of weeks in the past. The curve sparkline shows the actual data for the past 120 minutes using 1-minute data points.


Canary Response Times Metrics

The metrics available for display are canary queries.

Metric	Default Sparkline	Past Average Line	Description
System Heartbeat	Curve	Yes	The predefined metric that is available for all systems. This is the default canary query.
<i>canary query name</i> CQ	Curve	Yes	Canary queries defined by the Teradata Viewpoint Administrator appear on the Metrics tab.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting.

1. Select .
2. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
3. [Optional] Reformat exported data to match the format in the view.

Settings View

The **Settings** view provides the following tabs that allow you to customize the portlet to monitor key metric trends that are important to efficient operation of the system.

Metrics


Select a system and the metrics to monitor, and define the threshold and vertical axis range for the selected metrics. Values that exceed the defined thresholds are highlighted in red.

Past Averages

Specify the time period used to calculate and display the shaded skyline graph behind each metric line graph.

Managing Default Settings





In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.


Setting Metrics Display

1. Click  in the portlet frame and select **Settings**.
2. Click the **Metrics** tab.
3. Select a metric from the list.
You can add  or remove  metrics.
4. [Optional] Enter a **Threshold** value.
5. [Optional] Enter a **Vertical Axis Range** value.
6. [Optional] Click  and drag the row to reorder the metrics for display.
7. Click **OK**.

Setting Past Averages

Specify the time frame used to calculate and display the shaded skyline graph in the details view.

Each segment of the skyline graph is the average of data samples taken each week at the same time and day of the week.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Past Averages** tab.
3. Enter the number of weeks up to 99.
4. Click **OK**.

Completed Queries

Completed Queries Overview

Use the **Completed Queries** portlet to view the final status of processes that ran on an Aster system or in Teradata Machine Learning Engine or Teradata QueryGrid. Use the portlet to analyze the performance of longer-running queries, to investigate processes that did not complete successfully, and to see how the system is being used in general.

Each process is a SQL query, SQL command, or a block of SQL statements. The statements might originate from Analytics Database, Teradata Machine Learning Engine, or QueryGrid, or contain Aster SQL-MapReduce functions.

Monitor multiple systems by opening additional instances of the **Completed Queries** portlet.

Completed Queries View

The **Completed Queries** view displays summary information about completed processes and queries so you can monitor and locate issues.

System Selection

Shows the name of the system currently selected and allows you to select a different system.

Time Selection

Shows the time frame of the information displayed in the summary table and allows you to select a different time frame.

State Filter Bar

Displays a count of the processes or queries in each state. Click a state to show only the processes or queries in that state.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information for all processes and queries. The view is refreshed when new data is collected. The table is configured in the **Configure Columns** dialog box. Click a row in the table to see details.

☐ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Time Selection

State Filter Bar

Filters

Summary Table

System Selection

Table Actions

Completed Queries

asterDev 3:04:43 PM

Last 24 hours

1,242 All

1,237 Completed

4 Error

0 Canceled

USER	STATE ...	TYPE	DURATION	PROCESS ID	SES:
Technicians	✓	SQLMR	0:01:10	8044180737928971	6580
db_superuser	✓	SQLMR	0:01:09	3941920660702607	4827
Administrator	✓	SQLMR	0:01:09	6437921897534634	4888
Database Admin	✓	SQLMR	0:01:09	4058366320186212	2428
Developer	✓	SQLMR	0:01:09	6177068692500197	3786

Page 1 of 5 (1242 rows total)

Changing the System or Time Frame

You can choose a different system or time frame to view processes.

- Select ▼ in the portlet frame and select a system.
- Select ▼ next to the time frame selection and select a time frame.

The data appears in the summary table.

State Filter Bar

The state filter bar allows you to display specific states in the view.

2,406 All

2,406 Completed

0 Error

0 Canceled

The state filter buttons provide, in near-real-time, a count of processes in each state. The buttons for a Teradata Aster system are listed as following.

All

Number of processes that finished running

Completed

Number of processes that ran successfully

Error

Number of processes that encountered a problem when running

Canceled

Number of processes that were stopped

The state filter buttons for a Teradata QueryGrid are listed as following.

All

Number of processes that finished running

Successful

Number of processes that ran successfully

Failed




Number of processes that failed and did not complete successfully

Unknown

Number of processes whose completion status is not known

State Icons

The following icons appear in the **Completed Queries** portlet.

Icon	Name	Definition
	Canceled	Process was stopped
	Completed	Process finished running
	Error	Error encountered during processing

Completed Queries View Metrics**Teradata Aster**

Metric	Description
User	User name that submitted the process
State	Process state, such as error, canceled, completed

Metric	Description
State Icon	Symbol of the process state
Type	Values for process types are as follows: <ul style="list-style-type: none"> • SQLMR (SQL-MapReduce) • SQL • Teradata Import • Teradata Export
Duration	How long the process ran
Process ID	Process identifier
Session ID	Session identifier
Start	Time the process began running
End	Time the process stopped running
Workload	Workload name in which the process ran
Database	Database name against which the process ran

Teradata Machine Learning Engine

Metric	Description
User	User name that submitted the process
State	Process state, such as error, canceled, completed
State Icon	Symbol of the process state
Type	Values for process types are as follows: <ul style="list-style-type: none"> • SQLMR (SQL-MapReduce) • SQL • Teradata Import • Teradata Export
Duration	How long the process ran
Process ID	Process identifier
Session ID	Session identifier
Start	Time the process began running
End	Time the process stopped running
CPU Percent	(8.10 or later)
CPU Seconds	(8.10 or later)
Memory KB	(8.10 or later)

Metric	Description
Total I/O	(8.10 or later)

Teradata QueryGrid

Metric	Description	Type
Start Time	Time the query started running	Number
User	Initiating user	String
Link Name	Name of the link that defines the initiating and target connectors	String
Status	Indicates the final status of the query	String
Duration	How long the query has been running	Number
Rows Transferred	Number of rows transferred by the query	Number
Bytes Transferred	Number of bytes transferred by the query	Number
Transfer Rate	Rate at which data is being transferred	Number
Initiating System	System that initiated the query	String
Initiating Connector	Connector that initiated the query	String
Target System	System being accessed or to which data is being transferred	String
Target Connector	Connector through which the data is being accessed or transferred	String

Details View for All Systems

The details view displays statistics and information about the selected process or query. This view is accessed by clicking a process or query row in the summary table.

Tabs

Provides important details about the **Overview**, **SQL**, and **Explain** tabs.

Details

Displays details of the selected process or query in sections that are specific to each tab.

< > Previous/Next

Allows you to move through processes or queries without returning to the summary table.

Previous/Next

Completed Queries asterDev 3:43:16 PM

Process ID: 8044180737928971087

Overview SQL Explain

QUERY INFO

State: COMPLETED
 Duration: 00:01:10
 Start: 2/18/14 10:26:17 PM
 End: 2/18/14 10:27:27 PM
 State Details: -

WORKLOAD INFO

Workload Policy: WKD-test
 Service Class: SC-Default
 Priority: MEDIUM
 Weight: 100

SESSION INFO

Session ID: 65803937021310461
 User: db_superuser
 Database: beehive
 IP Address: 127.0.0.1

Process Details

Overview Tab for Teradata Aster

The **Overview** tab displays detailed information about key metrics for the selected process. The metric values show the status of a Teradata Aster process.

Query Information

Query Information	Description
State	Process state, such as error, canceled, completed
Duration	How long the process ran, displayed as hh:mm:ss
Start	Time the process began running
End	Time the process stopped running
State Details	Additional information about errors in operating condition

Workload Information

Workload Information	Description
Workload Policy	Name of the workload policy used to manage the process. Defines a set of related queries, and allows them to be prioritized in a similar manner.

Workload Information	Description
Service Class	Name of the service class used by the workload policy. Specifies the share of system resources allocated to a workload as a function of the class's priority and weight values.
Priority	Indicates the importance of the process. <ul style="list-style-type: none"> • High • Medium • Low
Weight	Number ranging from 1 to 100 that indicates the precedence of the process within the priority level, increasing the resources that can be allocated for the process. A higher value indicates a greater level of precedence.

Session Information

Session Information	Description
Session ID	Unique ID for the session
User	Name of user that submitted the query
Database	Name of the database against which the process ran
IP Address	IP address of the user

Overview Tab for Teradata Machine Learning Engine

The **Overview** tab displays detailed information about key metrics for the selected process.

Query Information

Query Information	Description
State	Process state, such as failed, canceled, or completed
Duration	How long the process ran, displayed as hh:mm:ss
Start	Time the process began running
End	Time the process stopped running
State Details	Additional information about errors in operating condition


Session Information

Session Information	Description
Session ID	Unique ID for the session
User	Name of user who submitted the query
IP Address	IP address of the user

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. This information is read-only.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.


Explain Tab

The **Explain** tab displays each Explain step of the query and allows you to export the Explain data.

To view all steps and phase details, you can select **All** from the **Show** list. To view less Explain data, select **Summary**.

Step Information	Description
Step Status Icon	Indicates the status of each Explain step. In the Completed Queries portlet, the icon is always black, indicating completed.
Time	Duration of the step
Type	Category of step such as Command, CommitPreparedTransaction, DataTransfer, Import, PrepareTransaction, and Query.
Phase/Detail	Summarized or full explanation of query that executes in this step. Can be used to diagnose the cause of the failure or slowness.

Exporting Data

1. From the **Explain** or **Message** tab, select .
2. Save the file to a location.

The data is exported as a .csv file.

Details View for Teradata QueryGrid

When you select a Teradata QueryGrid session from a database session or a Teradata QueryGrid manager in the portlet's summary view, the details for the query are displayed.

The top portion of the screen displays a graphical representation of the progress and status of the query, with icons representing the initiating system, the phase of the current operation being performed, and the target system. The initiating system is the data source from which the query is generated. The target system is the system against which the query is being run. Teradata QueryGrid queries go through a metadata and execution phase before they are completed. Teradata QueryGrid queries may also access data on target systems or transfer data from one system to another.

The lower part of the screen displays tabs with information about the system or operation. If you arrive at the details screen from a Teradata system session, details are displayed for the target or initiator system that is relevant for the selected step. If you arrive at the details screen while viewing a list of Teradata QueryGrid manager sessions, details for the selected operation are displayed. You can click on any of the icons in the graphical representation (initiator, operator, or target) to view the metrics for the selected system or operation.

Initiating System

The data source from which the query is generated.

Target System

The system from which data is accessed or to which it is transferred.

Operation

The phase of the operation being performed.

Tabs

Related metrics are organized into tabs for the system or operation selected in the top portion of the view.

System or Operation Details

The metrics related to the system or operation selected in the top portion of the view.

Query Monitor 10:45:17 AM

QueryGrid Operation Details

Initiating System: sdt04625 OPERATION: 1111 Target System: sdt04935

Operation Details

Tabs: Overview | Phases | Configuration


GENERAL		TRANSFER INFO	
Query ID:	4d32ec6b-659d-48ab-a159-0000000001c0	Type:	Import
Phase:	Data transfer	Rows:	14.2M
Duration:	1:04:57.098	Bytes:	176MB
Config version:	Active	Rate:	9.49 MB/sec
Link name:	I4625-4935-b1	Bottleneck:	-
		Compression ratio:	-

INITIATOR INFO		TARGET INFO	
System:	sdt04625	System:	sdt04935
Total nodes:	1	Total nodes:	1
Total CPU:	0	Total CPU:	0.15
CPU skew:	1.0	CPU skew:	1.0
Data skew:	1.0	Data skew:	1.0

[QueryGrid systems ≥ 2.10] If a query fails, an option to download a query support bundle appears.

Completed Queries sdt49644 11:47:15 AM

QueryGrid Operation Details

Initiator: sdt33816 OPERATION:  Target: sdt34738

Operation Details

Tabs: Overview | PHASES | CONFIGURATION | ERRORS

[Download Query Support Bundle](#)

GENERAL		TRANSFER INFO	
Query ID:	dfa91220-f032-4922-9233-000000000001	Type:	-
Phase:	Failure	Rows:	-
Duration:	0:00:00.063	Bytes:	-B
Config version:	Active	Rate:	-
Link name:	hkcon1_hkcon2	Bottleneck:	-
		Compression ratio:	-




INITIATOR INFO		TARGET INFO	
System:	sdt33816	System:	sdt34738
Total nodes:	-	Total nodes:	-
Total CPU:	-	Total CPU:	-
CPU skew:	1.0	CPU skew:	1.0
Data skew:	1.0	Data skew:	1.0

QueryGrid Icons






The details view for Teradata QueryGrid includes a graphical representation of the progress of the selected query. The icons here represent the initiating or target systems and the phases of the operations performed.


System Icons

The icons in the following table represent the query on the initiator and target systems. The icons inside the smaller circle represents the query state. In addition to the query states shown as following, additional possible states for the query are listed in [State Icons](#).

Icon	Name	Definition
	No query is running	No query is currently running
	Active	The query is running on the system
	Idle	The query is currently idle

Operation Phase Icons

Icon	Name	Definition
	Metadata Operation	A handshake between systems to negotiate data types and validate queries.
	Execution Operation	Query processing prior to transferring the data.
	Data transfer	The transferring of data from one system to another.
	Completed	Query has successfully finished processing.
	Failed	Query did not finish processing successfully. [QueryGrid systems ≥ 2.10] To download the support bundle, click Download Support Query Bundle .

Icon	Name	Definition
	Unknown	Query processing status is not currently available.

Initiator and Target Metrics

When an initiator or target system icon is selected in the details view for Teradata QueryGrid, two tabs display metrics for the system: an **Overview** tab and a **SQL** tab.

Overview Tab for Teradata System

The **Overview Tab** displays information for a QueryGrid query that was initiated by or is a target of a Teradata system.

QueryGrid Information

Metric	Description	Type
System	Name of the initiating or target system	String
Connector	Point from which the query originates or ends	String
Software version	The version of the connector	String
Network	The logical network used for this side of the data transfer	String

Session Information

Metric	Description	Type
User	User who submitted the query	String
Query ID	ID of the query	String
Session Number	Number of the session	String
Request Number	Number of the request	String
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session	String

Query Information


Metric	Description	Type
State	State of the query	String

Metric	Description	Type
Duration	Amount of time the query has been running (HH:MM:SS)	Number
Workload	Workload in which the query is running	String
Request CPU	Total request CPU (seconds)	Number
Request I/O	Total request I/O count	Number

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Operations Metrics

When an operation phase icon is selected in the details view for Teradata QueryGrid, three tabs display metrics for the operation: an **Overview** tab, a **Phases** tab, and a **Configuration** tab.

Overview Tab

The **Overview** tab displays general information about the selected operation for a Teradata QueryGrid query session.

General

Metric	Description	Type
Query ID	ID of the query	String
Phase	Phase of the operation	String
Duration	Duration of the QueryGrid operation	Number

Metric	Description	Type
Config version	Requested version of the QueryGrid operation	Number
Link Name	Name of the link used for this QueryGrid operation	String

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Total nodes	Number of nodes on the initiating system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Transfer Information

Metric	Description	Type
Type	Type of transfer performed	String
Rows	Number of rows transferred	Number
Bytes	Number of bytes transferred	Number
Rate	Rate that data is transferred for the QueryGrid operation	Number
Bottleneck	Limiting factor for the QueryGrid operation	String
Compression ratio	Uncompressed size divided by the compressed size	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Total nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Total Nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Phases Tab

The **Phases** tab is available when you click an icon that represents a Teradata QueryGrid query operation. It contains tabs that provide information about the metadata, execution, and data transfer phases of the query session. The **Phases** tab support up to five panes depending on the details associated with the operation, for example, Metadata1, Metadata2, Metadata3, Execution1, Execution2, Data Transfer.

Metadata Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	The node that initiated the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String

Metric	Description	Type
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Execution Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	Node from which the query initiated	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String

Metric	Description	Type
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Data Transfer Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	String
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Metric	Description	Type
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the max CPU	String
Data skew	Node that had the max data	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Total nodes	Number of nodes	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	Number
Data skew node	Node that had the maximum amount of data	Number
Bytes sent	Number of bytes sent	Number

Metric	Description	Type
Bytes received	Number of bytes received	Number
Compression ratio	Whether or not compression is enabled	String

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Configuration Tab

The **Configuration** tab contains provides configuration information about the selected operation for a QueryGrid query session.

Fabric

Metric	Description	Type
Name	Name of the fabric used in Teradata QueryGrid	String
Port	Port on which the fabric runs	String
Software version	Version of software that the fabric runs	String

Initiating Connector

Metric	Description	Type
System	Name of the initiating system	String
Connector	Point from which the query originated	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Communication Policy

Metric	Description	Type
Name	Name given to this communication policy	String
Encryption	Whether or not data encryption is enabled	String
Compression	Whether or not compression is enabled	String

Target Connector

Metric	Description	Type
System	Name of the target system	String
Connector	Destination point for the query	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Errors Tab

The **Errors** tab provides information about Teradata QueryGrid queries that had errors. It displays only if there are errors.

Metric	Description	Type
Time	Time the error occurred	Number
Phase	Phase when the error occurred	String
Message	The error message	String
Node	The hostname of the node that generated the error	String
Connector	The name of the connector that resulted in the error	String
System	Name of the system	String
Origin	Indicates whether the error occurred on the initiator or the target of the Teradata QueryGrid operation	String

Elastic Performance

Elastic Performance Overview

The **Elastic Performance** portlet allows you to view your Teradata system usage and identify periods where usage exceeds licensed capacity. You can export data to assist in verifying that pay per use charges reflected in the billing invoice from Teradata Corporation are accurate. This portlet is disabled by default and must be enabled using **Portlet Library**.

The **Settings** view allows you to set the capacity baseline that is specified in your TPERF Capacity on Demand Order Addendum. If you need additional information or assistance, contact Teradata Customer Services.

Elastic Performance View

The **Elastic Performance** view provides a graph of system usage over different time periods for each node type on the Teradata system, allowing you to see when usage exceeds baseline. The Elastic Usage collector must be enabled to see data in the view.

The following list describes the features in this view:

Selection Menus

Shows and allows you to change the system and the graph type being displayed. Use the selection menus to select the following graph types:

Average CPU Usage: Displays bars with the average CPU usage per period, depending on the option selected in **Monitored Systems > Collectors > Elastic Usage**. The percentage of CPU usage after baseline is highlighted in the graph.

CPU seconds above baseline: Displays CPU seconds after baseline per day in the month graph or per week in the quarter graph.

Toolbar

Displays the date range. The dates and times represent system usage displayed in the graph. The **Export** button allows you to create a .csv file containing data for a customized date range.

System Usage

Displays the actual system usage with one graph for each node type on the system. As you hover on the bars in the graph, information balloons display details about the data for the time period. The bar at the right side of the **CPU seconds above baseline** graph will be shorter when the entire time period is not completed and data is still being gathered.

System CPU Limit: The CPU limit set on the Teradata system, which displays after the Elastic Limit data collector is enabled. You can set the CPU limit using the **Workload Designer** portlet or the Schmon utility.

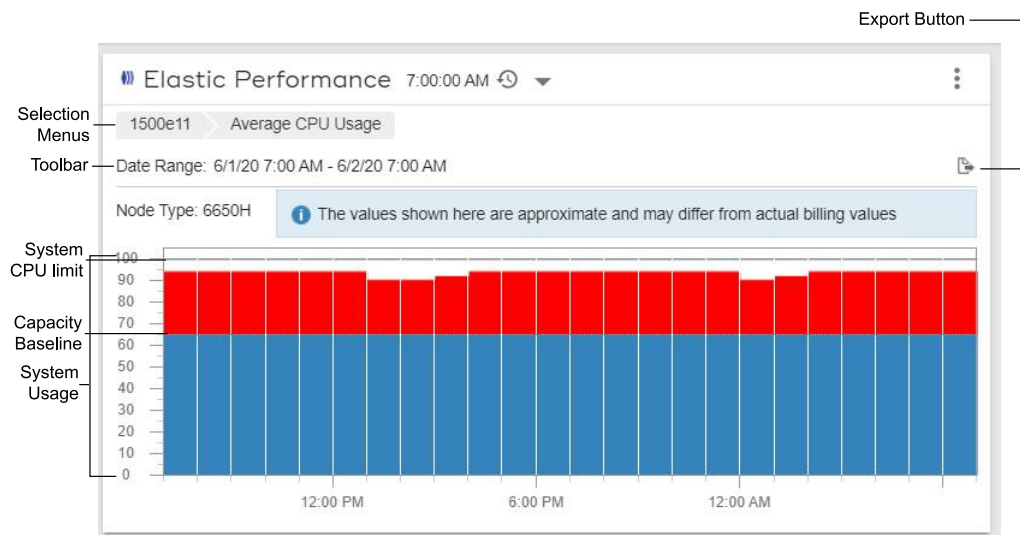
Capacity Baseline: The capacity baseline is set in the **Settings** view. Licensed capacity is shown under the line. Usage after the line, displayed in red, is subject to pay-per-use charges.

Data Gap: A period of time when no usage data was available. Bars containing data gaps are represented by a lighter color. The total amount of time during which data was not available is shown in the information balloon as you hover over the bar.

Information Balloon: The information balloon for the **CPU seconds above baseline** graph displays the following, depending on the type of system and the settings you have chosen:

- **CPU seconds above baseline:** The total number of CPU seconds used that were after the baseline in the day or week period.
- **Hours above baseline:** The count of all hours in the day or week period that had an average CPU usage after the baseline.
- **Average CPU when above baseline:** The average CPU usage for all hours in the day or week period that had an average CPU usage after the baseline.

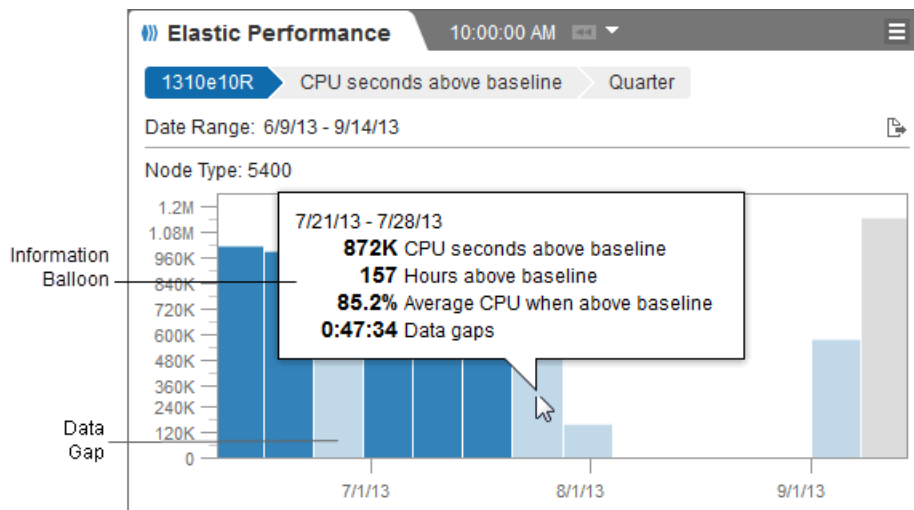
Following is an example of the **Average CPU Usage per hour** graph, where the bars represent the average percent of CPU usage per hour.



Note:

If **Elastic Performance Report** is enabled for Elastic Usage Collector, the **Elastic Performance** portlet displays the data as per the custom report enabled for the “Average CPU Usage” report. The bar graph shows the data for custom performance report that was enabled. An info message is added beside the node type “The values shown here are approximate and they may differ from the actual billing values”.

Following is an example of the **CPU seconds above baseline** graph, where the bars represent the CPU seconds after baseline for a quarter.




Selecting a System

You can use selection menus to choose a system and graph type. Menu choices are based on what you selected in the previous menu.

1. In the selection menu, click the currently selected system name to display a list of available systems.
2. Select a system.
Only systems that have the Elastic Usage collector enabled are included in the system selection.
3. Select >.
4. To select a graph type, do one of the following:
 - Select **Average CPU Usage**.
 - Select **CPU seconds above baseline**, select > , and then select **Month** or **Quarter**.
5. Select >.

Exporting Usage Data

You can export usage data for further analysis. The date range shown in the toolbar is the default period for which data is exported. Start and end dates are inclusive. The time zone used for the export date range is the Teradata system time zone.

1. In the selection menu, select a system and graph type.
2. Click .
3. [Optional] Change the start and end date.
4. Select one of the following export options:

Export File Type	Description
Summary Data (CSV)	Exports aggregate data from the graph for viewing.
Summary Data (ZIP)	Exports aggregate data from the graph for billing purposes.
All Data (ZIP)	Exports all available data points from the selected time period for billing purposes.


5. Click **Export**.
6. Save the file to a location.

Settings View

The **Settings** view allows you to set the capacity baseline, which should be based on your TPERF Capacity on Demand Order Addendum.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Setting Capacity Baseline

The percentage you enter in the **Settings** view is the capacity that you paid for and is displayed in the **Elastic Performance** view.

1. Click  in the portlet frame and select **Settings**.
2. Enter the capacity baseline as specified in your TPERF Capacity on Demand Order Addendum.
3. Click **OK**.

Hadoop Services

Hadoop Services Overview

The **Hadoop Services** portlet enables you to monitor service status, system use, and key metrics running on a Hadoop system (Cloudera or Hortonworks). You can monitor current and past performance over a set time frame.

Hadoop Services View

The **Hadoop Services** view displays summary information about the services on a Teradata Hadoop system so you can monitor and locate issues.

Summary View

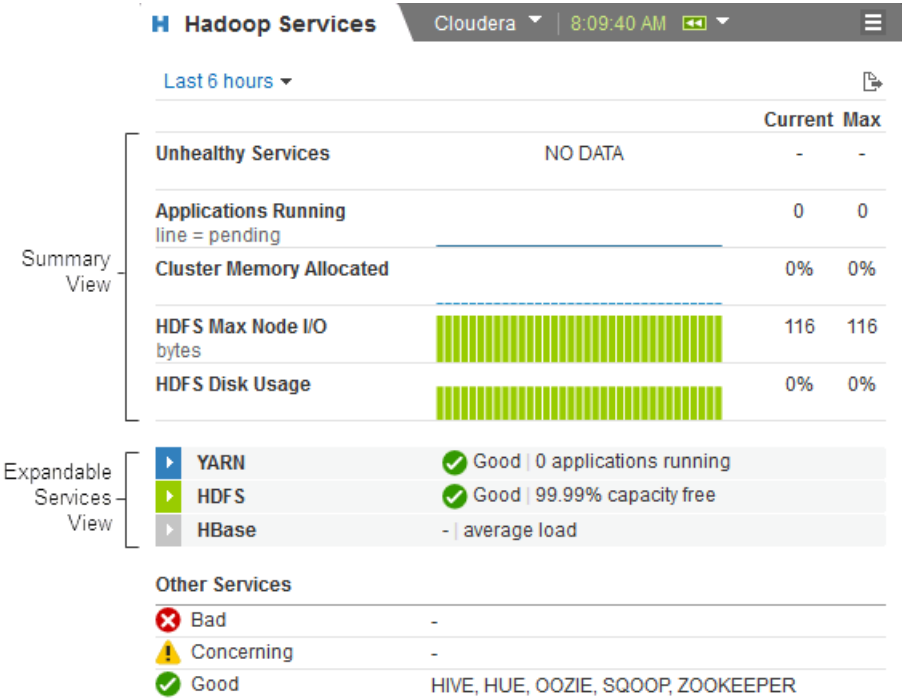
Displays summary information for all Hadoop services.

Expandable Services View

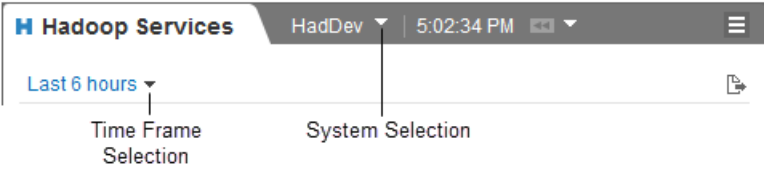
Expands the following services to see related metrics, which provide further information on the selected service:

- YARN (CDH and HDP 2.1 and later)
- MapReduce (HDP 1.3)
- HDFS
- HBase

Following is an example of the CDH **Hadoop Services** view.




Viewing Hadoop Services



- 1. From the **Hadoop Services** view, select ▼ to select a Hadoop system.
- 2. Select ▼ to select a time frame.

Time Frame	Intervals (Calculated Average)
Last hour	1 minute
Last 3 hours	5 minutes
Last 6 hours	10 minutes
Last 24 hours	30 minutes
Last 48 hours	1 hour
Last week	4 hours
Last month	12 hours

3. To view details about a particular service, select  next to any of the following:

Option	Description
YARN (CDH and HDP 2.1 and later)	Lists applications running, ResourceManager statistics, and NodeManager statistics.
MapReduce (HDP 1.3)	Lists jobs running, jobtracker statistics, map and reduce slots totals and live tasktrackers.
HDFS (CDH and all supported HDP versions)	Lists space usage, namenode statistics, datanode statistics and a total count of HDFS files and directories.
HBase (CDH and all supported HDP versions)	Lists load average, region servers statistics and the master server statistics.

4. [Optional] Select **Open in new window**.

The Hadoop web interface for the service appears, allowing you to view additional details about the service.

Understanding the Summary View


The summary view displays a high-level overview of the performance trends on the selected Hadoop system. Hover over a graph to see an information balloon containing detailed information about the data point for the interval. The right-most bar provides the current reading.

The **Current** column shows the current reading. If you are using the Rewind functionality, the column displays the most recent reading in the time frame defined by Rewind.

The **Max** column shows the maximum value in the selected time frame. This is the actual reading, not a calculated average.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains graph data for the time frame.

1. From the **Hadoop Services** view, select a system and a time frame.
2. Click .
3. Save the file using the browser options.

The file is saved to your download area or to a location that you specify, depending on the browser settings.

Teradata Hadoop Services Metrics

A 1024 byte kilobyte is used as the base for all memory-related and space-related metrics that display, such as spool space, disk space, memory usage, and so on. A 1000 byte kilobyte is used as the base for all other metrics.

CDH

Summary View Metrics

Metric	Description
Unhealthy Services	Aggregate number of all CDH services in a Bad or Concerning state (Unknown and Disabled are not included)
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

YARN Metrics

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Completed	Number of YARN applications that executed successfully
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
Cluster Memory Reserved	Percent of available memory reserved across all NodeManager instances
Cluster Memory Skew	Comparison of the largest NodeManager memory allocated to the average memory allocated
Containers Allocated	Number of YARN containers currently allocated across the cluster
Containers Pending	Number of YARN containers currently pending across the cluster
Containers Reserved	Number of YARN containers currently reserved across the cluster
NodeManagers	Number of nodemanagers in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states
ResourceManager Up Since	Timestamp when the ResourceManager service started
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
ResourceManager UI	Open web interface for the service in a new window

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes	Number of datanodes in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region Servers	Number of region servers in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states
Master Server Up since	Timestamp when the master server started
Master Server Heap	Percentage of heap space used in the master server JVM
Master Server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Bad	Services in a critical state
Concerning	Services in a degraded state
Disabled	Services in a disabled state when there are one or more in this state
Good	Services in a good state
Unknown	Services in a unknown state when there are one or more in this state

HDP 2.1 and Later

Summary View Metrics

Metric	Description
Components Down	Number of components not started
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

YARN Metrics

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Completed	Number of YARN applications that executed successfully
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
Cluster Memory Reserved	Percent of available memory reserved across all NodeManager instances
Cluster Memory Skew	Comparison of the largest NodeManager memory allocated to the average memory allocated
Containers Allocated	Number of YARN containers currently allocated across the cluster
Containers Pending	Number of YARN containers currently pending across the cluster
Containers Reserved	Number of YARN containers currently reserved across the cluster
NodeManagers Live	Number of NodeManagers currently running
NodeManagers Total	Number of NodeManagers configured on the system
ResourceManager Up Since	Timestamp when the ResourceManager service started
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
ResourceManager UI	Open web interface for the service in a new window

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes Live	Number of datanode processes currently running
Datanodes Dead	Number of datanode processes that are currently dead
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region servers Live	Number of region servers that are currently running
Region servers Total	Number of region servers that are configured
Master server Up since	Timestamp when the master server started
Master server Heap	Percentage of heap space used in the master server JVM
Master server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Up	Hadoop Services currently running
Down	Hadoop Services not running
Client Only	Services that do not execute a daemon process in the Hadoop cluster

HDP 1.3**Summary View Metrics**

Metric	Description
Components Down	Number of components not started

Metric	Description
Map Tasks Running	Number of map tasks running Number of map tasks waiting in the queue
Reduce Tasks Running	Number of reduce tasks running Number of reduce tasks waiting in the queue
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

MapReduce Metrics

Metric	Description
Jobs Running	Number of MapReduce jobs currently executing
Jobs Completed	Number of MapReduce jobs that completed successfully
Jobs Failed	Number of MapReduce jobs that failed to complete successfully
Jobtracker Up since	Timestamp when the jobtracker service started
Jobtracker Heap	Percentage of heap space used in the jobtracker JVM
Jobtracker UI	Open web interface for the service in new window
Map slots Occupied	Number of map slots occupied in the MapReduce cluster
Map slots Reserved	Number of map slots reserved in the MapReduce cluster
Reduce slots Occupied	Number of reduce slots occupied in the MapReduce cluster
Reduce slots Reserved	Number of reduce slots reserved in the MapReduce cluster
Tasktrackers Live	Number of tasktrackers currently running
Tasktrackers Total	Number of tasktrackers available on the system

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes Live	Number of datanode processes currently running
Datanodes Dead	Number of datanode processes that are currently dead
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM

Metric	Description
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region servers Live	Number of region servers that are currently running
Region servers Total	Number of region servers that are configured
Master server Up since	Timestamp when the master server started
Master server Heap	Percentage of heap space used in the master server JVM
Master server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Up	Hadoop Services currently running
Down	Hadoop Services not running

Lock Viewer

Lock Viewer Overview

The **Lock Viewer** portlet allows you to view the lock data such as transaction identifiers, session identifiers, lock object identifiers, and global deadlocks. Viewing lock information helps you determine whether system performance has been degraded by a database lock.

You can locate locks in a specific database, in a specific time frame, caused by a specific user, or that are blocking a specific user. You can also view details about selected sessions, including details about other queries that are blocking the session you select, what time the query started, and when the query became blocked.

Lock Viewer View

The **Lock Viewer** view displays a summary of information about the database locks that have occurred in the last 5 minutes, 1 hour, 24 hours, or 1 week. Lock data is displayed in the table only after the lock contention is resolved. Real-time lock contention data is not available. The Teradata Viewpoint Administrator must enable the appropriate collectors to display data in the **Lock Viewer** view.

The following list describes the features in this view:

System Selection

Select the system whose lock information you want displayed.

Time Frame Selection

Select the time frame for which you want the lock information displayed.

You can use the rewind feature with any time frame selection to retrieve lock data and compare it to data for a different date and time. You are less likely to see data for the last 5 minutes unless you use rewind.

Filters

Changes the displayed data to show only rows that match your filter criteria. You can use the filters to search for locks in a specific database, caused by a specific user, or that are blocking a specific user.

Summary Table

Displays lock data for transaction identifiers, session identifiers, lock object identifiers, and global deadlocks. The table is configured in the **Configure Columns** dialog box. You can click a row in the table to see statistics and information about the delayed query.

Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

BLOCK TIME	DATABASE	TABLE	HOST	MODE	SESSION
9/14/16 8:10:00 AM	DBC	EventLog	0	Write	0
9/14/16 10:34:59 AM	DBC	EventLog	0	Write	0
9/14/16 10:34:59 AM	DBC	EventLog	0	Write	0
9/14/16 2:11:41 PM	DBC	EventLog	0	Write	0
9/14/16 2:11:41 PM	DBC	EventLog	0	Write	0
9/14/16 8:23:23 PM	DBC	TempTables	0	Read	0
9/14/16 9:45:09 PM	DBC	EventLog	0	Write	0
9/14/16 11:08:06 PM	DBC	EventLog	0	Write	0
9/15/16 5:08:03 AM	DBC	EventLog	0	Write	0
9/15/16 5:19:30 AM	DBC	EventLog	0	Write	0
9/15/16 1:55:51 PM	DBC	EventLog	0	Write	0
9/15/16 2:45:51 PM	DBC	TVM	1	Read	2513517

87 rows total

Changing the System or Time Frame

You can choose a different system or time frame to view the lock data, such as transaction identifiers, session identifiers, lock object identifiers, and global deadlocks. Lock data displays in the table only after the lock contention is resolved. Real-time lock contention data is not available.

1. Select ▼ in the portlet frame and select a system.
2. Select ▼ next to the time frame selection and select a time frame.

The lock data, session identifiers, lock object identifiers, and global deadlocks for that system and time frame appears in the summary table.

Summary Locks

This topic describes the information that appears in the summary table.

Column	Description		Type
Blocked User	User name whose query is blocked by another user		String or Number
Blocking Host	Logical host ID of the transaction that imposed the lock		Number
Blocking Level	Lock level of the transaction that imposed the lock. The following values indicate the level.		String
	Database	Locked at the database level	
	Table	Locked at the table level	
	Row hash	Locked at the row level	
	Table partition range	Locked at the table partition level	
	Row hash in partition range	Locked at the row level in the partition range	
	Row hash in one partition	Locked at the row level in one partition	
	Row key range	Locked for a range of rows	
Blocking Mode	Lock level mode of the transaction that imposed the lock. The following values indicate the mode for Teradata Database 14.10 and later.		String
	Access	Access	
	IAccess	Intentional Access	
	Exclusive	Exclusive	
	IExclusive	Intentional Exclusive	
	Read	Read	
	IRead	Intentional Read	
	ReadHL	Read HUT (Host Utility) Lock	
	IReadHL	Intentional Read HUT (Host Utility) Lock	
	Read4X	Read for Exclusive	
	IRead4X	Intentional Read for Exclusive	
	Read4W	Read for Write	
	IRead4W	Intentional Read for Write	
	Write	Write	
IWrite	Intentional write		
Blocking Session	Session number of the transaction that imposed the lock. The following session numbers are used to indicate database internal sessions.		Number

Column	Description		Type														
	<table><tr><td>0</td><td>System User</td></tr><tr><td>1</td><td>System Accounting</td></tr><tr><td>2</td><td>System Recovery</td></tr><tr><td>3</td><td>Archive/Restore</td></tr></table>	0	System User	1	System Accounting	2	System Recovery	3	Archive/Restore								
0	System User																
1	System Accounting																
2	System Recovery																
3	Archive/Restore																
Blocking User	Username that is blocking the query of another user		String														
Block Time	Date and time the session became blocked, displayed as <i>mm/dd/yy hh:mm:ss</i>		Number														
Database	Database name on which the lock was requested		String or Number														
Deadlock	<p>Indicates whether the lock contention results in a deadlock. The following values indicate the type of deadlock.</p> <table><tr><td>N</td><td>No deadlock</td></tr><tr><td>Y</td><td>Local deadlock</td></tr><tr><td>G</td><td>Global deadlock</td></tr></table>		N	No deadlock	Y	Local deadlock	G	Global deadlock	String								
N	No deadlock																
Y	Local deadlock																
G	Global deadlock																
Delay	Total time the transaction waited for the block, displayed as <i>HH:MM:SS</i>		Number														
Host	Logical host ID of the transaction that was waiting for the lock		Number														
Level	<p>Lock level of the transaction that was waiting for the lock. The following values indicate the level.</p> <table><tr><td>Database</td><td>Locked at the database level</td></tr><tr><td>Table</td><td>Locked at the table level</td></tr><tr><td>Row hash</td><td>Locked at the row level</td></tr><tr><td>Table partition range</td><td>Locked at the table partition level</td></tr><tr><td>Row hash in partition range</td><td>Locked at the row level in the partition range</td></tr><tr><td>Row hash in one partition</td><td>Locked at the row level in one partition</td></tr><tr><td>Row key range</td><td>Locked for a range of rows</td></tr></table>		Database	Locked at the database level	Table	Locked at the table level	Row hash	Locked at the row level	Table partition range	Locked at the table partition level	Row hash in partition range	Locked at the row level in the partition range	Row hash in one partition	Locked at the row level in one partition	Row key range	Locked for a range of rows	String
Database	Locked at the database level																
Table	Locked at the table level																
Row hash	Locked at the row level																
Table partition range	Locked at the table partition level																
Row hash in partition range	Locked at the row level in the partition range																
Row hash in one partition	Locked at the row level in one partition																
Row key range	Locked for a range of rows																
Mode	<p>Lock level mode of the transaction that was waiting for the lock. The following values indicate the mode for Teradata Database 14.10 and later.</p> <table><tr><td>Access</td><td>Access</td></tr><tr><td>IAccess</td><td>Intentional Access</td></tr><tr><td>Exclusive</td><td>Exclusive</td></tr></table>		Access	Access	IAccess	Intentional Access	Exclusive	Exclusive	String								
Access	Access																
IAccess	Intentional Access																
Exclusive	Exclusive																

Column	Description		Type
	IExclusive	Intentional Exclusive	
	Read	Read	
	IRead	Intentional Read	
	ReadHL	Read HUT (Host Utility) Lock	
	IReadHL	Intentional Read HUT (Host Utility) Lock	
	Read4X	Read for Exclusive	
	IRead4X	Intentional Read for Exclusive	
	Read4W	Read for Write	
	IRead4W	Intentional Read for Write	
	Write	Write	
	IWrite	Intentional write	
Multiple Blocker	Indicates whether more than one transaction encountered the same lock contention. The following values indicate the status.		String
	Y	Two or more blocked transactions	
	N	One blocked transaction	
Processor ID	AMP identifier where the lock was requested		Number
Session	Session number of the transaction that was locked. The following session numbers are used to indicate database internal sessions.		Number
	0	System User	
	1	System Accounting	
	2	System Recovery	
Table	Table name on which the lock was requested		String or Number

Details View

The details view displays statistics and information about the selected session. This view can be accessed by clicking a session row in the summary table.

Tabs

Teradata systems have **Overview**, **SQL**, **Explain**, **Skew**, **Query Band**, and **Blocked By** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

Lock Viewer Dev1 10:03:17 AM

Session: **2496172**

Tabs: **Overview** SQL Explain Skew Query Band Blocked By

Query Details

QUERY INFO		WORKLOAD INFO	
State:	Blocked	Name:	WD-Default
Block time start:	9/14/16 5:46:43 AM	Method:	Timeshare
Time in state:	0:10:59	CPU decay:	Level 0
Query start time:	9/14/16 5:46:42 AM	Classification mode:	Auto
Total duration:	0:10:59	Virtual partition:	Standard
Spool space:	-	I/O decay:	Level 0
Hot AMP spool:	-		
Spool skew:	-	SESSION INFO	
Temp space:	0	User:	DBC
Request CPU:	0	Account:	DBC
Request I/O:	-	Partition:	DBC/SQL
PJI:	0	Requests:	2
Unnecessary I/O:	0	Source:	(TCP/IP) 912b 10.25.35.60 SDLL9643.LABS. TERADATA.COM 29573 ROOT BTEQ 01 LSS
SNAPSHOT INFO			
CPU use:	0%		
Impact CPU:	0		
Snapshot CPU skew:	0%		
Snapshot I/O skew:	0%		

Overview Tab for Terata System

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Block Time Start	Date and time the session became blocked, displayed as <i>mm/dd/yy hh:mm:ss</i> .
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Query Start Time	Date and time the query started, displayed as <i>mm/dd/yy hh:mm:ss</i> .
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using

Query Information	Description
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	(Teradata Database 15.0 and earlier) CPU impact on the system based on the highest utilized AMP
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running

Workload Information	Description
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11


Session Information

Session Information	Description
User	Name of the user that submitted the query
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system


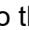
SQL Tab


The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.





Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Explain Tab for Teradata System


The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status. If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

If you have the correct permissions, use  to retrieve the latest Explain steps from the database. Click **Refresh** to update. The refresh screen data is not maintained once the refresh view closes. Repeatedly retrieving current Explain data can impact system performance.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	 - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step

Step Information	Description
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Exporting Explain Data

1. From the Explain tab, click .
2. Save the file to a location.
The Explain data is exported as a .csv file.

Skew Tab for Teradata System

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Query Band Tab for Teradata System

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Blocked By Tab for Teradata System

When a query is blocked, use the **Blocked By** tab to identify the other queries blocking the selected query. By reviewing the number and type of queries blocking the selected query, you can identify blocking issues on the system, determine the impact of this blocking on the selected query and, based on this information, decide the best course of action to resolve the issue.

The information in this tab is read-only. The tab is available only when the selected query is blocked. You can drill down into the session information for the blocking session.

The following information is available in the noted versions of Analytics Database.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example:		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
	<p>Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName</p> <p>The information displays in the following order:</p> <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock 		
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

MAPS Manager

MAPS Manager Overview

The **MAPS Manager** portlet allows you to incrementally redistribute data from tables across new AMPs after a system expansion instead of taking a system offline for an extended time. You can also actively monitor systems and provide recommendations to make sure objects are in optimal maps. Using appropriately sized maps increases parallelism and speed when accessing objects and provides greater efficiency for smaller objects.

MAPS Manager manages the following objects:

- Tables
- Hash indexes
- Join indexes

The portlet provides two methods to move objects to a new map. Both methods allow you to improve system performance without extended downtime. You can also view all maps defined on the system and to add sparse maps as needed.

System Monitoring

System Monitoring provides ongoing recommendations to optimize the location of objects. You can accept recommendations and schedule moves when convenient.

Object Workflows

Object workflows let you create ad-hoc workflows where you can select which objects to move and the order and time to move the objects.

Maps

Maps shows all available contiguous and sparse maps on the system. It also identifies the system default contiguous map and its sparse maps that are used for analysis.

You can also grant and revoke permissions for the maps on your system.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

MAPS Manager Exceptions

Object selection does not include objects from the DBC and TDMaps databases. The database tools exclude stored procedures, such as AddObjectListEntrySP.

MAPS Manager Best Practices

Method	Description
System Monitoring	<ul style="list-style-type: none"> • Set up System Monitoring to periodically evaluate map usage. • Use after changing the system default map. Set up the analysis to run again. • Make sure the user credentials have access to all objects and maps in the system. • Set up DBQL settings. • Schedule moves during system off-peak hours.
Object Workflows	<ul style="list-style-type: none"> • Use Object Workflows after system expansion to redistribute objects over new AMPs. • Make sure database user credentials have access to maps and objects needed. • Use to move objects from an obsolete sparse map before deleting the map. • Set up an exclusion list. • Set up DBQL settings. • Add a notification set. • Evaluate the actionlist. • Check Perm space.

MAPS Manager Concepts

Maps, analyze, move, and workers are concepts integral to the **MAPS Manager** portlet.

What are Maps?

Analytics Database is composed of parallel units known as “AMPs.” Each node in the configuration has a predefined number of AMPs, and typically the rows of a table are spread across all AMPs in the system.

Maps (Hash Maps) determine how to distribute data from individual table rows among the AMPs. Maps list one or more AMPs. Every table is associated with a map to distribute the table data (rows) among the AMPs in the map.

Other objects such as user defined functions (UDF) are also associated with maps. A UDF is placed on the AMP that stores the associated table's rows.

MAPS Manager uses two types of maps:

- *Contiguous maps* include all AMPs within a specified range. Analytics Database creates contiguous maps during a system initialization or reconfiguration.

- *Sparse maps* include a subset of AMPs from a contiguous map. Analytics Database provides two sparse maps: one sparse map for a subset of the AMPs and one sparse map for one AMP.

MAPS Manager selects the optimal map for each object and provides the recommendation in System Monitoring or lists the target map in Object Workflows.

The order of map selection is first the system default map and the associated system created sparse maps, then any user created map that the user has been granted permissions to use. For System Monitoring, the credentials entered in the settings determine the maps for use. For Object Workflows, the logon credentials determine the maps that are selected or available. The user created maps must have a unique number of AMPs. If there are two user created maps with the same number of AMPs, **MAPS Manager** selects the map that alphabetically comes first. For example, the map Ann03 is selected instead of the map Vin03.

You can change the target map after analysis using workflow options. For more information, see *Teradata® Database Design*.

For the maps defined for a system, see the **Maps** tab. You can also use this view to create new sparse maps and grant or revoke permissions for users and roles.

What is Analyze?

For System Monitoring, the analyze stage creates a list of recommendations.

The analyze stage for Object Workflows determines the order and grouping of objects and which maps to use. For groups and objects, it is always in order of largest to smallest. During the analyze stage, the **MAPS Manager** portlet creates an action list. You can make manual changes to the action list. An analysis occurs once per workflow, but you can run multiple analyses at the same time.

What is Move?

In the **MAPS Manager** portlet, you can schedule moves after analysis using either System Monitoring or Object Workflows.

A move is the process of redistributing object data over system AMPs based on the map assigned to the object. This determines how many AMPs are activated when the object is accessed. You can schedule moves to occur during low system usage to reduce impact.

Only one move can occur on a system at any given time.

The move stage is an iterative process that moves objects to a new map according to the action list generated during the analyze stage. You can schedule the move or manually start the move after the analyze stage is complete.

The speed of the move stage depends on how busy the system is, how many workers are assigned, and the size of the objects.

What are Workers?

Workers are database sessions that work in parallel to move objects to new maps. The workers follow the group order specified in the action list. For serial groups, one worker is needed. For parallel groups, more workers can be added. You can add or remove workers while a move is running.

User Permissions

The Roles Manager admin portlet controls the permissions for MAPS Manager. In the Roles Manager admin portlet, select the appropriate option.

Option	Description
View System	Determines which Teradata systems you have access to through the MAPS Manager portlet.
View Workflows	Determines if you can view the Object Workflows tab.
View System Monitoring	Determines if you can view the System Monitoring tab.
View Maps	Determines if you can view the Maps tab.
Manage Maps	Determines if you have access to maps management options, such as grant, revoke, create or drop maps.
Admin for All Workflows	Determines if you can manage workflows created by other users or roles.

For System Monitoring, you need credentials with access to all objects and maps on the system. The TDMaps user has the access required, but needs logon access enabled in the database. In System Monitoring, any user of the portlet can see any object and map where Monitoring Settings credentials have access.

Connecting to a System

1. In the **Connect to System** dialog box, select a system from the list.
If a Unity Environment was added in the Viewpoint **Monitored Systems** portlet, it can be selected from the **System** list. Teradata systems managed by the Unity Environment are not listed separately.
2. Type your Teradata system user ID and password.
3. [Optional] Do any of the following:
 - Enter an **Account String**
 - Select an **Authentication Mechanism** from the list
 - Select a **Character Set** from the list
4. Click **Connect**.

Disconnecting from a System

1. From the portlet frame, click ▼ next to the system name, and select **Disconnect**.

Changing Systems

1. From the portlet frame, click ▼ next to the system name, and select **Change System**.
2. Select a system from the list.
3. Enter your Teradata system user ID and other login information.
4. Click **Connect**.

MAPS Manager View

The **MAPS Manager** view shows instructions before first use.

System Monitoring tab

The instructions show the first steps you must complete for System Monitoring.

Object Workflows tab

The instructions show the first steps to complete for Object Workflows.

Maps tab

The view lists all maps defined on the system.

Global Settings

The **Settings** list allows you to enter settings for System Monitoring and Object Workflows:

- **Monitoring Settings** apply only to System Monitoring.
- **Manage Notification Sets** and **Edit Exclusion List** apply to both System Monitoring and Object Workflows.
- **Edit DBQL Settings** apply to both System Monitoring and Object Workflows.

Setting Up DBQL Options

MAPS Manager uses DBQL settings during the analysis stage to group objects based on join patterns in logged queries instead of grouping by database. MAPS Manager uses the settings to create objects lists with the **Objects associated with query bands** and **Objects accessed but not in the system-default map** options.

For effective use of DBQL data, make sure the correct level of DBQL logging is enabled.

1. Click **Settings** ▼ and select **Edit DBQL Settings**.
2. From the **DBQL log database** list, select a database.

3. Select a date range.
4. Click **Save**.

Creating an Exclusion List

Create an exclusion list for objects you do not want to move. An exclusion list allows you to filter objects from appearing in the object selection lists for Object Workflows or exclude the objects from System Monitoring.

1. Click **Settings** ▾ and select **Edit Exclusion List**.
2. Select objects.
3. Click **Save**.

Managing Notifications


Notifications let you monitor workflows for events or errors during analysis and moves. Before selecting notifications, set up alert actions in the **Alert Setup** portlet. View notifications in the **Alert Viewer** portlet.

You can add notification sets to specific Object Workflows or to System Monitoring.

Monitoring Notifications

Adding a Monitoring Notification

Monitoring notifications only apply to events on the **System Monitoring** tab.

1. Click **Settings** ▾ and select **Manage Notifications**.
2. Click **Monitoring Notifications** tab.
3. Click .
4. Select the alerts to include for system monitoring.
 - a. Select the level of **Severity**.
 - b. Select the **Action**.
The actions sets are defined in the **Alert Setup** portlet.
 - c. Enter a message.
5. Click **Save**.

Event Types for Monitoring Notifications

Notifications	Description
Analysis Error	The analysis did not complete due to an error.
Move failed to start	A move did not start because another move was running.

Notifications	Description
Move running long	A move is running past the allotted time.
Object cannot be moved	An object cannot be moved due to a permission or other issue.
Object delayed or locked	An object is delayed or locked for a time.
Invalid user credentials	The workflow could not run due to invalid credentials. Check the user permissions.

Workflow Notifications

Adding a New Notification Set

Add a notification set to use with workflow monitoring.

1. Click **Settings** ▾ and select **Manage Notifications**.
2. Click **Workflow Notifications** tab.
3. Click **+**.
4. Enter the **Notification set name**.
5. Select the alerts to include in the Notification set.
 - a. Select the level of **Severity**.
 - b. Select the **Action**.
The actions sets are defined in the **Alert Setup** portlet.
 - c. Enter a message.
6. Click **Save**.
See the **Workflows** details view to add an alert to a workflow.

Editing a Notification Set

1. Click **Settings** ▾ and select **Manage Notifications**.
2. Click **Workflow Notifications** tab.
3. Click the name of the set.
4. Edit the notifications.
5. Click **Save**.

Event Types for Notification Sets

Notifications	Description
Analysis Error	The analysis did not complete due to an error.

Notifications	Description
Move failed to start	A move did not start because another move was running.
Move running long	A move is running past the allotted time.
Object cannot be moved	An object cannot be moved due to a permission or other issue.
Object delayed or locked	An object is delayed or locked for a time.
Invalid user credentials	The workflow could not run due to invalid credentials. Check the user permissions.

Adding Notifications to a Workflow

1. From the **Object Workflows** tab, select a workflow.
2. Select a notification set.
3. Click **Close**.

Removing Notifications from a Workflow

1. From the **Object Workflows** tab, select a workflow.
2. Select **None** from the **Notification set** list.
3. Click **Close**.

System Monitoring

System Monitoring provides recommendations to move objects to new maps to improve system performance. The recommendations are divided into categories. Each category shows a reason for moving objects.

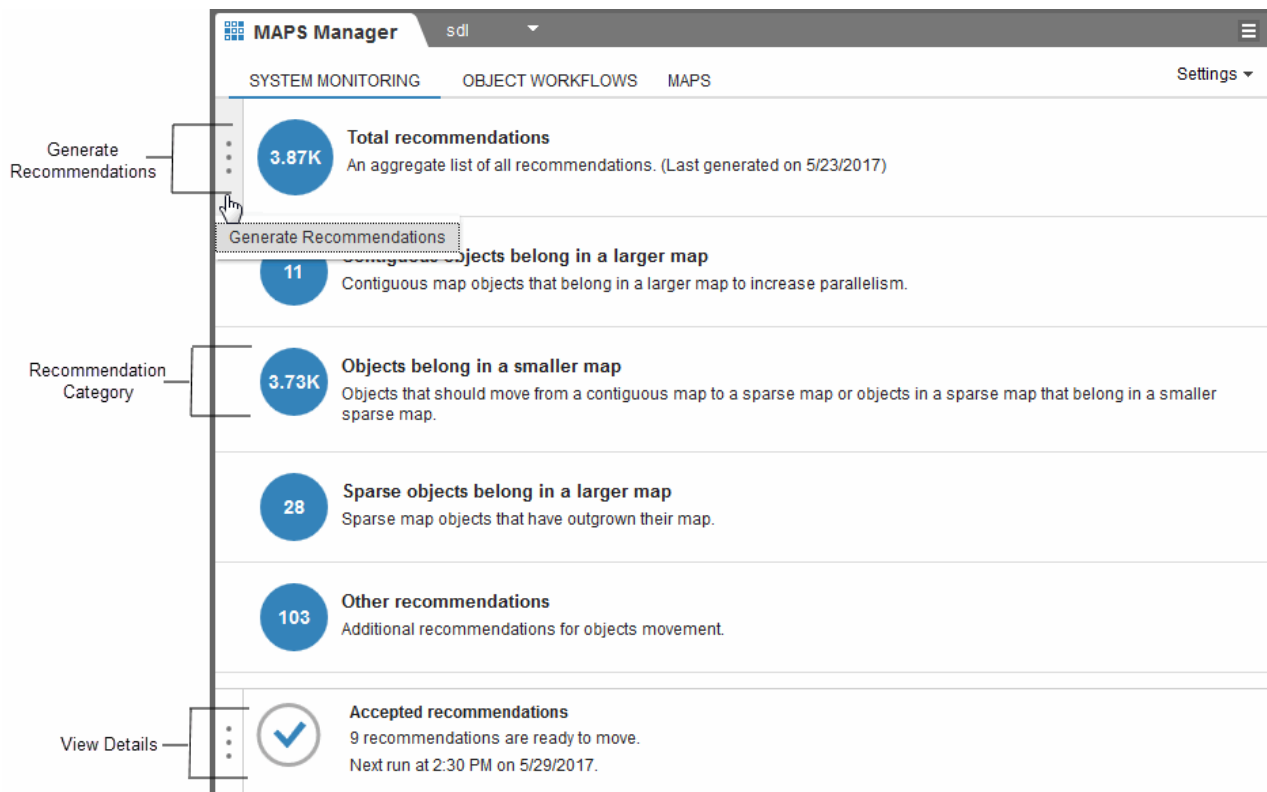
Before you use System Monitoring, you must do the following:

- Enter credentials that have access to all objects and maps in the system in the **Monitoring Settings**.
- [Optional] Select a schedule for Analyze.
- Select a schedule for Moves.
- [Optional] Populate DBQL settings to get better recommendations.

After you set up System Monitoring, the **MAPS Manager** portlet is ready to generate recommendations. Once the recommendations are available, you can view the details before accepting. After you accept the recommendations, the objects are moved according to the schedule.

System Monitoring View

The System Monitoring view shows the recommendations after analysis.



Monitoring Settings

The **Settings** option includes Monitoring Settings. You must enter credentials before getting recommendations.

Generate Recommendations

Generate Recommendations allows you to see recommendations for moving objects to new maps.

Recommendation Category

Each category shows a recommendation for movement.

View Details

The details show accepted recommendations.

Types of Recommendations

All recommendations are sorted into a category. The following categories are available:

Sparse objects that belong in a larger map

The recommendation is to move the objects from a sparse map to a contiguous map or larger sparse map.

Contiguous objects that belong in a larger map

The recommendation is to move objects to a larger map to increase parallelism.

Objects that belong in a smaller map

Either sparse or contiguous objects that need a smaller map.

Other Recommendations

Any move recommendation that does not fit the other categories appears here.

Total Recommendations

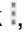
This category is a combination of all other categories.

Generating Recommendations

1. Click the **System Monitoring** tab.
2. Click **Settings** ▾ and select **Monitoring Settings**.
3. Enter credentials that have rights to access all maps and objects on the system.

Credential	Description
Username	Enter the username with permissions.
Password	Enter the associated password.
[Optional] Account String	Enter the account string.
Authentication Mechanism	Select an option.

4. [Optional] Click the **Analysis Schedule** tab.
 - a. Complete the schedule to indicate the repeat pattern, time, and day.
 - b. Click **Save**.

MAPS Manager has a default analysis schedule. Change the schedule to suit your needs.
5. After the analysis runs or to manually generate recommendations, click , then select **Generate Recommendations**.
You only need to set up System Monitoring once. In the future, select **Generate Recommendations** any time to receive recommendations.

Viewing Object Details for Recommendations

1. On the **System Monitoring** tab, select a category for recommendations.

2. Click an object to view details.
3. Click **Close**.


Recommendation Metrics

Metric	Description
Database	Database where object is located
Object	Object to be moved
Type	Type of object
Current Map	Current map location of object
Co-location	Colocation information
Recommended Map	Recommended new map
Recommendation	Description of object indicating reason for new map
Errors	Any error message
Recommendations bucket	Category for map recommendation

Accepting Recommendations

1. On the **System Monitoring** tab, select a category for recommendations.
2. Do one of the following:
 - Select the first box for all recommendations
 - Select boxes for single recommendations
3. Click **Accept Selected**.

Viewing Accepted Recommendations

1. On the **System Monitoring** tab, next to **Accepted Recommendations**, click .
2. Select **View Details**.

Accepted Recommendation Metrics

Metric	Description
Database	Database where object is located
Object	Object to be moved
Type	Type of object

Metric	Description
Current Map	Current map location of object
Recommended Map	Recommended new map
Recommendation	Description of object indicating reason for new map
Errors	Any error message
Action details	Command showing changes
Group	Group associated with object
Object Order	Object move order

Setting up a Move Schedule for System Monitoring

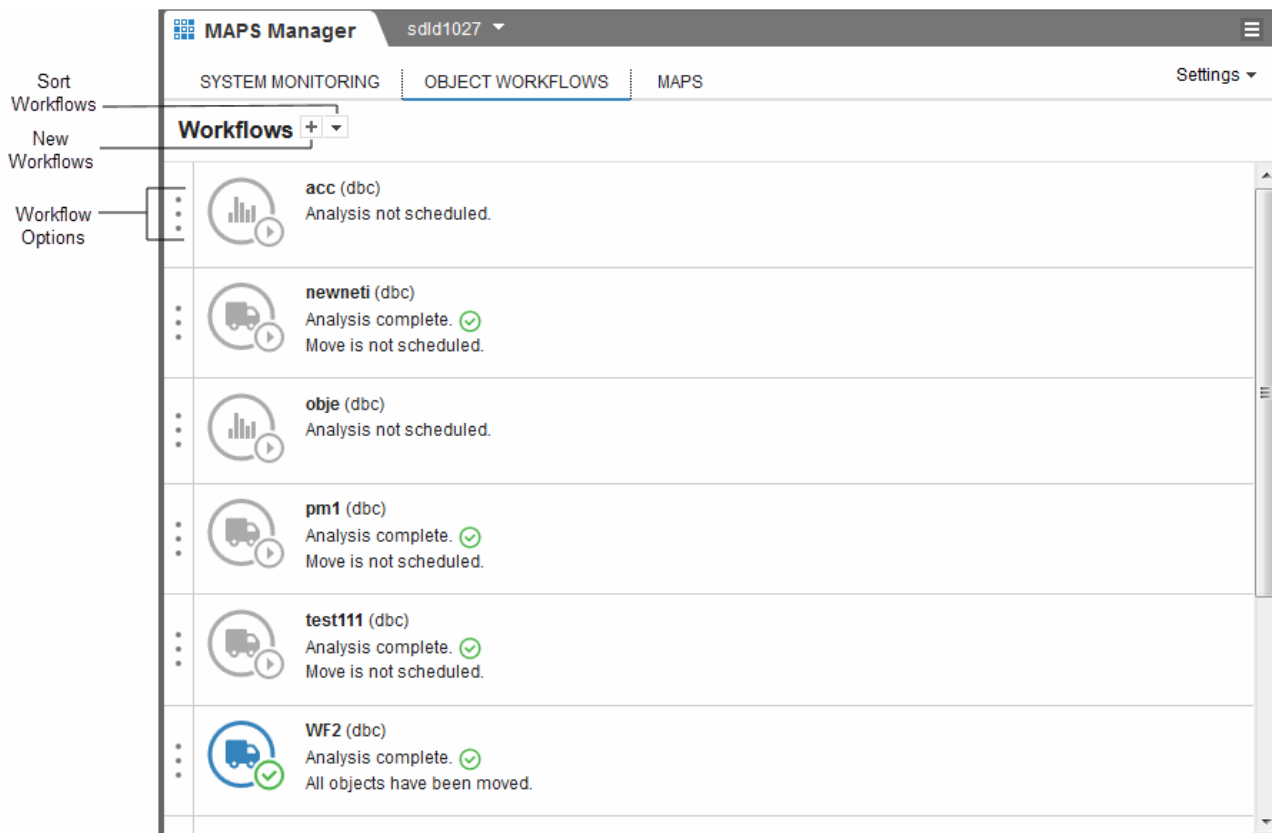
1. Click the **System Monitoring** tab.
2. Click **Settings** ▾ and select **Monitoring Settings**.
3. Click the **Move Schedule** tab.
4. Complete the schedule to indicate the repeat pattern, time, and day.
5. Click **Save**.

Viewing Previously Moved Objects in System Monitoring

1. On the **System Monitoring** tab, click **Accepted Recommendations**.
2. Next to **Accepted Recommendations**, click ▾ and select **View Previously Moved Objects**.

MAPS Manager Object Workflow Overview

The **Object Workflows** tab allows you to create ad hoc workflows.



Settings

Settings allow you to manage notifications, edit DBQL settings or add exclusion lists

New Workflows

Allows you to set up an analysis to determine objects to move.

Sort Workflows

Allows you to view workflows by owner and importance.

Workflow Options

Allows you to view details or perform management functions related to the current stage of the workflow.

Query Bands and MAPS Manager

MAPS Manager has the query band **WDClassification = MoveTableToMap**.

Query bands can help you troubleshoot issues, monitor processes, or specify workload management rules to manage system impact. You can view query bands in the **Query Monitor** portlet.

Workloads Associated with MAPS Manager

In the **Workload Designer** portlet, there is a default **MAPS Manager** workload named **WD-Maps-Mover** to assist with MAPS Manager workload management.

Creating a Pre-Expansion Workflow


Pre-expansion workflows help you plan for system expansions by allowing you view the redistribution of objects before the resources are available. When using SQL Engine 16.20 Update 2 and later, the option **Create New Pre-Expansion Workflow** appears in the **Object Workflows** tab. Only one pre-expansion workflow can be created per Teradata system.

1. Click the **Object Workflows** tab.
2. Next to **Workflows**, click ▾.
3. Select **Create New Pre-Expansion Workflow**.
4. Enter a unique name for the workflow, and then click **Next**.
The label "Pre-Expansion Workflow" is appended to the workflow name when the workflow is displayed in the **Object Workflows** tab and in the **Workflow Details** view.
5. From the **Object List type** list, select an option and then click **Next**.
6. Select the Workflow Options, and then click **Create Workflow**.
7. Analyze the pre-expansion workflow to prepare for conversion to a regular workflow.

Converting a Pre-Expansion Workflow to a Regular Workflow

1. After analysis of the pre-expansion workflow has been successfully completed, do the following:
 - a. Click **CONVERT WORKFLOW**.
 - b. Select a contiguous map, then click **CONVERT WORKFLOW** to confirm conversion.
 Viewpoint converts the workflow and removes the **Pre-Expansion Workflow** label from the workflow.

Creating a Workflow

1. Click the **Object Workflows** tab.
2. Next to **Workflows**, click .
3. Enter a name for the workflow, and then click **Next**.
The name must be unique.
4. From the **Object List type** list, select an option and then click **Next**.
5. Select the Workflow Options, and then click **Create Workflow**.

Object List Type Selection


You can select objects to include in a workflow using different methods.

Selecting the Entire System

Selecting the entire system allows you to add all available objects on the system at the time the analysis is run. Some objects are exempt from moving, such as system objects.

1. From the **Object list type** list, select **Entire system**.
2. Click **Next**.

Selecting a Database

1. Open **MAPS Manager** and select **Object Workflows**.
2. Next to **Workflows**, select .
3. In **New Workflow**, name the workflow and select **Next**.
4. In **Object list type**, select **Manually select databases**, select the target databases, and click **Next**.

Selecting an Object List from a File

To select an object list from a file, you must set up the file before creating the workflow. The file must be a .csv (database, object) on each line.

1. From the **Object list type** list, select **Import an object list from a file**.
2. Click **Import List**.
3. Click **Browse** to select the file.
4. Click **Import**.
5. Click **Next**.


Selecting Objects Manually

Manually selecting objects allows you to choose a set of objects from all available objects on the system.

1. From the **Object list type** list, select **Manually select objects**.
2. Do one of the following:
 - Select each object to include.
You can filter for databases or objects.
 - Select the top check box to select all available objects.
3. Click **Next**.

Selecting Objects Associated with Query Bands

This feature requires DBQL configuration and the appropriate level of logging enabled.

1. From the **Object list type** list, select **Objects associated with query bands**.
2. Enter the name and value.
3. [Optional] Click  to add query bands.
4. Click **Next**.

Selecting Objects Accessed But Not in the System-Default-Map

This option only appears when there are two or more contiguous maps on the system. Selecting **Objects Accessed But Not in the System-Default-Map** allows you to find and select all remaining objects that are not in the new system default contiguous map.

1. From the **Object list type** list, select **Objects Accessed But Not in the System-Default-Map**.
2. Click **Find Objects**.
3. Click **Next**.

Workflow Options

Set scheduling and grouping options on the last view of the workflow wizard.

Setting Workflow Options

1. Next to **Analysis Schedule**, select an option.
You can schedule the analysis to run later or wait to manually run the analysis after the workflow is saved. You can view currently scheduled analyses.
2. Next to **Object Grouping in the List of Actions**, select a method to group objects.
The join pattern option requires that DBQL is setup and the appropriate level of logging is enabled.

Note:

If you do not see the option, you need to set up DBQL.

3. [Pre-Expansion Workflow only] Next to **Planned Number of Nodes**, enter the number of nodes for future expansion planning.

Note:

If you do not see this option, you are creating a regular workflow, not a pre-expansion workflow.

4. [Optional] Next to **Automatically Move Objects**, select **Move objects immediately after the analysis without reviewing move actions**.
If you select this option, the move begins immediately after the analysis without allowing time for evaluation of the list of actions. If you skip this option, you can set up a schedule and modify the list of actions before moving any objects.
5. [Optional] From the **Notification set**, select a notification set.

You must create a notification set before adding.

6. Click **Create Workflow**.

Working with Object Workflows






After you create a new object workflow, the workflow appears in the **Object Workflows** list. From this list you can edit and monitor the workflow.

Object Workflows View

The **Object Workflows** view lists the status of current workflows. You can view the status icons for each workflow and access the **Workflow Details** view.

Workflow Icons











Analysis Status Icons for MAPS Manager

Icon	Status
	Analysis is ready, but not scheduled.
	Analysis was terminated and is not scheduled.
	Analysis is scheduled.
	Analysis is running.
	Analysis is complete.

Note:

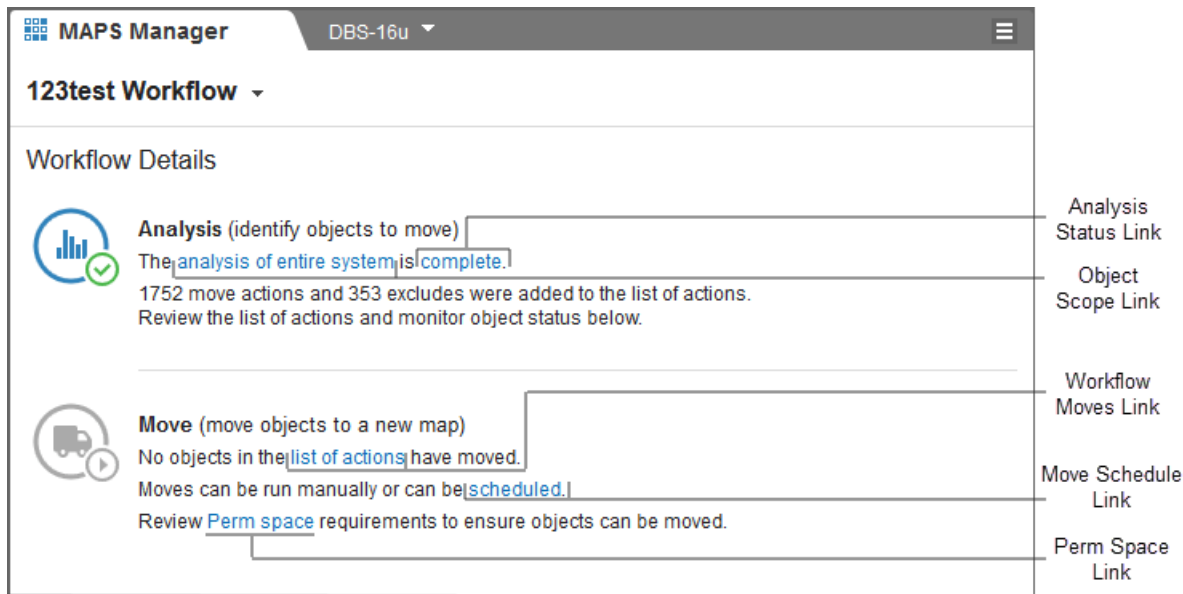
Hover over the icons for more information.

Move Status Icons for MAPS Manager

Icon	Status
	Moves are not available. Analysis is needed.
	No moves scheduled.
	Moves did not start due to conflict. Move is not scheduled.
	Moves are not complete and not scheduled.
	Moves are scheduled.
	Moves are scheduled to continue later.
	Moves did not start due to a conflict with another move. Another attempt is scheduled.
	Moves started, but did not complete due to a conflict with another move. Another attempt is scheduled.
	Moves are running.
	All objects have been moved.

Workflow Details View

The **Workflow Details** view includes links to let you manage the workflow.



Analysis Status Link

Allows you to view the analysis results.

Object Scope Link

Allows you to modify which objects are included.

Workflow Moves Link

Shows all objects that are identified to move and which objects are excluded. The list shows the status objects listed. It allows you to change object order, target maps, exclude, or un-exclude objects.

Move Schedule Link

Allows you to add or replace a move schedule.

Perm Space Link

Allows you to increase or decrease Perm Space.

Editing Objects to Analyze

You can change the scope of objects to analyze for a workflow. This option is not available after the analysis begins. You can reset the analysis if needed.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Analysis** section, click the analysis link.
3. Next to **Scope of Analysis**, click the objects selected.
4. In the **Object list type** list, select an option.

5. Select objects to include.
6. Click **OK**.

Object List Metrics

Metric	Description
Database	Name of Analytics Database
Object	Name of Teradata object
Type	Type of database object
Current Map	Map associated with object
Peak Perm	Largest amount of Perm space used by object
Current Perm	Perm space currently used by object
Map Type	Sparse or contiguous
Map Amps	Number of AMPs associated with the map
Fallback	Designation of fallback status

Adding or Editing an Analysis Schedule

You can only modify the schedule of an analysis before it begins.

1. From the **Object Workflows** tab, select a workflow.
2. Click **analysis of <number> of objects**.
3. Click one of the following:
 - **Not scheduled**
 - **scheduled**
4. Next to **Schedule**, select a date.
5. Click **OK**.

Viewing Analysis Status

1. From the **Object Workflows** tab, select a workflow.
2. In the **Analysis** section, click **complete** to view the results.
The results show the runtime information, the scope, the number of objects excluded, and the destination maps.

Workflow Actions

After analysis, each workflow has a list of actions. The list of actions defines which objects are moved to a different map and which objects are excluded from moving (already in an appropriate map). You can edit these actions. You can export the list of actions.

Viewing Workflow Actions

The **Workflow Actions** view is available after the analysis.


1. From the **Object Workflows** tab, select a workflow.
2. Click **list of actions**.
If the moves have completed, instead of **list of actions**, **<number of actions> move actions** appears.

Changing Object Order

You can move an object to a new position in a list of actions for a workflow.

Objects are moved based on their group number and order within the group. Change the object order to move the object sooner or later than specified after the analysis.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click **list of actions**.
3. Do one of the following to select objects to reorder:

Option	Description
Single object	<ol style="list-style-type: none"> a. Click  next to an object and select Change Object Order. b. From the Target Group list, select a new target group. c. From the Group position list, select a new position. d. Click Save.
Multiple objects	<ol style="list-style-type: none"> a. From the More options list, select Change Object Order. b. Select the objects. c. Click Change Order.

Changing Target Map

You can change a target map to override the map listed after analysis. To prevent issues with system performance, do not place objects that are too large into sparse maps. This option is available until all moves are complete.

1. From the **Object Workflows** tab, select a workflow.

2. Click **list of actions**.
3. Select objects to move to a different target map.

Option	Description
Single object	a. Click <input type="checkbox"/> next to an object and select Change Target Map .
Multiple objects	a. From the More options list, select Change Target Map .

4. In the **Target Map** list, select **contiguous** or **sparse**, then select the destination map.
5. Click **Change Map**.

Excluding Objects

After the analysis generates a list of actions, the list contains objects to move (have alter statement) and objects that do not need to move (excludes). To override the analysis and not move an object, you can change the object to exclude. When you exclude an object, the group and object order, and map recommendations are not saved.

1. From the **Object Workflows** tab, select a workflow.
2. Click **list of actions**.
3. Do one of the following to exclude objects:

Option	Description
Single object	a. Click <input type="checkbox"/> next to an object and select Exclude Object .
Multiple objects	a. From the More options list, select Exclude Objects .

4. Click **Exclude Objects**.

Un-Excluding Objects

During analysis, if an object is excluded or if you previously excluded an object from the scope of the workflow, you can un-exclude the object.

1. From the **Object Workflows** tab, select a workflow.
2. Click **list of actions**.
3. Under **Workflow Actions**, select **Display objects that are excluded from moving** from the list.
4. Do one of the following to un-exclude objects:

Option	Description
Single object	a. Click <input type="checkbox"/> next to an object and select Un-Exclude Object .
Multiple objects	a. From the More options list, select Un-exclude Objects .

5. In the **Target Map** list, select **contiguous** or **sparse**, then select the destination map.
6. In the the **Target Group** list, select the group number.
7. In the **Group position**, select the order of the object in the specified group.
8. Click **Un-exclude Objects**.

Workflow Actions Metrics

Metric	Description
Group Order	Order in which the group is moved
Group Type	Serial or Parallel
Database	Name of the Analytics Database
Object	Name of the Analytics Database object
Type	Type of Database object
Object Order	Order within the groups in which objects are moved into new maps
Size	Perm space used by object
Target Map	Destination map for object
Target Map AMPs	Number of AMPs associated with the target map
Source Map	Original map for the object
Edited	Flag that shows changes to the action by a user
Status	Move status: Pending, Moved, Moving, Warning, and Error
Action Details	Statement used to move or skip the object
Warning	Details of constraints that affect the movement of the object
Error Code	Error returned when a worker attempted to move the object
Description	Additional information about the completed or failed action
Move Duration	Time spent moving the object

Note:

Exporting the **Workflow Actions** metrics presents the **Move Duration** time in milliseconds.

Manage Groups

During analysis, **MAPS Manager** groups objects to move and defines the order of the moves. You must make sure there is enough Perm space allocated to move each object.

You can add parallel groups. Parallel groups move multiple objects at the same time depending on the number of workers assigned. If multiple objects are moved in parallel, which are in the same database, you need to make sure there is Perm space available to move all of the objects at once (based on workers).

Parallel groups are skipped if there are time constraints. Before running a group, **MAPS Manager** checks to see if there is enough time to finish moving the objects within the group. If there is not enough time, the group is skipped. **MAPS Manager** continues checking until a group that fits the available time slot is found. Skipped groups are retried in the next move.

If there are serial and parallel groups, the first worker is allocated to the serial group. Serial order moves one object at a time in order using one worker. The other workers are allocated to parallel groups. If the serial list finishes during a run and the parallel list still has actions remaining, during the next run, all workers are allocated to the parallel group.

When you delete a group, you select a new group to contain the objects. You can also change the order the objects appear in the new group.

Managing Groups

The **Workflow Actions** view allows you to manage groups for objects. If you delete a group, the objects are not deleted and you must select a new group.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click **list of actions**.
3. From the **More options** list, select **Manage Groups**.
4. From the **Action** list, select one of the following:

Option	Description
Add parallel group	<ol style="list-style-type: none"> a. Select Add parallel group. b. From the Group position list, select a position.
Add serial group	<ol style="list-style-type: none"> a. Select Add serial group.
Change group order	<ol style="list-style-type: none"> a. Select Change group order. b. From the Group to move list, select a group number. c. From the New position list, select a position.
Delete group	<ol style="list-style-type: none"> a. Select Delete group. b. From the Group to delete list, select a group number. c. From the Move objects to list, select a new group. d. From the Position in target group list, select a position.

5. Click **Save**.

Space Management

Space management lets you add or remove space from databases to make sure enough space is available to move objects.

When you move an object to a new map, the original object remains until a copy of the object is successfully created on the target AMPs. Temporarily, the database needs enough Perm space for two copies of the object.

Increasing the workers increases the number of objects that are moved at the same time. Make sure the available Perm space is enough to move the n largest objects in the database with n equal to the number of workers.

Reviewing Perm Space

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click **Perm space**.
The Review Perm Space Requirements view appears. You can change the Perm space available or the number of workers.
3. Click **Close**.

Review Space Requirements Metrics

Metric	Description
Database	Name of the Analytics Database
Objects	Number of objects in the Analytics Database that must be moved
Free Perm	Unused Perm space in the database
Recommended Free Perm	Estimated total Perm space needed to complete a move
Free Perm Needed	Estimated Perm space needed to complete a move

Increasing Perm Space

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click **Perm Space**.
3. Select one of the following:

Option	Description
Increase Perm space to all databases	<ol style="list-style-type: none"> a. Click Increase All to Recommended. b. From the Reserve Perm pool list, select the appropriate database. c. [Optional] From the Workers list, change the number of workers. The amount of space allocated is based on the estimated number of workers.

Option	Description
Increase Perm space to a specific database	<ol style="list-style-type: none"> Click <input type="checkbox"/> next to the appropriate database, and then click Allocate Space. Click <input type="checkbox"/> to select a reserve Perm pool. Click and drag the slider to change the amount of space or enter a value.

- Click **Save**.

Decreasing Perm Space

Use these steps to decrease the Perm space in databases with no pending move actions.

- From the **Object Workflows** tab, select a workflow.
- In the **Move** section, click **Perm Space**.
- From the **Review Perm Space Requirements** list, select **Databases with no pending move actions**.
- Click **Decrease All**.
- Click **Close**.

Moving Objects Using Move Now


- Click the **Object Workflows** tab.
- Click ☐ next to the workflow you want to move, then select **Move Now**.
The workflow must have completed analysis and have at least one object to move.
- [Optional] Specify the number of workers.
- From the **End** list, select one of the following:

Option	Description
When complete	Workflow continues until all scheduled moves are complete.
At <time>	<p>Workflows continues until the time scheduled.</p> <p>If you select Allow in progress groups or objects to finish if time expires, the workflow continues until groups or objects complete. If you do not select this option, you must reschedule moves that do not complete later.</p>

- Click **Move**.

Aborting a Move

Aborting moves gracefully allows in progress groups or objects to complete their move before stopping. Aborting moves immediately stops all sessions and you must retry any objects that were not completed later.

1. Click the **Object Workflows** tab, then select  next to a workflow performing a move.
2. Select **Abort move** and then choose an option.

Option	Description
Graceful stop	Allows groups or objects to finish. All objects that have been successfully moved before an abort remain in the target map. No roll-back exists for the job.
Immediate stop	Aborts job immediately. All sessions and objects end. Retry moves later.

3. Click **Abort**.

Move Schedule for Object Workflows

Schedule a move when it impacts the system the least. You can set up moves to occur at regular intervals until all objects in the list of actions have moved.

Two moves cannot occur at the same time on a given system. If a job is scheduled to start while another job is running, the second job does not start and is not queued. Make sure to schedule moves when they do not conflict with other moves.

Adding a Move Schedule to Object Workflows

Schedule a move to occur after analysis during a workflow.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click **scheduled**.
3. Next to **Move Schedules**, click **+**.
4. Review the **Currently Scheduled Move Jobs** to select a time when a move is not in progress. You cannot schedule a move at the same time as another move.
5. From the **Recurrence** list, select an interval.
6. From the **Start** list, select a date and time.
7. From the **End** list, select one of the following:

Option	Description
When complete	Workflow continues until all scheduled moves are complete.
At <time>	Workflows continues until the time scheduled. If you select Allow in progress groups or objects to finish if time expires , workflow continues until groups or objects complete. If this is not selected, moves must be rescheduled later.

8. [Optional] From the **Number of workers** list, select the number of workers to allocate to the move.
9. Click **Save**.

Adding or Removing Workers During a Move

During a move, you can add workers. This may increase the need for Perm space.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click the **<n> worker is allocated**.

Move (move objects to a new map)

2 of 4 objects in the [list of actions](#) have moved (2.5K of 7.5K).

[1 worker is allocated](#) to this move job.

3. From the **Number of workers** list, select a number.
4. Click **Change**.

Viewing the Previous Move Jobs List

You can view all move attempts including completed moves and scheduled moves that did not start due to a job already running.

1. From the **Object Workflows** tab, select a workflow.
2. In the **Move** section, click the **completed** link.

Previous Move Jobs Metrics

Metric	Description
Start	Time and date move job began
Preparation Time	Duration of preparation time
Move Time	Total time move ran
Objects Moved	Number of objects transferred during the move job
Size Moved	Total size of objects moved
Objects Remaining	Number of objects remaining in the list of actions
End Status	Status at the end of the move: <ul style="list-style-type: none"> • Abort • Completed • Completed with issues • Error
Average B/S	Average bytes per second transferred during the move job
Maximum Workers	Highest number of workers allocated to the move

Viewing Workflow History

1. Click the **Object Workflows** tab.
2. Next to **Workflows**, click ▾, and then select **View Workflow History**.
3. Select a workflow to view details.
4. Click **Close**.

Workflow History Metrics

Workflow History Metric	Description
Workflow	Name of workflow
Final Status	Status of last performed step in workflow
End Date	Last date the workflow was running
Objects Moved	Number of objects moved

Maps View

The **Maps** view allows you to see what maps are available on the Teradata system. You can add sparse maps, manage map access, and view map details in this view.

Maps View Metrics





Field	Description
Map	Name of map and system default indicator
Map Type	Contiguous or sparse
AMPS	Number of Amps allocated for the map
Objects	Number of objects in the map
Objects Rollup	Number of objects in contiguous and children sparse maps
Used for Analysis	Yes or No response to indicate if used during Object Workflow analysis

Viewing Map Details

1. From the **MAPS Manager** portlet, click ▾ next to the system default map and select **View Details**.

Managing or Revoking Grants

Grant or revoke access to contiguous or sparse maps.

1. From the **MAPS Manager** portlet, click ▾ next to the map.
2. [Optional] To grant access to the map, from the **Available users** list, select or enter the name of the user.
3. Click .
4. [Optional] To remove access from the map, from the **Users with access** list, select or enter the name of the user.
5. Click .
6. [Optional] To grant access to a role, from the Available Roles list, select or enter a role.
7. Click .
8. [Optional] To revoke grants from a role, from the Available Roles list, select or enter a role.
9. Click .
10. Click **Update**.

Creating a Sparse Map

1. From the **MAPS Manager** portlet, click ▾ next to the system default map and select **Create Sparse Map**.
2. Enter a map name.
3. Enter the number of AMPS.
Viewpoint requires maps have a unique number of AMPS. If a map exists with the same number, you cannot proceed.
4. Click **Create**.

Deleting a Sparse Map

The option to delete a sparse map is only available when the sparse map is empty. Move objects to a different map before attempting to delete a sparse map.

1. From the **MAPS Manager** portlet, click ▾ next to the sparse map and select **Drop Sparse Map**.
2. Click **Delete**.

Changes to the System Default Map

If you change the system default map, a warning notification appears next to all Object Workflows and in System Monitoring. Click the link in the message to re-run the analysis. Any objects in the old system default map need to be included in a new Object Workflow or System Monitoring.

Warning Notification		Rerun Analysis link
test**** (dbc)	Analysis not scheduled.	
test123 (dbc)	Analysis complete. Move is not scheduled. The default contiguous map for this system has changed to TD_Map2, causing workflow move actions to become invalid. Rerun the analysis to use the new default map.	
test44941 (dbc)	Analysis complete. Move is not scheduled. The default contiguous map for this system has changed to TD_Map2, causing workflow move actions to become invalid. Rerun the analysis to use the new default map.	
testAnalysisAnotherUser (dbc)	Analysis complete. Move is not scheduled. 2 move errors occurred during the last move.	
testanalysispercent (dbc)	Analysis complete. All objects have been moved.	

Metric Heatmap

Metric Heatmap Overview

The **Metric Heatmap** portlet allows you to monitor single system metrics and resource usage trends within a specified time frame to help identify periods when a system is either over-utilized or under-utilized. You can use this information to determine when to schedule resource-intensive jobs with minimal impact to other users. Data displays in either a heatmap or graph view.

Metric Heatmap Menus and Toolbar

The menus and toolbar allow you to specify the display attributes of the views.

Use the selection menus to define a metric to monitor by selecting a:

- **System**
- **Category**
- **Metric**
- **Source** or **Variant**

The first time you add the portlet to the page, the first available metric that you have permission to view is displayed. However, if you only have permission to view **Table** metrics, no data appears in the portlet until you select a metric using the selection menus.

The **1 mo** and **3 mo** buttons allow you to adjust the time frame displayed. As you adjust the time frame, the toolbar refreshes automatically to display the new time frame in the toolbar.

The **Export** button allows you to create a .csv file containing selected data.

Heatmap View Features

The **Display** button allows you to show or hide cell values and adjust high and low values that determine cell shading. The **Display cell values** check box is only available when the metric displayed is a percentage. If the selected time frame results in narrow cells, cell values are not displayed even though the **Display cell values** check box is selected.

The threshold button displays the current threshold value and allows you to set the threshold slider and adjust the threshold. The threshold button and slider only appear when the metric displayed is a percentage. If the threshold value is within the gray-scale range, defined when you select the **Display** button, the cells are displayed in a shade of the exception color equivalent to the corresponding gray-scale shade for that cell value.

Heatmap View

The heatmap view displays a heatmap of variations and trends in system resource usage. The heatmap plots days on the horizontal axis and hours along the vertical axis. Customize the view using the menus, toolbar, and date-range slider to select metrics and date ranges.

The following list describes the features in this view:

Selection Menus

Shows the selections that define the metric currently being displayed. You can change the metric by changing the selections.

Toolbar

Contains the time frame list. Select a time frame from the list to change the graph. The time frame start and end dates are displayed on the toolbar next to the time frame buttons. You can use the **Display** and threshold buttons to define how the cells are shaded. **Export** creates a .csv file containing selected data.

Date-Range Slider

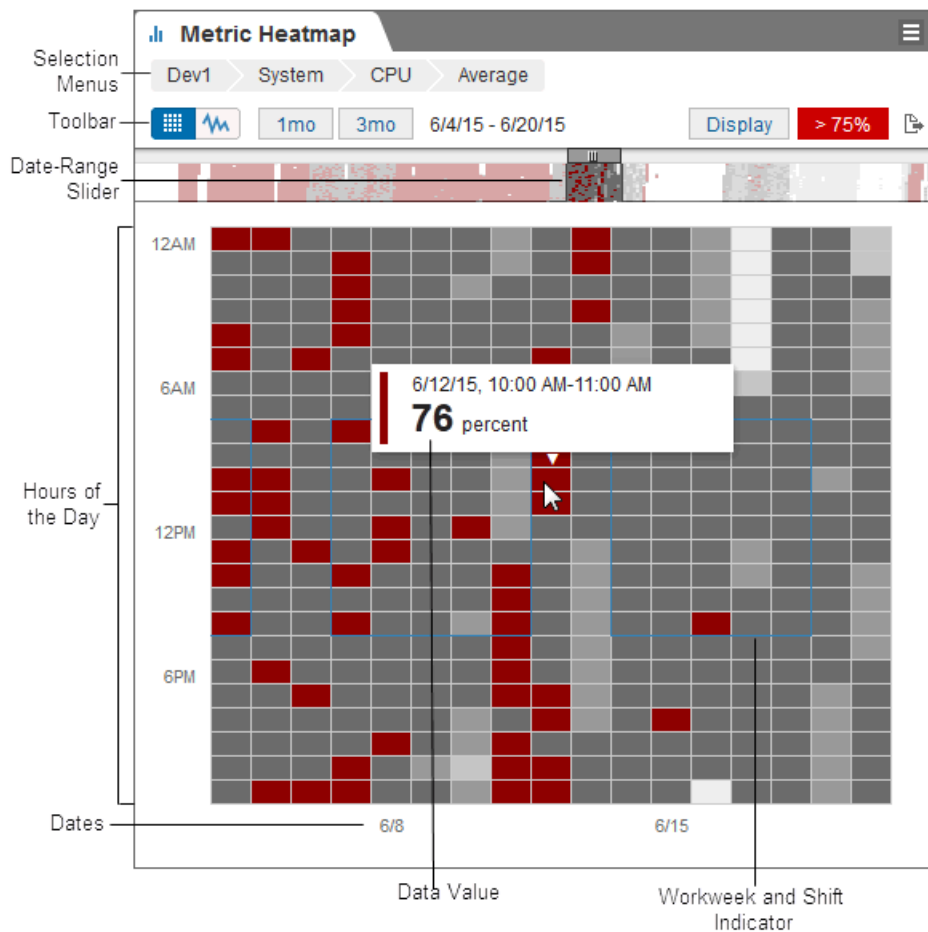
Controls the amount of historical data displayed in the view and adjusts to display from 7 days to approximately 4 months of historical data. Dates outside the date-range slider timeline are shaded lighter. The darker-shaded area between the slider handles represents the current date range being displayed.

Cell

Contains data about an hour on a specific date in the heatmap. Hover over any cell to display an information balloon containing the date, hour, and metric value for the 1-hour block.

Work Week and Shift Indicator

Highlights times when most users are normally accessing system resources with a box around the work week and work shift hours. Specify work day and shift hours using the **Settings** view.



Graph View

The graph view displays a graph of variations and trends in system resource usage. The graph plots the time frame on the horizontal axis and the metric values on the vertical axis. Customize the view using the menus, toolbar, and date-range slider to select metrics and date ranges.

The following list describes the features in this view:

Selection Menus

Shows the selections that define the metric currently being displayed. You can change the metric by changing the selections.

Toolbar

Contains the time frame list. Select a time frame from the list to change the graph. The time frame start and end dates are displayed on the toolbar next to the time frame buttons. **Export** creates a .csv file containing selected data.

Date-Range Slider

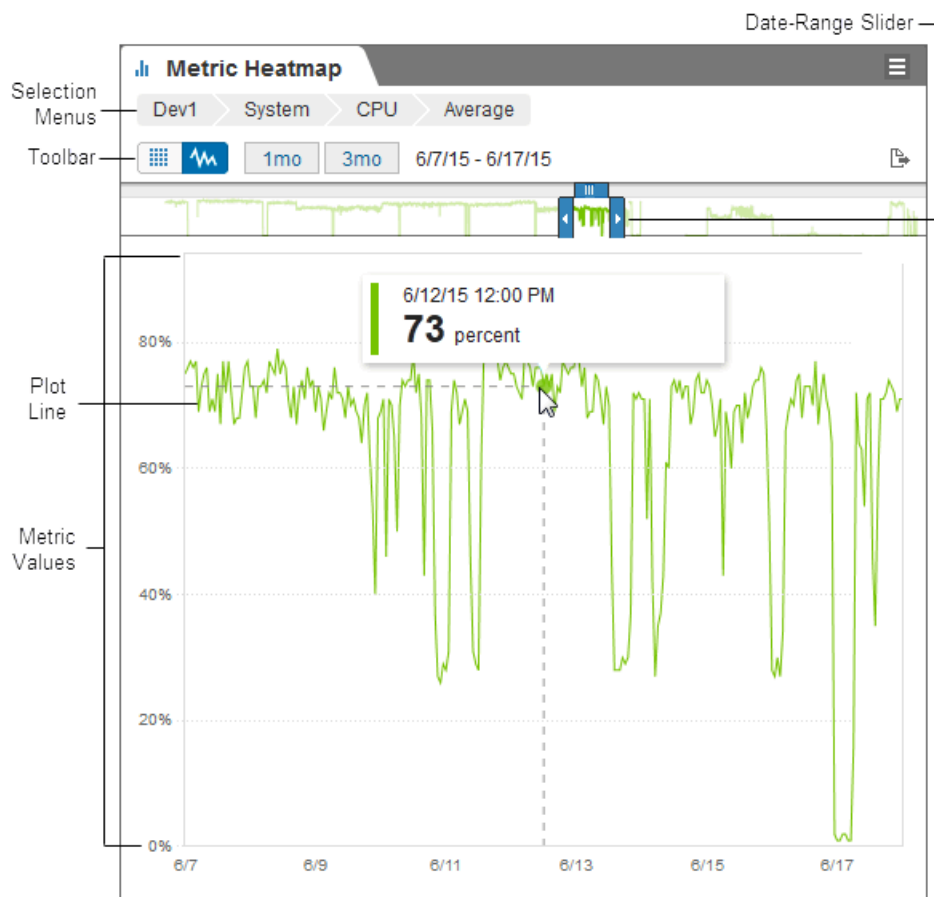
Controls the amount of historical data displayed in the view and adjusts to display from 7 days to approximately 4 months of historical data. Dates outside the date-range slider timeline are shaded lighter. The darker-shaded area between the slider handles represents the current date range being displayed.

Line Graph

Shows the numerical value of metric on the vertical axis. Hover over a number and drag it up or down to expand or compress the vertical scale.

Plot Line

Displays metric data values in the graph. Hover over a plot line to see an information balloon containing the date and time the value was collected and the metric value.





Selecting Metrics and Systems to Monitor

You can use selection menus to choose a metric and system to monitor in the portlet, starting with the highest level menu on the far left. Menu choices are available based on the previous menu, so not all menus apply to all metrics.

You are not required to make a selection from every menu each time you modify a metric selection.

1. In the selection menu, click the currently selected system name to display a list of available systems.
2. Select a system.
3. Select >.
4. Select a category from the menu.
5. Select >.
6. Select a metric from the menu.
7. Select >.
8. [Optional] Select a source or variant from the menu.
The menu for source or variant (sum, average, minimum, maximum, standard deviation of that metric) is only available for certain metric categories.
9. Select >.

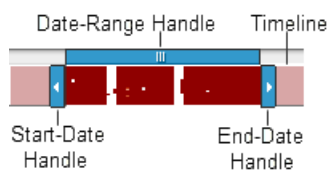
Selecting the View

1. Select a view:
 - Click  to see the graph view.
 - Click  to see the heatmap view.

Adjusting the Date Range

Use the toolbar to set the range of dates to display for the selected metric and system.

1. Hover over the date-range slider to activate the start-date and end-date handles.



2. Drag the start-date handle until the intended start date is displayed in the toolbar.
3. Drag the end-date handle until the intended end date is displayed in the toolbar.
4. [Optional] Drag the date-range slider left or right to move the start-date and end-date handles in unison.

5. [Optional] Click **1 mo.**
One month of historical data is displayed ending at the current end date.
6. [Optional] Click **3 mo.**
Slightly more than three months of historical data is displayed ending at the current end date.

Setting Cell Values and Shading in the Heatmap View

You can display cell values and set metric thresholds for applying cell shading for each system monitored. The shading provides a visual overview by highlighting interesting values.

You must repeat this procedure for each available metric to make sure appropriate shading is applied when the metric is selected.

1. From the heatmap view, click **Display**.
2. Do one of the following:
 - Select the **Display cell values** check box to display cell values.
 - Clear the **Display cell values** check box to display cells without values.

The **Display cell values** check box is only available when the metric displayed is a percentage. If the selected date range results in narrow cells, cell values are not displayed even though the **Display cell values** check box is selected.
3. Set the lower and upper threshold values in the **Apply shading to values between** boxes.
4. Click **OK**.

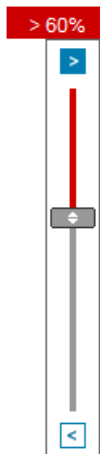
Values under the defined range appear in a very light shade of gray, while values after the range have a very dark shade. Values within the defined range appear in three shades of gray.




Setting the Threshold Value in the Heatmap View

The threshold button displays the current threshold value. Use the threshold button to access the threshold slider and adjust exception highlighting to show when the system is over-used or under-used.

The threshold button and slider only appear when the metric displayed is a percentage.


1. In the toolbar, click the threshold button.
The threshold slider appears.



2. Do one of the following:
 - Click  to set an upper threshold, highlighting all values greater than the selected value in shades of red.
 - Click  to set a lower threshold, highlighting all values less than the selected value in shades of gray.
3. Drag  up or down to adjust the threshold value.
The heatmap view highlights change to reflect values beyond or under the threshold setting.
4. Click anywhere off the threshold slider to close the slider and return to the portlet.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains up to the last 8 months of data, regardless of the selected date range. One variant (Avg, Max, or Min) is exported for the selected metric.

1. In the selection menu, select a system, category, metric, and variant.
2. Click .
3. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
4. [Optional] Reformat exported data to match the format in the view.

Trend-Reporting Metrics for Teradata System

Database Metrics

Metric	Description	Type
Current Perm	Sum of the current amount of permanent space (disk space) used across all vprocs	Number

Metric	Description	Type
Current Perm Max	Maximum permanent space used on a single vproc	Number
Current Perm Min	Minimum permanent space used on a single vproc	Number
Current Spool	Sum of the current amount of spool space (disk space) used across all vprocs	Number
Current Spool Max	Maximum spool space used on a single vproc	Number
Current Spool Min	Minimum spool space used on a single vproc	Number
Current Temp	Sum of the current amount of temporary space (disk space) used across all vprocs	Number
Current Temp Max	Maximum temporary space used on a single vproc	Number
Current Temp Min	Minimum temporary space used on a single vproc	Number
Peak Perm	Sum of the largest amount of permanent space used across all vprocs since the last reset of the peak perm value	Number
Peak Perm Max	Maximum of the largest amount of permanent space used on a single vproc since the last reset	Number
Peak Perm Min	Minimum of the largest amount of permanent space used on a single vproc since the last reset	Number
Peak Spool	Sum of the largest amount of spool used across all vprocs since the last reset of the peak perm value	Number
Peak Spool Max	Maximum of the largest amount of spool used on a single vproc since the last reset	Number
Peak Spool Min	Minimum of the largest amount of spool used on a single vproc since the last reset	Number
Peak Temp	Sum of the largest amount of temporary space used across all vprocs since the last reset of the peak perm value	Number
Peak Temp Max	Maximum of the largest amount of temporary space used on a single vproc since the last reset	Number
Peak Temp Min	Minimum of the largest amount of temporary space used on a single vproc since the last reset	Number
Perm Limit	Maximum space available for permanent space storage of all index tables, data tables, subtables, stored procedures, triggers, and permanent journals	Number
Spool Limit	Maximum space available for storage of spool database objects	Number
Temp Limit	Maximum space available for storage of temporary database objects	Number

Performance Metrics

Metric	Description	Type
Blocking Duration	Total time the heartbeat query is blocked	Number
Concurrency	Average number of active tasks running on Analytics Database while the heartbeat query is running	Number
Response Time	Average response time of the selected canary query	Number
Retrieve Time	Average time to retrieve all results	Number
Rows Returned	Average number of rows returned by the selected canary query	Number
Total Time	Combined response time and retrieve time	Number

Query Metrics

Metric	Description	Type
Abort Total	Total aborted queries	Number
Active AMPs Avg	Average number of active AMPs	Number
Active AMPs Max	Maximum number of active AMPs	Number
Active AMPs Min	Minimum number of active AMPs	Number
Active AMPs Stddev	Standard deviation of active AMPs	Number
AMP CPU Skew Avg	Average percent of AMP CPU skew	Percent
AMP CPU Skew Max	Maximum percent of AMP CPU skew	Percent
AMP CPU Skew Min	Minimum percent of AMP CPU skew	Percent
AMP CPU Skew Stddev	Standard deviation percent of AMP CPU skew	Percent
AMP I/O Skew Avg	Average percent AMP I/O skew	Percent
AMP I/O Skew Max	Maximum percent AMP I/O skew	Percent
AMP I/O Skew Min	Minimum percent AMP I/O skew	Percent
AMP I/O Skew Stddev	Standard deviation of AMP I/O skew	Percent
Cache Total	Total queries found in step cache	Number
Concurrency	Average number of queries executing simultaneously during the hour	Number
Delay Time Avg	Average query delay time	Number
Delay Time Max	Maximum query delay time	Number
Delay Time Min	Minimum query delay time	Number

Metric	Description	Type
Delay Time Stddev	Standard deviation of query delay time	Number
First Response Time Avg	Average time it takes to receive the first response	Number
First Response Time Max	Maximum time it takes to receive the first response	Number
First Response Time Min	Minimum time it takes to receive the first response	Number
First Response Time Stddev	Standard deviation of time it takes to receive the first response	Number
Max Steps in Parallelism Avg	Average number of level-2 query steps executed in parallel	Number
Max Steps in Parallelism Max	Maximum number of level-2 query steps executed in parallel	Number
Max Steps in Parallelism Min	Minimum number of level-2 query steps executed in parallel	Number
Max Steps in Parallelism Stddev	Standard deviation of number of level-2 query steps executed in parallel	Number
Number of Steps Avg	Average number of query steps	Number
Number of Steps Max	Maximum number of query steps	Number
Number of Steps Min	Minimum number of query steps	Number
Number of Steps Stddev	Standard deviation of number of query steps	Number
Query Count	Queries logged in dbc.QryLog, or the sum of Querycount from dbc.QryLogsummary during the sample period	Number
Query Seconds Avg	Average query run time	Number
Query Seconds Max	Maximum query run time	Number
Query Seconds Min	Minimum query run time	Number
Query Seconds Stddev	Standard deviation of query run time	Number
Result Rows Avg	Average number of result rows for the query	Number
Result Rows Max	Maximum number of result rows for the query	Number
Result Rows Min	Minimum number of result rows for the query	Number
Result Rows Stddev	Standard deviation of number of result rows for the query	Number
Spool Usage Avg	Average amount of spool the query used	Number
Spool Usage Max	Maximum amount of spool the query used	Number
Spool Usage Min	Minimum amount of spool the query used	Number
Spool Usage Stddev	Standard deviation of amount of spool the query used	Number
Steps with Parallelism Avg	Average number of query steps with parallel steps	Number

Metric	Description	Type
Steps with Parallelism Max	Maximum number of query steps with parallel steps	Number
Steps with Parallelism Min	Minimum number of query steps with parallel steps	Number
Steps with Parallelism Stddev	Standard deviation of number of query steps with parallel steps	Number
Total CPU Time Avg	Average CPU time. CPU time is calculated as the total AMP CPU time plus the total parser and dispatcher CPU time for the query.	Number
Total CPU Time Max	Maximum CPU time. CPU time is calculated as the total AMP CPU time plus the total parser and dispatcher CPU time for the query.	Number
Total CPU Time Min	Minimum CPU time. CPU time is calculated as the total AMP CPU time plus the total parser and dispatcher CPU time for the query.	Number
Total CPU Time Stddev	Standard deviation of CPU time. CPU time is calculated as the total AMP CPU time plus the total parser and dispatcher CPU time for the query.	Number
Total I/O Count Avg	Average I/O count for the query	Number
Total I/O Count Max	Maximum I/O count for the query	Number
Total I/O Count Min	Minimum I/O count for the query	Number
Total I/O Count Stddev	Standard deviation of I/O count for the query	Number
Warning Total	Total queries with warnings	Number

System Metrics

Metric	Description	Type
Active Sessions	Sessions with active queries	Number
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP	Percent
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP	Percent
Average Memory Available	Average approximate free memory, as reported by the Linux “free” command, across all nodes	Number
AWT	Average number of AMP worker tasks in use on each AMP	Number
Components	Components, such as BYNETs or AMPs, that are down	Number
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.	Percent

Metric	Description	Type
CPU AWT Processing	CPU resources spent processing an AMP worker task. Normalized by multiplying the value returned from Analytics Database by the number of physical CPUs on the associated node of the AWT.	Percent
CPU/Disk Ratio	Ratio of CPU use to disk use. Calculated as the node CPU time divided by the physical disk usage.	Percent
CPU Dispatcher Processing	CPU resources spent in PE Dispatcher processing	Percent
DBC Disk Space	Available DBC disk space in use	Percent
Disk I/O	Disk I/Os	Number
Disk Reads	Total physical disk reads per system during the sample period	Number
Disk Usage	Disk use on the system	Percent
Disk Writes	Total physical disk writes per system during the sample period	Number
FSG Cache Miss	Percentage of FSG cache misses. Calculated as physical I/Os divided by logical I/Os.	Percent
Host Block Reads	Total message blocks (one or more messages sent in one physical group) received from all clients	Number
Host Block Writes	Total message blocks sent to all hosts	Number
Index Ratio	Percentage of index hits. Calculated as the index I/Os divided by the data block I/Os.	Percent
Logical MB/Sec	Logical I/O of the system in megabytes per second	Number
Max Disk Space By AMP	Available disk space currently in use	Percent
Max Spool Space By AMP	Available spool space in use	Percent
Minimum Memory Available	Minimum approximate free memory, as reported by the Linux “free” command, across all nodes. The value is from the node with the least free memory.	Number
Net A Usage	Percent of BYNET A usage (BYNET receiver usage)	Percent
Net Reads	Total reads from the BYNET during the sample period	Number
Net Writes	Total messages written to the BYNET during the sample period	Number
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
Node CPU Usage	Average percent of CPU usage of all online nodes in the configuration	Percent
Node I/O Skew	Comparison of disk use on the busiest node to the average node	Percent
Nodes Down	Number of system nodes down	Number

Metric	Description	Type
Parallelism	Percentage of parallelism, calculated as the average CPU usage / maximum CPU usage x 100	Percent
PE CPU Usage	Average CPU usage of the PEs. Calculated as the average usage of the parser and the dispatcher (both normalized values). Normalized by multiplying the value by the number of physical CPUs on the associated node of the PE.	Percent
Read I/O	Percentage of I/O that are reads	Percent
Session Login Count	Sessions currently logged on to the system	Number
Swap Drops	Total pages or segments dropped from memory during the sample period due to swapping	Number
Swap Reads	Total pages or segments read into memory from the disk after a prior write or drop during the sample period	Number
Swap Writes	Total pages or segments written into swap area from memory during the sample period	Number
System CPU	Average CPU time spent in System mode	Percent
Total AMP CPU	Total AMP CPU use	Percent
Total Disk Space	Total disk space currently in use	Percent
Total Node CPU	Total node CPU use	Percent
User CPU	Average CPU time spent in User mode	Percent
Vproc CPU Usage	Average percent of the CPU usage of all online vprocs in the configuration. Normalized by multiplying the value returned from Analytics Database by the number of physical CPUs in the associated vproc.	Percent
Wait I/O CPU	Average CPU time spent waiting for I/O	Percent

Table Metrics

Metric	Description	Type
Index Count	Total indexes	Number
Table Size	Table size in bytes, including index, fallback, and journal subtables	Number

Trend-Reporting Metrics for Aster System

Metric	Description	Type
Active Processes	Processes that are currently running	Number

Metric	Description	Type
Components Down	Nodes that are not available	Number
Components Passive	Nodes not processing queries but can be made ready to process queries when needed	Number
CPU	Average node CPU use	Percent
Disk I/O	Data transfer rate, in bytes per second	Number
Disk In	Rate of disk writes, in bytes per second	Number
Disk Out	Rate of disk reads, in bytes per second	Number
Max Disk By Node	For the node with the most used disk space, the amount of used disk space on the node	Percent
Memory	Average node memory use	Percent
Network In	Rate of network traffic into the node, in bytes per second	Number
Network Out	Rate of network traffic out of the node, in bytes per second	Number
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
Node I/O Skew	Comparison of I/O use on the busiest node to the average node	Percent
Queen Disk Space	Used disk space on the queen node	Percent
Session Login Count	Sessions currently logged on to the system	Number
Total Space	Used disk space on the system	Percent

Trend-Reporting Metrics for Teradata Machine Learning Engine

Metric	Description	Type
Active Processes	Processes that are currently running	Number
CPU	Average CPU use	Percent
Pod CPU Skew	Comparison of CPU use on the busiest queen or worker to that of the average one	Percent
Session Login Count	Sessions currently logged on to the system	Number

Trend-Reporting Metrics for Hadoop System

System Metrics (CDH, HDP 2.1 and later, and HDP 1.3)

Metric	Description	Type
CPU Idle	CPU time not processing any commands and the system not having an outstanding disk I/O request	Percent
CPU Nice	CPU time spent executing at the user level with nice priority	Percent
CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
CPU System	CPU time spent running kernel code	Percent
CPU Usage	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.	Percent
CPU User	CPU time spent running non-kernel code	Percent
CPU Wait I/O	CPU time spent waiting for I/O	Percent
Disk Skew	Comparison of disk space on the most full node to the average node	Percent
Disk Use	Disk space being used on a system	Percent
Load average last 15 minutes	Average number of jobs in the job queue over the last 15 minutes	Number
Load average last 5 minutes	Average number of jobs in the job queue over the last 5 minutes	Number
Load average last minute	Average number of jobs in the job queue over the last minute	Number
Memory Usage	Average memory use of the system during a sample period	Percent
Network In	Rate of network traffic into the node in bytes per second	Number
Network Out	Rate of network traffic out of the node in bytes per second	Number

HDFS Metrics (CDH, HDP 2.1 and later, and HDP 1.3)

Metric	Description	Type
Blocks Corrupt	Blocks whose replicas are all corrupt	Number
Blocks Excess	Blocks that exceed their target replication for the file they belong to	Number
Blocks Missing	Blocks with no replicas anywhere in the cluster	Number
Blocks Pending Deletion	Blocks waiting for deletion	Number
Blocks Pending Replication	Blocks waiting to be replicated	Number

Metric	Description	Type
Blocks Scheduled for Replication	Blocks scheduled for replication	Number
Blocks Under Replicated	Blocks that do not meet their target replication for the file they belong to	Number
Datanode I/O	Disk use on the data node	Number
Disk Capacity Used	Bytes of disk space currently used by HDFS	Number
Disk Usage	Available disk space used by HDFS	Percent
Files + Directories	Total files and directories in HDFS	Number
Files Appended	Files appended	Number
Files Created	Files created	Number
Files Deleted	Files deleted	Number
Total Load	Connections to HDFS	Number

YARN Metrics (CDH and HDP 2.1 and later)

The following YARN metrics are available to analyze resource usage for a Hadoop system.

Metric	Description	Type
Applications Completed	YARN applications that were completed during the interval	Number
Applications Failed	YARN applications that failed during the interval	Number
Applications Running	Average YARN applications that were running during the interval	Number
Applications Submitted	YARN applications that were submitted during the interval	Number
Cluster Memory Allocated	Available cluster memory that was marked as allocated	Percent
Cluster Memory Reserved	Available memory that was marked as reserved	Percent
Cluster Memory Skew	Skew of the cluster memory across the different NodeManager instances	Percent
YARN Containers Allocated	Average YARN containers that were allocated/running during the interval	Number

MapReduce Metrics (HDP 1.3)

Metric	Description	Type
Jobs Completed	Jobs that completed successfully	Number
Jobs Failed	Jobs that failed before completion	Number


Metric	Description	Type
Jobs Running	Jobs currently executing in the system	Number
Jobs Submitted	Jobs queued to execute in the system	Number
Map Tasks Completed	Map tasks completed successfully	Number
Map Tasks Failed	Map tasks that failed before completion	Number
Map Tasks Launched	Map tasks opened	Number
Map Tasks Running	Map tasks currently executing in the system	Number
Map Tasks Waiting	Map tasks queued to run	Number
Reduce Tasks Completed	Reduce tasks completed successfully	Number
Reduce Tasks Failed	Reduce tasks that failed before completion	Number
Reduce Tasks Launched	Reduce tasks opened	Number
Reduce Tasks Running	Reduce tasks running	Number
Reduce Tasks Waiting	Reduce tasks queued to run	Number

Settings View

The **Settings** view allows you to specify the work days and hours during which most users are using system resources, which visually emphasizes the information in the heatmap view.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Configuring the Work Week and Shift

You can outline specific work days and hours with a dark blue border in the heatmap view. The work-week-and-shift-border helps you easily identify the days and hours when system resources are used most.

1. Click  in the portlet frame and select **Settings**.
2. Select the number of shifts (work days) per week from the **Total number of shifts per week** list. A single row displays for each shift.
3. For each shift, click the appropriate arrow, and then select the following:
 - **Start Day**
 - **Start Time**
 - **Shift Duration**
4. Click **OK**.

Overlapping Shifts

When multiple shifts are defined on the same day and the hours you specify overlap, they are combined into one shift and are enclosed in one set of shift-indicator lines in the heatmap view.

For example, assume the first shift starts at 6 AM and lasts for 8 hours and the second shift starts at 1 PM and lasts for 8 hours. After you save these settings, the work-week lines and shift-indicator lines show a single shift from 6 AM to 9 PM. The **Settings** view displays the **Total number of shifts per week** reduced by one, and the **Shift Duration** indicates a longer single shift.

Metrics Analysis

Metrics Analysis Overview

The **Metrics Analysis** portlet allows you to analyze resource usage trends for one or more systems. Trends are graphed according to metrics such as CPU, memory, and throughput within a specified time frame. The information is refreshed every 60 seconds.

The **Metrics Analysis** view displays a graph with time on the horizontal axis, metric values on the vertical axis, and a different color for each metric plot line. Plot lines show the average metric values as a line and have enhanced interactive features, such as performance envelopes and information balloons that show the actual, minimum, and maximum metric values. A time frame selection list allows you to monitor current usage within the last hour or view usage trends over the last 24 hours, 1 week, 1 month, 3 months, or 6 months.

Note:

Not all metrics display all of these features. Cumulative metrics, such as Jobs Completed and Applications Completed, do not display performance envelopes, minimum and maximum values, or average values.

The **Settings** view allows you to select and organize metrics to display, set thresholds, and adjust the vertical axis range for each metric.

Metrics Analysis View

The **Metrics Analysis** view displays a graph that represents system resource usage. One or more systems can be monitored in the time frame using the same or different metrics. Metrics can be selected with different thresholds and vertical axis ranges.

The following list describes the features in this view:

Toolbar

Contains the time frame list. Select a time frame from the list to change the graph. **Export** creates a .csv file containing selected data.

Line Graph

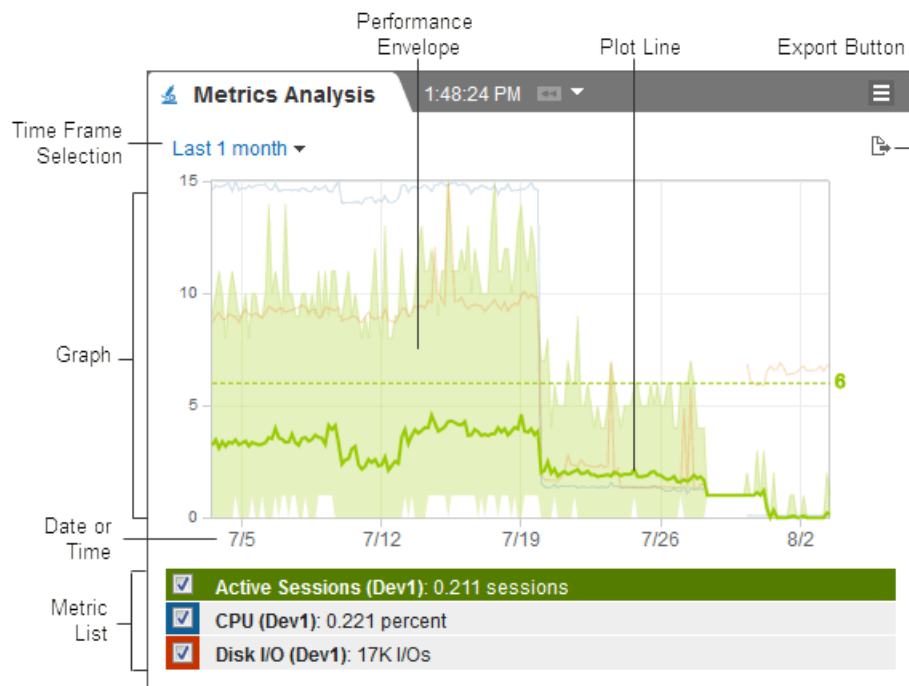
Plots the time frame on the horizontal axis and the metric values on the vertical axis. The *plot line* is a colored line that represents the average metric values. Click a plot line to see the highlighted plot line and performance envelope in a lighter shade of the assigned color. The *performance envelope* represents the upper and lower metric values.

Time Frame

Shows the frame of time selected from the time frame list on the toolbar. For example, if one hour is selected, times from the previous hour appear across the bottom of the graph.

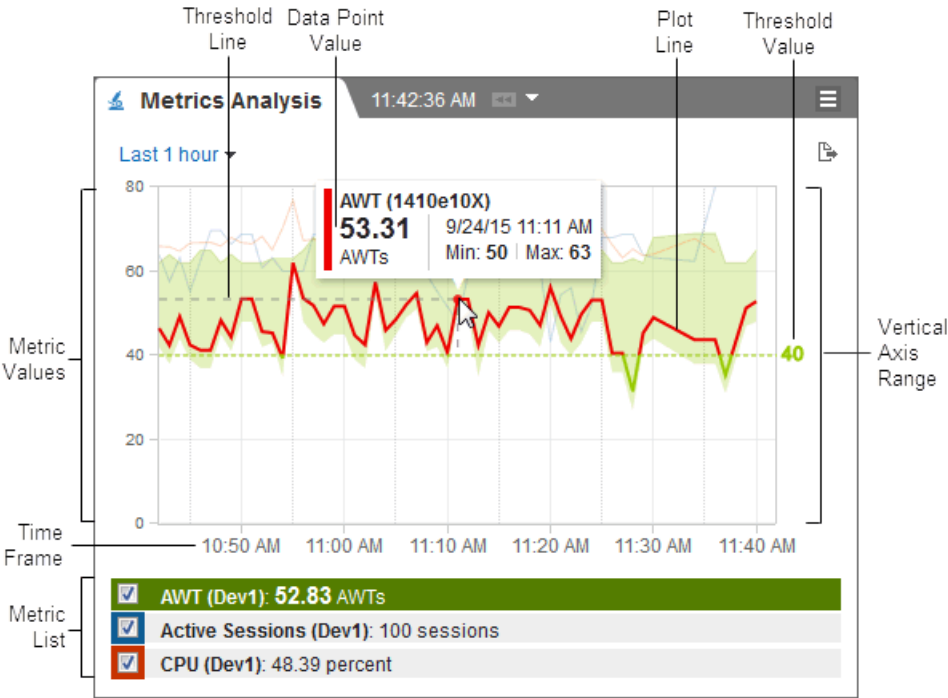
Metrics

Lists up to 10 metrics that are defined in the **Settings** view. Each list position is assigned a color and a check box. The metric name appears along with the name of the system being monitored, if there is more than one system enabled. Select a metric name to highlight the plot line and see the performance envelope in the graph.



Thresholds

Set thresholds to highlight metric values that are outside normal operating ranges. Set maximum vertical axis ranges to limit the range of values displayed for selected metrics. After a metric threshold is set in the **Settings** view, the threshold line is displayed across the graph with the threshold value to the right of the graph when you select a metric plot line. As you hover over the selected plot line, information balloons display the data point value and the minimum and maximum performance envelope values. The plot line and information balloons change to red when the threshold is exceeded during the time frame. Threshold settings are optional and can be set any time after the metric has been configured.



Setting the Time Frame

You can set the time frame that is used to plot the graph.

1. In the toolbar, select a time frame from the list.
The portlet refreshes, and the graph is redrawn.

Removing Metrics from the Graph

You can disable a metric plot line to remove it from the graph and the metric remains in the list.

1. Clear a colored check box in the metric list.

Toolbar

The toolbar allows you to choose the time frame to display in the graph. The graph displays the oldest data on the left and the most recent data on the right. For each metric, data is collected every 15 seconds and averaged according to the time frame chosen.

Option	Description
Last 1 hour	Displays the metric values for the last hour, plotted by minute
Last 24 hours	Displays the metric values for the last 24 hours, plotted by 15-minute periods


Option	Description
Last 1 week	Displays the metric values for the last week, plotted by hour
Last 1 month	Displays the metric values for the last month, plotted by 4-hour periods
Last 3 months	Displays the metric values for the last 3 months, plotted by 12-hour periods
Last 6 months	Displays the metric values for the last 6 months, plotted by day

Selecting a different time frame causes the portlet to immediately recalculate and update the graph based on the data points collected for the metrics.

The **Export** button allows you to create a .csv file containing selected data.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains data for the selected time period. Minimum, average, and maximum is exported for all metrics configured in the **Settings** view.

1. In the toolbar, select a time frame from the list and select 

Option	Frequency
Last 1 hour	Data points at 1-minute intervals
Last 24 hours	Data points at 15-minute intervals
Last 1 week	Data points at 1-hour intervals
Last 1 month (Export)	Data points at 4-hour intervals
Last 1 month (Export By Hour)	Data points at 1-hour intervals
Last 3 months (Export)	Data points at 12-hour intervals
Last 3 months (Export By Hour)	Data points at 1-hour intervals
Last 6 months (Export)	Data points at 1-day intervals
Last 6 months (Export By Hour)	Data points at 1-hour intervals

2. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
3. [Optional] Reformat exported data to match the format in the view.

Analysis Metrics for Teradata System

The following metrics are available to analyze resource usage for a Teradata system.

Metric	Description	Type
Active Sessions	Sessions with active queries	Number
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP	Percent
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP	Percent
Average Memory Available	Average approximate free memory, as reported by the Linux “free” command, across all nodes	Number
Average Outstanding Disk Reqs	Outstanding disk requests (disk queue size)	Number
AWT	Average AMP worker tasks in use on each AMP	Number
Components	Components, such as BYNETs or AMPs that are down	Number
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.	Percent
CPU AWT Processing	CPU resources spent processing an AMP worker task. Normalized by multiplying the value returned from Analytics Database by the number of physical CPUs on the associated node of the AWT.	Percent
CPU/Disk Ratio	Ratio of CPU use to disk use. Calculated as the node CPU time divided by the physical disk usage.	Percent
CPU Dispatcher Processing	CPU resources spent in PE Dispatcher processing	Percent
DBC Disk Space	Available DBC disk space in use	Percent
Disk I/O	Total number of physical disks read and write	Number
Disk Reads	Total physical disk reads per system during the sample period	Number
Disk Usage	Disk use on the system	Percent
Disk Writes	Total physical disk writes per system during the sample period	Number
Enabled CPU Usage	<p>CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled</p> <p>Tooltip content:</p> <ul style="list-style-type: none"> • TDEnabledCPUs % (of NCPUs) : Percentage of Teradata enabled CPUs of total NCPUs • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system 	Percent
FSG Cache Miss	Percentage of FSG cache misses. Calculated as physical I/Os divided by logical I/Os.	Percent
Host Block Reads	Message blocks (one or more messages sent in one physical group) received from all clients	Number
Host Block Writes	Message blocks sent to all hosts	Number

Metric	Description	Type
Index Ratio	Index I/Os divided by the data block I/Os	Percent
Logical MB/Sec	Logical I/O of the system in megabytes per second	Number
Max Disk Space By AMP	Available disk space currently in use	Percent
Max Spool Space By AMP	Available spool space in use	Percent
Minimum Memory Available	Minimum approximate free memory, as reported by the Linux “free” command, across all nodes. The value is from the node with the least free memory.	Number
Net A Usage	BYNET A usage (BYNET receiver usage)	Percent
Net Reads	Reads from the BYNET during the sample period	Number
Net Writes	Messages written to the BYNET during the sample period	Number
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
Node CPU Usage	Average CPU usage of all online nodes in the configuration	Percent
Node I/O Skew	Comparison of disk use on the busiest node to the average node	Percent
Nodes Down	System nodes down	Number
Parallelism	Percentage of parallelism, calculated as the average CPU usage divided by the maximum CPU usage x 100	Percent
PE CPU Usage	Average CPU usage of the PEs, calculated as the average usage of the parser and the dispatcher (both normalized values). Normalized by multiplying the value by the number of physical CPUs on the associated node of the PE.	Percent
Read I/O	Percent of I/O that are reads	Percent
Session Login Count	Sessions currently logged on to the system	Number
Swap Drops	Pages or segments dropped from memory during the sample period due to swapping	Number
Swap Reads	Pages or segments read into memory from the disk after a prior write or drop during the sample period	Number
Swap Writes	Pages or segments written into swap area from memory during the sample period	Number
System CPU	Average CPU time spent executing kernel system calls or servicing I/O and timer hardware interrupts	Percent
Total AMP CPU	Total AMP CPU use	Percent
Total Disk Space	Total disk space in use	Percent
Total Node CPU	Total node CPU	Percent

Metric	Description	Type
	Tooltip content for IFX 2.1 and higher systems with Elastic TCore enabled: <ul style="list-style-type: none"> • TDEnabledCPUs % (of NCPUs) : Percentage of Teradata enabled CPUs of total NCPUs • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system 	
User CPU	Average CPU time spent executing code on behalf of operating system user processes	Percent
Vproc CPU Usage	Average CPU usage of all online vprocs in the configuration	Percent
Wait I/O CPU	Average CPU time spent waiting for I/O	Percent

Analysis Metrics for Aster System

The following metrics are available to analyze resource usage for an Aster system.

Metric	Description	Type
Active Processes	Processes that are currently running	Number
Components Down	Nodes that are not available	Number
Components Passive	Nodes not processing queries but can be made ready to process queries when needed	Number
CPU	Average node CPU use	Percent
Disk I/O	Data transfer rate, in bytes per second	Number
Disk In	Rate of disk writes, in bytes per second	Number
Disk Out	Rate of disk reads, in bytes per second	Number
Max Disk By Node	For the node with the most used disk space, the amount of used disk space on the node	Percent
Memory	Average node memory use	Percent
Network In	Rate of network traffic into the node, in bytes per second	Number
Network Out	Rate of network traffic out of the node, in bytes per second	Number
Node CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
Node I/O Skew	Comparison of I/O use on the busiest node to the average node	Percent
Queen Disk Space	Used disk space on the queen node	Percent
Session Login Count	Sessions currently logged on to the system	Number

Metric	Description	Type
Total Space	Used disk space on the system	Percent

Analysis Metrics for Teradata Machine Learning Engine

The following metrics are available to analyze resource usage for a Teradata Machine Learning Engine system.

Metric	Description	Type
Active Processes	Processes that are currently running	Number
CPU	Average CPU use	Percent
Pod CPU Skew	Comparison of CPU use on the busiest queen or worker to the average one	Percent
Session Login Count	Sessions currently logged on to the system	Number

Analysis Metrics for Hadoop System

System Metrics (CDH, HDP 2.1 and later, and HDP 1.3)

The system metrics listed in the following table are available to analyze resource usage.

Metric	Description	Type
CPU Idle	CPU time not processing any commands and the system not having an outstanding disk I/O request	Percent
CPU Nice	CPU time spent executing at the user level with nice priority	Percent
CPU Skew	Comparison of CPU use on the busiest node to the average node	Percent
CPU System	CPU time spent running kernel code	Percent
CPU Usage	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.	Percent
CPU User	CPU time spent running non-kernel code	Percent
CPU Wait I/O	CPU time spent waiting for I/O	Percent
Disk Skew	Comparison of disk space on the most full node to the average node	Percent
Disk Use	Disk space being used on a system	Percent
Load average last 15 minutes	Average number of jobs in the job queue over the last 15 minutes	Number

Metric	Description	Type
Load average last 5 minutes	Average number of jobs in the job queue over the last 5 minutes	Number
Load average last minute	Average number of jobs in the job queue over the last minute	Number
Memory Usage	Average memory use of the system during a sample period	Percent
Network In	Rate of network traffic into the node in bytes per second	Number
Network Out	Rate of network traffic out of the node in bytes per second	Number

HDFS Metrics (CDH, HDP 2.1 and later, and HDP 1.3)

The HDFS metrics listed in the following table are available to analyze resource usage.

Metric	Description	Type
Blocks Corrupt	Blocks whose replicas are all corrupt	Number
Blocks Excess	Blocks that exceed their target replication for the file they belong to	Number
Blocks Missing	Blocks with no replicas anywhere in the cluster	Number
Blocks Pending Deletion	Blocks waiting for deletion	Number
Blocks Pending Replication	Blocks waiting to be replicated	Number
Blocks Scheduled for Replication	Blocks scheduled for replication	Number
Blocks Under Replicated	Blocks that do not meet their target replication for the file they belong to	Number
Datanode I/O	Disk use on the data node	Number
Disk Capacity Used	Bytes of disk space currently used by HDFS	Number
Disk Usage	Available disk space used by HDFS	Percent
Files + Directories	Total files and directories in HDFS	Number
Files Appended	Files appended	Number
Files Created	Files created	Number
Files Deleted	Files deleted	Number
Total Load	Connections to HDFS	Number

YARN Metrics (CDH and HDP 2.1 and later)

The YARN metrics listed in the following table are available to analyze resource usage.

Metric	Description	Type
Applications Completed	YARN applications that were completed during the interval	Number
Applications Failed	YARN applications that failed during the interval	Number
Applications Running	Average YARN applications that were running during the interval	Number
Applications Submitted	YARN applications that were submitted during the interval	Number
Cluster Memory Allocated	Available cluster memory that was marked as allocated	Percent
Cluster Memory Reserved	Available memory that was marked as reserved	Percent
Cluster Memory Skew	Skew of the cluster memory across the different NodeManager instances	Percent
YARN Containers Allocated	Average YARN containers that were allocated/running during the interval	Number

MapReduce Metrics (HDP 1.3)

The MapReduce metrics listed in the following table are available to analyze resource usage.

Metric	Description	Type
Jobs Completed	Jobs that completed successfully	Number
Jobs Failed	Jobs that failed before completion	Number
Jobs Running	Jobs currently executing in the system	Number
Jobs Submitted	Jobs queued to execute in the system	Number
Map Tasks Completed	Map tasks completed successfully	Number
Map Tasks Failed	Map tasks that failed before completion	Number
Map Tasks Launched	Map tasks opened	Number
Map Tasks Running	Map tasks currently executing in the system	Number
Map Tasks Waiting	Map tasks queued to run	Number
Reduce Tasks Completed	Reduce tasks completed successfully	Number
Reduce Tasks Failed	Reduce tasks that failed before completion	Number
Reduce Tasks Launched	Reduce tasks opened	Number
Reduce Tasks Running	Reduce tasks running	Number


Metric	Description	Type
Reduce Tasks Waiting	Reduce tasks queued to run	Number

Settings View

The **Settings** view allows you to select which systems and metrics to monitor and display. You can set thresholds and vertical axis ranges for selected metrics.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.





1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.



Adding Metrics

You can add metrics that are then plotted in the graph.

1. Click  in the portlet frame and select **Settings**.
2. [Optional] Select the **Enable** check box to enable the metric for display.
3. [Optional] Select a system from the list.
The **System** list is available when more than one Teradata Viewpoint server is enabled for the Data Collection Service.
4. Select a metric from the list.
You can add  or remove  metrics. Up to 10 metrics can be listed.
5. [Optional] Enter a **Threshold** value.
6. [Optional] Enter a **Vertical Axis Range** value.
7. [Optional] Click  and drag the row to reorder the metrics for display.
8. Click **OK**.


Changing the Metrics Display Order

You can change the order of the metrics that are listed under the graph. Reordering the metric list affects which color is assigned to the metric.

1. Click  in the portlet frame and select **Settings**.
2. On the metric row, click  and drag the row to its new location.
3. Click **OK**.



Enabling and Disabling Metrics

You can enable and disable metrics from the graph. When a metric is disabled, the metric plot line is not displayed in the graph. Data points are still being collected, so the metric is included in the list under the graph and can be reactivated at a later time.

1. Click  in the portlet frame and select **Settings**.
2. Do one of the following to enable or disable a metric:
 - Select **Enable** on the metric line to enable the metric.
 - Clear **Enable** on the metric line to disable the metric.
3. Click **OK**.

Deleting Metrics

You can delete metrics from the graph. When a metric is deleted in the **Settings** view, the metric plot line is not displayed in the graph, and the metric is not included in the metric list.

1. Click  in the portlet frame and select **Settings**.
2. Click  on the row of the metric you want to delete.
The metric row disappears.
3. Click **OK**.

My Queries

My Queries Overview

The **My Queries** portlet allows you to view and manage your queries across multiple Teradata systems and in Teradata QueryGrid. You can see if queries are queued or blocked, and you can see their impact on system resources.

Use the **My Queries** portlet to view information about queries in either the summary view or the details view. The summary view contains a table with one row allocated to each of the sessions logged on under one or more user names. Select a row in the summary view to see additional session and query information in the details view. Use the **SQL**, **Explain**, **Blocked By**, or **Query Band** tab in the details view to display information for the selected session.

The **Settings** view allows you to select Teradata systems and accounts to monitor, and select a format for the SQL that appears in the query details view.

Note:

- If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is not available. You cannot enter alternate credentials.
 - Support for TASM ARM is available for SQL Engine versions 17.10 and later.
-

My Queries View

After you add an account in the **Profile** portlet for each Teradata system you want to monitor, the **My Queries** view displays queries running on the selected systems under your login.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays information about each session, with columns configured specifically for the current view. These statistics are sampled and the table is refreshed every 30 seconds. Highlighted values in any row indicate that a session has exceeded threshold criteria for a specific metric. You can click anywhere in a row to see session details.

▼ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

The screenshot shows the 'My Queries' interface. At the top, there's a header bar with 'My Queries', a time '2:49:55 PM', and a menu icon. Below the header, there's a 'Filters' section with input fields for 'SESSION ID', 'USERNAME', and 'CPU USE'. To the right of the filters is a 'Table Actions' dropdown menu. The main area is a 'Summary Table' with columns: 'SESSION ID', 'USERNAME', 'CPU USE', 'STATE ICON', and 'IN STATE'. The table contains 8 rows of data. At the bottom of the table, it says '36 rows total'.

SESSION ID	USERNAME	CPU USE	STATE ICON	IN STATE
85505	DBC	0.003	⊞	4:30:08
1149114	VIEWPOINT	0	⊞	4:29:42
1188475	VIEWPOINT	0	⊞	4:29:42
1212259	VIEWPOINT	0	⊞	4:29:42
1230229	VIEWPOINT	0	⊞	4:29:42
1232488	VIEWPOINT	0	⊞	4:29:42
1250550	VIEWPOINT	0.001	⊞	4:29:42
1270834	VIEWPOINT	0	⊞	4:29:42
1410097	VIEWPOINT	0	⊞	4:29:42

36 rows total

Session View Metrics

Metrics available for monitoring and display are listed as following. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Account	Account from which a query was submitted	String
Blocked Time	How long the query has been blocked	Number
CPU Use	Available CPU seconds on the system used during the last sampling period	Percent
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number
Duration	How long the query has been running	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
In State	How long the query has been in the current state	Number

Metric	Description	Type
Partition	Partition in which the query is running	String
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String
Query Band	Set of name-value pairs for the session or transaction	String
Request CPU	CPU seconds needed to run the query	Number
Request I/O	Disk I/Os performed to run the query	Number
Request Count	Queries the session has executed	Number
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Session ID	Unique session identifier	Number
Snapshot CPU Skew	CPU skew during the last sample	Percent
Snapshot I/O Skew	I/O skew during the last sample	Percent
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.	Percent
Start	Time that the query started running on Analytics Database	Number
State	Text describing the current state of the query	String
State Icon	Icon representing the current state of the query	Icon
System	Full name of the system running the query	String
Temp Space	Temp space used by the query, using a 1024 byte kilobyte as the base	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time	Number
Username	Name of the user who submitted the query	String
Workload	Workload in which the query is running	String

State Icons

The following icons appear in the portlet.

Icons	Name	Definition
×	Aborting	Query has been aborted, and changes are being rolled back.
▶▶	Active	Query is running.
▶	Blocked	Query is waiting for a lock held by another query.
	Delayed	Query is in a delay queue waiting to run.
⌘	Deferred	Query is in a deferred queue waiting to run.
⏻	Host-Restart	Host associated with this session is currently restarting.
□	Idle	No query is running.
?	Other	Query is in an unknown state.
⋮	Parsing	Query is being parsed before running.
⋮ _{QT}	QTDelayed	Query delayed due to a queue table restriction.
↔	Response	Query is returning results to the user.
⋮ _{SES}	SESDelayed	Query is delayed because a utility limit has been exceeded in FastLoad, MultiLoad, FastExport, or ARC.

Details View for All Systems

The details view displays statistics and information about the selected session. This view can be accessed by clicking a session row in the summary table.

Tabs

Provides important query details on the **Overview**, **SQL**, **Explain**, **Skew**, **Blocked By**, **Query Band**, **Defer**, and **Delay** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

▼ Manage Queries

Manages the query or session for Teradata systems with **Abort**, **Change Priority**, **Change Workload**, and **Release Query** options.

< > Previous/Next

Allows you to move through sessions without returning to the summary table.

System: Prod1 | Session: 2824421

Manage Queries Previous/Next

Tabs: Overview SQL Explain Skew Query Band Blocked By

Query Details

QUERY INFO

State: » Active
Time in State: 0:01:00
Total Duration: 0:01:00

Spool Space: 159,232
Temp Space: 0

Request CPU: 0
Request I/O: -

PJI: 0
Unnecessary I/O: 2.73

WORKLOAD INFO

Name: WD_Mode...
Classification Mode: Auto

SESSION INFO

User: VIEWPOINT
Account: DBC
Partition: DBC/SQL
Requests: 11

SNAPSHOT INFO

CPU Use: 0%
Impact CPU: 5.62
Snapshot CPU Skew: 18.2%
Snapshot I/O Skew: 17.9%

Source:
(TCP/IP) b308 153.64.107.218 SKATE ;SKATECOP2/153.64.202.243:1025 C ID=79EF3A1A ROOT JDBC15.00.00.0 9;1.7.0_45 01 LSS

Overview Tab

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds

Query Information	Description
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11

Workload Information	Description
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	<p>How a query or session is assigned to a workload. Available values are:</p> <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload <p>This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.</p>
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Session Information


Session Information	Description
User	Name of the user that submitted the query
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a


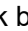
query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.





Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.


Explain Tab

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status.

If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	 - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Exporting Explain Data

1. From the Explain tab, click .
2. Save the file to a location.
The Explain data is exported as a .csv file.

Skew Tab

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Blocked By Tab

The **Blocked By** tab displays details about other queries that are blocking the selected query. This information is read-only. The tab is available only when the selected query is blocked.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table <code>databaseName.tableName</code> , Waiting on write lock on rowhash <code>databaseName.tableName</code> The information displays in the following order: 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock		Available
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Defer Tab

The Defer tab displays details about the rules that are deferring a query. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Defer Time	Duration of time request has been deferred by Arrival Rate Meter rules
User Name	User who submitted the query

Rule Information	Description
Rule Name	Name of rule causing request to defer
Rule Type	TASM type of rule causing request to defer
Overridable	Indicates if the Teradata DBA can abort or release the query

Delay Tab

The **Delay** tab displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Details View for Teradata QueryGrid

When you select a Teradata QueryGrid session from a database session or a Teradata QueryGrid manager in the portlet's summary view, the details for the query are displayed.

The top portion of the screen displays a graphical representation of the progress and status of the query, with icons representing the initiating system, the phase of the current operation being performed, and the target system. The initiating system is the data source from which the query is generated. The target system is the system against which the query is being run. Teradata QueryGrid queries go through a metadata and execution phase before they are completed. Teradata QueryGrid queries may also access data on target systems or transfer data from one system to another.

The lower part of the screen displays tabs with information about the system or operation. If you arrive at the details screen from a Teradata system session, details are displayed for the target or initiator system that is relevant for the selected step. If you arrive at the details screen while viewing a list of Teradata QueryGrid manager sessions, details for the selected operation are displayed. You can click on any of the icons in the graphical representation (initiator, operator, or target) to view the metrics for the selected system or operation.

Initiating System

The data source from which the query is generated.

Target System

The system from which data is accessed or to which it is transferred.

Operation

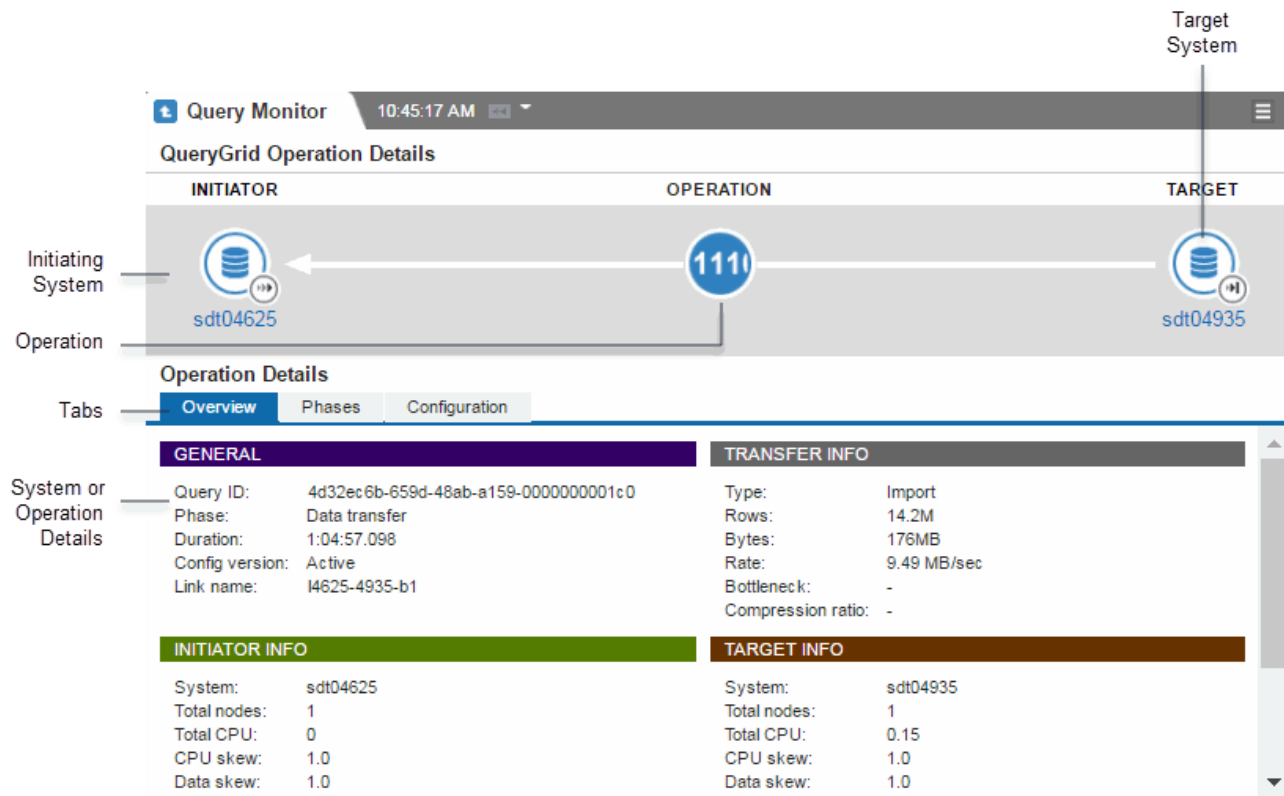
The phase of the operation being performed.

Tabs

Related metrics are organized into tabs for the system or operation selected in the top portion of the view.

System or Operation Details

The metrics related to the system or operation selected in the top portion of the view.




QueryGrid Icons

The details view for Teradata QueryGrid includes a graphical representation of the progress of the selected query. The icons here represent the initiating or target systems and the phases of the operations performed.







System Icons

The icons in the following table represent the query on the initiator and target systems. The icons inside the smaller circle represents the query state. In addition to the query states shown as following, additional possible states for the query are listed in [State Icons](#).

Icon	Name	Definition
	No query is running	No query is currently running
	Active	The query is running on the system

Icon	Name	Definition
	Idle	The query is currently idle

Operation Phase Icons

Icon	Name	Definition
	Metadata Operation	A handshake between systems to negotiate data types and validate queries.
	Execution Operation	Query processing prior to transferring the data.
	Data transfer	The transferring of data from one system to another.
	Completed	Query has successfully finished processing.
	Failed	Query did not finish processing successfully. [QueryGrid systems ≥ 2.10] To download the support bundle, click Download Support Query Bundle .
	Unknown	Query processing status is not currently available.

Initiator and Target Metrics

When an initiator or target system icon is selected in the details view for Teradata QueryGrid, two tabs display metrics for the system: an **Overview** tab and a **SQL** tab.

Overview Tab for Teradata System

The **Overview Tab** displays information for a QueryGrid query that was initiated by or is a target of a Teradata system.

QueryGrid Information

Metric	Description	Type
System	Name of the initiating or target system	String

Metric	Description	Type
Connector	Point from which the query originates or ends	String
Software version	The version of the connector	String
Network	The logical network used for this side of the data transfer	String

Session Information

Metric	Description	Type
User	User who submitted the query	String
Query ID	ID of the query	String
Session Number	Number of the session	String
Request Number	Number of the request	String
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session	String


Query Information

Metric	Description	Type
State	State of the query	String
Duration	Amount of time the query has been running (HH:MM:SS)	Number
Workload	Workload in which the query is running	String
Request CPU	Total request CPU (seconds)	Number
Request I/O	Total request I/O count	Number

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Operations Metrics

When an operation phase icon is selected in the details view for Teradata QueryGrid, three tabs display metrics for the operation: an **Overview** tab, a **Phases** tab, and a **Configuration** tab.

Overview Tab

The **Overview** tab displays general information about the selected operation for a Teradata QueryGrid query session.

General

Metric	Description	Type
Query ID	ID of the query	String
Phase	Phase of the operation	String
Duration	Duration of the QueryGrid operation	Number
Config version	Requested version of the QueryGrid operation	Number
Link Name	Name of the link used for this QueryGrid operation	String

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Total nodes	Number of nodes on the initiating system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Transfer Information

Metric	Description	Type
Type	Type of transfer performed	String
Rows	Number of rows transferred	Number
Bytes	Number of bytes transferred	Number
Rate	Rate that data is transferred for the QueryGrid operation	Number
Bottleneck	Limiting factor for the QueryGrid operation	String
Compression ratio	Uncompressed size divided by the compressed size	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Total nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Total Nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Phases Tab

The **Phases** tab is available when you click an icon that represents a Teradata QueryGrid query operation. It contains tabs that provide information about the metadata, execution, and data transfer phases of the query session. The **Phases** tab support up to five panes depending on the details associated with the operation, for example, Metadata1, Metadata2, Metadata3, Execution1, Execution2, Data Transfer.

Metadata Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	The node that initiated the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number

Metric	Description	Type
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Execution Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	Node from which the query initiated	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number

Metric	Description	Type
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Data Transfer Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	String
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number

Metric	Description	Type
CPU skew node	Node that had the max CPU	String
Data skew	Node that had the max data	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Total nodes	Number of nodes	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	Number
Data skew node	Node that had the maximum amount of data	Number
Bytes sent	Number of bytes sent	Number
Bytes received	Number of bytes received	Number
Compression ratio	Whether or not compression is enabled	String

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Configuration Tab

The **Configuration** tab contains provides configuration information about the selected operation for a QueryGrid query session.

Fabric

Metric	Description	Type
Name	Name of the fabric used in Teradata QueryGrid	String
Port	Port on which the fabric runs	String
Software version	Version of software that the fabric runs	String

Initiating Connector

Metric	Description	Type
System	Name of the initiating system	String
Connector	Point from which the query originated	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Communication Policy

Metric	Description	Type
Name	Name given to this communication policy	String
Encryption	Whether or not data encryption is enabled	String
Compression	Whether or not compression is enabled	String

Target Connector

Metric	Description	Type
System	Name of the target system	String
Connector	Destination point for the query	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Errors Tab

The **Errors** tab provides information about Teradata QueryGrid queries that had errors. It displays only if there are errors.

Metric	Description	Type
Time	Time the error occurred	Number
Phase	Phase when the error occurred	String
Message	The error message	String
Node	The hostname of the node that generated the error	String
Connector	The name of the connector that resulted in the error	String
System	Name of the system	String
Origin	Indicates whether the error occurred on the initiator or the target of the Teradata QueryGrid operation	String

Managing Queries and Sessions

You can manage queries and sessions to improve workload performance for Teradata systems.

Abort

Abort the selected query or session

Change Priority

Change the priority of the selected query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change priorities or workloads, or release queries. If you log off, close a portlet, or open a new portlet instance, you must log in again.

Note:

Change Workload is available only if Teradata Active System Management (TASM) is enabled. If TASM is disabled, **Change Priority** is available. If you do not see **Change Workload** or **Change Priority** in the list, the system you are monitoring does not support these features or you do not have permission to use them. If the query you are monitoring is delayed, only **Release Query** is available.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. Click the row of the query you want to abort.
2. Click ▼ to the right of the session number and select **Abort**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
6. Click **Next**.
7. Click **Next** to confirm your selection.
8. Click **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Priority**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.

- Type an account string.
6. [Optional] Select the check box to use the account string as the default for the selected session. The priority will be changed for the selected query and all subsequent queries in the current session.
 7. Click **Next**.
 8. Click **Next** to confirm your selection.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Workload**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Select a different workload from the list and click **Next**.
5. Click **Next** to confirm your selection.
6. Click **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Release Query**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Click **Next** to confirm your selection.

Settings View

The **Settings** view provides the following tabs that allow you to select systems and accounts to monitor and to change display options.

Systems


Select the Teradata systems and accounts to monitor.

Display

Choose formatted or unformatted SQL to display on the **SQL** tab of the query details view.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:


Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Selecting Systems and Accounts to Monitor


Use the **Systems** tab in the **Settings** view to select Teradata systems and accounts to monitor.

Add and associate accounts for each system using the **Profile** portlet. Only predefined accounts appear on the **Systems** tab.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Systems** tab.
3. For each Teradata system you want to monitor:
 - a. Select the check box next to the system.
 - b. Select the check boxes next to the accounts you want to monitor.
4. Click **OK**.

Selecting Display Options

The **Display** tab in the **Settings** view allows you to format the SQL that appears on the **SQL** tab of the query details view.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Display** tab.
3. Do one of the following for Teradata systems:
 - Select the **Format SQL** check box to display formatted SQL.
 - Clear the **Format SQL** check box to display unformatted SQL.
4. Click **OK**.

Node Monitor

Node Monitor Overview

The **Node Monitor** portlet enables you to monitor nodes on Teradata Aster and Teradata Hadoop systems and queens and workers in Teradata Machine Learning Engine.

For a Teradata Aster system, the portlet displays the operation and status of worker, loader, and queen nodes, as well as the amount of data stored on the node. You can view details such as status and type for each node on the system. Hardware statistics, including CPU and Memory, are available for selected time intervals.

For Teradata Machine Learning Engine, you can view CPU and network traffic information at the pod level.

For a Teradata Hadoop system, you can view node-level metrics, available Hadoop components (Hortonworks and Cloudera), and the status of components for each node on the system. You can view hardware statistics details such as CPU usage, memory usage, and network activity, as well as detailed service component and JVM metrics for the HDFS and YARN (HDP 2.1 and later) or MapReduce (HDP 1.3) services.

Node Monitor View

The **Node Monitor** view displays summary information about the nodes so you can monitor node performance and identify issues. The view displays on a Teradata Machine Learning Engine or in a Teradata Aster or Teradata Hadoop system.

System Selection

Shows the name of the system currently displayed.

State Filter Bar

Displays a count of the nodes in each state. Click a state to show only the nodes in that state.

Filters

Filters allow you to change displayed data by showing only rows that match your filter criteria. Sort on the column headers to group data in ascending or descending order.

Summary Table

Displays summary information for all nodes. The table is configured in the **Configure Columns** dialog box. The view is refreshed every 30 seconds. Click a row in the table to see details for that row.

☐ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

The screenshot shows the Node Monitor interface. At the top, there's a 'System Selection' dropdown set to 'Prod2' and a timestamp '11:21:03 AM'. Below this is the 'State Filter Bar' with buttons for 'All' (4), 'Active' (2), 'Failed' (1), 'Prepared' (0), and 'Passive' (1). The 'All' button is selected. Below the filter bar is a 'Filters' section with input fields for 'NAME', 'IP ADDRESS', 'TYPE', 'STATE ...', and '% FULL'. The 'Summary Table' displays four rows of data:

NAME	IP ADDRESS	TYPE	STATE ...	% FULL
10.25.32.225	10.25.32.225	QUEEN	»»	68.86
10.25.33.22	10.25.33.22	WORKER	»	56.53
10.25.33.227	10.25.33.227	WORKER	»»	60
10.25.33.28	10.25.33.28	LOADER	✖	44.41

At the bottom of the table, it says '4 rows total'.

State Filter Bar

The state filter bar allows you to display specific states in the view.

The diagram shows the 'State Filter Buttons' with the following counts: 'All' (5), 'Active' (1), 'Failed' (0), 'Prepared' (4), and 'Passive' (0). The 'All' button is highlighted.

The state filter buttons provide a count of nodes in each state. When you click a state in the state filter bar, only nodes in the selected state appear in the summary table.

For Teradata Aster systems, these states are available:

All

Total nodes on the selected system

Active

Nodes that are online and available for work

Failed

Nodes that are not participating in the nCluster

Prepared

Nodes that are ready to be incorporated in the nCluster and able to host a virtual worker

Passive

Nodes that contain copies of virtual worker data and that can be made active with a rebalance operation

For Teradata Hadoop systems, these states are available:

All

Total nodes on the selected Teradata Hadoop system

Master Node










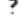
Nodes not running a datanode component

Data Node

Nodes running a datanode component

State Icons

The following icons appear in the **Node Monitor** portlet when a Teradata Aster system is selected.

Icons	Name	Definition
	Active	Node is online and available to process data
	Cleaning	Node is being cleared of all data
	Failed	Node is offline and not participating in the nCluster
	Installing	Node is being added to a cluster
	New	Node recently added to the system
	Passive	Node holds copies of virtual worker data
	Prepared	Node is ready to be incorporated in the cluster and can host a virtual worker
	Preparing	Node is preparing to host a virtual worker
	Stopped	Node is powered down
	Unknown	Node is not registered

Node Monitor Metrics

Available metrics differ depending on the type of system and selection menu options you choose.

Node Monitor Metrics for Aster System

Metric	Description
% Full	Percent of disk space used on the node
CPU	Node CPU use
Disk In	Rate of disk writes in bytes per second
Disk Out	Rate of disk reads in bytes per second
IP Address	IP address of the node
Name	Name of the node
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
Node ID	Unique ID for the node
State	Process state, such as active, installing, passive
State Icon	Symbol of the process state
Type	Values for node types are as follows: <ul style="list-style-type: none"> • Worker nodes hold the data and process the queries • Queen nodes manage the process, including the performance of the other nodes and delegation of queries • Loader nodes transfer data in and out of virtual workers

Node Monitor Metrics for Teradata Machine Learning Engine

Metric	Description
CPU	CPU usage
Name	Name of the queen or worker
Network In	Rate of network traffic into the queen or worker in bytes per second
Network Out	Rate of network traffic out of the queen or worker in bytes per second
State	Process state, such as active, installing, passive
State Icon	Symbol of the process state
Type	Values are as follows: <ul style="list-style-type: none"> • Worker nodes hold the data and process the queries • Queen nodes manage the process, including the performance of the workers and delegation of queries

Node Monitor Metrics for Hadoop System

CDH

Metric	Description
Containers Allocated	Number of YARN containers currently allocated for the NodeManager instance on this node
CPU	Node CPU use
HDFS I/O	Number of bytes read and written for Hadoop
IP Address	IP address of the node
Memory	Percentage of memory use on the node
Name	Name of the node
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
NM Memory Allocated	Percent of available memory allocated by the NodeManager instance on this node
Rack ID	User-defined name that specifies the physical location of the node
Type	Class of node, such as Master or Data: <ul style="list-style-type: none"> • Data nodes are those running the HDFS datanode service or the YARN NodeManager service • Master nodes are those running the HDFS namenode service or the YARN ResourceManager service • Nodes that are not running any HDFS or YARN services are classified as type other
Unhealthy Services	Count of services for the node that are in an unhealthy state

HDP 2.1 and later

Metric	Description
Containers Allocated	Number of YARN containers currently allocated for the NodeManager instance on this node
CPU	Node CPU use
HDFS I/O	Number of bytes read and written for Hadoop
IP Address	IP address of the node
Memory	Percentage of memory use on the node
Name	Name of the node

Metric	Description
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
NM Memory Allocated	Percent of available memory allocated by the NodeManager instance on this node
Rack ID	User-defined name that specifies the physical location of the node
Components Down	Count of components for the node that are not running
Type	Class of node, such as Master or Data: <ul style="list-style-type: none"> • Data nodes are those running the HDFS datanode service or the YARN NodeManager service • Master nodes are those running the HDFS namenode service or the YARN ResourceManager service • Nodes that are not running any HDFS or YARN services are classified as type other

HDP 1.3

Metric	Description
CPU	Node CPU use
HDFS I/O	Number of bytes read and written for Hadoop
IP Address	IP address of the node
Memory	Percentage of memory use on the node
Name	Name of the node
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
Rack ID	User-defined name that specifies the physical location of the node
Components Down	Count of components for the node that are not running
Type	Class of node, such as Master or Data: <ul style="list-style-type: none"> • Data nodes are those running the HDFS datanode service or the MapReduce tasktracker service • Master nodes are those running the HDFS namenode service or the MapReduce jobtracker service • Nodes that are not running any HDFS or MapReduce services are classified as type other

Details View

The details view displays statistics and information about the row you select in the summary table.

Tabs

Provides important query details on the **Overview** and **Hardware Statistics** tabs.

Node Details

Displays details in sections that are specific to each tab.

< > Previous/Next

Allows you to move through sessions without returning to the summary table.

The screenshot shows the Node Monitor interface. At the top, there's a header bar with 'Node Monitor', a dropdown menu showing 'HDP2.1R5', a clock showing '12:42:58 PM', and a 'Previous/Next' button. Below the header, the selected node is 'tdh035d2.labs.teradata.com'. There are two tabs: 'Overview' (selected) and 'Hardware Statistics'. The main content area is divided into two sections: 'NODE INFO' and 'COMPONENTS INFO'. The 'NODE INFO' section shows details like Type (DATA), Components (16), IP Address (39.64.8.4), Rack ID (-), and Up Since (8/20/15 10:48:32 PM). The 'COMPONENTS INFO' section is a table with two columns: 'COMPONENTS' and 'STATUS'. It lists 16 components, all of which are 'Up'. A scrollbar is visible on the right side of the table. At the bottom of the table, it says '16 rows total'.

COMPONENTS	STATUS
DATANODE	Up
FALCON_CLIENT	Up
GANGLIA_MONITOR	Up
HBASE_CLIENT	Up
HBASE_REGIONSERVER	Up
HCAT	Up
HDFS_CLIENT	Up
HIVE_CLIENT	Up
MAPREDUCE2_CLIENT	Up
NODEMANAGER	Up
OOZIE_CLIENT	Up
PIG	Up
SQOOP	Up
TEZ_CLIENT	Up
YARN_CLIENT	Up
ZOOKEEPER_CLIENT	Up

Node Overview Tab for Aster System

The **Overview** tab displays detailed information about the selected node. The metric values provide a snapshot of Aster system status.

Node Information

Node Info	Description
State	State of the node
Type	Values for node types are as follows: <ul style="list-style-type: none"> • Worker nodes hold the data and process the queries • Queen nodes manage the process, including the performance of the other nodes and delegation of queries • Loader nodes transfer data in and out of virtual workers
Virtual Workers	Number of virtual workers on the node
IP Address	IP address of the node
Rack ID	User-defined name that specifies the physical location of the node
Up Since	Date and time the node came online

Virtual Worker Information

Virtual Worker Info	Description
Index	The partition number corresponding to a virtual worker
Status	Current operating condition of the virtual worker
Type	Values for virtual worker types are as follows: <ul style="list-style-type: none"> • Primary is the virtual worker in which the first copy of data is inserted by Teradata Aster • Replica is the virtual worker in which a copy of the data is inserted by Teradata Aster
Data Size	Amount of data on the virtual worker
Nodes with Replicas	List of nodes that contain a copy of the virtual worker data

Node Overview Tab for Teradata Machine Learning Engine

The **Overview** tab displays detailed information about the selected row. The metric values provide a snapshot of Teradata Machine Learning Engine status.

Pod Information

Pod Info	Description
Type	Values are as follows:

Pod Info	Description
	<ul style="list-style-type: none"> Worker pods hold the data and process the queries Queen pods manage the process, including the performance of the workers and delegation of queries
Virtual Workers	Number of virtual workers on the pod
Pod Name	Name of the pod
Pod Status	Status of the pod

Virtual Worker Information

Virtual Worker Info	Description
Index	The partition number corresponding to a virtual worker
Data Size	Amount of data on the virtual worker

Node Overview Tab for Hadoop System

The **Overview** tab for a Hadoop system displays detailed information about the nodes. The metric values provide a snapshot of Hadoop system status.

Node Information

Node Info	Description
Type (HDP 2.1 and later)	Values for node types are as follows: <ul style="list-style-type: none"> Data nodes are composed of NodeManager processes and a datanode service, which store and retrieve blocks when directed and report back to the master node with lists of blocks they are storing. Master nodes handle the ResourceManager and namenode services, which submit and track applications and determine the location of data.
Type (HDP 1.3)	Values for node types are as follows: <ul style="list-style-type: none"> Data nodes are composed of tasktracker processes and a datanode service, which store and retrieve blocks when directed and report back to the master node with lists of blocks they are storing. Master nodes handle the jobtracker and namenode services, which submit and track MapReduce jobs and determine the location of data.
Components	Hadoop components available on the node
IP Address	IP address of the node
Rack ID	User-defined name that specifies the physical location of the node
Up Since	Date and time the node came online

Components Information

Components Info	Description
Components (HDP 2.1 and later)	<p>You can click one of the following Hadoop components to see details for the component:</p> <ul style="list-style-type: none"> • ResourceManager (master node) • NodeManager (data node) • Namenode (master node) • Datanode (data node) <p>Note: Other installed Hadoop components on the node are shown in this list with their current status.</p>
Components (HDP 1.3)	<p>You can click one of the following Hadoop components to see details for the component:</p> <ul style="list-style-type: none"> • Jobtracker (master node) • Tasktracker (data node) • Namenode (master node) • Datanode (data node) <p>Note: Other installed Hadoop components on the node are shown in this list with their current status.</p>
Status	Current operating condition of the component

Hadoop Components Overview Tab for Datanode

The **Overview** tab for Teradata Hadoop datanode displays detailed information about the component. The metric values provide a snapshot of Teradata Hadoop datanode status. The datanode **Overview** tab is accessed by clicking the DATANODE row under **Components Info**.

Bytes (Current)

Metrics	Description
Written	Total bytes written
Read	Total bytes read

Blocks (Cumulative - Last Hour)

Metrics	Description
Written	Number of blocks written over the last hour

Metrics	Description
Read	Number of blocks read over the last hour
Replicated	Number of blocks replicated over the last hour
Removed	Number of blocks removed over the last hour
Verified	Number of blocks verified over the last hour
Verification Failures	Number of block verification failures over the last hour

Hadoop Components Overview Tab for Tasktracker

The **Overview** tab for Teradata Hadoop tasktracker displays detailed information about MapReduce slots, tasks, and shuffle handler use. The tasktracker **Overview** tab is accessed by clicking the TASKTRACKER row under **Components Info**.

Slots (Current)

Slots	Description
Map Slots Occupied	Number of occupied map slots out of the total map slots available on the tasktracker
Reduce Slots Occupied	Number of occupied reduce slots out of the total reduce slots available on the tasktracker

Shuffle

Shuffle	Description	
Current	Handler Busy %	Percentage of time the handler is busy with the shuffle process
Cumulative - Last Hour	Output Bytes	Number of bytes output from shuffle operations in the last hour
	Failed Outputs	Number of failed shuffle outputs during the last hour
	Successful Outputs	Number of successful shuffle outputs during the last hour
	Exceptions Caught	Number of shuffle exceptions caught during the last hour

Hadoop Components Overview Tab for NodeManager

The **Overview** tab for Teradata Hadoop NodeManager displays detailed information about the current utilization of containers and memory. The NodeManager **Overview** tab is accessed by clicking the NODEMANAGER row under **Components Info**.

Utilization (Current)

Utilization (Current)	Description
Containers Allocated	Number of YARN containers currently allocated
Allocated Memory	Amount of the available memory currently allocated to containers

Hadoop Components Overview Tab for Namenode

The **Overview** tab for Teradata Hadoop namenode component displays detailed HDFS information about HDFS capacity, files and directories, and blocks in the Hadoop cluster. The namenode **Overview** tab is accessed by clicking the NAMENODE row under **Components Info**.

HDFS Capacity

Capacity	Description
Used	Capacity currently used by HDFS
Non-DFS Used	Space on the datanodes that is used by files that are not part of DFS
Remaining	Space on the datanodes that is not currently in use
Total	Total capacity available on the datanodes

HDFS Files and Directories

Files + Directories	Description
Total	The total number of files and directories in HDFS

HDFS Blocks

Blocks	Description
Total	Total number of blocks
Under-replicated	Number of blocks that do not meet their target replication for the file they belong to
Missing	Number of blocks with no replicas anywhere in the cluster

Blocks	Description
Pending Deletion	Number of blocks waiting for deletion
Pending Replication	Number of blocks waiting to be replicated

Hadoop Components Overview Tab for Jobtracker

The **Overview** tab for Teradata Hadoop jobtracker displays detailed information about MapReduce jobs, slots, and tasks. The **Overview** tab is accessed by clicking the JOBTRACKER row under **Components Info**.

Jobs

Jobs	Description	
Current	Running	Number of MapReduce jobs currently executing in the system
	Waiting	Number of MapReduce jobs queued to run
Cumulative (Last Hour)	Submitted	Number of MapReduce jobs submitted in the last hour
	Completed	Number of MapReduce jobs that completed successfully in the last hour

Slots (Current)

Slots	Description
Map Slots Occupied	Number of map slots occupied in the MapReduce cluster
Map Slots Reserved	Number of map slots reserved in the MapReduce cluster
Reduce Slots Occupied	Number of reduce slots occupied in the MapReduce cluster
Reduce Slots Reserved	Number of reduce slots reserved in the MapReduce cluster

Tasks (Cumulative - Last Hour)

Tasks	Description
Map Tasks Launched	Number of map tasks submitted to run in the MapReduce cluster in the last hour
Map Tasks Completed	Number of map tasks that completed successfully in the MapReduce cluster in the last hour
Map Tasks Failed	Number of map tasks that did not complete successfully in the MapReduce cluster in the last hour

Tasks	Description
Reduce Tasks Launched	Number of reduce tasks submitted to run in the MapReduce cluster in the last hour
Reduce Tasks Completed	Number of reduce tasks that completed successfully in the MapReduce cluster in the last hour
Reduce Tasks Failed	Number of reduce tasks did not complete successfully in the MapReduce cluster in the last hour

Hadoop Components Overview Tab for ResourceManager

The **Overview** tab for Teradata Hadoop ResourceManager displays detailed information about YARN applications and NodeManager cluster memory. The Overview tab is accessed by clicking the RESOURCEMANAGER row under **Components Info**.

Applications

Applications	Description	
Current	Running	Number of YARN applications currently executing
	Pending	Number of YARN applications that have not yet received any containers for execution
Cumulative (Last Hour)	Submitted	Number of YARN applications submitted in the last hour
	Completed	Number of YARN applications completed in the last hour

Cluster Memory

Memory	Description
Allocated	Percent of the available memory allocated across all NodeManager instances
Reserved	Percent of the available memory reserved across all NodeManager instances
Total	Available memory across all NodeManager instances

Hadoop Components JVM Tab

The **JVM** tab for Teradata Hadoop master node and data node displays detailed information about heap memory, non-heap memory, and thread counts.

Heap Memory

Heap Memory	Description
Used	Amount of memory currently used to store Java objects
Committed	Amount of allocated memory for Java object storage
Total	Configured maximum size of Java heap

Non-Heap Memory

Non-Heap Memory	Description
Used	Amount of memory currently used by Java to store loaded classes and other metadata
Committed	Amount of allocated memory for Java loaded classes and other metadata storage
Total	Configured maximum size of Java non-heap

Thread Counts

Thread Counts	Description
New	Amount of threads created but not started
Runnable	Amount of threads executing
Blocked	Amount of threads waiting for a monitor lock
Waiting	Amount of threads waiting indefinitely for another thread to perform an action
Timed Waiting	Amount of threads waiting for a specified wait time for another thread to perform an action
Terminated	Amount of threads that can no longer execute

Hardware Statistics Tab

The **Hardware Statistics** tab (or **Statistics**, for Teradata Machine Learning Engine) displays information over a specified time interval. You can view statistics for the last hour, 12 hours, or 24 hours and export the statistics. The metrics in the list beneath the graph are the current values. The plot lines on the graph represent values collected over the selected time frame. Select a metric name to see the plot line highlighted and surrounded by the performance envelope in the graph.

Teradata Aster System

List Metrics	Description
CPU	Node CPU use
Memory	Memory used on the node
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
Disk In	Rate of disk writes in bytes per second
Disk Out	Rate of disk reads in bytes per second

Teradata Machine Learning Engine

List Metrics	Description
CPU	Queen or worker CPU use
Network In	Rate of network traffic into the queen or worker in bytes per second
Network Out	Rate of network traffic out of the queen or worker in bytes per second

Teradata Hadoop System

List Metrics	Description
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages
HDFS I/O	Number of bytes read and written by the datanode Note: HDFS I/O is displayed only on nodes where there is a datanode service.
Memory Used	Memory used on the node
Memory Swap Used	Amount of memory used for swap space
Network In	Rate of network traffic into the node in bytes per second
Network Out	Rate of network traffic out of the node in bytes per second
Disk Space Used	Amount of used disk space on the node

Setting the Time Frame

You can set the time frame that is used to plot the graph.

1. In the toolbar, select a time frame from the list.
The portlet refreshes, and the graph is redrawn.


Removing Metrics from the Graph

You can disable a metric plot line to remove it from the graph and the metric remains in the list.

1. Clear a colored check box in the metric list.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains data for the time frame selected on the **Hardware Statistics** tab.

1. In the toolbar, select a time frame from the list.
2. Click .
3. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.

Node Resources

Node Resources Overview

The **Node Resources** portlet allows you to monitor physical and virtual resources to locate over- or under-utilized resources and isolate performance or system issues. You choose which system and resources to monitor on your Teradata system.

The **Node Resources** portlet provides metrics about the following:

- Percentage of CPU used by nodes or vprocs
- How system resource usage is spread across the vprocs
- How much physical disk I/O, BYNET traffic, or host reads and writes are occurring
- Whether congestion or excessive swapping is an issue on a single or group of nodes or vprocs

These metrics allow you to analyze data and detect issues, such as poor node or vproc parallel efficiency or higher than expected disk read/write ratio or swap I/O rate.

The **Node Resources** portlet provides controls that let you choose what information is displayed.

Summary Views

The summary views provide a status of the resources on your Teradata system so you can monitor and locate issues. The **Nodes** summary view displays the status of nodes, and the **Vprocs** summary view displays the status of vprocs.

The following list describes the features in these views:

Selection Menus

Shows the system and resources currently being displayed, including the system, type of node or vproc, and time frame.

Toolbar

Shows the metric that determines the data that is displayed. Click on the metric to select a different one from the list.

Metric Distribution Graph

Graphically shows distribution of the metric's value per node or vproc.

State Filter Bar

Displays a count of the resources in each state. Click any state in the bar to change the displayed data in the summary table to show only the resources in that state.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information about the selected nodes and vprocs.

Outlying Values

Values that are outside 1.5 times the interquartile range (IQR).

Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Following is an example of the **Vprocs** summary view when you select **All Nodes > Vprocs > All** from the selection menu. The **Nodes** summary view, which appears when you select **All Nodes > Nodes > All** from the selection menu, shows the same features. Only the data types and values are different.



Selecting a System to Monitor

You can use selection menus to choose a system and type of resource to monitor in the portlet, starting with the highest level menu on the far left. Menu choices are based on what you selected in the previous menu.

1. In the selection menu, click the currently selected system name to display a list of available systems.
2. Select a system.
3. Select >.
4. Select a node type.
5. Select >.
6. Select a resource.
7. Select >.
8. Select a resource type.
9. Select >.
10. Select a time frame.

11. Select >.

Selection Menu Choices

The **Node Resources** portlet displays data for a single Teradata system. The selection menu allows you to choose the type of data that appears in the summary views and details views. Depending on your system, you may have access to multi-generational systems.

Selection Menu Choices	Type of Data	Description
Nodes	All	Node data on all nodes
	With Amps	Node data on all nodes with AMPs
	Without Amps	Node data on all nodes without AMPs
Vprocs	All	Vproc data on all AMPs and PEs
	Amps	Vproc data on only AMPs
	PEs	Vproc data on only PEs
6700, 6690, 5650, and so on		Multi-generational system data For example, by selecting 6700, you can view and compare data for 6700 nodes within a system.

Node Resources Metrics

Select metrics to appear in the summary and details views for nodes, vprocs, AMPs, and PEs.

If the processing is skewed towards one or more AMPs (that is, the parallelism is not 100%), the metrics may display a value greater than 100% do to the way the data is normalized.

Some metrics are available only in certain views.

Resource Type	Graph Type	View
Nodes	All	N
	With AMPs	
	Without AMPs	
Vprocs	All	V
	AMPs	
	PEs	

A 1024 byte kilobyte is used as the base for all memory-related and space-related metrics that display, such as spool space, disk space, memory usage, and so on. A 1000 byte kilobyte is used as the base for all other metrics.

Metric	Description	Type	View
AMP Count	Number of AMPs on the node	Number	N
Available AWTs for New Work	Total unused AMP worker tasks available for new work (WorkNew/Work00). The WorkNew (Work00) work type is limited by default to 50 AWTs. If 50 AWTs are already in-use servicing WorkNew message types, there may still be AWTs in the unreserved pool that under this definition will not be considered available.	Number	V
Available Unreserved AWTs	(Teradata Database 16.0 and later) Total unused AMP worker tasks for the entire unreserved pool (all work types). This is the number of AWTs available in the unreserved pool able to be used by all work types in combination, not limited to WorkNew work types.	Number	V
AWTs In Use	Total in-use AMP worker tasks (all work types)	Number	V
Clique Number	The clique number of the node	Number	N
Cluster Number	Cluster to which the AMP is assigned	Number	V
CPU AWT	Percent of CPU usage by AMP worker tasks	Percent	V
CPU I/O Wait	Percent of CPU resources in idle and waiting for I/O completion	Number	N
CPU PE	Percent of CPU usage in PE dispatcher processing	Percent	V
CPU System	Percent of CPU usage spent in user service processing	Percent	N, V
CPU Usage	Percent of CPU usage	Percent	N, V
CPU Use Monitor Virtual Resource	Percent of CPU usage not spent being idle from Monitor Virtual Resource PM/API	Percent	V
CPU User	Percent of CPU usage spent in non-service user code processing	Percent	N
Disk I/O	Number of disk I/Os	Number	N, V
Disk Out Req Avg	Average number of outstanding disk requests	Number	N, V
Disk Reads	Total physical disk reads	Number	N, V
Disk Usage	Percentage of disk usage	Percent	N, V
Disk Writes	Total physical disk writes	Number	N, V
Host Block Reads	Message blocks received from all hosts	Number	N, V
Host Block Writes	Message blocks sent to all hosts	Number	N, V
Host ID	Identifier of the host or LAN associated with the PE	Number	V

Metric	Description	Type	View
Host I/O	Number of host I/Os	Number	N, V
Max I/O Response	Maximum I/O response time for the AMP	Number	V
Mem Allocate KB	Change in vproc memory from the previous sample time frame	Number	N, V
Mem Allocates	Segments allocated to memory resources	Number	V
Message Count	Messages currently queued for delivery to an AMP	Number	V
Message DQ Count	Messages dequeued by an AMP	Number	V
Net A Usage	Total BYNET utilization	Percent	N
Net I/O	Number of messages between the node and the BYNET	Number	N, V
Net Reads	Messages read from the BYNET and input into the node	Number	N, V
Net Writes	Messages written from the node and output to the BYNET	Number	N, V
Node ID	Node identifier (cabinet ID and node ID)	Number	N, V
Node Type	Type of node	Various	N
PE Count	Number of active PEs on the node	Number	N
Session Log Count	Number of sessions currently logged on to the PE	Number	V
Session Run Count	Number of current sessions whose Initiate Requests (TSR messages) are addressed to this vproc	Number	V
Status	Status of the vproc	Character	N, V
Swaps	Number of segments swapped between memory and disk	Number	N
Swap Reads	Segments read into node memory from the disk due to swapping	Number	N
Swap Writes	Segments written to disk from node memory due to swapping	Number	N
Vproc ID	Vproc identifier	Number	V
Vproc Type	Type of vproc	Character	V

Summary Table

The summary table displays statistics and information about the resources at one time, helping you to isolate performance or system issues. This table is useful for comparing multiple vprocs.

You can select, lock, and designate the order of columns from the **Table Actions**  list.

You can click a row to access the details view for a particular node or vproc, depending on the resource you chose from the selection menus.

Details View

The details view displays statistics and information about a particular node or vproc to help you isolate performance or system issues. This view can be accessed by clicking on a row in the summary view.

Node Resources

10:54:04 AM

Dev1

All Nodes

Vprocs

All

Last 1 hour

Vproc ID: 0

GENERAL		I/O	
Type:	AMP	Disk Usage:	1.44%
Status:	Up	Disk I/O:	5.24K
Cluster Number:	0	Disk Reads:	1.02K
Node ID:	1-01	Disk Writes:	4.22K
		Disk Out Req Avg:	0.0
Session Run Count:	0	Max I/O Resp:	1028.0
Available AWT:	50	Net I/O:	52K
AWT in Use:	1	Net Reads:	27.4K
		Net Writes:	24.6K
Message Count:	0		
Message DQ Count:	1.06M		
CPU		MEMORY	
		Mem Allocates:	33.3K
		Mem Allocate KB:	130K
CPU Usage:	0.13%		
CPU System:	0.08%		
CPU AWT:	0.04%		

Performance Data Collection

Performance Data Collection Overview

The **Performance Data Collection** portlet uses jobs to collect performance data from database sources. This data can then be used to generate performance analysis reports.

The following are just a few of the features of **Performance Data Collection**:

- Collect data automatically within specific time frames. For example, Session data is collected every 10 minutes, and Accounting data is collected every hour.
- Create alerts to notify when a job fails or a database threshold has been met.

The **Performance Data Collection** portlet displays data in two views. The summary view shows a summary of all jobs and current information for any jobs that are enabled. This top-level view for the portlet refreshes every 30 seconds. The details view displays detailed information on a selected job, such as status and information about the tables of data included in the job. This view shows completed information and is not refreshed unless the selected job is currently in progress.

Supported versions of Teradata Database are 15.0 and later. Teradata Aster and Teradata Hadoop systems are not supported.

Performance Data Collection Jobs

Data is collected and maintained by jobs implemented in Teradata Viewpoint. The collection start times are configurable.

Category	Job	Default Collection Times
Every 10 Minutes Job	Session	On the "0" minute, for example, 1:00, 1:10, 1:20, and so on.
Hourly Job	Accounting	On the "00" minute, for example, 1:00, 2:00, 3:00, and so on.
Daily Jobs	Info Tables	12:05 a.m. daily
	Disk Space	
	Logon/Logoff	
	Resource Usage	
	TDWM	
	Access Log	
	DBQL	
	Statistics	

Category	Job	Default Collection Times
Weekly Job	Maintenance	4:15 a.m. every Sunday

Performance Data Collection View

This summary view displays summary information on the **Performance Data Collection** jobs.

System Selection

Displays the name of the system currently selected and allows you to select a different system from the portlet frame.

Toolbar

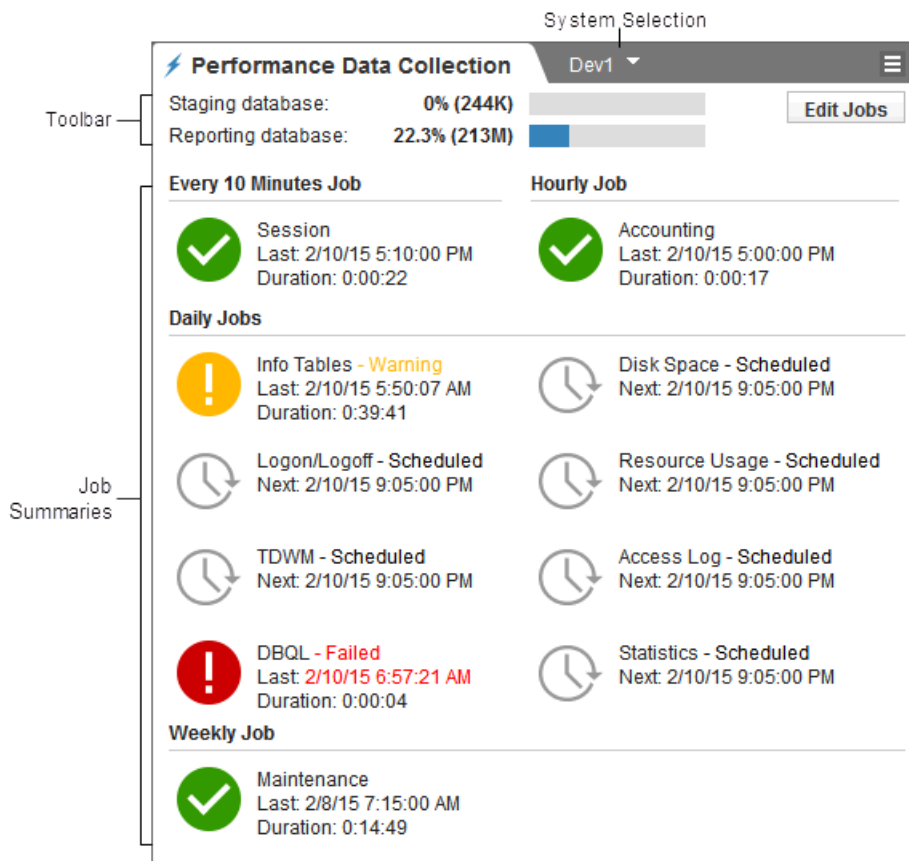
Staging Database: Shows the peak perm space used in the PDCRSTG database, using a 1024 byte kilobyte as the base.

Reporting Database: Shows the current perm space used in the PDCRDATA database, using a 1024 byte kilobyte as the base.

Edit Jobs: Opens the **Edit Jobs** view, where you can edit job settings, retention periods, and alerts.

Job Summaries

Shows information about the jobs, such as previous and next scheduled start times, status, and duration of the last run.



The status of each job is shown by colored text and icons.

Icon	Status	Definition
	In Progress	The job is currently running. The start time of the job is displayed.
	Completed	The job has successfully completed. The start time of the job and its duration are displayed.
	Warning	The job has run but not collected any data. <i>Warning</i> and the next time the job is scheduled to run are displayed.
	Failed	The job has failed. <i>Failed</i> , the last start time, and the duration of the job are displayed.
	Disabled	The job is disabled and will not run until it is enabled. When the portlet is first opened, all of the jobs, except Weekly , display this icon. When jobs are enabled, all of the icons change to the scheduled state icon.
	Scheduled	The job is scheduled to run at the specified time. <i>Scheduled</i> and the next time the job is scheduled to run are displayed. This is the state displayed the first time a job is run for a system.

Performance Data Collection Details View

The details view displays information about the selected job. This view can be accessed by clicking a job in the summary view or in the history view.

Toolbar

- Job Type
- Job Category
- ▼ **Actions**
 - **Run Job** allows you to manually start a job.
 - **View Job History** provides a view of previous runs of this job.

Job Information

- **Status** - The status of the job.
- **Start time** - The time the job started running.
- **Stop time** - The time that the job stopped running.
- **Duration** - The amount of time the job took to run.
- **Total row count** - The number of rows collected in the reporting database across all tables as a part of this job.

Filters

Filters allow you to change displayed data by showing only rows that match your filter criteria.

▼ Table Actions

- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display.
- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Summary Table

Displays information about each table loaded in the job.

Tabs

The **Table Details** tab displays information about each table loaded in the job, including its start time, duration, status, and row count.

The **SQL Log** tab displays the list of SQL statements that were executed for that job. You can click a row to display details about the SQL statement execution.

Performance Data Collection

Dev1

Session Job | Every 10 Minutes

Actions

Status: Completed

Start time: 9/20/16 5:10:00 PM

Stop time: 9/20/16 5:12:30 PM

Duration: 0:02:30

Total row count: 86

Table Details

SQL Log

TABLE

START

DURATION

STATUS

TDWMThrottleStats_

9/20/16 5:10:01 PM

0:00:07

Completed

TDWMUtilityStats_Hs

9/20/16 5:10:09 PM

0:00:08

Completed

MonitorSession_Hst

9/20/16 5:10:18 PM

0:02:10

Completed

3 rows total

SQL Log Metrics

The metrics on the SQL Log tab are listed as following.

Note:

SQL log data is not collected for job executions in Viewpoint versions earlier than 16.00.

Metric	Description	Type
Start	Time that the query started running	Number
Duration	How long the query ran	Number
SQL	The actual SQL statement	String
Parameters	Parameters passed to the query	String
Error Code	SQL error code, if an error was generated during execution of the query	String
Error Message	SQL exception message, if an error was generated during execution of the query	String

SQL Log Details

Click any of the item from the **SQL Log** tab to open its detailed information. Following are the metrics appear for the selected SQL statement:

Metric	Description
Start	When the SQL statement started running
Duration	How long the SQL statement ran
SQL	The actual SQL statement
Parameters Note: [Viewpoint V17.10.02.00 and Later] Parameters changes to Parameters (Output Parameter) for SQL statement of type Procedure	Parameters passed to the SQL statement
Error Message	SQL exception message, if an error generates during execution of the query

Table Metrics

Metric	Description	Type
Table	Name of the table	String
Start	Time that the query started running	Number
Duration	How long the query ran	Number
Status	Status of the query execution	String
Row Count	The number of rows affected, if the table had data in it	Number

Viewing the Details of a Job

You can view the details of jobs from both the summary view and the history view.

1. On the summary view or history view, click a job to access the details view for that job.

Note:

If a job is in progress, the details view refreshes every 30 seconds.

Running a Job

A job can be run manually even if it is currently scheduled or disabled, unless it is already in progress.

1. In the details view of a job, click **Actions > Run Job**.
The job starts to run and the job details view is refreshed.

Note:

If the job cannot run immediately, its status displays as *Queued* and the next run time is not displayed.

Viewing a Job History

You can view the information from previous runs of a job on the **Job History** view.

1. Click a job in the summary view.
2. In the details view that appears, click **Actions > View Job History**.
The run history for that job appears, showing its **Start Time**, **Duration**, **Status**, **Row Count**, and **Stop Time** for each run.
3. To see a particular time frame of the history, click the time frame list in this view and select one of the options.

The available time frames depend on the type of job.

Job Type	History Time Frames
Session Job History	<ul style="list-style-type: none"> • Last hour • Last 12 hours • Last 24 hours • Last 3 days • Last 7 days • Last 30 days
Accounting Job History	<ul style="list-style-type: none"> • Last 12 hours • Last 24 hours • Last 3 days • Last 7 days • Last 30 days
Daily Jobs History	<ul style="list-style-type: none"> • Last 7 days • Last 14 days • Last 30 days • Last 90 days
Weekly Job History	<ul style="list-style-type: none"> • Last month • Last 3 months • Last 6 months • Last 12 months

The information for the time frame appears in the summary table of the view.

Edit Jobs View

The **Edit Jobs** view allows you to enable jobs and set the start time for jobs. You access this view by clicking **Edit Jobs** in the summary view.

The **Edit Jobs** view has three tabs that allow you to edit all aspects of a job:

- **Jobs:** Shows the start times set for each job type (hourly, daily, weekly) and the time zone used to interpret the times specified in the portlet, such as job start times.
- **Retention Periods:** Shows the number of days data is retained before being purged for each job type and table within a job type.
- **Alerts:** Shows created alerts in a summary table. This tab also has an **Actions** list, which contains options for adding job alerts and space alerts.

Jobs Tab

The **Jobs** tab on the **Edit Jobs** view displays options you can set for jobs and a button to access the **Credentials** dialog box. If you want to run all of the jobs, select the **Enable jobs** check box.

Your selections on this tab include changing the time zone, selecting job categories, and setting job start times.

Configuring Jobs

1. In the summary view, click **Edit Jobs**.
2. In the **Edit Jobs** dialog box, select the **Enable jobs** check box.

Note:

Selecting this check box enables all of the jobs. If you want only certain jobs enabled, clear the check boxes of those jobs.

3. [Optional] You can change the time zone for all enabled jobs if needed.
4. Do the following to configure the jobs that you have enabled.

Option	Action
Every 10 Minutes Job	Type a value, from 0 to 9, to set the minute mark for each 10 minutes. For example, if you set the value to 3, and the current time is 10:00 a.m., the job will run at 10:03, 10:13, and so on.
Hourly Job	Type a value, from 0 to 59, to set the start time for each hour. For example, if you set the value to 15, and the current time is 10:00 a.m., the job will run at 10:15, 11:15, and so on.

Option	Action
Daily Jobs	Type a value to set the start time for the daily jobs. Select + or - buttons to change the default value, add a new schedule, or remove an existing schedule as required. The daily jobs will be triggered multiple times in a day as per the scheduled start times.
Weekly Job	Type a value for a start time and select a day from the list, or leave the defaults in place.



5. [Optional] Click **Restore Defaults** to return the values to their original states.
6. Click **Save**.

Credentials

Credentials are commonly a user name and a password. The **Performance Data Collection** portlet validates credentials to confirm that users have the right, or permissions, to execute jobs against the Analytics Database. For example, if you want to set retention periods or alerts for jobs, you need to have credentials that the system considers valid.

Logging In with a Valid PDCRAdmin Password

To run the **Performance Data Collection** jobs, you must log in with a valid PDCRAdmin password.

1. On the **Edit Jobs** view, click **Credentials**.
The **Credentials** dialog box appears, with **Username** already populated.
2. Enter the password for the PDCRAdmin user.
If you have previously saved a password, this field is already filled with placeholder text.
3. Click **Test** to authenticate your password.
If the operation is successful,  appears. If the operation fails,  appears.
4. [Optional] Enter an account string.
5. [Optional] Select an authentication mechanism from the list.
6. Click **Save**.

Retention Periods Tab

The **Retention Periods** tab on the **Edit Jobs** view displays retention information for all jobs and the tables associated with those jobs. A retention period is the number of days a table retains its data before that data is purged.

The **Weekly Job** purges data that exceeds the retention limits set for each table. Data is purged from a table on a First-In-First-Out basis.

The **Cleanup Schedule** in the **Monitored Systems** portlet allows you to schedule table cleanup to occur after **Performance Data Collection** daily jobs have completed successfully. This makes sure the data is archived to the **Performance Data Collection** database before the cleanup occurs.

Configuring Retention Periods

You can change the retention period for any job or table within a job by using the **Retention Periods** tab in the **Edit Jobs** view.

1. Click in the box of the job type and applicable table that you want to configure.
2. Type the number of days you want for the retention period.
Numbers 1 through 9999 are valid.
3. [Optional] Click **Restore Defaults** to return the values to their original states.
4. Click **Save**.

Alerts Tab

Alerts are used to notify interested parties when a job fails to complete successfully or when the space for a database reaches a specific threshold.

The **Alerts** tab in the **Edit Jobs** view displays a summary of existing alerts, including their enabled status and their type: job or space. In this tab, you can edit, delete, and add alerts.

Alerts created in the **Performance Data Collection** portlet can be displayed in the **Alert Viewer** portlet. This portlet displays alert type, severity, source, duration, error messages, and the action taken when the alert is triggered. The option to view **Performance Data Collection** alerts in the **Alert Viewer** portlet is set in the **Alert Setup** portlet.

Adding a Job Alert

1. In the summary view, click **Edit Jobs**.
2. In the **Alerts** tab, click **Actions > Add Job Alert**.
3. In the **Add Job Alert** dialog box, type a name for the alert.
Names must be unique, and characters are case-sensitive.
4. Do one of the following when adding an alert:
 - If you want the alert to be enabled immediately, select the **Enable alert** check box.
 - If you want to wait before enabling the alert, make sure the **Enable alert** check box is cleared.
The jobs check boxes remain selected, but the alert will not be triggered until the **Enable alert** check box is selected.
5. Clear the check box of any job you do not want this alert to act upon.
6. Select a severity for the alert from the list.
7. Under **Alert Action**, select an alert action from the list.

Alert actions specify the action the system takes in the event of an alert, for example, send an email or run a program. Alert actions are created in the **Alert Setup** portlet.


8. [Optional] Type a limit for **Do not run twice in** a number of minutes.
The alert action does not run twice during the time frame that you enter. Numbers between 0 and 1440 are valid.
9. [Optional] Under **Message**, type a message that appears when the alert action runs.
10. Click **Save**.

Adding a Space Alert


You can create an alert to notify interested parties if a database space threshold is met.

1. In the main view, click **Edit Jobs**.
2. In the **Alerts** tab, click **Actions > Add Space Alert**.
3. In the **Add Space Alert** dialog box, type a name for the alert.
Names must be unique, and the characters are case-sensitive.
4. [Optional] Select the **Enable alert** check box to enable the alert.
5. Select a severity for the alert from the list.
6. Under **Alert Rules**, select a database from the list.
7. In the box to the right of the selected database, type a threshold percentage.
The number must be between 0 and 100. If the database exceeds this threshold, an alert is triggered.
8. Under **Alert Action**, select an alert action from the list.
Alert actions specify the action the system should take in the event of an alert, for example, send an email or run a program. Alert actions are created in the **Alert Setup** portlet.
9. [Optional] Type a limit for **Do not run twice in** a number of minutes.
The alert action does not run twice during the time frame that you enter. Only numbers between 0 and 1440 are allowed.
10. [Optional] Under **Message**, type a message that appears when the alert action executes.
11. Click **Save**.

Editing a Job Alert


1. In the **Edit Jobs** view, click the **Alerts** tab.
2. Click  in the row of the job alert you want to edit.
3. Make your changes in the **Edit Job Alert** dialog box.
4. Click **Save**.

Editing a Space Alert

1. In the main view, click **Edit Jobs**.
2. In the **Alerts** tab, click  in the row of the space alert you want to edit.

3. Make your changes in the **Edit Space Alert** dialog box.
4. Click **Save**.

Deleting an Alert

1. In the main view, click **Edit Jobs**.
2. In the **Alerts** tab, click  in the row of the alert you want to delete.
A confirmation message appears.
3. Select **Delete**.

Performance Data Collection Alert Metrics

These metrics are available for Teradata **Performance Data Collection** alerts. The associated property names allow you to customize alert actions in the **Alert Setup** portlet or customize the message in the **Performance Data Collection** portlet.

Database Space Alert Types Metrics

Metric	Description	Property Name
Database Name	Name of the database	<i>databaseName</i>
Current Perm	Percentage of the total permanent disk space the database is currently using	<i>currentPerm</i>
Peak Perm	Highest percentage of the total permanent disk space the database has used since the last time the statistics were set by the Database Administrator	<i>peakPerm</i>

Job Alert Types Metrics

Metric	Description	Property Name
Job Type	Type of job	<i>jobType</i>
Job Status	Status of a job	<i>jobStatus</i>
Error	Error message displayed in Alert Viewer when an alert is triggered	<i>errorMessage</i>
Duration	Amount of time the job took to run (in milliseconds)	<i>duration</i>
Count	Number of rows collected in the reporting database across all tables as a part of a job	<i>rowCount</i>

Productivity

Productivity Overview

The **Productivity** portlet allows you to monitor system health and query performance trends of a Teradata system. You can analyze performance trends over a 24-hour or a 48-hour period. You can also compare current performance with average performance measured across one or more weeks of historical data. Additionally, you can monitor and compare the productivity of multiple systems by adding a **Productivity** portlet to the view for each Teradata system to be monitored.

The Teradata Viewpoint Administrator enables system health metrics, configures the degraded and critical thresholds, and customizes the states for the **System Health** portlet in the **Monitored Systems** portlet. These customized states also appear in the **Productivity** portlet. The metrics and thresholds are usually selected to highlight an unusual load on the system that has the potential to impact overall performance.

Productivity View

The **Productivity** view displays a high-level overview of the performance trends for a Teradata system during the preceding 24 hours or 48 hours. The **Export** button allows you to create a .csv file containing selected data.

Data trends are represented in a horizontal graph called a *sparkline*. Hover over a sparkline to see an information balloon containing detailed information about the data point. The sparkline types are:

Canary Response Times

Canary-response-time metrics show the time (in milliseconds) a query takes to complete. Canary response times, such as **System Heartbeat** are represented by a *curve* sparkline. Every 15 minutes, a new data point is added to the sparkline, representing the average canary response time recorded during the interval. A large dot and a number at the end of the sparkline indicate the last data point captured.

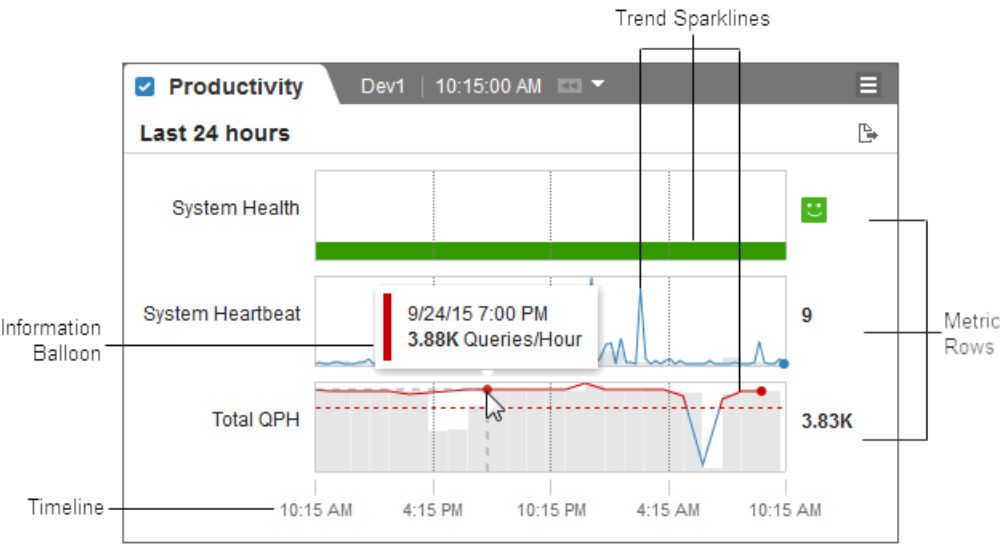
Queries Per Hour

Queries-per-hour (QPH) metrics show the number of queries against the database for each hour over the 24-hour or 48-hour period for the selected metric. For example, the **TOTAL QPH** sparkline is a *skyline* sparkline, used to represent QPH metrics. A flat, solid data point represents the number of queries per hour on the timeline.

System Health

The **System Health** metric provides an overview of the worst state of the monitored system, sampled at 15-minute intervals during the 24-hour or 48-hour period. **System Health** is represented by a *bar* sparkline and uses default or custom states and tooltips. The values on the sparkline represent the worst system state experienced during each 15-minute period.

A 24-hour period is shown in the minimized view and a 48-hour period is shown in the maximized view. You can change these views from:


















To display data in the **Productivity** view, the Teradata Viewpoint Administrator must first enable the Query Count, System Stats, and Canary Queries collectors. The Teradata Viewpoint Administrator also specifies the data collection rate for each collector.

Factors external to Teradata Viewpoint can also affect the collection and display of data. The Teradata System Administrator must enable query logging on the Teradata system.

Productivity Metrics


You can select and organize metrics to display, set thresholds, adjust the vertical axis range for each metric, and adjust the duration over which the averages are calculated using the **Settings** view.

Metric	Default Sparkline	Past Average Line	Default	Description												
Canary Response Times Metrics	Curve	Yes	No	Select individual canary query metrics												
System Metrics: System Health	Bar	No	Yes	Shows customized states or these default states: healthy, warning, critical, down, or unknown. The icon style cannot be changed.												
				<table><tr><th>Default Icon</th><th>Text Color</th><th>Default State</th><th>Definition</th></tr><tr><td></td><td>Green</td><td>Healthy</td><td>All metrics are within <i>healthy</i> ranges.</td></tr><tr><td></td><td>Yellow</td><td>Degraded</td><td>At least one metric exceeded a <i>degraded</i> threshold.</td></tr></table>	Default Icon	Text Color	Default State	Definition		Green	Healthy	All metrics are within <i>healthy</i> ranges.		Yellow	Degraded	At least one metric exceeded a <i>degraded</i> threshold.
				Default Icon	Text Color	Default State	Definition									
	Green	Healthy	All metrics are within <i>healthy</i> ranges.													
	Yellow	Degraded	At least one metric exceeded a <i>degraded</i> threshold.													

Metric	Default Sparkline	Past Average Line	Default	Description																
				<table><tr><th>Default Icon</th><th>Text Color</th><th>Default State</th><th>Definition</th></tr><tr><td></td><td>Red</td><td>Critical</td><td>At least one metric exceeded a <i>critical</i> threshold.</td></tr><tr><td></td><td>Black</td><td>Down</td><td>The selected system is <i>down</i>.</td></tr><tr><td></td><td>Gray</td><td>Unknown</td><td>Status of the selected system is <i>unknown</i>.</td></tr></table>	Default Icon	Text Color	Default State	Definition		Red	Critical	At least one metric exceeded a <i>critical</i> threshold.		Black	Down	The selected system is <i>down</i> .		Gray	Unknown	Status of the selected system is <i>unknown</i> .
Default Icon	Text Color	Default State	Definition																	
	Red	Critical	At least one metric exceeded a <i>critical</i> threshold.																	
	Black	Down	The selected system is <i>down</i> .																	
	Gray	Unknown	Status of the selected system is <i>unknown</i> .																	
System Metrics: System Heartbeat	Curve	Yes	Yes	Shows whether the Teradata system is responsive																
System Metrics: Total QPH	Skyline	Yes	Yes	Select total queries per hour metrics																
Queries Per Hour Metrics	Skyline	Yes	No	Select individual queries per hour metrics																

Exporting Metrics

You can export data to a .csv file for further analysis and formatting.

1. Select .
2. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
3. [Optional] Reformat exported data to match the format in the view.

Settings View

The **Settings** view provides the following tabs that allow you to customize the portlet to monitor key metric trends that are important to efficient operation of the system.

Metrics


Select a system and the metrics to monitor, and define the threshold and vertical axis range for the selected metrics. Values that exceed the defined thresholds are highlighted in red.

Past Averages

Specify the time period used to calculate and display the shaded skyline graph behind each metric line graph.

Managing Default Settings





In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.


Setting Metrics Display

1. Click  in the portlet frame and select **Settings**.
2. Click the **Metrics** tab.
3. Select a metric from the list.
You can add  or remove  metrics.
4. [Optional] Enter a **Threshold** value.
5. [Optional] Enter a **Vertical Axis Range** value.
6. [Optional] Click  and drag the row to reorder the metrics for display.
7. Click **OK**.

Setting Past Averages

Specify the time frame used to calculate and display the shaded skyline graph in the details view.

Each segment of the skyline graph is the average of data samples taken each week at the same time and day of the week.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Past Averages** tab.
3. Enter the number of weeks up to 99.
4. Click **OK**.

Query Groups

Query Groups Overview

The **Query Groups** portlet allows you to view information about session queries running on a Teradata system or in Teradata QueryGrid that match defined criteria.

A *query group* is any combination of one or more Analytics Database users, account strings, query bands, and workloads. For example, a query group that contains a group of users, could be configured for you to view queries based on the country of the user.

Queries appear in a group because they match certain criteria or characteristics. You can only view groups that have been assigned to you. If you identify a problem query, you can correct the problem by changing the priority or workload, releasing the query, or aborting the query or session. Take these actions for one query or session, or multiple queries or sessions at a time.

The summary view contains a table with one row allocated to each group running on the database. It contains all the information matching the criteria set for that group.

The details view contains session and query information.

The **Query Groups** portlet allows you to filter queries by group for all sessions. You can set thresholds for certain columns, and when the threshold is exceeded, the information is highlighted in the summary table.

The **Settings** view allows you to change display options.

You can monitor multiple Teradata systems by opening more than one instance of the **Query Groups** portlet.

Note:

Support for TASM ARM is available for SQL Engine versions 17.10 and later.

Query Groups View

The **Query Groups** view displays a summary of queries running on the selected system.

Selection Menus

Shows the system and query group currently being displayed. Choose a system and a query group.

State Filter Bar

Displays a count of the sessions in each of the states. Click any state in the bar to change the displayed data in the summary table to only show the sessions running a query in the selected state.

Filters

Filters allow you to change displayed data by showing only rows that match your filter criteria. You can also sort on the column headers to find the details you want.

Summary Table

Displays information about each session for the group of queries, with columns configured specifically for the current view. The view is refreshed when new data is collected. Click on a row in the table to see details.

Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Abort, **Change Priority**, **Change Workload**, and **Release Query** allow you to manage the selected queries and sessions.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information on **Clear Filters**, **Configure Columns**, and **Export** table actions, see [Summary Table Controls](#).

Overflow Menu

Displays a list of states. You can select another state to replace the current state.

The screenshot shows the 'Query Groups' interface. At the top, there's a header with 'Query Groups', a time '4:14:04 PM', and a menu icon. Below the header, there's a 'Selection' bar with 'Dev1' and 'TestGroup'. Under 'Selection', there's a 'State Filter' bar with buttons for '70 All', '35 Not Idle', '27 Active', '3 Block', '2 Delay', and '0 Abort'. Below the state filter, there's a 'Filters' section with input fields for 'SESSION ID', 'STATE', 'ΔCPU', 'CPU SKEW', and 'DURATION'. The main part of the interface is the 'Summary Table' which displays a list of sessions with columns: SESSION ID, STATE, ΔCPU, CPU SKEW, and DURATION. The table shows 10 rows of data. At the bottom of the table, it says '70 rows total'. To the right of the table, there's a 'Table Actions' and 'Overflow Menu' icon.

SESSION ID	STATE	ΔCPU	CPU SKEW	DURATION
2740331	ACTIVE	58.22	7.751	0:05:00
2740327	ACTIVE	55.71	25.4	0:01:00
2740322	ACTIVE	17.3	58.89	0:02:00
2740409	IDLE	14		0:00:00
2740317	ACTIVE	8.211	61.11	0:01:00
2740313	ACTIVE	3.531	62.58	0:05:00
2740335	ACTIVE	3.082	94.44	0:04:00
2740398	ACTIVE	2.781	61.11	0:01:00
2740417	ACTIVE	2.766	60.42	0:01:00
2740416	ACTIVE	2.424	64.81	0:03:00

70 rows total

Selecting a System to Monitor

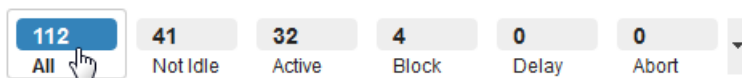
You can use selection menus to choose a system and a query group to monitor in the portlet, starting with the highest level menu on the far left. Menu choices are based on what you selected in the previous menu.

You are not required to make a selection from every menu each time you modify a selection.

1. In the selection menu, click the currently selected system name to display a list of available systems.
2. Select a system.
3. Select >.
4. Select a group from the menu.
5. Select >.

State Filter Bar

The state filter bar allows you to display specific states in the view. If more than six states are available in a Teradata system, you can select another state from the **Overflow Menu** to replace the state in the last position.



The state filter buttons provide a count of sessions for each state category. Click any button to filter for the selected state or select a state from the list. For example, click **Delay** to display the sessions running a query that is queued and waiting to run in the selected system. Hover over a state filter button to display the expanded content for large, abbreviated numbers or truncated state names.

All

Sessions currently running in the selected Teradata system

Not Idle

Sessions in any state except idle

Active

Sessions running a query that is in progress

Block

Sessions running a query that is waiting for a locked resource, such as a database table or view

Defer

Sessions running a query that is in deferred queue waiting to run.

Delay

Sessions running a query that is queued waiting to run

Abort

Sessions running a query that is in the process of aborting (rolling back changes made by the query)

Resp

Sessions running a query that has completed and is sending (responding) spooled data back to the user

Idle

Sessions not currently running a query

Parse

Sessions running a query that is being parsed. It has not begun to run.

Other

Sessions whose status is unknown and do not fall into any of the previous categories

QTDelayed

Sessions delayed due to a queue table restriction

SesDelayed

Utility sessions that are in the workload delay queue







Response-Held

In SLES 11, sessions running a query whose response is being held until a specified number of seconds elapses

State Icons

The following icons appear in the portlet.

Icons	Name	Definition
×	Aborting	Query has been aborted, and changes are being rolled back.
»»	Active	Query is running.
»I	Blocked	Query is waiting for a lock held by another query.
III	Delayed	Query is in a delay queue waiting to run.
⌘	Deferred	Query is in a deferred queue waiting to run.
🔄	Host-Restart	Host associated with this session is currently restarting.

Icons	Name	Definition
	Idle	No query is running.
	Other	Query is in an unknown state.
	Parsing	Query is being parsed before running.
	QTDelayed	Query delayed due to a queue table restriction.
	Response	Query is returning results to the user.
	SESDelayed	Query is delayed because a utility limit has been exceeded in FastLoad, MultiLoad, FastExport, or ARC.

Query Group Metrics

Metric	Description	Type
Account	Account from which a query was submitted	String
Blocked Time	How long the query has been blocked	Number
CPU Use	Percent of available CPU seconds on the system used during the last sampling period	Percent
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number
Duration	How long the query has been running	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
In State	How long the query has been in the current state	Number
Partition	Partition in which the query is running	String
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(For Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String

Metric	Description	Type
Query Band	Entire query band string. Query bands are a set of name-value pairs defined by the user to tag sessions or transactions with an ID through a SQL interface.	String
Request CPU	CPU seconds needed to run the query	Number
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Request I/O	Disk I/Os performed to run the query	Number
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Request Count	Number of queries the session has executed	Number
Session ID	Unique session identifier	Number
Snapshot CPU Skew	CPU skew during the last sample	Percent
Snapshot I/O Skew	I/O skew during the last sample	Percent
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.	Percent
Start	Time that the query started running on Analytics Database	Number
State Icon	Icon representing the current state of the query	Icon
State	Text describing the current state of the query	String
Temp Space	Temp space used by the query, using a 1024 byte kilobyte as the base	Number
Unnecessary I/O	Ratio of I/O to CPU for the current query	Number
User	Name of the user who submitted the query	String
Workload	Workload in which the query is running	String

Details View for All Systems

The details view displays statistics and information about the selected session. This view can be accessed by clicking a session row in the summary table.

Tabs

Provides important query details on the **Overview**, **SQL**, **Explain**, **Skew**, **Blocked By**, **Query Band**, **Defer**, and **Delay** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

▼ Manage Queries

Manages the query or session for Teradata systems with **Abort** , **Change Priority** , **Change Workload** , and **Release Query** options.

<> Previous/Next

Allows you to move through sessions without returning to the summary table.

System: Prod1 | Session: 2824421

Manage Queries Previous/Next

Tabs: Overview SQL Explain Skew Query Band Blocked By

Query Details

QUERY INFO

State: »» Active
Time in State: 0:01:00
Total Duration: 0:01:00

Spool Space: 159,232
Temp Space: 0

Request CPU: 0
Request I/O: -

PJI: 0
Unnecessary I/O: 2.73

WORKLOAD INFO

Name: WD_Mode...
Classification Mode: Auto

SESSION INFO

User: VIEWPOINT
Account: DBC
Partition: DBC/SQL
Requests: 11

Source:
(TCP/IP) b308 153.64.107.218 SKATE
;SKATECOP2/153.64.202.243:1025 C
ID=79EF3A1A ROOT JDBC15.00.00.0
9;1.7.0_45 01 LSS

SNAPSHOT INFO

CPU Use: 0%
Impact CPU: 5.62
Snapshot CPU Skew: 18.2%
Snapshot I/O Skew: 17.9%

Overview Tab

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using

Query Information	Description
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11


Session Information

Session Information	Description
User	Name of the user that submitted the query
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system


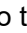
SQL Tab


The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.





Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Explain Tab


The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status. If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

If you have the correct permissions, use  to retrieve the latest Explain steps from the database. Click **Refresh** to update. The refresh screen data is not maintained once the refresh view closes. Repeatedly retrieving current Explain data can impact system performance.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	 - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step

Step Information	Description
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Exporting Explain Data

1. From the Explain tab, click .
2. Save the file to a location.
The Explain data is exported as a .csv file.

Skew Tab

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Blocked By Tab

The **Blocked By** tab displays details about other queries that are blocking the selected query. This information is read-only. The tab is available only when the selected query is blocked.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock 		Available
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	

Blocked By Information	Description	14.10 and earlier	15.00 and later
Locked	Name of the locked object	Available	

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Defer Tab

The Defer tab displays details about the rules that are deferring a query. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Defer Time	Duration of time request has been deferred by Arrival Rate Meter rules
User Name	User who submitted the query

Rule Information	Description
Rule Name	Name of rule causing request to defer
Rule Type	TASM type of rule causing request to defer
Overridable	Indicates if the Teradata DBA can abort or release the query

Delay Tab

The **Delay** tab displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Details View for Teradata QueryGrid

When you select a Teradata QueryGrid session from an Analytics Database session or a Teradata QueryGrid manager in the portlet's summary view, the details for the query are displayed.

The top portion of the screen displays a graphical representation of the progress and status of the query, with icons representing the initiating system, the phase of the current operation being performed, and the target system. The initiating system is the data source from which the query is generated. The target system is the system against which the query is being run. Teradata QueryGrid queries go through a metadata and execution phase before they are completed. Teradata QueryGrid queries may also access data on target systems or transfer data from one system to another.

The lower part of the screen displays tabs with information about the system or operation. If you arrive at the details screen from a Teradata system session, details are displayed for the target or initiator system that is relevant for the selected step. If you arrive at the details screen while viewing a list of Teradata QueryGrid manager sessions, details for the selected operation are displayed. You can click on any of the icons in the graphical representation (initiator, operator, or target) to view the metrics for the selected system or operation.

Initiating System

The data source from which the query is generated.

Target System

The system from which data is accessed or to which it is transferred.

Operation

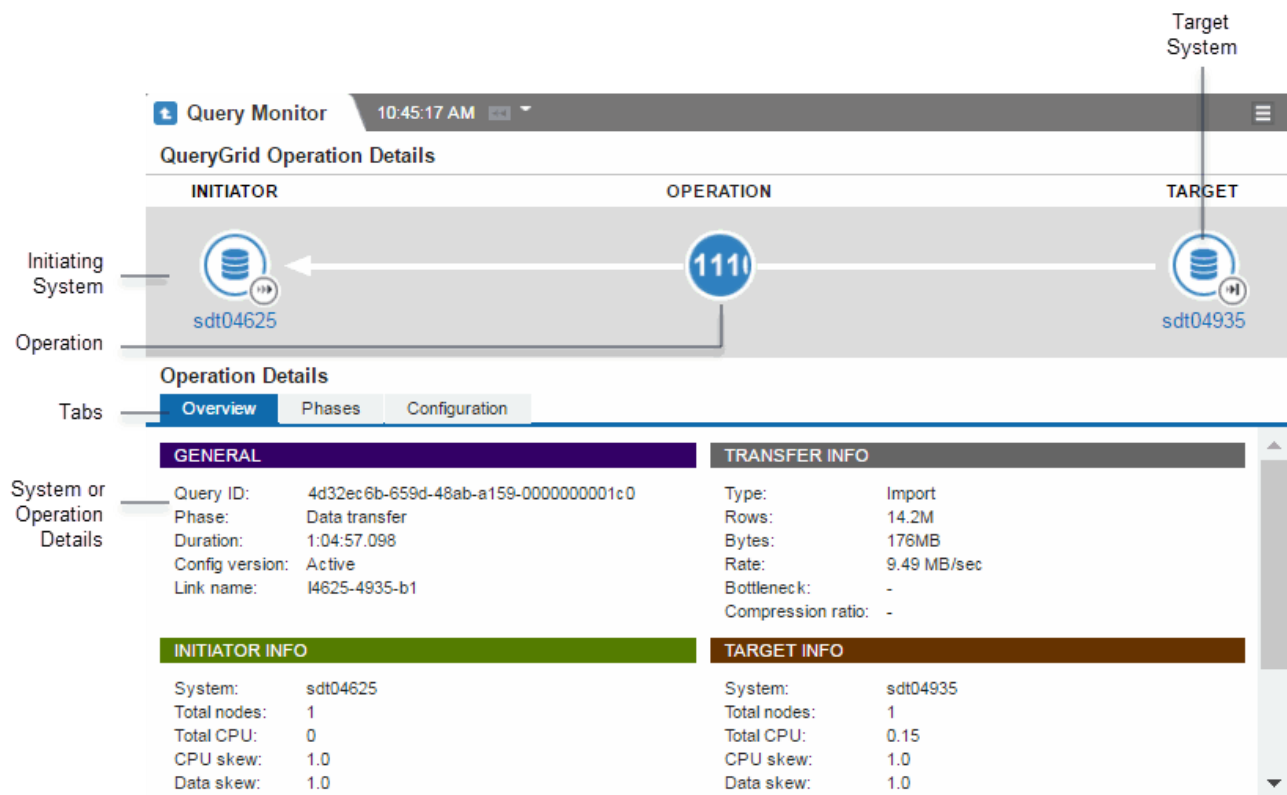
The phase of the operation being performed.

Tabs

Related metrics are organized into tabs for the system or operation selected in the top portion of the view.

System or Operation Details




The metrics related to the system or operation selected in the top portion of the view.

**QueryGrid Icons**







The details view for Teradata QueryGrid includes a graphical representation of the progress of the selected query. The icons here represent the initiating or target systems and the phases of the operations performed.

System Icons

The icons in the following table represent the query on the initiator and target systems. The icons inside the smaller circle represents the query state. In addition to the query states shown as following, additional possible states for the query are listed in [State Icons](#).

Icon	Name	Definition
	No query is running	No query is currently running
	Active	The query is running on the system
	Idle	The query is currently idle

Operation Phase Icons

Icon	Name	Definition
	Metadata Operation	A handshake between systems to negotiate data types and validate queries.
	Execution Operation	Query processing prior to transferring the data.
	Data transfer	The transferring of data from one system to another.
	Completed	Query has successfully finished processing.
	Failed	Query did not finish processing successfully. [QueryGrid systems ≥ 2.10] To download the support bundle, click Download Support Query Bundle .
	Unknown	Query processing status is not currently available.

Initiator and Target Metrics

When an initiator or target system icon is selected in the details view for Teradata QueryGrid, two tabs display metrics for the system: an **Overview** tab and a **SQL** tab.

Overview Tab for Teradata System

The **Overview Tab** displays information for a QueryGrid query that was initiated by or is a target of a Teradata system.

QueryGrid Information

Metric	Description	Type
System	Name of the initiating or target system	String
Connector	Point from which the query originates or ends	String
Software version	The version of the connector	String
Network	The logical network used for this side of the data transfer	String

Session Information

Metric	Description	Type
User	User who submitted the query	String
Query ID	ID of the query	String
Session Number	Number of the session	String
Request Number	Number of the request	String
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session	String


Query Information

Metric	Description	Type
State	State of the query	String
Duration	Amount of time the query has been running (HH:MM:SS)	Number
Workload	Workload in which the query is running	String
Request CPU	Total request CPU (seconds)	Number
Request I/O	Total request I/O count	Number

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Operations Metrics

When an operation phase icon is selected in the details view for Teradata QueryGrid, three tabs display metrics for the operation: an **Overview** tab, a **Phases** tab, and a **Configuration** tab.

Overview Tab

The **Overview** tab displays general information about the selected operation for a Teradata QueryGrid query session.

General

Metric	Description	Type
Query ID	ID of the query	String
Phase	Phase of the operation	String
Duration	Duration of the QueryGrid operation	Number
Config version	Requested version of the QueryGrid operation	Number
Link Name	Name of the link used for this QueryGrid operation	String

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Total nodes	Number of nodes on the initiating system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Transfer Information

Metric	Description	Type
Type	Type of transfer performed	String
Rows	Number of rows transferred	Number
Bytes	Number of bytes transferred	Number
Rate	Rate that data is transferred for the QueryGrid operation	Number
Bottleneck	Limiting factor for the QueryGrid operation	String
Compression ratio	Uncompressed size divided by the compressed size	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Total nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Total Nodes	Number of nodes on the target system	Number

Metric	Description	Type
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Phases Tab

The **Phases** tab is available when you click an icon that represents a Teradata QueryGrid query operation. It contains tabs that provide information about the metadata, execution, and data transfer phases of the query session. The **Phases** tab support up to five panes depending on the details associated with the operation, for example, Metadata1, Metadata2, Metadata3, Execution1, Execution2, Data Transfer.

Metadata Tab**Initiator Information**

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	The node that initiated the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String

Metric	Description	Type
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Execution Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	Node from which the query initiated	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String

Metric	Description	Type
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Data Transfer Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	String
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the max CPU	String
Data skew	Node that had the max data	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Total nodes	Number of nodes	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	Number
Data skew node	Node that had the maximum amount of data	Number
Bytes sent	Number of bytes sent	Number
Bytes received	Number of bytes received	Number
Compression ratio	Whether or not compression is enabled	String

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays up to four bridges.

Configuration Tab

The **Configuration** tab contains provides configuration information about the selected operation for a QueryGrid query session.

Fabric

Metric	Description	Type
Name	Name of the fabric used in Teradata QueryGrid	String
Port	Port on which the fabric runs	String
Software version	Version of software that the fabric runs	String

Initiating Connector

Metric	Description	Type
System	Name of the initiating system	String
Connector	Point from which the query originated	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Communication Policy

Metric	Description	Type
Name	Name given to this communication policy	String
Encryption	Whether or not data encryption is enabled	String
Compression	Whether or not compression is enabled	String

Target Connector

Metric	Description	Type
System	Name of the target system	String
Connector	Destination point for the query	String

Metric	Description	Type
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Errors Tab

The **Errors** tab provides information about Teradata QueryGrid queries that had errors. It displays only if there are errors.

Metric	Description	Type
Time	Time the error occurred	Number
Phase	Phase when the error occurred	String
Message	The error message	String
Node	The hostname of the node that generated the error	String
Connector	The name of the connector that resulted in the error	String
System	Name of the system	String
Origin	Indicates whether the error occurred on the initiator or the target of the Teradata QueryGrid operation	String

Managing Queries and Sessions

You can manage queries and sessions to improve workload performance for Teradata systems.

Abort

Abort the selected query or session

Change Priority

Change the priority of the selected query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change priorities or workloads, or release queries. If you log off, close a portlet, or open a new portlet instance, you must log in again.

Note:

Change Workload is available only if Teradata Active System Management (TASM) is enabled. If TASM is disabled, **Change Priority** is available. If you do not see **Change Workload** or **Change Priority** in the list, the system you are monitoring does not support these features or you do not have permission to use them. If the query you are monitoring is delayed, only **Release Query** is available.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. Click the row of the query you want to abort.
2. Click ▾ to the right of the session number and select **Abort**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
6. Click **Next**.
7. Click **Next** to confirm your selection.
8. Click **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▾ to the right of the session number and select **Change Priority**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.
 - Type an account string.
6. [Optional] Select the check box to use the account string as the default for the selected session. The priority will be changed for the selected query and all subsequent queries in the current session.
7. Click **Next**.
8. Click **Next** to confirm your selection.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Workload**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Select a different workload from the list and click **Next**.
5. Click **Next** to confirm your selection.
6. Click **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Release Query**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Click **Next** to confirm your selection.

Changing Multiple Sessions

In Teradata systems, you can change system resources for multiple sessions from the summary view.

1. Select a session action from the **Table Actions** list.
Check boxes appear next to the sessions.
2. Do one of the following:
 - Select the first check box located in the column heading to choose all sessions.
 - Select specific check boxes to choose only those sessions.
3. Click **Next**.
4. Log in to Analytics Database, if prompted.
5. Follow the instructions that appear on subsequent screens.

Settings View

The **Settings** view allows you to choose whether formatted or unformatted SQL is displayed in the **SQL** tab.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click ⚙ in the portlet frame and select **Settings**.


2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Selecting Display Options

In the **Settings** view, you can choose whether formatted or unformatted SQL is displayed in the **SQL** tab.

1. Click  in the portlet frame and select **Settings**.
2. Do one of the following:
 - Select the **Format SQL** check box to display formatted SQL.
 - Clear the **Format SQL** check box to display unformatted SQL.
3. Click **OK**.

Query Log

Query Log Overview

The **Query Log** portlet enables Teradata DBAs to view key reports based on the historical DBQL data in the PDCRDATA.DBQLogTbl_Hst table in Analytics Database. If the number of DBQL rows exceeds 10 million per day, you can specify the name of a database view in the **Monitored Systems** portlet to load a subset of the query log data to make sure only the most important data is collected.

You can review queries designated as suspect based on thresholds set for the Query Log data collector in the **Monitored Systems** portlet, and access system-wide, application specific, and user specific information. You can view key metrics for any date for which query log data has been collected and retained by Viewpoint, and drill down to see details about specific queries. You can also view the trends of key performance indicators, aggregated daily or weekly and going back over a time frame from the last six months to two years.

You can only see data in the **Query Log** portlet if the optional query logging (DBQL) feature has been enabled in Analytics Database and the query logs are archived using PDCR. For more information, see *Teradata® Database Administration*.

The **Query Log** portlet does not participate in the rewind feature. The portlet provides a date selector to select the date of the query log information you want to view.

Query Log View

The **Query Log** view displays the following information:

System Selection

Shows the name of the system currently selected and allows you to select a different system.

Date Selector

Shows the date of the information displayed in the **Applications** and **Users** counts and the summary table of logged and suspect queries, and allows you to select a different date.

Applications Count

Shows the number of applications that submitted queries on the selected date.

Users Count

Shows the number of users who submitted queries on the selected date.

Bar Chart

Displays a visual representation of the numbers of queries that fall into each category for the selected metric. When you hover over each bar, a balloon displays the actual number of queries the bar represents.

Trend Chart

Displays a visual representation of the magnitude of the selected metric over the selected time period. When you hover over any point on the trend line, a balloon displays the actual magnitude of the selected metric on that date.

Teradata System Version Change

If selected, displays a vertical orange line in the trend chart indicating the date of a change in the version of a Teradata system.

Ruleset Change

If selected, displays a vertical blue line in the trend chart indicating the date of a change in the active ruleset for the system.

Summary Table

Displays logged and suspect queries for the selected date.

Filters

Filters allow you to change displayed data by showing only rows that match your filter criteria. You can also sort on the column headers to find the details you want.

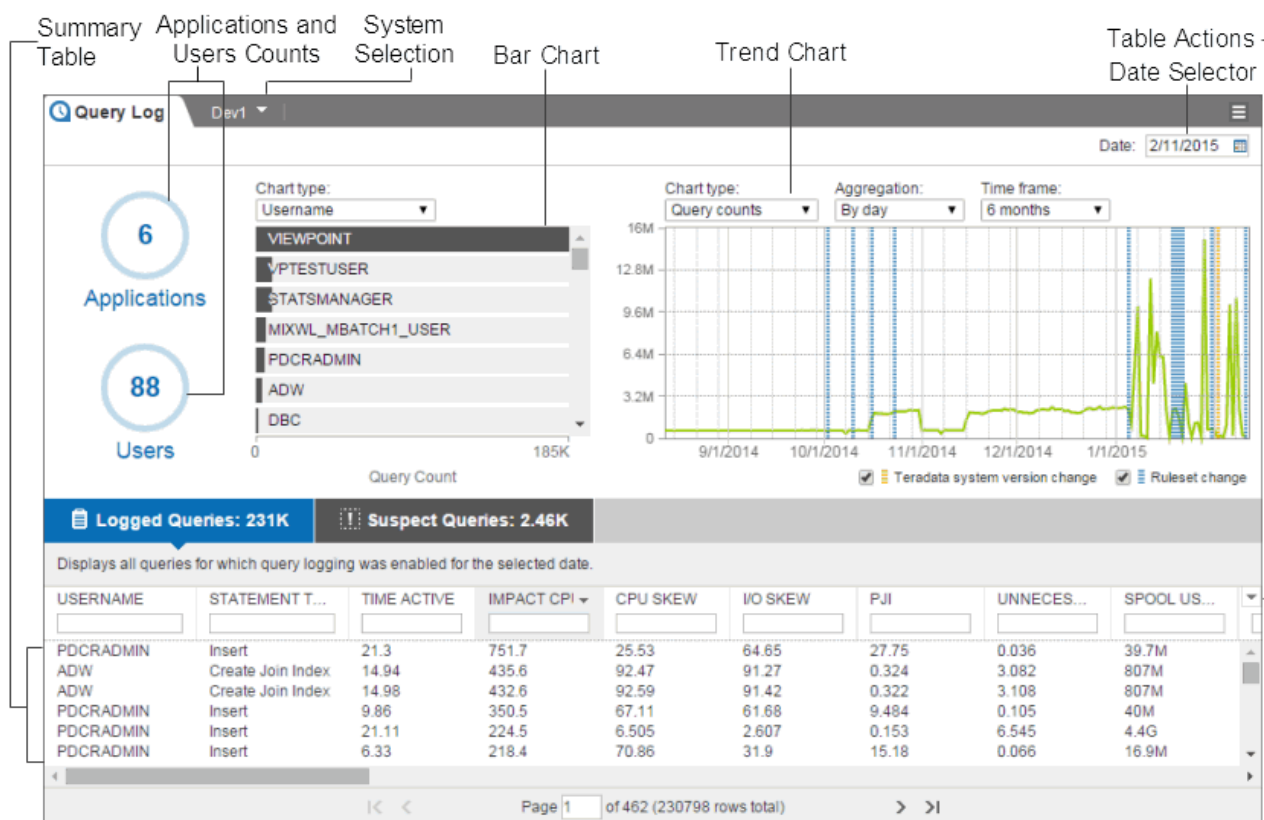
☐ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to select, lock, and order the displayed columns.

Export creates a .csv file containing all available data up to one million rows. If more than one million rows of data are available, then use the filters to display the data you want before exporting it.

For more information, see [Summary Table Controls](#).



Changing the Date

Use the date selector to designate the date of the information you want to view. The date you select determines the information displayed in the **Applications** and **Users** count, the bar chart, and the summary tables. It does not affect the values in the trend chart.

The time zone and time the data collector begins collecting data as well as the amount of data retained are determined by the values defined for the Query Log data collector in the **Monitored Systems** portlet.

1. Click the date selector box.
2. Select a date from the calendar.

Bar Chart Basics

The bar chart displays a visual representation of the numbers of queries that fall into each category for the selected metric. When you hover over each bar, a balloon displays the actual number of queries the bar represents. The horizontal **Query Count** line under the chart indicates the number of queries represented by the full width of the chart.

Bar Chart Metrics

Values for the following metrics can be displayed in the bar chart.

Metric	Description
Account String	Number of queries submitted by the account string represented by the bar
App ID	Number of queries associated with the displayed App ID. The App ID is the ID of the application used to access Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.
Cache Flag	Number of queries that used the cache flag type represented by the bar. The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan was generated. • G: The query is a parameterized query and a <i>generic</i> plan was generated. • A: The query is a parameterized query and a <i>specific always</i> decision was taken. Each time a query is submitted, the USING values are considered and the query is parsed.
Client Address	Number of queries submitted by the client IP address represented by the bar
ClientID	Number of queries submitted by the client ID represented by the bar
Error Code	Number of queries that have the error code number represented on the bar
Number of Active AMPs	Number of queries that used the number of AMPs represented by the bar
Statement Type	Number of queries that used the SQL statement type represented by the bar
Username	Number of queries submitted by the Teradata user represented by the bar
Workload End	Number of queries that finished execution in the workload represented by the bar
Workload Start	Number of queries that started execution in the workload represented by the bar
Workload Start -> End	Number of queries that started and finished in the workload represented by the bar

Trend Chart Basics

The trend chart displays a visual representation of the magnitude of the selected metric over the selected time period. The horizontal axis represents the date, and the vertical axis represents the value of the metric selected for display. The green trend line plots the metric values that occurred during the selected time frame, until the day before the current date. As you move your mouse pointer over any point on the trend line, a balloon displays the actual magnitude of the selected metric on that date.

You can choose to display the metrics aggregated by day or by week, as well as the time frame to display. You can also select the check boxes under the chart to display vertical lines indicating the dates of changes in Teradata system version or ruleset.

Trend Chart Metrics

Trends for the following metrics can be displayed.

Metric	Description
Burn rate	Percentage of the overall system capacity that was effectively consumed by all queries during their execution
Gating efficiency	Percentage indicating the overall query duration time spent executing versus being delayed. A value of 100% indicates all time was spent executing.
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the sum of the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.
Max spool	Maximum amount of temporary spool space used by any one query, using a 1024 byte kilobyte as the base
Query counts	Total number of logged queries
Suspect queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector
System utilization	Percentage of the system's CPU seconds that was utilized by logged queries
Time active	Amount of time in seconds that was recorded as utilized by logged queries. Excludes any time queries were delayed before execution.
Time delay	Amount of time in seconds that queries were delayed before being executed
Time expansion	Ratio of how long queries took to execute versus the ideal execution time. The larger the number, the more the execution time of the query was expanded. For example, a value of 10x to indicates an expansion of 10 times the ideal execution time.
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries
Total I/O	Total number of I/O requests
Total parser CPU	Total amount of CPU consumed by the parsing engine in processing queries

Configuring the Trend Chart

You can configure and display a trend chart that displays aggregate values for key metrics over a selected time period. You can choose to view the values aggregated by day or by week. You can also choose whether to display the values from the past six months, one year, or two years, through the day before the current date.

1. Select the metric you want to display from the **Chart type** list.
2. Select the aggregation period from the **Aggregation** list.
3. Select the time period display from the **Time frame** list.

Logged Queries Tab

The **Logged Queries** tab provides key metrics for queries logged on the selected date.

Suspect Queries Tab

The **Suspect Queries** tab displays information for all logged queries that are designated as suspect. Suspect queries are those whose values surpass thresholds defined for the Query Log data collector in the **Monitored Systems** portlet. Thresholds can be set for the following metrics:

- CPU Skew
- I/O Skew
- Product Join Indicator
- Unnecessary I/O

Any queries that result in more than one AMP CPU second consumed and that exceed the defined thresholds are listed in the **Suspect Queries** tab, and the value that exceeds the threshold is displayed in red.

Logged Queries and Suspect Queries Tabs Metrics

Metrics available for display are listed as following.

Metric	Description	Type
Account	Account of the user who submitted the query	String
Active AMPs	Number of AMPs involved in executing the query	Number
AMP CPU Time	Total number of CPU seconds consumed by AMPs in processing the query	Number
App ID	The ID of the application used to access the Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.	String
Burn Rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent
Cache Flag	The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated. • A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed. 	String
Client Address	IP address of the client	String

Metric	Description	Type
Client ID	ID of the client that submitted the query	String
CPU Skew	CPU skew of the entire query	Percent
Error Code	Error code associated with the query	String
Final Workload	Name of the workload under which the query finished	String
Flagged	Denotes whether the query has been flagged for further analysis	String
Gating Efficiency	The percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates that all time was spent executing.	Percent
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
I/O Skew	I/O skew of the entire query	Percent
Max AMP CPU Time	The maximum CPU consumed by a single AMP while executing the query	Number
Max AMP I/O	The maximum I/Os submitted by a single AMP while executing the query	Number
Parser CPU Time	Total amount of CPU seconds consumed by the parsing engine in processing the query	Number
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proc ID	ID of the parsing engine that run the query	String
Query ID	System-wide unique ID that identifies the query	String
Row Count	Number of result rows returned by the query	Number
Spool Usage	Amount of spool consumed by the query, using a 1024 byte kilobyte as the base	Number
Start Time	The time the query was submitted	Number
Starting Workload	Name of the workload under which the query started execution	String
Statement Type	SQL statement type the query used	String
Step Count	Number of steps in the query plan	Number
Time Active	Number of seconds the query spent running. Excludes any time the query was delayed before being executed.	Number
Time Delay	Number of seconds the query was delayed before running.	Number
Total Duration	Total elapsed time it took for the query to run once it was submitted	Number
Time Expansion	Ratio of how long the query took to run versus how long it could have ideally run. The larger the number, the more the execution time of the	Number

Metric	Description	Type
	query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	
Total I/O Count	Total number of I/O requests the query used	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring in proportion to the CPU consumed	Number
Username	Name of the user who submitted the query	String

Viewing Queries by Application

You can view summary information about the applications that submitted logged queries on the selected date.

Note:

Applications can be defined in the **Query Group Setup** portlet. Query Groups are automatically created when a query is logged with an "app" or "applicationName" query band.

1. Click the **Applications** count button.

Applications View Metrics

Metrics available for display are listed as following.

Metric	Description
Application	Application that submitted queries on the selected date
Burn rate	Percentage of the overall system capacity that was effectively consumed by all queries during their execution
Gating efficiency	Percentage indicating the overall query duration time spent executing versus being delayed. A value of 100% indicates all time was spent executing.
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.
Max spool	Total amount of temporary spool space used
Query counts	Total number of executed queries
Suspect Queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector
System Utilization	Percentage of the system's CPU that was utilized
Time active	Amount of time in seconds that queries spent executing. Excludes any time queries were delayed before being executed.

Metric	Description
Time delay	Amount of time in seconds that queries were delayed before being executed
Time Expansion	Ratio of how long the query took to execute versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded. For example, a value of 10x to indicates an expansion of 10 times the ideal execution time.
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries
Total I/O	Total number of I/O requests
Total parser CPU	Total amount of CPU consumed by the parsing engine in processing queries

Viewing Queries by User

You can view summary query information for users who submitted logged queries on the selected date.

1. Click the **Users** count button.

Users View Metrics

Metrics available for display are listed as following.

Metric	Description
Burn rate	Percentage of the overall system capacity that was effectively consumed by this query during its execution
Gating efficiency	Percentage of the overall system capacity that was effectively consumed by all queries during their execution. A value of 100 indicates all time was spent executing.
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.
Max spool	Total amount of temporary spool space used
Query counts	Total number of executed queries
Suspect Queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector
System Utilization	Percentage of the system's CPU utilized by the user or application
Time active	Amount of time in seconds the query or job spent executing. Excludes any time the query or job was delayed before being executed.
Time delay	Amount of time in seconds the query or job was delayed before being executed
Time Expansion	Ratio of how long the query took to execute versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded.

Metric	Description
	For example, a value of 10x to indicates an expansion of 10 times the ideal execution time.
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries
Total I/O	Total number of I/O requests
Total parser CPU	Total amount of CPU consumed by the parsing engine in processing queries
Username	Name of the user who submitted the query

Application and User Queries Details View

You can drill down on any of the applications listed in the **Applications** view. You can also drill down on any of the users listed in the **Users** view.

When you click a row for an application or user, details for that application or user are displayed on three tabs: **Overview**, **Queries**, and **Trends**.

If you are viewing application details, you can also select the application version of the information you want to view. Application version is always included when a query is logged with the "version" query band.

Overview Tab

The **Overview** tab includes a **Summary** section that lists key metric values and a **Counts** section that displays bar charts representing query counts for a selected metric.

Summary

The **Summary** section displays values for the following metrics.

Metric	Description	Type
Burn Rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent
Gating Efficiency	Percentage indicating the percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates all time was spent executing.	Percent
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
Logged Queries	Number of queries for which query logging was enabled for the selected date	Number
Max Spool	Total amount of temporary spool space used	Number

Metric	Description	Type
Suspect Queries	Number of queries designated as suspect based on thresholds set in the Query Log data collector	Number
System Utilization	Percentage of the system's CPU utilized by the user or application	Percent
Time Active	Number of seconds the query spent executing. Excludes any time the query was delayed before being executed.	Number
Time Delay	Number of seconds the query was delayed before being executed	Number
Time Expansion	Ratio of how long the query took to execute versus how long it could have ideally run. The larger the number, the more the execution time of the query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	Number
Total AMP CPU	Total amount of CPU consumed by AMPs in processing queries	Number
Total I/O Count	Total number of I/O requests the query used	Number
Total Parser CPU	Total amount of CPU consumed by the parsing engine in processing the query	Number

Counts

The **Counts** section displays bar charts representing total counts for the following metrics.

Metric	Description
Account String	Number of queries submitted by the account string represented by the bar
App ID	Number of queries associated with the displayed App ID. The App ID is the ID of the application used to access Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.
Cache Flag	Number of queries that used the cache flag type represented by the bar. The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated. • A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed.
Client Address	Number of queries submitted by the client IP address represented by the bar
ClientID	Number of queries submitted by the client ID represented by the bar
Error Code	Number of queries that have the error code number represented on the bar
Number of Active AMPs	Number of queries that used the number of AMPs represented by the bar
Statement Type	Number of queries that used the SQL statement type represented by the bar

Metric	Description
Username	Number of queries submitted by the Teradata user represented by the bar
Workload end	Number of queries that finished execution in the workload represented by the bar
Workload start	Number of queries that started execution in the workload represented by the bar
Workload start -> end	Number of queries that started and finished execution in the same workload represented by the bar

Queries Tab


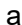
The **Queries** tab enables you to view two subtabs:

- **Logged Queries:** All queries that were logged for the selected user or application and application version.
- **Suspect Queries:** All suspect queries that were logged for the selected user or application and application version. The values that exceeded the defined thresholds are highlighted in red.

Trends Tab

The trend chart displays aggregated values for all of the queries submitted by the selected user or application. You can choose a daily or weekly aggregation period. You can view a time frame of the last six months, the last year, or the last two years. You can choose to display vertical lines indicating the date of ruleset and version changes using the check boxes on the upper right. Application version is recorded when a query is logged with the "version" query band.

Configured charts are shared among all similar views. For example, if you drill down into a user and set up the chart in the **Trends** tab, the **Trends** tab for applications will have the same configuration.

By default, one chart is displayed. You can add  and remove  charts.

Query Details View

You can view query details by clicking any of the queries listed in the **Query Log** portlet tables. Query details are displayed on an **Overview** and a **SQL** tab. If applicable, **Query Band** and **Errors** tabs are also displayed.

Viewing Query Details

You can view query details from several different views in the **Query Log** portlet.

1. Do one of the following:

Option	Description
From the Query Log view	Click a query in the Logged Queries or Suspect Queries tab.
From the Applications view	<ol style="list-style-type: none"> Click the row of the application. Click the Queries tab. Click a query in the Logged Queries or Suspect Queries tab.
From the Users view	<ol style="list-style-type: none"> Click the row of the username. Click the Queries tab. Click a query in the Logged Queries or Suspect Queries tab.

Overview Tab

The **Overview** tab on the query details view displays five sections containing information about the selected query. The sections and the metrics each displays are shown as following.

Query Information

Metric	Description	Type
Query ID	System-wide unique ID that identifies the query	String
Proc ID	ID of the parsing engine that runs the query	String
Start time	The time the query was submitted	Number
Statement type	SQL statement type the query used	String
Row count	Number of result rows returned by the query	Number
Active AMPs	Number of AMPs involved in executing the query	Number
Step count	Number of steps in the query plan	Number
Cache flag	The cache flag types are: <ul style="list-style-type: none"> • T: The flag is found in step cache • S: The query is a parameterized query and a <i>specific</i> plan is generated. • G: The query is a parameterized query and a <i>generic</i> plan is generated. • A: The query is a parameterized query and a <i>specific always</i> decision is taken. Each time a query is submitted, the USING values are considered and the query is parsed. 	String

Query Duration

Metric	Description	Type
Time active	Number of seconds the query spent running. Excludes any time the query was delayed before being executed.	Number

Metric	Description	Type
Time delay	Number of seconds the query was delayed before running	Number
Total duration	Total elapsed time it took for the query to run once it was submitted	Number

Query Performance

Metric	Description	Type
AMP CPU time	Total number of CPU seconds consumed by AMPs in processing the query	Number
Parser CPU time	Total number of CPU seconds consumed by the parsing engine in processing the query	Number
Total I/O count	Total number of I/O requests the query used	Number
Impact CPU	Impact of jobs on the system taking CPU skew into account. The value is equal to the sum of the highest CPU time in seconds consumed by any one AMP on behalf of a job or query multiplied by the number of AMPs.	Number
CPU skew	CPU skew for the entire query	Percent
I/O skew	I/O skew for the entire query	Percent
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time	Number
Spool usage	Amount of spool consumed by the query	Number
Gating efficiency	Percentage indicating the percentage of the overall query duration spent executing versus being delayed. A value of 100 indicates all time was spent executing.	Percent
Burn rate	Percentage of the overall system capacity effectively consumed by the query during its execution	Percent
Time expansion	Ratio of how long the query took to run versus how long it could have ideally run. The larger the number, the more the run time of the query was expanded. For example, a value of 8.8 means an expansion of 8.8 times the ideal execution time.	Number

Workload Information

Metric	Description	Type
Starting workload	Name of the workload under which the query started execution	String
Final workload	Name of the workload under which the query finished	String

Session Information

Metric	Description	Type
Username	Name of the user who submitted the query	String
Application	Name of the application that submitted the query	String
Version	Version of the application that submitted the query	String
Account	Account of the user who submitted the query	String
App ID	The ID of the application used to access the Analytics Database. Typically the App ID is the name and version of the Teradata Tool or Utility accessing Analytics Database.	String
Client ID	ID of the client that submitted the query	String
Client address	IP address of the client	String

SQL Tab

The **SQL** tab displays the SQL for the selected query. The SQL information is presented from the PDCRDATA.DBQLSqlTbl_Hst table, if available. If not available, the information from the PDCRDATA.DBQLLogTbl_Hst table is displayed.

The information is read-only.

By default, the SQL is formatted. Clear the **Format SQL** check box to display unformatted SQL.

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Errors Tab

The **Errors** tab has a text field that displays any errors for the selected query. The information is read-only. The tab is available only when the query resulted in an error.

Flagging a Query for Analysis

You can choose to flag a query for further analysis. This tags the query so that you can load it for analysis in Index Analyzer.

1. From the query details view, select the **Flag for analysis** check box at the top right.

Query Monitor

Query Monitor Overview

The **Query Monitor** portlet allows you to view information about queries running in the relevant system so you can target problem queries.

For Teradata systems, you can see the progress of submitted queries and target longer-running queries so you can optimize them. After you have identified a problem query, you can take action to correct the problem by changing the priority or workload, releasing the query, or aborting the query or session. You can take these actions for a single query or session, or multiple queries or sessions at a time. The session view provides a table listing the sessions, account strings, users, or utilities running on the database.

From the **Settings** view for Teradata systems, you can choose to display the top sessions graph and format SQL, and set the criteria values used to display sessions in the **My Criteria** view and **Skewed AMP** views. The **My Criteria** session list displays only sessions that have exceeded the defined thresholds.

For Aster systems, you can see the progress of queries. The session view provides a table listing the sessions, databases, and users running on the database. From the **Settings** view for Aster systems, you can choose to display the sessions list and state filter bar.

For Teradata Machine Learning Engine, you can see the progress of queries. The session view provides a table listing the users.

For Teradata QueryGrid queries, you can see information about the progress and status of Teradata QueryGrid requests. You can see details and a visualization of the query's progress in the [Details View for Teradata QueryGrid](#).

Note:

- If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is not available. You cannot enter alternate credentials.
 - Support for TASM ARM is available for SQL Engine versions 17.10 and later.
-

Query Monitor View

The **Query Monitor** view shows a list of sessions running on a system and allows you to categorize the displayed sessions, such as by account string, session, user, and so on.

Selection Menus


Enables you to select the system and report that defines the sessions displayed.

Top Sessions Graph

Displays the sessions with the largest value compared with the total number of sessions currently running in Analytics Database.

State Filter Bar

Displays a count of the sessions in each state. Click any state in the bar to change the displayed data in the summary table to show only the sessions running a query in the selected state.

 **Overflow Menu** Displays additional states. If the number of states on the **State Filter Bar** exceeds the available space, you can select another state to replace the current state.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information about sessions, account strings, users, or utilities in columns. The current view is configured in the **Configure Columns** dialog box and **Settings** view. The view is refreshed when new data is collected. Click a row in the table to see details.

Table Actions

For Teradata systems, table actions include:

- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display and set thresholds.
- **Abort**, **Change Priority**, **Change Workload**, and **Release Query** allow you to manage the selected queries and sessions.
- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For Teradata Aster systems, table actions include:

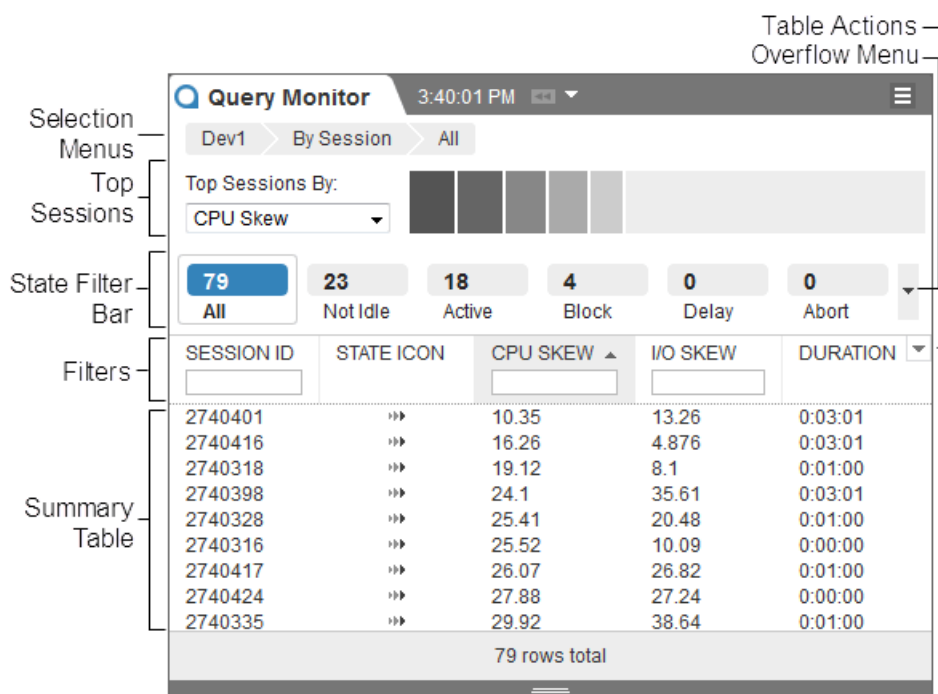
- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display.
- **Abort Process** allows you to abort the selected non-idle process.
- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information on **Clear Filters**, **Configure Columns**, and **Export** table actions, see [Summary Table Controls](#).

For Teradata Machine Learning Engine systems, table actions include:

- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display.
- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information on **Clear Filters**, **Configure Columns**, and **Export** table actions, see [Summary Table Controls](#).



Selecting a System to Monitor

Choose a system, type of view, and information to monitor in the portlet, starting with the highest-level menu on the far left. Menu choices are based on what you selected in the previous menu.

1. In the selection menu, click the currently selected system name to display a list of available systems.
2. Select a system.
3. Select >.
4. Select a type of view from the menu.
5. Select >.
6. If you selected **By Session** or **By Session (Exclude Utilities)** for a Teradata system, you can further define the criteria.
7. Select >.

Selection Menu Choices

The **Query Monitor** portlet displays data for one system at a time, based on the data type you select. The selection menu choices are specific to the type of system selected.

Each of the following menu choices displays data in the summary view.

Teradata System

Selection Menu Choices	Description
By Account String	Displays session information by account string
By Blocker	Displays sessions that are blocking other sessions
By Session > All	Displays all sessions
By Session > My Criteria	Displays sessions with metric values that exceed thresholds set in the Settings view
By Session (Exclude Utilities) > All	Displays all sessions other than utility sessions
By Session (Exclude Utilities) > My Criteria	Displays non-utility sessions with metric values that exceed thresholds set in the Settings view
By User	Displays session information by user
By Utility > By Job	Displays utility session information by job
By Utility > By Type	Displays active job count and current limit for all utilities
By Vproc > By PE	Displays PEs by most skewed sessions
By Vproc > By Skewed AMP	Displays AMPs by most skewed sessions that exceed the metric values that are set in the Settings view

Aster System

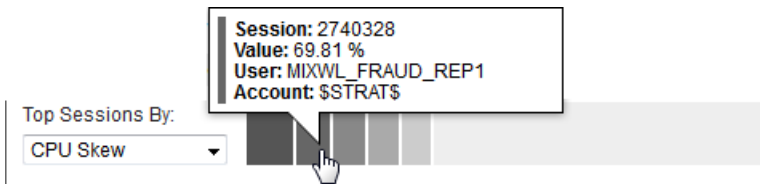
Selection Menu Choices	Description
By Database	Displays rows of databases
By Session	Displays rows of sessions
By User	Displays rows of users

Teradata Machine Learning Engine

Selection Menu Choices	Description
By User	Displays rows of users
By Session	Displays rows of sessions

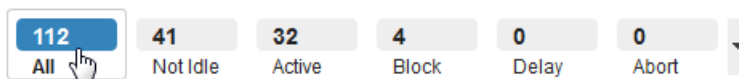
Top Sessions Graph

Select **Top Sessions Graph** in the **Settings** view to show or hide the **Top Sessions By** graph. Hover over a bar to see the session, value, user, and account.



State Filter Bar

The state filter bar allows you to display specific states in the view. For a Teradata system, if the number of states exceeds the available space, you can select another state from the **Overflow Menu**. The new state replaces the last state on the state filter bar.



The state filter buttons provide a count of sessions for each state category. Click a state in the state filter bar to display only sessions with a query in the selected state. For example, click **Delay** to display the sessions running a query that is queued and waiting to run. Hover over a state filter button to display the full state name or number.

Teradata System States

All

Sessions in the system

Not Idle

Sessions in any state except idle

Active

Sessions processing queries

Block

Sessions processing a query waiting for a locked resource, such as a database table or view

Deferred

Sessions running a query that is in deferred queue waiting to run.

Delay

Sessions waiting to be processed

Abort

Sessions in the process of aborting (rolling back changes made by the query)

Resp

Sessions running a query that has completed and is sending (responding) spooled data back to the user

Idle

Sessions not currently running a query

Parse

Sessions running a query that is being parsed and has not begun to execute

Other

Sessions whose status is unknown and do not fall into another category

QTDelayed

Sessions delayed due to a queue table restriction

SesDelayed

Utility sessions that are in the workload delay queue

Response-Held

In SLES 11, sessions running a query whose response is being held until a specified number of seconds elapses

Direct Blockers*

Sessions that immediately block the current session

For example, if session 1 is blocking session 2, which in turn is blocking session 3, the **Blocked By** tab for session 3 shows both sessions 1 and 2 under **All**. Under **Direct Blockers**, it shows only session 2.

Root Cause*

Sessions that are not blocked, but are causing other sessions to be blocked

Granted*

Sessions that are blocked and are blocking other sessions because they have been granted a lock

Waiting*

Sessions that are blocked and are ahead of other blocked sessions in the queue

*Appears when **By Blocker** has been selected from the menu and when connected to Teradata Database 15.0 and later.

Aster System States

All

Sessions in the system

Active

Sessions processing queries

Idle

Sessions not running

Canceled

Processes that were stopped

Pending

Sessions waiting to be processed

Teradata Machine Learning Engine States

All

Sessions in the system

Active

Sessions processing queries





Idle

Sessions not running




State Icons

Teradata Sytstem State Icons



Icons	Name	Definition
×	Aborting	Query has been aborted, and changes are being rolled back.
»»	Active	Query is running.
»I	Blocked	Query is waiting for a lock held by another query.
III	Delayed	Query is in a delay queue waiting to run.
⌚	Deferred	Query is in a deferred queue waiting to run.
🔌	Host-Restart	Host associated with this session is currently restarting.
▣	Idle	No query is running.

Icons	Name	Definition
?	Other	Query is in an unknown state.
	Parsing	Query is being parsed before running.
	QTDelayed	Query delayed due to a queue table restriction.
	Response	Query is returning results to the user.
	SESDelayed	Query is delayed because a utility limit has been exceeded in FastLoad, MultiLoad, FastExport, or ARC.

Aster System State Icons

Icons	Name	Definition
	Active	Process is running.
	Canceled	Process has been canceled.
	Pending	Process is waiting to run.

Teradata Machine Learning Engine State Icons

Icons	Name	Definition
	Active	Query is running.
	Idle	No query is running.

Query Monitor Metrics

Available metrics differ depending on the type of system and selection menu options you choose.

Teradata System Metrics

Session Metrics

The selection menu choice **By Session > All** or **My Criteria** or **By Session (Exclude Utilities) > All** or **My Criteria** makes the following metrics available. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Account	Account from which a query was submitted	String
Blocked Time	How long the query has been blocked	Number

Metric	Description	Type
CPU Use	Percent of available CPU seconds on the system used during the last sampling period	Percent
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number
Duration	How long the query has been running	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
In State	How long the query has been in the current state	Number
Partition	Partition in which the query is running	String
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(For Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String
Query Band	Entire query band string. Query bands are a set of name-value pairs defined by the user to tag sessions or transactions with an ID through a SQL interface.	String
Request CPU	CPU seconds needed to run the query	Number
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Request I/O	Disk I/Os performed to run the query	Number
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Request Count	Number of queries the session has executed	Number
Session ID	Unique session identifier	Number
Snapshot CPU Skew	CPU skew during the last sample	Percent
Snapshot I/O Skew	I/O skew during the last sample	Percent
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.	Percent

Metric	Description	Type
Start	Time that the query started running on Analytics Database	Number
State Icon	Icon representing the current state of the query	Icon
State	Text describing the current state of the query	String
Temp Space	Temp space used by the query, using a 1024 byte kilobyte as the base	Number
Unnecessary I/O	Ratio of I/O to CPU for the current query	Number
User	Name of the user who submitted the query	String
Workload	Workload in which the query is running	String

Account String and User Metrics

The selection menu choice **By Account String** or **By User** makes the following metrics available. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number
Δ Request Count	Total number of queries submitted by the user or account string since the last sample	Number
Aborting Count	Queries that have been aborted and changes are being rolled back	Number
Account	Account from which a query was submitted	String
Active Count	Queries that are running	Number
Blocked Count	Queries waiting for a lock held by another query	Number
CPU Use	Available CPU seconds on the system used during the last sampling period	Percent
Delayed Count	Queries in a delay queue waiting to run	Number
Idle Count	Queries that are not running	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
Parsing Count	Queries being parsed before running	Number
Response Count	Queries that are returning results to the user	Number
Spool	Current spool used by the request across all AMPs in bytes, using a 1024 byte kilobyte as the base	Number

Metric	Description	Type
Total	Sessions for the user or account (total of active, aborting, blocked, idle, parsing, response, delayed, and unknown)	Number
Unknown Count	Queries in an unknown state	Number
Username	Name of the user that submitted the query	String

Utility Metrics

The selection menu choice **By Utility > By Job** or **By Type** makes the following metrics available. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Active	Number of utilities of this utility type that are active	Number
AMP CPU	Total AMP CPU-seconds consumed by the job	Number
AMP I/O	I/O count across all AMPs for the job	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Limit	Current limit on the number of utilities of this type	Number
LSN	Login sequence number for the job	Number
Name	Utility names	String
Sessions	Session count for the job	Number
Username	Name of the user who submitted the query	String
Utility	Utility in which the session is running	String

Vproc Metrics

The selection menu choice **By VProc** makes the following metrics available.

Metric	Description	Type
AMP ID	ID of the skewed AMP, limited to those AMPs with sessions that exceed the metric values that are set in the Skew tab in the Settings view	Number
PE	PE ID	Number
PE CPU	PE usage in CPU-seconds	Number
Sessions	<ul style="list-style-type: none"> Sessions count on the PE Skewed session count on the AMP 	Number

Metric	Description	Type
	Note: This number is not the total number of sessions on an AMP	
Skewed CPU	CPU-seconds used by skewed sessions (sessions that exceed the skew threshold set in the Settings view) on an AMP Note: This number is not the total CPU usage on an AMP	Number

Blocker Metrics

The selection menu choice **By Blocker** makes the following metrics available by default. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Session ID	Unique session identifier of the query that is blocking	Number
Username	Name of the user that is running the query that holds the lock	String
State Icon	Icon representing the current state of the query	Icon
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 	String
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session	Number
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d+hh:mm:ss</i>	Number
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d+hh:mm:ss</i>	Number
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 	String

Metric	Description	Type
	3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock	

Aster System Metrics

Session Metrics

The selection menu choice **By Session** makes the following metrics available.

Metric	Description	Type
Database	Name of the database on which the process is running	String
Icon	Values for process types are as follows: <ul style="list-style-type: none"> • SQL-MR • SQL • Teradata Import • Teradata Export 	Number
In State	How long the process has been in the current state	Number
Process ID	Unique process identifier	Number
Session ID	Unique session identifier	Number
Start	Time the process began running	Number
State	Process state, such as active, canceled, pending	String
State Icon	Symbol of the process state	Icon
Username	Name of the user who submitted the process	String
Workload	Workload in which the process is running	String

Database and User Metrics

The selection menu choice **By Database** or **By User** makes the following metrics available.

Metric	Description	Type
Active Count	Active processes	Number
Canceling Count	Processes created by the user that were canceled and are stopping	Number

Metric	Description	Type
Database	Name of the database on which the process is running	String
Idle Count	Number of processes with no running processes on this database or for this user	Number
Pending Count	Number of processes waiting to run	Number
Total	Number of active, idle, and pending processes	Number
Username	Name of the user who submitted the processes	String

Teradata Machine Learning Engine Metrics

User Metrics

The selection menu choice **By User** makes the following metrics available.

Metric	Description	Type
In State	How long the process has been in the current state	Number
Process ID	Unique process identifier	Number
Session ID	Unique session identifier	Number
Start	Time the process began running	Number
State	State of the process	String
State Icon	Symbol of the process state	Icon
Type	Type of the phase, such as data transfer, import, export, install file, or prepare transaction.	String
Username	Name of the user who submitted the process	String

Session Metrics

The selection menu choice **By Session** makes the following metrics available.

Metric	Description	Type
Session ID	Unique session identifier	Number
Process ID	Process identifier	Number
Username	Name of user who submitted query	String
State	Text describing the current state of the query	String
Type	Statement type for this process	String

Metric	Description	Type
In State	How long the query has been in the current state	Number
Start	Time that the query started running	Number
CPU Percent	CPU use in percent	Double Precision
CPU Seconds	Total CPU seconds consumed at collection	Double Precision
Memory KB	Memory use in KBs	Number
Total I/O	Total I/O in use	Number

QueryGrid System Metrics

Metric	Description	Type
Start	Time the query started running	Number
User	Initiating user	String
Link Name	Name of the link that defines the initiating and target connectors	String
Phase	Indicates the current phase: metadata, data transfer, or execution	String
Duration	How long the query has been running	Number
Rows Transferred	Number of rows transferred by the query	Number
Bytes Transferred	Number of bytes transferred by the query	Number
Transfer Rate	Rate at which data is being transferred	Number
Initiating System	System that initiated the query	String
Initiating Connector	Connector that initiated the query	String
Target System	System being accessed or to which data is being transferred	String
Target Connector	Connector through which the data is being accessed or transferred	String

Details View for All Systems

The details view displays statistics and information about the selected session. This view can be accessed by clicking a session row in the summary table.

Note:

If you are displaying Teradata system queries and you click a row for a QueryGrid query, a blue informational message banner indicates that the query is a QueryGrid query. You can click the **View Details** button to see the [Details View for Teradata QueryGrid](#).

Tabs

Teradata systems have **Overview**, **SQL**, **Explain**, **Skew**, **Query Band**, **Blocking**, **Blocked By**, **Defer**, and **Delay** tabs.

Teradata Aster systems have **Overview**, **SQL**, and **Explain** tabs.

Teradata Machine Learning Engine systems have **Overview**, **SQL**, and **Explain** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

▼ Manage Queries

Manages the query or session for Teradata systems with **Abort**, **Change Priority**, **Change Workload**, and **Release Query** options.

< > Previous/Next

Allows you to move through sessions without returning to the summary table.

The screenshot displays the Teradata Query Monitor interface. At the top, it shows 'System: Prod1 | Session: 2824421'. Below this is a 'Manage Queries' button and a 'Previous/Next' navigation control. The 'Overview' tab is selected, showing a summary of query and session information. The interface is divided into two main columns: 'QUERY INFO' and 'WORKLOAD INFO'. The 'QUERY INFO' column includes details like State (Active), Time in State (0:01:00), Total Duration (0:01:00), Spool Space (159,232), Temp Space (0), Request CPU (0), Request I/O (-), PJI (0), and Unnecessary I/O (2.73). The 'WORKLOAD INFO' column includes Name (WD_Mode...), Classification Mode (Auto), and Session Info (User: VIEWPOINT, Account: DBC, Partition: DBC/SQL, Requests: 11). A 'SNAPSHOT INFO' section at the bottom left shows CPU Use (0%), Impact CPU (5.62), Snapshot CPU Skew (18.2%), and Snapshot I/O Skew (17.9%). A source information box on the right lists the connection details: (TCP/IP) b308 153.64.107.218 SKATE ;SKATECOP2/153.64.202.243:1025 C ID=79EF3A1A ROOT JDBC15.00.00.0 9;1.7.0_45 01 LSS.

Overview Tab for Teradata System

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted. Under SESSION INFO, you can click a user or account to see all sessions for that user or account string.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	(Teradata Database 15.0 and earlier) CPU impact on the system based on the highest utilized AMP
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Session Information

Session Information	Description
User	Name of the user that submitted the query

Session Information	Description
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system

Overview Tab for Aster System

The **Overview** tab displays detailed information about key metrics for the selected session. The metric values provide a view of the query status on the system.

Query Information

Query Information	Description
Process ID	Process identifier
State	Process state, such as active, canceled, pending
In State	How long the process has been in the current state
Start	Time the process started running

Workload Information

Workload Information	Description
Workload Policy	Name of the workload policy used to manage the process. Defines a set of related queries, and allows them to be prioritized in a similar manner.
Service Class	Name of the service class used by the workload policy. Specifies the share of system resources allocated to a workload as a function of the class's priority and weight values.
Priority	Indicates the importance of the process. <ul style="list-style-type: none"> • High • Medium • Low

Workload Information	Description
Weight	Number ranging from 1 to 100 that indicates the precedence of the process within the priority level, increasing the resources that can be allocated for the process. A higher value indicates a greater level of precedence.

Session Information

Session Information	Description
User	Name of user that submitted the query
Database	Name of the database on which the process ran
IP Address	IP address of the user that submitted the process

Overview Tab for Teradata Machine Learning Engine

The **Overview** tab displays detailed information about key metrics for the selected session. The metric values provide a view of the query status on the engine.

Query Information

Query Information	Description
CPU Percentage	Percentage of total system CPU currently being used by the session
CPU Seconds	Number of CPU seconds used by the query in the last minute
In State	Amount of time the query has been running in the same state (HH:MM:SS)
Memory (KB)	Total memory in KB
Process ID	Process identifier
Read Bytes	Total number of bytes read from the storage layer
Read I/O Count	Total read I/O operations (read syscalls)
Start	Time the process started running
State	Process state, such as active, canceled, pending
Total I/O	Total I/O available
Write Bytes	Total number of bytes written to the storage layer
Write I/O Count	Total write I/O operations (write syscalls)


Session Information

Session Information	Description
User	Name of user that submitted the query
IP Address	IP address of the user that submitted the process

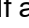
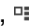
SQL Tab


The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Explain Tab for Teradata System

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status. If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

If you have the correct permissions, use  to retrieve the latest Explain steps from the database. Click **Refresh** to update. The refresh screen data is not maintained once the refresh view closes. Repeatedly retrieving current Explain data can impact system performance.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.

Step Information	Description
Confidence Level Indicator Icon	<ul style="list-style-type: none"> - No confidence in the estimate - Low confidence in the estimate - High confidence in the estimate - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Explain Tab for Aster System

The **Explain** tab displays the Explain text that shows the progress of a query. Each Explain step has a colored box that indicates the status of the step. This information is read-only. You can also export Explain data.

To view all steps and phase details, click **Show all**. Alternatively, click **Show summary** to view a shortened version of the Explain data.

Step Information	Description
Step Status Icon	Indicates the status of each Explain step: <ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black box. Active steps are indicated by a pulsating box (flashes gray and white). Steps to run are at the bottom of the list and indicated by a white box.
Time	Duration of the step
Type	Category of step, such as Command, CommitPreparedTransaction, DataTransfer, Import, PrepareTransaction, and Query
Phase/Detail	Summarized or full explanation of query that executes in this step. Can be used to diagnose the cause of the failure or slowness.


Explain Tab for Teradata Machine Learning Engine

The **Explain** tab displays the Explain text that shows the progress of a query. Each Explain step has a colored box that indicates the status of the step. This information is read-only. You can also export Explain data.

To view all steps and phase details, click **Show all**. Alternatively, click **Show summary** to view a shortened version of the Explain data.

Step Information	Description
Step Status Icon	Indicates the status of each Explain step: <ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black box. Active steps are indicated by a pulsating box (flashes gray and white). Steps to run are at the bottom of the list and indicated by a white box.
Time	Duration of the step
Type	Category of step, such as Command, Query, SQLMR, and DataTransfer
Phase/Detail	Summarized or full explanation of query that executes in this step. Can be used to diagnose the cause of the failure or slowness.

Exporting Explain Data

1. From the Explain tab, click .
2. Save the file to a location.
The Explain data is exported as a .csv file.

Skew Tab for Teradata System

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Query Band Tab for Teradata System

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Blocked By Tab for Teradata System

When a query is blocked, use the **Blocked By** tab to identify the other queries blocking the selected query. By reviewing the number and type of queries blocking the selected query, you can identify blocking issues on the system, determine the impact of this blocking on the selected query and, based on this information, decide the best course of action to resolve the issue.

The information in this tab is read-only. The tab is available only when the selected query is blocked. You can drill down into the session information for the blocking session. Use the **Session ID** link when connected to Teradata Database 14.10 and earlier, and click anywhere in a row in the table when connected to Teradata Database 15.0 and later.

The following information is available in the noted versions of Analytics Database.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	<p>A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName</p> <p>The information displays in the following order:</p> <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock 		Available
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Blocking Tab for Teradata System

When a query is holding locks that are blocking other queries from completing, use the **Blocking** tab to gauge the impact of this blocking. By analyzing the number and type of queries blocked by this query, you can decide the best course of action to resolve the issue.

This information in this tab is read-only. The tab is available only when the selected query is blocking other queries and when connected to Teradata Database 15.0 and later. From the tab, you can drill down to view the details of the blocked session.

Blocking Information	Description
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock
Session ID	Unique session identifier of the query that is blocking
Blocked Time	How long the query has been blocked, displayed as <i>d +hh :mm :ss</i>
Username	Name of the user that is running the query that holds the lock
Workload	Name of the workload where the query is actively running
Query Band	String of key/value pairs of the query band for the session or transaction

Defer Tab for Teradata System

The Defer tab displays details about the rules that are deferring a query. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Defer Time	Duration of time request has been deferred by Arrival Rate Meter rules
User Name	User who submitted the query

Rule Information	Description
Rule Name	Name of rule causing request to defer
Rule Type	TASM type of rule causing request to defer
Overridable	Indicates if the Teradata DBA can abort or release the query

Delay Tab for Teradata System

The **Delay** tab displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Details View for Teradata QueryGrid

When you select a Teradata QueryGrid session from a database session or a Teradata QueryGrid manager in the portlet's summary view, the details for the query are displayed.

The top portion of the screen displays a graphical representation of the progress and status of the query, with icons representing the initiating system, the phase of the current operation being performed, and the target system. The initiating system is the data source from which the query is generated. The target system is the system against which the query is being run. Teradata QueryGrid queries go through a metadata and execution phase before they are completed. Teradata QueryGrid queries may also access data on target systems or transfer data from one system to another.

The lower part of the screen displays tabs with information about the system or operation. If you arrive at the details screen from a Teradata system session, details are displayed for the target or initiator system that is relevant for the selected step. If you arrive at the details screen while viewing a list of Teradata QueryGrid manager sessions, details for the selected operation are displayed. You can click on any of the icons in the graphical representation (initiator, operator, or target) to view the metrics for the selected system or operation.

Initiating System

The data source from which the query is generated.

Target System

The system from which data is accessed or to which it is transferred.

Operation

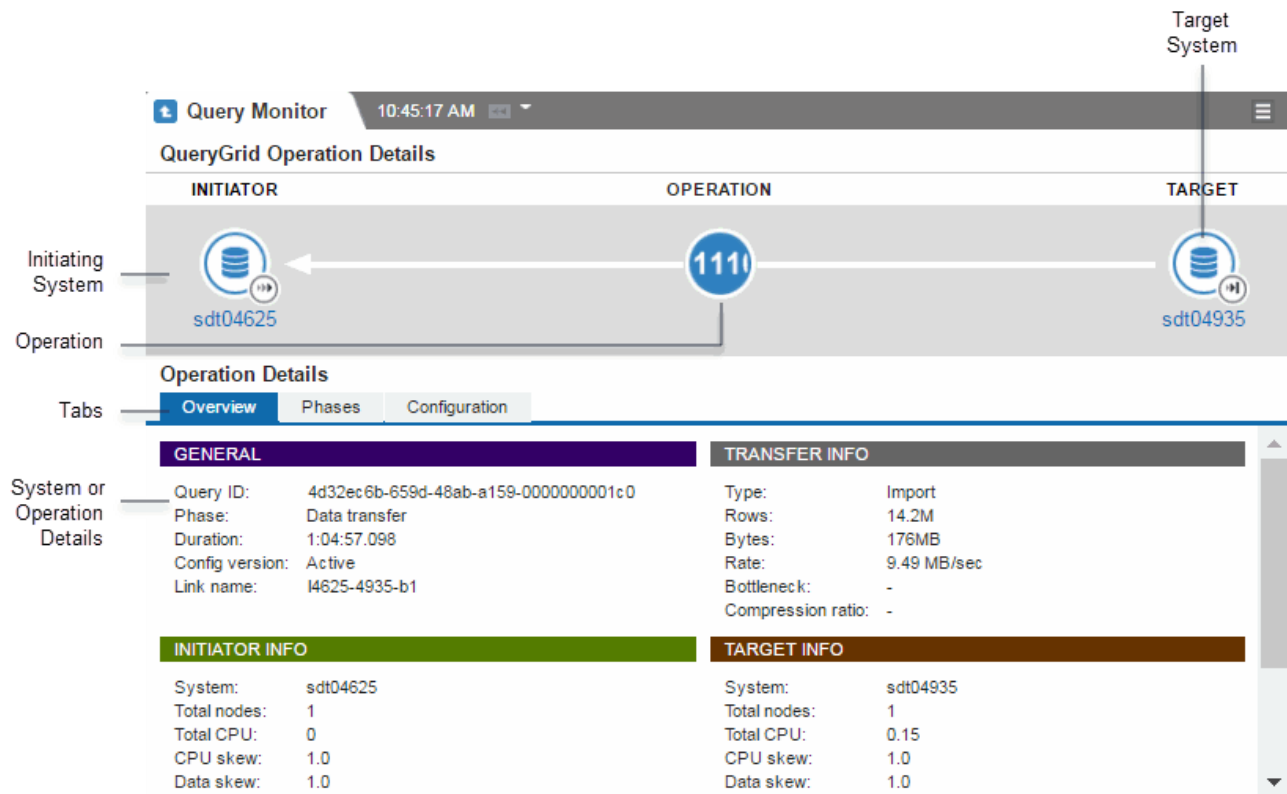
The phase of the operation being performed.

Tabs

Related metrics are organized into tabs for the system or operation selected in the top portion of the view.

System or Operation Details

The metrics related to the system or operation selected in the top portion of the view.




QueryGrid Icons

The details view for Teradata QueryGrid includes a graphical representation of the progress of the selected query. The icons here represent the initiating or target systems and the phases of the operations performed.







System Icons

The icons in the following table represent the query on the initiator and target systems. The icons inside the smaller circle represents the query state. In addition to the query states shown as following, additional possible states for the query are listed in [State Icons](#).

Icon	Name	Definition
	No query is running	No query is currently running
	Active	The query is running on the system

Icon	Name	Definition
	Idle	The query is currently idle

Operation Phase Icons

Icon	Name	Definition
	Metadata Operation	A handshake between systems to negotiate data types and validate queries.
	Execution Operation	Query processing prior to transferring the data.
	Data transfer	The transferring of data from one system to another.
	Completed	Query has successfully finished processing.
	Failed	Query did not finish processing successfully. [QueryGrid systems ≥ 2.10] To download the support bundle, click Download Support Query Bundle .
	Unknown	Query processing status is not currently available.

Initiator and Target Metrics

When an initiator or target system icon is selected in the details view for Teradata QueryGrid, two tabs display metrics for the system: an **Overview** tab and a **SQL** tab.

Overview Tab for Teradata System

The **Overview Tab** displays information for a QueryGrid query that was initiated by or is a target of a Teradata system.

QueryGrid Information

Metric	Description	Type
System	Name of the initiating or target system	String

Metric	Description	Type
Connector	Point from which the query originates or ends	String
Software version	The version of the connector	String
Network	The logical network used for this side of the data transfer	String

Session Information

Metric	Description	Type
User	User who submitted the query	String
Query ID	ID of the query	String
Session Number	Number of the session	String
Request Number	Number of the request	String
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session	String


Query Information

Metric	Description	Type
State	State of the query	String
Duration	Amount of time the query has been running (HH:MM:SS)	Number
Workload	Workload in which the query is running	String
Request CPU	Total request CPU (seconds)	Number
Request I/O	Total request I/O count	Number

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.

Operations Metrics

When an operation phase icon is selected in the details view for Teradata QueryGrid, three tabs display metrics for the operation: an **Overview** tab, a **Phases** tab, and a **Configuration** tab.

Overview Tab

The **Overview** tab displays general information about the selected operation for a Teradata QueryGrid query session.

General

Metric	Description	Type
Query ID	ID of the query	String
Phase	Phase of the operation	String
Duration	Duration of the QueryGrid operation	Number
Config version	Requested version of the QueryGrid operation	Number
Link Name	Name of the link used for this QueryGrid operation	String

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Total nodes	Number of nodes on the initiating system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Transfer Information

Metric	Description	Type
Type	Type of transfer performed	String
Rows	Number of rows transferred	Number
Bytes	Number of bytes transferred	Number
Rate	Rate that data is transferred for the QueryGrid operation	Number
Bottleneck	Limiting factor for the QueryGrid operation	String
Compression ratio	Uncompressed size divided by the compressed size	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Total nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Total Nodes	Number of nodes on the target system	Number
Total CPU	Total CPU usage time consumed	Number
CPU skew	CPU skew for the current query	Number
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Phases Tab

The **Phases** tab is available when you click an icon that represents a Teradata QueryGrid query operation. It contains tabs that provide information about the metadata, execution, and data transfer phases of the query session. The **Phases** tab support up to five panes depending on the details associated with the operation, for example, Metadata1, Metadata2, Metadata3, Execution1, Execution2, Data Transfer.

Metadata Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	The node that initiated the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number

Metric	Description	Type
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Execution Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Initiating node	Node from which the query initiated	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Driver node	The node that handled the metadata or execution operation	String
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number

Bridge Information

Note:

This information is not available for Fabric versions 2.15 and later.

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number

Metric	Description	Type
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Node	Number of nodes using the bridge	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number

Data Transfer Tab

Initiator Information

Metric	Description	Type
System	Name of the initiating system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	String
Data skew	Ratio between the size of the data on the maximum node versus the average node	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Target Information

Metric	Description	Type
System	Name of the target system	String
Start time	Start time of the phase	Number
Duration	Amount of time the phase has been running (HH:MM:SS)	Number
Total nodes	Total number of nodes	Number
CPU	Total CPU consumed by the fabric for this phase (seconds)	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number

Metric	Description	Type
CPU skew node	Node that had the max CPU	String
Data skew	Node that had the max data	Number
Data skew node	Node that had the maximum amount of data	String
Bytes transfered	Number of bytes transferred	Number

Bridge Information

Metric	Description	Type
Bridge	Name of the bridge	String
Start time	Start time of the bridge	Number
Duration	Amount of elapsed time the bridge has been active for this session, displayed as d+hh:mm:ss	Number
Total nodes	Number of nodes	Number
CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
CPU skew	Ratio between the CPU consumed on the most active node versus the average node	Number
CPU skew node	Node that had the maximum CPU	Number
Data skew node	Node that had the maximum amount of data	Number
Bytes sent	Number of bytes sent	Number
Bytes received	Number of bytes received	Number
Compression ratio	Whether or not compression is enabled	String

Note:

[Viewpoint V17.10.02.00 and Later] the bridge information for the QueryGrid Manager 2.18 and higher displays upto four bridges.

Configuration Tab

The **Configuration** tab contains provides configuration information about the selected operation for a QueryGrid query session.

Fabric

Metric	Description	Type
Name	Name of the fabric used in Teradata QueryGrid	String
Port	Port on which the fabric runs	String
Software version	Version of software that the fabric runs	String

Initiating Connector

Metric	Description	Type
System	Name of the initiating system	String
Connector	Point from which the query originated	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Communication Policy

Metric	Description	Type
Name	Name given to this communication policy	String
Encryption	Whether or not data encryption is enabled	String
Compression	Whether or not compression is enabled	String

Target Connector

Metric	Description	Type
System	Name of the target system	String
Connector	Destination point for the query	String
Software version	Version of the connector software	String
Network	Network that determines the interfaces to use for data transfer	String

Errors Tab

The **Errors** tab provides information about Teradata QueryGrid queries that had errors. It displays only if there are errors.

Metric	Description	Type
Time	Time the error occurred	Number
Phase	Phase when the error occurred	String
Message	The error message	String
Node	The hostname of the node that generated the error	String
Connector	The name of the connector that resulted in the error	String
System	Name of the system	String
Origin	Indicates whether the error occurred on the initiator or the target of the Teradata QueryGrid operation	String

Managing Queries and Sessions for Teradata System

You can manage queries and sessions to improve workload performance for Teradata systems.

Abort

Abort the selected query or session

Change Priority

Change the priority of the selected query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change priorities or workloads, or release queries. If you log off, close a portlet, or open a new portlet instance, you must log in again.

Note:

Change Workload is available only if Teradata Active System Management (TASM) is enabled. If TASM is disabled, **Change Priority** is available. If you do not see **Change Workload** or **Change Priority** in the list, the system you are monitoring does not support these features or you do not have permission to use them. If the query you are monitoring is delayed, only **Release Query** is available.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. Click the row of the query you want to abort.
2. Click ▾ to the right of the session number and select **Abort**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
6. Click **Next**.
7. Click **Next** to confirm your selection.
8. Click **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▾ to the right of the session number and select **Change Priority**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.

- Type an account string.
6. [Optional] Select the check box to use the account string as the default for the selected session. The priority will be changed for the selected query and all subsequent queries in the current session.
 7. Click **Next**.
 8. Click **Next** to confirm your selection.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Workload**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Select a different workload from the list and click **Next**.
5. Click **Next** to confirm your selection.
6. Click **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Release Query**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Click **Next** to confirm your selection.

Changing Multiple Sessions

For Teradata systems, you can change system resources for multiple sessions, users, or account strings in the following locations:

- All sessions table from the summary view
 - My Criteria sessions table from the summary view
 - User sessions table from the details view
 - Account strings table from the details view
 - Utility table from the details view
1. Select **By Session** or **By Session (Exclude Utilities)** in the selection menu and one of the following choices:
 - All
 - My Criteria
 2. [Optional] Access user or account string tables from the details view.
 - a. Click a session.
 - b. Click the **User** or **Account** link under **SESSION INFO**.
 3. Select a session action from the **Table Actions** list.
Check boxes appear next to the sessions.
 4. Do one of the following:
 - Select the first check box located in the column heading to select all sessions.
 - Select the check boxes for specific sessions.
 5. Click **Next**.
 6. Log in to Analytics Database, if prompted.
 7. Follow the instructions on subsequent screens.

Managing Sessions for Aster System

You can manage sessions to improve workload performance for Aster systems.

Abort

Abort the selected process.

Aborting a Process

For Teradata Aster systems, you can abort a non-idle process that is blocking other processes or consuming too many resources.

1. Select **By Session** in the selection menu.
2. Do one of the following:

Option	Description
Abort a single process	<ol style="list-style-type: none"> a. Click the row of the session ID you want to abort. b. Click ▾ to the right of the session number and select Abort Process. c. Click Next.
Abort multiple processes	<ol style="list-style-type: none"> a. From the Table Actions list, select Abort Process. Check boxes appear next to the session ID. b. Do one of the following: <ul style="list-style-type: none"> • Select the first check box located in the column heading to select all processes. • Select the check boxes for specific processes.

3. Click **Next**.
4. Log in to Teradata Aster, if prompted.
5. Click **Next** to confirm you want to abort the processes.
A confirmation appears indicating the abort process was successfully submitted.
6. Click **OK**.

Settings View

The **Settings** view provides the following tabs that allow you to change display options and customize the way query and session information displays in the portlet.

Display

For Teradata systems, choose whether or not to display the top sessions graph and whether to format the SQL to display in the **SQL** tab.

Criteria


For Teradata systems, set the criteria values used to filter sessions in the **My Criteria** view. If no thresholds are specified, all sessions display in the view. After setting the criteria value and enabling a threshold, any query exceeding the criteria value appears in the sessions table in the **My Criteria** view.

Skew

For Teradata systems, set the CPU skew threshold and CPU limit for sessions that are displayed in the **By Vproc > By Skewed AMP** view.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.


Selecting Display Options

You can select a SQL format and display or hide the **Top Sessions Graph** in the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Display** tab.
3. Do one of the following for Teradata systems:
 - Select the **Top Sessions Graph** check box to display the top session graph in the portlet.
 - Clear the **Top Sessions Graph** check box to hide the top sessions graph in the portlet.
4. Do one of the following for Teradata systems:
 - Select the **Format SQL** check box to display formatted SQL.
 - Clear the **Format SQL** check box to display unformatted SQL.
5. Click **OK**.

Selecting Threshold Criteria


For Teradata systems, you can specify thresholds to control which sessions appear in the **My Criteria** view.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Criteria** tab.
3. Select the check box for the threshold you want enabled.

4. Enter a threshold to influence which sessions display in the view.
5. [Optional] Select the **Save these thresholds as the highlight thresholds for the My Criteria view** check box.
Metrics exceeding at least one threshold will be highlighted in the view.
6. Click **OK**.
Sessions exceeding at least one threshold appear in the view.

Selecting Skew Settings

For Teradata systems, you can set the CPU skew threshold and CPU limit for sessions that are displayed in the **By Vproc > By Skewed AMP** view.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Skew** tab.
3. Enter threshold values for each session you want to display in the view.
4. Click **OK**.

Query Spotlight

Query Spotlight Overview

The **Query Spotlight** portlet allows you to view completed queries that exceeded the thresholds you set. You can view individual queries of interest as a line graph over time and against a variety of metrics. The **Snapshot** tab displays the query data at any point in time.

Use **Query Spotlight** to:

- View general information about a query.
- Inspect a query as a line graph using the date-range slider and time marker.
- Set a date and time in the rewind toolbar. This feature is useful for rewinding data in surrounding portlets to a point of time specified within a query. The **Query Spotlight** portlet can set time and date but it does not respond to changes made in the rewind toolbar.

The **Settings** view allows you to choose the query criteria that filter which metrics appear in the views.

You can monitor multiple systems by opening additional instances of the Query Spotlight portlet. Rewind control is given over to the last instance of **Query Spotlight** to use the time marker.

Query Spotlight View

The **Query Spotlight** view displays summary information about the completed queries so you can monitor and locate issues.

System Selection

Shows the name of the system currently selected and allows you to select a different system.

Time Selection

Shows the time frame of the information displayed in the summary table and allows you to select a different time frame.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information about each query. The current view is configured in the **Configure Columns** dialog box and **Settings** view. Click on a row in the table to see details.

☐ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

System Selection: Dev1

Time Selection: Last 1 week (2/12/14 4:30 PM - 2/19/14 4:30 PM)

Table Actions: [Menu Icon]

USERNAME	IMPACT CPU	REQ CPU	START
VIEWPOINT	0.288	0.676	2/19/14 3:00:28 PM
MIXWL_STATS_REP	0.288	27.18	2/19/14 2:40:16 PM
MIXWL_INV_REP5	0.288	0.448	2/19/14 1:20:05 PM
VIEWPOINT	0.36	1.46	2/19/14 7:10:40 AM
VIEWPOINT	0.432	1.8	2/19/14 7:10:40 AM
MIXWL_INV_REP5	0.432	5.692	2/18/14 11:27:12 AM
MIXWL_OLAP2	0.648	1.044	2/19/14 2:16:17 AM
VIEWPOINT	1.296	1.732	2/19/14 4:21:50 AM
MIXWL_RET_REP5	1.368	32.37	2/18/14 7:42:10 PM
VIEWPOINT	2.016	2.348	2/19/14 7:13:41 AM

18 rows total

Changing the System or Time Frame

You can choose a different system or time frame to view queries.

- Select ▼ in the portlet frame and select a system.
- Select a time frame:
 1. Select ▼ next to the time frame selection and select a time frame.
Hours, week, and month are all based on the current time. **Custom** is a fixed time period.
 2. If you chose **Custom**, enter a date and time frame, and click **OK**.

Query Spotlight View Metrics

Metrics available for monitoring and display are as following. For a list of metric sources for some of the metrics listed as following, see [Metric Sources Overview](#).

Metric	Description	Type
Duration	How long the query has been running	Number
End	Time that the query stopped running on Analytics Database	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number

Metric	Description	Type
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String
Req CPU	CPU seconds needed to run the query	Number
Req CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Req I/O	Disk I/Os performed to run the query	Number
Req I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Req Number	Number of the specific request	Number
Session ID	Unique session identifier	Number
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Spool Skew	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Percent
Start	Time that the query started running on Analytics Database	Number
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time	Number
Username	Name of the user who submitted the query	String

Details View

The details view displays a graph, statistics, and information about a selected query. Choose a metric from the metrics list to display one query as a line graph.

Metrics List

The list of metrics that can be graphed are as follows: CPU Skew, Impact CPU, I/O Skew, PJI, CPU Delta, I/O Delta, Spool, and Unnecessary I/O.

Date-Range Slider

A slider used to control the amount of data displayed in the graph. The slider appears for longer queries.

The darker-shaded area between the slider handles represents the current date range displayed in the graph.

Metrics Graph

Charts selected metric for a query using Explain steps. As you hover over the graph, information balloons display query values.

Time Marker

A marker that pinpoints information about a query at a particular point in time. Used to rewind time in the rewind toolbar when **Control Rewind** is enabled. Determines the data displayed in the **Snapshot** and **Blocked By** tabs.

Job Steps

Numbers that denote the Explain step that occurs between the dashed lines in the graph. Job steps are shown just after the graph area. Overlapping steps or steps-within-steps are not shown.

When there is not enough room to display individual steps, the numbers are grouped together. For example, you may see 1, 2, 3, 4, 8. Step 8 represents steps 5, 6, 7, and 8.

Tabs

Provides important query details on the **Overview** , **SQL** , **Explain** , **Blocked By** , and **Snapshot** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

< > Previous/Next

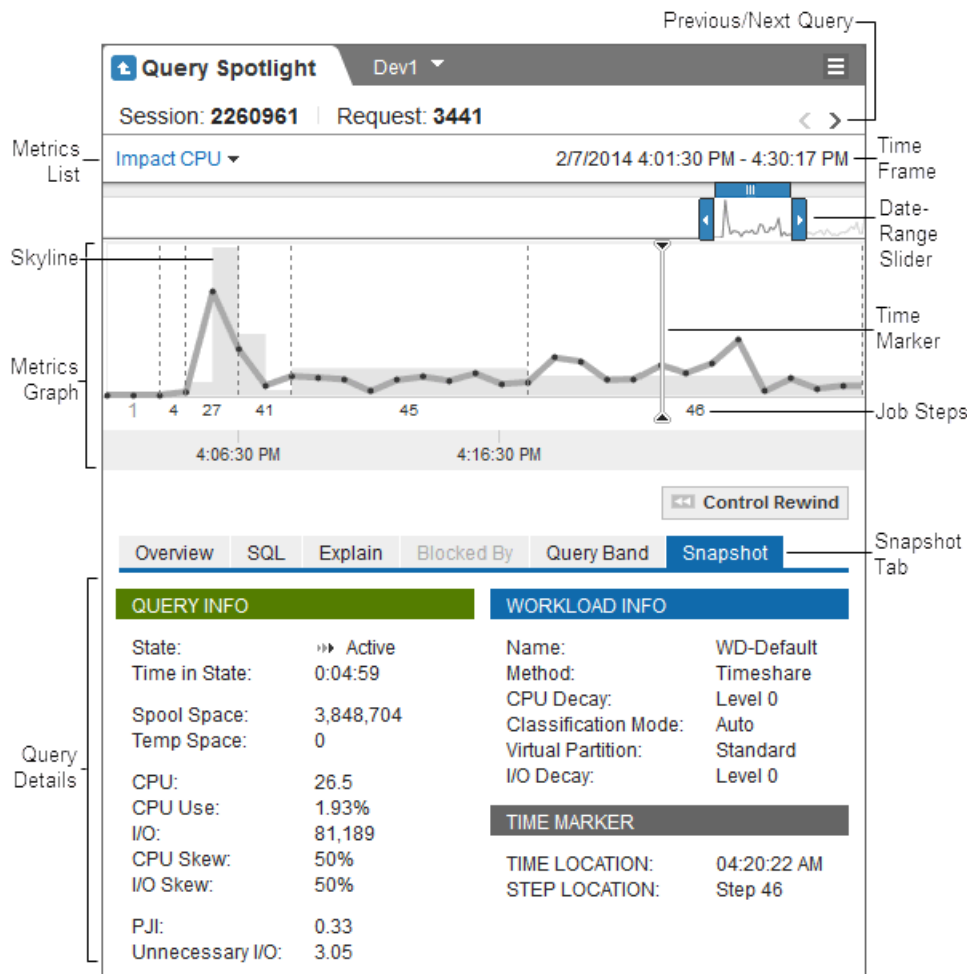
Allows you to move through sessions without returning to the summary table.

Time Frame

Displays the date and time duration of the query.

Skyline

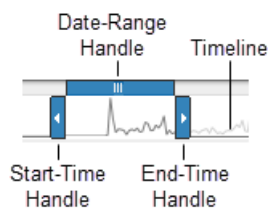
Skyline blocks represent a metric over the duration of the designated query step. For **Impact CPU** this metric is the average *Delta CPU*. For **Unnecessary I/O** it is the average *Delta IO*. No other metrics display skyline data.



Adjusting the Date Range

The date-range slider shows the most recent part of the query. Time and date ranges are displayed in the time frame before the slider on the right.

1. Hover over the date-range slider to activate the start-time and end-time handles.



2. Drag the start-time handle until the intended start time is displayed in the time frame.
3. Drag the end-time handle until the intended end time is displayed in the time frame.

4. [Optional] Drag the date-range handle left or right to move the start-time and end-time handles in unison.

Selecting a Metric to Graph

You can select different metrics to graph. The metrics contain historical details you can use to investigate queries.

1. Click ▼ next to the metric name.
2. Select a metric from the list.

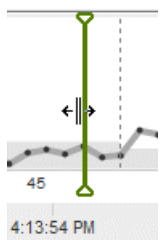
Query Spotlight and Rewind

The **Query Spotlight** portlet cannot rewind and does not use the rewind toolbar. The **Control Rewind** button in the **Query Spotlight** portlet is a toggle that enables portal rewind for portlets that support the rewind function on the same page. When rewind is enabled, the time marker in **Query Spotlight** sets the time and date in the rewind toolbar and controls the time and date for data in the other portlets. This allows you to investigate how metrics compare across portlets during the same time frame.

When **Control Rewind** is toggled off, the time marker no longer controls the date and time for data in the other portlets and portal rewind is controlled from the rewind toolbar.

Enabling and Using Rewind Control

The time marker is a vertical line in the graph with small triangles at the top and bottom. Its position represents a point in time and controls rewind for the portlets that support the rewind function. The time marker appears green when **Control Rewind** is enabled and you interact with it.



1. From the **Details View**, click **Control Rewind** located under the graph in the middle of the page. The **Control Rewind** button turns green and the rewind toolbar appears.
2. Click or drag the time marker to a specific point in time.
3. [Optional] Drag the time marker to another point in time.

Disabling Rewind Control

1. Click **Control Rewind**.

The **Control Rewind** button turns gray. The time marker no longer controls the time and date in the rewind toolbar, and portal rewind is controlled from the rewind toolbar.

Overview Tab

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted.


Metric	Description	Type
Total Duration	How long the query has run, displayed as <i>hh:mm:ss</i>	Number
User	Name of the user that submitted the query	Character
Account	Account of the user that submitted the query	Character
Req Number	Number of I/O operations the query performed	Number
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.	Number
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.	Percentage
Req CPU	Total CPU usage time, in seconds, consumed by the query	Number
Impact CPU	(Teradata Database 15.10 and later) CPU impact on the system based on the highest utilized AMP. Calculated at the request level.	Number
Req I/O	Number of I/O operations the query performed	Number
Req CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percentage
Req I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percentage
Req Max Spool	Amount of spool space the query requires	Number
Source	Source details, such as application name, IP address, and host user name	Character

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure



displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.





Exporting SQL

1. From the **SQL** tab, select .
2. Save the file to a location.
The SQL statement is exported as a .txt file.


Explain Tab

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status.

If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	 - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Exporting Explain Data

1. From the Explain tab, click .
2. Save the file to a location.
The Explain data is exported as a .csv file.

Blocked By Tab

The **Blocked By** tab displays details about other queries that are blocking the selected query. This information is read-only. The tab is available only when the selected query is blocked.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
	3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock		
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Snapshot Tab

The **Snapshot** tab displays detailed query data for any point in time. Metrics that exceed defined thresholds are highlighted.

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>

Query Information	Description
Spool Space	Amount of spool space the query is using
Temp Space	Amount of temp space the query is using
CPU	Total CPU seconds needed to run the query, in seconds
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
I/O	Total number of disk I/Os performed
CPU Skew	CPU skew during the last sample
I/O Skew	I/O skew during the last sample
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload

Workload Information	Description
	This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Time Marker	Description	Type
Time Location	Temporal location of the time marker within the Explain step	Number
Step Location	Explain step in which the time marker is located	Number

Settings View

The **Settings** view provides the following tabs that allow you to change display options and customize the way session information displays in the **Query Spotlight** view:

Display


Choose formatted or unformatted SQL to display on the **SQL** tab of the details view.

Criteria

Specify thresholds to control which sessions appear in the **Query Spotlight** view. After enabling the check boxes, sessions exceeding at least one threshold appear in the view.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:


Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.

4. Click **OK**.


Selecting Display Options

The **Display** tab in the **Settings** view allows you to format the SQL that appears on the **SQL** tab of the query details view.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Display** tab.
3. Do one of the following for Teradata systems:
 - Select the **Format SQL** check box to display formatted SQL.
 - Clear the **Format SQL** check box to display unformatted SQL.
4. Click **OK**.

Selecting Threshold Criteria

You can specify thresholds to control which sessions appear in the **Query Spotlight** view. By default, the portlet displays sessions that exceed the default thresholds configured in the **Settings** view.

1. Click  in the portlet frame and select **Settings**.
 2. Click the **Criteria** tab.
 3. Select the check box for the threshold you want enabled.
 4. Enter a threshold to influence which sessions display in the view.
 5. Click **OK**.
- Sessions exceeding at least one threshold appear in the view.

Remote Console

Remote Console Overview

The **Remote Console** portlet allows you to run many of the Vantage console utilities remotely from within the Teradata Viewpoint portal.

Using this portlet, you can:

- Select or search for a system.
- Select or search for a utility.
- Enter console utility commands.
- Display responses from the commands.

Teradata field engineers, database operators, system administrators, and system programmers use Vantage utilities to administer, configure, monitor, and diagnose issues with Analytics Database.

Remote Console View

The **Remote Console** view allows you to select and run Vantage console utilities. You can run only one utility at a time for each portlet instance. Permission to use each console utility is granted to individual users based on their role and portlet access. If a message appears in the console denying access to a utility, you are not authorized to use that utility.

The following list describes the features in this view:

Selection Menus

Shows the name of the system and utility currently displayed. Click to select another system or utility.

Clear Console

Clears the content from the console.

Disconnect

Disconnects the utility from Analytics Database.

Console

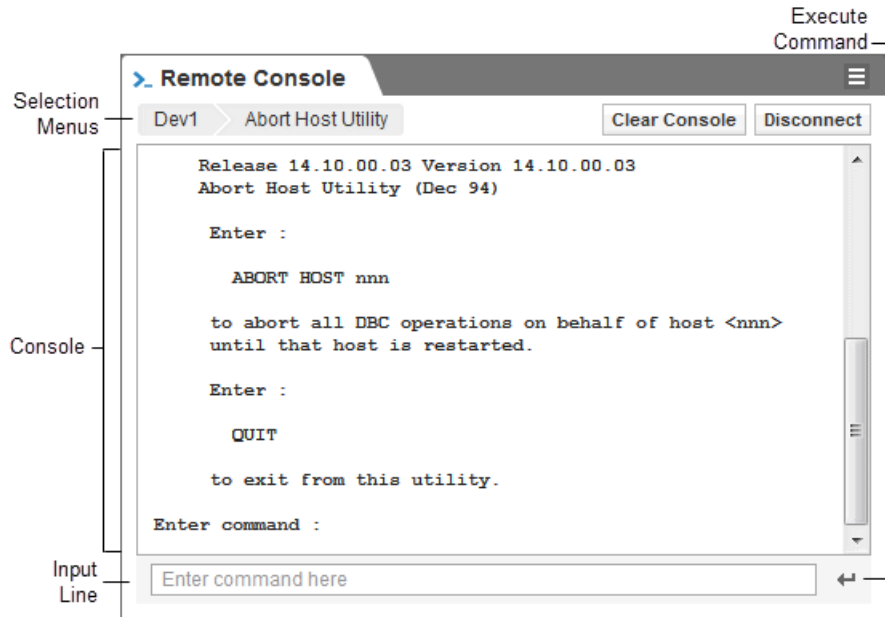
Displays responses from the commands.

Input Line

Use to enter commands.

Execute Command

Click to execute the command.



Remote Console Menus and Toolbar

The menus and toolbar allow you to specify the display attributes of the **Remote Console** view.



Use the selection menus to:

- Select or search for a system.
- Select or search for a utility.

Use **Clear Console** to clear the content from the console.

Use **Disconnect** to disconnect from the utility and Analytics Database.

Running a Utility

Note:

After 30 minutes of inactivity in the portlet, the Analytics Database connection times out.

1. Choose a system from the selection menu.

Option	Description
Select	Select the system from the list.
Search	Type the system name in the search box. As you type, the list shows the systems that match your filter criteria. Search is not case-sensitive. Note: Search is available only if the number of systems exceeds the length of the list.

2. Select >.
3. Choose a utility from the selection menu.

Option	Description
Select	Select the utility from the list.
Search	Type the utility name in the search box. As you type, the list shows the systems that match your filter criteria. Search is not case-sensitive. Note: Search is available only if the number of systems exceeds the length of the list.

4. Select >.
The name of the utility appears in the console.
5. Type a command on the input line.

Note:

The commands are temporarily saved and available as long as you are running the utility. After disconnecting from the utility, the commands are cleared.

6. Click ← to execute the command.
The command results appear in the console.
7. [Optional] From the input line, use your keyboard up and down arrows to access previous commands.

Changing a System or Utility

You can change the currently selected system, utility, or both.

- Change the system:
 1. Click the currently selected system name and replace it with the new system name.
 2. Select >.
 3. [Optional] Select a new utility.
 4. Select >.

- Change the utility:
 1. Click the currently selected utility and replace it with the new utility.
 2. Select **>**.

Clearing the Console

1. On the toolbar, click **Clear Console**.
The contents of the console no longer appear.

Disconnecting from a Utility

You can disconnect from a utility and Analytics Database.

1. On the toolbar, click **Disconnect**.
A confirmation message appears.

If you select **Disconnect**, the input line no longer accepts input and the portlet no longer responds to commands.

The Analytics Database connection closes when you remove the portlet from the portal page or click **Disconnect**. If you close the browser without clicking **Disconnect**, the Analytics Database connection times out in 30 minutes.

Console Utilities

Teradata field engineers, database operators, system administrators, and system programmers use Vantage utilities to administer, configure, monitor, and diagnose issues with Analytics Database.

Console Utilities

Information about available console utilities is listed in the following table. You can access only the following subset of Vantage utilities.

Utility	Description
Abort Host	Cancel all outstanding transactions running on a failed host until the system restarts the host.
Check Table	Check for inconsistencies between internal data structures, such as table headers, row identifiers, and secondary indexes.
Configure	Define AMPs, PEs, and hosts, and describe their interrelationships for Analytics Database.
DBS Control	Display and modify the DBS Control Record fields.
Ferret	Define the scope of an action, such as a range of tables or selected vprocs, display the parameters and scope of the action, and perform the action. The action is to

Utility	Description
	either move data to reconfigure data blocks and cylinders, or display disk space and cylinder free space percent in use.
Gateway Global	Monitor and control the sessions of Analytics Database LAN-connected users.
Lock Display	Display a snapshot capture of all real-time database locks and their associated currently running sessions.
Operator Console	Run supervisor commands to manage the programs that perform Analytics Database operations.
Priority Scheduler	Create, modify, and monitor Analytics Database process prioritization parameters. All processes have an assigned priority based on their Analytics Database session. Priority Scheduler allocates CPU and I/O resources based on the priority.
Query Configuration	Display the current Analytics Database configuration, including the node, AMP, and PE identification and status.
Query Session	Monitor the state of selected Analytics Database sessions on selected logical host IDs.
Recovery Manager	Display information used to monitor progress of a database recovery.
Show Locks	Display locks placed by archive and recovery, and by table rebuild operations on databases and tables.
Teradata DWM Dump	Display information about active Teradata Dynamic Workload Manager (Teradata DWM) rules on a Teradata system.
Update Space	Recalculates the permanent, temporary, or spool space used by a single database or by all databases in a system.
Vproc Manager	Manage the vprocs including obtaining the status of specified vprocs, initializing vprocs, forcing a vproc to restart, and forcing a database restart.

Information on how to use a majority of these utilities is located in *Teradata® Database Utilities*.

Information on how to use Teradata DWM Dump is in *Teradata® Dynamic Workload Manager User Guide*.

Information on how to use Operator Console is in *Graphical User Interfaces: Database Window and Teradata® MultiTool*.


You can access these documents at: <https://docs.teradata.com/>.

Settings View

The **Settings** view allows you to change the console text and background colors.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Setting the Display Background

Use the **Settings** view to change the console text and the display background to light or dark.

1. Click  in the portlet frame and select **Settings**.
2. Change the color of the text and background by selecting one of the following options:
 - **Dark text on a light background**
 - **Light text on a dark background**
3. Click **OK**.

Space Usage

Space Usage Overview

The **Space Usage** portlet allows you to monitor database disk space usage for Teradata, Aster, and Hadoop systems. The views and the options vary depending on the system you choose.

For Teradata systems, you can monitor database data, such as the perm space, temp space, and spool space, and reallocate permanent disk space. The details view displays permanent space usage over time in information balloons. You can reallocate permanent disk space from one database to another using either the **Space Usage** view or the details view.

For Aster systems, you can monitor database data, such as the table count and current space usage. The details view displays the space usage, schema and level of skew.

For Hadoop systems, you can monitor node data, such as the total space, current space, percent in use and available space.

The **Settings** view allows you to customize the display for summary and details views and set thresholds for filters. Settings are available only for Teradata systems.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is not available. You cannot enter alternate credentials.

Space Usage View

The **Space Usage** view displays information about databases and nodes on the selected system. For Teradata systems, the view displays a summary of perm, spool, and temp space for each database on the selected system. Use this information to see if you need to reallocate disk space to maximize space usage.

For Aster systems, the view displays the number of tables and amount of space used for each database on the selected system.

For Hadoop systems, the view displays the host names, total space, current space, percent in use and available space for each node on the selected system.

The following list describes the features in this view:

Selection Menus

Displays the system and space usage report. You can choose a different system and report.

Threshold Filter Bar

Displays the number of databases on the system and the number of databases whose space usage exceeds set thresholds. This feature is available only for Teradata systems. You can set the thresholds in the **Settings** view.

Filters

Shows only rows that match your filter criteria.

Table Actions

This feature is available only for Teradata systems.

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are applied, only filtered data is exported

For more information, see [Summary Table Controls](#).

Add Space

Reallocates permanent disk space from one database to another. This feature is available only for Teradata systems.

Summary Table

Displays summary information for the space usage report. The table is configured in the **Configure Columns** dialog box. Click a row in the table to see details.

Following is an example of the **Space Usage** view that appears when you select a Teradata system from the selection menu.

The screenshot shows the 'Space Usage' window for 'Dev1' with the view set to 'By Database'. The time is 12:38:56 PM. The 'Threshold Filter Bar' at the top shows a threshold of 45 and three filters: '> 97% Perm', '> 85% Spool', and '> 95% Temp'. The 'Table Actions' menu is visible in the top right. The 'Filters' section includes input fields for NAME, MAX PERM, MAX SPOOL, MAX TEMP, and PEA. The 'Summary Table' lists the following data:

NAME	MAX PERM	MAX SPOOL	MAX TEMP	PEA
All	0	0	0	0
Crashdumps	978M	66G	66G	0
DBC	41.7G	66G	66G	1.2
Default	0	0	0	0
EXTUSER	0	66G	66G	0
My_Business	4.66G	14.9G	66G	18
Add Space	23.8M	954M	66G	11
	0	0	0	0
SQLJ	572M	66G	66G	40
SYSLIB	38.1M	66G	66G	98
SYSSPATIAL	153M	66G	66G	3.3
labs_LG3	9.54M	954M	66G	4K
labs_LG4	4.77M	954M	66G	4K

At the bottom, it indicates '45 rows total'.

Following is an example of the **Space Usage** view that appears when you select an Aster system from the selection menu. The Hadoop system view shares the same features as the Aster system.

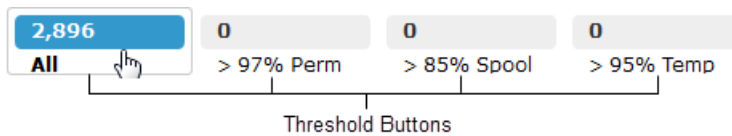
The screenshot shows the 'Space Usage' window for 'Aster44' with the view set to 'By Node'. The time is 12:43:11 PM. The table displays the following data:

NAME	TOTAL SPACE	CURRENT SP...	% IN USE	AVAILABL
10.80.44.20	20.4K	551.7	2.7	19.9K
10.80.44.21	20.4K	209.4	1.03	20.2K
10.80.44.22	20.4K	209.3	1.02	20.2K

At the bottom, it indicates '3 rows total'.

Threshold Filter Bar

For Teradata systems, the threshold filter bar allows you to view databases that exceed thresholds for space usage.



In the **By Database** report for Teradata systems, threshold buttons display the number of databases that use space in all categories and that exceed set thresholds. Click any threshold button to filter for the selected category. Define the thresholds in the **Settings** view.

In Teradata systems, these states are available:

All

List of databases on the system

Perm

Number of databases using permanent disk space in excess of the set threshold

Spool

Number of databases using temporary spool space in excess of the set threshold

Temp

Number of databases using temporary space in excess of the set threshold

Selecting a System and Report

You can use selection menus before the threshold filter bar to choose a system and report, starting with the highest level menu on the far left.

1. In the selection menu, click the currently selected system name to display a list of available systems.
 2. Select a system.
 3. Click **>**.
 4. Select a report from the **Select Report** dialog box.
 5. Click **>**.
- The report appears in the summary table.

Selection Menu Choices

The **Space Usage** portlet displays data for one system at a time. The selection menu allows you to choose the type of data that appears in the **Space Usage** view and in the details view.

Teradata System

Selection Menu Choices	Description
By Database	Displays details by database for permanent disk space, such as the total space being used and the amount of temporary spool space being used.
By Vproc	Displays details by vproc for permanent disk space, such as the total space being used and the amount of available space.

Aster System

Selection Menu Choices	Description
By Database	Displays details for the database, such as the name, table count, and space in use.
By Node	Displays details for each worker node, such as the percentage of space being used and the amount of available space.

Hadoop System

Selection Menu Choice	Description
By Node	Displays details for each node, such as the percentage of space being used and the amount of available space.

Adding Space

You can reallocate permanent disk space from the reserve Perm pool to the currently selected database. This feature is available only for Teradata systems. Permission to reallocate disk space is granted to individual users based on their Teradata Viewpoint role and portlet permissions.

- To reallocate space, do one of the following:
 - Click ☐ next to a row and click **Add Space**.
 - From the Space Usage details view, click **Add Space**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

2. Log in to the selected Analytics Database, if prompted.
3. Click **Next**.
4. From the **Reserve Perm pool** list, select the database from which you want to reallocate space.
5. For **Perm to transfer**, enter the number and select the units for the amount of space you want to add to the target database.
The amount of free Perm and max Perm after the allocation are displayed under the target database name.
6. Click **Add Space**.
A confirmation message appears.
7. Click **Transfer**.

Space Usage Metrics

Availability of metrics depends on the system and selection menu options you choose.

Teradata System Metrics

By Database

The selection menu choice **By Database** makes the following metrics available.

Column	Description	Type
Current Perm	<ul style="list-style-type: none"> In the Space Usage view: Amount of total permanent disk space the database is currently using In the details view: Amount of total permanent disk space the object is currently using 	Number
Current Perm %	Percentage of the total permanent disk space the database is currently using (Current Perm divided by Max Perm)	Percent
Current Perm Skew %	A measure of perm space data distribution across all AMPs for the selected database Note: A perfectly distributed database (same amount of data on each AMP) has a skew of 0. A database which has all the data on one AMP has a skew of 100.	Percent
Current Spool	Amount of total temporary spool space the database is currently using	Number
Current Spool %	Percentage of total temporary spool space the database is currently using (Current Spool divided by Max Spool)	Percent
Current Temp	Amount of total temporary space the database is currently using	Number
Current Temp %	Percentage of total temporary space the database is currently using (Current Spool divided by Max Spool)	Percent
Database Creator	Name of the user that submitted the Create Database statement	String

Column	Description	Type
Effective Perm Space %	(Teradata Database 16.0 and later) Percentage of Perm space used on the highest used AMP, accounting for the configured skew limit	Percent
Effective Spool Space %	(Teradata Database 16.0 and later) Percentage of Spool space used on the highest used AMP, accounting for the configured skew limit	Percent
Effective Temp Space %	(Teradata Database 16.0 and later) Percentage of Temp space used on the highest used AMP, accounting for the configured skew limit	Percent
Max Perm	Total amount of permanent disk space available for the database	Number
Max Spool	Total amount of temporary spool space available for the database	Number
Max Temp	Total amount of temporary space available for the database	Number
Name	Database name	String
Parent Database	Parent database name	String
Peak Perm	<ul style="list-style-type: none"> In the Space Usage view: Greatest amount of permanent disk space the database has used In the details view: Greatest amount of permanent disk space the object has used 	Number
Peak Perm %	Highest percentage of total permanent disk space the database has used (Peak Perm divided by Max Perm)	Percent
Perm Skew Limit	Permissible skew limit percent beyond the per-AMP quota of Perm space for a given user/database	Number
Peak Spool	Greatest amount of temporary spool space the database has used	Number
Peak Spool %	Highest percentage of total temporary spool space the database has used (Peak Spool divided by Max Spool)	Percent
Peak Temp	Greatest amount of temporary space the database has used	Number
Peak Temp %	Highest percentage of total temporary space the database has used (Peak Temp divided by Max Temp)	Percent
Skew %	A measure of data distribution among the nodes (details view only)	Number
Spool Skew Limit	Permissible skew limit percent beyond the per-AMP quota of spool space for a given user/database	Number
Temp Skew Limit	Permissible skew limit percent beyond the per-AMP quota of temp space for a given user/database	Number
Unused Perm	Unused perm space of the user (Max Perm - Current Perm)	Number
Unused Perm %	Percentage of total permanent disk space currently not being used (100% minus Current Perm %)	Percent

By Vproc

The selection menu choice **By Vproc** makes the following metrics available.

Column	Description	Type
% In Use	Percentage of perm space used by the vproc (Current Perm divided by Max Perm)	Percent
Allocated Perm	(Teradata Database 16.0 and later) Total perm space currently allocated to AMPs for the user/database	Number
Available Perm	Available perm space for this vproc (Max Perm divided by Current Perm)	Number
Current Perm	Current perm space used by the user. Current perm space depends on the volume of data already stored in the database.	Number
Max Perm	(Teradata Database 15.10 and earlier) Maximum perm space used by the user. Maximum perm space depends on the volume of data to be stored in the database.	Number
Vproc ID	Vproc number	Number

Aster System Metrics

By Database

The selection menu choice **By Database** makes the following metrics available.

Column	Description	Type
Current Space	<ul style="list-style-type: none"> In the Space Usage view: Amount of space used by the database In the details view: Amount of space used by the tables 	Number
Database Name	Name of the database	String
Schema	Schema where the table is stored (details view only)	String
Skew %	A measure of data distribution across the nodes (details view only)	Percent
Table	Table name (details view only)	String
Table Count	Number of tables in the database	Number

By Node

The selection menu choice **By Node** makes the following metrics available for each worker node.

Column	Description	Type
% In Use	Percentage of space being used	Percent
Available Space	Amount of unused space on each node	Number

Column	Description	Type
Compression Ratio	Ratio between the uncompressed size of the data and the compressed size of the data	Number
Current Space	Total amount of space occupied on the node (Current User plus Replica plus System Space)	Number
Total Space	Storage capacity	Number
Name	Name of the node	String
Node ID	ID of the node	Number
Replica Space	Amount of space used by replica copies of data	Number
Size On Disk	Amount of space required to store the data	Number
System Space	Amount of space used by non-data files, such as system files	Number
Uncompressed Size	Amount of uncompressed user data	Number
User Space	Amount of space used by primary copies of data	Number

Hadoop System Metrics

The **By Node** selection menu choice makes the following metrics available for each data node on a Hadoop system (Hortonworks and Cloudera).

Column	Description	Type
% In Use	Percentage of space being used	Percent
Available Space	Amount of unused space on each node	Number
Current Space	Total amount of space occupied on the node	Number
Total Space	Storage capacity	Number
Name	Name of the node	String

Space Usage Details View

Click a database name in the **Space Usage** view to access the details view. The details view displays space usage information about the selected database and the child databases and objects.

Use the sparkline to determine if the space usage has increased or decreased over time.

The following list describes the features in this view:

Database Hierarchy

Displays the database hierarchy of the currently selected database whose details are displayed. The hierarchy refers to the parent database that contains the currently selected database, the parent of the parent, and so on. Click any database name displayed in the database hierarchy to browse to and view the details for that database.

Add Space button

Reallocates permanent disk space from one database to another. This feature is available only for Teradata systems.

▾ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Sparkline

For the selected Analytics Database, displays the perm space usage over the last three months in the sparkline, where each data point on the sparkline is the amount of space usage recorded during a single day.

For the selected Aster database, displays the current space usage.

Hover over the sparkline to display detailed information about the space usage in an information balloon.

To the right of the sparkline, view the current, peak, and maximum perm for Analytics Database and current space usage for the Aster database.

Filters

Shows only rows that match your filter criteria.

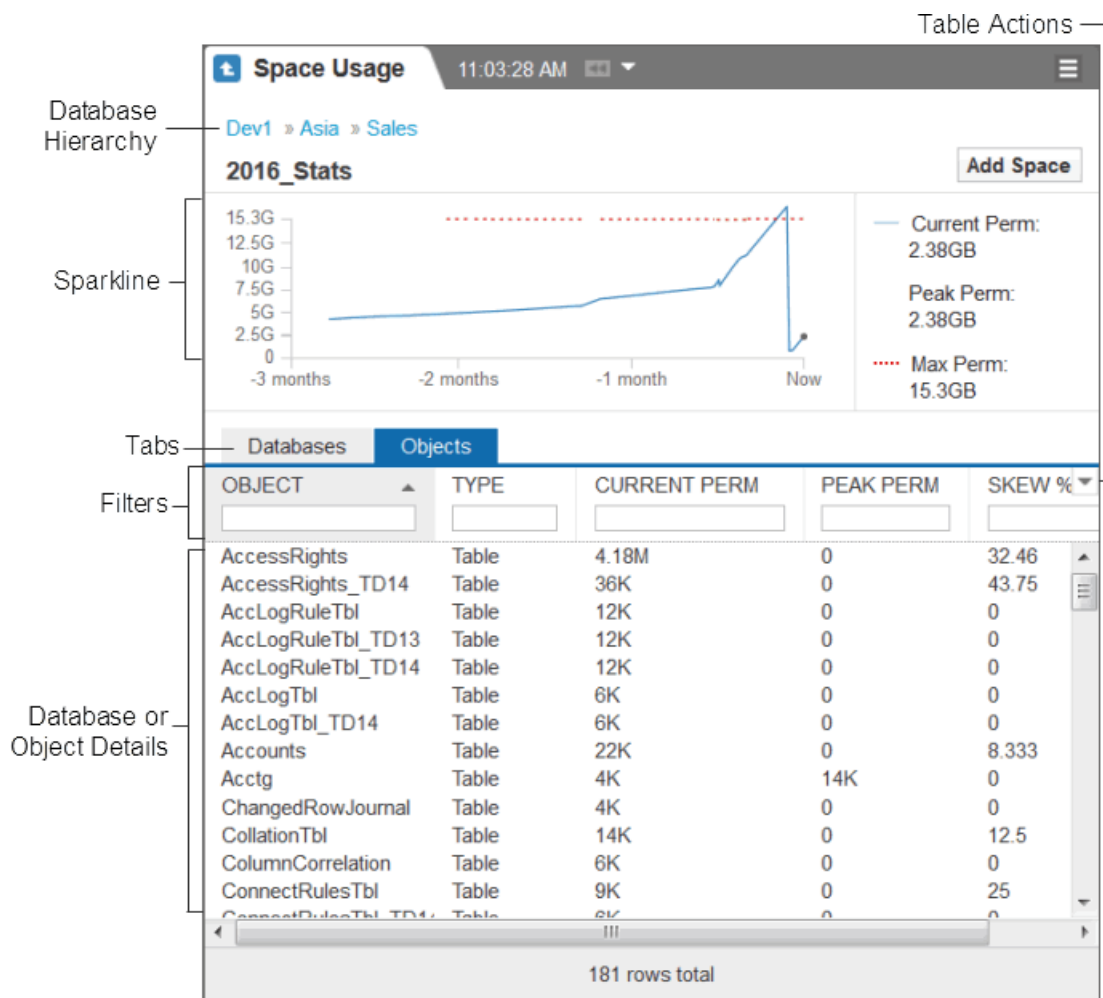
Tabs

The **Databases** tab displays information for all databases that are direct children of the currently selected database. You can click a child database in the list to view the details.

The **Objects** tab displays information for objects contained in the currently selected database.

Database or Object Details

Displays metrics about the child databases of (**Database** tab) or objects contained by (**Objects** tab) the currently selected database.

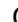


Settings View

You can access the **Settings** view from the portlet frame. Set the threshold values for the perm, spool, and temp space buttons in the toolbar of this view. Settings are applied to the **By Database** report for Teradata systems. Customize the way information is displayed in the **Space Usage** and details views.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.

Option	Description
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Setting Filter Thresholds and Display Options

For Teradata systems, you can set thresholds for the **By Database** report and customize the way information is displayed.

1. Click  in the portlet frame and select **Settings**.
2. Enter threshold values for the buttons in the **Threshold Filter Bar** for the **By Database** report.
3. Select the **Define KB, MB and GB as multiples of 1000 instead of 1024** check box to change the display in the **Space Usage** and details views.
4. Click **OK**.

SQL Scratchpad

SQL Scratchpad Overview

The **SQL Scratchpad** portlet allows you to enter SQL queries and retrieve query results from a Teradata system. You can connect to a system to run a single- or multi-statement query, view and export the query results, and save one or more queries and *pin*, or temporarily save, results. With the object browser, you can view a list of database objects and insert object names into queries.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

Query Box

The query box allows you to build queries and submit them to a Teradata system. The results are available under the query box after you run the queries.


System Menu

Displays the name of the currently connected system. Allows you to connect to or disconnect from a Teradata system, or switch to a different system.


Query Box

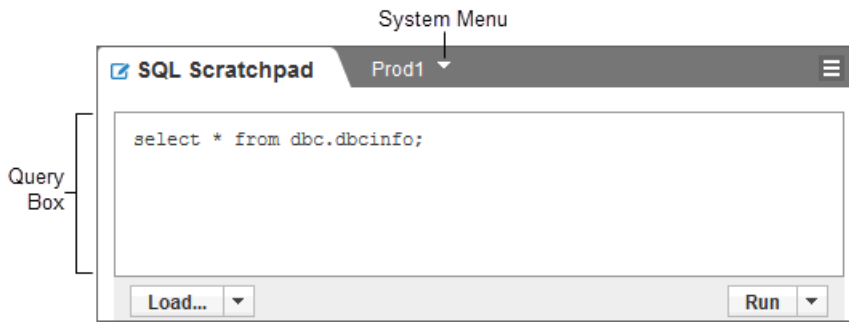
Allows you to enter and edit SQL queries.

Load

Shows the query history from the current session or saved queries from any session. The  list displays the **OBJECT BROWSER**.

Run

Runs statements in the query box and provides results in the query tabs. The  list allows you to run a selected query and to save all queries or selected queries.



Connecting to a System

You can connect to a Teradata system.

1. Select ▼ in the portlet frame and select a system.
2. In the **Connect to System** dialog box, enter the log in information.
3. Click **Connect**.

Disconnecting from a System

You can disconnect from a Teradata system to avoid idle connections.

1. Click ▼ next to the system name in the portlet frame.
2. Select **Disconnect**.
The system is disconnected, and the results tab is cleared.

Running a Query

You can enter and send multiple queries to Analytics Database.

1. Enter one or more queries into the query box.
2. Click **Run** to submit all the queries you entered against the currently connected system.

Running a Selected Query

You can run one or more adjacent queries if you do not want to run all queries that you typed in the query box.

1. In the query box, highlight one or more complete query statements you want to run.
When you select more than one adjacent query statement to run, highlight each statement, including the final semicolon.
2. Click the ▼ list next to **Run** and select **Run Selected Query**.

Query Statements

You can enter single- or multi-statement SQL queries using Data Definition Language (DDL), Data Control Language (DCL), or Data Manipulation Language (DML) statements in the query box.

DDL statements

Begin with SQL keywords, such as CREATE, ALTER, DROP, DELETE, MODIFY, HELP, SHOW, SET, REPLACE.

DCL statements

Begin with SQL keywords GRANT, REVOKE, GIVE.

DML statements

Begin with SQL keywords, such as SELECT, INSERT, UPDATE.

For single-statement queries, enter the query, optionally followed by a semicolon. For multi-statement queries, each statement must end with a semicolon.


Canceling a Long-Running Query

Occasionally, a query may take longer to execute than you expected. You can cancel the query before it completes.

1. Click **Cancel** in the results tab to stop a single, long-running query.

Saving All Queries


You can save all queries that you enter in the query box for reuse in current or future sessions.

1. Click the  list next to **Run** and select **Save All Queries**.
2. In the **Enter Query Name** dialog box, enter a name for your set of queries.
3. Click **Save**.
4. [Optional] In the query box view, click the **Load** button (not the **Load** list).
The **Saved** tab contains your set of saved queries.

Saving Selected Query

You can save one or more adjacent queries that you typed in the query box for reuse in current or future sessions.


1. In the query box, highlight one or more complete query statements you want to save. When you select more than one adjacent query statement to save, highlight each statement, including the final semicolon.

2. Click the  list next to **Run** and select **Save Selected Query**.
3. In the **Enter Query Name** dialog box, enter a name for the selection.
4. Click **Save**.
5. [Optional] In the query box view, click the **Load** button (not the **Load** list).
The **Saved** tab contains your set of saved queries.

Query Results

Query results are displayed in a table under the query box. You can view, pin, or export query results, as well as view SQL statements, Explain steps, and blocked query information.

Query Tabs

Display detailed query results in separately numbered tabs. Pinned query tabs show stored results temporarily for the current session. You can pin or unpin  query results to be stored in the current session.

Submitted Statement

Displays a portion of the SQL statement.

Results Tab

Displays a table of query results that can be exported.

SQL Tab

Displays the entire SQL statement for the selected query.

Explain Tab

Displays Explain steps for the query, including step statistics and explain text. Only appears when a query meets certain thresholds established by the Teradata Viewpoint Administrator.

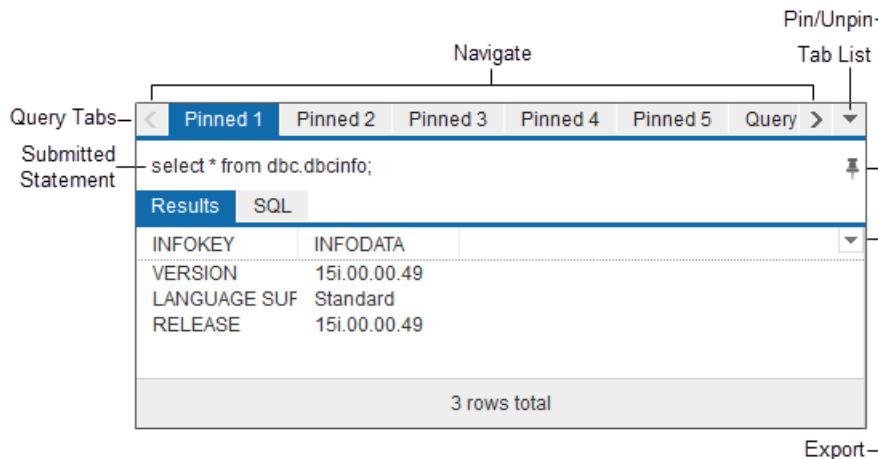
Blocked By Tab

Displays details about other queries that are blocking this query. Only appears when a query is blocked.

Export

Saves results of the currently selected **Query** tab to a .csv file.

For more information, see [Summary Table Controls](#).



Results for each statement in the query appear in a separately numbered **Query** tab. Results for a single statement query appear in the **Query 1** tab. Results for multi statement queries appear in additional tabs. For example, if you submit three statements, results appear in the **Query 1**, **Query 2**, and **Query 3** tabs.



Viewing Query Results

You can view results for queries that you submit in the **Query** tabs.

1. Click the **Query** tab corresponding to your query listed in the query box.
A table with the results of your query appears in the **Results** tab.

Pinning Query Results

You can pin query results for later reuse in the session. If you do not pin the **Query** tab, new queries overwrite results currently displayed in the tab.

1. Click the **Query** tab that contains the query you want to pin and click .
The tab name changes from **Query** to **Pinned**. The tab retains its number. Query results remain pinned until you unpin the tab or disconnect from the Teradata system.
2. [Optional] Click  to unpin your query.
The tab name changes from **Pinned** to **Query**. The results displayed in the results tab are cleared when you submit another query.

Results Tab

The **Results** tab provides detailed information about your query results. To view more columns in a wide results tab, drag the scroll box to the right. This information is read-only.

SQL Tab

The **SQL** tab displays the SQL for the selected query. This information is read-only.

Explain Tab

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session. Each Explain step is uniquely identified with a number where the background of the number box indicates status. This information is read-only. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator.

Step Information	Description
Step Number	<ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black number box. Active steps are indicated by a pulsating number box (flashes gray and white). Steps to run are at the bottom of the list and indicated by a white number box.
Est. Time	Estimated execution time for the step
Est. Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Blocked By Tab



The **Blocked By** tab displays details about other queries that are blocking this query. This information is read-only. The tab is available as long as other queries continue to block this query. When the blocking query conditions are resolved, this query can then complete, and the tab is no longer displayed.

Blocked By Information	Description
Username	Name of the user that is running the query that holds the lock
Session ID	Session ID of the blocking query
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access.
Locked	Name of the locked object

Object Browser


The object browser allows you to view a list of objects on a Teradata system. You can reduce the time required to build a query and reduce errors by inserting object names in a query.

The object browser displays database objects in the connected system as a hierarchically organized tree. Types of database objects include databases, users, tables, views, macros, functions, and procedures. You can use filtering to search for object types in the tree.

- Expand a branch of objects in the tree by clicking  next to the object type. The object browser totals the objects found on each expanded branch by object type. If no objects of the type are found, the branch displays **0 items**.
- Refresh the list of objects in the tree by clicking  next to the Teradata system name.

Filtering Objects

Filtering allows you to display only objects that match your filter criteria in the object browser. Filtering is not case-sensitive.

Filter boxes are located at each branch and sub-branch of the directory tree, under the name of the database object. When you click  next to the object, the filter box appears.

- Filter by entering text in the box.

For example, if you type *dbc* in the filter box under an Analytics Database name, objects containing *dbc* are displayed in the tree. If you use a wildcard, such as **dbc*, objects beginning with *dbc* and *DBC* are displayed.

- Filter for all other database object types by entering text in the box under the object type. Object types include **Tables**, **Views**, **Macros**, **Functions**, and **Procedures**.

For example, if you type *c1ear* in the filter box under a **Macros** sub-branch of the tree, macros containing the string *C1ear* are displayed in the tree.





Change the filter criteria by entering new text in the filter box. Unless you type new filter criteria or open other branches, the object browser maintains the state of the tree, including opened branches and applied filter criteria.




You cannot use *?* or *** as wildcards to match single or multiple characters in filter criteria.

The object browser can display a maximum of 500 objects at each branch of the tree. To display fewer results, enter filter criteria to narrow the list of database objects shown in the browser.

Object Types



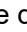
The database objects and associated icons that display in the object browser are listed as following.

Database Object	Icon
Database	
User	
Table	
View	


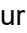


Database Object	Icon
Macro	
Function	
Procedure	

Inserting an Object into a Query

You can insert the qualified name of a database object into a query in the query box using the object browser. A *qualified object name* is a name composed of more than one identifier, joined with the dot character (.). The identifiers together specify a database object with more precision than a single name.

1. Position the cursor in the query box where you want to insert the object.
2. Click the  list next to the **Load** button and select **Insert Object**.
The **Object Browser** opens, displaying the name of the connected Teradata system.
3. Click  next to the Teradata system name to expand the directory tree.
4. [Optional] Enter filter criteria in the box under the Teradata system to find databases or users containing the criteria.
5. Click  next to the database object type in the tree to expand that sub-branch.
6. [Optional] Enter filter criteria in the box under the object type to find objects containing the criteria.
7. Click to highlight the object name.
8. Click **Insert Object** to copy the qualified object name to the query box.
The object name is inserted in the query box after the cursor location.
9. Click **Close**.

Example: Inserting an Object into a Query

1. Click  next to **Load** and select **Insert Object**.
The Object Browser opens.
2. Click  next to your database system name to expand the directory tree.
3. Click  next to DBC to expand the sub-branch in the directory tree.
4. Click  next to Tables to see all the table objects under the DBC sub-branch.
5. Click **Accounts** in the sub-branch to select the table.
Accounts is the identifier of a particular table in the DBC database.
6. Click **Insert Object**.
DBC.Accounts appears in the query box.

History Tab


The **History** tab provides the following information in the table or in balloon text for queries submitted in the current session:

- Date and time of the query
- Complete SQL statement
- Amount of time to complete the query
- Number of rows processed

You can also load queries from the **History** tab.

Viewing Previously Run Queries

You can view a list of queries submitted in the current session. You can load a query from the table for reuse in the same session.

1. In the query box view, click **Load** and click the **History** tab.
2. [Optional] For longer SQL statements, click  to display the complete query.
3. Click **Close**.

Loading a Previously Run Query

You can copy a query submitted in the current session to the query box for reuse or editing.

1. In the query box view, click **Load** and click the **History** tab.
2. In the query history table, click **Load** for the query you want to run.
The query is copied to the query box and the **History** tab closes.
3. [Optional] Edit the query before running it.
4. Click **Close**.

Saved Tab


The **Saved** tab provides the following information on queries saved in the current or previous sessions:

- Name of the saved query
- Complete SQL statement

You can also copy and delete queries using the **Saved** tab.

Viewing Saved Queries

1. In the query box view, click **Load** and click the **Saved** tab.
The name of the query and the SQL statement appear in the table of saved queries.

2. [Optional] For longer SQL statements, click  to display the complete query.
3. Click **Close**.


Loading a Saved Query

You can load a saved query into the query box for reuse or editing.

1. In the query box view, click **Load** and click the **Saved** tab.
2. In the table of saved queries, click **Load** for the query you want to run.
The query is copied to the query box and the **Saved** tab closes.
3. [Optional] Edit the query in the query box before running it.
4. Click **Close**.

Deleting a Saved Query

You can delete a saved query that you no longer need.

1. In the query box view, click **Load** and click the **Saved** tab.
2. In the table of saved queries, click  for the query you want to delete.
A confirmation message appears.
3. Click **OK**.
4. Click **Close**.

Stats Manager

Stats Manager Overview

The **Stats Manager** portlet allows you to manage Analytics Database statistics collection, which includes the ability to collect and analyze statistics, create and control jobs, and manage recommendations.

Use the portlet to do the following:

- View statistics on your system
- Identify statistics to be collected and schedule statistic collection
- Identify and collect missing statistics
- Detect and refresh stale statistics
- Identify and discontinue collecting unused statistics
- View when statistics were last collected and are scheduled for collection again

The **Stats Manager** portlet is available for systems running Teradata Database 14.10 and later.

Stats Manager Concepts

What Are Automate and Collect?

Automate

The automate option allows you to approve statistics for collection by collect jobs. You can determine whether to automate statistics on a single object or table, multiple objects or tables, entire databases, or entire systems.

New statistics collected on a system outside of the **Stats Manager** portlet are not automatically approved for collection by collect jobs. To approve new statistics for collection, set the automate options that are available on the **Statistics** tab.

Collect

A collect job submits COLLECT STATISTICS statements to the Analytics Database for statistics that are approved for automation. You can control the scope and schedule of each collect job, as well as assign a priority to individual COLLECT STATISTICS statements.

Running a collect job can be resource-intensive on the Analytics Database so you may want to run it during periods of lower activity.

What Is Analyze?

An analyze job evaluates statistics usage, identifies stale statistics, and generates recommendations on which statistics to collect, automate, reactivate, or deactivate.

Running an analyze job with the **Use DBQL** option enabled can be resource-intensive.

Working with Stats Manager

Statistics Integration

The **Stats Manager** portlet works with the Analytics Database and Viewpoint server to help reduce the impact of statistics management. The following describes the major components and how they work together.

Teradata

All analyze and collect job processing and statistics management action processing is done by the Analytics Database. This database also contains the TDStats database used by the **Stats Manager** portlet.

TDStats

This database contains additional information related to statistics management, such as the statistics that are automated, object scope lists for jobs, collect lists, job reports, and statistic recommendations. The TDStats database is updated when a job runs, a job completes, or when an action, such as automating a statistic, occurs.

Teradata Viewpoint Data Collection Service (DCS)

The DCS takes information from the Analytics Database and stores it in the local cache database for use by the **Stats Manager** portlet. The DCS is also used for job scheduling and initiating jobs. If the DCS is not running, scheduled jobs will not run.

Local Cache Database

This database contains all information that can be displayed in the **Stats Manager** portlet. The database is updated when a job completes, when an action such as Automate is performed, or when the Stats Manager collector runs. The local cache database makes data available quickly without impacting the Analytics Database.

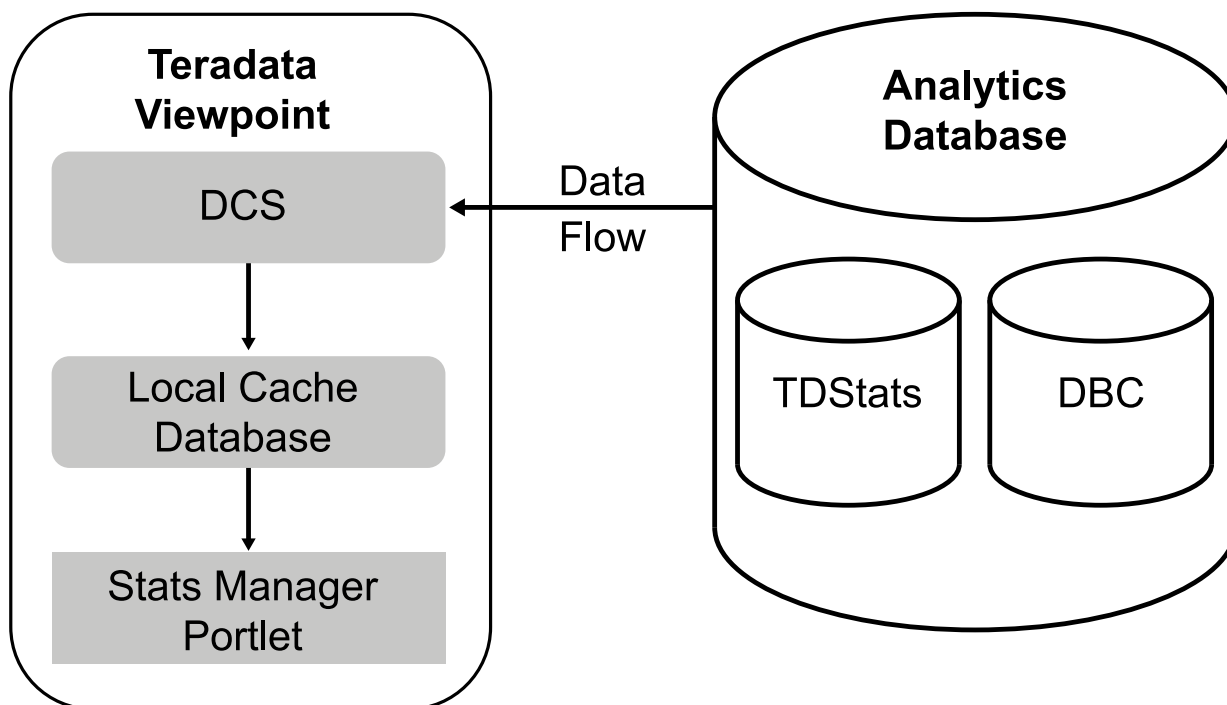
Stats Manager portlet

The **Stats Manager** portlet lets you view the status of statistics, view the status of jobs that analyze and collect statistics, and configure settings for the statistics management process.

Monitored Systems portlet

The **Monitored Systems** portlet lets you specify the login identifier used for managing statistics, the frequency of data collection from TDStats, and the frequency of data clean up for TDStats.

The following diagram shows how data flows from Analytics Database to the **Stats Manager** portlet.



Approaches to Automating Statistics Collection

You can use a submit-based approach, analyze-based approach, or a combination of both approaches, to automate statistics collection using the **Stats Manager** portlet.

Approach	Description
Submit-Based	Use this approach to: <ul style="list-style-type: none"> • Submit all automated statistics for frequent collection • Employ the Optimizer to determine if a statistic needs to be collected
Analyze-Based	Use this approach to: <ul style="list-style-type: none"> • Analyze the importance and staleness of statistics • Prioritize statistics submitted for collection based on the analyze results
Combination	Use a combination of both submit-based and analyze-based approaches to make sure data is collected on stale statistics on a timely basis.

Prior to implementing a statistics collection approach, you must do the following:

- Use the automate options to approve the statistics you want to automate.
- Create collect jobs to make sure all automated statistics are associated with a collect job.

All automated statistics are associated with a collect job when **100% of automated stats have collect jobs** displays on the **Statistics by Database** view on the **Statistics** tab.

Submit-Based Approach

The submit-based approach employs the Optimizer to determine if a statistic needs to be collected. You can enable collect thresholds to prevent repeat collection of statistics that were recently collected or have not changed significantly.

Following is an overview of the tasks employed in this approach:

1. Change the **Collect age threshold** and **Data change threshold** on the **Edit Collect Settings** dialog box to skip the collection of statistics that were recently collected or have not changed significantly.
2. Run the collect jobs frequently.

When a job is run, the Optimizer determines which statistics to collect, based on the threshold settings.

Review the **Collect Job report** for a list of collected and skipped statistics.

3. Verify that all the automated statistics are submitted for collection.

In the **Statistics by Database** view, examine the **Max Submit Age** metric for databases or **Submit Age** metric for objects. If the value is greater than the expected frequency, check the following:

- Statistics are associated with a collect job
- Collect job is scheduled to run with the expected frequency
- End state of the collect job is **Complete** or **Partially Complete**.

If the job is **Partially Complete** and not scheduled to resume, the job duration may be too short.

Related Information:

[Editing Collect Settings for a Single Statistic](#)

[Running a Job](#)

[Viewing Collect Reports](#)

Analyze-Based Approach

Use the analyze-based approach to evaluate the importance and staleness of statistics, and to prioritize the statistics submitted for collection.

During the collection process, statistics most in need of collection are submitted first. This makes sure that if a collect job does not complete in the allotted duration, then the higher priority statistics are more likely to have been collected.

If using the analyze-based approach, make sure the collect lists are automatically generated.

You can prioritize individual statistics by setting the user importance.

Following is an overview of the tasks employed in this approach:

1. Enable database query logging features (Requires DBQL option with STATSUSAGE. XMLPLAN is optional.) at least one week prior to running analyze jobs.
2. Create an analyze job.

Analyze all statistics on the system and select the **Use DBQL** option.

3. Schedule the analyze job to run at least once a week, so data changes are captured.
4. Access the **Analyze job report** once the analyze job has run.

Look for stale statistics. Verify that all stale statistics are scheduled for collection again in the near future.

5. From the **Statistics** tab, do the following:

- Review the **STALE STATS** column to see which statistics are considered stale.
- For collect jobs, make sure **STALE STATS** is nearing zero.

If the value is not nearing zero, increase the duration of the collect job or schedule the collect job to run again in the near future.

Related Information:

[Database Prerequisites for Analyze Jobs](#)

[Creating an Analyze Job](#)

[Viewing Analyze Reports](#)

Preventing Statistic Overcollection

The **Stats Manager** portlet can help prevent unnecessary statistic collection which may be contributing to an inefficient use of system resources. You may be collecting:

- Statistics that are not needed or provide no value
- Statistics that are needed but collected too frequently

Ensuring Unneeded Statistics Are Not Collected

Using the **Stats Manager** portlet, you can identify statistics that are not used by the Optimizer and stop collection of those statistics.

1. Enable database query logging with the USECOUNT option.
2. Create an analyze job.
 - Enable the option to evaluate statistics to determine if they should be collected or deactivated.
 - Set the length of time a statistic is not used in order for it to be considered inactive and recommended for deactivation.
3. Manually run or schedule the analyze job to run.

The analyze job generates deactivate recommendations and these statistics are no longer collected by collect jobs.

Ensuring Statistics Are Collected at the Right Frequency

You can specify collect settings thresholds to decrease the frequency of statistics collection.

1. Determine the statistics that are collected too frequently. For example, these could be statistics for data that is updated infrequently but the statistics are collected daily or are lower priority statistics that require additional time to collect.
2. From the **Statistics** tab, drill down to locate the statistic.
3. Edit the collect settings for the statistic.
 - Change the collect age threshold, data change thresholds, or both.
 - Select system-defined or user-defined thresholds to determine when to collect or skip collecting a statistic. System-defined thresholds are recommended initially, however user-defined thresholds provide more control over collection thresholds.
4. [Optional] Review the collect job reports to verify that the statistics are skipped as expected based on the thresholds you set.

Related Information:

[Creating an Analyze Job](#)

[Running a Job](#)

[Viewing Analyze Reports](#)

[Statistics Views](#)

[Editing Collect Settings for a Single Statistic](#)

Defining Object Lists and Exclusions

This topic provides general information on how to define object lists to use in job scopes, and how to prevent objects from being included in the jobs.

Creating Object Lists

You can create an object list in Analytics Database and use the list when specifying the scope of jobs.

1. Create an object list using the TDStats.CreateObjectList stored procedure.
2. Populate TDStats.ObjectListEntry with appropriate values.
3. Wait for the Stats Manager data collector to run for the object list to be recognized by the portlet. Once recognized, it appears in the existing object list on the **Scope** tab when editing a job.
4. Select the object list to have the job use for the job scope.

Excluding Objects from Jobs

Define Global Excludes to prevent objects from being included in Stats Manager jobs and from appearing in the portlet.

1. Use the TDStats.AddExcludedObject stored procedure to identify objects for exclusion.
2. Wait for the Stats Manager data collector to run for the excluded object list to take effect.

For more information, see *Teradata® Database Application Programming Reference*.

Related Information:

[Job Scope](#)

Working with Systems

You can set statistics management options in Analytics Database and use other Viewpoint portlets to support statistics management.

Changing Systems

1. Select ▼ in the portlet frame and select a system.
Only systems that have the Stats Manager collector enabled are included in the system selection menu.
The systems that appear in the system selection menu are based on the permissions assigned to your role.

Statistics Support in Other Portlets

In conjunction with the **Stats Manager** portlet, you can use other Viewpoint portlets to support statistics management.

Workload Designer

You can use the **Workload Designer** portlet to regulate the resource usage from COLLECT STATISTICS statements. From the **Classification** tab for a workload, filter, or throttle, add a **Query Characteristics** criteria type, and select the **COLLECT STATISTICS** statement type. This statement type is available in Teradata Database 14.10 and later.

Metric Heatmap

Before running resource-intensive jobs in the **Stats Manager** portlet, you can use the **Metric Heatmap** portlet to find hours in the day where system resources are not fully used. Use the **Stats Manager** portlet to schedule jobs to run during those hours.

Query Monitor

You can track the sessions used for statistics jobs using the **Query Monitor** portlet. All jobs that run in the **Stats Manager** portlet use the login that is specified when configuring the Stats Manager collector in the **Monitored Systems** portlet. If this login is reserved only for use by the **Stats Manager** portlet, you can use this unique identifier to see which sessions

are associated with statistics jobs. Also, the name of the job and job type appear in the query band information for the session in the **Query Monitor** portlet.

Database Prerequisites for Analyze Jobs

Some analyze job options require that database query logging (DBQL) features be enabled. For the analyze job options that you plan to use, you must enable the associated database feature prior to running the job.

Before running an analyze job that has the **Evaluate stat usage and staleness and find missing stats to collect** check box selected and the **Use DBQL** option selected, enable query logging with the STATSUSAGE option (XMLPLAN is optional) for the objects being analyzed. These options enable the analyze job to identify statistic usage and make recommendations about new statistics that should be collected as reported by the Optimizer and to identify stale statistics. The query logging features should be enabled for at least one week prior to running the analyze job.

Before running an analyze job that has the **Evaluate stats to determine if they should be collected or deactivated** check box selected, enable query logging with the USECOUNT option for the objects being analyzed. This option enables the analyze job to identify inactive statistics. The query logging feature should be enabled, prior to running the analyze job, for a period of time that is equal to the period specified in the **Recommend stats for deactivation when inactive for** option.

For more information, see *Teradata® Database Release Summary* and *Teradata® Database Administration*.

Working with Jobs

You can create jobs for analyzing and collecting statistics.

- An analyze job enables you to select database objects to analyze, define types of analysis to perform, and schedule when the job runs. You can evaluate statistics usage and staleness, find missing statistics to collect, evaluate statistics to determine if they should be collected or deactivated, and automate existing statistics.
- A collect job enables you to select the database objects to collect statistics on and schedule when the job runs.

Job Definitions View

The **Job Definitions** view displays the list of user-defined collect and analyze job definitions. From this view, you can create jobs, manage existing jobs, and review job reports.

▼ Actions

New Collect Job enables you to define a job to collect statistics, which includes specifying the scope, defining a collect list, and setting the schedule.

New Analyze Job enables you to define a job to evaluate statistic use and make recommendations. You can specify the scope, select functions, and set the schedule.

View History lists the run status and reports for collect and analyze jobs over time.

Edit System Analyze Settings manages all long running analyze jobs in a selected system for Teradata Database 15.10 and later.

Configure Alerts enables you to add a new alert for Collect/Analyze jobs, which includes specifying the severity, action, jobs, an optional field to provide message, and a checkbox to enable/disable the defined alert. The defined alert list can be viewed, edited, and deleted.

State Filter Bar

Displays a count of jobs in each state. Click a state to show only the jobs in that state.

▾ **Overflow Menu** displays additional states. If the available space on the state filter bar is exceeded, you can select another state to replace the last state.

Filters

Shows only rows that match your filter criteria.

▾ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

Abort allows you to stop the selected job.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Job Definitions Table

Displays summary information about jobs. The table is configured in the **Configure Columns** dialog box. Click a row in the table to see details.

Job Schedule

Displays a nine-day view of jobs that are running, scheduled to run, or have already run. Hover over a date to see a list of jobs.

The screenshot shows the 'Stats Manager' interface with the 'Jobs' tab selected. The top navigation bar includes 'Statistics' and 'Jobs' tabs. Below the navigation bar, there's a 'Home' link and a 'Dev1 | Job Definitions' section. A 'State Filter Bar' shows counts for various job states: 59 All, 8 Last 4..., 2 Next 4..., 1 Running, 21 Error, and 0 Partiall... An 'Actions' dropdown menu is visible. Below the filter bar is a 'Filters' section with input fields for 'JOB', 'LAST START', 'TYPE', and 'STA'. The main area displays a 'Job Definitions Table' with columns for Job Name, Last Start, Type, and Status. The table lists several jobs, including 'A_CltSCollectLst_3112043' and 'A_CltSCollectLst_311338'. Below the table, a 'Job Schedule' calendar for the period '2/22/14 - 3/2/14' is shown, displaying a nine-day view with dates and job counts for each day.

Viewing Job Definitions

You can view the state of analyze and collect jobs.

1. Click the **Jobs** tab.
2. [Optional] Click a row to view the job report.

Viewing the Job Schedule

The job schedule is a calendar that displays a nine-day view of jobs that are running, scheduled to run, or have already run.

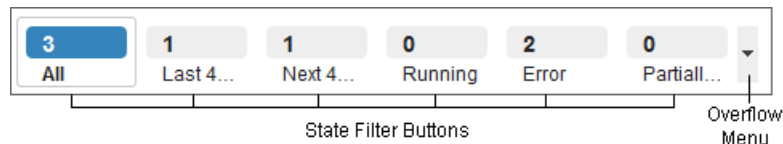
1. From the **Jobs** tab, hover over the **Job Schedule** calendar to view job names, job types, and start times for all jobs on a day.
2. Click the **Previous week** or **Next week** button to view the previous or next nine-day

This image shows a detailed view of the 'Job Schedule' calendar for the period '2/22/14 - 3/2/14'. The calendar displays a nine-day view with dates from Saturday 22 to Sunday 2. Each date cell contains a job count. Navigation buttons for 'Previous Week' and 'Next Week' are visible at the bottom. The 'Previous Week' button is highlighted, and the 'Next Week' button is also visible.

- Click **Show today** to return to the current nine-day period.
Show today appears only if viewing a nine-day period other than the current one.

State Filter Bar

The state filter bar allows you to display specific states in the view. If the number of states exceeds the available space, you can select another state from the **Overflow Menu** to replace the last state.



The state filter buttons show a count of jobs for each state category. Click a button to display only jobs in that state in the **Job Definitions** view. For example, click **Running** to display only the jobs that are currently running.

All

All jobs in the system.

Last 48 hours

Jobs that completed in the last 48 hours.

Next 48 hours

Jobs scheduled to start in the next 48 hours.

Running

Jobs that are currently running.

Error

Jobs that encountered an error during processing and did not complete. If an analyze job, only functions completed prior to the error are returned for the job report. If a collect job, any statistics collected prior to the error remain collected.

Partially Complete

Collect jobs that did not finish running in the time allocated for them.

Complete

Jobs that ran successfully.

Abort

Jobs that were stopped by the user. If an analyze job, only functions completed prior to the abort will appear in the job report. Incomplete functions will have no result. If a collect job, any statistics collected prior to the abort remain collected.

Analyze Jobs

Analyze jobs allow you to evaluate statistics status and get statistic-related recommendations. Analyzing objects enables you to determine where additional statistics may be useful and identify existing statistics that are used frequently or are stale.

You can create new jobs, edit existing jobs, and cancel jobs.

When you create an analyze job, you:

- Specify the database objects the job analyses.
You can define the scope to include all objects or just a few.
- Select functions to be performed when the job is run, such as:
 - Evaluate statistics usage and staleness and find missing statistics to collect
 - Evaluate statistics to determine if they need to be collected or deactivated
 - Automate existing statistics
- Specify when and how often the job runs.

The performed functions when the job runs are described in the following table.

Function	Description
Evaluate statistics usage and staleness and find missing statistics to collect	Determines which statistics are used, which statistics are stale, and which new statistics should be collected to help the Optimizer make better decisions. The job report lists statistics that are stale and statistics that do not exist but the Optimizer would have used them if they had existed. If the Use DBQL option is enabled, the job report lists statistics that were used by the Optimizer. After the job runs, collect recommendations are created.
Evaluate statistics to determine if they should be collected or deactivated	Determines statistics that are useful to the Optimizer and recommends if statistics should be activated or deactivated. The job report lists statistics that have not been used by the Optimizer for the specified period of time. These statistics can be flagged as inactive so that they are no longer collected by collect jobs. Also reports on statistics that have been flagged as inactive but that have been used by the Optimizer over the specified period. After the job runs, deactivate and reactivate recommendations are created.
Automate existing statistics	Reports on statistics that have not yet been automated. The job report lists statistics that have not been automated with an option to include only new statistics since a particular date or only those statistics that are used as determined by the first function listed before. After the job runs, automate recommendations are created.

When you specify the objects the job analyse, you can use existing object lists and custom object lists. An object list is a list of database objects used for a job scope. An existing list is a custom list from another job or a list created outside of the **Stats Manager** portlet.

Creating an Analyze Job

Prerequisite:

You must enable the DBQL option before selecting either of the evaluate stats options on the **Functions** tab.

1. From the **Jobs** tab, click **Actions > New Analyze Job**.
2. Enter a job name, up to 30 alphanumeric characters.
The name cannot be the same as an existing job name.
3. From the **Scope** tab, click one of the following to specify the objects that will be analyzed by this job:

Option	Description
All objects	Select to analyze all objects across the entire system.
Existing object list	Select to analyze statistics using an existing list that was created for other jobs. Object lists that were externally created outside of the portlet display first.
Custom object list	Select to analyze statistics using a custom object list. <ol style="list-style-type: none"> a. Click Edit list to customize the object list. b. From the Objects tab, select to display Automated objects or All objects. Then use the object browser to either select databases to include all related objects in those databases or select individual objects within those databases. c. From the Wildcards tab, use wildcard patterns to add objects to the job scope. Click Add Patterns to take the filter combinations you entered in the filter boxes and create a wildcard in the Wildcard Patterns pane. d. Click OK.

4. Click the **Functions** tab and select one or more of the following options to be performed when this job is run:

Option	Description
Evaluate stat usage and staleness and find missing stats to collect	Select to have Analytics Database evaluate statistics usage and staleness and find missing statistics to collect. Do any of the following: <ul style="list-style-type: none"> • [Optional] Select Require review before applying recommendations if you want to manually approve or reject recommendations. • Select Edit Settings to use DBQL (default) or rule-based analysis options to determine statistics usage and staleness, and find new collection opportunities, and then click OK. With DBQL, advanced options allow you to limit the queries to a specific user, application, account string, or query band. You can also override the global query log analysis time limit.
Evaluate stats to determine if they should be collected or deactivated	Select to have Analytics Database evaluate statistics to determine if they should be collected or deactivated. Do any of the following: <ul style="list-style-type: none"> • [Optional] Select Require review before applying recommendations if you want to manually approve or reject recommendations.

Option	Description
	<ul style="list-style-type: none"> Accept the default value or enter a different number to specify a frequency to recommend statistics be deactivated when they have been inactive for the specified amount of time.
Automate existing stats	<p>[Optional] Select to automate existing statistics. Do any of the following:</p> <ul style="list-style-type: none"> Select Require review before applying recommendations if you want to manually approve or reject recommendations. Select Limit to stats to enable the following options: <ul style="list-style-type: none"> Select Initially collected within the last to limit to statistics that were initially collected within the specified days, weeks, or months. Select That are actually used to limit to statistics that are actually used.

- Click the **Schedule** tab and do one of the following:
 - Click **Add Schedule** to add a new schedule.
 - Select an existing schedule to edit.
 - Click **OK**.
- Click **Save**.
The job appears in the list on the **Jobs** tab.
- After the job is run, review the results in the analyze job report.
 - If you selected **Evaluate stat usage and staleness and find missing stats to collect**, review the results in the **Missing**, **Usage**, and **Stale** tabs.
 - If you selected **Evaluate stats to determine if they should be collected or deactivated**, review the results in the **Deactivate** and **Reactivate** tabs.
 - If you selected **Automate existing stats**, review the results in the **Automate** tab.

Related Information:

[Database Prerequisites for Analyze Jobs](#)

[Using the Object Browser](#)

[Using Wildcard Patterns](#)

[Editing DBQL or Rule-Based Analysis Settings](#)

Collect Jobs

The collect job collects statistics that have been automated. A collect list is an ordered list of COLLECT STATISTICS statements used by a collect job.

When you create a collect job, you:

- Specify the database objects to be used to generate the collect list.

When you specify the objects whose automated statistics are collected by the job, you can use existing object lists and custom object lists. An object list is a list of database objects used for a job scope. An existing list is a custom list from another job or a list created outside of the **Stats Manager** portlet.

- Specify if you want the collect list to be generated manually or automatically.

When a collect list is automatically generated, it orders the COLLECT STATISTICS statements according to system-determined and user-specified importance. When a collect list is generated manually, you can specify a custom collect list that is either for the next collect only or for all collections in the future.

- Specify when and how long the job will run. You can set the collection time frame, recurrence, and duration.

Consider scheduling collect jobs when it does not negatively impact production workloads, such as at night or during the weekends to avoid peak hours when the system is busiest.

Creating a Collect Job

The collect job collects automated statistics. You can preview and customize the list of collected statistics and their collect order. The order in which statistics are collected is automatically prioritized by default, but you can change the priority before the collection is run.

- From the **Jobs** tab, click **Actions > New Collect Job**.
- Enter a job name, up to 30 alphanumeric characters.
The name cannot be the same as an existing job name.
- From the **Scope** tab, click one of the following to specify the objects that will be used to generate collect lists:

Option	Description
All automated objects	Select to collect automated statistics on all objects across the entire system.
Existing object list	Select to collect automated statistics using an existing list that was created for other jobs. Object lists that were externally created outside of the portlet display first.
Custom object list	Select to collect automated statistics using a custom object list. <ol style="list-style-type: none"> Click Edit list to customize the object list. From the Objects tab, use the object browser to either select databases to include all related objects in those databases or select individual objects within those databases. From the Wildcards tab, use wildcard patterns to add objects to the job scope. Click Add Patterns to take the filter combinations you entered in the filter boxes and create a wildcard in the Wildcard Patterns pane. Click OK.

- Click the **Collect List** tab and select one of the following:

Option	Description
Automatically generate the collect list	<ol style="list-style-type: none"> Select if you want COLLECT STATISTICS statements to be automatically ordered according to system-determined and user-specified importance.

Option	Description
	<p>This option is recommended.</p> <p>b. Select Preview list to generate a collect list.</p>
Use the list below for the next collect only	<p>a. Select if you want to customize the order of COLLECT STATISTICS statements the next time the job is run.</p> <p>This option automatically switches to a system-generated collect list after the job runs once. Use this option when you need to temporarily adjust collection priority to favor statistics that need immediate attention.</p> <p>b. Select Regenerate list to clear the current collect list and regenerate a new list.</p> <p>c. Click Continue.</p>
Use the list below for all collections	<p>a. Select if you want to customize the order of COLLECT STATISTICS statements every time the job is run in the future.</p> <p>b. Select Regenerate list to clear the current collect list and regenerate a new list.</p> <p>c. Click Continue.</p> <p>Choosing this option prevents changes made from the Scope tab and in the Edit Collect Settings dialog box from taking effect, because the COLLECT STATISTICS statement has already been generated.</p>

5. After the collect list is generated, click **Close**.
6. Click the **Schedule** tab and do one of the following:
 - Click **Add Schedule** to add a new schedule.
 - Select an existing schedule to edit.
7. Under **Schedule**, specify when this job will run.
8. [Optional] Select the following options:

Option	Description
Limit to	Under Duration , specify the maximum amount of time you want the job to run.
Resume collect list if not yet complete	<p>Under Options, select to enable the following options:</p> <ul style="list-style-type: none"> • Select If complete, restart job to start a job from the beginning if all items on the previous collect list of the job have already been submitted. • Select If complete, do nothing to prevent the job from running if all items on the previous collect list of the job have already been submitted.

9. Click **OK**.
10. Click **Save**.

The job appears in the list on the **Jobs** tab.

Related Information:

[Using the Object Browser](#)

[Using Wildcard Patterns](#)

Editing System Analyze Settings

The System Analyze Settings lets you manage all long running analyze jobs in a selected system.

1. From the **Jobs** tab, click **Actions > Edit System Analyze Settings**.
2. Click an option.

Option	Description
Enable progress monitoring for analyze jobs	When enabled, shows progress bar for analyze jobs and messages.
Set the default query log analysis time limit to	When enabled, shows time frame to limit query log analysis time. The default time is 0 minutes.

3. Click **OK**.

Configuring Alerts

Configuring alerts helps to add a new alert edit an existing alert or delete an existing alert.

1. From the **Jobs** tab, select **Actions > Configure Alerts**
2. An option to Add Alert is displayed with a list of already defined alerts for the selected system if any.

The Alert Definitions contains the summary of the alerts configured for a given system.

☐ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Adding a New Alert

Configure an alert by selecting the jobs, severity, and action. The alert is triggered when the selected jobs are run manually or run on schedule. This ends up in ERROR state and can be viewed in Alert Viewer portlet based on the preferences given in the selected action.

1. From the **Jobs** tab, click **Actions > Configure Alerts**


2. Click on **Add Alert** .
3. Select the checkbox **Enable Alert** to enable it.
4. Enter the **Alert Name**, up to 30 alphanumeric characters. The name cannot be same as an existing alert name.
5. Select **Severity** from Critical, High, Medium, and Low.
6. Select **Action** from the dropdown. The actions sets are defined in the **Alert Setup** portlet.
7. Enter a message for the alert (Optional).
8. Select one or more Jobs for which alerts must be triggered if those jobs end in ERROR state.


Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

For more information, see [Summary Table Controls](#).

Editing an Existing Alert

To modify an existing alert, select  of the **alert row** and select the **Edit** option

Deleting an Existing Alert

To modify an existing alert, select  of the **alert row** and select the **Delete** option.

In **Alert Viewer** portlet, alert has the source as StatsManager, AlertName shows as the alert name configured, and on drilling, the alert details will be shown. The alert criteria will be “Job did not complete due to an error”. The properties will have the details of Job Name, System Name, Alert Name, and Error Message.

Managing Jobs

You can manage jobs from the **Job Definitions** view.

Run Now

Enables you to manually run a job without waiting for its next scheduled start time.

Edit

Enables you to modify the scope, options, or schedules for a job.

Abort

Stops a running job.

Delete

Removes a job from the **Job Definitions** list. Existing history of the job is not lost.

Running a Job

Any job can be either scheduled or run immediately. The manual method runs a job immediately and is available only for jobs not in a running state.

1. From the **Jobs** tab, click ☐ next to a job and click **Run Now**.

Job Type	Description
Analyze	The job runs immediately and no further action is required.
Collect	<ol style="list-style-type: none"> a. [Optional] Select Limit to and enter a value up to three days to limit the duration of a job. b. [Optional] Select Resume Collect list to continue a partially complete job. c. Click Run.

Editing a Job

You can edit job parameters. If editing a collect job, you can edit the scope, collect list, and schedule. If editing an analyze job, you can edit the scope, functions, and schedule.

1. From the **Jobs** tab, do one of the following:
 - Click ☐ next to a job and click **Edit**.
 - Click a job and click **Actions > Edit**.

Aborting a Job

Aborting a collect job stops the statistics that are currently being collected. All of the statistics that were collected prior to the abort remain collected. Aborting an analyze job stops functions that are in progress, but any functions completed prior to the abort will return data for the job report.

The **Abort** option is available only for jobs with the state **Running**.

1. From the **Jobs** tab, do one of the following:
 - Click ☐ next to the job and click **Abort**.
 - Click the job and click **Actions > Abort**.

Deleting a Job

If you no longer need a job, you can delete it.

1. From the **Jobs** tab, do one of the following:

- Click ☐ next to a job and click **Delete**.
- Click the job and click **Actions > Delete**.

The job is removed from the **Job Definitions** list and is no longer available to run.

Job Scope

You can define the scope of collect and analyze jobs. The scope of a collect job defines the objects the job will use to generate a collect list. The scope of an analyze job defines the objects that will be analyzed by the job.

- A collect list is an ordered list of COLLECT STATISTICS statements used by a collect job.
- An object list is a list of database objects used for a job scope.
- An existing object list is a custom object list defined in another job or an object list created outside of the **Stats Manager** portlet.

You can select from the following options to specify the scope for a job.

All objects

Uses all objects in the system for analyze jobs.

All automated objects

Uses all automated objects in the system for collect jobs.

Existing object list

Enables you to select an object list from a number of already defined object lists. Object lists that were externally created outside of the portlet display first. After you select an object list, the objects in that list are displayed on the **Scope** tab.

Custom object list

Enables you to define object lists and use wildcard patterns to add objects to the scope.

Using Wildcard Patterns

You can avoid manually selecting what to include in a job scope by using wildcard patterns. Wildcards enable you to select all objects, databases, or both whose names match the criteria you specify. The wildcards are then evaluated at run time to determine the scope of the job.

1. From the **Jobs** tab, click ☐ next to a job and click **Edit**.
2. From the **Scope** tab, click the option **Custom object list** and then click **Edit list**.
3. From the **Wildcards** tab, enter a pattern for a database, object, or both in the filter boxes.
If you select only a database pattern, all objects in the databases that meet that criteria become the scope.
If you select only an object pattern, all objects that meet that criteria in any database are added to the scope.

4. Click **Add pattern**.


The pattern is displayed in the **Wildcard patterns** pane and is saved for reuse.

5. [Optional] Click  next to a pattern, and click **Delete** to delete the wildcard pattern.

6. Click **OK**.


Object Browser

The object browser allows you to view a list of objects on a Teradata system. You can select objects from this list to include when creating or editing the scope of a job. The object browser is populated with all database and object selections in effect when the tab is selected.

The object browser displays database objects in the connected system as a hierarchically organized tree. Expand a branch of objects in the tree by clicking  next to the object type. The object browser totals the objects found on each expanded branch by object type. If no objects of the type are found, the branch displays **0 items**.

Filtering Objects

Filtering allows you to display only objects that match your filter criteria in the object browser. Filtering is not case-sensitive.






Filter boxes are located under the name of the database object. When you click  next to the object, the filter box appears. Filter by entering text in the box. For example, if you type *dbc* in the filter box under an Analytics Database name, objects containing *dbc* are displayed in the tree. If you use a wildcard, such as **dbc*, objects beginning with *dbc* and *DBC* are displayed.

Change the filter criteria by entering new text in the filter box. You can use *?* or *** as wildcards to match single or multiple characters in filter criteria.

The object browser can display a maximum of 100 rows. If there are greater than 100 matches to the filter, a paging control is displayed with the total count. To display fewer results, enter filter criteria to narrow the list of database objects shown in the browser.



Object Types

The database objects and associated icons that display in the object browser are listed as following.

Database Object	Icon
Database	
User	
Table	
Hash Index	
Join Index	




Using the Object Browser

The object browser allows you to view a list of objects on a Teradata system.

1. From the **Jobs** tab, do one of the following:
 - Click  next to a job and click **Edit**.
 - Click a job and click **Actions > Edit**.
2. From the **Scope** tab, select **Custom object list** and then click **Edit list**.
3. Click  next to the Teradata system name to expand the directory tree.
4. If you are modifying an analyze job, from the **Objects** tab select either **Automated objects** or **All objects** to filter the objects that display in the tree.
5. Do one of the following:
 - Choose a database by selecting its check box.
If the database includes children, these objects are also selected.
 - Choose a child by selecting the check box next to the child.

Note:


If you select a child, the check box next to the database is cleared.

6. [Optional] Enter filter criteria in the box under an Analytics Database to find objects containing the criteria.
7. [Optional] Click  to the right of the filter box to select or clear matching objects.
 appears only if a database has been expanded but not selected.
8. [Optional] Click  to the right of the filter box to display more filters.
You can select any or all of these filters to further limit the objects displayed.

Editing Job Functions for Analyze Jobs

Prerequisite:

You must enable DBQL with the STATSUSAGE option (XMLPLAN is optional) before selecting either of the evaluate stats options on the **Functions** tab.

1. From the **Jobs** tab, do one of the following:
 - Click  next to an analyze job and click **Edit**.
 - Click an analyze job and click **Actions > Edit**.
2. Click the **Functions** tab and select one or more of the following options to be performed when this job is run:

Option	Description
Evaluate stat usage and staleness and find missing stats to collect	<p>Select to have Analytics Database evaluate statistics usage and staleness and find missing statistics to collect so they can be displayed in the portlet. The following selections are optional.</p> <ol style="list-style-type: none"> Select Require review before applying recommendations if you want to manually approve or reject recommendations. Select Edit Settings to use DBQL or rule-based analysis options to determine statistics usage and staleness, and find new collection opportunities, and then click OK.
Evaluate stats to determine if they should be collected or deactivated	<p>Select to have Analytics Database evaluate statistics to determine if they should be collected or deactivated. The following selections are optional.</p> <ol style="list-style-type: none"> Select Require review before applying recommendations if you want to manually approve or reject recommendations. Enter a number and specify a frequency if you want to recommend statistics be deactivated when they have been inactive for the specified amount of time.
Automate existing stats	<p>Select to automate existing statistics. The following selections are optional.</p> <ol style="list-style-type: none"> Select Require review before applying recommendations if you want to manually approve or reject recommendations. Select Limit to stats to enable the following options: <ul style="list-style-type: none"> Select Initially collected within the last to limit to statistics that were initially collected within the specified days, weeks, or months. Select That are actually used to limit to statistics that are actually used.

- Click **Save**.
- After the job is run, review the results in the analyze job report.
 - If you selected **Evaluate stat usage and staleness and find missing stats to collect**, review the results in the **Missing**, **Usage**, and **Stale** tabs.
 - If you selected **Evaluate stats to determine if they should be collected or deactivated**, review the results in the **Deactivate** and **Reactivate** tabs.
 - If you selected **Automate existing stats**, review the results in the **Automate** tab.

Related Information:

[Database Prerequisites for Analyze Jobs](#)

Editing DBQL or Rule-Based Analysis Settings

Prerequisite:

You must enable DBQL with the STATSUSAGE option (XMLPLAN is optional) before selecting **Use DBQL** from the **Functions** tab.

- From the **Jobs** tab, do one of the following:
 - Click ☐ next to a job and click **Edit**.

- Click a job and click **Actions > Edit**.
2. From the **Functions** tab, click **Edit Settings**.
 3. Choose one of the following options:

Option	Description
Use DBQL	<ol style="list-style-type: none"> Enter the name of the database containing log tables. [Optional] Select Limit by query log date and enter an interval or duration for collections. [Optional] Select Override global query log analysis time limit to change the time limit. [Optional] Click <input type="checkbox"/> to select Advanced Options.
Use rule-based analysis	[Optional] Select Limit to objects created in the last and enter a number and specify a frequency.

4. Click **OK**.

Related Information:

[Database Prerequisites for Analyze Jobs](#)

Editing the Collect List

The **Collect List** enables you to preview and customize the statistics that a job collects. You can select whether to automatically generate a collect list or use an existing list, and you can prioritize the objects in the list.

1. From the **Jobs** tab, do one of the following:
 - Click ☐ next to a collect job and click **Edit**.
 - Click a collect job and then click **Actions > Edit**.
2. Click the **Collect List** tab and select one of the following:

Option	Description
Automatically generate the collect list	<ol style="list-style-type: none"> Select if you want COLLECT STATISTICS statements to be automatically ordered according to system-determined and user-specified importance. This option is recommended. Select Preview list to generate a collect list.
Use the list below for the next collect only	<ol style="list-style-type: none"> Select if you want to customize the order of COLLECT STATISTICS statements the next time the job is run. This option automatically switches to a system-generated collect list after the job runs once. Use this option when you need to temporarily adjust collection priority to favor statistics that need immediate attention. Select Regenerate list to clear the current collect list and regenerate a new list. Click Continue.

Option	Description
Use the list below for all collections	<ol style="list-style-type: none"> Select if you want to customize the order of COLLECT STATISTICS statements every time the job is run in the future. Select Regenerate list to clear the current collect list and regenerate a new list. Click Continue. <p>Choosing this option prevents changes made from the Scope tab and in the Edit Collect Settings dialog box from taking effect, because the COLLECT STATISTICS statement has already been generated.</p>

- After the collect list is generated, click **Close**.

Setting the Job Schedule

Scheduling jobs enables you to collect and analyze statistics at regular intervals.

When scheduling a job, you can set the following options:

- Frequency of job runs
- Interval at which the job repeats
- Date or day of the week on which the job runs
- Time the job starts

For collect jobs, you can also limit job duration and set a partially complete job to resume where it left off.

Deleting a Job Schedule

A job can have one or more schedules. When a schedule is no longer needed, you can remove it from the schedule list for that job. This process deletes only the schedule, not the job itself.

- From the **Jobs** tab, click ☐ next to a job and click **Edit**.
- Click the **Schedule** tab.
- Click ☐ next to a schedule, and click **Delete** to delete the job schedule.
- Click **Save**.

Adding a Job Schedule

You can add a job schedule to a collect or analyze job.

- From the **Jobs** tab, click ☐ next to a job and click **Edit**.
- Click the **Schedule** tab.
- Click **Add Schedule**.
- Under **Schedule**, specify when this job will run.
- [Optional] If a collect job, select the following options:


Option	Description
Limit to	Under Duration , specify the maximum amount of time you want the job to run.
Resume collect list if not yet complete	Under Options , select to enable the following options: <ul style="list-style-type: none"> • Select If complete, restart job to start a job from the beginning if all items on the previous collect list of the job have already been submitted. • Select If complete, do nothing to prevent the job from running if all items on the previous collect list of the job have already been submitted.

6. Click **OK**.

Job Reports

Job reports show the status and results of a job. There are two report types: collect job reports and analyze job reports.


From the **Jobs** tab, you can browse to the following reports for collect and analyze jobs, and to review job history.

Report	Description	Navigation
Analyze Job Report	Shows job status and information about missing statistics, statistics usage, stale statistics, and statistics that should be reactivated, deactivated, or automated	Do either of the following: <ul style="list-style-type: none"> • Click a row from the Job Definitions view to see the current report for that job • Click a row from any job history view to see the associated report
Collect Job Report	Shows the overall job status and the status of the individual COLLECT STATISTICS statements found in the collect list of a job	
Analyze Report Details	Shows details associated with an individual statistic listed in the Analyze Job Report	From the Job Report view, click a row to see details of that job run
Collect Report Details	Shows COLLECT STATISTICS statements and details	
Job History (All jobs)	Lists all runs of collect and analyze jobs for a specified time frame, even if the job definition no longer exists	From the Job Definitions view, click Actions > View History
Job History (Single job)	Lists all runs of a single job for a specified time frame, even if the job definition no longer exists	Do either of the following: <ul style="list-style-type: none"> • From the Job Definitions view, click  next to a job and click View History • From the Job Report view, click Actions > View History

Note:

The **Actions** list is not available if you access a job report from any of the job history views.

Viewing Analyze Reports

- From the **Jobs** tab, do one of the following:
 - To view the current report for an analyze job, click the row of the job.
 - To view a past report for an analyze job, click  next to the job, click **View History**, and then click the job instance.

The **Analyze job report** view appears.

The info bar at the top of the view, displays the progress of the **Query log analysis** and **Staleness analysis**, as well as the **Start time**, **Duration**, and **CPU** time.

Note:

The info bar only displays if the **Enable progress monitoring for analyze jobs** option is checked in the **Edit System Analyze** settings.

The tabs in the report view show information about the statistics that were included in the analysis. The tabs displayed depend on the options selected from the **Functions** tab when you created or edited an analyze job.

- Select from the following tabs to view information about the job:

Tab	Displays
Usage	Information about the statistics that were used
Missing	Statistics that are recommended for collection
Stale	Statistics in the database whose status is stale
Reactivate	Statistics that were deactivated but are now recommended for reactivation
Deactivate	Statistics that were active but are now recommended for deactivation
Automate	Statistics that should be automated

- [Optional] Click a row of the table to see details for that job instance.

Related Information:

[Recommendations](#)

Viewing Collect Reports

- From the **Jobs** tab, do one of the following:

- To view the current report for a collect job, click the row of the job.
- To view a past report for a collect job, click ▾ next to the job, click **View History**, and then click the job instance.

The **Collect job report** view appears.

- From the **Show details for** list, select one of the following:
 - **Current** displays only results for the most recent run of the job.
 - **Cumulative** displays cumulative results for all runs of the job since starting collection from the beginning of a collect list.
- [Optional] Click a row of the table to see details for that job instance.

Viewing Job History

- From the **Jobs** tab, do one of the following:
 - To view the history of all jobs, click **Actions > View History**.
 - To view the history of a single job, click ▾ next to the job and click **View History**.
- Use the filter at the upper right corner to select a time frame.
- [Optional] To view the report of a job instance from this view, click the row of the instance.
Depending on the type of job, either an **Analyze job report** or a **Collect job report** appears.

Working with Statistics

Statistics Views

There are four statistics views. The descriptions and how to browse to each view are listed in the following table.

View	Description	Navigation from Home
Statistics by Database	View statistic information on all databases on a system and access system-generated recommendations.	Click the Statistics tab to see the Statistics by Database view.
Statistics by Object	View statistic information on objects, such as tables, for a single database.	From the Statistics by Database view, click a row in the table to see the Statistics by Object view.
Statistics	View statistics information for a single object, such as a table or index.	From the Statistics by Object view, click a row in the table to see the Statistics view.
Statistic Details	View details and collect settings for a single statistic.	From the Statistics view, click a row in the table to see the Statistic Details view.

Depending on the view, the following features are available:

▼ Actions

Automate enables statistics to be collected by collect jobs. You can automate a system, database, object, or single statistic.

Deautomate stops statistics from being collected by collect jobs. You can deautomate a system, database, object, or single statistic.

Edit Collect Settings allows you to edit thresholds, sampling, and histogram settings, where you can either choose to allow Analytics Database to decide the settings or select no thresholds or sampling.

Information Bar

% of stats approved for automation displays the percentage of statistics that are approved for automation, allowing you to determine if more statistics need to be approved for automation.

% of automated stats have collect jobs displays the percentage of automated statistics that have collect jobs, allowing you to determine if additional collect jobs are needed.

Recommendations Link

Displays a list of the recommendations when the link is clicked, allowing you to approve or reject recommendations.

Filters

Shows only rows that match your filter criteria.

▼ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display.

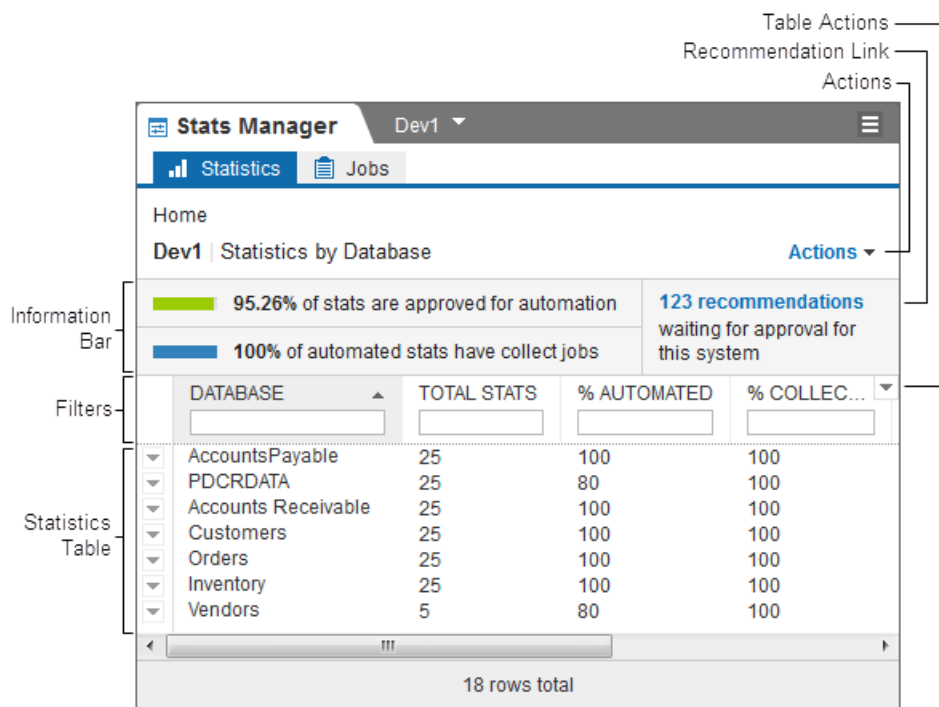
Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information, see [Summary Table Controls](#).

Statistics Table

Displays all objects with at least one statistic, or one recommendation that has not been approved or rejected. The table is configured using **Configure Columns** from the **Table Actions** list.

Following is an example of the **Statistics by Database** view.



Automating Statistics

When automated, statistics are approved for collection by collect jobs. You can automate existing statistics using an analyze job or you can manually automate statistics as described in the following table.

1. From the **Statistics** tab, do any of the following:

Scope	Description
System	To automate all statistics within a Teradata system: <ul style="list-style-type: none"> • From the Statistics by Database view, click Actions > Automate System.
Database	To automate all statistics within a single database, do one of the following: <ul style="list-style-type: none"> • From the Statistics by Object view, click Actions > Automate Database. • From the Statistics by Database view, click <input type="checkbox"/> next to the database name and click Automate. To automate all statistics within multiple databases: <ol style="list-style-type: none"> a. From the Statistics by Database view, click the Table Actions list and click Automate. b. Do one of the following: <ul style="list-style-type: none"> • Select the check box in the column heading to select all databases. • Select the check boxes for specific databases.
Object	To automate all statistics for a single object, do one of the following: <ul style="list-style-type: none"> • From the Statistics view, click Actions > Automate <object type>. • From the Statistics by Object view, click <input type="checkbox"/> next to the database name and click Automate. To automate all statistics for multiple objects:

Scope	Description
	<ol style="list-style-type: none"> From the Statistics by Object view, click the Table Actions list and click Automate. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all objects. Select the check boxes for specific objects.
Statistic	<p>To automate a single statistic, do one of the following:</p> <ul style="list-style-type: none"> From the Statistics view, click <input type="checkbox"/> next to the database name and click Automate. From the Statistic Details view, click Actions > Automate Statistic. <p>To automate multiple statistics for the same object:</p> <ol style="list-style-type: none"> From the Statistics view, click the Table Actions list and click Automate. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all statistics. Select the check boxes for specific statistics.

Deautomating Statistics

When you deautomate statistics, they stop being collected by collect jobs they were associated with and they cannot be added to jobs. Essentially, deautomating puts statistics in a state as if they were never automated. You can deautomate all statistics within a Teradata system, all statistics within one or more databases, all statistics for one or more objects, or one or more statistics for an object.

Deautomated statistics on an object will remain in a custom collect job scope until they are manually removed, but they will not be collected and the **Statistics** tab does not report an associated collect job.

Deautomated statistics continue to be collected if they were part of a custom collect list before the statistics were deautomated. The collect job is listed with the statistic, object, database, or system on the **Statistics** tab.

- From the **Statistics** tab, do any of the following:

Deautomate	Description
System	<p>To deautomate all statistics within a Teradata system:</p> <ul style="list-style-type: none"> From the Statistics by Database view, click Actions > Deautomate System.
Database	<p>To deautomate all statistics within a single database, do one of the following:</p> <ul style="list-style-type: none"> From the Statistics by Object view, click Actions > Deautomate Database. From the Statistics by Database view, click <input type="checkbox"/> next to the database name and click Deautomate. <p>To deautomate all statistics within multiple databases:</p> <ol style="list-style-type: none"> From the Statistics by Database view, click the Table Actions list and click Deautomate. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all databases. Select the check boxes for specific databases.
Object	<p>To deautomate all statistics for a single object, do one of the following:</p> <ul style="list-style-type: none"> From the Statistics view, click Actions > Deautomate <object type>.

Deautomate	Description
	<ul style="list-style-type: none"> From the Statistics by Object view, click <input type="checkbox"/> next to the database name and click Deautomate. <p>To deautomate all statistics for multiple objects:</p> <ol style="list-style-type: none"> From the Statistics by Object view, click the Table Actions list and click Deautomate. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all objects. Select the check boxes for specific objects.
Statistic	<p>To deautomate a single statistic, do one of the following:</p> <ul style="list-style-type: none"> From the Statistics view, click <input type="checkbox"/> next to the database name and click Deautomate. From the Statistic Details view, click Actions > Deautomate Statistic. <p>To deautomate multiple statistics for the same object:</p> <ol style="list-style-type: none"> From the Statistics view, click the Table Actions list and click Deautomate. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all statistics. Select the check boxes for specific statistics.

Editing Importance

Importance is a ranking assigned to each statistic that enables the Analytics Database order COLLECT STATISTICS statements when a collect list is generated for a collect job. When the Analytics Database automatically assigns the ranking it is called system importance. However, you can override the ranking of the statements and set it to low, medium, high, critical, or none by specifying the user importance.

- From the **Statistics** tab, select a database row.
- Select an object row.
The **Statistics** view appears.
- Do one of the following to edit the importance:
 - Click ☐ next to the statistic and click **Edit User Importance**.
 - Click the **Table Actions** list, and click **Edit User Importance**.
 - Click the statistic and click **Actions > Edit User Importance**.

- Select an option from the **User importance** list.

- Click **OK**.

View the updated ranking on the **Statistic details** tab in the **Statistic Details** view or in the **User Importance** column of the **Statistics** view.

Collect Settings

The collect settings allow you to:

- Set thresholds for collecting statistics based on how much time has passed or percentage the data set has changed since the last collection. This prevents over-collecting statistics.
- Set the sampling method to generate a histogram based on a representative sample of a subset of table rows instead of all table rows. This saves collection time. Sampling is useful if you are collecting statistics for very large tables with a lot of non-unique data. You can also set the maximum interval and maximum value length of the histogram.
- Set the maximum number of histogram intervals and maximum size for histogram values.

When setting the **Max interval**, the larger the number, the more optimal the granularity of the statistical data in the histogram. A finer granularity enables better single-table and join selectivity estimates for non-uniform data; however, the larger the number of intervals, the larger the size of the histogram, which can increase query optimization time.

When setting the **Max value length**, a larger maximum value size causes Analytics Database to retain the value until the specified maximum is reached. This can enable better single-table and join selectivity estimates for skewed columns; however, increasing the maximum value size also increases the size of the histogram which can increase query optimization time. For important benefits and limitations, see *Teradata® Database SQL Data Definition Language - Detailed Topics*.


The following describes the settings.

Setting	Description
System-defined	Use if you want the Analytics Database to set the thresholds, sampling methods, and histogram interval and length.
User-defined	Use if you want to specify these settings.
None	Use if you do not want to set thresholds or a sampling method.

There are two ways to change collect settings. You can either edit a single statistic or multiple statistics. Editing a single statistic gives you full control of all settings. Editing multiple statistics limits the settings you can change to **System-defined** and **None**; however, you can edit collect settings for all databases, objects on a single database, or a single object.

Editing Collect Settings for a Single Statistic

You can only edit collect settings for a single statistic.

1. From the **Statistics** tab, select a database row.
2. Select an object row.
The **Statistics** view appears.
3. Do one of the following to edit a single statistic:
 - Click  and click **Edit Collect Settings**.
 - Click a statistic row which displays the **Statistic Details** view, and click **Actions > Edit Collect Settings**
4. To define the collect settings, do any of the following:

- Select **System-defined** if you want Analytics Database to decide the settings.
- Select **User-defined** if you want to specify the settings.
- Select **None** if you do not want thresholds or a sampling method.

5. If you select **User-defined**, edit the following settings:

Option	Setting	Description
Threshold	Collect age threshold	A collection is not performed if the last collection was done within the specified number of days.
	Data change threshold	A collection is not performed unless the data set has changed at or beyond the specified percentage since the last collection.
Sampling	Sampling method	A histogram is generated using the specified percentage of the actual data. For more information about sampling, see <i>Teradata® Database Administration</i> .
Histogram	Max interval	Specify the maximum number of histogram intervals to be used for the collected statistics.
	Max value length	Specify the maximum size for histogram values. For single-character statistics on CHARACTER and VARCHAR columns, specify the number of characters, otherwise specify the number of bytes.

6. Click **OK**.

Editing Collect Settings for Multiple Statistics

You can edit collect settings for multiple statistics within a Teradata system, multiple statistics within one or more databases, or multiple statistics for one or more objects.

1. From the **Statistics** tab, do any of the following:

Settings	Description
System	To edit collect settings for multiple statistics within a Teradata system: <ul style="list-style-type: none"> • From the Statistics by Database view, click Actions > Edit Collect Settings.
Database	To edit collect settings for multiple statistics within a single database, do one of the following: <ul style="list-style-type: none"> • From the Statistics by Object view, click Actions > Edit Collect Settings. • From the Statistics by Database view, click <input type="checkbox"/> next to the database name and click Edit Collect Settings. <p>To reset collect settings for multiple statistics within multiple databases:</p> <ol style="list-style-type: none"> From the Statistics by Database view, click the Table Actions list and click Edit Collect Settings. Do one of the following: <ul style="list-style-type: none"> • Select the check box in the column heading to select all databases. • Select the check boxes for specific databases.

Settings	Description
Object	<p>To edit collect settings for a single object, do one of the following:</p> <ul style="list-style-type: none"> From the Statistics view, click Actions > Edit Collect Settings. From the Statistics by Object view, click <input type="checkbox"/> next to the database name and click Edit Collect Settings. <p>To edit collect settings for multiple statistics for multiple objects:</p> <ol style="list-style-type: none"> From the Statistics by Object view, click the Table Actions list and click Edit Collect Settings. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all objects. Select the check boxes for specific objects.

- To define the collect settings, do any of the following:
 - Select **System-defined** if you want Analytics Database to decide the settings.
 - Select **None** if you do not want thresholds, a sampling method, or histogram.
- Click **OK**.

Recommendations

Statistics recommendations are suggestions, produced by analyze jobs, about actions to be taken on statistics. There are four types of recommendations: collect, automate, reactivate, and deactivate.

Collect recommendations

Suggestions about columns or indexes that the Analytics Database will benefit from having statistics collected on.

Reactivate recommendations

Suggestions about statistics that were deactivated but are now being used again by Analytics Database.

Deactivate recommendations

Suggestions about statistics that are no longer being used by the Analytics Database and should no longer be collected.

Automate recommendations

Suggestions about statistics to approve for collection by collect jobs.

After an analyze job runs, you can approve or reject recommendations resulting from that job depending on the options you selected when you created or edited an analyze job. If you selected **Require review before applying recommendations**, you can manually approve or reject recommendations. If you did not select this option, the recommendations are automatically approved.

To see recommendations, click the recommendations link from the **Statistics** tab. The **Recommendations for System** view appears. You can drill down three levels to see recommendations in databases, for objects, or for a single statistic.

The screenshot shows the 'Stats Manager' interface. At the top, there's a header with 'Stats Manager' and a dropdown menu set to 'Dev1'. Below the header, there are two tabs: 'Statistics' (active) and 'Jobs'. The main content area shows a breadcrumb 'Home »' followed by 'Dev1 | Recommendations for System'. To the right of this is a dropdown menu set to 'Pending'. Below this, there are five tabs: 'All (10)' (active), 'Collect (2)', 'Reactiv... (0)', 'Deactiv... (0)', and 'Automate (1)'. A message states 'There are 0 total recommendations in the databases below.' At the bottom, there are four input fields labeled 'DATABASE', 'ALL', 'COLLECT', and 'AUTOMATE', each with a dropdown arrow.

Recommendations display in each view as described in the following table.

Recommendations View	Description
System	All databases on the system
Database	All objects on the database
<i>object</i>	Object, such as a table or index, by statistic
Statistic	Details for a single statistic


Use the filter, located in the upper right corner, to show either the pending recommendations that are awaiting approval or the rejected recommendations.

The **All** tab displays the total number of collect, automate, reactivate, and deactivate recommendations combined. The remaining tabs displays the collect, reactivate, deactivate, and automate recommendations in the databases or in the objects, depending on the view.

Approving Recommendations

After an analyze job runs, you can approve recommendations resulting from that job.

1. From the **Statistics** tab, click the recommendations link.
2. Do one of the following:

Option	Description
Approve one recommendation	Do one of the following: <ul style="list-style-type: none"> • From any tab except the All tab, click  next to a recommendation and click Approve. • From the Recommendations for statistic view, click Actions > Approve.
Approve one or more recommendations	<ol style="list-style-type: none"> a. From any tab except the All tab, click the Table Actions list and click Approve. b. Do one of the following: <ul style="list-style-type: none"> • Select the check box in the column heading to select all statistics.

Option	Description
	<ul style="list-style-type: none"> Select the check boxes for specific statistics. c. Click Next . d. Click Approve .
Approve all recommendations	a. From any tab except the All tab, click Approve all . b. Click Approve .

- Click **Close**.
The recommendations are approved.

Rejecting Recommendations

After an analyze job runs, you can reject recommendations resulting from that job.

- From the **Statistics** tab, click the recommendations link.
- Do one of the following:

Option	Description
Reject one recommendation	Do one of the following: <ul style="list-style-type: none"> From any tab except the All tab, click <input type="checkbox"/> next to a recommendation and click Reject. From the Recommendations for statistic view, click Actions > Reject.
Reject one or more recommendations	a. From any tab except the All tab, click the Table Actions list and click Reject . b. Do one of the following: <ul style="list-style-type: none"> Select the check box in the column heading to select all statistics. Select the check boxes for specific statistics c. Click Next . d. Click Reject .

- Click **Close**.
The recommendation is rejected.

Stats Manager Metrics

You can enable metrics that appear in the **Statistics** and **Jobs** views. Metrics available for selection and display are listed as following.

Note:

In this table, **Recommendations for Object** is used to represent all of the object types available in statistics metrics, such as database, table, index, and so on.

Metric	Description	View
% Automated	Percentage of automated statistics associated with the associated object, such as system, database, or table	Statistics by Database Statistics by Object
% Collect Coverage	Percentage of automated statistics associated with a collect job	Statistics by Database
% Data Change	Maximum estimated percent data change since the last collection for automated statistics as determined by an analyze job	Statistics by Object Statistics Statistics Details
Active	Indicates if a statistic has been deactivated as the result of an analyze job because it was not being used by the Optimizer	Statistics
All	Total number of recommendations	Recommendations for Object
Alert Name	Name of the alert	Alert definitions
Alert Enabled	<ul style="list-style-type: none"> true - alert is enabled. Alert is triggered for the selected jobs with Error state. false - alert is disabled. 	Alert definitions
Alert Severity	The severity of the alert can be: Critical, High, Medium, and Low.	Alert definitions
Alert Action Name	The name of the action. These actions are defined in the Alert Setup portlet.	Alert definitions
Automate	Total number of automate recommendations	Recommendations for Object
Automated	Indicates whether or not the statistic can be added to a collect job	Statistics Analyze Job Report
Automated Stats	Number of automated statistics associated with the object	Statistics by Database Statistics by Object
Collect	Total number of collect statistics recommendations	Recommendations for Object
Collect Age	Number of days since the statistic was last collected	Statistics Analyze Job Report
Collect Age Threshold	Number of days since the last collection before the statistic is considered stale	Analyze Job Report
Collect Jobs	List of collect jobs that include at least one statistic within the database or object	Statistics by Database Statistics by Object
	Associated collect jobs responsible for collecting this statistic	Analyze Job Report
Collect List Total	Number of COLLECT STATISTICS statements in the collect list	Collect Job History

Metric	Description	View
Collect Text	COLLECT STATISTICS statement	Collect Report Details
Columns	Columns associated with the statistic	Statistics Statistic Details Analyze Job Report Recommendations for Object
CPU Secs	Processing power, in CPU seconds, used to run the job	Collect Job History
Cumulative CPU Secs	Total processing power used to process collections since the job started collecting from the beginning of the collect list	Collect Job History
Cumulative Duration	Total time used to process collections since the job started collecting from the beginning of the collect list	Collect Job History
Cumulative Error Count	The total error count for Collect Job. This count is only displayed for collect jobs with state - RUNNING, PARTIALLY COMPLETE or COMPLETE.	Job Definitions
Cumulative Submitted Collects	Count of COLLECT STATISTICS statements submitted since the job started collecting from the beginning of the collect list	Collect Job History
Database	Name of the database	Statistics by Database Recommendations for Object Collect Job Report Analyze Job Report
Data Change Threshold	Percent the data can change before being considered stale	Analyze Job Report
Deactivate	Total number of deactivate recommendations	Recommendations for Object
Duration	How long the COLLECT STATISTICS statement took to run	Job History Collect Job Report Analyze Job Report
	Time allocated to run a collect job	Job Definitions
End	Date and time the job ended	Job History
	Time the statistic collection ended	Collect Job Report
Estimate Error Frequency	Number of times the lack of this statistic resulted in a cardinality estimation error	Recommendations for Object Analyze Job Report
Inactive Stats	Number of inactive statistics in the associated object or database	Statistics by Database Statistics by Object

Metric	Description	View
Index	Indicates whether a statistic was collected on an index	Statistics Statistics Details Analyze Job Report Recommendations for Object
Job	Name of the job	Job Definitions
Last Access	Last time the statistic was accessed	Statistics Recommendations for Object
Last Collect	Date and time the statistic was last collected	Statistics Statistic Details
Last Duration	Amount of time it took for the most recent instance of the associated job to run	Job Definitions
Last End	Time the most recent instance of the associated job ended	Job Definitions
Last Start	Time the most recent instance of the associated job started	Job Definitions
Last Submit	Last time statistics were submitted for collection	Statistics by Object Statistics Statistic Details
Max Collect Age	Number of days since the statistic with the oldest Last Collect date was collected	Statistics by Database Statistics by Object
Max Submit Age	Number of days since the automated statistic with the oldest Last Submit date was submitted	Statistics by Database
Missing Frequency	Number of times the statistic would have been used if it ran	Recommendations for Object
Missing Rank	Rank of this missing statistic compared to all other missing statistics in terms of Missing Frequency	Analyze Job Report
Most Recent	Time of the most recent recommendation	Recommendations for Object
Next Start	Time the job is scheduled to run again	Job Definitions Analyze Job Report
Null Count	Number of null values in a column	Statistic Details
Number of Jobs	The total number of jobs selected for the alert.	Alert definitions
Object	Name of the object	Statistics by Object Collect Job Report Analyze Job Report
Oldest Collect	Oldest collect time of any statistic on the object	Statistics by Object

Metric	Description	View
Origin	How the statistic became automated. Values for origin are as follows: <ul style="list-style-type: none"> Automate Analyze 	Statistics Statistic Details
Query Log Analysis	(Teradata Database 15.10 and later) Progress bar	Job Definitions Analyze Job Report
Reactivate	Total number of reactivate recommendations	Recommendations for Object
Recommendation	Indicates whether the event resulted in a recommendation	Analyze Job Report
Recommendation Time	Time the recommendation was made	Statistics Recommendations for Object
Recommendation Type	Type of recommendation made	Statistics Recommendations for Object
Row Count	Number of rows in the object as determined by the last collect	Statistics by Object Statistic Details
Skipped Collections	Number of consecutive times the COLLECT STATISTICS statement was submitted but the collect was skipped because the collect thresholds were not met	Statistics by Object Statistics Statistic Details
Stale stats	Number of automated statistics in the object with a state of stale or missing data	Statistics by Database Statistics by Object
Staleness Analysis	(Teradata Database 15.10 and later) Progress bar	Job Definitions Analyze Job Report
Start	Date and time the job started	Job History Collect Job Report Analyze Job Report
	Time the statistic collection started	Collect Job Report
	Time a job is scheduled to start	Job Definitions
State	Current state of the statistic. Values for state are as follows: <ul style="list-style-type: none"> Current Missing Stale 	Statistics
	Collect state of the statistic. Values for state are as follows: <ul style="list-style-type: none"> Collected Not Collected 	Collect Job Report

Metric	Description	View
	<ul style="list-style-type: none"> Error Skipped Pending 	
	<p>Current state of the job. Values for state are as follows:</p> <ul style="list-style-type: none"> Aborted - Jobs that were aborted by user Complete - Jobs that ran to completion successfully Error - Jobs that encountered an error during processing Partially Complete - Collect jobs that did not complete submitting all COLLECT STATISTICS statements in their associated collect list before their allotted time expired, but were otherwise successful Running - Jobs that have been started and have not yet ended 	<p>Job Definitions</p> <p>Collect Job History</p> <p>Collect Job Report</p>
Statistic	Name of the statistic	<p>Statistics</p> <p>Analyze Job Report</p> <p>Recommendations for Object</p>
Submit Age	Number of days since an automated statistic was last submitted for collection	<p>Statistics by Object</p> <p>Statistics</p>
Submit Order	Order in which collect list items are submitted	Collect Job Report
Submitted Collections	Progress bar	<p>Collect Job History</p> <p>Job Definitions</p>
System Importance	Ranking automatically assigned to each statistic that enables Analytics Database to order COLLECT STATISTICS statements when a collect list is generated for a collect job. This ranking can be overridden by editing importance.	<p>Statistics</p> <p>Statistic Details</p> <p>Analyze Job Report</p> <p>Recommendations for Object</p>
Total Stats	Number of statistics associated with the object	<p>Statistics by Database</p> <p>Statistics by Object</p>
Type	<p>Type of database object. Values for type are as follows:</p> <ul style="list-style-type: none"> Database User Table Join Index Hash Index 	Statistics by Object
	Type of job. Values for type are as follows:	Job Definitions

Metric	Description	View
	<ul style="list-style-type: none">AnalyzeCollect	
Unique Value Count	Number of unique values in the expression list	Statistic Details
User Importance	Ranking set by the user that overrides the ranking that is automatically assigned by the System Importance	Statistics Statistic Details

System Health






System Health Overview

The **System Health** portlet monitors the status and key metrics of all monitored systems against defined thresholds. This portlet reports status with customized states or the following default states: healthy, degraded, critical, down, or unknown. The portlet also uses customized or default tooltips. These states and tooltips enable you to investigate metrics exceeding healthy thresholds.

The Teradata Viewpoint Administrator enables system health metrics, configures the degraded and critical thresholds, and customizes the states and tooltips for your system in the **Monitored Systems** portlet. The metrics and thresholds are usually selected to highlight an unusual load on the system that has the potential to impact overall performance.

System Health View

The **System Health** view displays a summary of states for one or more systems, using colored text and icons to represent system health. From this view, click anywhere on a status icon or text to investigate metrics for that system in the **System Health** details view.

Default Icon	Text Color	Default States	Definition
	Green	Healthy	All metrics are within <i>healthy</i> ranges.
	Yellow	Degraded	At least one metric exceeded a <i>degraded</i> threshold.
	Red	Critical	At least one metric exceeded a <i>critical</i> threshold.
	Black	Down	The selected system is <i>down</i> .
	Gray	Unknown	Status of the selected system is <i>unknown</i> .

If your system does not use the default states and tooltips, the Teradata Viewpoint Administrator customized the states and tooltips. You can customize the icon style using the **Icons** tab in the **Settings** view.

Note:




Metrics that are configured by the Teradata Viewpoint Administrator as View Only are not used in determining system health.

System Health Details View

The **System Health** details view displays detailed statistics and information about each metric to evaluate the overall health of a system. If the Teradata Viewpoint Administrator defines metrics as View Only, these metrics are displayed with a gray background and are not used in determining system health. If any connection issue happens with the system because of a wrong configuration, then a relevant error message appears with the system status.

Status Icons

Show the status of the metrics:

-  represents metrics that are defined as View Only
-  represents metrics that reached a critical state
-  represents metrics that are degraded

Metric

Lists the names of the monitored metrics, using color to indicate metric health. Displays metric descriptions when you hover over the name.

You can click the following metrics to access further details when at least one component is down or when one service is in a bad or concerning state:

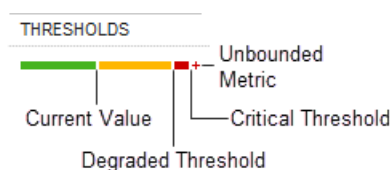
- **Components Down** (HDP)
- **Services Bad** (CDH)
- **Services Concerning** (CDH)

Value

Shows the metric value at the last system sampling.

Thresholds

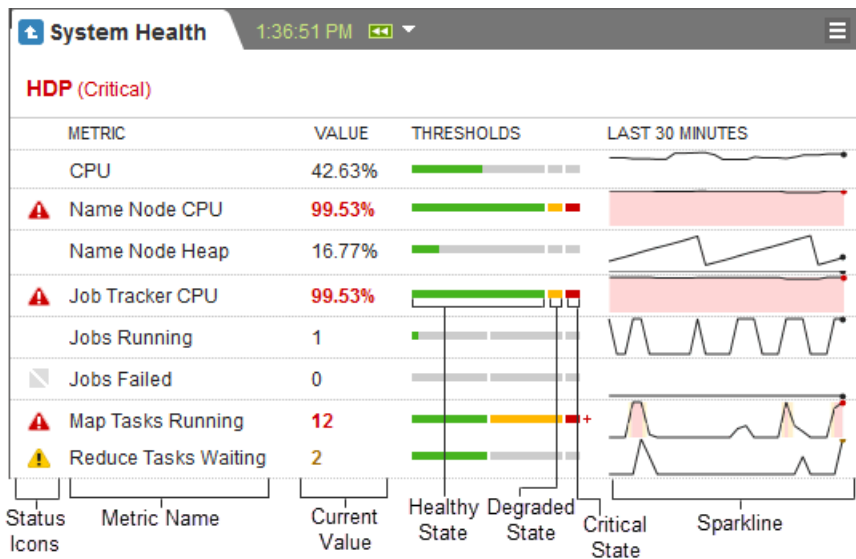
Shows the current metric value in terms of degraded and critical thresholds, indicated by gaps in the bar. Displays threshold values and default or custom tooltips when you hover over the bar.



The + at the end of some bar graphs indicates an unbounded metric, which has numerical values with no upper limit. The + is red when the unbounded metric value exceeds 1.1 times the critical threshold.

Last *nn* Minutes

Displays a sparkline with the trend in metric values over the interval selected in the **Settings** view of 15, 30, or 60 minutes.



Viewing Component Down Details

When the Component Down metric has a value of one or greater, you can access details that show the extent of your exposure to components that are down.

This option is not available on an Aster system.

- From the **System Health** details view, click the **Component Down** metric name or value. A list of down components and their host names appears.

Viewing Services Bad or Services Concerning Details

When the Services Bad or Services Concerning metric (CDH) has a value of one or greater, you can access details of the services in these states.

- From the **System Health** details view, click any of the following:
 - Services Bad** metric name or value
 - Services Concerning** metric name or value
 A list of bad or concerning services and their host names appears.

System Health Details View Metrics

Metrics appearing in your **System Health** details view are selected for you by the Teradata Viewpoint Administrator. Available metrics differ depending on the type of system you choose.

Teradata System Metrics

Metric	Description
Active Sessions	Number of sessions with active queries
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP
AMP Worker Tasks	Average number of AMP worker tasks in use on each AMP
Component Down	Number of components, such as BYNETs or AMPs, that are down
(CPU) System	Average CPU time spent executing kernel system calls or servicing I/O and timer hardware interrupts
(CPU) User	Average CPU time spent executing code on behalf of operating system user processes
CPU Utilization	Average node CPU use. CPU is calculated as the sum of the user CPU, system CPU, and wait I/O usage percentages, depending on which of these are enabled by the Teradata Viewpoint Administrator in the Monitored Systems portlet
(CPU) Wait I/O	Average percentage of CPU time spent waiting for I/O
DBC Disk Space	Available DBC disk space in use
Enabled CPU Usage Note: In Viewpoint Mobile view this metric is currently not available.	CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled
Max Disk by AMP	Available disk space currently in use
Max Spool by AMP	Available spool space in use
Memory Used	Represents an increase or decrease in the node-level memory. A null value or negative number reports as 0.
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of disk use on the busiest node to the average node
System Heartbeat	Canary query showing the response time of the most recent system heartbeat in milliseconds

Metric	Description
Total Disk Space	Percentage of total disk space currently in use
Additional Canaries	Additional canary queries that are defined for the system and used as part of the monitoring of system health

Aster System Metrics

Metric	Description
Active Sessions	Number of users and applications currently connected to database
Component Down	Number of nodes that are not available
Component Passive	Number of nodes not processing queries but can be made ready to process queries when needed
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Memory	Average node memory use
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Queen Disk Space	Percentage of used disk space on the queen node
Replication Factor	Number of copies of the user data
Total Space	Percentage of used space to overall storage capacity

Teradata Machine Learning Engine Metrics

Metric	Description
Active Sessions	Number of users and applications currently connected to the engine
CPU	Average worker or queen CPU use
Pod CPU Skew	Comparison of CPU use on the busiest worker or queen to that of the average one

Hadoop System Metrics

CDH

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Services Bad	Number of services in a critical state
Services Concerning	Number of services in a degraded state
Total Space	Percentage of used space to overall storage capacity

HDP 2.1 and later

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances

Metric	Description
Components Down	Number of services not started
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Total Space	Percentage of used space to overall storage capacity

HDP 1.3

Metric	Description
Blocks Corrupt	Number of blocks whose replicas are all corrupt
Blocks Missing	Number of blocks with no replicas anywhere in the cluster
Components Down	Number of service components not running
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.
Jobs Failed	Number of jobs that failed
Jobs Running	Number of jobs currently executing in the system
Job Tracker CPU	CPU use for the node running the jobtracker service
Map Tasks Running	Number of map tasks executing in the system
Map Tasks Waiting	Number of map tasks waiting to execute
Max Disk by Node	Amount of used disk space on the node with the most disk space in use
Name Node CPU	Node CPU use for the node running the namenode service
Name Node Heap	Percentage of heap space used in the namenode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node

Metric	Description
Reduce Tasks Running	Number of reduce tasks executing in the system
Reduce Tasks Waiting	Number of reduce tasks waiting to execute
RPC Latency JT	Average wait time in queue for jobtracker service calls
RPC Latency NN	Average wait time in queue for namenode service calls
Total Space	Percentage of used space to overall storage capacity

Settings View

The **Settings** view allows you to customize the portlet to display information that is most useful to you.

The following tabs allow you to customize your settings.

Systems

Select the systems to monitor in the **System Health** and **System Health** details views.

Icons


Select an icon style to use in the **System Health** view.

Trend Interval

Select the time frame for the **System Health** sparkline. Choose from among 15, 30, or 60 minutes of historical data.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.


Selecting Systems to Monitor

Use the **Systems** tab in the **Settings** view to select the systems to monitor in the summary and details views.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Systems** tab.
3. Select the check boxes of systems to display.
4. Click **OK**.

Selecting the System Health Icon Style


Use the **Icons** tab in the **Settings** view to select the icon style to display in the **System Health** view.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Icons** tab.
3. Select the icon style to use in the **System Health** view.
4. Click **OK**.

Setting the System Health Trend Interval

Use the **Trend Interval** tab in the **Settings** view to set the length of time used to plot the sparklines in the **System Health Details** view. This interval applies to all metrics selected for this view.

The **Example** pane provides a sample of how the **System Health** view might look after the change.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Trend Interval** tab.
3. From the **Interval** list, select an interval.
4. Click **OK**.

Today's Statistics

Today's Statistics Overview

The **Today's Statistics** portlet allows you to monitor current resource-use statistics and compare them to statistics collected during the same period one or more weeks in the past. Additionally, you can monitor and compare the performance of multiple Teradata systems by adding a **Today's Statistics** portlet to the portal page for each monitored system.

The Teradata Viewpoint Administrator typically uses the **Today's Statistics** portlet to determine if users are getting value from their Teradata system. Data is presented in a table, allowing you to compare metric values collected recently (today and as far in the past as yesterday) to values collected during the same time period up to 999 weeks in the past. This view is refreshed every 2 minutes with displayed values updated for each user-selected metric.

Today's Statistics View

The **Today's Statistics** view displays a statistical analysis of system-resource use over a user-defined time period. A table lists the monitored metrics and compares current statistics to statistics from the same time period 1 or more weeks in the past. Metrics are grouped into categories so that metric data collected at the same frequency is displayed together in the view. The view is refreshed every 2 minutes.

Table columns provide the following information:

STATISTIC

The name of the monitored metric.

LAST X HOUR(S)

The value of the metric from a user-specified number of hours ago, up to 999 hours in the past, until the current time.

SAME PERIOD X WEEK(S) AGO

The value of the metric over the same time period as **LAST X HOUR(S)**, up to 999 weeks in the past.

SINCE X AM/PM TODAY/YESTERDAY

The value of the metric calculated from a user-specified hour of the day, either today or yesterday, until the current time.

SAME PERIOD X WEEK(S) AGO

The value of the metric over the same time period as **SINCE X AM/PM TODAY/YESTERDAY**, up to 999 weeks in the past.

Today's Statistics				
Dev1 12:13:03 PM				
STATISTIC	LAST HOUR	SAME PERIOD 1 WEEK AGO	SINCE 12 AM TODAY	SAME PERIOD 1 WEEK AGO
Active Sessions	0.533	1.167	0.627	1.548
AMP CPU Skew	4.752%	5.325%	14.95%	3.321%
AMP I/O Skew	4.501%	7.645%	16.97%	4.442%
AWT	1.029	1.292	4.705	1.285
CPU	62.87%	67.62%	32.99%	72.31%
System CPU	5.058%	4.878%	2.751%	5.549%
Total Disk Sp...	1.855%	1.99%	2.413%	1.92%
User CPU	57.82%	62.74%	30.24%	66.76%
Wait I/O CPU	2.643%	2.086%	2.316%	1.468%
9 rows total				

To display statistical data in the **Today's Statistics** view, the Teradata Viewpoint Administrator must first enable the Query Count, System Stats, and Canary Queries data collectors. The Teradata Viewpoint Administrator also specifies the data collection rate for each collector.

Factors external to Teradata Viewpoint can also affect the collection and display of statistical data:

- *Query Logging.* The Teradata System Administrator must enable query logging on the Teradata system.
- *Data Collection.* Query data must have been collected during all time periods selected on the **Sampling Interval** tab in the **Settings** view. Teradata Viewpoint collectors and query logging must have been enabled during all selected time periods in order for **Today's Statistics** data to display correctly.

Use the **Settings** view to select the system and metrics category to monitor, set metrics thresholds, and set the sampling interval. The column headings depend on the selections made in the **Settings** view and how frequently data for each metric is collected. For example, the heading **LAST 2 HOURS** appears when the sampling interval is 2 hours and data for the previous 2 hours is available.

Today's Statistics View Metrics

Information about available metrics is provided in the tables. Use the **Metrics** tab in the **Settings** view to select metrics for display in the **Today's Statistics** view.

The Teradata Viewpoint Administrator must enable the appropriate collectors to calculate and display statistical data in the **Today's Statistics** view. When the Canary Queries Collector is the source for a metric, the data displayed is always the heartbeat query for the selected system.

System Metrics

The System Stats collector gathers system metrics. System metrics are calculated as an average.

Metric	Description	Type
Active Sessions	Sessions with active queries	Number
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP	Percent

Metric	Description	Type
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP	Percent
AWT	Average AMP worker tasks (AWTs) in use on each AMP	Number
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.	Percent
System CPU	Average CPU time spent in system mode	Percent
Total Disk Space	Total disk space currently in use	Number
User CPU	Average CPU time spent in user mode	Percent
Wait I/O CPU	Average CPU time spent waiting for I/O	Percent

Query Metrics

The Query Count collector gathers query metrics.

Metric	Description	Type
Concurrency	Average number of queries executing simultaneously during the hour	Number
First Response Time Avg	Average time in milliseconds it takes to receive the first response	Number
First Response Time Max	Maximum time in milliseconds it takes to receive the first response	Number
First Response Time Min	Minimum time in milliseconds it takes to receive the first response	Number
First Response Time Stddev	Standard deviation time in milliseconds it takes to receive the first response	Number
Query Count	Queries logged in dbc.QryLog, or the sum of Querycount from dbc.QryLogsummary during the sample period	Number
Spool Usage Avg	Average amount of spool in bytes the query used	Number
Spool Usage Max	Maximum amount of spool in bytes the query used	Number
Spool Usage Min	Minimum amount of spool in bytes the query used	Number
Spool Usage Stddev	Standard deviation of amount of spool in bytes the query used	Number

Performance Metrics

The Canary Queries collector gathers performance metrics. Performance metrics are calculated as an average.

Metric	Description	Type
Response Time	Average response time in milliseconds of the selected canary query	Number
Rows Returned	Average number of rows returned by the selected canary query	Number

Settings View

The **Settings** view allows you to customize the portlet to monitor key metric trends to determine if you are getting value from your Teradata system. Only one system can be selected for each portlet instance.

The following tabs allow you to customize your settings.

Metrics


Select appropriate metrics, by category, for display in the view, and set values for highlighting metrics beyond a specified threshold. Metric values that exceed the threshold are displayed in the exception color.

Sampling Interval

Specify the time periods from which statistics are displayed. Set the number of hours (up to 999) prior to the current time, and define a comparison period up to 999 weeks in the past.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.

1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Setting Metric Thresholds


You can select a metrics category to monitor and set thresholds to trigger the highlighting of values that exceed the thresholds.

1. Click  in the portlet frame and select **Settings**.

2. Click the **Metrics** tab.
3. Select a metrics category.
4. In the **Metric** column, select check boxes for the metrics you want to monitor in the view.
5. [Optional] For each metric selected, enter a **Threshold** value.
Values exceeding the threshold are highlighted in the view using the exception color.
6. Click **OK**.

Setting the Sampling Interval

You can set the length of time used to calculate the statistics shown in the **Today's Statistics** view. The settings are reflected in the column headings. The **Sampling Interval** applies to all metrics selected for display.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Sampling Interval** tab.
3. Enter the number of hours (up to 999) for **Statistics will be sampled during the last X hours**.
4. Enter the **since** sampling interval (up to 48 hours) as follows:
 - a. Select the time of day that begins the interval.
 - b. Select **Today** or **Yesterday** from the list.
5. Enter a value for the comparison period (up to 999 weeks) for **Statistics will be compared to the same period X weeks ago**.
6. Click **OK**.

TVS Monitor

TVS Monitor Overview

The **TVS Monitor** portlet allows you to view statistics on data temperature and storage grade of cylinders allocated in Teradata Virtual Storage (TVS). Statistics reflect current relationships between data temperature and storage grade, and historical trends in the management of storage grades based on data temperature. This portlet is disabled by default and must be enabled using **Portlet Library**.

Temperature and grade metrics help assess storage requirements for your data, as well as fine tune the balance between disk performance and data-access demands. Suppose that statistics indicate an increasing proportion of hot data allocated to slow-grade storage. If fast storage is currently full, the Teradata System Administrator might consider upgrading to faster storage. If fast storage is not full, the Teradata System Administrator should verify that data migration and metrics collection have been enabled to make sure the slow migration of hot data over time.

TVS Monitor Concepts

Storage Grade

Storage grade refers to the access rate capabilities of disk storage. Access speed is graded as slow, medium, or fast. Performance criteria for storage grade include physical speed of disk read-write access, electronics, and buffer memory.

Slow-grade storage involves more I/O time than fast-grade storage. Cylinders on physical disk drives are often associated with slow- and medium-grade storage. Cylinders on solid-state drives can be associated with fast-grade storage.

Typically, Teradata Virtual Storage optimizes cylinder allocation by reserving fast-grade storage for frequently accessed data, and slow-grade storage for infrequently accessed data. Medium-grade storage is allocated for all other data-access frequencies.

Data Temperature

Data temperature refers to the frequency of access to data stored in space types managed by Teradata Virtual Storage (TVS). Data is considered hot, warm, or cold.

Data temperature is not static, but evolves over time, based on data-access patterns tracked by TVS. Data temperature increases the more often it is accessed, and decreases as access becomes less frequent.




Space Types

A *space type* is a classification of how disk space is allocated in Analytics Database. Each space type, such as perm, spool, or temp, is conventionally associated with a particular kind of stored data. For example, the perm space type is disk space allocated to storage of tables, indexes, and permanent journals in the database, and typically is the largest space type in the system.

Teradata Virtual Storage reports on how space types are dynamically managed, according to patterns of data access and the grade of the storage device.

Selecting a System and View


You can select a system and then select different views of the relationships among data temperature, storage grade, and space type.

1. Select ▼ in the portlet frame and select a system.
2. Select a view:
 - Click  to see the **Cylinder Distribution by Grade** view.
 - Click  to see the **Historical Cylinder Distribution by Grade** view.
 - Click  to see the **Cylinder Allocation by Temperature and Grade** view.

Viewing Cylinder Distribution by Grade

The **Cylinder Distribution by Grade** view displays graphs of the current proportional allocation of cold, warm, and hot data in slow-, medium-, and fast-grade storage, as well as the total number of cylinders assigned to each storage grade. The dashed line in each graph marks peak allocation for the storage grade in number and percent of cylinders, aggregated over the past 30 days.

The percent values in the graphs and balloon text that are displayed for this view indicate the percentage of the total number of cylinders in the selected grade.



1. From the toolbar, click .
2. Click a graph to display its metrics in the table beneath the graphs.
 - Click the **Slow Storage** graph to display metrics on total slow cylinders, data temperatures, and peak allocation for slow storage.
 - Click the **Medium Storage** graph to display metrics on total medium cylinders, data temperatures, and peak allocation for medium storage.
 - Click the **Fast Storage** graph to display metrics on total fast cylinders, data temperatures, and peak allocation for fast storage.

Viewing Historical Cylinder Distribution by Grade

The **Historical Cylinder Distribution by Grade** view displays data temperature trends over time in each storage grade. The horizontal axis marks days and the vertical axis marks number of cylinders, in millions.

You can change how temperature allocation metrics and total number of cylinders are displayed for a specific point in time.


The percent values that are displayed under the graph for this view indicate the percentage of the total number of cylinders in the selected grade.

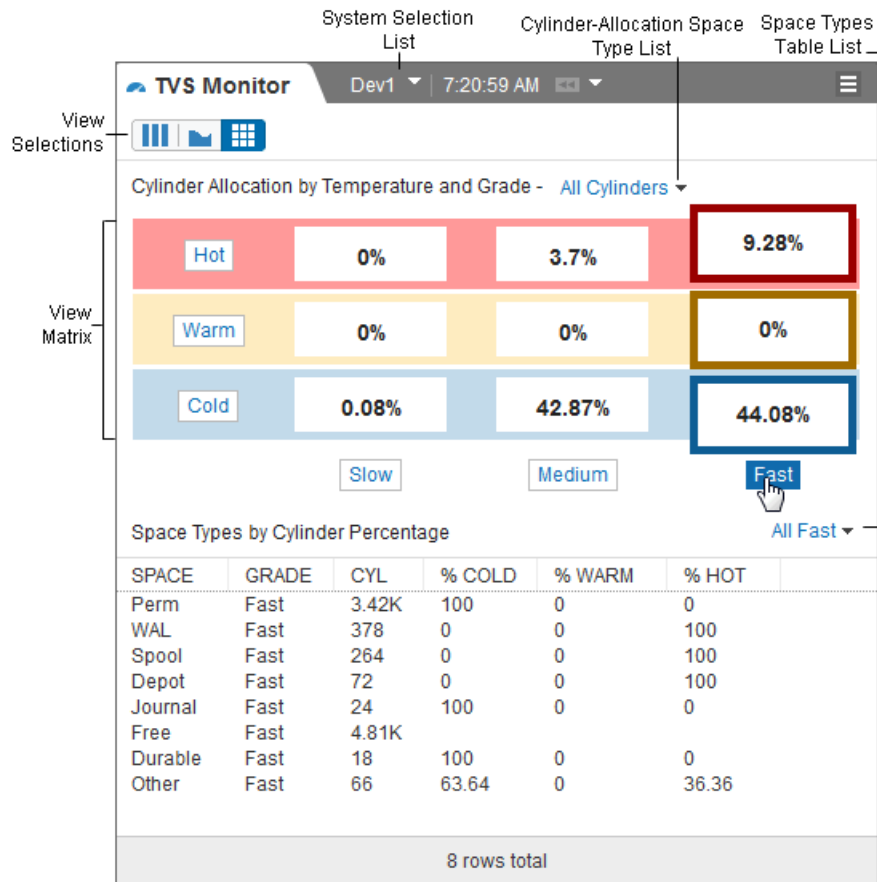
1. From the toolbar, click .
2. Hover over the graph to display and move the time marker.
As you move the time marker, the display of data in the table beneath the graph changes.
3. [Optional] Select a grade from the list before the graph to show the proportion of cylinders, allocated by temperature, for the grade.
4. [Optional] Select a time period from the list on the right to show the proportion of cylinders, allocated by temperature, over the previous week or month.
5. [Optional] Click  in the toolbar to export the on-screen data.

Viewing Cylinder Allocation by Temperature and Grade

The **Cylinder Allocation by Temperature and Grade** view displays the current relationships among data temperature, storage grade, and space type. Change the display of the matrix and the **Space Types by Cylinder Percentage** table to see allocation metrics for a specific combination of temperature, grade, or space.

The percent values that are displayed in the matrix for this view indicate the percentage of the total of non-allocated space of the selected cylinder type.

1. From the toolbar, click .
2. Select a cylinder allocation space type from the list before the matrix to display a space type in the matrix and in the space-types table.
3. To filter the data shown in the space-types table, do one of the following:
 - Click a row button (**Hot**, **Warm**, or **Cold**) in the matrix to filter the space-types table by data temperature allocation.
 - Click a column button (**Slow**, **Medium**, or **Fast**) in the bottom row of the matrix to filter the space-types table by storage grade.
 - Click any cell in the matrix to filter the space-types table by a temperature-grade combination.
4. Select a temperature, grade, or temperature-grade combination from the list before the space-types table to filter space types by temperature, grade, or a temperature-grade combination.



Space Types by Cylinder Percentage Table

The **Space Types by Cylinder Percentage** table in the **Cylinder Allocation by Temperature and Grade** view provides metrics about space types, storage grades, and data temperatures managed by Teradata Virtual Storage.

SPACE

Lists the conventional space types (**Perm**, **WAL**, **Spool**, **RSpool**, **Depot**, **Journal**, **Durable**, **Other**) as well as the unused cylinder space (**Free**).

GRADE

Lists the storage grade (**Slow**, **Medium**, **Fast**).

CYL

Lists the count of cylinders for a given space type, storage grade, or type-grade combination, in millions of cylinders or thousands (K) of cylinders.

TEMP

Lists the temperature (**Cold**, **Warm**, **Hot**).

% SPACE

Lists the percentage of cylinders for a storage grade in a given space type. When you filter the matrix by space type, the sum of this column for the space type is 100%. For example, if the number of cylinders (**CYL**) for **WAL** is 200 for slow grade, 400 for medium grade, and 200 for fast grade, the **% SPACE** column lists 25 for slow, 50 for medium, and 25 for fast, for a total of 100%. When you filter the table by **All Zones**, the sum of this column for all storage grades for a given space type is 100%. This column is displayed only if you filter the matrix by space type, or if you filter the table by **All Zones**.

% COLD

Lists the percentage of cylinders for a given space type-grade combination allocated to cold data. The sum of **% COLD**, **% WARM**, and **% HOT** is 100% for a given space type-grade combination. This column is displayed only if you filter the table by storage grade or by **All Zones**.

% WARM

Lists the percentage of cylinders for a given space type-grade combination allocated to warm data. The sum of **% COLD**, **% WARM**, and **% HOT** is 100% for a given space type-grade combination. This column is displayed only if you filter the table by storage grade or by **All Zones**.

% HOT

Lists the percentage of cylinders for a given space type-grade combination allocated to hot data. The sum of **% COLD**, **% WARM**, and **% HOT** is 100% for a given space type-grade combination. This column is displayed only if you filter the table by storage grade or by **All Zones**.

% SLOW

Lists the percentage of cylinders for a given space type-temperature combination allocated to slow grade storage. The sum of **% SLOW**, **% MEDIUM**, and **% FAST** is 100% for a given space type-temperature combination. This column is displayed only if you filter the table by data temperature or by **All Zones**.

% MEDIUM

Lists the percentage of cylinders for a given space type-temperature combination allocated to medium grade storage. The sum of **% SLOW**, **% MEDIUM**, and **% FAST** is 100% for a given space type-temperature combination. This column is displayed only if you filter the table by data temperature or by **All Zones**.

% FAST

Lists the percentage of cylinders for a given space type-temperature combination allocated to fast grade storage. The sum of % **SLOW**, % **MEDIUM**, and % **FAST** is 100% for a given space type-temperature combination. This column is displayed only if you filter the table by data temperature or by **All Zones**.

% OF ZONE

Lists the percentage of cylinders for a given space type in the temperature-grade combination. The sum of cylinder percentages in the temperature-grade-combination zone is 100%. This column is displayed only if you filter the table by a temperature-grade combination.

Viewpoint Calendar

Viewpoint Calendar Overview

The **Viewpoint Calendar** portlet allows you to schedule and communicate events that might impact one or more Teradata systems. You must have appropriate permissions to manage system events. A single instance of the **Viewpoint Calendar** portlet can be added to a Teradata Viewpoint portal page. Each instance is identical, displaying the same information.

The **Viewpoint Calendar** portlet provides the following views:

Week View

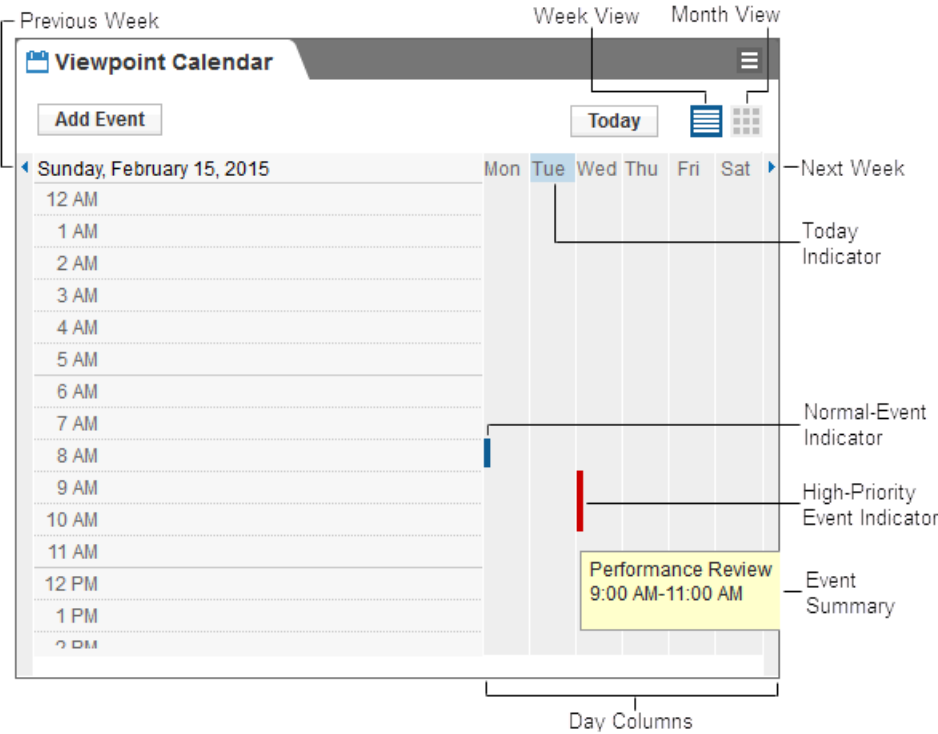
The **Week** view displays all events between Sunday and Saturday of the selected week. In this view, only one day of the week is expanded at a time and event indicators on other days of the week are displayed at the scheduled time in the appropriate day column.

Month View



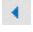

The **Month** view displays all events with start times falling between Sunday prior to the start of the selected month and Saturday after the end of the month. In this view, event indicators are color-coded squares.

Calendar Week View

The **Week** view displays all events falling between Sunday and Saturday of the selected week. In this view, only one day of the week can be expanded at a time. The selected day can include one or more event summaries. The summary shows the event start and end times, and the event title. The remaining days are collapsed into narrow columns. Click in any day column to expand and display a different day of the week. Today's date is shaded darker if it falls within the selected week.

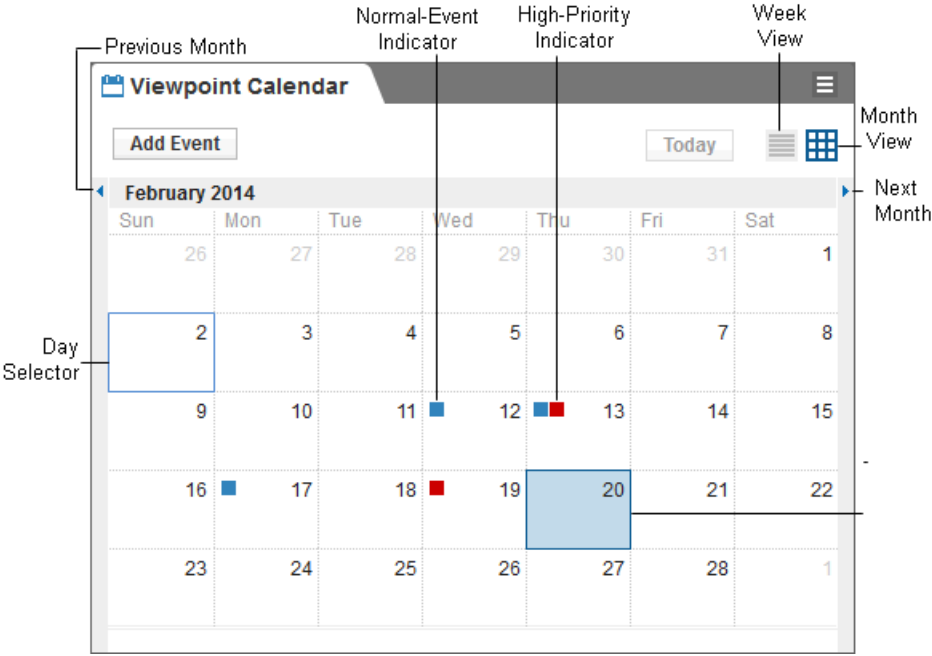


From the **Week** view, you can see event details by clicking anywhere in the event summary or event indicator. Events scheduled on other days of the week are indicated by a vertical line located at the scheduled time in the appropriate day column.





Buttons	Description
Add Event	Opens the Add Event dialog box to add an event to the calendar.
Today	Opens the monthly calendar that includes today.
	Displays the Week view. If you click the date (number), the selected day is expanded in the Week view.
	Displays the Month view.
	Displays the previous week.
	Displays the next week.

Calendar Month View

The **Month** view displays all events with start times that occur during the selected month. Events shown in the selected month that start either in the previous or following month are indicated by a lighter color. Today is shaded darker if it falls within the selected month.

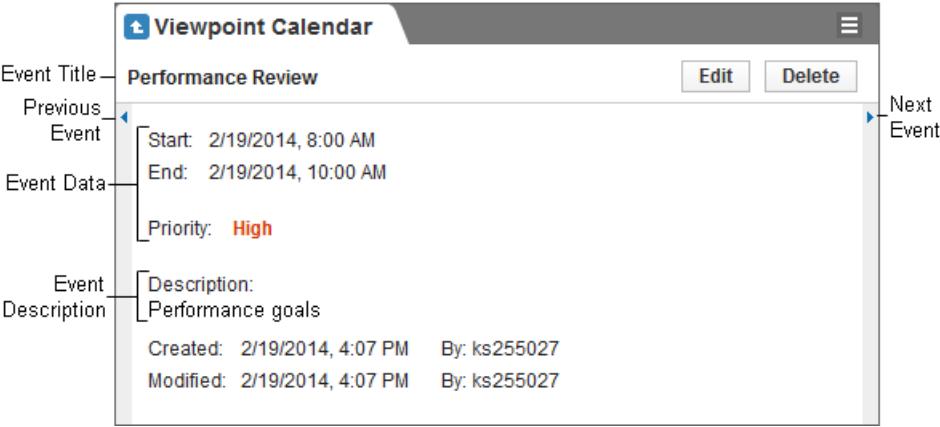


From the **Month** view, you can see event details by clicking on a date (number) or an event indicator.

Buttons	Description
Add Event	Opens the Add Event dialog box to add an event to the calendar.
Today	Opens the monthly calendar that includes today.
	Displays the Week view. If you click the date (number), the selected day is expanded in the Week view.
	Displays the Month view.
	Displays the previous month.
	Displays the next month.

Calendar Event Details View

This view displays the details of the selected event, including event priority, who created it (the author), whether it was modified, and when it is scheduled to occur. The event title and description are also included in this view.



From this view, you can manage the calendar and scheduled events.

Buttons	Description
	Return to the last displayed calendar view.
Edit	Open the Edit Event dialog box to edit the event details.
Delete	Delete this event.
	Displays the Week view containing the event.
	Displays the Month view containing the event.
	Displays the previous event.
	Displays the next event.

Calendar Events

An event cannot last longer than 24 hours. You can define a time and priority to an event and, when you save it, an event indicator (a colored box) appears in the **Week** view or **Month** view. A blue indicator means normal and a red indicator means high-priority events. Hover over an event indicator to display details about the event. All times shown are for the time zone defined in the user profile.

Adding an Event

Adding an event changes the system event calendar and affects all users. By default, new events are added to the currently selected day.

- From either the **Calendar Week** or **Month** view, click **Add Event**.
The **Add Event** dialog box appears with the current date in the **Start** and **End** date fields.

Current Date

Add Event

Start: 9/5/2013 08:00 AM
End: 9/5/2013 09:00 AM
Priority: ☐ High ☒ Normal
Title:
Description:
Author: ks255027

Add Cancel

2. Set the **Start** date:
 - a. Click inside the **Start** date field.
A calendar appears.

Previous Month Currently Selected Date Next Month

Viewpoint Calendar

Add Event

Start: 2/20/2014 08:00 AM
End:



< February 2014 >

Priority:
Title:
Description:
Author: ks255027

Add Cancel

New-Event Date

Su	Mo	Tu	We	Th	Fr	Sa
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	

- b. Click  or  to browse to the appropriate month.
 - c. Select the date from the calendar.
3. Set the **End** date.
4. Select the event **Start** and **End** times from the list.

5. Select **AM** (morning) or **PM** (afternoon) from the list.
6. Select a priority:
 - **Normal**
 - **High**
7. Enter the event **Title**.
8. [Optional] Enter the event **Description**.
9. Click **Add**.

Editing an Event

Editing an event changes the event instance affecting all users.

1. From either the **Calendar Week** or **Month** view, select the event you want to edit.
The event details view appears.
2. Click **Edit**.
The **Edit Event** dialog box appears, showing event details.
3. Make the applicable changes and click **OK**.

Deleting an Event

Deleting an event removes the event instance from the database and affects all users.

1. From either the **Calendar Week** or **Month** view, select the event you want to delete.
2. In the **Event** details view, click **Delete**.
A confirmation message appears.
3. Select **Delete**.

Viewpoint Dashboard

Viewpoint Dashboard Overview

Viewpoint Dashboard allows you to view the current state of all Teradata, Aster, and Hadoop systems. It provides a high-level view to monitor overall system activity allowing you to rapidly address issues.

Trends

Monitor system metrics and analyze system resource usage trends.

System Health

View the current state of monitored systems and key metrics against defined thresholds and then investigate metrics exceeding healthy thresholds.

Workloads

Monitor the flow of queries to identify any unusual conditions related to workload performance.

Queries

View and manage queries running on Teradata and Aster systems so you can target problem queries and take action to correct problems.

Alerts

View alerts triggered for systems.

Hadoop Services

View service status and monitor key performance indicators on a Hadoop system.

Note:

Support for TASM ARM is available for SQL Engine versions 17.10 and later.

Viewpoint Dashboard View

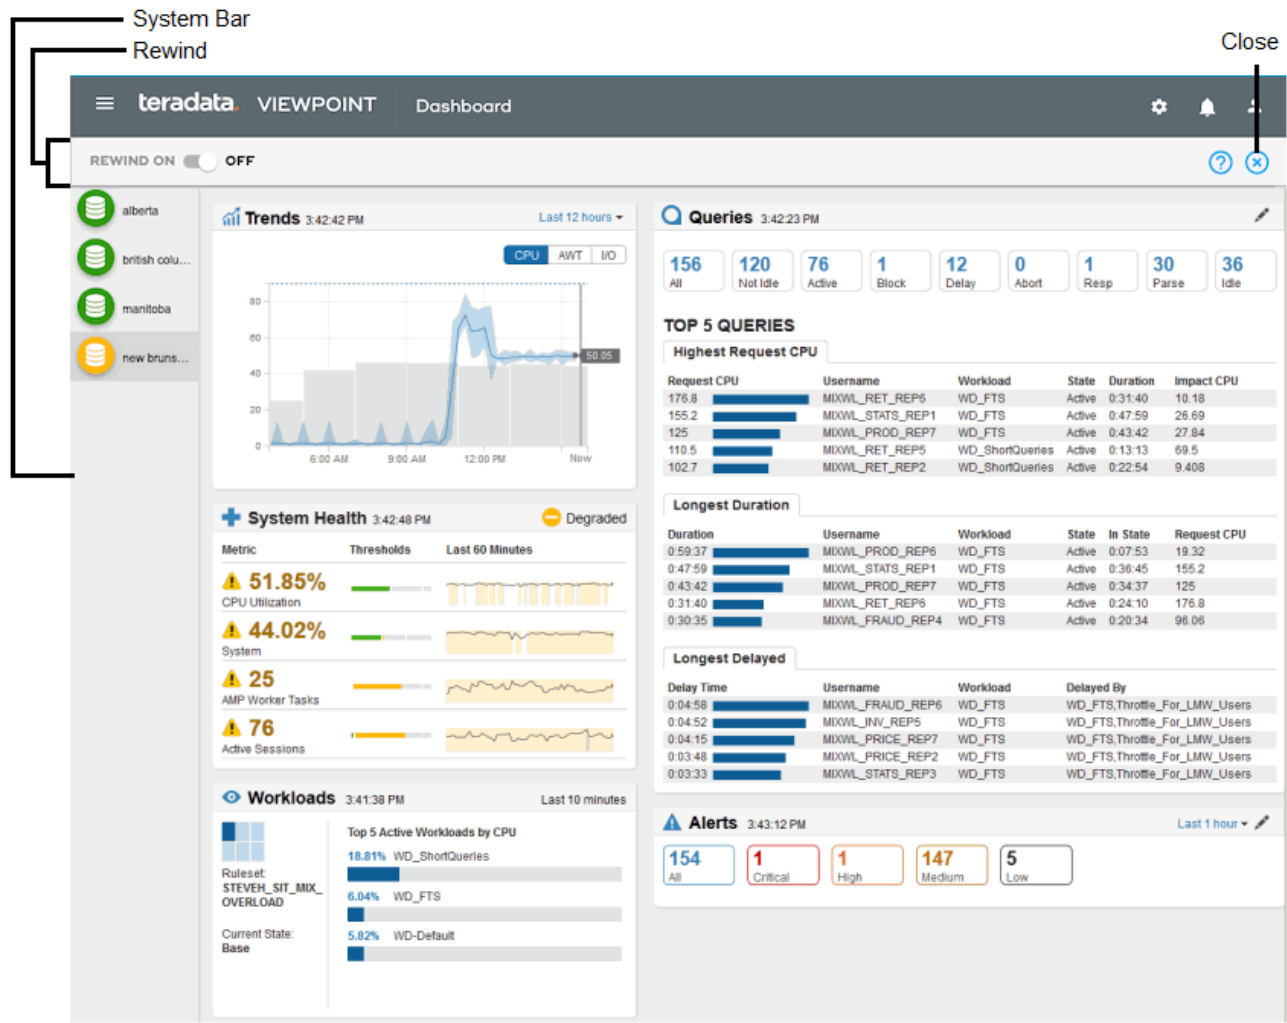
The **Viewpoint Dashboard** view is a top-level view that allows you to access the **Trends**, **System Health**, **Workloads**, **Queries**, **Alerts**, and **Hadoop Services** views.

From this view you can:

- Access all views
- Access all systems
- Open the **Viewpoint Dashboard** view
- Close the **Viewpoint Dashboard** view from two locations, enabling you to return to the portlet view

The **Viewpoint Dashboard** view supports the rewind feature. When the rewind feature is turned on in the portal frame, the open and close toggle button in the upper left corner displays a green rewind icon.

Following is an example of a Teradata system in the **Viewpoint Dashboard** view.




Opening the Viewpoint Dashboard

If the **Viewpoint Dashboard** is not available, the Teradata Viewpoint Administrator can assign permissions. For more information, see [Viewpoint Dashboard Permissions](#).

1. Do one of the following to open the **Viewpoint Dashboard**:

Option	Description
Open and select a system	a. Hover over the system icon list in the left pane in the Teradata Viewpoint portal page.

Option	Description
	<p>Status icons enable you to see the state of the system and whether it needs attention.</p> <p>b. Click a system icon in the list.</p> <p>The Viewpoint Dashboard opens to display a list of all systems.</p>
Open and return to the previously selected system	<p>a. Click  located before the system bar in the upper left corner in the Viewpoint Dashboard.</p> <p>The system you were previously viewing during the current session is automatically selected.</p>

The following views display depending on the type of system you select:

- **Trends**
- **System Health**
- **Workloads**
- **Queries**
- **Alerts**
- **Hadoop Services**

Selecting a View

1. From the **Viewpoint Dashboard** view, click in any of the following areas:

- System Health
- Workloads
- Queries
- Alerts

The navigation pane appears and shows a high-level status of the system health, workloads, queries, and alerts on the system that may need attention.


2. [Optional] Select any heading in the navigation pane to get all information for that view.
For example, click **System Health** to see the **System Health** view.
3. [Optional] Select a link underneath the heading in the navigation pane to get detailed or filtered information:

Option	Description
System Health	Displays detailed information of up to three metrics that are in a critical or degraded state, or that have exceeded a threshold. If the state is healthy, detailed information of up to two metrics displays.
Workloads	Selects and displays detailed information for the workload selected.
Queries	Displays only sessions with queries in an Active, Block, or Delay state depending on the state selected.

Option	Description
Alerts	Displays only Critical, High, or Medium alerts of the selected severity.

4. [Optional] Click the system name from the system bar to return to the **Viewpoint Dashboard** view.

Closing the Viewpoint Dashboard

1. From the upper right corner in the **Viewpoint Dashboard**, click .

Trends View

The **Trends** view allows you to analyze resource usage trends for the selected Teradata, Aster, or Hadoop system. Trends are graphed according to CPU, AWT, I/O, MEM, or DISK within a specified time frame, depending on the type of system. The information is refreshed every 60 seconds.

A time frame selection list allows you to monitor current resource usage within the last hour or view resource usage trends over the last 6 hours, 12 hours, 24 hours, or 1 week.

A graph displays with time on the horizontal axis and metric values on the vertical axis. The following list describes the features in this graph:

Plot Line

The *plot line* is a blue line that represents the average metric value. The plot line appears red where critical thresholds are exceeded. Hover over the plot line to see an information balloon that shows the actual, minimum, and maximum metric value, and a past average.

Past Averages

Past averages show the average of data samples taken for the last 2 weeks at the same time and day of the week indicated on the horizontal axis.

For example, if you set the time frame to **Last 1 hour**, the last gray segment displays the average of the data sample of the same 10 minutes (aggregation period) for last week and the week prior. The data sample for each week is calculated by averaging the aggregated data over the aggregation period.

Time Frame Selection	Aggregation Period
Last 1 hour	Every 10 minutes
Last 6 hours	Every 10 minutes
Last 12 hours	Every 15 minutes
Last 24 hours	Every 1 hour
Last 1 week	Every 6 hours

Performance Envelope

The *performance envelope* represents the upper and lower metric values.

Skyline

Use the skyline to compare the current performance with the system performance during the same time period or in the future.

Now

The **Now** point on the horizontal axis indicates the value of the last data point captured. A number to the right of the vertical line also indicates the **Now** data-point value.

Future Past Average

Use the future past average to estimate system workload.

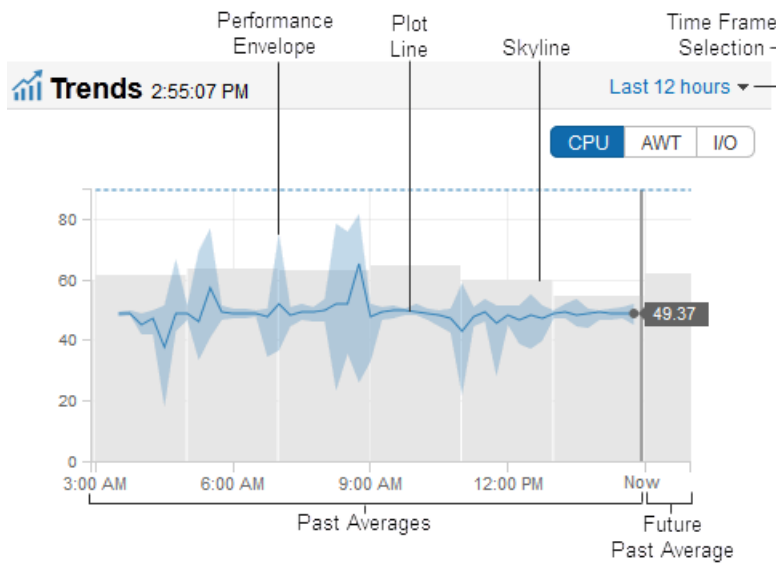
The future past average is based on the following time selections.

Time Frame Selection	Displayed Past Averages	Displayed Future Past Average
Last 1 hour	1 hour	10 minutes
Last 6 hours	6 hours	60 minutes
Last 12 hours	12 hours	2 hours
Last 24 hours	24 hours	4 hours
Last 1 week	1 week	24 hours

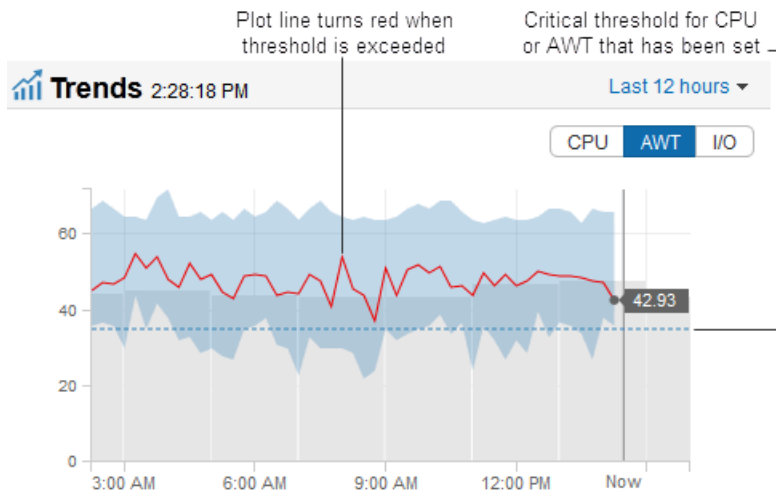
Thresholds

The thresholds settings for CPU, ENABLED CPU, and AWT cannot be set independently for the **Trends** view in **Viewpoint Dashboard**. The settings are inherited from the critical thresholds defined for calculating system health for each system. For more information, see [Configuring System Health Values](#). Thresholds cannot be set for I/O.

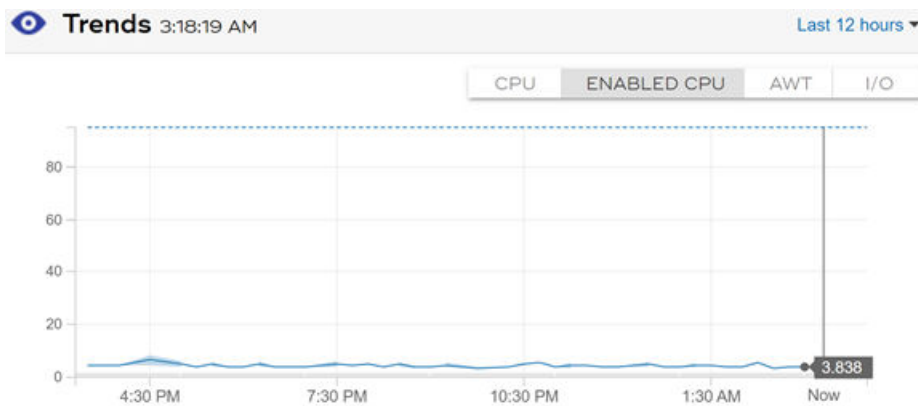
Following is an example with the **CPU** option selected.



Following is an example with the **AWT** option selected that shows the plot line turning red when critical thresholds are exceeded.



For IFX 2.1 and higher systems with enabled Elastic TCore, **ENABLED CPU** option appears in **Trends** view on **Viewpoint Dashboard**. It displays CPU Utilization based on enabled CPUs. Following is an example with the **ENABLED CPU** option selected.



Changing the Trends Display for Teradata System

You can change the **Trends** view display for a Teradata system.

1. From the **Trends** view, click one of the following:

Option	Description
CPU Note: Tool tip contents are not applicable for Viewpoint Mobile view.	Displays the percentage of average node CPU use. CPU is calculated as the sum of CPU User and CPU System. Tooltip content for IFX 2.1 and higher systems with Elastic TCore enabled: <ul style="list-style-type: none"> • TDEnabledCPUs % (of NCPUs) : Percentage of Teradata enabled CPUs of total NCPUs • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system
ENABLED CPU Note: In Viewpoint Mobile view this metric is currently not available.	Displays the CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled Tooltip content: <ul style="list-style-type: none"> • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system
AWT	Displays the average number of AMP worker tasks in use on each AMP.
I/O	Displays the number of disk I/Os. I/O is calculated as the sum of physical disk reads and writes per system during the sample period.

The graph changes to represent the chosen metric.

Changing the Trends Display for Aster System

You can change the **Trends** view display for an Aster system.

1. From the **Trends** view, click one of the following:

Option	Description
CPU	Displays the percentage of the average node CPU use.
MEM	Displays the percentage of the average node memory use.
I/O	Displays the data transfer rate, in bytes per second.

The graph changes to represent the chosen metric.

Changing the Trends Display for Hadoop System

You can change the **Trends** view display for a Hadoop system.

1. From the **Trends** view, click one of the following:

Option	Description
CPU	Displays the percentage of the average node CPU use. CPU is calculated as the sum of CPU User and CPU System.
MEM	Displays the percentage of the average memory use of the system during a sample period.
DISK	Displays the percentage of disk space being used on a system.

The graph changes to represent the chosen metric.

Setting the Time Frame

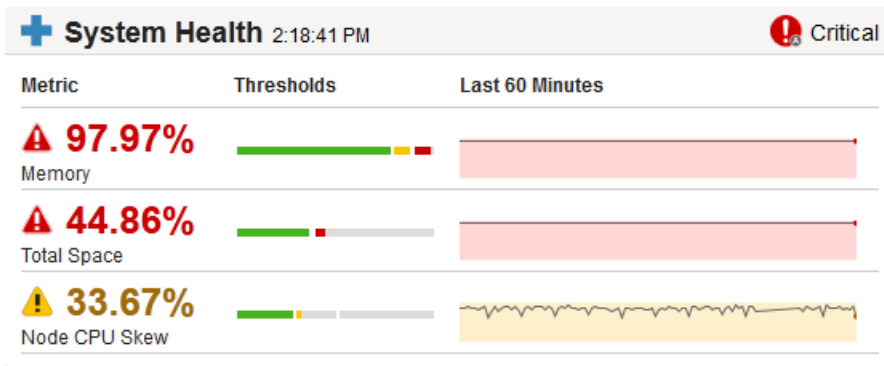
You can set the time frame that is used to plot the graph.

1. From the **Trends** view, select a time frame from the list.
The view refreshes, and the graph is redrawn.

System Health View

The **System Health** view displays icons to indicate the overall system health for the selected system. If the system is in a degraded or critical state, information displays to give you a better understanding of the issues with the system. The timestamp displays the date and time the values were collected.

Following is an example of a Teradata Aster system in a critical state.



Default Icon	Text Color	Default States	Definition
	Green	Healthy	All metrics are within <i>healthy</i> ranges.
	Yellow	Degraded	At least one metric exceeded a <i>degraded</i> threshold.
	Red	Critical	At least one metric exceeded a <i>critical</i> threshold.
	Black	Down	The selected system is <i>down</i> .
	Gray	Unknown	Status of the selected system is <i>unknown</i> .

If your system does not use the default states and tooltips, the Teradata Viewpoint Administrator customized the states and tooltips. Unlike in the **System Health** portlet, you cannot customize the icon style.

Status Icons

Show the status of the metrics in the **Metric** column:

- ! represents metrics that reached a critical state
- ! represents metrics that are degraded

From this view, click any row to investigate metrics in the **System Health** details view.

Details View

The **System Health** details view displays detailed statistics and information about each metric to evaluate the overall health of a system. If the Teradata Viewpoint Administrator defines metrics as View Only, these metrics are displayed with a gray background and are not used in determining system health. If any connection issue happens with the system because of a wrong configuration, then a relevant error message appears with the system status.

Metric




Lists the names of the monitored metrics, using color to indicate metric health. Displays metric descriptions when you hover over the name.

You can click the following metrics to access further details when at least one component is down or when one service is in a bad or concerning state:

- **Components Down** (HDP)
- **Services Bad** (CDH)
- **Services Concerning** (CDH)

Status Icons

Show the status of the metrics:

-  represents metrics that are defined as View Only
-  represents metrics that reached a critical state
-  represents metrics that are degraded

Note:

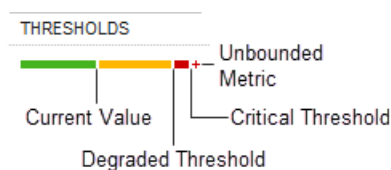
Metrics that are configured by the Teradata Viewpoint Administrator as View Only are not used in determining system health.

Value

Shows the metric value at the last system sampling.

Thresholds

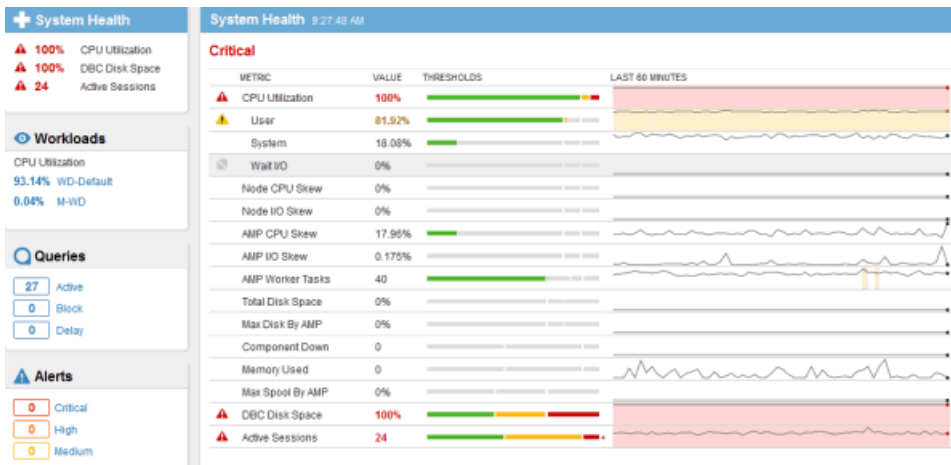
Shows the current metric value in terms of degraded and critical thresholds, indicated by gaps in the bar. Displays threshold values and default or custom tooltips when you hover over the bar.



The **+** at the end of some bar graphs indicates an unbounded metric, which has numerical values with no upper limit. The **+** is red when the unbounded metric value exceeds 1.1 times the critical threshold.

Last *nn* Minutes

Displays a sparkline with the trend in metric values over the interval selected in the **Settings** view in the **System Health** portlet of 15, 30, or 60 minutes.

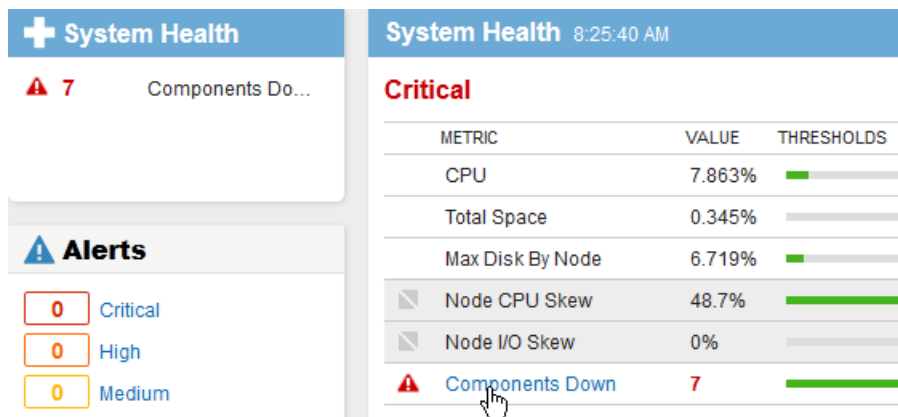


Viewing Component Down Details

When the Component Down metric has a value of one or greater, you can access details that show the extent of your exposure to components that are down.

This option is not available on an Aster system.

- From the **System Health** details view, click **Components Down**.



- Review the list of component types and the affected components, and click **Close** to return to the **System Health** details view.

Viewing Services Bad or Services Concerning Details

When the Services Bad or Services Concerning metric (CDH) has a value of one or greater, you can access details of the services in these states.

- From the **System Health** details view, click any of the following:
 - Services Bad** metric name or value

- **Services Concerning** metric name or value

A list of bad or concerning services and their host names appears.

System Health Metrics for Teradata System

Metric	Description
Active Sessions	Number of sessions with active queries
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP
AMP Worker Tasks	Average number of AMP worker tasks in use on each AMP
Component Down	Number of components, such as BYNETs or AMPs, that are down
(CPU) System	Average CPU time spent executing kernel system calls or servicing I/O and timer hardware interrupts
(CPU) User	Average CPU time spent executing code on behalf of operating system user processes
CPU Utilization	Average node CPU use. CPU is calculated as the sum of the user CPU, system CPU, and wait I/O usage percentages, depending on which of these are enabled by the Teradata Viewpoint Administrator in the Monitored Systems portlet
(CPU) Wait I/O	Average percentage of CPU time spent waiting for I/O
DBC Disk Space	Available DBC disk space in use
Enabled CPU Usage Note: In Viewpoint Mobile view this metric is currently not available.	CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled
Max Disk by AMP	Available disk space currently in use
Max Spool by AMP	Available spool space in use
Memory Used	Represents an increase or decrease in the node-level memory. A null value or negative number reports as 0.
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of disk use on the busiest node to the average node
System Heartbeat	Canary query showing the response time of the most recent system heartbeat in milliseconds
Total Disk Space	Percentage of total disk space currently in use

Metric	Description
Additional Canaries	Additional canary queries that are defined for the system and used as part of the monitoring of system health

System Health Metrics for Aster System

Metric	Description
Active Sessions	Number of users and applications currently connected to database
Component Down	Number of nodes that are not available
Component Passive	Number of nodes not processing queries but can be made ready to process queries when needed
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Memory	Average node memory use
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Queen Disk Space	Percentage of used disk space on the queen node
Replication Factor	Number of copies of the user data
Total Space	Percentage of used space to overall storage capacity

System Health Metrics for Hadoop System

CDH

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node

Metric	Description
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Services Bad	Number of services in a critical state
Services Concerning	Number of services in a degraded state
Total Space	Percentage of used space to overall storage capacity

HDP 2.1 and later

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
Components Down	Number of services not started
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls

Metric	Description
Total Space	Percentage of used space to overall storage capacity

HDP 1.3

Metric	Description
Blocks Corrupt	Number of blocks whose replicas are all corrupt
Blocks Missing	Number of blocks with no replicas anywhere in the cluster
Components Down	Number of service components not running
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.
Jobs Failed	Number of jobs that failed
Jobs Running	Number of jobs currently executing in the system
Job Tracker CPU	CPU use for the node running the jobtracker service
Map Tasks Running	Number of map tasks executing in the system
Map Tasks Waiting	Number of map tasks waiting to execute
Max Disk by Node	Amount of used disk space on the node with the most disk space in use
Name Node CPU	Node CPU use for the node running the namenode service
Name Node Heap	Percentage of heap space used in the namenode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Reduce Tasks Running	Number of reduce tasks executing in the system
Reduce Tasks Waiting	Number of reduce tasks waiting to execute
RPC Latency JT	Average wait time in queue for jobtracker service calls
RPC Latency NN	Average wait time in queue for namenode service calls
Total Space	Percentage of used space to overall storage capacity

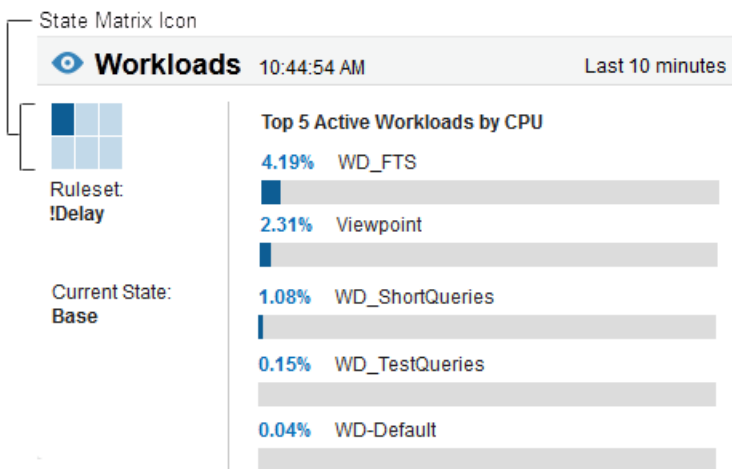
Workloads View

The **Workloads** view allows you to monitor workload management activity in Analytics Database. The timestamp displays the date and time the values were collected. The available options are different if Analytics Database is running on SUSE Linux Enterprise Server version 10 (SLES 10) or lower, or SLES 11. For conceptual information, see the **Workload Monitor** portlet documentation.

The **Workloads** view displays the following for the last 10 minutes:

- Name of the currently active ruleset and the current state. Information about the current state, planned environment, and health condition is displayed by mousing over the graphic before the ruleset name.
- Toggle buttons allow you to change the **Active Workloads** display.
- Percentage of resources being consumed by the **Top 5 Active Workloads** or **Top 5 Active Workloads by CPU**, depending on your system type.

Following is an example of the **Workloads** view with no I/O metrics.



Changing the Active Workloads Display

You can change the **Active Workloads** display to view the active workload resource consumption for the last 10 minute sampling period for the top 5 active workloads for a ruleset.

The CPU and I/O toggle buttons only display for SLES 10 and 11 systems.

1. From the **Workloads** view, click one of the following to show the top five active workloads within the sampling period:

Option	Description
CPU	SLES 10 or 11: Displays the percentage of system CPU consumed by each workload.
I/O	SLES 11: Displays the percentage of I/O consumed by each workload.

Viewing Workload Details

1. Click any of the workloads listed under **Top 5 Active Workloads**.
The workloads details view appears.

Details View

The **Workloads** details view helps you analyze workload performance.

The upper left corner of the view, shows the state matrix, name of the active ruleset, and the current state.

Current

- **Active:** Shows the number of requests that are currently active.
- **Delayed:** Shows the number of requests that are currently delayed by system and workload throttles.

Cumulative

- **Arrivals:** Shows the cumulative number of requests arriving.
- **Requests Delayed:** Shows the cumulative number of requests that are delayed by system and workload throttles.
- **Rejected Requests:** Shows the cumulative number of requests that were rejected and not allowed to continue processing.
- **Exceptions:** Shows the cumulative number of requests that triggered an exception.
- **Aborted:** Shows the cumulative number of requests that have been aborted because of exceptions.

Show Details

Shows the trends sparkline graphs.

Hide Details

Hides the trends sparkline graphs.

Sampling Period

Indicates the sampling period of the displayed information is from the last 10 minutes.

Workload List

Shows a list of active workloads on the system. You can click any row to display the workload metrics.

☐ Table Actions

Configure Columns allows you to choose the columns to display and set thresholds.

Note:

Changing the **Configure Columns** settings on one system changes the settings on all systems.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

Trends Sparkline Graphs

Displays trends sparkline graphs that illustrate workload performance metrics for the selected workload.

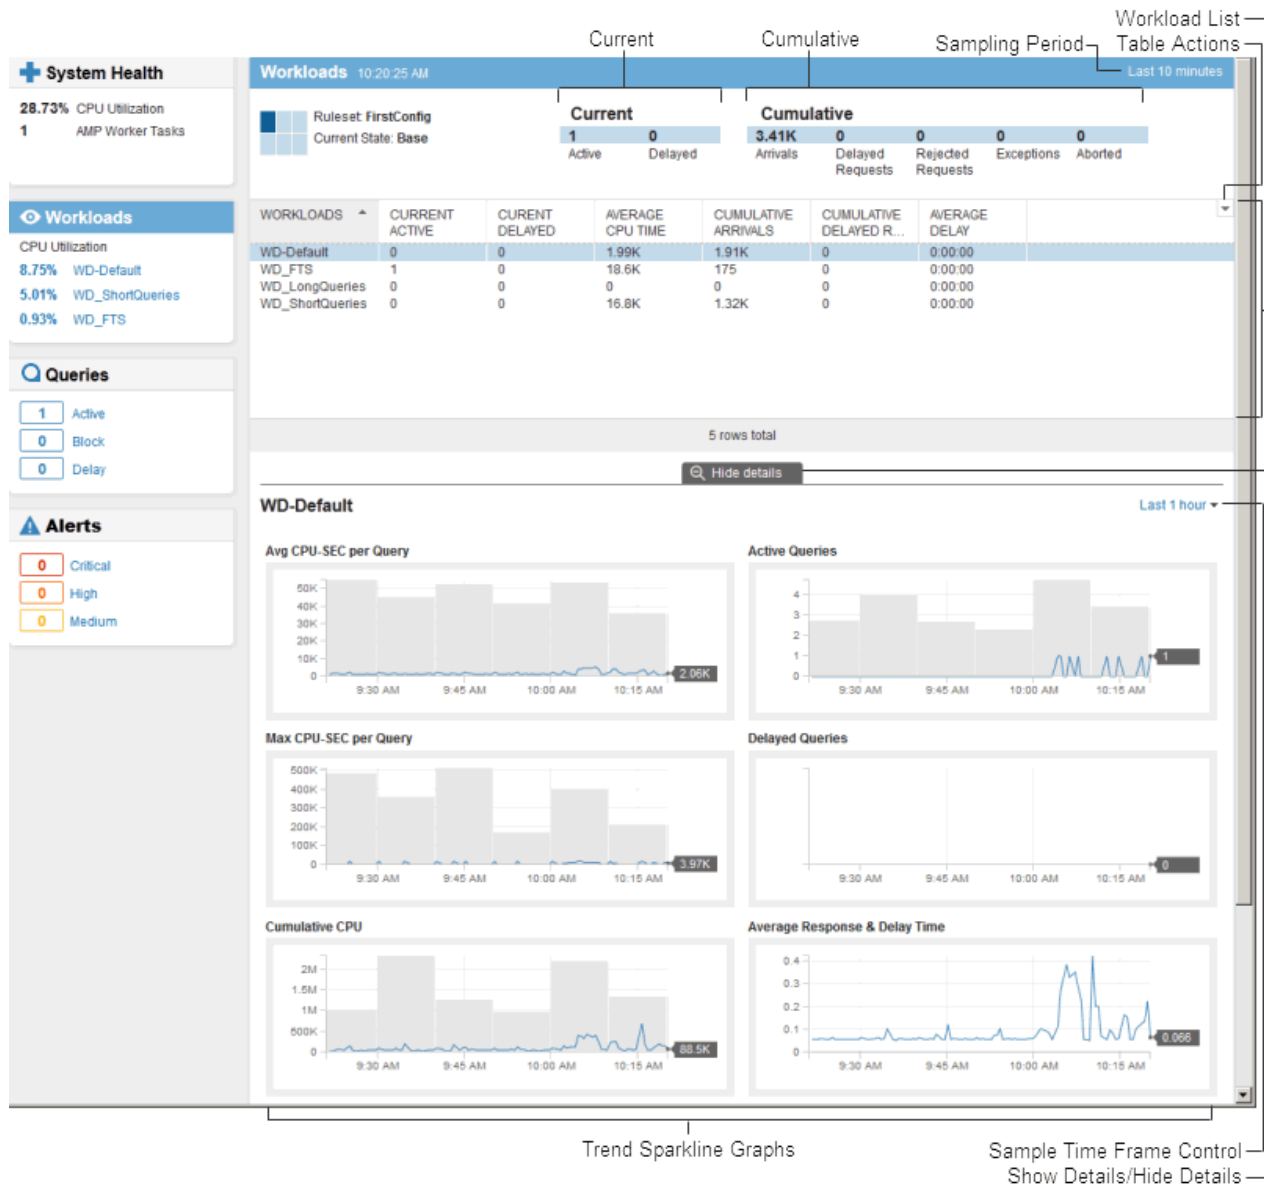
The solid line represents current data for the metric and the grayed skyline represents the past averages for that same metric. When you hover over the sparkline, the end data points and past averages appear in the information balloon. Past averages shows the averages for the time indicated on the horizontal axis.

You can set the default trend interval that displays on the horizontal axis of the graphs. Use the **Trend Interval** tab in the **Settings** view in the **Workload Monitor** portlet.

For metric descriptions of the trends sparkline graphs, see [Workloads Metrics for Teradata System](#).

Sampling Time Frame Control

Changes the sampling time frame for the metrics displayed in the graphs.



Workloads Metrics for Teradata System

Workload metric descriptions for the trends sparkline graphs shown in the **Workloads** details view are listed as following.

Metric	Description	Operating System
Avg CPU-SEC per Query	Average CPU time, in seconds, of all workload queries	SLES 10 and 11
Max CPU-SEC per Query	Maximum CPU time, in seconds, of all queries that completed	SLES 10 and 11

Metric	Description	Operating System
Cumulative CPU	Cumulative CPU seconds per AMP	SLES 10
Average Physical KB I/O per Query	Average physical I/O usage, in kilobytes, of all workload queries	SLES 11
Throughput	Number of queries arriving and completing every minute	SLES 10 and 11
Active Queries	Number of active queries	SLES 10 and 11
Delayed Queries	Number of queries in the delay queue	SLES 10 and 11
Average Response Time & Delay Time	Average response time and delay time for all completed queries	SLES 10 and 11
Maximum Response Time	Longest response time of all requests that completed	SLES 10 and 11

Queries View

The **Queries** view displays a state filter bar and a **Top 5 Queries** area. The timestamp displays the date and time the values were collected.

State Filter Bar

Displays a count of the sessions in each state. Click any state in the bar to show only the sessions running a query in the selected state.

Top 5 Queries

Displays the top five queries on a Teradata system in the following categories:

- **Highest Request CPU** displays the top five queries with the highest request CPU seconds consumed in running the query.
- **Highest CPU Skew Overhead** displays the top five queries that have the highest request CPU skew (applies to Teradata Database 15.10 and later and SLES 10 and 11).

Highest CPU Skew Overhead is calculated as the **Impact CPU** minus the **Request CPU**. It is the difference between how much CPU the query effectively consumed, based on how parallel it was executed, versus how much CPU the request actually consumed. It is a measure of skew for a query, measured over the lifetime of the query. For example, a value of 0 means the query was not skewed. A higher number means the query was both skewed and consumed higher amounts of system resources.

- **Longest Duration** displays the top five queries that have been running the longest.
- **Longest Delayed** displays the top five queries currently in the delay queue that have been waiting the longest time to run.

For Teradata systems, you can set thresholds for each of the categories to control which queries are highlighted in the **Top 5 Queries** view. For information on setting thresholds, see [Setting Thresholds for Top 5 Queries](#).

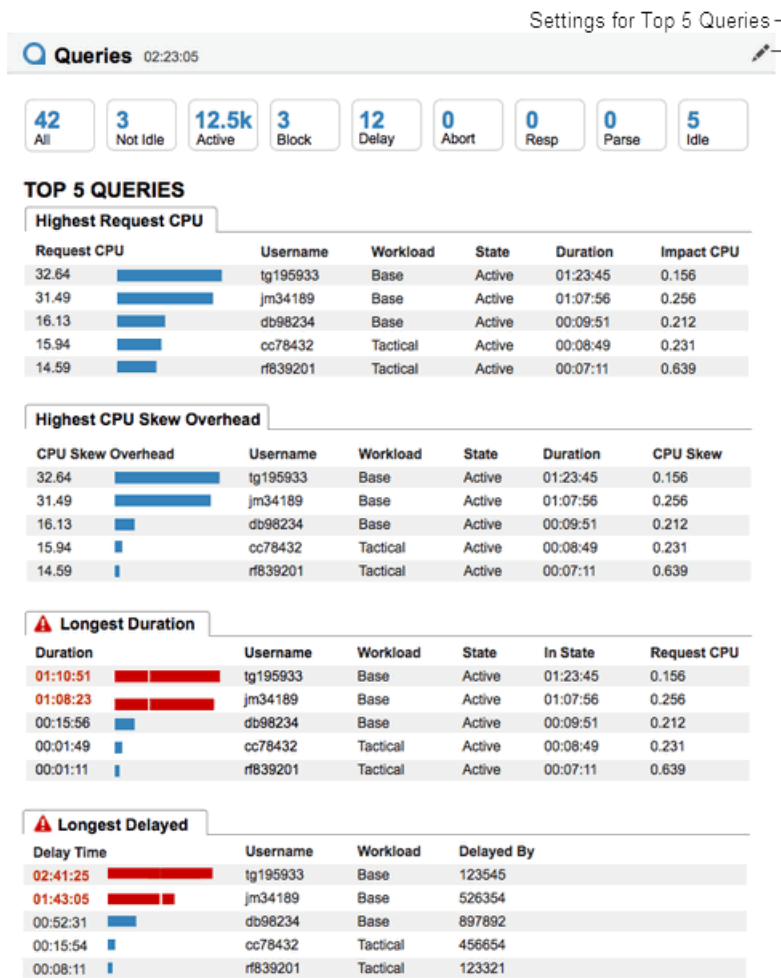
The bar graphs allow you to visually compare one query to another query within each category in the **Top 5 Queries** view. A vertical white line represents the threshold value. Queries exceeding the threshold display a red bar graph. Queries not exceeding the threshold, or if no thresholds are set, display a blue bar graph.

The following describes the meaning of the icons or if there is no icon:

- A green icon displays next to the heading to indicate no thresholds have been exceeded and all queries are in a healthy state.
- A red icon displays next to the heading and the metric and bar graph are highlighted in red to visually emphasize when at least one query exceeds a threshold. If you click the query row, details display for the query that exceeded a threshold.
- No icon displays if no thresholds are set.

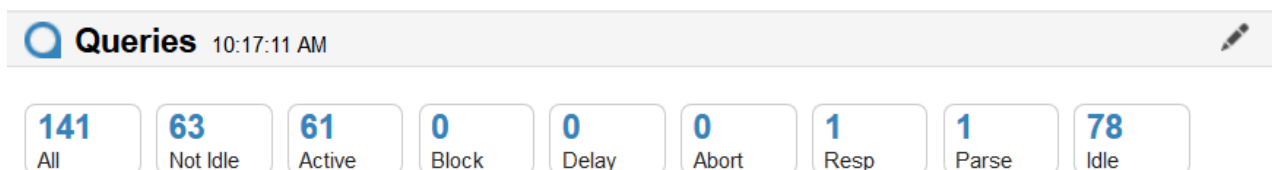
Note:

Some information displayed for the selected system is obtained from the **Query Monitor** portlet.



Queries State Filter Bar

The state filter bar allows you to display specific states by drilling down to the details view.



The state filter buttons provide a count of sessions for each state category. Click a state in the queries state filter bar to drill down to the details view. The details view displays only sessions with a query in the selected state.

Teradata System States

All	Sessions in the system
Not Idle	Sessions in any state except idle
Active	Sessions processing queries
Block	Sessions processing a query waiting for a locked resource, such as a database table or view
Defer	Sessions waiting under deferred queue to be processed
Delay	Sessions waiting to be processed
Abort	Sessions in the process of aborting (rolling back changes made by the query)
Resp	Sessions running a query that has completed and is sending (responding) spooled data back to the user
Idle	Sessions not currently running a query
Parse	Sessions running a query that is being parsed and has not begun to execute
Other	Sessions whose status is unknown and do not fall into another category
QTDelayed	Sessions delayed due to a queue table restriction
SesDelayed	Utility sessions that are in the workload delay queue

Response-Held

In SLES 11, sessions running a query whose response is being held until a specified number of seconds elapses

Aster System States**All**

Sessions in the system

Active

Sessions processing queries

Idle

Sessions not running

Canceled

Processes that were stopped

Pending

Sessions waiting to be processed

Details Views

There are two **Query** details views:

Query Details by State

Displays summary information about sessions in columns.

Filters

Shows only rows that match your filter criteria.

☐ **Table Actions**

For Teradata systems, table actions include:

- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display.

Note:

Changing the **Configure Columns** settings on one system changes the settings on all systems.

- **Abort**, **Change Priority**, **Change Workload**, and **Release Query** allow you to manage the selected queries and sessions.

- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For Aster systems, table actions include:

- **Clear Filters** removes any content in the filter boxes.
- **Configure Columns** allows you to choose the columns to display.

Note:

Changing the **Configure Columns** settings on one system changes the settings on all systems.

- **Abort Process** allows you to abort the selected non-idle process.
- **Export** creates a .csv file containing all available data. If filters are used, only filtered data is exported.

For more information on **Clear Filters**, **Configure Columns**, and **Export** table actions, see [Summary Table Controls](#).

Query Details by Session

Displays details of the selected query in sections that are specific to each tab.

Tabs

Teradata systems have **Overview**, **SQL**, **Explain**, **Skew**, **Query Band**, **Blocking**, **Blocked By**, **Defer**, and **Delay** tabs.

Teradata Aster systems have **Overview**, **SQL**, and **Explain** tabs.

Teradata Machine Learning Engine systems have **Overview**, **SQL**, and **Explain** tabs.

Show Details

Shows the **Query Details by Session** view.

Hide Details

Hides the **Query Details by Session** view.

|| Pause Refresh

Pauses the refreshing of data, allowing you to view the data in all tabs in the **Query Details by Session** view. The **Query Details by State** view is not affected by using this control.

► Resume Refresh

Resumes bringing the most recent data from the Viewpoint server to all tabs in the **Query Details by Session** view.

System Health

- 1479 System Heartbeat
- 45 Active Sessions

Workloads

- CPU Utilization
- 8.88% WD_FTS
- 6.86% WD_ShortQueries
- 4.22% WD-Default

Queries

122 rows total

Session: 4286193

Overview | SQL | Explain | Skew | Query Band

QUERY INFO

- State: Idle
- Time in state: 0:01:35
- Total duration: 0:00:00
- Spool space: -
- Temp space: 0
- Request CPU: -
- Request I/O: -
- PJL: -
- Unnecessary I/O: 0

WORKLOAD INFO

- Name: -
- Classification mode: -

SESSION INFO

- User: VIEWPOINT
- Account: DBC
- Partition: MONITOR
- Requests: 1,025,214
- Source: (TCP/IP) a44e 153.64.107.75 CALVIN2:CALVIN2COP1/153.64.207.8 8:1025 CID=5F35B4B ROOT JDBC14.10.00.17:1.6.0_30 01 LSS

SNAPSHOT INFO

- CPU use: 0%
- Impact CPU: 0
- Snapshot CPU skew: -
- Snapshot I/O skew: -

Table Actions: Query Details by Session, Query Details by State

Pause/Resume Refresh
Show Details/Hide Details

Queries Metrics for Teradata System

Metrics for a Teradata system that display in the **Viewpoint Dashboard Queries** view are listed following.

Metric	Description	Type
Account	Account from which a query was submitted	String
Blocked Time	How long the query has been blocked	Number
CPU Skew	Percentage of how inefficiently a query or job was executed in terms of CPU parallelism. A value of 0% indicates the query was executed perfectly in parallel, meaning each AMP consumed identical amounts of CPU. A value of 50% indicates a single AMP consumed twice as much CPU as the average. As CPU Skew	Percent

Metric	Description	Type
	approaches 100%, it indicates the query or job was not executed in parallel.	
CPU Use	Percent of available CPU seconds on the system used during the last sampling period	Percent
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules	Number
Delayed By	Names of the throttles causing the query to be delayed	String
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number
Duration	How long the query has been running	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
In State	How long the query has been in the current state	Number
Partition	Partition in which the query is running	String
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String
Query Band	Entire query band string. Query bands are a set of name-value pairs defined by the user to tag sessions or transactions with an ID through a SQL interface.	String
Req CPU	CPU seconds needed to run the query	Number
Req CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Req Count	Number of queries the session has executed	Number
Req I/O	Disk I/Os performed to run the query	Number
Req I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Session ID	Unique session identifier	Number
Snapshot CPU Skew	CPU skew during the last sample	Percent
Snapshot I/O Skew	I/O skew during the last sample	Percent

Metric	Description	Type
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Start	Time that the query started running on Analytics Database	Number
State	Text describing the current state of the query	String
State Icon	Icon representing the current state of the query	Icon
Temp Space	Temp space used by the query, using a 1024 byte kilobyte as the base	Number
Unnecessary I/O	Ratio of I/O to CPU for the current query	Number
Username	Name of the user who submitted the query	String
Workload	Workload in which the query is running	String

Queries Metrics for Aster System

Metrics for an Aster system that display in the **Viewpoint Dashboard Queries** view are listed as following.

Metric	Description	Type
Database	Name of the database on which the process is running	String
In State	How long the process has been in the current state	Number
Process ID	Unique process identifier	Number
Session ID	Unique session identifier	Number
Start	Time the process began running	Number
State	Process state, such as active, canceled, pending	String
State Icon	Symbol of the process state	Icon
Type	Values for process types are as follows: <ul style="list-style-type: none"> • SQLMR (SQL-MapReduce) • SQL • Teradata Import • Teradata Export 	String
Username	Name of the user who submitted the process	String
Workload	Workload in which the process is running	String

Viewing Query Details by State

1. From the **Viewpoint Dashboard**, do any of the following:

- Click any state in the **Queries** state filter bar.
- Click any row under **Top 5 Queries**.

The query results that display are filtered based on what you selected.

Viewing Query Details by Session

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.

The query details by session displays after the query details by state. The information is filtered based on the row you selected.

Overview Tab for Teradata System

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system. Metrics that exceed defined thresholds are highlighted. Under SESSION INFO, you can click a user or account to see all sessions for that user or account string.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.

Query Information	Description
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	(Teradata Database 15.0 and earlier) CPU impact on the system based on the highest utilized AMP
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare

Workload Information	Description
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Session Information

Session Information	Description
User	Name of the user that submitted the query
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system

Overview Tab for Aster System

The **Overview** tab displays detailed information about key metrics for the selected session. The metric values provide a view of the query status on the system.

Query Information

Query Information	Description
Process ID	Process identifier
State	Process state, such as active, canceled, pending
In State	How long the process has been in the current state
Start	Time the process started running

Workload Information

Workload Information	Description
Workload Policy	Name of the workload policy used to manage the process. Defines a set of related queries, and allows them to be prioritized in a similar manner.
Service Class	Name of the service class used by the workload policy. Specifies the share of system resources allocated to a workload as a function of the class's priority and weight values.
Priority	Indicates the importance of the process. <ul style="list-style-type: none"> • High • Medium • Low
Weight	Number ranging from 1 to 100 that indicates the precedence of the process within the priority level, increasing the resources that can be allocated for the process. A higher value indicates a greater level of precedence.

Session Information

Session Information	Description
User	Name of user that submitted the query
Database	Name of the database on which the process ran
IP Address	IP address of the user that submitted the process

SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a


query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Exporting SQL



1. From the **Viewpoint Dashboard**, do any of the following:


Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.





The query details view appears.

2. From the **SQL** tab, click .
3. Save the file to a location.
The SQL statement is exported as a .txt file.

Explain Tab for Teradata System

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status. If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

If you have the correct permissions, use  to retrieve the latest Explain steps from the database. Click **Refresh** to update. The refresh screen data is not maintained once the refresh view closes. Repeatedly retrieving current Explain data can impact system performance.

Step Information	Description
Step Number	<ul style="list-style-type: none"> • Completed steps are at the top of the list and indicated by a black number box. • Active steps are indicated by a pulsating number box (flashes blue). • Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	<ul style="list-style-type: none">  - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step

Step Information	Description
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Explain Tab for Aster System

The **Explain** tab displays the Explain text that shows the progress of a query. Each Explain step has a colored box that indicates the status of the step. This information is read-only. You can also export Explain data.

To view all steps and phase details, click **Show all**. Alternatively, click **Show summary** to view a shortened version of the Explain data.


Step Information	Description
Step Status Icon	Indicates the status of each Explain step: <ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black box. Active steps are indicated by a pulsating box (flashes gray and white). Steps to run are at the bottom of the list and indicated by a white box.
Time	Duration of the step
Type	Category of step, such as Command, CommitPreparedTransaction, DataTransfer, Import, PrepareTransaction, and Query
Phase/Detail	Summarized or full explanation of query that executes in this step. Can be used to diagnose the cause of the failure or slowness.

Exporting Explain Data

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.

The query details view appears.

2. From the **Explain** tab, click .
3. Save the file to a location.

The Explain data is exported as a .csv file.

Skew Tab for Teradata System

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Query Band Tab for Teradata System

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Blocked by Tab for Teradata System

When a query is blocked, use the **Blocked By** tab to identify the other queries blocking the selected query. By reviewing the number and type of queries blocking the selected query, you can identify

blocking issues on the system, determine the impact of this blocking on the selected query and, based on this information, decide the best course of action to resolve the issue.

The information in this tab is read-only. The tab is available only when the selected query is blocked. You can drill down into the session information for the blocking session. Use the **Session ID** link when connected to Teradata Database 14.10 and earlier, and click anywhere in a row in the table when connected to Teradata Database 15.0 and later.

The following information is available in the noted versions of Analytics Database.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row 		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
	hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock		
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Blocking Tab for Teradata System

When a query is holding locks that are blocking other queries from completing, use the **Blocking** tab to gauge the impact of this blocking. By analyzing the number and type of queries blocked by this query, you can decide the best course of action to resolve the issue.

This information in this tab is read-only. The tab is available only when the selected query is blocking other queries and when connected to Teradata Database 15.0 and later. From the tab, you can drill down to view the details of the blocked session.

Blocking Information	Description
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>

Blocking Information	Description
Blocking Locks	<p>A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName</p> <p>The information displays in the following order:</p> <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock
Session ID	Unique session identifier of the query that is blocking
Blocked Time	How long the query has been blocked, displayed as <i>d + hh :mm :ss</i>
Username	Name of the user that is running the query that holds the lock
Workload	Name of the workload where the query is actively running
Query Band	String of key/value pairs of the query band for the session or transaction

Defer Tab for Teradata System

The Defer tab displays details about the rules that are deferring a query. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Defer Time	Duration of time request has been deferred by Arrival Rate Meter rules
User Name	User who submitted the query

Rule Information	Description
Rule Name	Name of rule causing request to defer
Rule Type	TASM type of rule causing request to defer
Overridable	Indicates if the Teradata DBA can abort or release the query

Delay Tab for Teradata System

The **Delay** tab displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Managing Queries and Sessions for Teradata System

You can manage queries and sessions to improve workload performance for Teradata systems.

Abort

Abort the selected query or session

Change Priority

Change the priority of the selected query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change priorities or workloads, or release queries. If you log off, you must log in again.

Note:

Change Workload is available only if Teradata Active System Management (TASM) is enabled. If TASM is disabled, **Change Priority** is available. If you do not see **Change Workload** or **Change Priority** in the list, the system you are monitoring does not support these features or you do not have permission to use them. If the query you are monitoring is delayed, only **Release Query** is available.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.

The session number of the query you want to abort appears in the details view.

2. Click ▼ to the right of the session number and select **Abort**.
3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
6. Click **Next**.
7. Click **Next** to confirm your selection.
8. Click **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> Click any row under Top 5 Queries.

The query details view appears.

- Click ▼ to the right of the session number and select **Change Priority**.
- Log in to Analytics Database, if prompted.
- Click **Next**.
- Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.
 - Type an account string.
- [Optional] Select the check box to use the account string as the default for the selected session.
The priority will be changed for the selected query and all subsequent queries in the current session.
- Click **Next**.
- Click **Next** to confirm your selection.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

- From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> Click any row under Top 5 Queries.

The query details view appears.

- Click ▼ to the right of the session number and select **Change Workload**.
- Log in to Analytics Database, if prompted.
- Select a different workload from the list and click **Next**.
- Click **Next** to confirm your selection.
- Click **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.

The query details view appears.

2. Click ▼ to the right of the session number and select **Release Query**.
3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Click **Next** to confirm your selection.

Managing Sessions for Aster System

You can manage sessions to improve workload performance for Aster systems.

Abort

Abort the selected process.

Aborting a Process

For Teradata Aster systems, you can abort a non-idle process that is blocking other processes or consuming too many resources.

1. From the **Viewpoint Dashboard**, do any of the following:

Option	Description
Queries filtered by state	<ol style="list-style-type: none"> a. Click any state in the Queries state filter bar. The query results that display are filtered based on what you selected. b. Click the row in the table for the session you want to see.
Top 5 queries	<ol style="list-style-type: none"> a. Click any row under Top 5 Queries.

The query details view appears.


2. Do one of the following:

Option	Description
Abort a single process	<ol style="list-style-type: none"> Click ▼ to the right of the session number and select Abort Process. Click Next.
Abort multiple processes	<ol style="list-style-type: none"> From the Table Actions list, select Abort Process. Check boxes appear next to the session ID. Do one of the following: <ul style="list-style-type: none"> Select the first check box located in the column heading to select all processes. Select the check boxes for specific processes.

- Click **Next**.
- Log in to Teradata Aster, if prompted.
- Click **Next** to confirm you want to abort the processes.
A confirmation appears indicating the abort process was successfully submitted.
- Click **OK**.

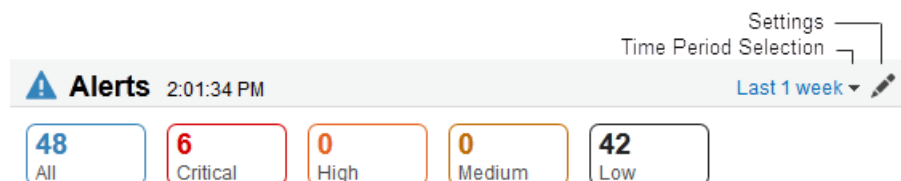
Setting Thresholds for Top 5 Queries

For Teradata systems, you can specify thresholds to control which queries are highlighted in the **Top 5 Queries** view. Configure these settings for each system.

- From the **Top 5 Queries** view, click .
- Enter one or more threshold values.
Metrics exceeding at least one threshold value will be highlighted in the **Top 5 Queries** view.
- Click **OK**.

Alerts View

The **Alerts** view allows users to view triggered alerts and are specific to your system. The information in the view is refreshed every 30 seconds.



An alert is an event that the Teradata System Administrator defines as being significant. The Teradata System Administrator assigns alert severity levels to rank alerts, and can also include an explanatory message. Following are the severity levels:

- Critical

- High
- Medium
- Low

Alert Example

The Teradata System Administrator can define that a database exceeding a certain amount of storage usage triggers an alert. After the usage level is exceeded, an alert appears in the **Alerts** view to inform the user that the threshold was exceeded.

You can change a setting so hidden alerts display. For more information, see [Displaying Hidden Alerts](#).

Details Views

The **Alerts** view contains the following details views:

Alert Details by Severity

View the alert details filtered by severity, including the timestamp, alert type, alert name, and source.

Alert Details for an Alert

View the alert details for the selected alert, including the type, source, severity, and any relevant messages.

You can filter the alerts by severity, time period, type, or name. You can also combine the filters to narrow the results further.

Severity Filter Bar

Displays a count of the alerts for each severity. Click any severity in the bar to change the displayed data in the summary table to show only the alerts of that severity.

☒ **Last 1 hour** Sets the time period for the alerts in the **Severity Filter Bar**.

Filters

Shows only rows that match your filter criteria.

Summary Table

Displays summary information about alerts in columns. The current view is configured in the **Columns** dialog box. The view is refreshed every 30 seconds. Click a row in the table to see details.

☒ Table Actions

Clear Filters removes any content in the filter boxes.

Configure Columns allows you to choose the columns to display and set thresholds.

Note:

Changing the **Configure Columns** settings on one system changes the settings on all systems.

Export creates a .csv file containing all available data. If filters are used, only filtered data is exported.

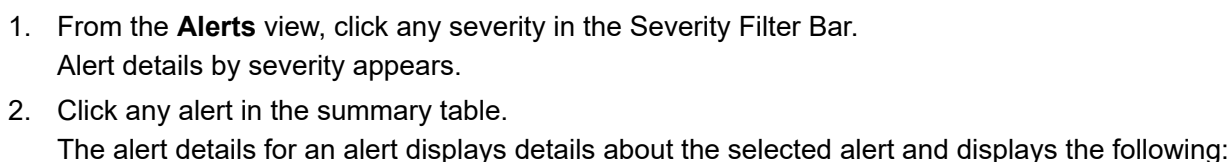
Hide Alerts enables you to select and hide alerts so that they are no longer displayed in the summary list. Hidden alerts can be viewed and then displayed again in the summary list.

Show Details

Shows the **Alert Details for an Alert** view.

Hide Details

Hides the **Alert Details for an Alert** view.



- **General:** Alert type, severity, source, timestamp, resulting action, and alert criteria.
- **Properties :** Alert properties and their associated values. For metric descriptions, see [Alerts](#).
- **Messages:** Any messages issued.

Alert Metrics

Metric	Description
Alert Name	Name of the alert
Alert Type	Type of alert
Message	Message providing more information about the alert
Severity	Severity level assigned to the alert
Severity Icon	Colored icon representing the severity level
Source	Source (Viewpoint, Teradata Alerts, Server Management, Performance Data Collection, or Teradata Ecosystem Manager)
Timestamp	Date and time when the alert was issued

Hiding Alerts Basics

By default, the **Alerts** view displays all alerts that have been generated. You can hide some alerts so they no longer appear in the summary list. For example, you might want to hide all alerts generated for database space conditions or alerts associated with conditions that were subsequently resolved.

You can also display hidden alerts, and then unhide them. When you display hidden alerts, the alert rows appear with a line through them, and their associated severity indicators appear as unfilled squares. If you click the alert to view its details, the severity indicator next to the alert name also indicates the hidden state of the alert.

If you hide alerts in one of the following, the alerts are hidden in the other two:




- **Alert Viewer** portlet
- **Viewpoint Dashboard**
- **Viewpoint Mobile**


The **Settings** view allows you to display previously hidden alerts.

Hiding Alerts

You can select and hide alerts from the **Viewpoint Dashboard Alerts** view.

- To hide only one alert, do the following:


Option	Description
From the Alert Details by Severity view	<ol style="list-style-type: none"> 1. Click  in the row of the alert you want to hide. 2. Select Hide Alert. <p>The alert no longer appears in the view.</p>
From the Alert Details by Severity view	<ol style="list-style-type: none"> 1. Click  in the table header to open the Table Actions list. 2. Select Hide Alerts. 3. Select the check box in the row of the alert you want to hide. 4. Click Save. <p>The alert no longer appears in the view.</p>
From the Alert Details by Alert Name view	<ol style="list-style-type: none"> 1. Click  next to the alert name. 2. Select Hide Alert. <p>The colored severity indicator square appears unfilled to indicate that the alert is hidden. The alert no longer appears in the Alert Details by Severity view.</p>

- To hide one or more alerts, do the following:
 1. From the **Alert Details by Severity** view, click  in the table header to open the **Table Actions** list.
 2. Select **Hide Alerts**.
 3. Select the check boxes in the rows of the alerts you want to hide.
 4. Click **Save**.

The alerts no longer appear in the **Alert Details by Severity** view.

Displaying Hidden Alerts

You can designate that hidden alerts be displayed. Hidden alerts are displayed with a line through them, and their associated severity indicators appear as unfilled squares.

1. From the **Alerts** view, click .
2. Click the **Display hidden alerts** check box.
3. Click **OK**.

Unhiding Alerts

Prerequisite:

You must display alerts that were previously hidden before you can unhide them.

You can select and unhide alerts from the **Viewpoint Dashboard Alerts** view. When you unhide an alert, it is displayed again in the **Alert Details by Severity** view and the severity indicator associated with it is displayed as a filled square.

- To unhide only one alert, do the following:

Option	Description
From the Alert Details by Severity view	<ol style="list-style-type: none"> 1. Click <input type="checkbox"/> in the row of the alert you want to unhide. 2. Select Unhide Alert.
From the Alert Details by Severity view	<ol style="list-style-type: none"> 1. Click <input type="checkbox"/> in the table header to open the Table Actions list. 2. Select Hide Alerts. 3. Clear the check box in the row of the alert you want to unhide. 4. Click Save.
From the Alert Details by Alert Name view	<ol style="list-style-type: none"> 1. Click <input type="checkbox"/> next to the alert name. 2. Select Unhide Alert.

The alert displays in its normal active state again in the summary view and the severity indicator associated with it is displayed as a filled square.

- To unhide one or more alerts, do the following:
 1. From the **Alert Details by Severity** view, click ☐ in the table header to open the **Table Actions** list.
 2. Select **Hide Alerts**.
 3. Clear the check boxes in the rows of the alerts you want to unhide.
 4. Click **Save**.

The alert displays in its normal active state again in the summary view and the severity indicator associated with it is displayed as a filled square.

Related Information:

[Displaying Hidden Alerts](#)

Hadoop Services View

The **Hadoop Services** view allows you to monitor service status, system use, and key metrics for Hadoop services running on a Hadoop system. The timestamp displays the date and time the values were collected.

Summary View

Displays summary information for all Hadoop services.

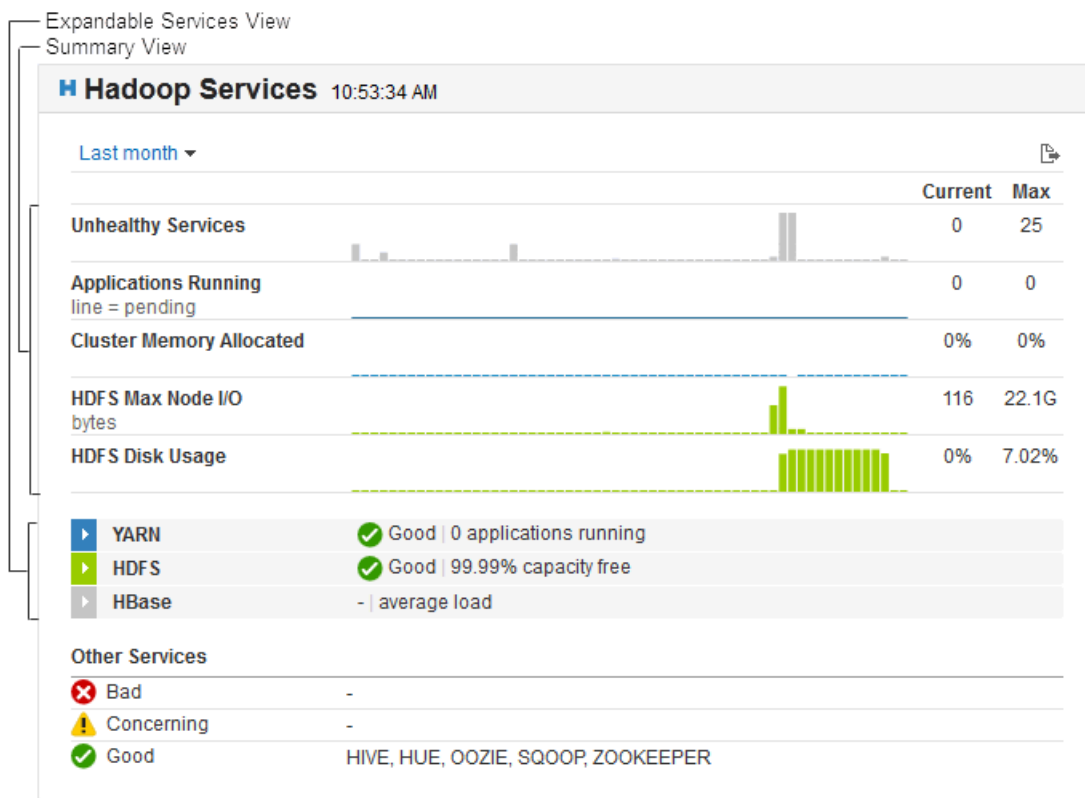
Expandable Services View

Expands the following services to see related metrics, which provide further information on the selected service:

- YARN (CDH and HDP 2.1 and later)

- MapReduce (HDP 1.3)
- HDFS
- HBase

Following is an example of the CDH **Hadoop Services** view.




Viewing Hadoop Services

1. From the **Hadoop Services** view, click ▾ to select a time frame.

Time Frame	Intervals (Calculated Average)
Last hour	1 minute
Last 3 hours	5 minutes
Last 6 hours	10 minutes
Last 24 hours	30 minutes
Last 48 hours	1 hour
Last week	4 hours

Time Frame	Intervals (Calculated Average)
Last month	12 hours

2. To view details about a particular service, select  next to any of the following:

Option	Description
YARN (CDH and HDP 2.1 and later)	Lists applications running, ResourceManager statistics, and NodeManager statistics.
MapReduce (HDP 1.3)	Lists jobs running, jobtracker statistics, map and reduce slots totals and live tasktrackers.
HDFS (CDH and all supported HDP versions)	Lists space usage, namenode statistics, datanode statistics and a total count of HDFS files and directories.
HBase (CDH and all supported HDP versions)	Lists load average, region servers statistics and the master server statistics.

3. [Optional] Select **Open in new window**.

The Hadoop web interface for the service appears, allowing you to view additional details about the service.

Understanding the Summary View


The summary view displays a high-level overview of the performance trends on the selected Hadoop system. Hover over a graph to see an information balloon containing detailed information about the data point for the interval. The right-most bar provides the current reading.

The **Current** column shows the current reading. If you are using the Rewind functionality, the column displays the most recent reading in the time frame defined by Rewind.

The **Max** column shows the maximum value in the selected time frame. This is the actual reading, not a calculated average.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains graph data for the time frame.

1. From the **Hadoop Services** view, select a system and a time frame.
2. Click .
3. Save the file using the browser options.

The file is saved to your download area or to a location that you specify, depending on the browser settings.

Hadoop Services Metrics

A 1024 byte kilobyte is used as the base for all memory-related and space-related metrics that display, such as spool space, disk space, memory usage, and so on. A 1000 byte kilobyte is used as the base for all other metrics.

CDH

Summary View Metrics

Metric	Description
Unhealthy Services	Aggregate number of all CDH services in a Bad or Concerning state (Unknown and Disabled are not included)
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

YARN Metrics

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Completed	Number of YARN applications that executed successfully
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
Cluster Memory Reserved	Percent of available memory reserved across all NodeManager instances
Cluster Memory Skew	Comparison of the largest NodeManager memory allocated to the average memory allocated
Containers Allocated	Number of YARN containers currently allocated across the cluster
Containers Pending	Number of YARN containers currently pending across the cluster
Containers Reserved	Number of YARN containers currently reserved across the cluster
NodeManagers	Number of nodemanagers in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states

Metric	Description
ResourceManager Up Since	Timestamp when the ResourceManager service started
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
ResourceManager UI	Open web interface for the service in a new window

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes	Number of datanodes in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region Servers	Number of region servers in a bad (critical), concerning (degraded), and good state Unknown and disabled states display when there are one or more in those states
Master Server Up since	Timestamp when the master server started
Master Server Heap	Percentage of heap space used in the master server JVM
Master Server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Bad	Services in a critical state
Concerning	Services in a degraded state
Disabled	Services in a disabled state when there are one or more in this state

Metric	Description
Good	Services in a good state
Unknown	Services in a unknown state when there are one or more in this state

HDP 2.1 and Later

Summary View Metrics

Metric	Description
Components Down	Number of components not started
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

YARN Metrics

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Completed	Number of YARN applications that executed successfully
Applications Running	Number of YARN applications currently executing
Cluster Memory Allocated	Percent of available memory allocated across all NodeManager instances
Cluster Memory Reserved	Percent of available memory reserved across all NodeManager instances
Cluster Memory Skew	Comparison of the largest NodeManager memory allocated to the average memory allocated
Containers Allocated	Number of YARN containers currently allocated across the cluster
Containers Pending	Number of YARN containers currently pending across the cluster
Containers Reserved	Number of YARN containers currently reserved across the cluster
NodeManagers Live	Number of NodeManagers currently running
NodeManagers Total	Number of NodeManagers configured on the system
ResourceManager Up Since	Timestamp when the ResourceManager service started
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM

Metric	Description
ResourceManager UI	Open web interface for the service in a new window

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes Live	Number of datanode processes currently running
Datanodes Dead	Number of datanode processes that are currently dead
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region servers Live	Number of region servers that are currently running
Region servers Total	Number of region servers that are configured
Master server Up since	Timestamp when the master server started
Master server Heap	Percentage of heap space used in the master server JVM
Master server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Up	Hadoop Services currently running
Down	Hadoop Services not running
Client Only	Services that do not execute a daemon process in the Hadoop cluster

HDP 1.3**Summary View Metrics**

Metric	Description
Components Down	Number of components not started
Map Tasks Running	Number of map tasks running Number of map tasks waiting in the queue
Reduce Tasks Running	Number of reduce tasks running Number of reduce tasks waiting in the queue
HDFS Max Node I/O	Highest I/O level in bytes on any node in the HDFS system
HDFS Disk Usage	Percentage of space being used

MapReduce Metrics

Metric	Description
Jobs Running	Number of MapReduce jobs currently executing
Jobs Completed	Number of MapReduce jobs that completed successfully
Jobs Failed	Number of MapReduce jobs that failed to complete successfully
Jobtracker Up since	Timestamp when the jobtracker service started
Jobtracker Heap	Percentage of heap space used in the jobtracker JVM
Jobtracker UI	Open web interface for the service in new window
Map slots Occupied	Number of map slots occupied in the MapReduce cluster
Map slots Reserved	Number of map slots reserved in the MapReduce cluster
Reduce slots Occupied	Number of reduce slots occupied in the MapReduce cluster
Reduce slots Reserved	Number of reduce slots reserved in the MapReduce cluster
Tasktrackers Live	Number of tasktrackers currently running
Tasktrackers Total	Number of tasktrackers available on the system

HDFS Metrics

Metric	Description
Capacity Usage	Percentage of used space to overall storage capacity
Datanodes Live	Number of datanode processes currently running

Metric	Description
Datanodes Dead	Number of datanode processes that are currently dead
Files + Directories Total	Total number of files and directories in HDFS
Namenode Up since	Timestamp when the namenode service started
Namenode Heap	Percentage of heap space used in the namenode JVM
Namenode UI	Open web interface for the service in new window

HBase Metrics

Metric	Description
Load Average	Average region load per region server
Region servers Live	Number of region servers that are currently running
Region servers Total	Number of region servers that are configured
Master server Up since	Timestamp when the master server started
Master server Heap	Percentage of heap space used in the master server JVM
Master server UI	Open web interface for the service in new window

Other Services Metrics

Metric	Description
Up	Hadoop Services currently running
Down	Hadoop Services not running

Viewpoint Mobile

Viewpoint Mobile Overview


Viewpoint Mobile enables you to view key Teradata, Aster, and Hadoop system information on your mobile device and take certain actions. It displays the overall health of systems and enables you to drill down to view trends, system health, queries, and alert information for each system. You can abort, change priority, change workload, or release a query using Viewpoint Mobile, and you can load, enter, and run SQL queries.

Viewpoint Mobile Basics

Log In

You access **Viewpoint Mobile** by setting your mobile device's Internet browser to the same location you use for desktop Viewpoint. Viewpoint recognizes you are using a mobile device and automatically takes you to the mobile login screen to enter your username and password.

View Desktop Site

You can access the full Viewpoint desktop by selecting **View Desktop Site** from the **Viewpoint Mobile** menu . You can be logged in to **Viewpoint Mobile** and the desktop site at the same time.

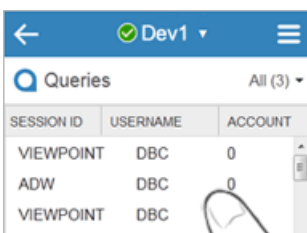
Log Out

You log out from **Viewpoint Mobile** by using the **Log Out** option in the menu .

Working with Summary Tables

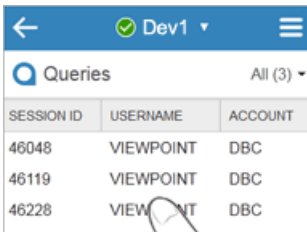
You can perform a number of actions to customize and filter the summary tables displayed in the **Queries** and **Alerts** views:

- Sort the contents of a column by tapping the column header name. Tap again to reverse the sort order.
- Filter a column by pressing the title to expose a filter box. Then press inside the box to display the keyboard for entering text.



SESSION ID	USERNAME	ACCOUNT
VIEWPOINT	DBC	0
ADW	DBC	0
VIEWPOINT	DBC	0

- Change the width of a column by pressing the column heading separator line as you drag your finger to the right or left.



SESSION ID	USERNAME	ACCOUNT
46048	VIEWPOINT	DBC
46119	VIEWPOINT	DBC
46228	VIEWPOINT	DBC

Settings View

Systems to monitor

Enables you to select the systems to monitor.

Format SQL

When selected, automatically wraps the SQL code displayed in the **SQL** view of **Query Details**.

Configure columns

Enables you to select which columns to display in the **Queries** and **Alerts** view tables and their order.

View settings in last

Determines the time period for which to display alerts in the **Alerts** view table.


Display hidden alerts

When selected, hidden alerts are displayed. A hidden alert displays with a line through it.



Restore defaults

Restores alerts settings to their default state.


Selecting Systems to Monitor

1. Tap .
2. Tap **Settings**.
3. Tap the names of the systems you want to monitor.
4. Tap a selected name to clear the selection.

Configuring Columns to Display


1. Tap .
2. Tap **Settings**.
3. Swipe to browse to the **Queries** and **Alerts** settings.
The column names are listed in the order in which they appear in **Viewpoint Mobile** tables.
4. Tap the names of the columns you want to display.
Tap a selected name to clear the selection.
5. [Optional] Press  and drag a column row to reorder the column.

Displaying Hidden Alerts

1. Tap .
2. Tap **Settings**.
3. Swipe to browse to the **Alerts** settings.
4. Tap **Display hidden alerts** at the bottom.
In the **Alerts** view, the alert appears with a line through it and the colored severity icon appears unfilled.




Setting the Alerts Time Period



You can designate the time period for which to display alerts.

1. Tap .
2. Tap **Settings**.
3. Swipe to browse to **Alerts** settings.
4. Under **View settings in last**, tap the time period for which to display alerts.



Home View

The **Home View** is a top-level view that displays a list of the systems for which you have permissions and an icon indicating the overall health of each system.

Default Icon	Default States	Definition
	Healthy	All metrics are within <i>healthy</i> ranges.
	Degraded	At least one metric exceeded a <i>degraded</i> threshold.
	Critical	At least one metric exceeded a <i>critical</i> threshold.

Default Icon	Default States	Definition
	Down	The system is <i>down</i> .
	Unknown	Status of the system is <i>unknown</i> .

In addition, if a system has breached thresholds, the first three threshold names and values are displayed, with any critical metrics displayed first.

-  represents metrics that reached a critical state
-  represents metrics that are degraded

Selecting a System

In order for you to access system information, the Teradata Viewpoint Administrator must assign permissions. For more information, see [Viewpoint Mobile Permissions](#).

You access a system from the **Home** view. If you have drilled down to any other level, you can change the system using the system selection list at the top of the screen.

1. From the **Home** view, tap a listed system to drill down and view its details.

System Overview

The **System Overview** provides a top-level summary view of the trends, system health, queries, and alerts for the selected system.

System Overview Section	Description
Trends	Displays a trends chart that plots the CPU, AWT, MEM, DISK, or I/O during the last 1, 6, or 12 hour period.
System Health	If there are any breached thresholds, the metric name and value are displayed. Critical state metrics are displayed first.
Queries	Displays the number of active, blocked, and delayed queries for a Teradata or Aster system.
Alerts	Displays the number of Critical and High severity level alerts

Selecting a View

The **Trends** view is displayed directly in the top section of the System Overview; there is no need to drill down to see further details. The other views are displayed when you drill down.

1. From **System Overview**, tap inside one of the following sections:

Option	Description
System Health	<ul style="list-style-type: none"> • Tap View system health to display a list of system health metrics, scrolled to the top of the view • Tap a breached system threshold to display a list of system health metrics with the selected metric at the top
Queries	<ul style="list-style-type: none"> • Tap View total number of queries to display a list of all queries • Tap Active, Block, or Delay to display a list of queries in those states
Alerts	<ul style="list-style-type: none"> • Tap View total number of alerts in the last time period to display a list of the alerts generated in the stated time period • Tap Critical or High to display a list of alerts of those severity levels

Trends View

The **Trends** view allows you to analyze resource usage trends for the selected Teradata, Aster, or Hadoop system. Trends are graphed according to CPU, AWT, I/O, MEM, or DISK within a specified time frame, depending on the type of system. The information is refreshed every 60 seconds.

A time frame selection list allows you to monitor current resource usage within the last hour or view resource usage trends over the last 6 or 12 hours.

A graph displays with time on the horizontal axis and metric values on the vertical axis. The following list describes the features in this graph:

Plot Line

The *plot line* is a blue line that represents the average metric value. The plot line appears red where critical thresholds are exceeded. If you press the line and drag your finger to the right or left, information balloons pop up to display the average value of the metric at the time indicated.

Past Averages

Past averages show the average of data samples taken for the last 2 weeks at the same time and day of the week indicated on the horizontal axis.

For example, if you set the time frame to **Last 1 hour**, the last gray segment displays the average of the data sample of the same 10 minutes (aggregation period) for last week and the week prior. The data sample for each week is calculated by averaging the aggregated data over the aggregation period.

Time Frame Selection	Aggregation Period
Last 1 hour	Every 10 minutes
Last 6 hours	Every 10 minutes

Time Frame Selection	Aggregation Period
Last 12 hours	Every 15 minutes

Performance Envelope

The *performance envelope* represents the upper and lower metric values.

Skyline

Use the skyline to compare the current performance with the system performance during the same time period or in the future.

Now

The **Now** point on the horizontal axis indicates the value of the last data point captured. A number to the right of the vertical line also indicates the **Now** data-point value.

Future Past Average

Use the future past average to estimate system workload.

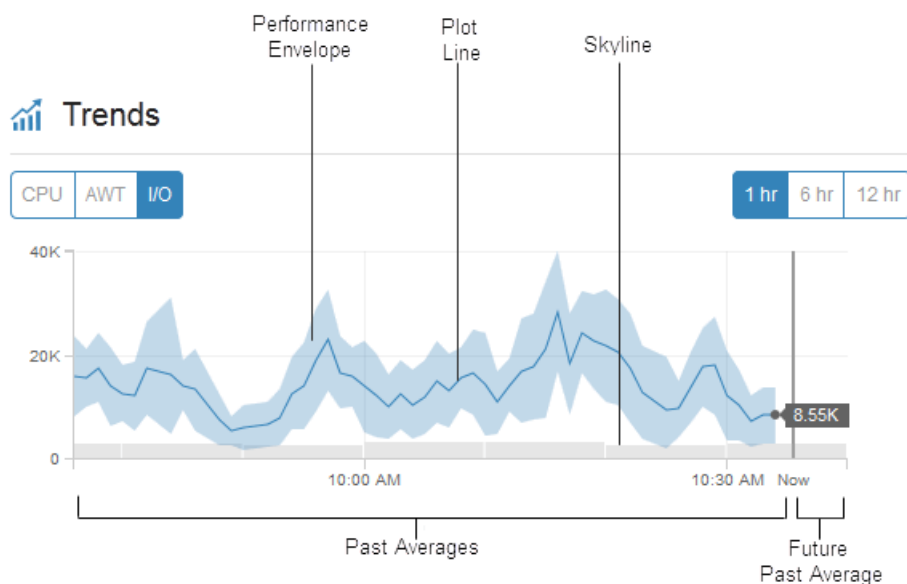
The future past average is based on the following time selections.

Time Selection	Displayed Past Averages	Displayed Future Past Average
Last 1 hour	1 hour	10 minutes
Last 6 hours	6 hours	60 minutes
Last 12 hours	12 hours	2 hours

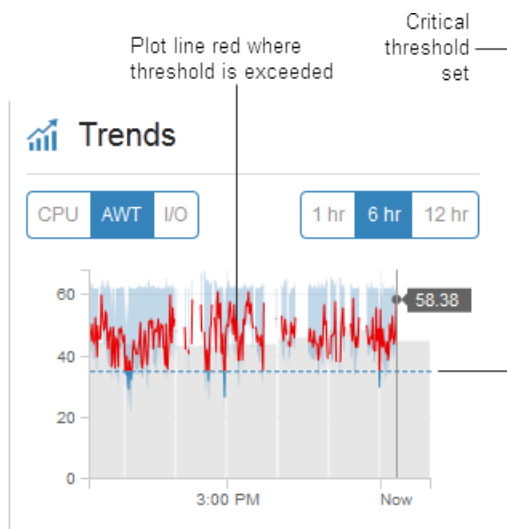
Thresholds

The thresholds settings for CPU and AWT cannot be set independently for the **Trends** view in **Viewpoint Mobile**. The settings are inherited from the critical thresholds defined for calculating system health for each system. For more information, see [Configuring System Health Values](#). Thresholds cannot be set for I/O.

In the following example the I/O for a system in the last hour is displayed.



In the following example the red portion of the plot line shows when the AWT has exceeded the threshold during the last six hours.



Changing the Trends Display for Teradata System

You can change the **Trends** view display for a Teradata system.

1. In the **Trends** view, tap one of the following:

Option	Description
CPU	Displays the percentage of average node CPU use. CPU is calculated as the sum of CPU User and CPU System.

Option	Description
Note: Tool tip contents are not applicable for Viewpoint Mobile view.	Tooltip content for IFX 2.1 and higher systems with Elastic TCore enabled: <ul style="list-style-type: none"> • TDEnabledCPUs % (of NCPUs) : Percentage of Teradata enabled CPUs of total NCPUs • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system
ENABLED CPU Note: In Viewpoint Mobile view this metric is currently not available.	Displays the CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled Tooltip content: <ul style="list-style-type: none"> • MaxTCore: Maximum TCore that the system supports • CurrentTCore: Current enabled TCore setting on the system
AWT	Displays the average number of AMP worker tasks in use on each AMP.
I/O	Displays the number of disk I/Os. I/O is calculated as the sum of physical disk reads and writes per system during the sample period.

The graph changes to represent the chosen metric.

Changing the Trends Display for Aster System

You can change the **Trends** view display for an Aster system.

1. In the **Trends** view, tap one of the following:

Option	Description
CPU	Displays the percentage of the average node CPU use.
MEM	Displays the percentage of the average node memory use.
I/O	Displays the data transfer rate, in bytes per second.

The graph changes to represent the chosen metric.

Changing the Trends Display for Hadoop System

You can change the **Trends** view display for a Hadoop system.

1. In the **Trends** view, tap one of the following:

Option	Description
CPU	Displays the percentage of the average node CPU use. CPU is calculated as the sum of CPU User and CPU System.
MEM	Displays the percentage of the average memory use of the system during a sample period.

Option	Description
DISK	Displays the percentage of disk space being used on a system.

The graph changes to represent the chosen metric.

Setting the Time Frame



You can set the time frame that is used to plot the graph.

1. From the **Trends** view, select a time frame.
The view refreshes, and the graph is redrawn.

System Health View

The **System Health** view lists the values of key metrics used in determining the health of a system. In addition, tapping a system metric opens a graph that displays its trend.

Metric values that have not breached any thresholds are displayed in green. For any breached thresholds the following icons display, and the metric's value is listed in the color indicated as following. Trends graphs for any breached metrics are displayed automatically; you can tap a graph to hide it.

-  represents metrics that reached a critical state. Metric value is displayed in red.
-  represents metrics that are degraded. Metric value is displayed in brown.

You can tap the following metrics to access further details when at least one component is down or when one service is in a bad or concerning state:

- **Components Down** (HDP)
- **Services Bad** (CDH)
- **Services Concerning** (CDH)

Viewing Component Down Details

When the Component Down metric has a value of one or greater, you can access details that show the extent of your exposure to components that are down.

This option is not available on an Aster system.

1. From the **System Health** view, tap **Components Down**.
The Components Down zone is expanded.
2. Tap **View Details**.
A list of down components and their component type appears.

Viewing Services Bad or Services Concerning Details

When the Services Bad or Services Concerning metric (CDH) has a value of one or greater, you can access details of the services in these states.

- From the **System Health** details view, tap any of the following:
 - Services Bad** metric name or value
 - Services Concerning** metric name or value

A list of bad or concerning services and their host names appears.

System Health Metrics for Teradata System

Metric	Description
Active Sessions	Number of sessions with active queries
AMP CPU Skew	Comparison of CPU use on the busiest AMP to the average AMP
AMP I/O Skew	Comparison of disk use on the busiest AMP to the average AMP
AMP Worker Tasks	Average number of AMP worker tasks in use on each AMP
Component Down	Number of components, such as BYNETs or AMPs, that are down
(CPU) System	Average CPU time spent executing kernel system calls or servicing I/O and timer hardware interrupts
(CPU) User	Average CPU time spent executing code on behalf of operating system user processes
CPU Utilization	Average node CPU use. CPU is calculated as the sum of the user CPU, system CPU, and wait I/O usage percentages, depending on which of these are enabled by the Teradata Viewpoint Administrator in the Monitored Systems portlet
(CPU) Wait I/O	Average percentage of CPU time spent waiting for I/O
DBC Disk Space	Available DBC disk space in use
Enabled CPU Usage Note: In Viewpoint Mobile view this metric is currently not available.	CPU utilization based on enabled CPUs for IFX 2.1 and higher systems with Elastic TCore enabled
Max Disk by AMP	Available disk space currently in use
Max Spool by AMP	Available spool space in use
Memory Used	Represents an increase or decrease in the node-level memory. A null value or negative number reports as 0.

Metric	Description
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of disk use on the busiest node to the average node
System Heartbeat	Canary query showing the response time of the most recent system heartbeat in milliseconds
Total Disk Space	Percentage of total disk space currently in use
Additional Canaries	Additional canary queries that are defined for the system and used as part of the monitoring of system health

System Health Metrics for Aster System

Metric	Description
Active Sessions	Number of users and applications currently connected to database
Component Down	Number of nodes that are not available
Component Passive	Number of nodes not processing queries but can be made ready to process queries when needed
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Memory	Average node memory use
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Queen Disk Space	Percentage of used disk space on the queen node
Replication Factor	Number of copies of the user data
Total Space	Percentage of used space to overall storage capacity

System Health Metrics for Hadoop System

CDH

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing

Metric	Description
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Services Bad	Number of services in a critical state
Services Concerning	Number of services in a degraded state
Total Space	Percentage of used space to overall storage capacity

HDP 2.1 and later

Metric	Description
Applications Failed	Number of YARN applications that failed to execute successfully
Applications Running	Number of YARN applications currently executing
Blocks Corrupt	Amount of corrupt blocks in HDFS
Blocks Missing	Amount of missing blocks in HDFS
Cluster Memory Allocated	Percent of the available memory allocated across all NodeManager instances
Components Down	Number of services not started
CPU	Average node CPU use
Max Disk by Node	Largest percentage of used disk space on a node
Name Node CPU	Average node CPU use for nodes running NameNode services
Name Node Heap	Percentage of heap space used in the NameNode JVM

Metric	Description
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
ResourceManager Heap	Percentage of heap space used in the ResourceManager JVM
RPC Latency - RM	Average wait time in queue for ResourceManager service calls
RPC Latency - NN	Average wait time in queue for NameNode service calls
Total Space	Percentage of used space to overall storage capacity

HDP 1.3

Metric	Description
Blocks Corrupt	Number of blocks whose replicas are all corrupt
Blocks Missing	Number of blocks with no replicas anywhere in the cluster
Components Down	Number of service components not running
CPU	Average node CPU use. CPU is calculated as the sum of the user CPU and system CPU usage percentages.
Jobs Failed	Number of jobs that failed
Jobs Running	Number of jobs currently executing in the system
Job Tracker CPU	CPU use for the node running the jobtracker service
Map Tasks Running	Number of map tasks executing in the system
Map Tasks Waiting	Number of map tasks waiting to execute
Max Disk by Node	Amount of used disk space on the node with the most disk space in use
Name Node CPU	Node CPU use for the node running the namenode service
Name Node Heap	Percentage of heap space used in the namenode JVM
Node CPU Skew	Comparison of CPU use on the busiest node to the average node
Node I/O Skew	Comparison of I/O use on the busiest node to the average node
Reduce Tasks Running	Number of reduce tasks executing in the system
Reduce Tasks Waiting	Number of reduce tasks waiting to execute
RPC Latency JT	Average wait time in queue for jobtracker service calls
RPC Latency NN	Average wait time in queue for namenode service calls
Total Space	Percentage of used space to overall storage capacity

Queries View

The **Queries** view displays a summary table listing the sessions running on the system and allows you to drill down to see more details about each session.

Queries States

The **Queries** list before the summary table allows you to select a state so that only queries in that state are displayed.

Teradata System States

All

Sessions in the system

Not Idle

Sessions in any state except idle

Active

Sessions processing queries

Block

Sessions processing a query waiting for a locked resource, such as a database table or view

Defer

Sessions waiting under deferred queue to be processed

Delay

Sessions waiting to be processed

Abort

Sessions in the process of aborting (rolling back changes made by the query)

Resp

Sessions running a query that has completed and is sending (responding) spooled data back to the user

Idle

Sessions not currently running a query

Parse

Sessions running a query that is being parsed and has not begun to execute

Other

Sessions whose status is unknown and do not fall into another category

QTDelayed

Sessions delayed due to a queue table restriction

SesDelayed

Utility sessions that are in the workload delay queue

Response-Held

In SLES 11, sessions running a query whose response is being held until a specified number of seconds elapses

Aster System States**All**

Sessions in the system

Active

Sessions processing queries

Idle

Sessions not running

Canceled

Processes that were stopped

Pending

Sessions waiting to be processed

Details View

You can tap on any row in the queries summary table in the **Queries** view to see **Query Details** views. When you drill down, an **Overview** view is displayed. Other views may be available, depending on the kind of Teradata system and the nature and state of the query. If other details views are available, you can view them by swiping left and right or using the arrow keys at the top. The query details views include:

Overview view

Provides detailed information about key metrics for the selected session and its queries.

SQL view

Displays the SQL for the selected query.

Explain view

Displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in a SQL session.

Skew view

Displays details about the level of skew in the query or session.

Query Band view

Displays the profile, session, and transaction query bands for the selected query.

Blocked By view

Identifies the other queries blocking the selected query.

Blocking view

Indicates the number of sessions the selected session is blocking, and the significance and impact of the blocking.

Delay view

Displays details about all rules that are delaying a query.

Queries Metrics for Teradata System

Metrics for a Teradata system that display in the **Viewpoint Mobile Queries** view are listed in the following table.

Metric	Description	Type
Account	Account from which a query was submitted	String
Blocked Time	How long the query has been blocked	Number
CPU Skew	Percentage of how inefficiently a query or job was executed in terms of CPU parallelism. A value of 0% indicates the query was executed perfectly in parallel, meaning each AMP consumed identical amounts of CPU. A value of 50% indicates a single AMP consumed twice as much CPU as the average. As CPU Skew approaches 100%, it indicates the query or job was not executed in parallel.	Percent
CPU Use	Percent of available CPU seconds on the system used during the last sampling period	Percent
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules	Number
Delayed By	Names of the throttles causing the query to be delayed	String
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample	Number
Δ I/O	I/O count since the last sample	Number

Metric	Description	Type
Duration	How long the query has been running	Number
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Number
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.	Number
In State	How long the query has been in the current state	Number
Partition	Partition in which the query is running	String
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation	Number
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used	String
Query Band	Entire query band string. Query bands are a set of name-value pairs defined by the user to tag sessions or transactions with an ID through a SQL interface.	String
Req CPU	CPU seconds needed to run the query	Number
Req CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query	Percent
Req Count	Number of queries the session has executed	Number
Req I/O	Disk I/Os performed to run the query	Number
Req I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query	Percent
Session ID	Unique session identifier	Number
Snapshot CPU Skew	CPU skew during the last sample	Percent
Snapshot I/O Skew	I/O skew during the last sample	Percent
Spool	Spool space the query requires, using a 1024 byte kilobyte as the base	Number
Start	Time that the query started running on Analytics Database	Number
State	Text describing the current state of the query	String
State Icon	Icon representing the current state of the query	Icon
Temp Space	Temp space used by the query, using a 1024 byte kilobyte as the base	Number
Unnecessary I/O	Ratio of I/O to CPU for the current query	Number
Username	Name of the user who submitted the query	String

Metric	Description	Type
Workload	Workload in which the query is running	String

Queries Metrics for Aster System

Metrics for an Aster system that display in the **Viewpoint Mobile Queries** view are listed in the following table.

Metric	Description	Type
Database	Name of the database on which the process is running	String
In State	How long the process has been in the current state	Number
Process ID	Unique process identifier	Number
Session ID	Unique session identifier	Number
Start	Time the process began running	Number
State	Process state, such as active, canceled, pending	String
State Icon	Symbol of the process state	Icon
Type	Values for process types are as follows: <ul style="list-style-type: none"> • SQLMR (SQL-MapReduce) • SQL • Teradata Import • Teradata Export 	String
Username	Name of the user who submitted the process	String
Workload	Workload in which the process is running	String

Viewing Query Details by State

- From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** to view all queries.
 - Tap **Active**, **Block**, or **Delay** to view the queries in those states.

Viewing Query Details by Session

- From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** to display all queries.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
- Tap the row in the table for the session you want to see.

Overview View for Teradata System

The **Overview** view provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Snapshot Information

Snapshot Information	Description
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Impact CPU	(Teradata Database 15.0 and earlier) CPU impact on the system based on the highest utilized AMP
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM • Request - Query is assigned manually to a workload using Change Workload • Session - Queries initiated in that session are assigned manually to a workload using Change Workload This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Session Information

Session Information	Description
User	Name of the user that submitted the query

Session Information	Description
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system

Overview View for Aster System

The **Overview** view displays detailed information about key metrics for the selected session. The metric values provide a view of the query status on the system.

Query Information

Query Information	Description
Process ID	Process identifier
State	Process state, such as active, canceled, pending
In State	How long the process has been in the current state
Start	Time the process started running

Workload Information

Workload Information	Description
Workload Policy	Name of the workload policy used to manage the process. Defines a set of related queries, and allows them to be prioritized in a similar manner.
Service Class	Name of the service class used by the workload policy. Specifies the share of system resources allocated to a workload as a function of the class's priority and weight values.
Priority	Indicates the importance of the process. <ul style="list-style-type: none"> • High • Medium • Low

Workload Information	Description
Weight	Number ranging from 1 to 100 that indicates the precedence of the process within the priority level, increasing the resources that can be allocated for the process. A higher value indicates a greater level of precedence.

Session Information

Session Information	Description
User	Name of user that submitted the query
Database	Name of the database on which the process ran
IP Address	IP address of the user that submitted the process



SQL View

The **SQL** view displays the SQL for the selected query. If the SQL is from a stored procedure on a Teradata Database 15.10 and later system, the name of that stored procedure is also displayed. This information is read-only.

The **SQL** view is available only when a query is active, blocked, or delayed. For Teradata systems, the view is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

By default, the SQL is displayed on one long line. You can press and move your finger slowly to the right to read the text. If you want the SQL display to wrap automatically, tap the **Wrap SQL** check box before the text display.

Explain View for Teradata System

The **Explain** view displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session. The view is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status. If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

Step Information	Description
Step Number	<ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black number box. Active steps are indicated by a pulsating number box (flashes blue). Steps to run are at the bottom of the list and indicated by a white number box.

Step Information	Description
Confidence Level Indicator Icon	<ul style="list-style-type: none"> - No confidence in the estimate - Low confidence in the estimate - High confidence in the estimate - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Explain View for Aster System

The **Explain** view displays the Explain text that shows the progress of a query. Each Explain step has a colored box that indicates the status of the step. This information is read-only.

By default, all steps and details are displayed. Tap **Show summary only** to view a shortened version of the Explain data.

Step Information	Description
Step Status Icon	<p>Indicates the status of each Explain step:</p> <ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black box. Active steps are indicated by a pulsating box (flashes gray and white). Steps to run are at the bottom of the list and indicated by a white box.
Time	Duration of the step
Type	Category of step, such as Command, CommitPreparedTransaction, DataTransfer, Import, PrepareTransaction, and Query
Phase/Detail	Summarized or full explanation of query that executes in this step. Can be used to diagnose the cause of the failure or slowness.

Skew View for Teradata System

The **Skew** view displays details about the level of skew in the query or session. The **Skew** view does not display when the **Delay** view is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count

Skew Information	Description
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Query Band View for Teradata System

The **Query Band** view displays the profile, session, and transaction query bands for the selected query. This information is read-only. The view is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Blocked by View for Teradata System

When a query is blocked, use the **Blocked By** view to identify the other queries blocking the selected query. By reviewing the number and type of queries blocking the selected query, you can identify blocking issues on the system, determine the impact of this blocking on the selected query and, based on this information, decide the best course of action to resolve the issue.

The information in this view is read-only. The view is available only when the selected query is blocked. You can drill down into the session information for the blocking session.

The following information is available in the noted versions of Analytics Database.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		Available
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock 		Available
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	

Blocked By Information	Description	14.10 and earlier	15.00 and later
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Blocking View for Teradata System

When a query is holding locks that are blocking other queries from completing, use the **Blocking** view to gauge the impact of this blocking. By analyzing the number and type of queries blocked by this query, you can decide the best course of action to resolve the issue.

This information in this view is read-only. The view is available only when the selected query is blocking other queries and when connected to Teradata Database 15.0 and later.

Blocking Information	Description
Blocking Type	Indicates the significance of the blocker. Available values are: <ul style="list-style-type: none"> • Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked • Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock • Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>
Blocking Locks	A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName The information displays in the following order: <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock

Blocking Information	Description
	5. Name of the table object over which a lock conflict is preventing the session from being granted a lock
Session ID	Unique session identifier of the query that is blocking
Blocked Time	How long the query has been blocked, displayed as <i>d +hh :mm :ss</i>
Username	Name of the user that is running the query that holds the lock
Workload	Name of the workload where the query is actively running
Query Band	String of key/value pairs of the query band for the session or transaction

Delay View for Teradata System

The **Delay** view displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Managing Queries and Sessions for Teradata System

You can manage queries and sessions to improve workload performance for Teradata systems.

Abort

Abort the selected query or session

Change Priority

Change the priority of the selected query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change priorities or workloads, or release queries. If you log off, you must log in again.

Note:

Change Workload is available only if Teradata Active System Management (TASM) is enabled. If TASM is disabled, **Change Priority** is available. If you do not see **Change Workload** or **Change Priority** in the list, the system you are monitoring does not support these features or you do not have permission to use them. If the query you are monitoring is delayed, only **Release Query** is available.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** and then use the query state selection list or column filters to narrow the list.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
2. Tap the row of the query you want to abort.
3. Tap **Actions** and select **Abort**.
4. Log in to Analytics Database, if prompted.
5. Tap **Next**.
6. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
7. Tap **Next**.
8. Tap **Next** to confirm your selection, and tap **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** and then use the query state selection list or column filters to narrow the list.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
2. Tap the row of the query you want to abort.
3. Tap **Actions** and select **Change Priority**.
4. Log in to Analytics Database, if prompted.
5. Tap **Next**.
6. Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.
 - Type an account string.
7. [Optional] Select the check box to use the account string as the default for the selected session. The priority will be changed for the selected query and all subsequent queries in the current session.
8. Tap **Next**.
9. Tap **Next** to confirm your selection, and tap **OK**.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

1. From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** and then use the query state selection list or column filters to narrow the list.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
2. Tap the row of the query whose workload you want to change.
3. Tap **Actions** and select **Change Workload**.
4. Log in to Analytics Database, if prompted.
5. Tap **Next**.
6. Select a new workload from the list, and tap the check box to use this as the default workload for this session.
7. Tap **Next**.

8. Tap **Next** to confirm your selection, and tap **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** and then use the query state selection list or column filters to narrow the list.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
2. Tap the row of the query you want to release.
3. Tap **Actions** and select **Release Query**.
4. Log in to Analytics Database, if prompted.
5. Tap **Next**.
6. Tap **Next** to confirm your selection, and tap **OK**.

Managing Sessions for Aster System

You can manage sessions to improve workload performance for Aster systems.

Abort

Abort the selected process.

Aborting a Process

For Teradata Aster systems, you can abort a non-idle process that is blocking other processes or consuming too many resources.

1. From the **Queries** section of **System Overview**, do one of the following:
 - Tap **View total number of queries** and then use the query state selection list or column filters to narrow the list.
 - Tap **Active**, **Block**, or **Delay** to display the queries in those states.
2. Tap the row of the process you want to abort.
3. Tap **Actions** and select **Abort Process**.
4. Log in to Teradata Aster, if prompted.
5. Tap **Next** to confirm you want to abort the process.
A confirmation appears indicating the abort process was successfully submitted.
6. Tap **OK**.

Alerts View

The **Alerts** view allows users to view triggered alerts and are specific to your system. The summary table displays a list of all alerts generated in a designated time period, or can be filtered to display only the alerts of a selected severity level. The information in the view is refreshed every 60 seconds.

An alert is an event that the Teradata System Administrator defines as being significant. The Teradata System Administrator assigns alert severity levels to rank alerts, and can also include an explanatory message. Following are the severity levels:

- Critical
- High
- Medium
- Low

The timestamp in this view updates when the number of alerts changes.

Alert Example

The Teradata System Administrator can define that a database exceeding a certain amount of storage usage triggers an alert. After the usage level is exceeded, an alert appears in the **Alerts** view to inform the user that the threshold was exceeded.

You can change a setting so hidden alerts display. For more information, see [Displaying Hidden Alerts](#).

Viewing Alert Details

1. From the **Alerts** section of **System Overview**, do one of the following:

Option	Description
View all alerts first	<ol style="list-style-type: none"> a. Tap View total number of alerts in the last time period to display all alerts generated in that time period. b. Tap ▼ before the table and select the severity level.
View critical or high severity alerts	Tap Critical or High to display a list of alerts of those severity levels.

2. Tap the row in the table for the alert you want to see.

The alert details for an alert displays details about the selected alert and displays the following:

- **General:** Alert type, severity, source, timestamp, resulting action, and alert criteria.
- **Properties :** Alert properties and their associated values. For metric descriptions, see [Alerts](#).
- **Messages:** Any messages issued.

Alert Metrics

Metric	Description
Alert Name	Name of the alert
Alert Type	Type of alert
Message	Message providing more information about the alert
Severity	Severity level assigned to the alert
Severity Icon	Colored icon representing the severity level
Source	Source (Viewpoint, Teradata Alerts, Server Management, Performance Data Collection, or Teradata Ecosystem Manager)
Timestamp	Date and time when the alert was issued

Hiding Alerts Basics

By default, the **Alerts** view displays all alerts that have been generated. You can hide some alerts so they no longer appear in the summary list. For example, you might want to hide all alerts generated for database space conditions or alerts associated with conditions that were subsequently resolved.

You can also display hidden alerts, and then unhide them. When you display hidden alerts, the alert rows appear with a line through them, and their associated severity indicators appear as unfilled squares. If you tap the alert to view its details, the severity indicator next to the alert name also indicates the hidden state of the alert.

If you hide alerts in one of the following, the alerts are hidden in the other two:

- **Alert Viewer** portlet
- **Viewpoint Dashboard**
- **Viewpoint Mobile**

The **Settings** view allows you to display previously hidden alerts.

Hiding Alerts

1. In the **Alerts** view, tap the row of the alert you want to hide.
2. In the **Alerts Details** view, tap **Actions** and select **Hide Alert**.
The alert does not appear, or appears with a line through it if hidden alerts have been set to display in **Settings** for alerts.

Unhiding Alerts

Prerequisite:


You must display alerts that were previously hidden before you can unhide them.

You can select and unhide alerts that were previously hidden. When you unhide an alert, it is displayed again and the severity indicator associated with it is displayed as a filled square.

1. In the **Alerts** view, tap the row of the alert you want to unhide.
2. In the **Alerts Details** view, tap **Actions** and select **Unhide Alert**.
The alert displays in its normal active state again in the summary view and the severity indicator associated with it is displayed as a filled square.

SQL Mobile View

The **SQL Mobile** view enables you to enter SQL queries and retrieve query results from a Teradata system. You can connect to a system to run a single statement query that you enter or load a query that you saved using the **SQL Scratchpad** portlet in the Viewpoint desktop site.

You access the **SQL Mobile** view from the **Viewpoint Mobile** menu .

You enter and run queries using the query box section at the top of the view, and the query results are displayed in the bottom section of the view.

You must use the **Connect** button at the top of the view to log in to a system before you can load or run a query.

Query Box

The query box allows you to build queries and submit them to a Teradata system. The results are available under the query box after you run the queries.

Query Box

Allows you to enter and edit SQL queries.

Load

Displays and enables you to load queries saved when using the **SQL Scratchpad** portlet or queries run during the current session.

Run

Runs statements in the query box and provides results in the query results area.

Clear

Clears all text entered in the query box and displayed in the query results area of the **SQL Mobile** view.

Cancel

Cancels the query that is currently running.

Connecting to a System

1. In the **SQL Mobile** view, tap **Connect**.
2. Enter the login information and tap **Connect**.

Disconnecting from a System

1. In the **SQL Mobile** view, tap **Disconnect**.

Running a Query

You can enter and submit a single statement query to an Analytics Database.

1. Enter a single statement query in the query box.
2. Tap **Run** to submit the query you entered against the currently connected system.
The results are displayed in a table in the lower section of the **SQL Mobile** view.

Query Statements

You can enter single statement SQL queries using Data Definition Language (DDL), Data Control Language (DCL), or Data Manipulation Language (DML) statements.

DDL statements

Begin with SQL keywords, such as CREATE, ALTER, DROP, DELETE, MODIFY, HELP, SHOW, SET, REPLACE.

DCL statements

Begin with SQL keywords GRANT, REVOKE, GIVE.

DML statements

Begin with SQL keywords, such as SELECT, INSERT, UPDATE.

Enter the query, optionally followed by a semicolon, in the query box.

Canceling a Long-Running Query

Occasionally, a query may take longer to execute than you expected. You can cancel the query before it completes.

1. While the query is running, tap **Cancel**.

Query Results

The query results section in the lower part of the **SQL Mobile** view displays the results of the query that is run, or a message if the query entered is not valid.

Viewing Query Results

1. When the query box has a query entered, tap **Run** to submit the query.
The results are displayed in a table in the lower part of the **SQL Mobile** view.

History View

The **History** view in **SQL Mobile** contains a list of queries submitted in the current session. The following is displayed for each query:

- The date and time of the query
- The SQL statement submitted
- An icon indicating whether the query completed successfully
- A **Load** button that can be tapped to load the query so that it can be run again

Viewing Previously Run Queries

You can view a list of queries submitted in the current session. You can load a query from the table for reuse in the same session.

1. In the **SQL Mobile** view, tap **Load**.
Saved queries, if any, are listed.
2. Swipe to browse to the **History** view.
The queries run during the current session are listed.

Loading a Previously Run Query

You can copy a query submitted in the current session to the query box for reuse or editing.

1. In the **SQL Mobile** view, tap **Load**.
Saved queries, if any, are listed.

2. Swipe to browse to the **History** view.
The queries run during the current session are listed.
3. Tap **Load** next to the query you want to run again.
The query is copied into the query box in the **SQL Mobile** view.

Saved View

The **Saved** view in **SQL Mobile** displays a list of queries saved using the **SQL Scratchpad** portlet. The following is displayed for each query:

- The name given to the query when it was saved using the **SQL Scratchpad** portlet
- A **Load** button that can be tapped to load the query so that it can be run again

Viewing Saved Queries

You can view a list of queries saved when using the **SQL Scratchpad** portlet.

1. In the **SQL Mobile** view, tap **Load**.
Saved queries, if any, are listed.

Loading a Saved Query

1. In the **SQL Mobile** view, tap **Load**.
Saved queries, if any, are listed.
2. Tap **Load** next to the query you want to run again.
The query is copied into the query box in the **SQL Mobile** view.

Viewpoint Monitoring

Viewpoint Monitoring Overview

The **Viewpoint Monitoring** portlet allows you to analyze resource usage trends on one or more Teradata Viewpoint servers or other managed servers. Trends are graphed according to metrics such as CPU, memory, and throughput within a specified time frame.

The **Viewpoint Monitoring** view displays a graph with time on the horizontal axis, metric values on the vertical axis, and a different color for each metric plot line. Plot lines show the average metric values as a line and have enhanced interactive features, such as performance envelopes and information balloons that show the actual, minimum, and maximum metric values. A time frame selection list allows you to monitor current usage within the last hour or view usage trends over the last 24 hours, 1 week, 1 month, 3 months, or 6 months.

The **Settings** view allows you to select and organize metrics to display, set thresholds, and adjust the vertical axis range for each metric.

Viewpoint Monitoring View

The **Viewpoint Monitoring** view displays a graph that shows Teradata Viewpoint server resource usage. Monitor one or more Teradata Viewpoint servers in the time frame using the same or different metrics. Select metrics with different thresholds and vertical axis ranges.

The following list describes the features in this view:

Toolbar

Contains the time frame list. Select a time frame from the list to change the graph. **Export** creates a .csv file containing selected data.

Line Graph

Plots the time frame on the horizontal axis and the metric values on the vertical axis. The *plot line* is a colored line that represents the average metric values. Click a plot line to see the highlighted plot line and performance envelope in a lighter shade of the assigned color. The *performance envelope* represents the upper and lower metric values.

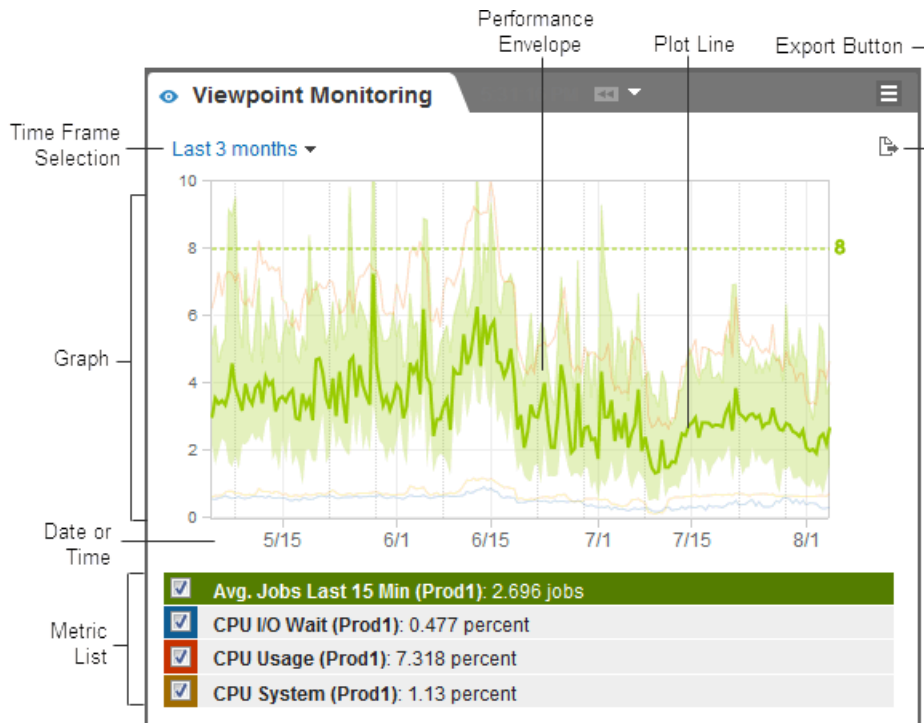
Time Frame

Shows the frame of time selected from the time frame list on the toolbar. For example, if one hour is selected, times from the previous hour appear across the bottom of the graph.

Metrics

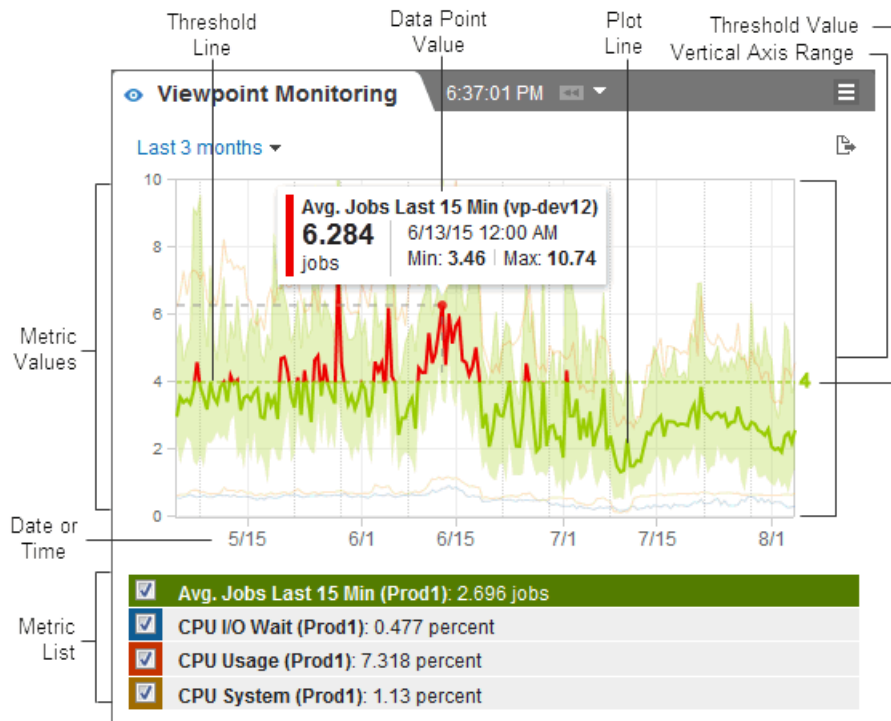
Lists up to 10 metrics that are defined in the **Settings** view. Each list position is assigned a color and a check box. The metric name appears along with the name of the system being

monitored, if there is more than one system enabled. Select a metric name to highlight the plot line and see the performance envelope in the graph.



Metric Thresholds

Set thresholds to highlight metric values that are outside normal operating ranges. Set maximum vertical axis ranges to limit the range of values displayed for selected metrics. After a metric threshold is set in the **Settings** view, the threshold line is displayed across the graph with the threshold value to the right of the graph when you select a metric plot line. As you hover over the selected plot line, information balloons display the data point value and the minimum and maximum performance envelope values. The plot line and information balloons change to red when the threshold is exceeded during the time frame. Threshold settings are optional and can be set any time after the metric has been configured.



Setting the Time Frame

You can set the time frame that is used to plot the graph.

1. In the toolbar, select a time frame from the list.
The portlet refreshes, and the graph is redrawn.

Removing Metrics from the Graph

You can disable a metric plot line to remove it from the graph and the metric remains in the list.

1. Clear a colored check box in the metric list.

Toolbar

The toolbar allows you to choose the time frame to display in the graph. The graph displays the oldest data on the left and the most recent data on the right. For each metric, data is collected every 15 seconds and averaged according to the time frame chosen.

Option	Description
Last 1 hour	Displays the metric values for the last hour, plotted by minute
Last 24 hours	Displays the metric values for the last 24 hours, plotted by 15-minute periods

Option	Description
Last 1 week	Displays the metric values for the last week, plotted by hour
Last 1 month	Displays the metric values for the last month, plotted by 4-hour periods
Last 3 months	Displays the metric values for the last 3 months, plotted by 12-hour periods
Last 6 months	Displays the metric values for the last 6 months, plotted by day

Selecting a different time frame causes the portlet to immediately recalculate and update the graph based on the data points collected for the metrics.


The **Export** button allows you to create a .csv file containing selected data.

Exporting Metrics

You can export data to a .csv file for further analysis and formatting. The exported .csv file contains data for the selected time period. Minimum, average, and maximum is exported for all metrics configured in the **Settings** view.

1. In the toolbar, select a time frame from the list.

Option	Frequency
Last 1 hour	Data points at 1-minute intervals
Last 24 hours	Data points at 15-minute intervals
Last 1 week	Data points at 1-hour intervals
Last 1 month	Data points at 4-hour intervals
Last 3 months	Data points at 12-hour intervals
Last 6 months	Data points at 1-day intervals

2. Click .
3. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
4. [Optional] Reformat exported data to match the format in the view.

Viewpoint Monitoring Metrics

The following metrics can be used to analyze resource usage.

CPU Metrics

Metric	Description	Type
CPU Idle	Percent of total CPU time spent idle	Percent
CPU Involuntary Wait	Percent of total CPU time spent waiting involuntarily	Percent
CPU I/O Wait	Percent of total CPU time spent waiting for I/O	Percent
CPU System	Percent of total CPU time used by system processes	Percent
CPU Usage	Percent of CPU time used by user and system processes	Percent
CPU User	Percent of total CPU time used by user processes	Percent

Memory Metrics

Metric	Description	Type
Memory Buffers	Amount of memory currently used as a disk buffer cache	KB
Memory Cache	Amount of cache memory	KB
Memory Free	Amount of idle memory	KB
Memory Total	Total amount of physical memory in the system	KB
Memory Total Swap	Amount of memory used for swap space	KB
Memory Usage	Total amount of memory used by the system, calculated as the total memory minus the amount of free memory	KB
Memory Virtual	Amount of virtual memory used	KB
Procs Sleeping	Number of processes in uninterruptible sleep	Number
Procs Waiting	Number of processes waiting for run time	Number
Swap In	Amount of memory swapped in from disk per second	KB/sec
Swap Out	Amount of memory swapped to disk per second	KB/sec

Throughput Metrics

Metric	Description	Type
Avg. Jobs Last 15 Min	Average number of processes running or able to run during the last 15 minutes	Number
Avg. Jobs Last 5 Min	Average number of processes running or able to run during the last 5 minutes	Number
Avg. Jobs Last Min	Average number of processes running or able to run during the last minute	Number


Metric	Description	Type
Blocks Received	Number of blocks received from a block device per second	Number
Blocks Sent	Number of blocks sent to a block device per second	Number
Context Switches Per Second	Number of context switches per second	Number
Interrupts Per Second	Number of interrupts per second	Number

Settings View

The **Settings** view allows you to select which systems and metrics to monitor and display. You can set thresholds and vertical axis ranges for selected metrics.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.




1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:


Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Adding Metrics



You can add metrics that are then plotted in the graph.

1. Click  in the portlet frame and select **Settings**.
2. [Optional] Select the **Enable** check box to enable the metric for display.
3. [Optional] Select a system from the list.
The **System** list is available when more than one Teradata Viewpoint server is enabled for the Data Collection Service.
4. Select a metric from the list.
You can add  or remove  metrics. Up to 10 metrics can be listed.
5. [Optional] Enter a **Threshold** value.
6. [Optional] Enter a **Vertical Axis Range** value.

7. [Optional] Click  and drag the row to reorder the metrics for display.
8. Click **OK**.


Changing the Metrics Display Order

You can change the order of the metrics that are listed under the graph. Reordering the metric list affects which color is assigned to the metric.

1. Click  in the portlet frame and select **Settings**.
2. On the metric row, click  and drag the row to its new location.
3. Click **OK**.



Managing Metrics

You can enable and disable metrics from the graph. When a metric is disabled, the metric plot line is not displayed in the graph. Data points are still being collected, so the metric is included in the list under the graph and can be reactivated at a later time.

1. Click  in the portlet frame and select **Settings**.
2. Do one of the following to enable or disable a metric:
 - Select **Enable** on the metric line to enable the metric.
 - Clear **Enable** on the metric line to disable the metric.
3. Click **OK**.

Deleting Metrics

You can delete metrics from the graph. When a metric is deleted in the **Settings** view, the metric plot line is not displayed in the graph, and the metric is not included in the metric list.

1. Click  in the portlet frame and select **Settings**.
2. Click  on the row of the metric you want to delete.
The metric row disappears.
3. Click **OK**.

Workload Management

Workload Management Scenarios

How Can I Use Throttles to Limit Concurrent Queries to Improve Throughput?

Queries complete faster when there are fewer running on the system at the same time. Using throttles, you can control the amount of work you allow onto the system, and thereby increase the overall system throughput.

The following are examples of ways you can use throttles to improve the flow of traffic on your system:

- Limit the number of queries that can run concurrently
- Limit the number of queries a specific user can run concurrently
- Limit the number of users that can run queries concurrently

The following is an overview of the process:

1. [Edit a ruleset.](#)
2. [Create a system throttle.](#)
3. [Adding Request Source Classification Type](#)
4. [Apply throttle limits.](#)
5. [Review the delay queue.](#)

Scenario assumptions




This scenario assumes you already have the following set up:

- Teradata Database version: 16.00
- Operating system: SLES 11
- System: PROD1
- Ruleset: FILTERS/THROTTLES

Example: Limiting the Number of Queries That Can Run Concurrently



Suppose you want to improve the system throughput by limiting the number of queries that can run concurrently on a system to never exceed three queries. To do this, place any queries that exceed this limit in a delay queue. When one query completes, the next query in the queue starts executing.


1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.

- b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **FILTERS/THROTTLES**.
2. Create a new system throttle.
 - a. From the ruleset toolbar, click **Throttles**.
 - b. Next to **System Throttles**, click .
 - c. Enter the name of the throttle, in this case **Concurrent Queries**.
 - d. Select the **Rule Type** option **Collective**.
 - e. Click **Save**.
3. Add classification criteria for the system throttle to apply to all users.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Request Source**.
 - c. Click **Add**.
The **Edit Request Source Criteria** dialog box appears.
 - d. From the **Source Type** list, select **Username**.
 - e. From the **Usernames** list, select **All**.
 - f. Click **Include**.
 - g. Click **OK**.
 - h. Click **Save**.
4. Add throttle limits to allow only three queries to execute at a time.
 - a. Click the **State Specific Settings** tab.
 - b. Under **Default Settings**, type 3 in the box to set the concurrency limit.
 - c. Select **Delay** to place queries exceeding the limit in a delay queue.
 - d. Click **Save**.
5. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **FILTERS/THROTTLES** ruleset, click  and select **Make Active**.
6. [Optional] Review the queries in the delay queue.
 - a. Open the **Workload Monitor** portlet.
 - b. Click  in the portlet frame and select **Settings**.
 - c. Click the **Systems** tab.
 - d. From the **Systems** list, select **PROD1**.
 - e. Click **OK**.
 - f. From the toolbar, click on **Delayed Requests** to view metrics for requests that were delayed by throttles.
 - g. Click on the Session ID to view more details.

Example: Limiting the Number of Queries That a Specific User Can Run Concurrently

Suppose you want to prevent overloading the system and improve the system throughput by limiting a single user's queries to process only one at a time. For example, John Smith has 25 queries that he submits before he leaves from work at night and he needs the results only when he returns. You can place any queries that exceed John's limit in a delay queue. When one query completes, the next query in the queue starts executing.

1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.
 - b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **FILTERS/THROTTLES**.
2. Create a new system throttle.
 - a. From the ruleset toolbar, click **Throttles**.
 - b. Next to **System Throttles**, click .
 - c. Enter the name of the throttle, in this case John Smith Concurrent Queries.
 - d. Select the **Rule Type** option **Collective**.
 - e. Click **Save**.
3. Add classification criteria to limit the rule to John Smith.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Request Source**.
 - c. Click **Add**.
The **Edit Request Source Criteria** dialog box appears.
 - d. From the **Source Type** list, select **Username**.
 - e. Under **Username**, enter the user ID for the database session: johnsmith.
 - f. Click **Include**.
 - g. Click **OK**.
 - h. Click **Save**.
4. Add throttle limits to allow only one query to execute at a time.
 - a. Click the **State Specific Settings** tab.
 - b. Under **Default Settings** at the bottom of the **Throttles** view, type 1 in the box to set the concurrency limit.
 - c. Select **Delay** to place queries exceeding the limit in a delay queue.
 - d. Click **Save**.
5. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **FILTERS/THROTTLES** ruleset, click  and select **Make Active**.
6. [Optional] Review the queries in the delay queue.
 - a. Open the **Workload Monitor** portlet.

- b. Click  in the portlet frame and select **Settings**.
- c. Click the **Systems** tab.
- d. From the **Systems** list, select **PROD1**.
- e. Click **OK**.
- f. From the toolbar, click on **Delayed Requests** to view metrics for requests that were delayed by throttles.
- g. Click on the Session ID to view more details.

How Can I Prioritize Tactical Queries and Define Their SLGs?

Tactical queries are short, critical queries that have defined service-level goals (SLGs). Workload management SLGs define criteria for the response time goals for the queries. Many times, tactical queries need to complete in less than a second. When a query is classified as tactical, all available resources are used to process it, enabling it to meet the SLGs.

The following is an overview of the process:

1. [Edit a ruleset.](#)
2. [Create a workload.](#)
3. [Define the service-level goal.](#)
4. [Assign tactical priority.](#)
5. [Review and analyze missed service-level goals.](#)

Scenario assumptions



This scenario assumes you already have the following set up:

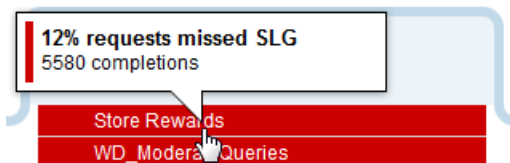
- Teradata Database version: 16.00
- Operating system: SLES 11
- System: PROD1
- Ruleset: SLG

Example: Prioritizing Tactical Queries and Defining Their SLGs

Suppose a grocery store chain wants to give personalized coupons. At checkout, the customers slide their rewards cards. The database is queried for suitable coupons, and the results need to return within one second. You want to set the SLG to have a response time of no more than one second, 98 percent of the time.

1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.
 - b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **SLG**.
2. Create a workload.

- a. From the ruleset toolbar, click **Workloads**.
 - b. Next to **Workloads**, click .
 - c. In the **General** tab, enter a name for the workload, in this case Store Rewards.
 - d. Under **Workload Management Method**, select **Tactical**.
 - e. Click **Save**.
3. Add classification criteria to limit the rule to the Store Rewards application, that submits the queries to the Teradata system.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Request Source**.
 - c. Click **Add**.
The **Edit Request Source Criteria** dialog box appears.
 - d. From the **Source Type** list, select **Application**.
 - e. Under **Application**, enter the name of the application that queries the database: StoreRewards_Application
 - f. Click **Include**.
 - g. Click **OK**.
 - h. Click **Save**.
 4. Define the service-level goals.
 - a. Click the **Service Level Goals** tab.
 - b. Under **Default Settings**, select **Response Time Goal**.
 - c. In the **Response Time** field, enter 1 for a response time goal of maximum 1 second.
 - d. In the **Service Percent** field, enter 98.
This means the response time goal is met at least 98% of the time.
 - e. Click **Save**.
 5. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **SLG** ruleset, click  and select **Make Active**.
 6. [Optional] Review requests that missed the SLG criteria you defined.
 - a. Open the **Workload Monitor** portlet.
 - b. In the **Dynamic Pipes View**, hover over the Store Rewards workload in the workload pane.
An information balloon appears, showing the percentage of requests that missed the SLG.



- c. For more details, click the name of the workload.
The **Workload Details View** appears, providing details about, for example, request arrivals and completions.

How Can I Use Filters to Prevent Poorly Written Queries from Reducing My Performance?

A poorly written query can return too much data, making the results unusable. These types of queries often produce meaningless results, are very resource-intensive, and reduce the performance of your system.

Using filters, you can protect your system from poorly written queries by preventing these queries from even starting to execute. Rejecting bad queries also helps educate the users to write better and more specific queries, which can reduce the number of bad queries over time.

For information about how to write good queries, see the Teradata SQL documentation.

For example, you can prevent the following queries from executing:

- Queries that return too many rows
- Queries that take too long to run
- Queries that scan very large tables

The following is an overview of the process:

1. [Edit a ruleset.](#)
2. [Create a filter.](#)
3. [Set filter classifications.](#)
 - [Set filter classifications based on query characteristics.](#)
 - [Set filter classifications based on the target.](#)
4. [Review rejected queries.](#)

Scenario assumptions



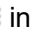
This scenario assumes you already have the following set up:

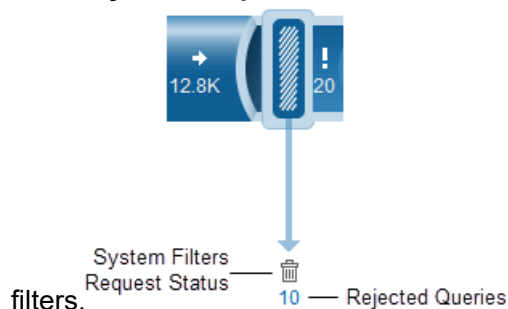
- Teradata Database version: 16.00
- Operating system: SLES 11
- System: PROD1
- Ruleset: BadQueries

Example: Filtering to Reject Queries with a Final Row Count of More than a Specified Number

Suppose you want to protect your system from bad queries by preventing the execution of queries that have an estimated final row count of over one billion rows.

1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.
 - b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **BadQueries**.




2. Create a filter.
 - a. From the ruleset toolbar, click **Filters**.
 - b. Next to **Filters**, click .
 - c. Enter the name of the filter, in this case `Final Rows Over 1000000000`.
 - d. Click **Save**.
3. Add classification criteria for the filter to apply to queries that return an estimated one billion rows or more.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Query Characteristics**.
 - c. Click **Add**.
The **Edit Query Characteristics Criteria** dialog box appears.
 - d. Select the **Estimated Final Row Count** check box.
 - e. Select the **Row Count \geq** check box, and enter `1000000000` in the field.
 - f. Select the **Join Type** check box.
 - g. From the **Include Only** list, select **Unconstrained Product Join**.
 - h. Click **OK**.
 - i. Click **Save**.
4. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **BadQueries** ruleset, click  and select **Make Active**.
5. [Optional] Review the queries rejected according to the criteria you specified.
 - a. Open the **Workload Monitor** portlet.
 - b. Click  in the portlet frame and select **Settings**.
 - c. Click the **Systems** tab.
 - d. From the **Systems** list, select **PROD1**.
 - e. Click **OK**.
 - f. In the **Dynamic Pipes View**, click on the number under the request status icon for system



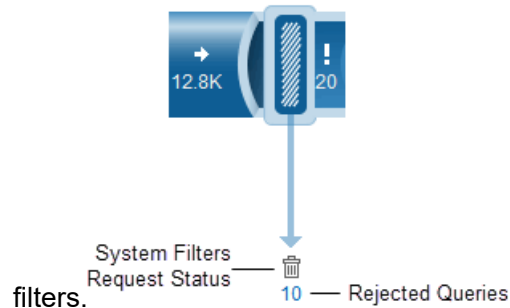
- g. Review the summary table to identify the queries rejected by this filter.
The Reason column includes For Rule Name 'Final Rows Over 1000000000'.

Example: Preventing Queries that Take Too Long to Run

Suppose you want to protect your system from bad queries by preventing queries that take five minutes or more to complete.

1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.
 - b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **BadQueries**.
2. Create a filter.
 - a. From the ruleset toolbar, click **Filters**.
 - b. Next to **Filters**, click .
 - c. Enter the name of the filter, in this case **Processing Time Over 5 min**.
 - d. Click **Save**.
3. Add classification criteria for the filter to apply to queries that are estimated to take five minutes or more to complete.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Query Characteristics**.
 - c. Click **Add**.
The **Edit Query Characteristics Criteria** dialog box appears.
 - d. Select the **Estimated Total Processing Time** check box.
 - e. Select the **Processing Time ≥** check box and enter **300** in the field.
 - f. Click **OK**.
 - g. Click **Save**.
4. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **BadQueries** ruleset, click  and select **Make Active**.
5. [Optional] Review the queries rejected according to the criteria you specified.
 - a. Open the **Workload Monitor** portlet.
 - b. Click  in the portlet frame and select **Settings**.
 - c. Click the **Systems** tab.
 - d. From the **Systems** list, select **PROD1**.
 - e. Click **OK**.

- f. In the **Dynamic Pipes View**, click on the number under the request status icon for system





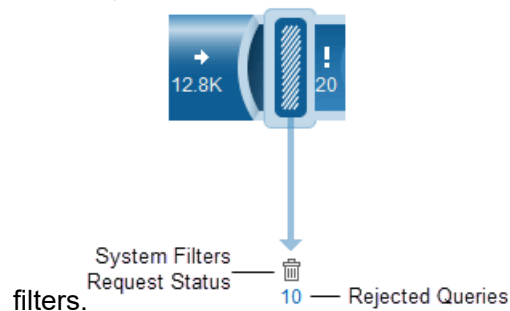
- g. Review the summary table to identify the queries rejected by this filter.
The Reason column includes For Rule Name 'Processing Time Over 5 min'.

Example: Preventing Full Table Scans on a Specific Table

Suppose you want to protect your system from bad queries by preventing users from running full table scans. For example, your company call center logs details about the calls in a call history table, which has grown very large. You want to prevent users from running full table scans on this particular table.

1. Select a ruleset to edit.
 - a. Open the **Workload Designer** portlet and the **Workload Designer** view.
 - b. From the **For System** list, select **PROD1**.
 - c. In the local **Working** section, click the ruleset name, **BadQueries**.
2. Create a filter.
 - a. From the ruleset toolbar, click **Filters**.
 - b. Next to **Filters**, click **+**.
 - c. Enter the name of the filter, in this case **Full Table Scan**.
 - d. Click **Save**.
3. Add classification criteria for the filter to apply to queries that attempt to run full table scans on the table **CallHistoryTable**.
 - a. Click the **Classification** tab.
 - b. From the **Add Classification Criteria** list, select **Target**.
 - c. Click **Add**.
The **Edit Target Criteria** dialog box appears.
 - d. From the **Target Type** list, select **Table**.
 - e. From the **Database** list, select the database that has the call history table: **History**.
 - f. Under **Table**, type the name of the table you want to prevent full table scans on: **CallHistoryTable**.
 - g. Click **Include**.
The table shows up in the **Selected** list under **Included** with the name **History.CallHistoryTable**.
 - h. From the **Selected** list, click **✎** next to **History.CallHistoryTable**.
The **Edit Criteria** dialog box appears.

- ☒ Full Table Scan
☒ Include
☐ Exclude
- i. Select the criteria **Full Table Scan**, click **Include**, and click **OK**.
 - j. Click **OK**.
 - k. Click **Save**.
 4. Activate the ruleset.
 - a. Go back to the **Workload Designer** view.
 - b. For the **BadQueries** ruleset, click  and select **Make Active**.
 5. [Optional] Review the queries rejected according to the criteria you specified.
 - a. Open the **Workload Monitor** portlet.
 - b. Click  in the portlet frame and select **Settings**.
 - c. Click the **Systems** tab.
 - d. From the **Systems** list, select **PROD1**.
 - e. Click **OK**.
 - f. In the **Dynamic Pipes View**, click on the number under the request status icon for system



- g. Review the summary table to identify the queries rejected by this filter.
The Reason column includes For Rule Name 'Full Table Scan'.

Workload Designer Overview

The **Workload Designer** portlet allows you to create rulesets for workload management. The features vary depending on whether you are connected to a Teradata system or an Aster system.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure displayed data is current in the portlet.

Workload Designer (Teradata System)

The available options are different if Analytics Database is running on SUSE Linux Enterprise Server version 10 (SLES 10) or lower, or SLES 11. After you select a system, you can create, edit, view, or perform other actions on rulesets on that system. The ruleset views provide controls that allow you to define, control, balance, and refine rules for managing workloads. Use the following features to manage workloads:

Rulesets

Collections of related filters, throttles, events, states, and workload rules.

States

Cause actions when a specific combination of planned environment and health condition occur.

Classification criteria

Determine which queries use which rules.

Sessions

Define limits for query, utility, and utility sessions.

Filters

Reject queries that meet the filter criteria.

Meters

Limit the number of queries that are allowed to run based on time interval (time unit).

Throttles

Limit the number of queries that are allowed to run concurrently.

Workloads

Group queries that have common characteristics or common requirements.

Exceptions

Cause actions when one or more specific events occur.

Note:

Support for TASM ARM is available for SQL Engine versions 17.10 and later.

TASM and IWM Systems in SLES 10 and SLES 11

Teradata Workload Management offers different strategies, depending on the platform and license you purchase:

TASM

Teradata Active System Management (TASM) is included in the enterprise tier license to perform full workload management in a Teradata system.

IWM

Teradata Integrated Workload Management (IWM) is a Teradata feature that performs basic workload management. Customers on base EDW platforms can optionally purchase TASM.

Note:

If you have any questions about your TASM license, contact your Teradata Account Manager.

TASM offers a full workload management feature set. IWM workload management features are more limited and have different options in SLES 10 and SLES 11, depending on the system you select.

States

In Teradata Database 15.00 and later, the IWM systems have options for planned environments, but not for unplanned events or health conditions. In Teradata Database 14.10 and earlier, the IWM systems have no options for custom states, health conditions, planned environments, or state-specific values for query sessions, filters, and throttles. The SLES 10 IWM system has no options for workloads.

Virtual Partitions

The IWM system has no options for virtual partitions.

Workload Management Methods

The IWM system has no options for the SLG Tiers workload management method. The IWM system uses only the tactical and timeshare management methods.

Exceptions

Neither SLES 10 nor SLES 11 IWM systems provide the ability to define exceptions at the ruleset level. SLES 11 IWM systems do provide the ability to define tactical exceptions for tactical workloads.

Workload Designer View

The **Workload Designer** view shows summary information about rulesets for a system or model system. The **Working**, **Ready**, and **Active** sections contain the following options:

Working

Names and descriptions of rulesets that are being edited. If a ruleset is locked by someone else, you have fewer options than if you are the ruleset owner. In **Working**, you can create and import rulesets. Rulesets in **Ready** can be copied to **Working** for editing. Rulesets in **Working** can also appear in **Ready** and **Active**.

Ready

Rulesets that have been saved to the Teradata system, but are not active. A ruleset must be in **Ready** before it is copied to **Active**. The **Active** ruleset cannot be deleted from **Ready**.

Active

Active ruleset on the Teradata system. The only option available in the options list, if you have permissions, is to deactivate the ruleset.

Ruleset background color indicates synchronization between the **Working** and **Ready** sections. Gray in both sections indicates that the same version of the ruleset is in both sections. Blue indicates that the ruleset has been modified in the **Working** section since it was last saved to the **Ready** section.

TASM ruleset migration is performed automatically during the Teradata system upgrade.

Selecting System Settings

In Teradata Database 14.10 and later, you can specify values or use the default system values to throttle queries consuming excessive memory. In system settings, you can create memory groups that are available for query characteristic classification. The settings apply to the whole system, not individual rulesets. Changes take effect after the next ruleset is activated.

1. From the **Workload Designer** view, click **System Settings**.
2. Select the **Specify custom values** check box and enter a value for one or more of the following threshold options to set the estimated memory use:

Option	Description
Very Large	Must be greater than the Large threshold
Large	Must be greater than the Increased threshold
Increased	Must be less than the Large threshold

3. Click **OK**.

Ruleset Options

The **Workload Designer** view shows individual rulesets and specific options for each ruleset based on the permissions assigned to your role and the section in which the ruleset is located.

Working

The **Working** section contains rulesets on the Viewpoint server that are editable. Any ruleset listed in **Working** can also be listed in **Ready** and **Active**.

Option	Description
Make Ready	Copies the ruleset to the Teradata server and to the Ready section.

Option	Description
Make Active	Makes the ruleset the active ruleset on the Teradata server. Copies the ruleset to the Ready and Active sections.
View/Edit	Opens the ruleset for viewing and editing.
Show All	Lists all ruleset attributes on one page.
Lock	Locks the ruleset so only the lock creator can edit the ruleset.
Unlock	Unlocks the ruleset so others can edit the ruleset.
Clone	Creates a copy of the ruleset. This option is useful if you want to use an existing ruleset as a base or template to create a ruleset.
Export	Exports the ruleset as an XML file. Use with the Import button to copy a ruleset from one system to another.
Delete	Removes the ruleset from the Working section.

Ready

The **Ready** section lists rulesets saved to the Teradata server. Authorized users can create a ruleset and add it to the **Ready** section. Then, from the **Ready** section, you can select **Copy to Working Rulesets** and edit the ruleset.

Option	Description
Activate	Makes the ruleset the active ruleset on the Teradata server.
Copy to Working Rulesets	Copies the ruleset to the Working section.
Delete	Removes the ruleset from the Teradata server.


Active

The **Active** section contains the active ruleset on the Teradata server. If you have permission, the only available option is to **Deactivate** the current ruleset. In SLES 11, the deactivate option is not available because there must always be an active ruleset.

Creating a Ruleset

A *ruleset* is a complete collection of related filters, throttles, events, states, and workload rules. You can create multiple rulesets, but only one ruleset is active on the Teradata server at a time. After creating a ruleset, you can use the toolbar buttons to specify settings, such as states, sessions, and workloads. New rulesets are automatically locked so only the owner can edit the ruleset.



1. From the **Workload Designer** view, select a system from the list.

2. Click .
3. Enter a ruleset name.
4. [Optional] Enter a description.
5. Click **Save**.
6. [Optional] Specify settings using the tabs in the **General** view, such as **Intervals** and **Bypass**.
7. [Optional] Click **States** and create a state matrix.
8. [Optional] Click **Sessions** and create any of the following:
 - **Query Sessions**
 - **Utility Limits**
 - **Utility Sessions**
9. [Optional] Click **Filters** and create a filter.
10. [Optional] Click **Meters** and create an Arrival Rate Meter.
11. [Optional] Click **Throttles** and create a system throttle.
12. [Optional] Click **Workloads** and create a workload.
13. [Optional] Click **Exceptions** and create an exception.

Editing a Ruleset

The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. Without permission to edit the ruleset, the option is **View** and the ruleset view is read-only. You can edit rulesets only in the **Working** section.

You can only edit rulesets that are unlocked or locked by you.


1. From the **Workload Designer** view, select a system from the list.
 2. In the local **Working** section, click the ruleset  and select **View/Edit**.
 3. Specify settings using the toolbar buttons and tabbed views.
 4. Click **Save** after making changes in each view.
 5. Click  to return to the **Workload Designer** view.
- The date and time the ruleset was modified and by whom appears under the ruleset name in the **Working** section.


Related Information:

[Ruleset Locks](#)

Cloning a Ruleset


You can clone a ruleset to make an exact copy of the ruleset, except for the name. Cloning is a convenient way to create a ruleset using the settings of an existing ruleset as a base.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Clone**.
3. Enter a name.

4. [Optional] Enter a description.
5. Click **Save**.
6. Specify additional settings using the toolbar buttons and tabbed views.
7. Click **Save**.
8. Click  to return to the **Workload Designer** view.
The cloned ruleset appears in the **Working** section.


Deleting a Ruleset

Deleting a ruleset removes the ruleset and all associated information. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section or the system **Ready** section, click the ruleset  and select **Delete**.
3. Select **Delete**.
The ruleset is deleted from the section.


Importing a Ruleset

The import and export options can be used to copy a ruleset from one Viewpoint system to another. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. Only rulesets exported from **Workload Designer** and a database of the same release can be imported.

1. From the **Workload Designer** view, select a system from the list.
2. Click .
3. Enter a name for the imported ruleset.
4. Click **Browse**.
5. Locate and select the saved ruleset file.

Note:


Exported ruleset files might be stored in the download area configured for your browser.

6. Click **Save**.
7. Click  to return to the **Workload Designer** view.
The imported ruleset appears in the **Working** section.

Exporting a Ruleset

The import and export options can be used to copy a ruleset from one Viewpoint system to another. Only rulesets exported from Workload Designer and a database of the same release can be imported.






1. From the **Workload Designer** view, select a system from the list.

2. In the local **Working** section, click the ruleset  and select **Export**.
3. Click **Save**.

The ruleset file is saved to your download area or the location you specify, depending on your browser settings.

Showing All Criteria in a Ruleset

You can display a read-only summary of all settings and state-specific values for a single ruleset.


1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Show All**.
3. [Optional] Do any of the following:
 - Click **Collapse All** to hide all sections of the view.
 - Click **Print** to print the full summary.
 - Click  to show or  to hide the main sections of the view.
 - Click  to show or  to hide individual sections of the view, or click the section name which toggles the show or hide.

Copying a Ruleset to the Teradata System

When you are finished editing a ruleset in the **Working** section, copy the ruleset to the **Ready** section on the Teradata system. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. After copying a ruleset to the **Ready** section, you can make the ruleset the active ruleset on the Teradata system.


Note:

Copying a ruleset to a Teradata system that is under heavy load could take longer than expected.

1. From the **Workload Designer** view, select the Teradata system containing the ruleset you want to copy.
2. In the **Working** section, click the ruleset  and select **Make Ready**.

Activating a Ruleset


Activating a ruleset copies the ruleset to the active state on the selected system. Only one ruleset is active on the system at a time. A ruleset activated from the local **Working** section is copied to the **Ready** section before being made active. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.

1. From the **Workload Designer** view, select a system from the list.
2. To make the ruleset active, do one of the following:
 - From a ruleset in the **Working** section, click the ruleset  and select **Make Active**.

- From a ruleset in the **Ready** section, click the ruleset ▼ and select **Activate**.
3. Click **Activate** to confirm that you want to activate this ruleset.

Deactivating a Ruleset



Deactivating a ruleset removes the ruleset from the active state on the selected system. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.

1. From the **Workload Designer** view, select a system from the list.
2. In the system **Active** section, click the ruleset  and select **Deactivate**.
In SLES 11, the deactivate option is not available because there must always be an active ruleset.




Ruleset Locks

An exclusive lock can be placed on a ruleset so that the ruleset cannot be edited, deleted, or otherwise modified except by the owner of the lock. A ruleset is automatically locked when it is created and each time changes to the ruleset are saved. Use the **Workload Designer** view to lock and unlock rulesets. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. The Teradata Viewpoint Administrator can also grant your role permission to unlock any ruleset.

The **Workload Designer** view displays the ruleset lock status.

Option	Description
	Locked
	Unlocked


The ruleset views display the ruleset lock status.

Option	Description
	Unlocked
	Locked by the current user
	Locked by another user

Locking or Unlocking a Ruleset

Locking a ruleset prevents others from modifying the ruleset. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. The Teradata Viewpoint

Administrator can also grant your role permission to unlock any ruleset. A ruleset that is locked may be unlocked by the lock owner, or by any user in a role with permissions to unlock any ruleset.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Lock** or **Unlock**.

General View

The **General** view on the ruleset toolbar displays general attributes of a ruleset. The following tabs in the **General** view allow you to change settings when creating, editing, cloning, or viewing a ruleset:

General

Ruleset name and description.

Bypass

Users, accounts, and profiles whose queries are allowed to bypass system filters and throttles.

Option	Description
Source Type	Type of query source. You can select one or more items from each available source type. The dbc and tdwm user names cannot be removed.
Filter	Limits your choices in the source list.
Items	Lists the names of all available query sources of the selected type.
Bypass	Selected sources, listed by source type. Queries from each of the listed query sources bypass system filters and throttles for this ruleset.

Limits/Reserves (TASM, Advanced Tier License for Teradata IntelliFlex, and Advanced Tier License for EDW)

In SLES 11, you can limit the resources available to the system per planned environment. Limits are defined as a percentage of CPU or I/O. You can also reserve AWTs for use by the system per planned environment.

Option	Description
CPU Limit	Percentage of CPU that the system can use.
I/O Limit	Percentage of I/O that the system can use.
Reserved AWTs	Maximum number of AWTs to reserve for tactical workloads.
Max AWTs for new work	Maximum number of AWTs for new tactical work.

Note:

If you are using Workload Management Capacity on Demand (WM COD) with a contractual agreement in place, a special form of enforcement may have been configured outside the **Workload Designer** portlet by Teradata Customer Services, using special hardware COD packages. For Advanced Tier license, the COD flag must be set to **yes** in the Teradata system. These special enforcement COD settings can only be set at a fixed percentage of 12.5% units. After they are set, these hardware-enforced limits appear in parentheses on the **Limits/Reserves** tab. The values selected by the administrator for WM COD CPU and I/O limits within **Workload Designer** must be lower than or equal to the hardware COD settings. **Workload Designer** will not allow WM COD settings that are higher than the enforcement COD settings put in place by CS. If a pre-existing ruleset has WM COD limits set higher than these enforcement COD limits, the ruleset can be activated, but the lower hardware COD limits will override the specified WM COD values within Viewpoint.

Reserves (IWM)

You can reserve AWTs for use by the system per planned environment.

Option	Description
Reserved AWTs	Maximum number of AWTs to reserve for tactical workloads.
Max AWTs for new work	Maximum number of AWTs for new tactical work.

Other

You can change the settings for intervals, blockers, activation, timeshare decay, throttle delay prevention, throttle delay queue order, increased job and resource limits, and available AWTs definition.

Intervals

Collection and reporting intervals.

Option	Description
Event Interval	How often event thresholds are checked. Note: If system or workload events have Simple or Averaging qualification times, then the event interval must be less than or equal to the lowest qualification time for any of the events.
Dashboard Interval	How often workload statistics are collected.
Logging Interval	How often workload and exception logs are written.

Option	Description
Exception Interval	How often exception thresholds are checked.
Flex Throttle Action Interval	(Teradata Database 16.0 and later for SLES 11 EDW systems) How often the availability of system resources is checked. Must be a multiple of the Event Interval.

Blocker

Settings for responding to throttled blockers. The log is located at DBC.TDWMEventLog. Selecting **Log** only logs the blocker. **Abort** or **Release** logs the action after the abort or release occurs.

Option	Description
Block Cycles	Number of intervals over which the query must be blocked before the specified Block Action is taken.
Block Action	Action to perform in response to query blocking.

Activation

Options available when the ruleset is activated. The feature is available when the option is selected. The Events and States feature is available by default and cannot be disabled.

Option	Description
Filters and Utility Sessions	Filters reject user queries. Utility Sessions limit the number of sessions a specific utility can use.
System Throttles and Session Control	System throttles delay or reject queries at the system level. Session Control makes it possible to create session parameters.
Workloads	Workloads group queries that share characteristics so that a set of workload management controls can be applied to the group. Workloads are always available in SLES 11.
Events and States	Events are occurrences that trigger a response from the system. States are intersections of health conditions and planned environments.

Timeshare Decay

In SLES 11, you can enable timeshare decay, which automatically decreases the resource access rate of long-running queries in a Timeshare workload over time.

Prevent Mid-Transaction Throttle Delays

For Teradata Database 15.0 and later, select this option to allow subsequent queries within a transaction to run if there is a lock higher than an Access lock on a database object, thus allowing requests to not be delayed by workload throttles.

Order the Throttle Delay Queue

In Teradata 15.10 SLES 11, you can choose the order in which queries are released from the throttle delay queue.

Option	Description
By time delayed	The longer a query has been delayed, the sooner it will be executed.
By workload priority	The higher the workload priority of a query, the sooner it will be executed.

Utility Limits

Teradata Database 15.10 and later supports increased MLOADX job limits and increased AWT resource limits. When the check box is selected, the limit for MLOADX jobs increases from 30 to 120 and the limit for AWT resources increases from 60 to 70.

Clearing the check box decreases the maximum allowable resource and job limits for MLOADX jobs. The system uses the lower limits for job limits and AWT resources regardless of the ruleset configuration in this circumstance. If you clear the check box, make sure the limits for MLOADX jobs are less than the maximum (**Sessions > Utility Limits > Limit Details > State Specific Settings**) and that resource limits are less than the maximum (**Throttles > AWT Resource Limits > Limit Details > State Specific Settings**).

Define Available AWTs as

If you are using Teradata Database 16.0 and later, you can designate the definition used in determining available AWTs.



Option	Description
AWTs available for the WorkNew (Work00) work type	AWTs available for the WorkNew (Work00) work type. The WorkNew (Work00) work type is limited to 50 AWTs by default. If AWTs that can support WorkNew message are already in use servicing WorkNew message types, there may still be

Option	Description
	AWTs in the unreserved pool that will not be considered available.
AWTs available in the unreserved pool for use by any work type	AWTs available in the unreserved pool for use by any work type. This is number of AWTs available in the unreserved pool able to be used by all work types, not limited to WorkNew work types.

Defining System-Level Bypass Settings

A *system-level bypass* is a collection of users, accounts, and profiles that are not filtered or throttled at the system level. The **dbc** and **tdwm** usernames are listed automatically and cannot be changed, renamed, or removed.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Bypass** tab is current.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **General**.
3. Click the **Bypass** tab.
4. Select a **Source Type** from the list to exclude it from system-level filters and throttles.
5. [Optional] Enter a filter string in the **Filter** box to limit your choices.
6. Select a source from the list or select multiple sources by using **Ctrl** or **Shift**.
7. Click  to add your selections to the bypass list.
8. [Optional] Repeat steps 4 through 7 to add sources to the bypass list.
9. [Optional] In the bypass list, click  to include it in system-level filters and throttles.
10. Click **Save**.

Enabling Timeshare Decay

In SLES 11, you can enable timeshare decay.

Timeshare decay automatically decreases the resource access rate of queries in a Timeshare workload over time. Enabling decay prevents problematic queries from monopolizing system resources.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **General**.
3. Click the **Decay** tab.
4. Click the **Enable Timeshare Decay** check box.
5. Click **Save**.

Setting Limits and Reserves

In SLES 11, you can set limits and reserves for each planned environment.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **General**.
3. Click the **Limits/Reserves** tab (TASM, Advanced Tier License for Teradata IntelliFlex, and Advanced Tier License for EDW) or **Reserves** tab (IWM).
4. Select a planned environment.
5. Define limits on **CPU Utilization** and **I/O** for the planned environment.
This feature is not available on IWM systems, Teradata IntelliBase, or Appliance.
6. Define tactical reserves for tactical workloads in the planned environment.
7. Click **Save**.

States View

The **States** view on the ruleset toolbar shows a matrix that allows you to create and organize states for a ruleset, so one ruleset can respond to a range of different system conditions.

State

A set of operating rules that can be used to control how your system allocates resources. A state is the intersection of a health condition and a planned environment.

Health Condition

The condition or health of the system. For example, system conditions include system performance and availability considerations, such as nodes down at system startup. Unplanned events such as system, workload, user-defined, and combination events activate the health conditions.

Planned Environment

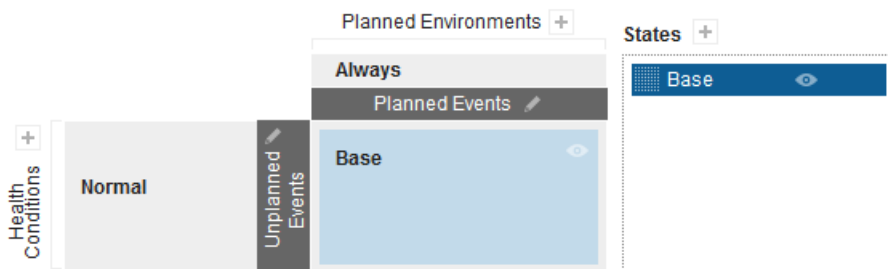
The kind of work the system is expected to perform, usually representing time periods or workload windows when critical applications, such as a load or month end, are running. For example, if you have many system users on weekdays, but run batch jobs on the weekends, allocate system resources differently during the week than you do on weekends by creating two planned environments: Weekdays and Weekends. Planned events such as period, user-defined, and combination events activate the planned environment.

In Teradata Database 15.00 and later, the IWM systems have options for planned environments, but not for unplanned events or health conditions. In Teradata Database 14.10 and earlier, the IWM systems have no options for custom states, health conditions, planned environments, or state-specific values for query sessions, filters, and throttles.

In the state matrix, create planned events, unplanned events, health conditions, planned environments, and corresponding states specific to your business situation. Update the state matrix at any time to reflect business and system requirements, or trigger priority changes and other system changes.

The **Normal** health condition, **Always** planned environment, and **Base** state are defaults. The defaults apply unless planned or unplanned events occur, triggering other configured states. The defaults cannot be deleted or moved within the state matrix. Like any state, the **Base** state can be used in multiple cells of the matrix.

Following is an example of the **States** view on a SLES 10 TASM system.



Any states you create use the default settings. You can set new values on the state-specific settings tabs in workloads, filters, throttles, query sessions, and utility limits to override the default settings.

Using only a few states in the state matrix reduces maintenance time. However, consider adding states to the matrix to manage the following situations:

- Consistent, peak workload hours or days where priority management must be strictly assigned and enforced
- Load or query times where priority tasks must finish within a specific time frame
- Conditions where resources must be managed in a different way, such as giving higher priority to critical work when system health is degraded

Editing a State Matrix

On TASM systems, the state matrix defines which state to use when a specific combination of planned environment and health condition exists. On IWM systems, the state matrix defines which state and operating rules will be used when a planned environment is in effect.




The default health condition, planned environment, and state give you the basic framework to allocate resources for workloads. However, you can create additional health conditions, planned environments, and states to fine-tune resource allocation.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
The state matrix appears. For a new ruleset, the default state matrix consists of one cell.
3. [Optional] If you are on a TASM system, do the following:
 - a. Define additional health conditions.

- b. Define one or more unplanned events.
An *unplanned event* is a system event, user-defined event, workload event, or a combination of these events.
- c. Drag an unplanned event to the appropriate health condition so the event triggers the health condition.
4. [Optional] Define additional planned environments.
5. [Optional] Define additional planned events.
A *planned event* is a period event, user-defined event, or a combination of these events.
6. [Optional] Drag a planned event to the appropriate planned environment so the event triggers the environment.
7. [Optional] Define additional states.
8. For each cell in the state matrix, drag and drop a state from the **States** list into the appropriate cell of the matrix.
When the defined combination occurs, the state is triggered.
9. Click **Save**.

Defining Health Conditions (TASM)



Health conditions define levels of system health and are used to reallocate system resources when an event degrades the system. When at least one unplanned event occurs, a health condition can be triggered. The default health condition is **Normal**, and it is used if no other health conditions are triggered. The **Normal** health condition always remains at the top. The lowest severity is listed at the top. The highest severity is at the bottom. If multiple unplanned events are active at the same time, the health condition with the highest severity is triggered.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Health Conditions**, click  .
A health condition is added with the default name **newCond**.
4. Next to the health condition, click  .
5. Enter a name for the health condition.
6. Enter a minimum duration, in seconds, for the health condition.
The event that triggers the change to the health condition must remain before the trigger threshold for the minimum duration you enter before the health condition goes into effect. Setting minimum duration prevents short incidents of an event from triggering a change in the health condition.
7. Click **OK**.
8. [Optional] Next to a health condition, click  to delete a health condition.
9. Click **Save**.

System Events (TASM)

A *system event* is an unplanned event, such as a down node. Incorporating system events into health conditions within the state matrix gives you greater control over the actions Analytics Database takes when unexpected events occur. To create an event that only sends out a notification, create the event, but do not assign it to any health condition. When the event occurs, the notification action you specified is triggered.

Defining System Events in the State Matrix (TASM)

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Unplanned Events**, click .
4. Next to **Available Events**, click .
5. Select **System Event**.
6. Enter a name.
7. [Optional] Enter a description.
8. Under **Configure Event Trigger**, select a **System Event Type** from the list.
Maximum and minimum limits must be positive integers.
9. [Optional] Select a qualification option to prevent very short incidents from triggering an event.
10. [Optional] Under **Configure Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

11. Click **OK**.
12. [Optional] Drag the event under a specific health condition.
When the event occurs, it triggers the health condition.
13. Click **Close**.
14. Click **Save**.

System Event Type Options (TASM)

The following are the available system event type options.

Event Types	Description
Node Down	Maximum allowed number of nodes down in a clique, defined as a percentage between 0 and 100. The default is 24%.
AMP Fatal	Maximum number of AMPs reported as fatal at system startup.

Event Types	Description
Available AWTs	Number of AWTs available on the specified number of AMPs as defined in the Other tab of the General view.
Gateway Fatal	Maximum number of gateways reported as fatal at system startup.
PE Fatal	Maximum number of PEs reported as fatal at system startup.
Flow Control	Maximum allowed number of AMPs in flow control.
CPU Utilization	Defines when the system CPU values are consistently outside defined CPU utilization threshold values.
CPU Skew	Maximum system-wide skew.
I/O Usage	(Teradata Database 16.0 and later) Identifies I/O bandwidth bottlenecks and assesses the scope of the bottleneck.

Note:

Setting qualification times can prevent very short incidents from triggering events.

Qualification Method Options	Description
Simple	Specifies how long an event threshold must be met before an event is triggered.
Immediate	Specifies that an event is triggered immediately after an event threshold is met.
Averaging	Specifies the following: <ul style="list-style-type: none"> Under Qualification Method, the time period during which the data sample is averaged. Under Qualification Time, how long the rolling average of the metric value must meet the event threshold before an event is triggered.

Workload Events (TASM)



A *workload event* is an unplanned event, such as a high number of active requests. Use workload events to transition to different health conditions when the events occur. For example, select **Active Requests** as the event type and set a threshold for the number of queries that can be active at one time. When the threshold is met or exceeded, the workload event triggers a change in the health condition.

To create an event that only sends out a notification, create the event but do not assign it to any health condition. When the event occurs, the notification action you specified is triggered.

Note:

When events have been defined for a workload, the workload cannot be deleted until the events have been manually deleted.

Defining Workload Events in the State Matrix (TASM)

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Unplanned Events**, click .
4. Next to **Available Events**, click .
5. Select **Workload Event**.
6. Enter a name.
7. [Optional] Enter a description.
8. Under **Configure Event Trigger**, do the following:
 - a. Select a **Workload**.
 - b. Select a **Workload Event Type** from the list.
 - c. Define a threshold for the workload event type.
Maximum and minimum limits must be positive integers.
 - d. [Optional] Select a qualification option to prevent very short incidents from triggering an event.
9. [Optional] Under **Configure Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

10. Click **OK**.
11. [Optional] Drag the event under a specific health condition.
When the event occurs, it triggers the health condition.
12. Click **Close**.
13. Click **Save**.

Workload Event Type Options (TASM)

You can define thresholds for workload event types. When a threshold is met or exceeded, the associated workload event occurs. Maximum and minimum limits must be positive integers. The following event types are available.

Event Types	Description
Active Requests	Defines maximum or minimum number of queries that can be active at one time. Active Requests are not available for utility workloads.



Event Types	Description
Arrivals	Defines maximum or minimum per-second arrival rate for queries. Arrivals are not available with utility workloads.
AWT Wait Time	Defines minimum time a step in a request can wait to acquire an AWT.
CPU Utilization	Defines maximum or minimum CPU usage for a query.
Delay Queue Depth	Defines minimum number of queries in the delay queue.
Delay Queue Time	Defines a minimum for the time a request can be in the delay queue.
SLG Response Time	Triggers an event based on the Response Time SLG set for the workload. Available only if a Response Time SLG is set for the workload.
SLG Throughput	Triggers an event based on the Throughput SLG set for the workload. Available only if a Throughput SLG is set for the workload.


Some workload event types allow you to set a qualification time. Setting qualification times can prevent very short incidents from triggering events. The qualification time is checked at the end of each interval.

Qualification Method Options	Description
Simple	Specifies how long an event threshold must be met before an event is triggered.
Immediate	Specifies that an event is triggered immediately after an event threshold is met.
Averaging	Specifies the following: <ul style="list-style-type: none"> Under Qualification Method, the time period during which the data sample is averaged. Under Qualification Time, how long the rolling average of the metric value must meet the event threshold before an event is triggered.

Defining Planned Environments



Planned environments are used to reallocate system resources during scheduled times. Planned environments are triggered when at least one associated planned event occurs. The default planned environment is **Always** and it cannot be deleted or moved. The order of precedence is from lowest to highest, reading from left to right. The planned environment with the highest precedence is activated if multiple planned events are active at the same time.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Planned Environments**, click .
A planned environment is added with the default name **newEnv**.
4. Next to the planned environment, click .

5. Enter a name.
6. Click outside the name.
7. If multiple planned environments exist, click a planned environment name and drag the name to the left or right in the list to change the order of precedence.
8. [Optional] Next to the environment name, click  to delete a planned environment.
You cannot delete the **Always** environment.
9. Click **Save**.

Defining Period Events in the State Matrix

A *period event* is a planned event occurring on specific days and times, such as month-end financial processing. Use periodic events to manage events that happen on a recurring basis. To create an event that only sends out a notification, create the event, but do not assign it to any planned environment. When the event occurs, the notification action you specified is triggered.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Planned Events**, click .
4. Next to **Available Events**, click .
5. Select **Period Event**.
6. Enter a name.
7. [Optional] Enter a description.
8. [Optional] Select **Create New Corresponding Planned Environment** to automatically create a planned environment that this event triggers.
9. Select **Day of Week** or **Day of Month**, and click a single day or multiple days.
10. [Optional] Select **Month of Year**, and click one or more months.
11. [Optional] Select **Start Time**, and enter a start and end time.
12. [Optional] Select **Wrap around midnight** to have a time frame spanning midnight for a period event.
13. [Optional] Under **Configure Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

14. Click **OK**.
15. [Optional] Drag the event under a planned environment.
16. Click **Close**.

17. Click **Save**.

Example: Wrap around Midnight

When creating period events, use the **Wrap around midnight** option to have a time frame spanning midnight.

If the **From** time of a period event is later than the **To** time, two time segments are available: midnight until the **To** time and the **From** time until midnight. When **Wrap around midnight** is not selected, the period event is in effect for segment 1 and segment 2 on each specified day. When **Wrap around midnight** is selected, the event is in effect for segment 2 on each specified day and for segment 1 on each day following the specified day.



For example, specify that a period event occurs on Mondays and Tuesdays with a **From** time of 17:00 and a **To** time of 08:00.

If Wrap around midnight is not selected:			
	Monday	Tuesday	Wednesday
Midnight—08:00 (time segment 1)	Yes	Yes	No
08:00—17:00	No	No	No
17:00—23:59 (time segment 2)	Yes	Yes	No

If Wrap around midnight is selected:			
	Monday	Tuesday	Wednesday
Midnight—08:00 (time segment 1)	No	Yes	Yes
08:00—17:00	No	No	No
17:00—23:59 (time segment 2)	Yes	Yes	No

Defining User-Defined Events in the State Matrix

A *user-defined event* is a SQL request that you define. User-defined events can be planned or unplanned. To create an event that only sends out a notification, create the event, but do not assign it to any planned environment or health condition. When the event occurs, the notification action you specified is triggered.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Unplanned Events** (with TASM license only) or **Planned Events**, click .
4. Next to **Available Events**, click .
5. Select **User Defined Event**.
6. Enter a name.

7. [Optional] Enter a description.
8. [Optional] If you are creating the user-defined event as a planned event, you can select **Create New Corresponding Planned Environment** to automatically create a planned environment that this event triggers.
9. Under **Activate/Deactivate Event**, copy the appropriate SQL request text.
10. Paste the text into an SQL script.
11. [Optional] Under **Configure Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.



12. Click **OK**.
13. [Optional] Do one of the following:

Event Type	Description
Unplanned	Drag the event to a health condition and click Close (with TASM license only).
Planned	Drag the event to a planned environment and click Close .

14. Click **Save**.

Defining Event Combinations in the State Matrix

An *event combination* is a mix of two or more different events, such as period, system, and user-defined events. Event combinations can be planned or unplanned. To create an event that only sends out a notification, create the event, but do not assign it to any planned environment or health condition. When the event occurs, the notification action you specified is triggered.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **Unplanned Events** (with TASM license only) or **Planned Events**, click .
4. Next to **Available Events**, click .
5. Select **Event Combination**.
6. Enter a name.
7. [Optional] Enter a description.
8. [Optional] If you are creating the combination event as a planned event, you can select **Create New Corresponding Planned Environment** to automatically create a planned environment that this event triggers.

9. Add available events and operators to build the formula representing the combination of events that must occur to trigger the event. For example, Node Down OR Batch Processing, or NOT PEI. When creating event combinations, avoid placing two operators or two events next to each other. When an event combination is valid, the background of the **Event Combination Formula** box is white. If a combination is invalid, the background is orange.
10. [Optional] Under **Configure Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.




11. Click **OK**.
12. [Optional] Do one of the following:

Event Type	Description
Unplanned	Drag the event to a health condition and click Close (with TASM license only).
Planned	Drag the event to a planned environment and click Close .

13. Click **Save**.

Defining States in the State Matrix

You can create a state to control how resources are allocated in the state matrix.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to **States** at the top of the states list, click .
A state is added to the list with the default name **NewState**.
4. Next to the state name, click  to edit the state name.
5. Enter a name.
6. Click outside the name.
7. [Optional] Next to the state name, click  to delete a state name.
8. Click **Save**.

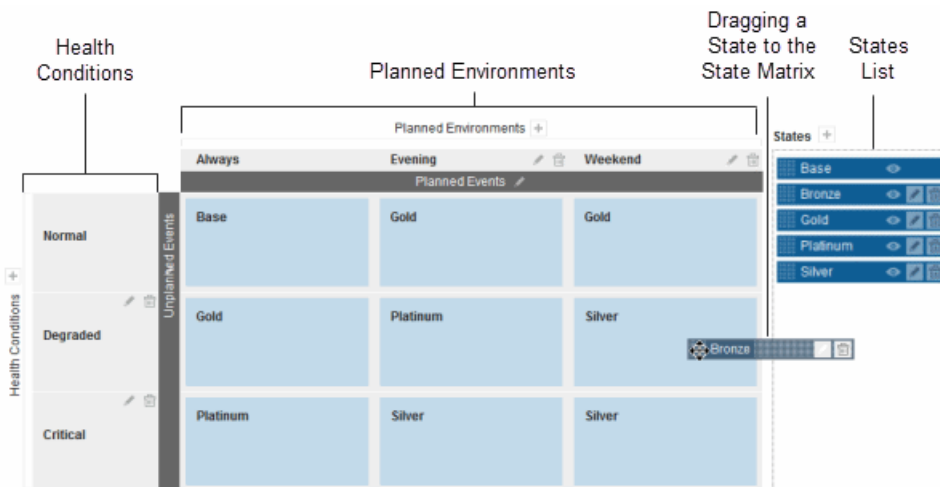
Mapping States in the State Matrix

After creating a state, you can map it in the state matrix to control how resources are allocated.

By default, the cell in the upper left corner of the state matrix is assigned the **Normal** health condition (with TASM license only), the **Always** environment, and the **Base** state. This cell cannot be changed. All other cells in the matrix must be associated with a single state. Any state, including the **Base** state, can be used in multiple cells.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Drag and drop a state name from the **States** list to a matrix cell to map a state in the state matrix.

Following is an example of a state matrix for a TASM system.



4. [Optional] Next to the state name in the **States** list, click to view details about a state.
5. Click **Save**.

Mapping Events in the State Matrix


If you have created events in the state matrix, you can combine the events with health conditions and planned environments for greater control of Analytics Database. Create the health conditions and planned environments that you need before mapping events in the state matrix.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. To map a planned event:
 - a. Next to **Planned Events**, click .
 - b. In **Available Events**, drag and drop an event from the list into one or more of the **Planned Environments** columns.
 - c. [Optional] Next to **Available Events**, click to create a planned event.
4. If you are on a TASM system, do the following to map an unplanned event:
 - a. Next to **Unplanned Events**, click .
 - b. In **Available Events**, drag an event from the list and drop it into one or more of the **Health Conditions** rows.
 - c. [Optional] Click next to **Available Events** to create an unplanned event.

5. Click **Close**.
6. Click **Save**.

Deleting States from the State List

If a state is not used in the state matrix, you can remove it from the **State** list.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **States**.
3. Next to a state name in the **State** list, click .
4. Click **Save**.

Classification Settings

Workload Designer provides a common classification process for workloads, filters, meters, throttles, query sessions, and utility sessions on the **Classification** tab. Classification determines which queries use which rules. Analytics Database detects classification criteria before executing queries. The goal in creating a useful classification scheme is to meet business goals and fine-tune control of Analytics Database. The following classification criteria types are available:

Request source

Where the query comes from, such as username, account name, account string, profile, application, client IP address, or client ID.

Target

What the query is acting on, such as database, table, view, macro, stored procedure, function, or method.

Query characteristics

What the query is composed of, such as statement type, AMP limits, any step time exceeds, step row count, final row count, estimated processing time, minimum step time, join type, full table scan, memory usage, or incremental planning.

Query band

What metadata is attached to the query, such as user location or application version.

Utility

Which utility submitted the query, such as FastLoad, FastExport, or Backup Utilities.

You can modify the classification settings in response to data monitoring, regular historical analysis, or changes. For example, classification groups may need to be created, or existing groups modified, if an application is added, two Teradata systems are consolidated, or service-level goals are missed.

A good approach to using classification is to first use request source to determine where the query is coming from. Often the account string is selected, but other options include username, account name, or client IP address. If you need a more detailed level of classification, establish where the query data is located, such as a database, table, or view. To narrow classification further, select query characteristics, query bands, or utilities. For utilities, use the check boxes to select the specific utilities you want to include.

For example, you could create a filter and add the request source classification to reject all queries from the Finance department when the Red state is in effect. To further refine the filter, add the query characteristic classification to filter out all requests from the Finance department that are estimated to run longer than 10 seconds. You include classification items in a filter to reject those items. The query characteristic setting is added to the request source setting already in place.

Generally *AND* logic is used when you add more than one criterion, so that the definition becomes more specific as more criteria are added. However, if you add both the Username and Profile request sources, the relationship between the two can be specified as *AND* or *OR*.

The screenshot shows the 'Workload: Monday' configuration page. It has tabs for 'General', 'Classification' (selected), 'Throttles', 'Service Level Goals', and 'Exceptions'. Below the tabs, a message states: 'Classification criteria determine the workload into which each query will be classified.' There is an 'Add Classification Criteria' section with a dropdown menu set to 'Request Source' and an 'Add' button. Below this, two criteria are listed, connected by an 'AND' operator. The first criterion is 'Request Source' with 'Include: (Usernames: PUBLIC)' and 'Exclude: (Usernames: Crashdumps)'. The second criterion is 'Query Characteristics' with 'Include Only: DDL OR DML statements' and 'Include only queries that use all AMPs'. At the bottom, there are 'Save' and 'Reset' buttons.

Request Source Classification Type

The request source classification type establishes which username, account name, account string, profile, application, client IP address, or client ID is making the request.

Consider the following when using request source to classify information:

- A source type can only be used once in each rule and can have multiple parameters.
- When source and target groups exist together, they are joined by an AND.
- A match string must be an exact match.

Adding Request Source Classification Type

You can classify filters, meters, throttles, workloads, utility sessions, and query sessions by request sources such as account name or client IP address.

1. Do one of the following:
 - From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.
 - From the **Sessions** tab, click **Query Sessions** or **Utility Sessions**.
2. Select an item or create one.
3. Click the **Classification** tab.
4. Select a request source criteria or add one.
5. Select a **Source Type** from the list.
 - Username
 - Account Name
 - Account String
 - Profile
 - Application
 - Client IP Address

Note:


[SQLE 17.20.00.00 and Higher] Select the **Use actual client IP address** check-box to use the actual client IP address instead of the network address.

- Client ID
6. Do at least one of the following:
 - Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.
 7. Click **OK**.

Target Classification Type

The target classification type specifies the query data location. You can classify filters, meters, throttles, or workloads by targets such as database, table, or stored procedure, and any subcriteria. For example, if you select a database as the target, you could add subcriteria so that it only applies if you are performing a full table scan.

Consider the following when using target to classify information:

- A target type can only be used once in each rule and can have multiple parameters, such as table or stored procedure. Multiple parameters are joined by AND.
- You can use targets parameter function and method.
- When source and target groups exist together in Analytics Database, they are joined by an AND.
- You can add subcriteria for each selected target parameter. Next to target items containing subcriteria,  appears. If you select multiple subcriteria, they must all be present for the classification setting to be used.
- A match string must be an exact match.
- In Teradata Database 15.0 and later, you can set the estimated minimum percent of table blocks accessed during the query.


Adding Target Classification Type

You can classify filters, meters, throttles, or workloads by targets such as database, table, view, macro, stored procedure, function, or method. If some objects are not included in the list for a database, they can be manually included as a match string.

For Teradata Database 15.0 and later, you can also select targets from external QueryGrid servers, such as Aster or Hadoop servers.

1. From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.
2. Select an item or create one.
3. Click the **Classification** tab.
4. Select **Target** as the classification criterion and click **Add**.
5. Select a **Target Type** from the list.
6. Do one of the following, depending on the target type:

Target Type	Action
Table, View, Macro, Stored Procedure, Function, or Method	Select the database that contains the target.
QueryGrid Server (for Teradata Database 15.0 and later)	Select the QueryGrid server type.


7. Do at least one of the following:
 - Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.
8. [Optional] For databases, tables, views, macros, and stored procedures, click  next to an item in the **Selected** list, and choose from the following:

Subcriteria	Description
Full Table Scan	Includes or excludes all-row scans
Join Type	Selects no join or a join type
Estimated Step Row Count \geq	Specifies the estimated minimum rows at each step
Estimated Step Row Count \leq	Specifies the estimated maximum rows at each step
Estimated Step Processing Time \geq	Specifies the estimated minimum processing time at each step
Estimated Step Processing Time \leq	(Teradata Database 15.0 and later) Specifies the estimated maximum processing time at each step
Estimated percent of table blocks accessed during the query \geq	(Table target databases and Teradata Database 15.0 and later) Specifies the estimated minimum percent of table blocks accessed during the query

If multiple subcriteria are added to a single item, all subcriteria conditions must be true in order for the query to be classified into the rule.

a. [Optional] In Teradata Database 15.0 and later, select a statement type (applies only to Table target databases).

b. Click **OK**.

Next to target items containing subcriteria,  appears.

9. Click **OK**.

Query Characteristic Classification Type

The query characteristic classification type describes a query by answering such questions as what does the query do and how long will the query run.

Consider the following when using query characteristics to classify information:

- A query characteristic type can only be used once in each rule. After it is used, it no longer appears in the list.
- The query characteristic and utility types are mutually exclusive. If you use one, the other option is not available.
- A query characteristic type can have multiple parameters. The **Statement Type** parameter is joined to any additional parameters with AND. Multiple parameters are joined by AND.
- Some characteristics have less than (or equal to) and greater than (or equal to) values that can be set independently.

Adding Query Characteristic Classification Type

You can classify filters, meters, throttles, or workloads by query characteristics.

1. From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.

2. Select an item or create one.
3. Click the **Classification** tab.
4. Select a query characteristics criteria or add one.
5. Choose from the following query characteristics criteria:

Option	Description
Statement Type	(Teradata Database 14.10 and later) Options are: DDL , DML , and SELECT . COLLECT STATISTICS . Any additional parameters are joined with AND.
AMP Limits	Includes or excludes queries that use all AMPs for filter rules.
Multi-Statement Request	Allows you to set the minimum number of SQL statements a request must have to be classified as a <i>multi-statement request</i> (a request with two or more SQL statements in a single request separated by semicolons). A request with that number of or more SQL statements is not considered equivalent to a single statement request. You can select the Exclude option to disallow a multi-statement request from getting into certain workloads.
Estimated Step Row Count	Specifies estimated minimum and maximum rows at each step.
Estimated Final Row Count	Specifies estimated minimum and maximum rows in the result set.
Estimated Step Processing Time	Specifies an estimated minimum processing time at each step. (Teradata Database 15.0 and later) Specifies an estimated maximum processing time at each step. (Teradata Database 15.0 and later) After entering a value, click the Tab key to see a conversion of the seconds into days, hours, minutes, seconds, and milliseconds (<i>d+hh:mm:ss:ms</i>).
Estimated Total Processing Time	Specifies estimated minimum and maximum total processing time. Longer or more complex queries have less accurate estimates. (Teradata Database 15.0 and later) After entering a value, click the Tab key to see a conversion of the seconds into days, hours, minutes, seconds, and milliseconds (<i>d+hh:mm:ss:ms</i>).
Join Type	Options in the Include Only list are: No Join , Any Join , Product Join , No Product Join , Unconstrained Product Join , or No Unconstrained Product Join .
Full Table Scan	Includes or excludes full table scans for all rows.
Memory Usage	(Teradata Database 14.10 and later) Specifies a minimum threshold for queries using excessive memory. Options are Increased , Large , or Very Large . These thresholds are defined in System settings in the Workload Designer view.
Incremental Planning	Includes or excludes queries that use incremental planning. Incremental Planning and Execution (IPE) executes single-row and scalar subqueries and uses the incremental results to optimize processing of the main query. For more information, see <i>Teradata® Database SQL Request and Transaction Processing</i> .

In Analytics Database, multiple parameters are joined by AND.

6. Click **OK**.

Query Band Classification Type

The query band classification type specifies the query band name and value pair data attached to a query.

Consider the following when using query band to classify information:

- A query band type can only be used once in each rule. After it is used, it no longer appears in the list.
- A query band type requires a name and one or more values.
- Multiple values can be selected for the same name.
- Multiple included query band key and value pairs are joined with AND.
- Multiple excluded query band key and value pairs are joined with OR.

Adding Query Band Classification Type

A query band contains name and value pairs that use Analytics Database predefined names or custom names to specify metadata, such as user location or application version.

1. From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.
2. Select an item or create one.
3. Click the **Classification** tab.
4. Select a query band criteria or add one.
5. Select a query band name from the list or add one.
6. Select a **Previously Used Value** or enter a **New Value**.
You must select a name and a value.
7. Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria.
8. Click **OK**.

Utility Classification Type

The utility classification type identifies which utility submitted the query.

Consider the following when using utility criteria to classify information:

- A utility type can only be used once in each rule. After it is used, it no longer appears in the list.
- The utility and query characteristic types are mutually exclusive. If you use one, the other option is not available.
- The utility type is not available for a workload that has **Arrivals** or **Active Request** as the associated event type, or for a workload that has an associated exception.

- Available utility types include top level utilities such as **FastLoad**, **MultiLoad**, **FastExport**, and **Backup Utilities**, as well as specific implementations of a utility such as **JDBC FastExport** or **Archive/Restore (ARC)**.


Adding Utility Classification Type

You can classify filters, meters, throttles, workloads, or sessions for utilities such as FastLoad or MultiLoad.

1. Do one of the following:
 - From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.
 - From the **Sessions** tab, click **Utility Limits** or **Utility Sessions**.
2. Select an item or create one.
3. Click the **Classification** tab.
4. Select a utility criteria or add one.
5. Select any combination of **FastLoad**, **FastExport**, **Backup Utilities**, and **MultiLoad** utilities. Select a specific implementation of a utility, such as **JDBC FastLoad** or **Archive/Restore (ARC)**.
6. Click **OK**.

Deleting a Classification Type

You can delete classification criteria from filters, meters, throttles, workloads, query sessions, utility sessions, or utility limits at any time.

1. Do one of the following:
 - From the ruleset toolbar, click **Filters**, **Meters**, **Throttles**, or **Workloads**.
 - From the **Sessions** tab, click **Query Sessions**, **Utility Limits**, or **Utility Sessions**.
2. Select an item.
3. Click the **Classification** tab.
4. To the far right of the classification criterion name, click .
5. Click **Save**.

Sessions View

The **Sessions** view on the ruleset toolbar allows you to specify session limit information for a ruleset. The following tabs in the **Sessions** view control ruleset session limits from different perspectives:

Query Sessions

Limits on the number of query sessions that can be logged on at one time. You can create, enable, clone, and delete query sessions.

Query Sessions by State

Limits on the number of query sessions for each state. The default session limit for a state is listed, along with each state you have created and its assigned, state-specific session limit.

Utility Limits

Limits on the number of jobs per utility that can run at the same time. You can create, enable, clone, and delete utility limits.

Utility Limits by State

Limits on the number of jobs per utility in each state. The default utility limit is listed, along with each state you have created and its assigned, state-specific utility limit. The system default utility limits are also displayed.

Utility Sessions


Limits on the number of sessions a specific utility can use at one time. You can create, enable, clone, and delete utility sessions.

Utility Sessions Evaluation Order

Precedence of utility session rules, from highest to lowest. Evaluation order determines the rule in which the utility job is placed if a utility job matches more than one utility session rule.

Creating a Query Session

A *query session* limits the number of sessions that can be logged on at one time. After the query session is created, additional controls in the **Query Sessions** tab allow you to clone, delete, enable, or disable the query session. View all created query sessions on the **Query Sessions by State** tab.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Query Sessions** tab.
4. Next to **Query Sessions**, click .
5. Enter a name.
6. [Optional] Enter a description.
7. Select a **Rule Type** from the list:
 - Select **Collective** if you want all users that meet the classification criteria treated as a group, with the group allowed a maximum number of queries.
 - Select **Individual** if you want to apply limits to each user individually.
 - Select **Member** if you want accounts or profiles that represent user groups used as the classification criteria for the rule. Limits are placed on individuals in the group and no limit is placed on the account or group.
8. Click **Save**.
9. Click the **Classification** tab.

10. Add classification criteria.
11. Click **OK**.
12. Click **Save**.
13. Click the **State Specific Settings** tab.
14. Click a state to set a session limit.
15. Click **OK**.
16. Click **Save**.

Setting Classification for Query Sessions

Add classification settings to existing query sessions or when creating a query session. Classification determines which sessions are associated with the query session.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Query Sessions** tab.
4. Select a query session or create one.
5. Click the **Classification** tab.
6. Click **Add Criteria**.
7. From the **Source Type** list, select one of the following:

Option	Description
Username	Teradata system username
Account Name	Teradata system account name
Account String	Teradata system account identification string
Profile	Teradata system profile name
Application	Application on the network client
Client IP Address Note: [SQLE 17.20.00.00 and Higher] Select the Use actual client IP address check-box to use the actual client IP address instead of the network address.	IP address of the network client
Client ID	Login name on the network client

8. Do at least one of the following:

- Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
- Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.

9. Click **OK**.

10. [Optional] To add additional classification criteria:


- Click **Add Criteria**.
- Select the **Source Type**.
- Select the items you want to include.
- Click **OK**.

If you set both Username and Profile as request sources, you can choose **AND** or **OR** logic. If you set any other request source type combination, the additional criteria are evaluated using **AND** logic.

Setting State-Specific Values for Query Sessions

Add state-specific settings to existing query sessions or when creating a query session.

You can override the default by specifying session limits on a per-state basis. For example, you might want to have session limits during high-traffic states and no session limits during low-traffic states. View all created query sessions on the **Query Sessions By State** tab.


- Edit or create a ruleset.
- From the ruleset toolbar, click **Sessions**.
- Click the **Query Sessions** tab.
- Select a query session or create one.
- Click the **State Specific Settings** tab.
- Next to a state, click .
- Select **Create State Specific Settings**.
- Select **Unlimited**, or enter a session limit in the box.
- Click **OK**.

Your selection is applied to each cell having that state, and overrides the setting specified in **Default Settings**.

- [Optional] Change the default setting by selecting **Unlimited**, or enter a session limit number in the box.
- Click **Save**.

Creating a Utility Limit

A *utility limit* determines the number and type of utility jobs that can be run at one time. After the utility limit is created, additional controls in the **Utility Limits** tab allow you to clone, delete, and enable or disable the utility limit.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Utility Limits** tab.
4. Next to **Utility Limits**, click .
5. Enter a name.
6. [Optional] Enter a description.
7. Click **Save**.
8. Click the **Classification** tab.
9. Select the utilities to which the limit should be applied.
When a utility limit is created, several utilities are selected by default. The default utilities can be cleared.
10. Click **Save**.
11. Click the **State Specific Settings** tab.
12. Define utility limits for specific states.
13. Click **Save**.

Setting Classification for Utility Limits


Add classification settings to existing utility limits or when creating a utility limit. Classification determines which sessions are associated with a utility limit.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Utility Limits** tab.
4. Select a utility limit or create one.
5. Click the **Classification** tab.
6. Do one of the following:
 - Select a utility type, such as **FastLoad** or **FastExport**.
 - Select multiple utility types.
 - Select a specific version, such as **TPT Load Operator**.
 - Select a combination of utility types and versions.
7. Click **Save**.

Setting State-Specific Job Limits for Utility Limits

Set state-specific job limits for existing utility limits or when creating a utility limit. You can override the default by setting job limits on a per-state basis. For example, you might want to raise the job limit during a low-traffic state, and lower the job limit during a high-traffic state.


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Utility Limits** tab.
4. Select a utility limit or create one.
5. Click the **State Specific Settings** tab.
6. [Optional] Do any of the following:

Option	Description
Set a default job limit for the utility limit in all states	<ol style="list-style-type: none"> a. Under Default Settings, select the Job Limit check box and enter a number. b. Select Delay or Reject.
Set a job limit for the utility limit in a specific state	<ol style="list-style-type: none"> a. Next to a state, click . b. Select Create State Specific Settings. c. Select the Job Limit check box and enter a number. d. Select Delay or Reject. e. Click OK.

7. Click **Save**.

Creating a Utility Session Rule

A utility session rule defines the number of sessions that will be concurrently logged on to each utility. After the utility session rule is created, additional controls in the **Utility Sessions** tab allow you to clone, delete, and enable or disable the utility session rule.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Sessions**.
3. Click the **Utility Sessions** tab.
4. Next to **Utility Sessions**, click .
5. Enter a name.
6. [Optional] Enter a description.
7. Do one of the following:

Option	Description
Session-Based Utilities	<ol style="list-style-type: none"> Select Session Based Utilities. Select the utilities to which this session rule applies. From the list, select the data size. In the Number of Sessions box, enter the number of sessions that will be concurrently logged on for each instance of the selected utilities. <p>Note: When setting the number of sessions for the CSP Save Dump utility, the number of sessions cannot be less than the number of nodes.</p>
Backup/Archive/Restore (BAR)	<ol style="list-style-type: none"> Select Backup/Archive/Restore. From the list, select the data size. In the Build Processes box, enter the number of build processes that will be built concurrently for each BAR job.

- Click **Save**.

Setting Classification for Utility Sessions

Add classification settings to existing utility sessions or when creating a utility session. Classification determines which sessions are associated with a utility session.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

- Edit or create a ruleset.
- From the ruleset toolbar, click **Sessions**.
- Click the **Utility Sessions** tab.
- Select a utility session or create one.
- Click the **Classification** tab.
- In the **Add Classification Criteria** list, select **Request Source** or **Query Band**.
- Click **Add**.
- Specify options based on the classification criteria you selected.

Option	Description
Request Source	<p>Select the source type and do one of the following:</p> <ul style="list-style-type: none"> Enter a match string and use the Include and Exclude buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters. Select items from the list and use the Include and Exclude buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.

Option	Description
Query Band	<ol style="list-style-type: none"> Select a predefined query band name from the list. Select a Previously Used Value or enter a New Value. Use the Include and Exclude buttons to create your classification criteria.

- Click **OK**.
- [Optional] To add additional classification criteria:
 - Select the criteria from the **Add Classification Criteria** list.
 - Click **Add**.
 - Specify the items you want to include or exclude.
 - Click **OK**.

For Teradata Database version 15.10 and later, if you set both Username and Profile as a request source combination, you can choose **AND** or **OR** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.

- Click **Save**.

Setting Evaluation Order for Utility Sessions

If a utility job matches more than one utility session rule, the evaluation order determines the rule in which the utility job is placed. The rule in the highest position on the **Utility Sessions Evaluation Order** tab is applied. You cannot change the order of the system rules located at the bottom of the list.

- Edit or create a ruleset.
- From the ruleset toolbar, click **Sessions**.
- Click the **Utility Sessions Evaluation Order** tab.
- Drag rules to reorder the list of utility rules.
- Click **Save**.

Meters View

Meters limit the number of queries that can run based on the time interval for all queries in the system. Available time intervals are per second, per minute, and per hour. The default time interval is set to per minute.

Use **Meters** to control the number and types of query requests that you want processed and use the state matrix to indicate under which conditions you want certain limits with time unit to apply. For example, you can create a meter to:

- Limit a specific user to run no more than ten queries per minute.
- Limit a specific department to run no more than four queries per hour.

The following tabs on the **Meters** view display all meters information for a ruleset from different perspectives.


Arrival Rate Meters

List of all the Meters, their rates, and classification criteria.

Meter Rates by State

List of all the states to which specific meters are applied.


Creating Meters

1. Edit or create a ruleset.
2. From the ruleset toolbar, select **Meters**
3. Next to **Arrival Rate Meter**, click 
4. Enter a name
5. [Optional] Enter a description
6. [Optional] Select **Warning only** to have a warning message for the **Meter** logged by the database. The queries still run when a warning is logged.
7. Select a **Rule Type** option.

Option	Description
Disable Manual Release or Abort	Select Disable Manual Release or Abort to prevent Teradata DBAs from aborting or releasing Metered queries in the queue.

8. Select **Save**.
9. Select the **Classification** criteria.
10. **Add** and **Save** classification criteria.
11. Select the **State Specific Settings** tab.
12. Under the **Default Settings**, do one of the following

Limit	Description
Unlimited	Under Limit , select Unlimited .
Limited	Under Limited, enter the limit, select the time unit from the drop-down box and then select Delay or Reject .

13. [Optional] To override the default settings for a state do the following:
 - a. Next to a state name, click 
 - b. Select **Create State Specific Settings**
 - c. If you do not want to limit the queries for this state, select **Unlimited** under **Limit**.
 - d. To limit the queries for this state, enter a limit number in the box, select the time unit from the drop-down, and select either **Delay** or **Reject**.

Cloning Meters


To keep most of the settings from a Meter, clone and change any settings that do not apply.

1. From the ruleset toolbar, select **Meters**.
2. Click ▼ and select **Clone** in the row of a **Meter**.
3. Enter a name.
4. [Optional] Enter a description.
5. [Optional] Select **Warning Only** to have a warning message for the **Meter** logged by the database. The queries still run when a warning is logged.
6. Select a **Rule Type** option.

Option	Description
Disable Manual Release or Abort	Select Disable Manual Release or Abort to prevent Teradata DBAs from aborting or releasing Metered queries in the queue.

7. Select **Save**.
8. Select the **Classification** criteria.
9. **Add** and **Save** classification criteria.
10. Select the **State Specific Settings** tab.
11. Under the **Default Settings**, do one of the following

Limit	Description
Unlimited	Under Limit , select Unlimited .
Limited	Under Limited, enter the limit, select the time unit from the drop-down box and then select Delay or Reject .

12. [Optional] To override the default settings for a state do the following:
 - a. Next to a state name, select 
 - b. Select **Create State Specific Settings**
 - c. If you do not want to limit the queries for this state, select **Unlimited** under **Limit**.
 - d. If you want to limit the queries for this state, enter a limit number in the box, select the time unit from the drop-down, and select either **Delay** or **Reject**.

Setting Classification for Meters

Meters limit the number of user queries that can run at the given time interval, by rejecting them or putting them in a deferred queue. Classification determines which queries are associated with a Meter. Analytics Database detects classification criteria before executing queries.

Enabling the **Dictionary** data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab are current. If the collector is not enabled, enter usernames manually by entering a match string and selecting **Include**.

1. Edit or Create a ruleset.
2. From the ruleset toolbar, select **Meters**
3. Select a meter or create one.
4. Select the **Classification** tab.
5. In the **Add Classification Criteria** list, select one of the following:

Classification Type	Description
Request Source	Where the query comes from
Target	Where the query is actioned on
Query Characteristics	What the query is composed of
Query Band	What metadata is attached to the query
Utility	Which utility submitted the query

6. Select **Add**.
7. Specify options based on the classification options selected.
8. For **Request source** or **Target**, do at least one of the following:
 - Enter a match string and use the Include and Exclude buttons to add the match string. A match string can contain a ? to match exactly one character or an * to match zero or more characters
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select and select **Ok**.
9. [Optional] To add additional classification criteria:
 - a. Select the criteria from the **Add Classification Criteria** list.
 - b. Select **Add**
 - c. Specify the items to include or exclude, and select **Ok**.

For Teradata Database version 15.10 or later, if both **Username** and **Profile** is set as a request source combination, choose **And** or **Or** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.


Setting State Specific Values for Meters

You can set state-specific limits for **Meters** to override the default limit for a state. For example, to raise the limit of a **Meter** during a low-traffic state, lower the limit during a high-traffic state.

1. Edit or Create a ruleset.
2. From the ruleset toolbar, select **Meters**.
3. Select a **Meter** or create one.

4. Click the **State Specific Settings** tab.
5. Under **Default Settings**, do one of the following:

Limit	Description
Unlimited	Under Limit , select Unlimited .
Limited	Under Limited , enter the limit, select the time unit from the drop-down box and then select Delay or Reject .

6. [Optional] To override the default settings for a state do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**
 - c. If you do not want to limit the queries for this state, select **Unlimited** under **Limit**.
 - d. To limit the queries for this state, enter a limit number in the box, select the time unit from the drop-down, and select either **Delay** or **Reject**.

Disabling Meters

Disabling a Meter removes limits from a system, while maintaining the configured settings. You can enable the Meter as needed. For example, you can disable a Meter for testing or to measure and trend normal throughput for a system without a hard limit.

1. From the ruleset toolbar, select **Meters**.
2. Under **Arrival Rate Meters**, clear the **Enabled** check box and select **Save**.

Deleting Meters

Deleting a Meter removes limits from a system. You must be the ruleset owner to delete Meters.

1. From the ruleset toolbar, select **Meters**.
2. Select **Delete** in the row of a Meter.
3. Select **Yes** to confirm that you want to delete this Meter and **Save** the changes.

Filters View

The **Filters** view on the ruleset toolbar allows you to create a filter for a ruleset, using classification that identifies which queries should be affected. A *filter* rejects a query before the query starts running. Following are examples of using filters:

- Create a filter that prohibits a specific user from running a query with an estimated processing time of longer than 15 minutes.
- Create a filter to limit all members of a specific department that runs large reports from accessing the database during peak work hours.

The following tabs on the **Filters** view display all filter information for a ruleset from different perspectives:

Filters


List of all the filters, their status, and classification criteria.

Enabled by State

List of all the states to which specific filters are applied.

Creating a Filter

A filter rejects user queries.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Filters**.
3. Next to **Filters**, click .
4. Enter a name.
5. [Optional] Enter a description.
6. [Optional] Select **Warning Only** to have a warning message for the filter logged by the database.
The queries still run when a warning is logged.
7. Click **Save**.
8. Click the **Classification** tab.
9. Add and save classification criteria.
Include classification items in a filter to reject those items. For example, create a filter and add classification criteria to reject all queries from the finance department.
10. Click **Save**.
11. Click the **State Specific Settings** tab.
12. Set state-specific values for the filter.
You can enable or disable filters on a per-state basis.
13. Click **OK**.
14. Click **Save**.

Setting Classification for Filters

Filters reject user queries. Classification determines which queries are associated with a filter. Analytics Database detects classification criteria before executing queries.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Filters**.
3. Select a filter or create one.
4. Click the **Classification** tab.

5. In the **Add Classification Criteria** list, select one of the following:

Classification Type	Description
Request Source	Where the query comes from
Target	What the query is acting on
Query Characteristics	What the query is composed of
Query Band	What metadata is attached to the query
Utility	Which utility submitted the query

6. Click **Add**.
7. Specify options based on the classification options you selected.
8. For **Request Source** or **Target**, do at least one of the following:
- Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.
9. Click **OK**.
10. [Optional] To add additional classification criteria:
- a. Select the criteria from the **Add Classification Criteria** list.
 - b. Click **Add**.
 - c. Specify the items you want to include or exclude.
 - d. Click **OK**.


For Teradata Database version 15.10 and later, if you set both Username and Profile as a request source combination, you can choose **AND** or **OR** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.

11. Click **Save**.

Setting State-Specific Values for Filters

Add state-specific settings to existing filters or when creating a filter.

You can override the default by enabling or disabling the filter on a per-state basis. For example, you may want to leave a filter enabled under all circumstances except when a specific state occurs.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Filters**.
3. Select a filter or create one.
4. Click the **State Specific Settings** tab.
5. Next to a state, click .
6. Select **Create State Specific Settings**.

7. Select **Filter is Enabled** or **Filter is Disabled**.
8. Click **OK**.
Your selection is applied to each cell having that state, and overrides the settings specified in **Default Settings**.
9. [Optional] Change the default setting by selecting **Filter is Enabled** or **Filter is Disabled**.
10. Click **Save**.

Throttles View

The **Throttles** view on the ruleset toolbar allows you to control the concurrent queries by state or other criteria. A *throttle* limits the number of user queries that can be active at the same time based on a specified number, the query request source, target, or other query characteristics you select. When the throttle limit is reached, you can specify that new queries are placed in a delay queue or rejected.

In general, use throttles to control the number and types of query requests that you want processed and use the state matrix to indicate under which conditions you want certain limits to apply. For example, you can create a throttle to:

- Limit a specific user to running no more than 10 queries at a time.
- Limit a specific department to no more than 4 simultaneous queries.
- Limit a Base state to run 12 concurrent queries but allow a Midnight state to run an unlimited number.
- Limit queries from a specific request source, such as a load utility, to a specific target, such as a database or a table.
- Limit queries based on statement type, step row count, join properties, and more.

The following tabs in the **Throttles** view display throttle information for a ruleset from different perspectives:

Throttles

List of all the throttles, their status, and classification criteria. You can click a throttle name to drill down to detailed information about the throttle or to change its properties. These are several types of throttles:

- **System Throttles** limit the number of concurrent queries running on a Teradata system.
- **Virtual Partition Throttles** limit the number of concurrent queries running on a virtual partition.
- **Workload Group Throttles** limit the total number of concurrent queries allowed for all workloads in a specified group. Only workloads with delay throttles defined in the **Workloads** view can be added to a workload group throttle. Workload group throttles are available in Teradata Database 14.10 and later.
- **Workloads with Throttles Defined** limit the number of concurrent queries for an individual workload.

Throttle Limits by State

List of all the states to which specific throttles are applied. The FLEX attribute is displayed for workloads that have flex throttles enabled for a state.

Resource Limits

List of all the resource limits, their status, and classification criteria. You can click a resource limit name to drill down to detailed information or to change its properties.


Resource Limits by State

Lists the system default AWT resource limits, and the maximum percentage of AWTs available for each limit rule in each state. The lowest applicable limit shown in either section is enforced.

System Throttles

System throttles limit the number of queries that can run concurrently for all queries in the system. The concurrency limit you define represents an aggregate number that applies to all queries in the system.

Creating System Throttles


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Next to **System Throttle**, click .
4. Enter a name.
5. [Optional] Enter a description.
6. Select a **Rule Type** option:

Option	Description
Collective	Creates one queue for all matched queries. Select Collective if you want all users that meet the classification criteria treated as a group, with the group allowed a maximum number of queries.
Individual	Creates one queue for each object. Select Individual if you want to apply limits to each user individually.
Member	Creates one queue for each user. Select Member if you want accounts or profiles that represent user groups used as the classification criteria for the rule. Limits are placed on individuals in the group and no limit is placed on the account or group.
Disable Manual Release or Abort	Select Disable Manual Release or Abort to prevent Teradata DBAs from aborting or releasing throttled queries in the queue.

7. Click **Save**.


8. Click the **Classification** tab.
9. Add and save classification criteria.
10. Click the **State Specific Settings** tab.
11. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Under Concurrency Limit , enter a concurrency limit number in the box, and select either Delay or Reject .

12. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. Click **OK**.
13. Click **Save**.

Cloning System Throttles

If you want to keep most of the settings from a system throttle, you can clone it and change any settings that do not apply.


1. From the ruleset toolbar, click **Throttles**.
2. Click  and select **Clone** in the row of a system throttle.
3. Enter a name.
4. [Optional] Enter a description.
5. Select a **Rule Type** option:

Option	Description
Collective	Creates one queue for all matched queries. Select Collective if you want all users that meet the classification criteria treated as a group, with the group allowed a maximum number of queries.
Individual	Creates one queue for each object. Select Individual if you want to apply limits to each user individually.
Member	Creates one queue for each user. Select Member if you want accounts or profiles that represent user groups used as the classification criteria for the rule. Limits are placed on individuals in the group and no limit is placed on the account or group.
Disable Manual	Select Disable Manual Release or Abort to prevent Teradata DBAs from aborting or releasing throttled queries in the queue.

Option	Description
Release or Abort	

6. Click **Save**.
7. Click the **Classification** tab.
8. Add and save classification criteria.
9. Click the **State Specific Settings** tab.
10. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Under Concurrency Limit , enter a concurrency limit number in the box, and select either Delay or Reject .

11. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. Click **OK**.
12. Click **Save**.

Virtual Partition Throttles

If you are using Teradata Database 15.0 and later on SLES 11, you can limit the number of queries that can run concurrently in a particular virtual partition.


Virtual partitions are used to divide a system for a higher level of control over resource allocation and workload distribution. You can create virtual partitions and assign workloads to the partitions. Theoretically, one partition could use all the AWTs, rendering the other partitions useless. By creating a virtual partition throttle, you can make sure each partition is guaranteed access to the AWTs.

As with other throttles, you can specify virtual partition throttle limits for different states. The concurrency limit you define represents an aggregate number that applies to all queries in the partition.


Related Information:

[SLES 11 Virtual Partitions \(TASM\)](#)

Creating Virtual Partition Throttles

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Next to **Virtual Partition Throttle**, click .
4. Enter a name.
5. [Optional] Enter a description.
6. Select the virtual partition that you want to assign to this workload.
A virtual partition can have only one throttle.
7. Select **Disable manual Release or Abort** to prevent Teradata DBAs from aborting or releasing throttled queries in the queue.
8. Click **Save**.
9. Click the **State Specific Settings** tab.
10. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Under Concurrency Limit , enter a concurrency limit number in the box, and select either Delay or Reject .

11. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. Click **OK**.
12. Click **Save**.

Workload Group Throttles



Workload group throttles, available in Teradata Database 14.10 and later, limit the number of queries that can run concurrently for all workloads grouped in the throttle. The concurrency limit you define represents an aggregate number that applies to all workloads that are members of the group throttle.

Consider the following when adding workloads to a group throttle:


- Only workloads with delay throttles defined can be members of a group throttle.
- A workload with a reject throttle defined cannot be part of a group throttle.
- A workload cannot be a member of more than one group throttle.
- If you are using a SLES 11 system, the workloads in a group throttle must all belong to the same virtual partition.

Creating Workload Group Throttles

Workload group throttles are available in Teradata Database 14.10 and later.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Next to **Workload Group Throttles**, click .
4. Enter a name.
5. [Optional] Enter a description.
6. Select the virtual partition to which the workloads belong (SLES 11 systems only).
7. In the **Available Workloads** box, select one or more workloads and click  to move them to the **Workloads in this Group** box.
8. Click **Save**.
9. Click the **State Specific Settings** tab.
10. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Select Delay and enter a concurrency limit number in the box.

11. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, select **Delay** and enter a concurrency limit number in the box.
 - e. Click **OK**.
12. Click **Save**.

Flex Throttles

Workload throttles limit the number of queries that can run concurrently for all queries in a workload to provide a means of controlling overutilization of system resources.

Note:

Flex throttles are available in Teradata Database 16.0 and later for SLES 11 systems for TASM.

In some cases, regular workload throttles can prevent full system utilization. Flex throttles can detect situations where current throttle limits are preventing full system utilization and release work from the delay queue to utilize those resources. When flex throttles are enabled for a workload, they can be enabled for all states or else set for each state defined for that workload.

When setting flex throttles, you define:

- The time interval at which the availability of system resources is checked (the flex throttle action interval)
- System availability, based on the definition defined for the ruleset
- The events that would trigger flex throttle actions
- Actions to take place when the triggering event occurs

When you enable the flex throttle feature you can choose to enable evaluation mode to see how flex throttles would affect your system. In evaluation mode no queries are actually released when the trigger event occurs. Instead, entries detailing which queries would have been released are posted in the TDWM event log.

Enabling and Configuring Flex Throttles

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. In the **Workloads with Throttles Defined** section, click **Flex Throttles**.
4. Click the **Enable Flex Throttles** check box.
5. [Optional] To use the flex throttles in evaluation mode, select the **Enable Evaluation Mode** check box.

In evaluation mode flex throttles have no actual effect on your system, but the hypothetical results are reported in the TDWM event log.

6. Under **Triggering Events**, designate the **Available AWT** conditions to trigger an event.

Option	Description
Number of AMPs with Available AWTs	Specify the minimum number of AMPs with available AWTs. The event will be triggered if this number or more are available.
Number of Available AWTs	Specify the minimum number of available AWTs. The event will be triggered if this number or more are available.

Option	Description
Qualification Method	Specify the number of minutes for which both conditions in first two rows must be met to trigger the event.

7. [Optional] To further restrict the event definition to occur only when there is a minimum level of **CPU Utilization**, select the check box and configure the following:

Option	Description
System CPU	Specify the minimum system CPU utilization percentage. The event will be triggered if the CPU utilization is less than or equal to this value.
Qualification Method	Specify the number of minutes for which the minimum system CPU must be maintained to trigger the event. The average value of this metric must exceed the threshold for this time period.

8. [Optional] To further restrict the event definition to occur only when there is an I/O bottleneck, select the **I/O Usage** check box and configure the following:

Option	Description
Bandwidth	Specify the percentage of total I/O bandwidth that must be in use to trigger the event.
Monitored LUNs	TASM determines which LUNs are most vulnerable to throughput issues. Specify the percentage of these to monitor.
Triggered LUNs	The percentage of monitored LUNs that must be at the used Bandwidth percentage to trigger the event.
Qualification Method	Specify the interval at which an average Bandwidth is evaluated.
Qualification Time	Specify how long the condition must persist to qualify as an event.

9. Under **Actions**, enter the **Number of Queries to Release** on the flex throttle action interval. The number of queries cannot be more than twice the number specified for **Number of Available AWTs**. The flex throttle action interval is defined in the **Other** tab of the **General** category on the ruleset toolbar.
10. [Optional] Select actions to occur when a flex throttle event starts or ends.

Action	Description
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

11. Click **Save**.

Workload Throttles


Workload throttles limit the number of queries that can run concurrently for all queries in a workload.

Creating Workload Throttles

Define a throttle for a workload when you create or edit a workload in the **Workloads** view. When you set a throttle for a workload, any query that is classified into this workload is subject to the throttle you create.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Edit or create a workload.
4. Click the **Throttles** tab.
5. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	<ol style="list-style-type: none"> a. Under Concurrency Limit, enter a concurrency limit number in the box, and select either Delay or Reject. b. [Optional] If you selected Delay, have a SLES 11 EDW system, and are using Teradata Database 16.0 and later, you can select Enable Flex Throttles to apply flex throttles to the workload. <p>Under Concurrency Limit, enter a concurrency limit number in the box, and select either Delay or Reject.</p>


6. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. [Optional] If you selected **Delay**, have a SLES 11 EDW system, and are using Teradata Database 16.0 and later, you can select **Enable Flex Throttles** to apply flex throttles to when the workload is in that state.
 - f. Click **OK**.
7. [Optional] If the workload is a member of a workload group throttle, view throttle limits for the group throttle and for each workload that is a member of the group.
8. Click **Save**.

Cloning Workload Throttles

If you want to keep most of the settings from a workload throttle, you can clone it and change any settings that do not apply.

1. From the ruleset toolbar, click **Workloads**.
2. Click ▾ and select **Clone** in the row of a workload.
3. Enter a name.
4. [Optional] Enter a description.
5. Click **Save**.
6. Click the **Throttles** tab.
7. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	<p>a. Under Concurrency Limit, enter a concurrency limit number in the box, and select either Delay or Reject.</p> <p>b. [Optional] If you selected Delay, have a SLES 11 EDW system, and are using Teradata Database 16.0 and later, you can select Enable Flex Throttles to apply flex throttles to the workload.</p> <p>Under Concurrency Limit, enter a concurrency limit number in the box, and select either Delay or Reject.</p>

8. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. [Optional] If you selected **Delay**, have a SLES 11 EDW system, and are using Teradata Database 16.0 and later, you can select **Enable Flex Throttles** to apply flex throttles to when the workload is in that state.
 - f. Click **OK**.
9. [Optional] If you have group throttles defined, select a different group throttle from the list to make this workload a member of a workload group throttle.
The group throttle of the original workload are copied to the clone rule.
10. Click **Save**.

Setting Classification for Throttles

Throttles limit the number of user queries that can run at the same time by rejecting them or putting them in a delay queue. Classification determines which queries are associated with a throttle. Analytics Database detects classification criteria before executing queries.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Select a throttle or create one.
4. Click the **Classification** tab.
5. In the **Add Classification Criteria** list, select one of the following:

Classification Type	Description
Request Source	Where the query comes from
Target	What the query is acting on
Query Characteristics	What the query is composed of
Query Band	What metadata is attached to the query
Utility	Which utility submitted the query

6. Click **Add**.
7. Specify options based on the classification options you selected.
8. For **Request Source** or **Target**, do at least one of the following:
 - Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.
9. Click **OK**.
10. [Optional] To add additional classification criteria:
 - a. Select the criteria from the **Add Classification Criteria** list.
 - b. Click **Add**.
 - c. Specify the items you want to include or exclude.
 - d. Click **OK**.

For Teradata Database version 15.10 and later, if you set both Username and Profile as a request source combination, you can choose **AND** or **OR** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.


11. Click **Save**.

Setting State-Specific Values for Throttles

You can set state-specific concurrency limits for system throttles, virtual partition throttles, or workload group throttles to override the default concurrency limit for a state. For example, you may want to raise the concurrency limit of a system throttle during a low-traffic state, and lower the concurrency limit during a high-traffic state.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Select a system throttle, virtual partition throttle, or a workload group throttle or create one.
4. Click the **State Specific Settings** tab.
5. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Under Concurrency Limit , enter a concurrency limit number in the box, and select either Delay or Reject .

6. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. Click **OK**.
7. Click **Save**.

Disabling Throttles

Disabling a throttle removes concurrency limits from a system or workload group, while maintaining the configured settings. You can enable the throttle as needed.

For example, you can disable a throttle for testing or to measure and trend normal throughput for a system without a hard limit.


A disabled workload can be a member of a workload group throttle. When you disable a group throttle, you disable the aggregate limit for all workload members, but not for individual workloads. Throttles defined for individual member workloads remain active until you remove them individually.

1. From the ruleset toolbar, click **Throttles**.
2. Under **System Throttles**, clear the **ENABLED** check box.
3. Under **Virtual Partition Throttles**, clear the **ENABLED** check box.

4. Under **Workload Group Throttles**, clear the **ENABLED** check box.
5. Click **Save**.

Deleting Throttles

Deleting a throttle removes concurrency limits from a system, virtual partition, or workload group. You must be the ruleset owner in order to delete throttles.



1. From the ruleset toolbar, click **Throttles**.
2. Click  and select **Delete** in the row of a throttle.
3. Click **Yes** to confirm that you want to delete this throttle.
4. Click **Save**.

Resource Limits

Resource limits, available for Teradata Database 15.10 and later, restrict utility jobs based on the availability of AMP worker task (AWT) resources. Teradata System imposes system default AWT resource limits (viewable on the **Resource Limits by State** tab), but you can restrict resource usage further by designating limits for the utility types and by the request source or query band you designate.

AWT resource limits are applied to new utility jobs. If a limit is exceeded, any new jobs are delayed or rejected, depending on the option you set.


Creating Resource Limits

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Click the **Resource Limits** tab.
4. Next to **AWT Resource Limits**, click .
5. Enter a name.
6. [Optional] Enter a description.
7. Select the utilities for which the limit should be applied.
8. Click **Save**.
9. Click the **Classification** tab.
10. [Optional] Add and save classification criteria.
11. Click the **State Specific Settings** tab.
12. Under **Default Settings**, enter a percent of AWTs in the box, and select either **Delay** or **Reject**.
13. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. Enter a percent of AWTs in the box, and select either **Delay** or **Reject**.

- d. Click **OK**.
- 14. Click **Save**.

Cloning Resource Limits

If you want to keep most of the settings from an AWT resource limit, you can clone it and change any settings that do not apply.

1. From the ruleset toolbar, click **Throttles**.
2. Click the **Resource Limits** tab.
3. Click ▾ and select **Clone** in the row of the AWT resource limit you want to clone.
4. Enter a name.
5. [Optional] Enter a description.
6. [Optional] Add and save classification criteria.
7. Click the **State Specific Settings** tab.
8. Under **Default Settings**, enter a percent of AWTs in the box, and select either **Delay** or **Reject**.
9. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. Enter a percent of AWTs in the box, and select either **Delay** or **Reject**.
 - d. Click **OK**.
10. Click **Save**.

Setting Classification for Resource Limits

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Click the **Resource Limits** tab.
4. Select a resource limit or create one.
5. Click the **Classification** tab.
6. In the **Add Classification Criteria** list, select **Request Source** or **Query Band**.
7. Click **Add**.
8. Specify options based on the classification criteria you selected.

Option	Description
Request Source	Select the source type and do one of the following:


Option	Description
	<ul style="list-style-type: none"> Enter a match string and use the Include and Exclude buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters. Select items from the list and use the Include and Exclude buttons to create your classification criteria.
Query Band	<ol style="list-style-type: none"> Select a predefined query band name from the list. Select a Previously Used Value or enter a New Value. Use the Include and Exclude buttons to create your classification criteria.

9. [Optional] To add additional classification criteria:
 - Select the criteria from the **Add Classification Criteria** list.
 - Click **Add**.
 - Specify the items you want to include or exclude.
 - Click **OK**.

For Teradata Database version 15.10 and later, if you set both Username and Profile as a request source combination, you can choose **AND** or **OR** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.

10. Click **OK**.
11. Click **Save**.

Setting State-Specific Values for Resource Limits

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Click the **Resource Limits** tab.
4. Select a resource limit throttle or create one.
5. Click the **State Specific Settings** tab.
6. Under **Default Settings**, enter a percent of AWTs in the box, and select either **Delay** or **Reject**.
7. [Optional] To override the default settings for a state, do the following:
 - Next to a state name, click .
 - Select **Create State Specific Settings**.
 - Enter a percent of AWTs in the box, and select either **Delay** or **Reject**.
 - Click **OK**.
8. Click **Save**.

Disabling Resource Limits

Disabling a resource limit throttle removes the throttle while maintaining the configured settings. You can enable the throttle as needed.

1. From the ruleset toolbar, click **Throttles**.
2. Click the **Resource Limits** tab.
3. Under **AWT Resource Limits**, clear the **ENABLED** check box next to the resource limit you want to disable.
4. Click **Save**.

Deleting Resource Limits

1. From the ruleset toolbar, click **Throttles**.
2. Click the **Resource Limits** tab.
3. Click ☐ and select **Delete** in the row of a throttle.
4. Click **Yes** to confirm that you want to delete this throttle.
5. Click **Save**.

Workloads View

A *workload* is a group of queries that share characteristics so that a set of workload management controls can be applied to the group. A workload has working values and defining characteristics that are evaluated during the classification phase of system management. High-quality workload management can improve response times and provide more consistent response times for critical work. The IWM system has no options for the SLG Tiers workload management method or virtual partitions.

The following requests are examples of workloads:

- Batch jobs further subdivided by region or organization for reporting
- Weekly or monthly reports that follow the calendar or a regular schedule
- Jobs that are always critical whenever they occur

The **Workloads** view on the ruleset toolbar lists workload names, method or type, throttle limits, group throttles, and enabled status. For each workload, you can specify one or more of the following:

Classification

Criteria that determine if a query is assigned to the workload

Throttles

Limits on the number of concurrent active queries that can run in a workload

Service Level Goals

Goals for workload query performance in terms of response time or throughput

Hold Query Responses

For Teradata Database 15.10 SLES 11, enforces a delay so query responses are not returned sooner than an amount of time you specify

Exceptions

Thresholds and actions to take if a query exceeds exception criteria while executing

The following are the default workloads.

Default Workload	Description
H-WD	Uses \$H classification criteria, which has been deprecated by SLES 11.
L-WD	Uses \$L classification criteria, which has been deprecated by SLES 11.
M-WD	Uses \$M classification criteria, which has been deprecated by SLES 11.
T-WD	Uses \$R classification criteria, which has been deprecated by SLES 11.
WD-Default	Contains queries not classified into a specific workload. This workload cannot be deleted or disabled.
WD-MapsMover	(Teradata Database 16.10 and later) Use in conjunction with the Maps Mover operations. Supports table relocations across maps. This workload can be deleted or disabled, but it cannot be cloned or edited. This workload can be moved to any timeshare access method.

Creating and Editing a Workload


You can group queries that share characteristics into a workload so that a set of workload management controls can be applied to the group. Different options are available in SLES 10 and SLES 11 systems:

SLES 10 TASM system

Uses an *enforcement priority* (EP) or a label to assign weights for the queries running in a workload.

SLES 11 system

Uses a *workload management method* or a label to assign access levels for the queries running in a workload.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Next to **Workloads**, click .
4. In the **General** tab, enter a name.
5. [Optional] Enter a description.
6. [Optional] On a SLES 11 TASM system, select a virtual partition to assign the workload.
If only one virtual partition is available, it is selected by default.
7. Under **Enforcement Priority** or **Workload Management Method**, select one of the following options for queries in this workload:

Operating System	Description
SLES 10 (with TASM license only)	<p>Select an enforcement priority for the type of queries you expect to run in this workload.</p> <ul style="list-style-type: none"> • Tactical queries are short, critical queries with defined service level goals. They are typically single or few-AMP queries or all-AMP queries that consume less than 1 CPU-second per node. • Priority queries are important and have higher priority than most other work. • Normal queries are the average priority work running on the system. • Background queries run for work that does not have a response-time requirement.
SLES 11 (with TASM license only)	<p>Select a workload management method.</p> <ul style="list-style-type: none"> • Tactical queries are short and high-priority. They have access to all available resources. • SLG Tier queries have access to the resources that remain after tactical resource demands are met. Multiple tiers can be prioritized so higher level tiers have higher resource allocation. SLG Tier workloads are assigned to the lowest available tier. • Timeshare queries have access to the resources that remain after tactical and SLG resource demands are met. Timeshare workloads are assigned to the Medium access level, but can be changed.
SLES 11 IWM	<p>Select a workload management method.</p> <ul style="list-style-type: none"> • Tactical queries are short and high-priority. They have access to all available resources. • Timeshare queries have access to the resources that remain after tactical resource demands are met. Timeshare workloads are assigned to the Medium access level, but can be changed.

8. Click **Save**.
9. [Optional] Click the **Classification** tab and add criteria to classify incoming queries into this workload.
10. [Optional] Click the **Throttles** tab.
 - a. Click on a state to set a throttle to limit concurrent queries classified into this workload.
If you have a SLES 11 EDW system and are using Teradata Database 16.0 and later, you can designate that flex throttles be applied. See [Enabling and Configuring Flex Throttles](#) for more information.
 - b. Select a group throttle from the list to make this workload a member of a workload group throttle.
11. [Optional] On a SLES 11 TASM system, click the **Service Level Goals** tab and set service level goals for this workload.
12. [Optional] Click the **Exceptions** tab and create an exception with criteria that trigger actions or notifications for this workload.
For example, if a query in this workload is taking too long, create an exception that moves the query to a different workload.
13. [Optional] In a SLES 11 tactical workload, click the **Tactical Exception** tab and modify the tactical exception.
14. Click **Save**.

Setting Classification for Workloads

Workloads group queries so that a set of workload management controls can be applied to the group. Classification determines which queries are associated with a workload. Analytics Database detects classification criteria before executing queries.

You cannot add additional classification options to any of the default workloads.

Enabling the Dictionary data collector in the **Monitored Systems** portlet makes sure the lists of usernames in the **Classification** tab is current. If the collector is not enabled, you can enter usernames manually by entering a match string and clicking **Include**.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select or create a workload.
4. Click the **Classification** tab.
5. In the **Add Classification Criteria** list, select one of the following:

Classification Type	Description
Request Source	Where the query comes from
Target	What the query is acting on
Query Characteristics	What the query is composed of
Query Band	What metadata is attached to the query
Utility	Which utility submitted the query

6. Click **Add**.
7. Specify options based on the classification options you selected.
8. For **Request Source** or **Target**, do at least one of the following:
 - Enter a match string and use the **Include** and **Exclude** buttons to add the match string. A match string can contain ? to match exactly one character or * to match zero or more characters.
 - Select items from the list and use the **Include** and **Exclude** buttons to create your classification criteria. If there are more than 500 items, use the filter box to locate the items you want to select.
9. Click **OK**.
10. [Optional] To add additional classification criteria:
 - a. Select the criteria from the **Add Classification Criteria** list.
 - b. Click **Add**.
 - c. Specify the items you want to include or exclude.
 - d. Click **OK**.

For Teradata Database version 15.10 and later, if you set both Username and Profile as a request source combination, you can choose **AND** or **OR** logic. If you set any other classification criteria or request source type combination, the additional criteria are evaluated using **AND** logic.


11. Click **Save**.

Setting State-Specific Throttles for Workloads

Define a throttle for a workload when you create or edit a workload in the **Workloads** view. When you set a throttle for a workload, any query that is classified into this workload is subject to the throttle you create. You can set state-specific concurrency limits for workload throttles to override the default concurrency limit for a state.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select or create a workload.
4. Click the **Throttles** tab.
5. Under **Default Settings**, do one of the following:

Concurrency Limit	Description
Unlimited	Under Concurrency Limit , select Unlimited .
Limited	Under Concurrency Limit , enter a concurrency limit number in the box, and select either Delay or Reject .

6. [Optional] To override the default settings for a state, do the following:
 - a. Next to a state name, click .
 - b. Select **Create State Specific Settings**.
 - c. If you do not want to limit the concurrent queries for this state, select **Unlimited** under **Concurrency Limit**.
 - d. If you want to limit the concurrent queries for this state, enter a concurrency limit number in the box, and select either **Delay** or **Reject**.
 - e. Click **OK**.
7. [Optional] If you have Teradata Database 14.10 with group throttles defined, select a group throttle from the list to change the group throttle membership.
 - This workload must have a delay throttle defined.
 - This workload cannot have a reject throttle defined.
 - This workload can belong to only one group throttle.

Throttle limits appear for the group throttle and for each workload that is a member of the group throttle.

8. Click **Save**.

Setting Service-Level Goals

Service-level goals (SLGs) help determine if workload management is meeting expectations. You can set throughput and response time service-level goals.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select or create a workload.
4. Click the **Service-Level Goals** tab and select one of the following options:

Option	Description
Response Time Goal	Enter a time in seconds, and set a service percentage. For example, if you want 90% of queries to finish within 4 seconds, set a response time of 4 and a service percent of 90. You can have different goals in different planned environments. For example, you can set a response time of 4 seconds for a daytime environment and 12 seconds for a nighttime environment. Metrics graphs show response times.
Throughput Goal	Enter a throughput number per hour. You can set different goals for different planned environments. For example, type 200 if you expect 200 or more queries to be processed per hour.
No Goal	Sets no service-level goal for this workload. This option may be appropriate, for example, for very low priority background workloads. Note: You cannot remove a service-level goal from a workload if there is an associated workload event. To remove the service-level goal, delete the associated event.

5. [Optional] Click a planned environment and set environment-specific goals to override the default settings.
6. Click **Save**.

SLES 11 Holding Query Responses

If you are using SLES 11, you can designate that a delay be enforced so query responses are not returned sooner than an amount of time you specify. This option enables you to provide consistency and set users' expectations about when they will receive query results.

Held query responses do not count against throttle limits.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select or create a workload.
4. Click the **Hold Query Responses** tab.
5. Select one of the following options under **Default Settings** for the *Always* planned environment.

Option	Description
Do not hold query responses	Return query responses as soon as they are available.
Hold responses	Prevent query responses from being returned before the specified number of seconds elapses.

6. [Optional] Click another planned environment and designate the appropriate query response settings.
7. Click **Save**.

Workload Exceptions

A workload exception is a definition that includes thresholds and actions to take when a threshold is exceeded by an executing query in a workload. You can add multiple criteria to an exception. Workload exceptions are triggered when all exception criteria are satisfied.

Some exception criteria require that you specify the qualification time, or an amount of time the criteria must take place to be recognized. Setting qualification time prevents very short incidents of criteria from being acknowledged.

Neither SLES 10 nor SLES 11 IWM systems provide the ability to define exceptions at the ruleset level. SLES 11 IWM systems do provide the ability to define tactical exceptions for tactical workloads.

Exception Criteria Settings



Exceptions result when threshold values for the designated criteria are exceeded. Disk and skew criteria also require that the condition persist for a specified amount of time before an exception is triggered.

You can define an exception by specifying the values (and qualification time, if applicable) for the following criteria.

Name	Description
Maximum Spool Rows	The maximum number of rows that a query step can spool. Enter the maximum allowable number.
Spool Usage	The amount of spool space, per step, that a query uses. Enter the maximum value and select B (bytes), KB (kilobytes), MB (megabytes), or GB (gigabytes).
Number of AMPs	The number of AMPs that participate in a query. Enter the maximum number.
Blocked Time	The length of time a query is blocked by another query. Enter the maximum time in the format hhh:mm:ss (hours:minutes:seconds).
Elapsed Time	The length of time a query has been running (the response time). Enter the maximum amount in the format hhh:mm:ss (hours:minutes:seconds). By default the amount of time you enter includes blocked and throttle delay time. (Teradata Database 15.10 and later) You can exclude them by selecting the Exclude blocked time and Exclude throttle delay time check boxes.

Name	Description
I/O Count	The number of disk I/Os performed for the query. Enter the maximum number.
CPU Time (sum over all nodes)	The amount of CPU processing time consumed by a query, summed over all the nodes used. Enter the maximum time in CPU-seconds.
CPU Skew Percent	The percentage difference in CPU consumption between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum percentage and the qualification time in CPU-seconds.
I/O Skew Percent	The percentage difference in disk I/O counts between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum percentage and the qualification time in CPU-seconds.
CPU Skew Difference	The difference in CPU consumption between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum amount of time and the qualification time in CPU-seconds.
I/O Skew Difference	The difference in disk I/O counts between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum value and the qualification time in CPU-seconds.
CPU Disk Ratio	The ratio of CPU consumption to disk I/O during the last exception interval. Enter the maximum ratio and the qualification time in CPU-seconds.
I/O Physical Bytes	The amount of I/O space, summed over all nodes, that a query consumes. Enter the maximum I/O physical bytes, select KB (kilobytes), MB (megabytes), or GB (gigabytes), and enter the qualification time in CPU-seconds.

Defining a Workload Exception

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select or create a workload.
4. Click the **Exceptions** tab.
5. Next to **Exceptions**, click .
6. Enter a name.
7. [Optional] Enter a description.
8. Under **Criteria**, select an option from the list, click , and specify the additional required information.
You can select and add multiple criteria. If you do so, all criteria must be satisfied for the exception to be triggered.
9. For qualified criteria marked with *, enter a value in the **Qualification Time** box.
10. Under **Actions**, select one of the following actions to be performed when the specified criteria are met:

Action	Description
Notification Only	Sends notification and takes no other action.
Abort	Stops the query.
Abort Selects Only	Stops the query only if it is a SELECT.
Change Workload to	Changes the query to the workload you select from the list. This action is unavailable for Blocked Time or Elapsed Time .

11. [Optional] Under **Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Notify Database User	Sends an email alert to database user when the query stops. This option is available if you have selected only Abort/Abort Selects from actions.
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

12. Click **OK**.

SLES 11 Adjusting Tactical Workload Exceptions

In SLES 11, when you create a tactical workload, a tactical exception is automatically created to prevent a high-consuming query from depriving tactical workloads of system resources. Tactical exceptions cannot be deleted, but you can modify them if your environment has different needs.

Tactical exceptions occur when both threshold criteria are exceeded for either Tactical CPU Time or Tactical I/O Physical Byte. Tactical exceptions trigger an action that moves a high-consuming query to another workload and can trigger other actions, such as issuing an alert or running a program. Tactical exceptions apply to the tactical workloads they are associated with, across all environments.

- From the ruleset toolbar, click **Workloads**.
- Select or create a tactical workload.
- Click the **Tactical Exception** tab.
- Under **Criteria**, enter thresholds for **Tactical CPU Time** and **Tactical I/O Physical Byte**.
- Under **Action**, select a workload.
The query is moved to this workload when a query exceeds both **Tactical CPU Time** and **Tactical I/O Physical Byte** thresholds.
- [Optional] Under **Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Notify Database User	Sends an email alert to database user when the query stops. This option is available if you have selected only Abort/Abort Selects from actions.
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

7. Click **Save**.

SLG Summary

The **SLG Summary** tab displays a summary of service level goals that have been set for each workload and planned environment.

Service level goals are set for workload query performance from the **Service Level Goals** tab in a workload. You can set service level goals for all planned environments or set different goals for some planned environments.

- **Response Time Goal** combines the query response time with the service percent achieved.
- **Throughput Goal** is set for the queries that are processed in an hour.

Evaluation Order

Evaluation order determines the placement of queries into workloads and utility jobs into rules. Setting evaluation order is useful when you have created several workloads or utility sessions.

For workloads, if classification criteria can sort the same query into different workloads, evaluation order determines the workload in which the query is placed. For example, account finance users are classified into one workload and users in a specific geographical location are classified into a different workload. If an individual is a member of both groups, the workload highest on the evaluation order list is the workload in which the individual is sorted.

For utility sessions, if a utility job matches more than one utility session rule, evaluation order determines the rule in which the utility job is placed.

Setting Evaluation Order for Workloads

If classification criteria can sort the same query into different workloads, evaluation order determines the workload in which the query is placed. The workload in the highest position on the Evaluation Order tab is applied.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Evaluation Order** tab.

4. In the list of workloads, do one of the following to order the list and place higher priority workloads at the top:
 - Drag workload names to reorder.
 - Click in a number box, highlight the existing number, and enter a new number.
5. Click **Save**.

SLES 10 Workload Management

Setting Enforcement Priority for Workloads

On a SLES 10 TASM system, an *enforcement priority* (EP) is a label given to assign weights for the queries running in a workload.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Select a workload or create one.
4. Under **Enforcement Priority**, select an option for the type of queries you expect to run in this workload:
 - **Tactical** queries are short, critical queries with defined service level goals. They are typically single or few-AMP queries or all-AMP queries that consume less than 1 CPU-second per node.
 - **Priority** queries are important and have higher priority than most other work.
 - **Normal** queries are the average priority work running on the system.
 - **Background** queries run for work that does not have a response-time requirement.
5. Click **Save**.

Allocation Groups

In SLES 10, the priority scheduler uses enforcement priorities and allocation groups to manage resources. An *allocation group* (AG) is a group of workloads that share an enforcement priority (EP). The AG manages access to CPU resources based on the assigned weight from the EP. An EP is a label given to assign weights for the queries running in a workload. Each AG takes the EP of the first workload associated with it and after that, only workloads with the same EP can be associated to that AG. As a result, each AG is a collection of one or more workloads that share the same EP. Multiple AGs can use the same EP. You can group and manage AGs in resource partitions. The default allocation groups and any groups you create are assigned to resource partitions.

In SLES 10, Analytics Database uses default allocation groups L (Low), M (Medium), H (High), R (Rush) and the following EPs:

- **Tactical** queries are short, critical queries with defined service level goals. They are typically single or few-AMP queries or all-AMP queries that consume less than 1 CPU-second per node.
- **Priority** queries are important and have higher priority than most other work.



- **Normal** queries are the average priority work running on the system.
- **Background** queries run for work that does not have a response-time requirement.

You can change the default mapping by using the diagram on the **Workload Mapping** tab, where you can reassign workloads to other allocation groups and reassign allocation groups to other resource partitions. Allocation groups that contain workloads cannot be deleted.

Carefully consider the allocation groups you drag in and out of the tactical resource partition because this can have a significant effect on the amount of CPU available to the queries running in the allocation group. The tactical resource partition normally has a much higher weight than other resource partitions. If you move an allocation group containing long, resource-intensive queries into the tactical resource partition, those queries can consume a large amount of the available CPU, impacting queries running in other resource partitions. If you move a tactical allocation group out of the tactical resource partition, the queries running in the tactical allocation group are assigned a lower priority and may not meet their response-time goals.

Creating an Allocation Group






Create allocation groups using the **Workload Mapping** tab.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Workload Mapping** tab.
4. Next to the name of the resource partition, such as **Tactical** or **Standard**, to which you want to add an allocation group, and click .
5. Enter a name.
6. Click outside the name.
7. [Optional] To move the allocation group to a different resource partition, hover over the name of the group, and drag the group to a different column.
8. [Optional] To change the name of the allocation group, click  next to the name of the group.
9. Click **Save**.

Mapping Allocation Groups to Workloads

If you have created workloads, you can view the **Workload Mapping** tab to see how the workloads map to allocation groups and resource partitions. Allocation groups and resource partitions determine how much CPU is available to the queries that are running in the listed workloads. The default allocation groups are named L (Low), M (Medium), H (High), R (Rush) and represent levels of priority. These groups are used by the system for console utilities and cannot be deleted.


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Workload Mapping** tab.
4. Do any of the following:

Option	Description
Create allocation groups	<ul style="list-style-type: none"> Next to a resource partition name, such as Tactical or Standard, and click .
Create a resource partition	<ul style="list-style-type: none"> Next to Resource Partitions, click .
Change resource partition for an AG	<ul style="list-style-type: none"> Drag an allocation group to a different resource partition.
Change AG for a workload	<ul style="list-style-type: none"> Drag a workload to a different allocation group of same priority in any resource partition.
Rename a resource partition	<ul style="list-style-type: none"> To the right of a resource partition name, click .
Delete an empty allocation group	<ul style="list-style-type: none"> To the right of an allocation group name, click .
Delete an empty resource partition	<ul style="list-style-type: none"> To the right of a resource partition name, click .

- Click **Save**.

Deleting an Allocation Group

You cannot delete the L (Low), M (Medium), H (High), R (Rush) allocation groups or allocation groups that contain workloads.

- Edit or create a ruleset.
- From the ruleset toolbar, click **Workloads**.
- Click the **Workload Mapping** tab.
- To the right of an allocation group name, click .
- Click **Save**.




Adjusting Priority Distribution for Workloads

On the **Priority Distribution** tab, each cell in the diagram is an allocation group. The size of each cell is proportional to the relative weight and can be adjusted to change the percent of CPU available to each AG. Each column in the diagram is a resource partition. The width of each column is proportional to the relative weight and can be adjusted to change the percent of CPU available to each resource partition.




If you modify the base state on the **Priority Distribution** tab, the changes are applied to other states provided that no change has been made to any of the other states. After any of the states are changed, the only values that are applied are the changes to system values.

- Edit or create a ruleset.
- From the ruleset toolbar, click **Workloads**.
- Click the **Priority Distribution** tab.
- From the **State** list, select a state.

5. For resource partitions, do any of the following:

Option	Description
View partition details	<ul style="list-style-type: none"> • Hover over any resource partition name.
Edit partition details	<ul style="list-style-type: none"> • Click a resource partition name, and click . • Enter a system percentage allocation and, optionally, a CPU percentage limit.
Reallocate space between resource partitions	<ul style="list-style-type: none"> • Click a resource partition name, and drag the column resize bar.
Lock or unlock a column	<ul style="list-style-type: none"> • Click a resource partition name and click  to lock this column in the diagram, while other cells can be dynamically resized around this column. • Click a resource partition name and click  to unlock this column in the diagram so the allocation changes dynamically with other columns around it.

6. For allocation groups, do any of the following:

Option	Description
View AG details	<ul style="list-style-type: none"> • Hover over any allocation group.
Edit AG details	<ul style="list-style-type: none"> • Click in a cell and click . • Enter a system percentage allocation and, optionally, a CPU percentage limit.
Reallocate space in a resource partition	<ul style="list-style-type: none"> • Click in a cell and drag the cell resize bar.
Lock or unlock a cell	<ul style="list-style-type: none"> • Click in a cell and click  to lock this cell in the diagram, while other cells can be dynamically resized around this cell. • Click in a cell and click  to unlock this cell in the diagram so the allocation changes dynamically with other cells around it.

7. [Optional] Under **System Values**, do any of the following:
- Select **CPU Limit** and enter a CPU limit to set the maximum percentage of system CPU that Analytics Database can use.
 - Select **Reserved AWTs** and enter a number to set the maximum number of AMP worker tasks that are reserved for tactical workloads.
 - Select **Max Expedited AWTs** and enter a number to set the maximum number of AMP worker tasks that are created on each node.
8. Click **Save**.

SLES 11 Workload Management

Workload Distribution

In SLES 11, you can configure a virtual partition by allocating resources to workloads in the partition. Resources can be allocated differently for each planned environment. The resource consumption may differ from the allocation you set, based on activity on the system. Resources are allocated to management methods as follows:

Tactical

Tactical workloads have access to all available resources. A tactical exception is automatically created when you create a tactical workload in SLES 11.

SLG Tier

SLG Tier workloads have access to remaining resources after tactical demands are met. If multiple SLG Tiers exist, resources are allocated first to Tier 1, with workloads in each tier consuming resources and passing remaining resources down to the next tier.

A workload that is moved into a fully allocated tier receives minimal tier resources (5%). The resources are taken from the workload with the highest allocation in the tier.

Timeshare

Timeshare workloads have access to remaining resources after tactical and SLG resource demands are met.

- **Top** queries are given an access rate eight times the Low access level.
- **High** queries are given an access rate four times the Low access level.
- **Medium** queries are given an access rate two times the Low access level.
- **Low** queries are given the lowest access rate for the Timeshare method.

Understanding SLG Tier and Timeshare Workload Information

System Workload Report (SLES 11 only)

Supplemental SLG tier and timeshare workload information is displayed in the **System Workload Report**, accessed from the **Workload Distribution** tab. This report shows resource distribution percentages for the SLG tier and timeshare workloads. Reported values are based upon no activity in tactical workloads.

The following columns display:

Workload

Displays workloads by name, except for tactical workloads.

Virtual Partition

Virtual partition to which each workload belongs.

Tier

Tier to which each workload belongs. The following naming convention for timeshare workloads enables better sorting results:

- TS1 = Top timeshare
- TS2 = High timeshare
- TS3 = Medium timeshare
- TS4 = Low timeshare

% of Tier

Percent of resources, out of 100% of the tier resources, that is available to each workload.

% of System

Percent of resources, out of 100% of total system resources, that is available to each workload. This percentage changes with the addition of or change to WM COD.

% of System

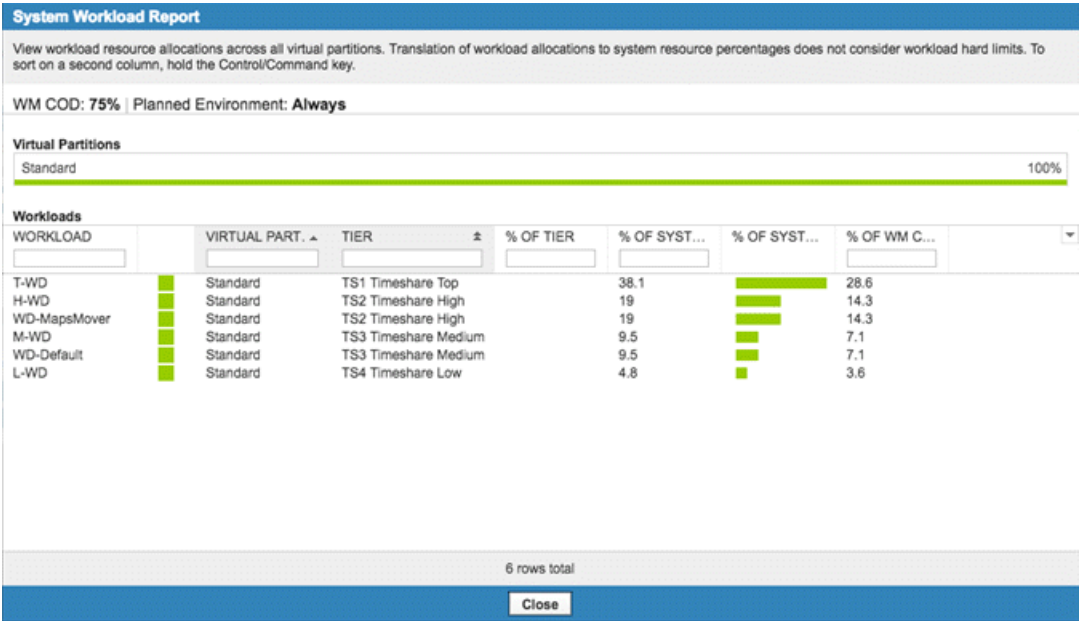
Bar charts representing the value in the **% of System** column. The scale of the bars is relative to the tier with the highest percentage. The highest percentage bar uses the full width of the column.

% of WM COD

Percent of resources within the amount that WM COD makes available. This setting does not change when WM COD is introduced or undergoes change.

Note:

If WM COD is set to 100%, WM COD percentage information will not appear in the **System Workload Report**.



Filters and Table Actions

Filters

Shows only rows that match your filter criteria. A second click reverses the sort order. Sort on two columns consecutively by sorting the first column and then using **Ctrl+Click** once in the column header of the second column. Using **Ctrl+Click** one time sets the secondary sort order. Primary sort order is indicated by a single arrow, and secondary sort order is indicated by a double arrow.

Upon opening the report, a primary ascending sort order is placed, by default, on the **Virtual Partition** column and a secondary ascending sort order on the **Tier** column. The secondary sort order on the **Tier** column is removed when the primary or secondary sort column has been modified. After modification, you see a single arrow to indicate the secondary sort order has been removed.

Table Actions

- **Clear Filters** removes any content in the filter boxes.
- **Export** creates a .csv file containing all available data.

For more information, see [Summary Table Controls](#).

Workload Distribution View Balloons

An information balloon appears when you hover over a workload or unallocated resources in the **Workload Distribution** view.

Workload balloons show information about the resources available to the workload:

Percent of tier

Percent of resources, out of 100% of the tier resources, that is available to that workload

Percent of system

Percent of resources, out of 100% of total system resources, that is available to that workload. This percentage changes with the addition of or change to WM COD

Percent of WM COD

Percent of resources within the amount that WM COD makes available. This setting does not change when WM COD is introduced or undergoes change

Hard limit

Maximum percent of the system resources allowed for this workload so that it cannot exceed its allocation and deprive lower tiers and workloads of system resources

SLG expedite

Specifies whether **SLG expedite** is enabled for a workload. This appears in the balloon only when the feature is enabled

Balloons for unallocated resources show information about the resources available to lower tiers and Timeshare workloads:

Percent of tier

Percent of resources, out of 100% of tier resources not allocated to workloads, that is available to workloads in lower tiers or timeshare. Ten percent of tier resources is always reserved to prevent SLG Tier workloads from depriving Timeshare workloads of all resources

Percent of system


Percent of system resources that are not allocated within the tier, that is available to workloads in lower tiers or timeshare

Percent of WM COD

Percent of resources within the amount that WM COD makes available and that is not allocated within the tier, that is available to workloads in lower tiers or timeshare

Distributing Workload Resources

In SLES 11, you can allocate resources to workloads for each planned environment.




1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Workload Distribution** tab.
4. [Optional] If multiple planned environments or virtual partitions exist, click  to select the area you want to manage.



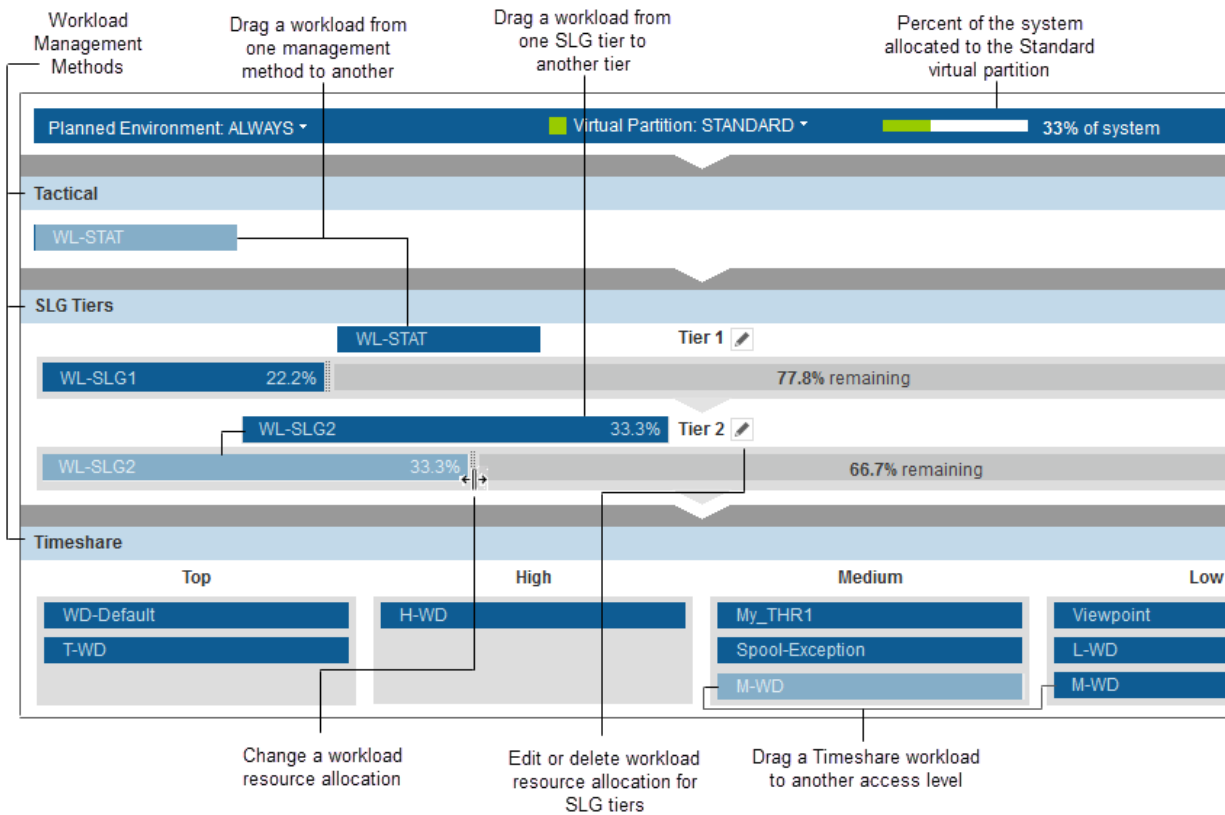
5. [Optional] Drag a workload from one management method to another to change resource allocation across planned environments.

Direction	Results
Tactical to SLG Tier	<ul style="list-style-type: none"> You can select the SLG Tier in the current planned environment. The workload is assigned to the lowest SLG Tier in all other planned environments. Minimum allocation percentage (5%) is assigned to the workload. The allocation comes from unallocated resources remaining for the tier, or from the workload with the highest allocation in the tier. Tactical exception is deleted for the workload.
Tactical to Timeshare	<ul style="list-style-type: none"> You can select the access level for the workload in the current environment. The access level is changed to medium in all other planned environments. Tactical exception is deleted for the workload.
SLG Tier to Tactical	<ul style="list-style-type: none"> Tactical exception is created for both Tactical CPU Time and Tactical I/O Physical Byte thresholds in the workload. Hard limits are removed for a workload if they were enabled.
SLG Tier to Timeshare	<ul style="list-style-type: none"> You can select the access level for the workload in the current environment. The access level is changed to medium in all other planned environments. Hard limits are removed for a workload if they were enabled.
Timeshare to Tactical	Tactical exception is created for both Tactical CPU Time and Tactical I/O Physical Byte thresholds in the workload.
Timeshare to SLG Tier	<ul style="list-style-type: none"> You can select the SLG Tier in the current planned environment. The workload is assigned to the lowest SLG Tier in all other planned environments. Minimum allocation percentage (5%) is assigned to the workload. The allocation comes from unallocated resources remaining for the tier, or from the workload with the highest allocation in the tier.

6. [Optional] Under **SLG Tier**, do any of the following:

Option	Description
Add a tier	Next to SLG Tiers , click  .
Change tier for a workload	<ul style="list-style-type: none"> If multiple tiers exist, drag workloads from one tier to another. (Teradata Database 14.10 and later) If you drag a workload from Tier 1 to a lower tier, SLG expedite is removed if it was enabled.
Change resource allocation for a workload	Drag the border of a workload horizontally.
Edit workload resource allocation for the tier (Teradata Database 14.10 and later)	<ul style="list-style-type: none"> Next to a tier name, click . Enable SLG expedite for Tier 1 workloads. Enable Hard limits for workloads in all tiers. <p>Note: It is recommended that hard limits be set no lower than 3% if only a single workload has a limit or no lower than 5% if multiple workloads have limits.</p>
Delete the tier	<ul style="list-style-type: none"> Next to the lowest tier name, click . <p>Tier 1 can never be deleted. Only tiers that do not contain workloads can be deleted. Only the lowest tier can be deleted. If the lowest tier does not contain workloads, the tier is automatically deleted when you save the tab settings.</p>

7. [Optional] Under **Timeshare**, drag a workload from one access level to another to assign the priority of resource allocation for queries in the workload.
8. Click **Save**.




Allocating Workload Resources in SLG Tiers

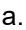
In SLES 11, you can modify resources allocated to workloads using the SLG Tiers workload management method.


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.

Note:

If you are creating a ruleset, at least one of the workloads in the ruleset must be set to use the SLG tier workload management method.

3. Click the **Workload Distribution** tab.
4. In the **SLG Tiers** section, click  next to the tier name.
The **Workload Resource Allocation** dialog box appears. The editable options in this dialog box are the following:

Option	Description
Tier allocation	Enter an allocation percent for each workload.
[Optional] Hard limits	<ol style="list-style-type: none"> a. Click  next to the workload. b. Select the Enable Hard Limit check box.

Option	Description
(Teradata Database 14.10 and later)	<p>The field is already populated with the % of System value, but this can be modified.</p> <p>c. Enter a number for the hard limit.</p> <p>Note: It is recommended that hard limits be set no lower than 3% if only a single workload has a limit or no lower than 5% if multiple workloads have limits.</p>
[Optional] SLG expedite (Teradata Database 14.10 and later)	<p>a. Click  next to the Tier 1 workload.</p> <p>b. Select the Enable SLG Expedite check box.</p>

5. Click **OK**.
6. Click **Save**.

SLES 11 Virtual Partitions (TASM)

In SLES 11, the priority scheduler uses virtual partitions to manage resources. The IWM system has no options for virtual partitions. A virtual partition is used to divide a system for a higher level of control over resource allocation and workload distribution. You can create virtual partitions and assign workloads to the partitions. A workload can be assigned to only one partition at a time. A maximum of 10 virtual partitions can exist.

The following tabs on the Workloads view are closely related and display partition information for a ruleset from different perspectives:

Virtual Partitions

List of the partitions on a system and the workloads that are assigned to each partition. You can add and remove partitions and reassign workloads to other partitions.

Partition Resources

List of the resource allocation for each virtual partition in each planned environment. You can specify CPU and I/O fixed limits.




Workload Distribution

View of the resource distribution to workloads in each virtual partition in each planned environment. In SLG Tiers you can specify system hard limits, and for Tier 1 you can enable SLG expedite.

Assigning a workload to a different partition does not change the workload management method. Assigning an SLG Tier workload to a different partition does not change the workload tier assignment if both partitions have the same tiers. If the target partition does not have the same tiers as the source partition, then the workload is assigned to the lowest SLG tier in the target partition. For example, if a



workload is moved from SLGTier3 in Virtual Partition A to Virtual Partition B, the workload is assigned to Tier3 in Virtual Partition B. If Tier3 does not exist in Virtual Partition B, then the workload is assigned to the lowest SLG Tier in Virtual Partition B.

An icon to the right of each workload represents the workload tier type.

Icon	Tier Type
	Tactical
	SLG Tier
	Timeshare

SLES 11 Adding Virtual Partitions

In SLES 11, you can create virtual partitions and use them to manage workload resource allocations. The IWM system has no options for virtual partitions.


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Virtual Partitions** tab.
4. Next to **Virtual Partitions**, click .
5. Enter a name.
6. Click outside the name.
7. Drag a workload to a virtual partition to assign the workload to the virtual partition.
A workload can exist in only one virtual partition.
8. Click **Save**.
9. [Optional] To delete a virtual partition, do the following:
You must move workloads to other virtual partitions before you can delete a partition.
 - a. Next to the virtual partition name, click .
 - b. Click **Save**.


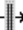
Managing Partition Resources


In SLES 11, you can allocate resources to a virtual partition in a planned environment. The IWM system has no options for virtual partitions.

You can change the resource allocation between multiple partitions and set fixed limits to prevent the partition from exceeding its allocation. If there are multiple planned environments, you can allocate a different level of resources for each planned environment. Resource consumption may differ from the allocation, based on activity on the system.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.

- Click the **Partition Resources** tab.
- Drag the border of a partition  in either direction to adjust the resource allocation for the partition.

Planned Environments		Virtual Partition Resource Allocation	
Always		Standard 35% 	Finance 40%

- To the right of a planned environment, click  to edit the resource allocation for the planned environment.
 - Set a percentage of resource allocation for each virtual partition in this planned environment. Total resource allocation must equal 100%.
 - In Teradata Database 14.10, select the **Enforce as a Fixed Limit** check box for any virtual partition in this planned environment. Fixed limits restrict CPU and I/O in this virtual partition in this planned environment.
 - Click **OK**.
- [Optional] Adjust the resource allocation for additional planned environments.
- Click **Save**.

Console Utilities

The **Console Utilities** tab maps Vantage console utilities to the workload in which they will run.

Console Utility to Workload Mapping

Maps console utilities to the default workloads in which to run.

Performance Group to Workload Mapping

Maps performance groups to the workloads in which to run. Many console utilities have the option to set the priority of the utility using the statement "PRIORITY=Performance Group Name". The **Performance Group to Workload Mapping** section of the **Console Utilities** tab maps the Performance Group Name used in the PRIORITY statement to a workload. The utility will then run in the mapped workload when the PRIORITY statement with the Performance Group Name is used.

Mapping Console Utilities to Workloads

Map console utilities to the default workloads in which to run.

- Edit or create a ruleset.
- From the ruleset toolbar, click **Workloads**.
- Click the **Console Utilities** tab.
- Under **Console Utility to Workload Mapping**, select a workload name for each console utility. The lists contain the workloads you have created and a predefined workload called WD-Default. WD-Default is used when no other workload is specified.

5. Click **Save**.




Mapping Performance Groups to Workloads

Map the default performance group names (H, L, M, R) to workloads.

The performance groups named L, M, H, and R stand for Low, Medium, High, and Rush. These groups represent levels of priority and are used by the system for console utilities.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Console Utilities** tab.
4. Under **Performance Group to Workload Mapping**, select a workload from the list for each performance group.
The lists contain the workloads you have created and a predefined workload called WD-Default. WD-Default is used when no other workload is specified.
5. Click **Save**.

Adding and Mapping a Performance Group Name

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Console Utilities** tab.
4. Next to **Performance Group to Workload Mapping**, click .
5. To the right of the new performance group, click .
6. Enter a name.
7. Click outside the name.
8. Select a workload from the list for the performance group.
9. Click **Save**.
10. [Optional] To delete a performance group, do the following:
 - a. Next to the performance group name, click .
 - b. Click **Yes** to confirm that you want to delete this performance group.
 - c. Click **Save**.

Exceptions View

The **Exceptions** view on the ruleset toolbar allows you to create exception thresholds and define actions to take when a threshold is exceeded for an executing query in each planned environment. Neither SLES 10 nor SLES 11 IWM systems provide the ability to define exceptions at the ruleset level. SLES 11 IWM systems do provide the ability to define tactical exceptions for tactical workloads.

The following tabs in the **Exceptions** view display all exception information for a ruleset from different perspectives:

By Planned Environment

Exceptions organized by planned environment. After selecting a planned environment, all exceptions are listed alphabetically by name. All workloads are listed in workload evaluation order. Select or clear the workload check boxes to enable or disable the use of individual exceptions for individual workloads. If the **All-WD** check box is selected, all check boxes in the row are selected and cannot be individually controlled.

By Workload

Exceptions are organized by workloads. After selecting a workload, all exceptions are listed by name. All planned environments are listed by name and by how they are arranged in the state matrix. Select or clear the planned environment check boxes to enable or disable individual exceptions for individual planned environments. Within a workload, if an exception is enabled for All-WD in a planned environment, the check box for the planned environment is selected and disabled. You can change the exception setting on the **Planned Environment** tab.

By Exception

List of all exceptions created for planned environments and workloads. After selecting an exception, all workloads are listed in workload evaluation order. All planned environments are listed by name. Select or clear the planned environment check boxes to enable or disable individual workloads for individual planned environments. If the All-WD check box is selected, all check boxes in the column are selected and cannot be individually controlled.

Exception Criteria Settings

Exceptions result when threshold values for the designated criteria are exceeded. Disk and skew criteria also require that the condition persist for a specified amount of time before an exception is triggered.



You can define an exception by specifying the values (and qualification time, if applicable) for the following criteria.

Name	Description
Maximum Spool Rows	The maximum number of rows that a query step can spool. Enter the maximum allowable number.
Spool Usage	The amount of spool space, per step, that a query uses. Enter the maximum value and select B (bytes), KB (kilobytes), MB (megabytes), or GB (gigabytes).
Number of AMPs	The number of AMPs that participate in a query. Enter the maximum number.
Blocked Time	The length of time a query is blocked by another query. Enter the maximum time in the format hhh:mm:ss (hours:minutes:seconds).

Name	Description
Elapsed Time	The length of time a query has been running (the response time). Enter the maximum amount in the format hhh:mm:ss (hours:minutes:seconds). By default the amount of time you enter includes blocked and throttle delay time. (Teradata Database 15.10 and later) You can exclude them by selecting the Exclude blocked time and Exclude throttle delay time check boxes.
I/O Count	The number of disk I/Os performed for the query. Enter the maximum number.
CPU Time (sum over all nodes)	The amount of CPU processing time consumed by a query, summed over all the nodes used. Enter the maximum time in CPU-seconds.
CPU Skew Percent	The percentage difference in CPU consumption between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum percentage and the qualification time in CPU-seconds.
I/O Skew Percent	The percentage difference in disk I/O counts between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum percentage and the qualification time in CPU-seconds.
CPU Skew Difference	The difference in CPU consumption between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum amount of time and the qualification time in CPU-seconds.
I/O Skew Difference	The difference in disk I/O counts between the busiest AMP and the average of all involved AMPs during the last exception interval. Enter the maximum value and the qualification time in CPU-seconds.
CPU Disk Ratio	The ratio of CPU consumption to disk I/O during the last exception interval. Enter the maximum ratio and the qualification time in CPU-seconds.
I/O Physical Bytes	The amount of I/O space, summed over all nodes, that a query consumes. Enter the maximum I/O physical bytes, select KB (kilobytes), MB (megabytes), or GB (gigabytes), and enter the qualification time in CPU-seconds.

Defining an Exception

An *exception* is a defined set of actions to take when a threshold is exceeded. If you add multiple criteria, all must be true in order for the exception to be triggered. Only one qualification time can be specified even if multiple qualified criteria are added. If qualified criteria are specified, the exception is triggered only if all of the qualified criteria are satisfied for the duration of the qualification time. After an exception is created, additional controls in the tabs of the **Exceptions** view allow you to manage the exception.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Exceptions**.
3. From any of the tabs in the **Exception** view, click  next to **Exceptions**.
4. Enter a name.
5. [Optional] Enter a description.
6. Under **Criteria**, select an option from the list, click , and specify the additional required information.

You can select and add multiple criteria. If you do so, all criteria must be satisfied for the exception to be triggered.

7. For qualified criteria marked with *, enter a value in the **Qualification Time** box.
8. Under **Actions**, select one of the following actions to be performed when the specified criteria are met:

Action	Description
Notification Only	Sends notification and takes no other action.
Abort	Stops the query.
Abort Selects Only	Stops the query only if it is a SELECT.
Change Workload to	Changes the query to the workload you select from the list. This action is unavailable for Blocked Time or Elapsed Time .

9. [Optional] Under **Notifications**, enable any of the following options for the start or the end of the event:

Notification	Description
Notify Database User	Sends an email alert to database user when the query stops. This option is available if you have selected only Abort/Abort Selects from actions.
Send Alert	Specifies the action to trigger.
Run Program	Specifies the Alerts registered programs to trigger.
Post to QTable	The string you enter is posted to the QTable.

10. Click **OK**.


SLES 10 Tactical CPU per Node Exceptions




If you select the **Tactical CPU per Node** exception criteria, the following rules apply:

- The exception can be enabled for workloads that use the tactical enforcement priority. Tactical workloads are in the top row of the **Workload Mapping** tab.
- The exception must have a **Change Workload to** action to another workload in the same resource partition column.
- The exception must have **Tactical CPU per Node** set to 5 CPU-seconds or less.
- The exception must have **CPU Time** set to a number greater than **Tactical CPU per Node**.
- Using **Tactical CPU per Node** criteria disables the **Blocked Time** and **Elapsed Time** criteria for the exception.




Managing Exceptions

Ruleset exceptions are managed from the ruleset **Exceptions** view.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Exceptions**.
3. [Optional] Create exceptions by clicking  next to **Exceptions** from any tab.
4. After exceptions have been created, click the **By Planned Environment**, **By Workload**, or **By Exception** tab to manage exceptions.
5. Depending on the tab, select **Planned Environment**, **Workload**, or **Exception** from the list.
6. [Optional] Edit an exception by doing any of the following:

Option	Description
Planned Environment	From the By Planned Environment tab, click  and select Edit .
Workload	From the By Workload tab, click  and select Edit .
Exception	From the By Exception tab, click  .

7. [Optional] Delete an exception by doing any of the following:

Option	Description
Planned Environment	From the By Planned Environment tab, click  and select Delete .
Workload	From the By Workload tab, click  and select Delete .
Exception	From the By Exception tab, click  .

8. In the selected view, select or clear the appropriate check boxes to enable or disable the exceptions for the listed workloads, planned environments, or workload and planned environment combinations.
9. Click **Save**.

Model Systems

You can create a model system and sample rulesets based on a Teradata version that is not yet available or installed. These models allow you to preview workload and ruleset options for a new Teradata release before you actually install the new version, such as upgrading from SLES 10 to SLES 11. You can work with these model rulesets without a physical system. Rulesets on a model system cannot be made ready or active.

Creating a Model System (Teradata System)


Create a model Teradata system to preview workload and ruleset options in an environment that is not yet available or installed.

1. From the **Workload Designer** view, select **Create Model System** from the **For System** list.
2. Enter a model system name up to 24 characters.
You can use spaces and alphanumeric characters.
3. Select **Teradata** for the system type.
4. Select the Teradata version.
5. Select your license type as follows:

Type	Description
Enterprise	TASM license <ul style="list-style-type: none"> • All workload management features are available, including states, sessions, filters, throttles, workloads, and exceptions.
Advanced	IWM license <ul style="list-style-type: none"> • States, sessions, filters, and throttles are available in SLES 10. • States, sessions, filters, throttles, and workloads are available in SLES 11.


6. Click **Create**.
7. Create rulesets to preview the options available in that environment.

Editing a Model System Name (Teradata System)

1. From the **Workload Designer** view, select the model system from the **For System** list.
2. Click .
3. Enter a name up to 24 characters.
You can use spaces and alphanumeric characters.
4. Click **OK**.

Deleting a Model System

When you delete a model system, all rulesets are also deleted.

1. From the **Workload Designer** view, select the model system from the **For System** list.
2. Click .
3. Select **Delete**.

Workload Designer (Aster System)

After selecting an Aster system, you can create, edit, view, or perform other actions on rulesets on that system. The ruleset views provide controls that allow you to define rules for managing workloads.

Workloads are known as *workload policies* in the *Teradata Aster® Database User Guide*.

Use the following features to manage workloads:

Rulesets

Collections of related throttles and workload rules.

Throttles

Limit the number of queries that are allowed to run concurrently.

Workloads

Prioritize and allocate resources among different sets of queries that are grouped by classification criteria, and set the service class to control query prioritization.

Note:

Reading the *Teradata Aster® Database User Guide* before creating workloads on an Aster system in Teradata Viewpoint is highly recommended.

Workload Designer View

The **Workload Designer** view shows summary information about rulesets for an Aster system.

Working

The **Working** section contains rulesets on the Viewpoint server. In **Working**, you can create, edit, and import rulesets, as well as other options. If a ruleset is locked by someone else, you have fewer options than if you are the ruleset owner. Rulesets in **Working** can also appear in **Active**.

Active

The **Active** section contains the active ruleset that is currently being used by the Teradata Aster system for workload management.

Ruleset background color indicates synchronization between the **Working** and **Active** sections. In the **Working** section, blue indicates the ruleset is not active and gray indicates the ruleset is an exact copy of the currently active ruleset. In the **Active** section, green indicates the ruleset is active on the target system.

If no exact copy of the active ruleset exists in the **Working** section, the active ruleset displays as Unknown Ruleset in the **Active** section.

Ruleset Options

The **Workload Designer** view shows individual rulesets and specific options for each ruleset based on the permissions assigned to your role and the section in which the ruleset is located.

Working

The **Working** section contains rulesets on the Viewpoint server that are editable.


Option	Description
Make Active	Makes the ruleset the active ruleset on the Aster system. Copies the ruleset to the Active section.
View/Edit	Opens the ruleset for viewing and editing.
Show All	Lists all ruleset attributes on one page.
Lock	Locks the ruleset so only the lock creator can edit the ruleset.
Unlock	Unlocks the ruleset so others can edit the ruleset.
Clone	Creates a copy of the ruleset. This option is useful if you want to use an existing ruleset as a base or template to create a ruleset.
Export	Exports the ruleset as an XML file. Use with the Import button to copy a ruleset from one system to another.
Delete	Removes the ruleset from the Working section.


Active

The **Active** section contains the active ruleset on the Aster server. When this ruleset does not match a working ruleset, an option is available to copy the ruleset to the **Working** section.

Creating a Ruleset

A *ruleset* for an Aster system is a collection of related throttles and workload rules. You can create multiple rulesets, but only one ruleset is active on the Aster system at a time. After creating a ruleset, you can use the toolbar buttons to specify settings, such as workloads. New rulesets are automatically locked so only the owner can edit the ruleset.



1. From the **Workload Designer** view, select a system from the list.
2. Click  to the right of **Local**.
3. Enter a name.
4. [Optional] Enter a description.
5. Click **Save**.

6. [Optional] Click **Throttles** and create a throttle.
7. [Optional] Click **Workloads** and create a workload.
8. Click  to return to the **Workload Designer** view.
The new ruleset appears in the **Working** section.

Editing a Ruleset

The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. Without permission to edit the ruleset, the option is **View** and the ruleset view is read-only. You can edit rulesets only in the **Working** section.

You can only edit rulesets that are unlocked or locked by you.



1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **View/Edit**.
3. Specify settings using the toolbar buttons and tabbed views.
4. Click **Save** after making changes in each view.
5. Click  to return to the **Workload Designer** view.
The date and time the ruleset was modified and by whom appears under the ruleset name in the **Working** section.

Related Information:

[Ruleset Locks](#)

Cloning a Ruleset


You can clone a ruleset to make an exact copy of the ruleset, except for the name. Cloning is a convenient way to create a ruleset using the settings of an existing ruleset as a base.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Clone**.
3. Enter a name.
4. [Optional] Enter a description.
5. Click **Save**.
6. Specify additional settings using the toolbar buttons and tabbed views.
7. Click **Save**.
8. Click  to return to the **Workload Designer** view.
The cloned ruleset appears in the **Working** section.

Deleting a Ruleset

You can delete a ruleset from the **Working** section which removes the ruleset from the Viewpoint system. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.

You cannot delete the active ruleset from the **Active** section. You cannot delete a locked ruleset unless you created it.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Delete**.
3. Select **Delete**.

The ruleset no longer appears in the **Working** section.


Importing a Ruleset

The import and export options can be used to copy a ruleset from one Viewpoint system to another. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. Only rulesets exported from **Workload Designer** and a database of the same release can be imported.

1. From the **Workload Designer** view, select a system from the list.
2. Click .
3. Enter a name for the imported ruleset.
4. Click **Browse**.
5. Locate and select the saved ruleset file.


Note:

Exported ruleset files might be stored in the download area configured for your browser.

6. Click **Save**.
 7. Click  to return to the **Workload Designer** view.
- The imported ruleset appears in the **Working** section.

Exporting a Ruleset




The import and export options can be used to copy a ruleset from one Viewpoint system to another. Only rulesets exported from Workload Designer and a database of the same release can be imported.

1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Export**.
3. Click **Save**.

The ruleset file is saved to your download area or the location you specify, depending on your browser settings.

Showing All Criteria in a Ruleset

You can display a read-only summary of all settings for a single ruleset.


1. From the **Workload Designer** view, select a system from the list.
2. In the local **Working** section, click the ruleset  and select **Show All**.
3. [Optional] Do any of the following:
 - Click **Collapse All** to hide all sections of the view.
 - Click **Print** to print the full summary.
 - Click  to show or  to hide the sections of the view.

Activating a Ruleset

Activating a ruleset allows you to begin enforcing your workload management rules on the target system.

Activating a ruleset automatically copies the ruleset to the **Active** section on the selected Aster system. Only one ruleset is active on the system at a time. When you activate a ruleset, all existing workload management settings on the Aster system are replaced and cannot be recovered unless you have a working copy.

The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.


1. From the **Workload Designer** view, select a system from the list.
2. For the ruleset that you want to activate in the **Working** section, click the ruleset  and select **Make Active**.
3. Click **Activate** to confirm that you want to activate this ruleset.
Your settings replace the current workload management settings on the Aster system. The ruleset appears in the **Active** section, displaying the date and time it was activated, and the user who activated it.

Related Information:


[Copying the Active Ruleset to the Working Section](#)

Copying the Active Ruleset to the Working Section

When the active ruleset does not match a working ruleset, you can copy the ruleset to the **Working** section. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action.

1. From the **Workload Designer** view, select the Aster system containing the ruleset you want to copy.
2. In the **Active** section, click the ruleset  and select **Copy to Working Ruleset**.



If the **Copy to Working Ruleset** option is not available, a copy of the ruleset already exists. Look for this gray ruleset in the **Working** section.

3. Enter a name.
 4. [Optional] Enter a description.
 5. Click **Save**.
 6. [Optional] Click **Throttles** and create a throttle.
 7. [Optional] Click **Workloads** and create a workload.
 8. Click  to return to the **Workload Designer** view.
- The new ruleset appears in the **Working** section.




Ruleset Locks

An exclusive lock can be placed on a ruleset so that the ruleset cannot be edited, deleted, or otherwise modified except by the owner of the lock. A ruleset is automatically locked when it is created and each time changes to the ruleset are saved. Use the **Workload Designer** view to lock and unlock rulesets. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. The Teradata Viewpoint Administrator can also grant your role permission to unlock any ruleset.

The **Workload Designer** view displays the ruleset lock status.

Option	Description
	Locked
	Unlocked


The ruleset views display the ruleset lock status.

Option	Description
	Unlocked
	Locked by the current user
	Locked by another user

Locking or Unlocking a Ruleset

Locking a ruleset prevents others from modifying the ruleset. The Teradata Viewpoint Administrator must grant your role permission to edit rulesets so you can complete this action. The Teradata Viewpoint Administrator can also grant your role permission to unlock any ruleset. A ruleset that is locked may be unlocked by the lock owner, or by any user in a role with permissions to unlock any ruleset.

1. From the **Workload Designer** view, select a system from the list.

2. In the local **Working** section, click the ruleset  and select **Lock** or **Unlock**.

General View

The **General** view on the ruleset toolbar enables you to change the name or description of a ruleset after it has been created.

Throttles View

The **Throttles** tab on the ruleset toolbar for an Aster system displays a list of throttles and their attributes, and allows you to create throttles that control concurrent queries.

The **Global Throttle** tab allows you to limit the maximum number of concurrent queries regardless of the attributes of the queries.

Throttles

A *throttle* limits the number of queries running at one time on an Aster system.

You specify the predicate when you create the throttle and it must be a valid SQL WHERE clause. You can use existing UDFs, SQL operators, workload management attributes, and AND and OR conjunctions.

You set a concurrency limit by specifying the maximum number of queries matching the predicate that can run at a given moment. Each query must pass all predicates before being admitted to the system. Queries submitted after the limit is reached are queued. *Throttles* are known as *admission limits* in the *Teradata Aster® Database User Guide*.


You can create a throttle to limit queries based on:

- A specific database to which the client is connected.
- A set of database roles of which the current user is a member, such as users who have been granted the admin role.
- A statement type, such as SELECT, DROP, EXPLAIN, REVOKE, SHOW, and others.
- Elapsed time, such as limiting a specific statement that has executed over 30 minutes by a user that is part of a group.

Global throttles limit the maximum number of concurrent queries to a value you specify regardless of the attributes of the queries. If the number of running queries reaches the value you set, new queries are queued. If the global throttle limit is not reached, the regular throttles setting determines if and when a query is admitted to the system. *Global throttles* are known as *global admission thresholds* in the *Teradata Aster® Database User Guide*.

Creating Throttles

You can use throttles to limit the number of queries running at one time.


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Next to **Throttles**, click .
4. Enter a name.
5. Create a WHERE clause and select attributes from the list for the workload predicate.
When specifying the predicate, you can use existing UDFs, SQL operators, workload management attributes, and AND and OR conjunctions.
6. Do one of the following to specify a concurrency limit:

Option	Description
Maximum concurrency	Allows the number of concurrent queries you specify to run. Queries submitted after the limit is reached are queued.
Deny access to new queries	Rejects new queries.

7. Click **Save**.

Deleting Throttles

You must be the ruleset owner or the ruleset must be unlocked in order to delete throttles.

1. From the ruleset toolbar, click **Throttles**.
2. Click  and select **Delete** in the row of a throttle.
3. Click **Delete** to confirm that you want to delete this rule.

Editing Global Throttles

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Throttles**.
3. Click the **Global Throttles** tab.
4. Do one of the following to specify a concurrency limit:

Option	Description
No limit	Applies no limit to the number of concurrent queries that can run.
Maximum concurrency	Allows the number of concurrent queries you specify to run. Queries submitted after the limit is reached are queued.
Deny access to new queries	Rejects new queries.

5. Click **Save**.

Workloads View

A *workload* is a group of queries that share characteristics so that a set of workload management controls can be applied to the group. A workload has working values and defining characteristics that are evaluated during the classification phase of system management. High-quality workload management can improve response times and provide more consistent response times for critical work.

The following requests are examples of workloads:

- Batch jobs further subdivided by region or organization for reporting
- Weekly or monthly reports that follow the calendar or a regular schedule
- Jobs that are always critical whenever they occur

The **Workloads** view, accessed from the ruleset toolbar, provides the following tabs:

Workloads

Displays a list of workloads, their attributes, and enabled status. You can add or delete a workload, or click a workload to view additional settings.

WD-Default is the default workload that is applied to queries not classified into a specific workload. The default workload appears last in the evaluation order to make sure it is applied only if no other workloads apply. The default workload cannot be deleted or disabled.

Evaluation Order

Evaluation order helps to determine the placement of queries into workloads. Queries are mapped to the first workload in the evaluation order whose predicate they match.

Service Classes

Service classes are an attribute you assign to each workload to control query prioritization. In a larger sense, service classes define how the database divides the available hardware resources among all running queries and other activities in the system, such as physical backups. For more information, see the *Teradata Aster® Database User Guide*.


SC-Default is the default service class applied to workloads without an assigned service class and is initially used by the default workload. SC-Default has a priority of medium and a weight of 100 by default, but you can change these values.

Creating and Editing a Workload

You can group queries that share characteristics into a workload for an Aster system so that a set of workload management controls can be applied to the group.

Service classes are an attribute you assign to each workload to control query prioritization. Consider creating service classes prior to starting this procedure so you can specify the service class when you create a workload.

WD-Default is the default workload that is applied to queries not classified into a specific workload. If you edit the default workload, you can change the service class, but not the predicate.

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Next to **Workloads**, click .
4. Enter a name.
The workload is enabled by default.
5. Create a WHERE clause and select attributes from the list for the workload predicate.
When specifying the predicate, you can use existing UDFs, SQL operators, workload management attributes, and AND and OR conjunctions.
6. Specify a service class.
7. Click **Save**.
8. [Optional] Set the evaluation order for workloads.
9. [Optional] Create services classes for workloads.
After you create service classes, consider editing existing workloads to apply those service classes.

Related Information:

[Setting Evaluation Order for Workloads](#)
[Creating Service Classes](#)

Evaluation Order

Evaluation order helps to determine the placement of queries into workloads. Queries are mapped to the first workload in the evaluation order whose predicate they match. Setting evaluation order is useful when you have created several workloads.

WD-Default is the default workload that is applied to queries not classified into a specific workload. The default workload appears last in the evaluation order to make sure it is applied only if no other workloads apply.

Setting Evaluation Order for Workloads

1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Evaluation Order** tab.
4. In the list of workloads, do one of the following to order the list and place higher priority workloads at the top:
 - Drag workload names to reorder.
 - Click in a number box, highlight the existing number, and enter a new number.

5. Click **Save**.

Service Classes

The service class controls query prioritization. Service classes define how an Aster system divides the available hardware resources among all active queries and other activities, such as physical backups, in the system.


When you create a service class for an Aster system, you can set the processing priority and memory limits.

- Processing priority represents the overall importance of the workload compared to other workloads. Weight is the relative resource allocation of the CPU and disk I/O within the priority. The weight value dictates the ratio of resource allocation. For example, if two statements execute with the same priority, but with weight values of 80 and 20, the system will aim to allocate resources in a 4:1 ratio, with most of the resources allocated to the statement with higher weight.
- Memory limits allows you to specify the maximum total memory per node for all queries associated with this service class. *Memory soft limit* is the highest percentage of memory the workload can occupy when there is contention for memory in the system. *Memory hard limit* is the highest percentage of memory the workload can occupy even if there is no contention for memory in the system.

These settings apply to all queries that map to that service class. Each service class can be used by one or more workloads.

Reading more about these limits in the *Teradata Aster® Database User Guide* is recommended before creating service classes in Teradata Viewpoint.

Creating Service Classes


1. Edit or create a ruleset.
2. From the ruleset toolbar, click **Workloads**.
3. Click the **Service Classes** tab.
4. Next to **Service Classes**, click .
5. Enter a name.
6. Do one of the following to set the processing priority:
 - Select **Prioritize queries**, select a priority from the list, and enter a weight from 1 (low) to 100 (high).
 - Select **Deny access to queries** to deny new queries from being admitted into the system for processing.
7. [Optional] Do one or more of the following to set memory limits:
The memory soft limit cannot be greater than the memory hard limit.

Option	Description
Memory soft limit check box	Select and enter the highest percentage of memory the workload can occupy if there is contention for memory in the system.
Memory hard limit check box	Select and enter the highest percentage of memory the workload can occupy if there is no contention for memory in the system.

8. Click **Save**.

Deleting Service Classes

You cannot delete a service class if it is associated with a workload; however, you can modify associated workloads to use a different service class or delete the workload.

1. From the ruleset toolbar, click **Workloads**.
2. Click the **Service Classes** tab.
3. Click  and select **Delete** in the row of a service class.
4. Click **Delete** to confirm that you want to delete this service class.

Workload Health Overview

The **Workload Health** portlet displays workload management activity in Analytics Database for one system at a time. The available options are different if Analytics Database is running on SUSE Linux Enterprise Server version 10 (SLES 10) or lower, or SLES 11. You choose which system, workloads, and metrics to monitor. **Filter** and **Sort** options allow you to customize the information displayed. Data in the **Workload Health** portlet is refreshed every minute and displays workloads that have the following characteristics:

- Completed processing according to their response-time Service Level Goals
- Missed their response-time Service Level Goals
- Are inactive
- Are disabled
- No defined response-time Service Level Goals

The health details view displays detailed metrics for an individual workload.

Note:

Support for TASM ARM is available for SQL Engine versions 17.10 and later.

Workload Health Concepts

What Is a Workload?

A *workload* is a group of queries that share characteristics so that a set of workload management controls can be applied to the group. A workload has working values and defining characteristics that are evaluated during the classification phase of system management. Queries are assigned to workloads based on classification criteria and exceptions.

What Is a Service Level Goal?

A *Service Level Goal (SLG)* measures whether queries are meeting defined workload management criteria. SLGs are defined for logging and reporting purposes. There are no provisions in Analytics Database to enforce these goals.

The **Workload Health** portlet displays only response-time SLGs.

What Is an Enforcement Priority?

When a workload is created in SLES 10, it is assigned an *enforcement priority* (EP) and an *allocation group* (AG). An EP is a label given to assign weights for the queries running in a workload. Each AG takes the EP of the first workload associated with it and after that, only workloads with the same EP can be associated to that AG. As a result, each AG is a collection of one or more workloads that share the same EP. Multiple AGs can use the same EP. Analytics Database uses the following EPs:

- **Tactical** queries are short, critical queries with defined service level goals. They are typically single or few-AMP queries or all-AMP queries that consume less than 1 CPU-second per node.
- **Priority** queries are important and have higher priority than most other work.
- **Normal** queries are the average priority work running on the system.
- **Background** queries run for work that does not have a response-time requirement.

What Are Tiers and Access Levels?

When a workload is created in SLES 11, it is assigned to one of the following workload management methods:

Tactical

Manages high-priority, critical queries with a short response-time requirement.

SLG Tier

Allocates resources according to the SLG tier on which the workload resides. Multiple tiers can be prioritized so higher-level tiers have higher resource allocation. This method is not used in systems with an IWM license.

Timeshare

Assigns the following access levels that indicate the priority of resource allocation:

Low

This is the baseline and lowest resource access rate for the Timeshare access method.

Medium

Queries are given an access rate 2 times the Low access level.

High

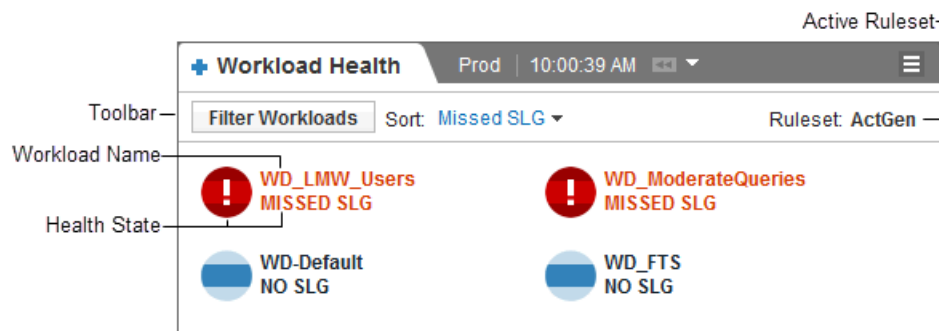
Queries are given an access rate 4 times the Low access level.

Top

Queries are given an access rate 8 times the Low access level.

Workload Health View

The **Workload Health** main view displays the health status of one or more workloads. Workload health is determined in relation to a response-time Service Level Goal (SLG).



System

Shows the name of the system currently displayed.

Active Ruleset

Shows the ruleset currently enabled on the Teradata system.

Filter Workloads

Opens a dialog box that allows you to filter workloads by group and name.

Sort

Enables you to sort workloads by **Missed SLG** or by **Name**.

Toolbar

Displays the active ruleset, the **Filter Workloads** button, and the **Sort** options.

Health State

Shows the workload health using color, icons, and predefined states.

Workload Name







Shows the name of a workload.

Click a workload icon to go to the workload health details view for that workload. This view is not available for workloads having a health state of NO DATA.

Workload Health States

Workload health is described using a set of icons and predefined states. Health states are listed in the default sort order.






Health states are determined by comparing the measured query response time with the expected response time and percent SLG metrics.

Icon	Text Color	State	Description
	Red	Missed SLG	Workload missed its response-time SLG for a period of time. Time periods for enforcement priorities in SLES 10 or lower: <ul style="list-style-type: none"> • Tactical: last minute • Priority: last 30 minutes • Normal and Background: last 60 minutes Time periods for tier/access levels in SLES 11: <ul style="list-style-type: none"> • Tactical: last minute • SLG Tiers: last 30 minutes • Timeshare: last 60 minutes
	Black	Met SLG	Workload met its enforcement priority SLGs in SLES 10 or lower, or tier/access level SLGs in SLES 11.
	Black	No SLG	No response time or percent SLGs have been set for this workload.
	Gray	No Data	No data is available for this workload.
	Gray	Disabled	Workload has been disabled.
	Gray	Inactive	Workload is inactive.

Selecting Filter Criteria

Use the **Filters Workloads** button to display only the workloads in which you are interested.

1. Click **Filter Workloads**.
2. Select at least one of the following filtering options:

Operating System	Description
SLES10 and earlier	<ul style="list-style-type: none"> • All Enforcement Priorities expands  to display the list. Select at least one enforcement priority. • All Health States expands  to display the list. Select at least one health state.
SLES11	<ul style="list-style-type: none"> • All Virtual Partitions expands  to display the list. Select at least one virtual partition. • All Tier Types/Access Levels expands  to display the list. Select at least one tier type or access level. • All Health States expands  to display the list. Select at least one health state.

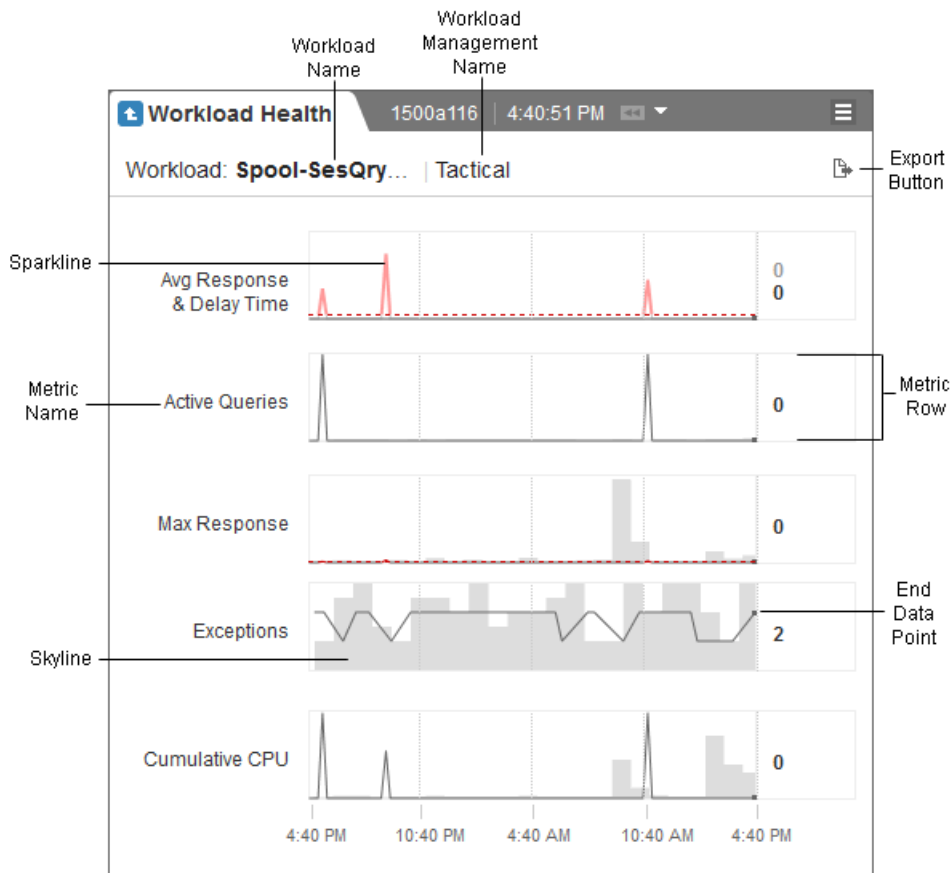
3. [Optional] Select the **Name Filters** tab and deselect a workload name check box to hide it from view.
4. Click **OK**.

Setting Sort Criteria

1. Click the **Sort** list to display the sort options.
2. To specify the sort criteria, do one of the following:
 - Click **Missed SLG** to sort workloads first by missed SLG health state and then alphabetically by workload name.
 - Click **Name** to sort workloads alphabetically by workload name, regardless of workload health state.

Workload Health Details View

The **Workload Health** details view displays metrics for a single workload. Use the **Settings** view to select the metrics. This details view appears after you click the workload icon or name for a workload in the **Workload Health** view. The **Workload Health** details view is not available for workloads with a health state of NO DATA.



Metrics are arranged in a two-column format:

- Metric names
- Sparkline graphs

As you hover over the sparklines, information balloons display event information. Each sparkline ends with a percentage or a value, depending on the metric.

The **Export** button allows you to create a .csv file containing selected data.

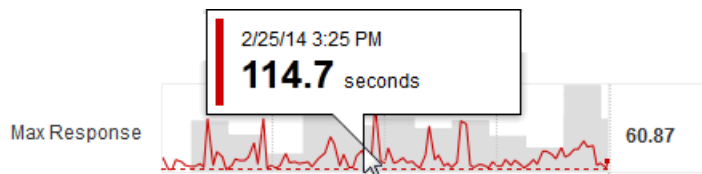
Workload Metrics Graphs

Sparkline graphs illustrate workload performance metrics.

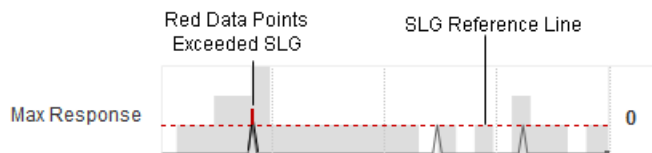
Solid lines represent the current data for their respective metric.

On a single-sparkline graph (for example, the Max Response metric), the solid line represents current data for the metric, and the grayed skyline represents the past averages for that same metric.

On a multi-sparkline graph, a metric can be expressed in more than one value. When you hover over the sparkline, the color of the metric name, end data points, and information balloons match.



A dotted line in the graph is called the *SLG reference line*. Data points on the SLG reference line are color-coded in red.




Workload Metrics

Metric	Description	Operating System
Aborted Queries	Number of aborted queries in the last 60 seconds	SLES 10 and 11
Active Queries	Number of active queries in the last 60 seconds	SLES 10 and 11
Arrivals	Number of queries arriving in the last 60 seconds. Historical data is not displayed for this metric.	SLES 10 and 11
Average CPU-Sec Per Query	Average CPU seconds per completed queries	SLES 10 and 11
Average Physical KB I/O Per Query	Average physical I/O usage, in kilobytes, of all queries in this workload	SLES 11
Average Response and Delay Time	Average response time and delay time for all active queries	SLES 10 and 11
Cumulative CPU	Cumulative CPU seconds per amp	SLES 10 and 11
Decayed Active Queries	Total number of queries that have decayed due to CPU or I/O (Timeshare workloads only)	SLES 11
Delayed Queries	Number of queries in the delay queue	SLES 10 and 11
Deferred Queries	Number of queries in deferred queue	SLES 11
Exceptions	Number of queries classified as exceptions in the last 60 seconds	SLES 10 and 11
Maximum CPU-Sec Per Query	Maximum CPU seconds used by completed queries is a default metric for the Tactical EP	SLES 10 and 11
Maximum Physical KB I/O Per Query	Maximum physical I/O usage, in kilobytes, for all queries in this workload	SLES 11

Metric	Description	Operating System
Maximum Response Time	Maximum response time for all active requests	SLES 10 and 11
Queries Throttled	Number of queries throttled in the workload in the last 60 seconds	SLES 10 and 11
Rejected Queries	Number of rejected queries in the last 60 seconds	SLES 10 and 11
Tactical Queries With Exceptions	Total number of queries reclassified to another workload as a result of a tactical exception based on CPU per node and I/O per node values. This metric is available only on the Metrics-Tactical tab in the Settings view.	SLES 11
Throughput	Number of completed queries in the last 60 seconds. This metric does not display historical data.	SLES 10 and 11

Exporting Metrics

You can export data to a .csv file for further analysis and formatting.

1. Select .
2. Save the file using the browser options.
The file is saved to your download area or to a location that you specify, depending on the browser settings.
3. [Optional] Reformat exported data to match the format in the view.

Settings View

The **Settings** view allows you to customize the portlet.

The following tabs allow you to customize your settings.

Systems

Choose a system to monitor.

Metrics

In SLES 10, select metrics for workloads in the Tactical, Priority, Normal, or Background enforcement priorities.

In SLES 11, select metrics for workloads in the Tactical, SLG Tier, or Timeshare management methods.

Trend Interval


Set the time frame displayed for the metric graphs.

Past Averages

Choose the time frame for which a historical average is calculated. This average is shown by the skyline plots in the details views.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.


Setting Metrics for Enforcement Priority or Tier/Access Level

For each enforcement priority group in SLES 10 or tier/access level in SLES 11, you can select and order metrics for display in the workload details view. Choose from a defined list of metrics.

1. Click  in the portlet frame and select **Settings**.
2. Click one of the metrics tabs for an enforcement priority group or tier/access level.

The display pane shows how metric rows appear in the portlet.


3. Do any of the following to change the way metric rows display:

Option	Description
Add a metric row	Select the check box of the metric in the Select metrics for display pane. That metric is added to the display pane beneath the rows currently displayed.
Remove a metric row	Clear the check box of the metric in the Select metrics for display pane.
Change metric row order	Use  to drag a metric already in the display pane to a new location in the pane (up or down).

4. Click **OK**.

Setting the Workload Trend Interval


Use the **Trend Interval** tab in the **Settings** view to set the length of time represented in sparkline graphs.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Trend Interval** tab.
3. Select an interval value from the list.
The **Sample view** pane shows an example of how data would appear when viewed at the selected interval.
4. Click **OK**.

Setting Past Averages

Specify the time frame used to calculate and display the shaded skyline graph in the details view.

Each segment of the skyline graph is the average of data samples taken each week at the same time and day of the week.

1. Click  in the portlet frame and select **Settings**.
2. Click the **Past Averages** tab.
3. Enter the number of weeks up to 99.
4. Click **OK**.

Workload Monitor Overview

The **Workload Monitor** portlet allows you to monitor workload management activity in Analytics Database. The available options are different if Analytics Database is running on SUSE Linux Enterprise Server version 10 (SLES 10) or lower, or SLES 11.

The **Workload Monitor** portlet displays the following:

- Multiple views of workloads and related data
- Allocation of system resources per workload or associated groupings
- Information about how queries are being processed
- Information about active queries

Note:

Support for TASM ARM is available for SQL Engine versions 17.10 and later.

Workload Monitor Concepts

What Is a Workload?

A *workload* is a group of queries that share characteristics so that a set of workload management controls can be applied to the group. A workload has working values and defining characteristics that are evaluated during the classification phase of system management. Queries are assigned to workloads based on classification criteria and exceptions.

What Is a Service Level Goal?

A *Service Level Goal (SLG)* is defined criteria against which workload management performance is measured. SLGs are defined for logging and reporting purposes. Only response time SLGs are used to indicate workload status. There are no provisions in Analytics Database to enforce these goals.

What Is an Exception?

A workload management *exception* is a collection of one or more thresholds that cause a defined action to occur. The **Workload Monitor** portlet allows you to view details of the workload exceptions as well as queries that are reclassified to another workload.

What Is an Enforcement Priority?

When a workload is created in SLES 10, it is assigned an *enforcement priority* (EP) and an *allocation group* (AG). An EP is a label given to assign weights for the queries running in a workload. Each AG takes the EP of the first workload associated with it and after that, only workloads with the same EP can be associated to that AG. As a result, each AG is a collection of one or more workloads that share the same EP. Multiple AGs can use the same EP. Analytics Database uses the following EPs:

- **Tactical** queries are short, critical queries with defined service level goals. They are typically single or few-AMP queries or all-AMP queries that consume less than 1 CPU-second per node.
- **Priority** queries are important and have higher priority than most other work.
- **Normal** queries are the average priority work running on the system.
- **Background** queries run for work that does not have a response-time requirement.

What Are Tiers and Access Levels?

When a workload is created in SLES 11, it is assigned to one of the following workload management methods:

Tactical

Manages high-priority, critical queries with a short response-time requirement.

SLG Tier

Allocates resources according to the SLG tier on which the workload resides. Multiple tiers can be prioritized so higher-level tiers have higher resource allocation. This method is not used in systems with an IWM license.

Timeshare

Assigns the following access levels that indicate the priority of resource allocation:

Low

This is the baseline and lowest resource access rate for the Timeshare access method.

Medium

Queries are given an access rate 2 times the Low access level.

High

Queries are given an access rate 4 times the Low access level.

Top

Queries are given an access rate 8 times the Low access level.

TASM and IWM Systems in SLES 10 and SLES 11

Teradata Workload Management offers different strategies, depending on the platform and license you purchase:

TASM

Teradata Active System Management (TASM) is included in the enterprise tier license to perform full workload management in a Teradata system.

IWM

Teradata Integrated Workload Management (IWM) is a Teradata feature that performs basic workload management. Customers on base EDW platforms can optionally purchase TASM.

Note:

If you have any questions about your TASM license, contact your Teradata Account Manager.

TASM offers a full workload management feature set. IWM workload management features are more limited and have different options in SLES 10 and SLES 11, depending on the system you select.

States

The IWM system has no health conditions.

Virtual Partitions

The IWM system has no virtual partitions.

Workload Management Methods

The SLES 11 IWM system uses only the tactical and timeshare management methods.

There is no Metrics-SLG Tier tab in Settings because the SLG Tiers management method is not available.

Exceptions

The SLES 11 IWM system has only tactical workload exceptions.

Workload Monitor Views

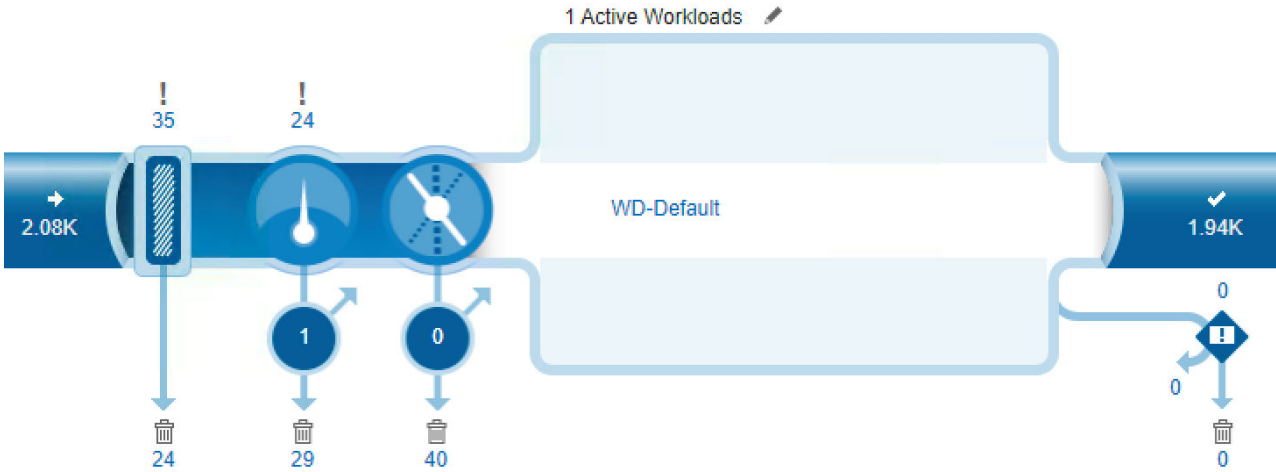
This topic lists the views available in the **Workload Monitor** portlet. More details are available from the request status icons on many views.

View Name	Description
Dynamic Pipes View	Displays workload performance in near real-time at each system management point of control
Static Pipes View	Displays workload metrics for all workloads and system resources by workload in the EP or tier/access level
Distribution View	Displays active workload resource consumption: <ul style="list-style-type: none"> • SLES 10: Displays CPU consumption per workload and relative weight per allocation group • SLES 11: Displays system resource consumption by virtual partition or workload
Distribution Details View	In SLES 11, displays system resource consumption by tier/access level or workload
Workload Details View	Displays details for a workload: <ul style="list-style-type: none"> • Cumulative: Cumulative values for the workload generated during the selected sampling period • Snapshot: Values for the workload at the current point in time • Trends: Sparkline graphs for the workload generated during the trend interval set in Settings
Workload Group View	Displays the group of workloads that belong to the AG in SLES 10 or tier/access level in SLES 11: <ul style="list-style-type: none"> • SLES 10: Percentage of AG CPU and system CPU consumed by each workload in the group • SLES 11: Percentage of partition CPU and system CPU consumed by each workload in the group






View Name	Description
State Details View	Displays information about the current and previous states, planned environments, health conditions, and events
Active Requests View	Displays metrics for requests that are active
Deferred Requests View	Displays metrics for requests that were deferred by Arrival Rate Meter
Delayed Requests View	Displays metrics for requests that were delayed by system and workload throttles
Request Details View	Displays statistics and information about the session

Request Status Icons

You can click the values associated with request status icons to see more detail.






Request Status	Icon	Description
Arrivals		Requests arriving into system management
Rejections Aborted		Requests that were not allowed to continue processing. This icon appears with the following values: <ul style="list-style-type: none"> Number of requests rejected by system filters Number of requests rejected by system throttles and workload throttles Number of requests aborted because of exceptions Number of requests rejected by Arrival Rate Meter



Request Status	Icon	Description
Warnings		Requests that have been allowed to continue processing after a warning message has been written to the system log file by System filters and Arrival Rate Meter
Delays		Requests that were delayed and completed or are currently delayed. Does not include requests that were delayed but are still active
Deferred Requests		Requests that were deferred and completed or are currently deferred. Does not include requests that were deferred but are still active.
Exceptions Reclassified		Requests that triggered an exception. This icon appears with these values: <ul style="list-style-type: none"> • Total number of workload exceptions in the sampling period • Number of requests reclassified to a different workload
Completions		Requests that executed successfully and no longer require system resources

Selecting Views

You can select a summary view and a sampling time frame to identify issues in workload management. Sampling time frame totals are updated during the portlet refresh, which is every 60 seconds. Workload details views are available from each summary view.

1. Choose one of the following summary views:

View	Description
Dynamic Pipes 	<ol style="list-style-type: none"> Click the Dynamic Pipes icon. Select a sampling time frame from the Cumulative System Data list. [Optional] In SLES 11, select a virtual partition to display data for only that partition.
Static Pipes 	<ol style="list-style-type: none"> Click the Static Pipes icon. Select a sampling time frame from the Cumulative System Data list. [Optional] In SLES 11, select a virtual partition to display data for only that partition.
SLES 10 Distribution 	<ol style="list-style-type: none"> Click the Distribution icon. Select a sampling time frame from the list. Click any workload or allocation group (AG) name to go to another view: <ul style="list-style-type: none"> • Click a workload name (other than TASM Unused/Idle or System) under CPU by Workload to access the workload details view. • Click an AG name under Relative Weight to access the workload group view.
SLES 11	<ol style="list-style-type: none"> Click the Distribution icon. Select CPU Consumption.

View	Description
Distribution Using CPU Consumption 	<ul style="list-style-type: none"> c. Select a sampling time frame from the list. d. Select a distribution details view: <ul style="list-style-type: none"> • Click a virtual partition name (other than TASM Unused/Idle or Internal) under CPU by Virtual Partition to access the distribution details view. • Click a workload name (other than TASM Unused/Idle or Internal) under CPU by Workload to access the workload details view. e. Select a group or workload details view: <ul style="list-style-type: none"> • Click a tier/access name under CPU by Tier/Access Level to access the group details view. • Click a workload name under CPU by Workload to access the workload details view.
SLES 11 Distribution Using I/O Consumption 	<ul style="list-style-type: none"> a. Click the Distribution icon. b. Select I/O Consumption. c. Select a sampling time frame from the list. d. Select a distribution details view: <ul style="list-style-type: none"> • Click a virtual partition name (other than TASM Unused/Idle or Internal) under I/O by Virtual Partition to access the distribution details view. • Click a workload name (other than TASM Unused/Idle or Internal) under I/O by Workload to access the workload details view. e. Select a group or workload details view: <ul style="list-style-type: none"> • Click a tier/access name under I/O by Tier/Access Level to access the group details view. • Click a workload name under I/O by Workload to access the workload details view.

2. From any of the workload summary views, click a workload name to display the workload details view, and click one of the following time frame buttons:

Time Frame	Description
Cumulative	<ul style="list-style-type: none"> a. Click Cumulative to display aggregated values for the selected sampling time frame. b. Select a sampling time frame from the Cumulative Workload Data list. c. Click the values associated with the request status icons to view details.
Snapshot	Click Snapshot to display current workload values.
Trends	Click Trends to display sparkline graphs for the trend interval set in Settings .

State Matrix

The Teradata system state matrix icon in the toolbar shows changes in health condition (with TASM license only), state, or planned environment during the cumulative sampling time frame.

The state matrix icon uses color to show the following:

Dark blue

Active-state cell

Medium blue

Previously active-state cell

Light blue

Inactive-state cell

Note:

During a state change, the cell representing the previous state changes from dark blue to medium blue. If there was a second state change during the sampling time frame, the previous state cell is shown in light blue.

The number of cells in the state matrix icon depends on the monitored system. If a one-by-one state matrix is configured, the state matrix icon appears as one active cell. When you hover over the state matrix icon, information balloons show detail information and the last state change if there was a state change within the sampling time frame.

In SLES 11, click the state matrix to view detailed information about the state.

State Details View

In SLES 11, the state details view provides additional information about the current and previous states, planned environments, health conditions, and events. Highlighted rows indicate the current planned environment or health condition that is in effect.

State	Description	
Active Planned Environments	Lists all planned environments by precedence that are in an active state. The first environment is used to determine the state.	
Active Health Conditions	Lists all health conditions by severity that are in an active state. The first health condition is used to determine the state.	
All Events	Lists all events and whether they are currently active. (Teradata Database 16.0 and later) Flex events are listed when the flex feature is configured.	
	FlexThrottle-AwtEvent	Flex throttle event was triggered because there were a sufficient number of AWTs or AMPs with available AWTs.
	FlexThrottle-CpuEvent	Flex throttle event was triggered because there were a sufficient number of AWTs or AMPs with available AWTs <i>and</i> CPU utilization was less than the threshold designated when the flex throttle was set.

You can sort on the **EVENT NAME** column in the **All Events** table by clicking on the column heading. A second click reverses the sort order.

Dynamic Pipes View

The Dynamic Pipes view helps you analyze workload performance.

Toolbar

Navigation Icons: You can click them to access the summary views.

State Matrix: Shows the changes in state, planned environment, or health condition during the cumulative sampling time frame. In SLES 11, you can click it to see details about the current state and previous states.

Request Icons: Shows the active and delayed requests. You can click them to see more detail.

Ruleset: Displays the name of the active ruleset.

Sampling Time Frame Control

Filters the system information for the selected sampling time frame.

Virtual Partition List

In SLES 11, filters the system information for the selected virtual partition or all virtual partitions.

Filter

Modifies the active workload display. If there are more than 18 workloads, they are grouped by EP in SLES 10 and by tier/access level in SLES 11.

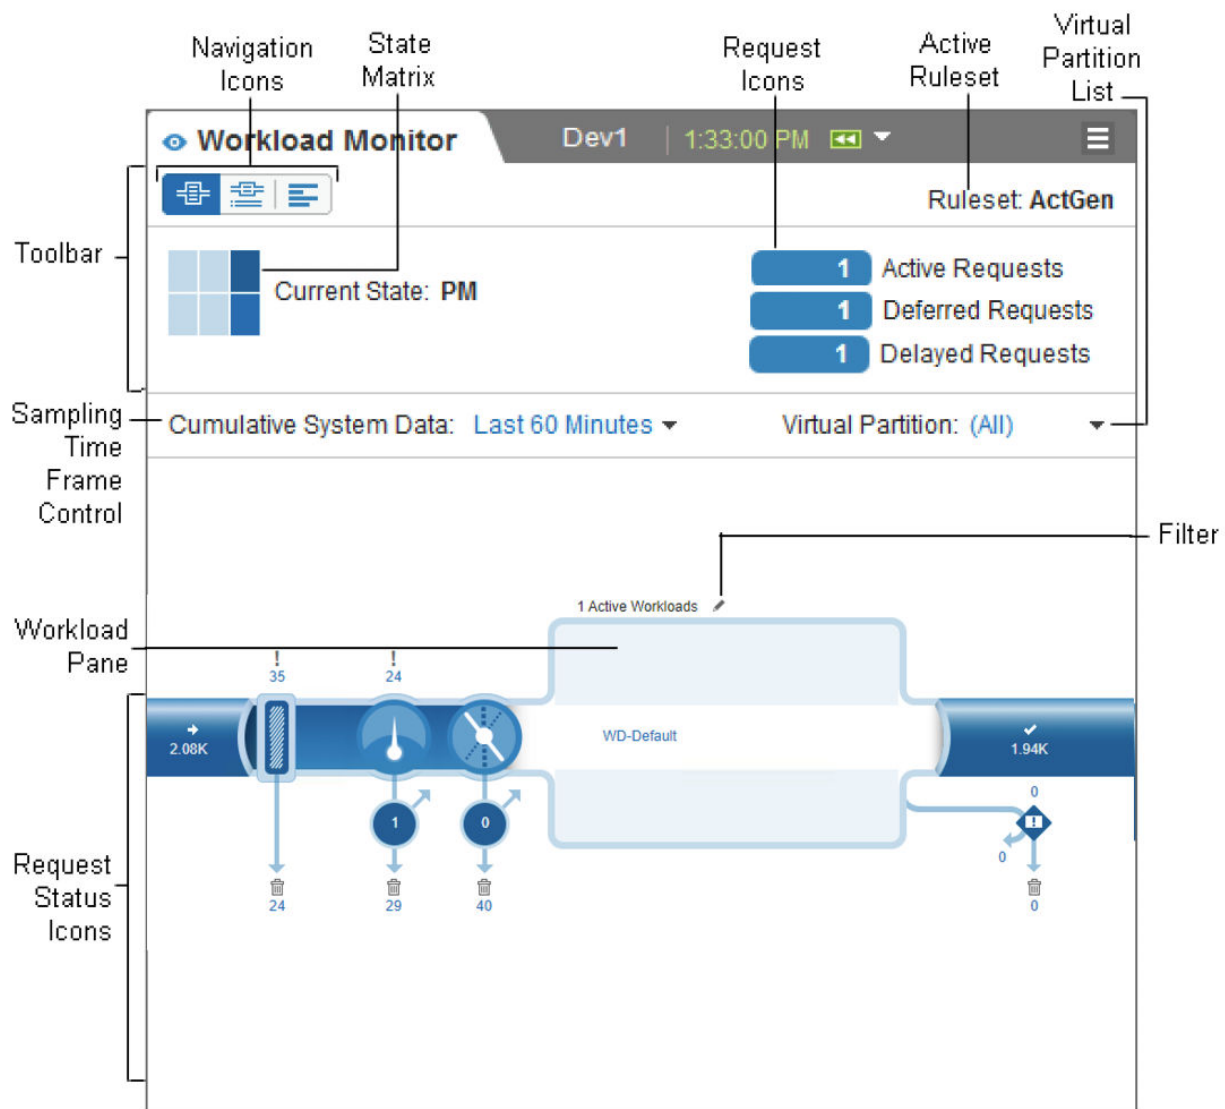
Workload Pane

Displays workloads, based on the filter criteria, and several interactive indicators:


- Red workload pipes highlight workloads that missed their SLGs.
- You can click the workload name to see the workload details view and access cumulative, snapshot, and trend data for each workload.
- A pulsing circle near the workload name indicates that a workload is adding requests to the delay queue.

Request Status Icons

Shows the number of requests processed at each step, including at system filters and throttles. You can click the values associated with these icons to see more detail.





Selecting Filter Criteria

1. In the Dynamic Pipes view, click  before the image. The **Filter Workloads** dialog box appears.
2. Select at least one of the following filtering options:

Note:

For SLES 10, at least one option must be selected from **All Enforcement Priorities** as well as at least one from the other categories. For SLES11, at least one option must be selected from **All Tier Types/Access Levels** as well as at least one from the other categories.

Operating System	Description
SLES 10	<ul style="list-style-type: none"> • Group by Enforcement Priority groups workloads by their enforcement priority. • All Enforcement Priorities expands  to display the list. Select at least one enforcement priority. • Missed SLG displays only those workloads which have missed their SLGs. • Active Workloads displays only those workloads which were active within the sampling period.
SLES 11	<ul style="list-style-type: none"> • Group by Tier Type/Access Level groups workloads by their tier type or access level. • All Tier Types/Access Levels expands  to display the list. Select at least one tier type or access level. • Missed SLG displays only those workloads which have missed their SLGs. • Active Workloads displays only those workloads which were active within the sampling period.

3. [Optional] Select the **Name Filters** tab and deselect a workload name check box to hide it from view.
4. Click **OK**.

Static Pipes View

The Static Pipes view displays all workloads and indicates how many are active.

Toolbar

Navigation Icons: You can click them to access the summary views.

Ruleset: Displays the name of the active ruleset.

State Matrix: Shows the changes in state, planned environment, or health condition during the cumulative sampling time frame. In SLES 11, you can click it to see details about the current state and previous states.

Request Icons: Shows the active and delayed requests. You can click them to see more detail.

Sampling Time Frame Control

Filters the system information for the selected sampling time frame.

Virtual Partition List

In SLES 11, filters the system information for the selected virtual partition or all virtual partitions.

Request Status Icons

Shows the number of requests processed at each step, including at system filters and throttles. You can click the values associated with these icons to see more detail.

Tabs

Workloads: Displays metrics for all workloads within the last sampling time frame. You can click a workload name to access the Workload Pipe view and see cumulative, snapshot, and trend data for each workload.

Enforcement Priorities: In SLES 10, displays summary data about all active workloads, aggregated across all EPs. Available metrics include the total number of active requests classified into one or more workloads in the EP, number of completed requests within the EP, average CPU time, and percentage of workloads that met their SLGs.

Tier/Access Level: In SLES 11, displays summary data about all active tiers and access levels for the selected virtual partition. Available metrics for each tier/access level include total number of active requests, number of completed requests, average CPU time, and the percentage of queries in the tiers/access levels that met their SLGs.

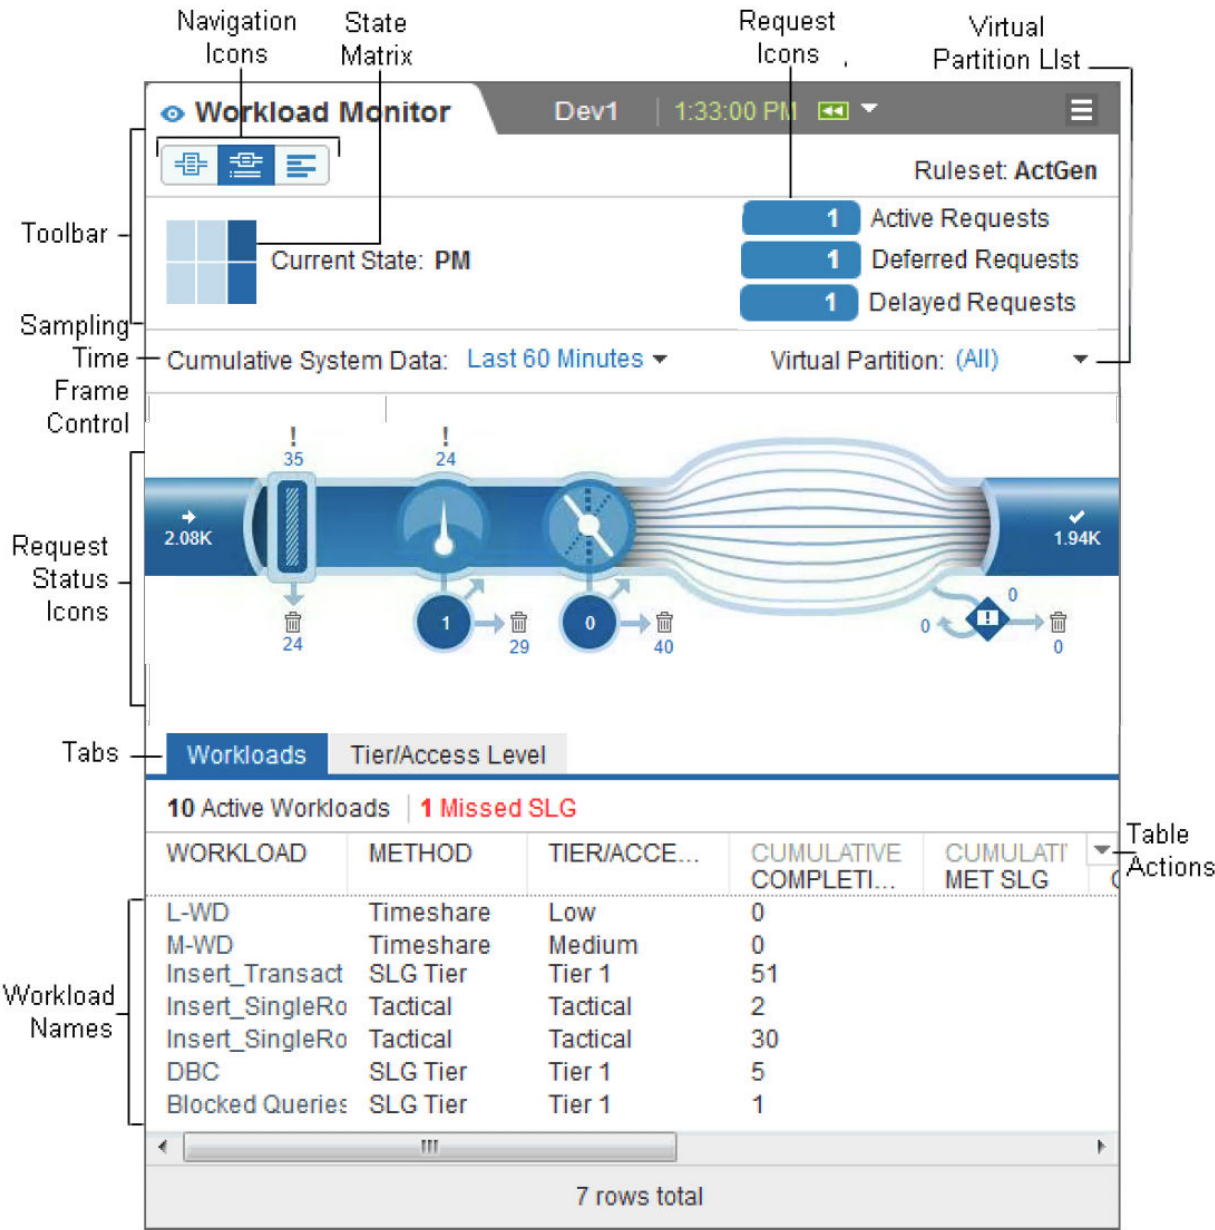
☐ Table Actions

Click ☐ to select the columns displayed on the **Workloads**, **Enforcement Priorities**, and **Tier/Access Level** tabs. You can also select **Export** to create a .csv file containing the information displayed in the summary table.

For more information, see [Summary Table Controls](#).

Workload Name

Shows the workload name. You can click it to show the workload details view and access cumulative, snapshot, and trend data for each workload.



Static Pipes Metrics

Workload metrics are displayed for each workload on the Static Pipes view.

Metric	Description	Operating System
Active Sampling Period	Sparkline displaying the number of active requests	SLES 10 and 11
Average CPU Time	Average CPU time, in seconds, of workload requests	SLES 10 and 11
Cumulative Aborted	Number of requests that have been aborted	SLES 10 and 11
Cumulative Arrivals	Number of requests arriving	SLES 10 and 11
Cumulative Arrivals Flex Throttles	Number of queries released as a result of flex throttles.	SLES 11
Cumulative Completion Flex Throttles	(Teradata Database 16.0 and later) Number of completed queries released as a result of flex throttles	SLES 11
Cumulative Completions	Number of requests completing	SLES 10 and 11
Cumulative Exceptions	Number of workload exceptions	SLES 10 and 11
Cumulative Met SLG	Percentage that a workload is meeting its defined SLG criteria	SLES 10 and 11
Current Active	Number of active requests	SLES 10 and 11
Current Delayed	Number of requests in the delay queue	SLES 10 and 11
Delay (AVG)	Average delay time for all active requests	SLES 10 and 11
Delay Queue	Sparkline displaying the number of requests in the delay queue	SLES 10 and 11
Defer Queue	Sparkline displaying the number of requests in the deferred queue	SLES 11
Enforcement Priority	Assigned enforcement priority	SLES 10
Method	Name of the workload management method, such as tactical, SLG tier, timeshare	SLES 11
Tier/Access Level	Tier and access level	SLES 11
Workload	Name of the workload	SLES 10 and 11

Active Requests Metrics

The following metrics can be displayed for active requests.

Metric	Description
Account	Account from which a query was submitted
Blocked Time	How long the query has been blocked
Snapshot CPU Skew	CPU skew during the last sample

Metric	Description
CPU Time	The amount of CPU processing time consumed by a query
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample
Duration	How long the query has been running
Host ID	Host ID or LAN ID associated with the PE that processed the login request for the session
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
In State	How long the query has been in the current state
Snapshot I/O Skew	I/O skew during the last sample
Δ I/O	I/O count since the last sample
Partition	Partition in which the query is running
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Proxy Username	(Teradata Database 15.10 and later) Name of the proxy user when a trusted session is used
Request CPU	CPU seconds needed to run the query
Request Count	Number of queries the session has executed
Request CPU Skew	CPU skew for the current query
Request I/O	Disk I/Os performed to run the query
Request I/O Skew	I/O skew for the current query
Session ID	Unique session identifier
Spool	Spool space the query requires
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Start	Time that the query started running on Analytics Database
State	Text describing the current state of the query
State Icon	Icon representing the current state of the query
Temp Space	Temp space used by the query

Metric	Description
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Username	Name of the user who submitted the query
Virtual Partition	(SLES 11) Name of the virtual partition to which the workload belongs
Workload	Workload to which the query belongs

Deferred Requests Matrix

The following metrics can be displayed for deferred requests. The metrics are listed by the tabs on which they appear.

By Workload

This tab displays all queries that are currently deferred.

Metric	Description
Account	Account from which a query was submitted
Blocked Time	How long the query has been blocked
Blocking Count	Number of consecutive times a session has blocked at least one other session
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample
Delay Time	Average time requests spends in the delay queue
Duration	How long the query has been running
Host ID	Host identifier or LAN identifier associated with the PE that processed the login request for the session
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
In State	How long the query has been in the current state
Δ I/O	I/O count since the last sample
Meter Deferred	Indicates whether any requests are deferred due to Arrival Rate Meter .
Partition	Partition in which the query is running

Metric	Description
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Proxy Username	(Teradata Database 15.10 and later) Name of the proxy user when a trusted session is used
Request CPU	CPU seconds needed to run the query
Request Count	Number of queries the session has executed
Request CPU Skew	CPU skew for the current query
Request I/O	Disk I/Os performed to run the query
Request I/O Skew	I/O skew for the current query
Session ID	Session identifier in which the query is running
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample
Spool	Spool space the query requires
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Start	Time that the query started running on Analytics Database
State	Text describing the current state of the query
State Icon	Icon representing the current state of the query
Temp Space	Temp space used by the query
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Username	Name of the user who submitted the query
Virtual Partition	Name of the virtual partition to which the workload belongs, for a SLES 11 system
Workload	Workload to which the query belongs

By Meter

This tab displays all queries currently included in a Meter counter. A query that is included in a meter counter might still be executing; it is deferred only if the limit is exceeded.

Metric	Description
Host ID	Host identifier or LAN identifier associated with the PE that processed the login request for the session
IFP ID	Identifier of the parsing engine vproc that initiated the request

Metric	Description
Meter Deferred	Indicates whether any requests are deferred due to Arrival Rate Meter .
Overridable	Indicates whether the query can be aborted or released
Request	Request number associated with the delayed request or session
Throttle Name	Name of the throttle managing the query
Session ID	Session identifier in which the query is running
Time Held	The total number of seconds the request has been held
Username	Name of the user who submitted the query

Meter Counts

This tab lists all currently active meters and their limits.

Metric	Description
Active Query Count	Number of active requests
Limit	Concurrency limit set by the throttle
Meter Deferred Count	Number of deferred requests.
Throttle Name	Name of the active throttle

Delayed Requests Metrics

The following metrics can be displayed for delayed requests. The metrics are listed by the tabs on which they appear.

By Workload

This tab displays all queries that are currently delayed.

Metric	Description
Account	Account from which a query was submitted
Blocked Time	How long the query has been blocked
Blocking Count	Number of consecutive times a session has blocked at least one other session
CPU Use	Percent of available CPU seconds on the system used during the last sampling period
Δ CPU	Total CPU usage time consumed, in seconds, since the last sample
Delay Time	Average time requests spends in the delay queue

Metric	Description
Duration	How long the query has been running
Host ID	Host identifier or LAN identifier associated with the PE that processed the login request for the session
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
In State	How long the query has been in the current state
Δ I/O	I/O count since the last sample
Partition	Partition in which the query is running
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Proxy Username	(Teradata Database 15.10 and later) Name of the proxy user when a trusted session is used
Request CPU	CPU seconds needed to run the query
Request Count	Number of queries the session has executed
Request CPU Skew	CPU skew for the current query
Request I/O	Disk I/Os performed to run the query
Request I/O Skew	I/O skew for the current query
Session ID	Session identifier in which the query is running
Snapshot CPU Skew	CPU skew during the last sample
Snapshot I/O Skew	I/O skew during the last sample
Spool	Spool space the query requires
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Start	Time that the query started running on Analytics Database
State	Text describing the current state of the query
State Icon	Icon representing the current state of the query
System Throttle Delay	Indicates whether any requests were delayed due to a system throttle
Temp Space	Temp space used by the query

Metric	Description
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Username	Name of the user who submitted the query
Utility Throttle Delay	Indicates whether any requests were delayed due to a utility throttle
Virtual Partition	Name of the virtual partition to which the workload belongs, for a SLES 11 system
Workload	Workload to which the query belongs
Workload Throttle Delay	Indicates whether any requests were delayed due to a workload throttle

By Throttle

This tab displays all queries currently included in a throttle counter. A query that is included in a throttle counter might still be executing; it is delayed only if the throttle limit is exceeded.

Metric	Description
Blocking Count	Number of consecutive times a session has blocked at least one other session
Flex Throttle Eligible	(Teradata Database 16.0 and later and a SLES EDW system) Indicates that the query is running in a workload that has flex throttles enabled for the current state
Host ID	Host identifier or LAN identifier associated with the PE that processed the login request for the session
IFP ID	Identifier of the parsing engine vproc that initiated the request
Overridable	Indicates whether the query can be aborted or released
Request	Request number associated with the delayed request or session
Throttle Name	Name of the throttle managing the query
Session ID	Session identifier in which the query is running
Throttle Type	Type of throttle impacting the query
Time Held	The total number of seconds the request has been held
Username	Name of the user who submitted the query

Throttle Counts

This tab lists all currently active throttles and their concurrency limits. A throttle might report queries as delayed even if its own concurrency limit is not met; this could occur due to the presence of other applicable throttles rather than because the throttle has placed queries in the delay queue itself.

Metric	Description
Active Query Count	Number of active requests
Delayed Query Count	Number of delay requests
Limit	Concurrency limit set by the throttle
Throttle Name	Name of the active throttle
Throttle Type	Type of throttle impacting the query

Distribution View

The Distribution view displays active workload resource consumption in an interactive graphical format.

Toolbar

Navigation Icons: You can click them to access the summary views.

Ruleset: Displays the name of the active ruleset.

State Matrix: Shows the changes in state, planned environment, or health condition during the cumulative sampling time frame. In SLES 11, you can view details about the current state and previous states.

Bars for SLES 10

You can see two bar charts that display actual system resource consumption compared to allocated system resources. Hover over any workload or AG name to see the relationship between the two. Mouse over the bar to see a balloon with more AG data. You can click any workload or AG name other than **TASM Unused/Idle** or **System** to go to another view:

- **CPU By Workload:** Displays the percentage of system CPU used by each workload that was active within the sampling time frame. Click a workload to access the workload details view.
- **Relative Weight:** Displays the system resources allocated to an AG. Every AG is associated with a relative weight that tells Teradata Priority Scheduler the priority of access to CPU for the requests running under that AG.

Bars for SLES 11

You can see two bar charts that compare system resource consumption by virtual partition or by workload.

- **CPU By Virtual Partition:** Displays the percentage of system CPU used by each virtual partition that was active within the sampling time frame. Click a virtual partition to access the Distribution details view.
- **CPU By Workload:** Displays the percentage of system CPU used by each workload that was active within the sampling time frame. Hover over a bar to see the percentage

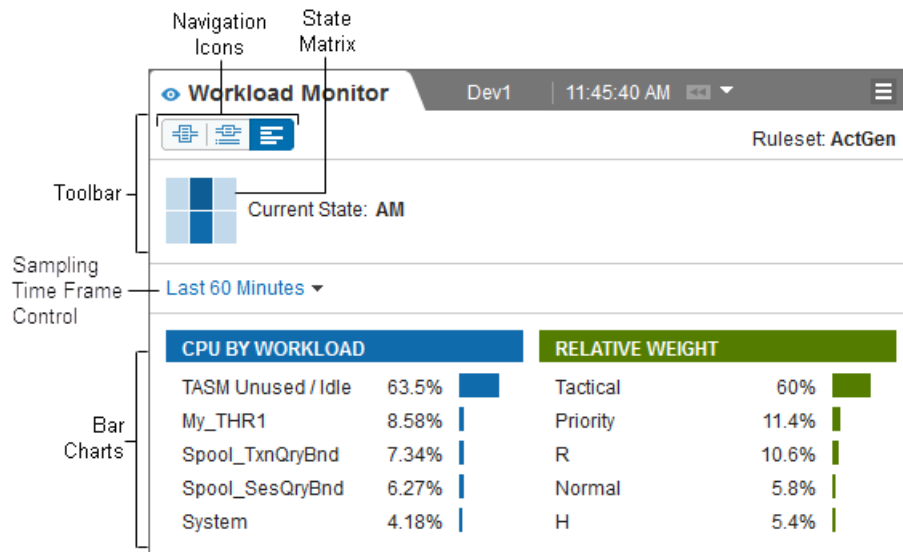
of partition CPU used by that workload. Click a workload to access the workload details view.

CPU Consumption or **I/O Consumption** toggles the consumption data for the sampling time frame.

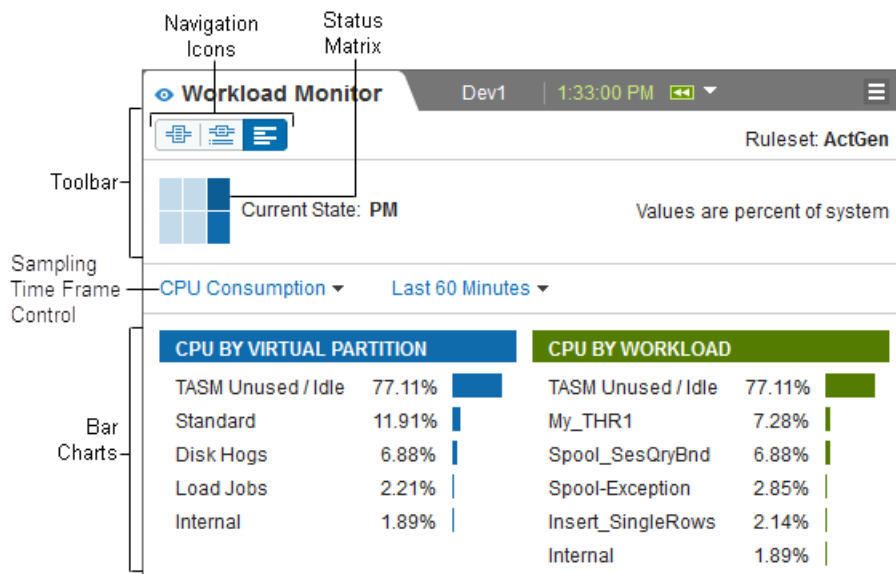
Sampling Time Frame Control

Filters the system information for the selected sampling time frame.

Following is an example Distribution view that you can see in SLES 10.



Following is an example Distribution view that you can see in SLES 11.



Distribution Details View

In SLES 11, the Distribution details view displays resource consumption for a virtual partition by tier/access level or workload.

State Matrix

Shows the changes in state, planned environment, or health condition during the cumulative sampling period. In SLES 11, you can click it to see details about the current state and previous states.

Toolbar

CPU Consumption or I/O Consumption: Toggles the consumption data for the sampling period.

Sampling Time Frame Control: Filters the system information for the selected sampling time frame.

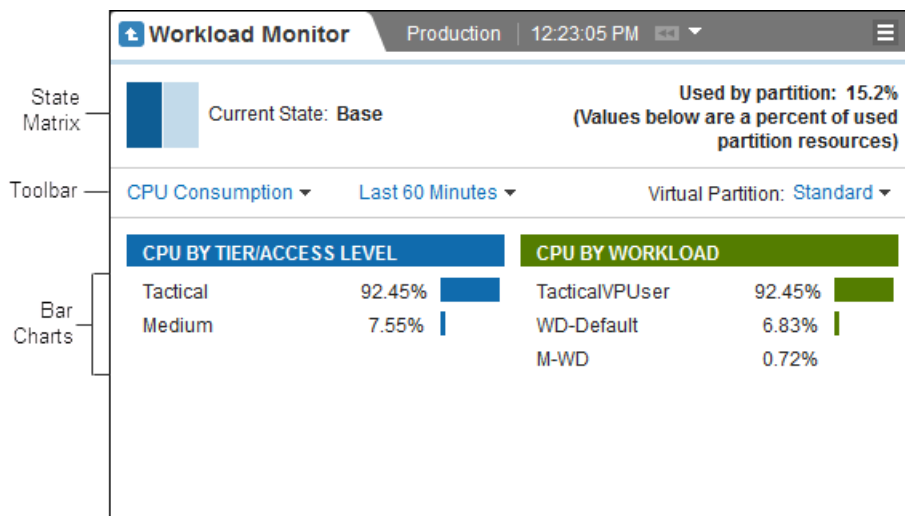
Virtual Partition List: In SLES 11, filters the system information for the selected virtual partition or all virtual partitions.

Bar Charts

Compare system resource consumption by tier/access level or by workload:

- **By Tier/Access Level:** Displays the percentage of partition CPU used by each tier/access level. Click a tier/access level to access the workload group view.
- **By Workload:** Displays the percentage of partition CPU used by each workload. Click a workload to access the workload details view.

Following is an example Distribution details view that you can see in SLES 11.



Workload Details View

The Workload details view helps you analyze individual workloads.

Toolbar

Time Frame Buttons

- **Cumulative:** Displays aggregated workload values for the selected sampling time frame.
- **Snapshot:** Displays current workload values.
- **Trends:** Displays sparkline graphs for the workload, using the trend interval set in **Settings**.

State Matrix: Shows the changes in state, planned environment, or health condition during the cumulative sampling time frame. In SLES 11, you can click it to see details about the current state and previous states.

Breadcrumbs

In SLES 10, the enforcement priority (EP) and allocation group (AG) are listed.

In SLES 11, the virtual partition, tier/access level, and workload are listed.

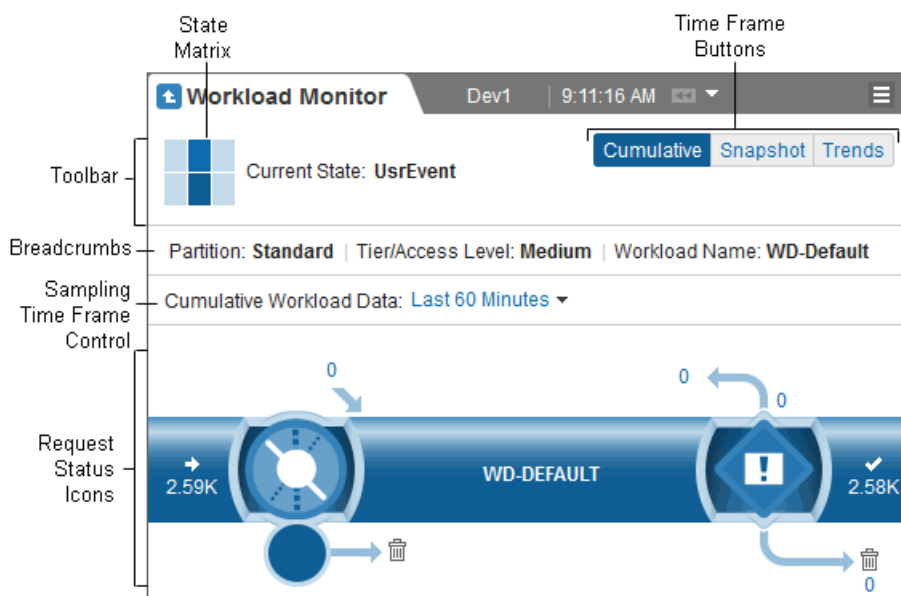
Sampling Time Frame Control

Filters the system information for the selected sampling time frame.

Request Status Icons

Shows the number of requests processed at each step, including at system filters and throttles. You can click the values associated with these icons to see more detail.

Following is an example Workload details view that you can see in SLES 11.



Trends View Metrics

Workload metrics are displayed for each workload on the Trends view. Trend intervals and past averages are set in **Settings**, and metrics are enabled and disabled in **Settings**.

Metric	Description	Operating System
Active Queries	Number of active queries	SLES 10 and 11
Avg CPU-SEC per Query	Average CPU time, in seconds, of all workload queries	SLES 10 and 11
Average Physical KB I/O per Query	Average physical I/O usage, in kilobytes, of all workload queries	SLES 11
Average Response Time and Delay Time	Average response time and delay time for all completed queries	SLES 10 and 11
Cumulative CPU	Cumulative CPU seconds per AMP	SLES 10 and 11
Decayed Active Queries	Total number of queries that have decayed due to CPU or I/O consumption (Timeshare workloads only)	SLES 11
Delayed Queries	Number of queries in the delay queue	SLES 10 and 11
Deferred Queries	Number of queries in deferred queue.	SLES 11
Impact CPU-SEC	Maximum amount of CPU time used by a workload compared to the total CPU used by all workloads on the system	SLES 10 and 11
Max CPU-SEC per Query	Maximum CPU time, in seconds, of all queries that completed	SLES 10 and 11
Maximum Physical KB I/O per Query	Maximum physical I/O usage, in kilobytes, for all queries in this workload	SLES 11
Maximum Response Time	Longest response time of all requests that completed	SLES 10 and 11
Tactical Queries with Exceptions	Total number of queries reassigned to another workload due to CPU or I/O per node tactical exception (Tactical workloads only)	SLES 11
Throughput	Number of queries arriving and completing every minute	SLES 10 and 11

Workload Group View

The Workload group view shows the group of workloads that belong to an AG or tier/access level.

- The AG Workload group view displays resource consumption as a percent of total CPU consumption and is accessed from the Distribution view by clicking an allocation group (for SLES 10).

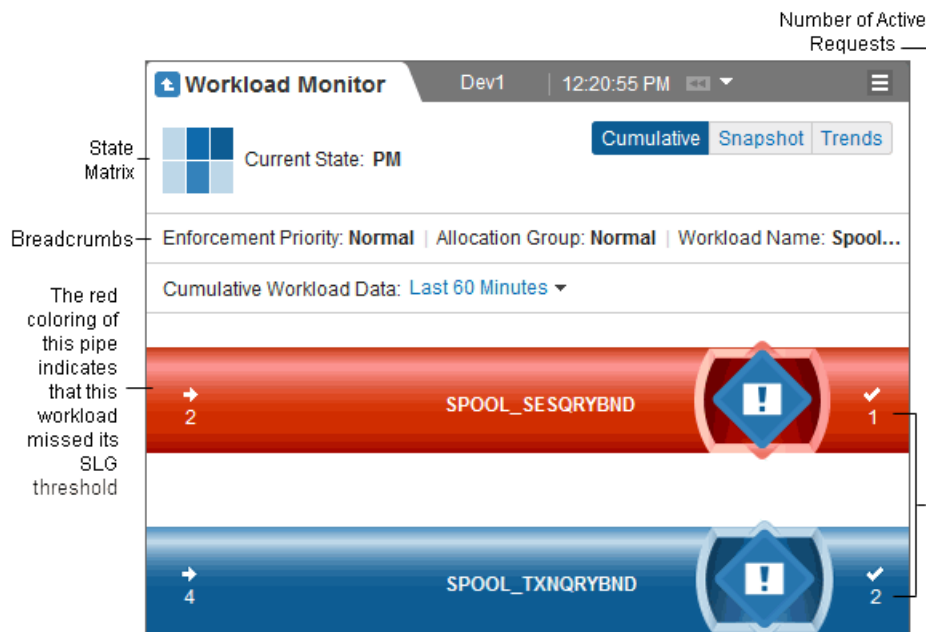
- The tier/access level Workload group view displays tier resource consumption as a percent of system or virtual partition and is accessed from the Distribution view by clicking a tier/access level (for SLES 11).

State Matrix

Shows the changes in state, planned environment, or health condition during the cumulative sampling time frame. In SLES 11, you can click it to see details about the current state and previous states.

Breadcrumbs

In SLES 10, the enforcement priority (EP), allocation group (AG), and workload are listed. In SLES 11, the virtual partition and tier/access level are listed.



Request Details View

The request details view displays statistics and information about the selected session. You can access this view by clicking a request icon and then one of the session rows.

Tabs

Provides important query details on the **Overview**, **SQL**, **Explain**, **Skew**, **Blocked By**, **Query Band**, **Defer**, and **Delay** tabs.

Query Details

Displays details of the selected query in sections that are specific to each tab.

▼ Manage Queries

Manages the query or session for Teradata systems with **Abort** , **Change Priority** , **Change Workload** , and **Release Query** options.

State Icons

The following icons appear in the portlet.

Icons	Name	Definition
×	Aborting	Query has been aborted, and changes are being rolled back.
▶▶	Active	Query is running.
⏸	Blocked	Query is waiting for a lock held by another query.
⏸	Delayed	Query is in a delay queue waiting to run.
⌚	Deferred	Query is in a deferred queue waiting to run.
🔄	Host-Restart	Host associated with this session is currently restarting.
⏸	Idle	No query is running.
?	Other	Query is in an unknown state.
🔍	Parsing	Query is being parsed before running.
⏸ QT	QTDelayed	Query delayed due to a queue table restriction.
↔	Response	Query is returning results to the user.
⏸ SES	SESDelayed	Query is delayed because a utility limit has been exceeded in FastLoad, MultiLoad, FastExport, or ARC.

Overview Tab

The **Overview** tab provides detailed information about key metrics for the selected session and its queries. The metric values provide a view of the query status on the system.

Query Information

Query Information	Description
State	Query state, such as active, blocked, terminate
Time in State	How long the query has been in the current state, displayed as <i>hh:mm:ss</i>
Total Duration	Total elapsed time it took for the query to execute once it was submitted
Spool Space	Amount of spool space the query is using

Query Information	Description
Hot AMP Spool	(Teradata Database 16.0 and later) Current spool value of the highest spool utilized AMP for the query. If no request is running on the session, no value displays.
Spool Skew	(Teradata Database 16.0 and later) Current spool skew for the query. If no request is running on the session, no value displays.
Temp Space	Amount of temp space the query is using
Request CPU	Total CPU seconds needed to run the query, in seconds
Impact CPU	CPU impact on the system based on the highest utilized AMP. (Teradata Database 15.0 and earlier) Calculated at the snapshot level. (Teradata Database 15.10 and later) Calculated at the request level.
Request I/O	Total number of disk I/Os performed
Request CPU Skew	(Teradata Database 15.10 and later) CPU skew for the current query
Request I/O Skew	(Teradata Database 15.10 and later) I/O skew for the current query
PJI	Ratio of the CPU milliseconds per I/O for the query, where a larger Product Join Index number indicates system performance degradation
Unnecessary I/O	All AMP I/O divided by all AMP CPU, displayed in milliseconds, to reveal large amounts of I/O occurring over a short period of time
Remote Data Imported	(Teradata Database 15.0 and later) Total bytes imported from a remote server for this query
Data Exported Remotely	(Teradata Database 15.0 and later) Total bytes exported to a remote server for this query

Workload Information

Workload Information	Description
Name	Name of the workload where the query is actively running
Method	Name of the workload management method in SLES 11. Available values are: <ul style="list-style-type: none"> • Tactical • SLG Tier (not used by IWM systems) • Timeshare
CPU Decay	Most severe level of CPU resource access restriction for queries in a Timeshare workload in SLES 11
CPU Exception Nodes	Number of nodes that exceeded the tactical CPU time exception criteria for the session in a Tactical workload in SLES 11
Classification Mode	How a query or session is assigned to a workload. Available values are: <ul style="list-style-type: none"> • Auto - Query is assigned automatically by TASM

Workload Information	Description
	<ul style="list-style-type: none"> Request - Query is assigned manually to a workload using Change Workload Session - Queries initiated in that session are assigned manually to a workload using Change Workload <p>This field is empty if this is not a DBC/SQL session or if Teradata Workload Management Category 3 is disabled.</p>
Virtual Partition	Name of the virtual partition in SLES 11
I/O Decay	Most severe level of I/O resource access restriction for queries in a Timeshare workload in SLES 11
I/O Exception Nodes	Number of nodes that exceeded the tactical I/O usage exception criteria for the session in a Tactical workload in SLES 11

Session Information



Session Information	Description
User	Name of the user that submitted the query
Proxy Username	(Teradata Database 15.0 and later) Name of the proxy user when a trusted session is used
Account	Account of the user that submitted the query
Source	Source details, such as application name, IP address, and host user name
Partition	Partition in which the query is running
Requests	Number of queries submitted by the session
Request Admission Time	Timestamp when the query was admitted to the system





SQL Tab

The **SQL** tab displays the SQL for the selected query and allows you to export the SQL. If the SQL is from a stored procedure on Teradata Database 15.10 and later, the name of the stored procedure displays. On Teradata Database 16.0 and later, the default database name of the session at the start of the request displays. If the SQL is from a stored procedure, the default database name at the time the stored procedure was compiled displays. This information is read-only. The tab is available only when a query is active, blocked, or delayed. For Teradata systems, the tab is available only when a query meets thresholds established by the Teradata Viewpoint Administrator.

Explain Tab

The **Explain** tab displays an abbreviated version of the Step statistics and Explain text that result from an Explain request in an SQL session and allows you to export Explain data. The tab is available only when a query meets certain thresholds established by the Teradata Viewpoint Administrator. The information is read-only. Each Explain step is uniquely identified with a number, where the background color of the number box indicates status.

If a query uses incremental planning,  appears with two gray bars next to the running steps, indicating that the total number of steps can change as additional steps are generated. When all steps are generated,  appears with three black bars next to the completed steps.

Step Information	Description
Step Number	<ul style="list-style-type: none"> Completed steps are at the top of the list and indicated by a black number box. Active steps are indicated by a pulsating number box (flashes blue). Steps to run are at the bottom of the list and indicated by a white number box.
Confidence Level Indicator Icon	 - No confidence in the estimate  - Low confidence in the estimate  - High confidence in the estimate  - High confidence in the estimate due to a join index
Estimated Time	Estimated execution time for the step
Estimated Rows	Estimated number of rows for the step
Actual Time	Actual CPU time consumed by the step, or blank if the step has not run
Actual Rows	Actual number of rows for the step, or blank if the step has not run

Blocked By Tab

The **Blocked By** tab displays details about other queries that are blocking the selected query. This information is read-only. The tab is available only when the selected query is blocked.

Blocked By Information	Description	14.10 and earlier	15.00 and later
Session ID	Unique session identifier of the query that is blocking	Available	Available
Username	Name of the user that is running the query that holds the lock	Available	Available
State Icon	Icon representing the current state of the query		Available
Blocking Type	Indicates the significance of the blocker. Available values are:		Available

Blocked By Information	Description	14.10 and earlier	15.00 and later
	<ul style="list-style-type: none"> Root Cause - Sessions that are not blocked, but are causing other sessions to be blocked Granted - Sessions that are blocked and are blocking other sessions because they have been granted a lock Waiting - Sessions that are blocked and are ahead of other blocked sessions in the queue 		
Blocking Count	Total number of sessions that are either directly or indirectly blocked by this session		Available
Blocking Time	Amount of elapsed time that this session has been blocking other sessions, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Impact	Blocked time summed together of each session that is either directly or indirectly blocked by this session, displayed as <i>d +hh :mm :ss</i>		Available
Blocking Locks	<p>A comma-separated list of the locks that are responsible for the blocking, for example: Granted write lock on table databaseName.tableName, Waiting on write lock on rowhash databaseName.tableName</p> <p>The information displays in the following order:</p> <ol style="list-style-type: none"> 1. Status of the lock causing the block (Waiting or Granted) 2. Mode (severity) of the lock involved in causing a block (exclusive, write, read, or access) 3. Type of object whose lock is causing the session described by the associated row to be blocked (database, table, row hash, table partition range, row hash partition range, row hash partition, and row key range) 4. Name of the database object over which a lock conflict is preventing the session from being granted a lock 5. Name of the table object over which a lock conflict is preventing the session from being granted a lock 		Available
Host	Host ID or LAN ID associated with the PE that processed the login request for the session	Available	
Lock Type	Type of lock. Type can be Exclusive, Read, Write, or Access	Available	
Status	Lock status. Status can be Waiting or Granted	Available	
Locked	Name of the locked object	Available	

Query Band Tab

The **Query Band** tab displays the profile, session, and transaction query bands for the selected query. This information is read-only. The tab is available only when a query band is included in the query.

Profile query bands are associated with Teradata user profiles and are automatically set on the session when that user logs in. Session query bands are set for an Analytics Database session.

Transaction query bands are set only for the current transaction. For more information, see *Teradata® Database Administration*.

Query Band Information	Description
Name	Name of the query band for the session or transaction
Value	Value of the query band for the session or transaction

Defer Tab

The Defer tab displays details about the rules that are deferring a query. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session
Defer Time	Duration of time request has been deferred by Arrival Rate Meter rules
User Name	User who submitted the query

Rule Information	Description
Rule Name	Name of rule causing request to defer
Rule Type	TASM type of rule causing request to defer
Overridable	Indicates if the Teradata DBA can abort or release the query

Delay Tab

The **Delay** tab displays details about all rules that are delaying a query. A scroll bar appears if there are more than two rules. This information is read-only.

Statistics Information	Description
Blocking Count	Number of consecutive times this session has blocked at least one other session

Statistics Information	Description
Delay Time	Duration of time request has been delayed by utility, system throttle, workload throttle, or workload group throttle rules
Utility Throttle	Indicates if request has been delayed by a utility throttle rule
System Throttle	Indicates if request has been delayed by a system throttle rule
Workload Throttle	Indicates if request has been delayed by a workload throttle rule
Workload Group Throttle	Indicates if request has been delayed by a workload group throttle rule

Rule Information	Description
Rule Name	Name of rule causing request delay
Rule Type	TASM type of rule causing request delay
Overridable	Indicates if the Teradata DBA can abort or release the request

Skew Tab

The **Skew** tab displays details about the level of skew in the query or session. The **Skew** tab does not display when the **Delay** tab is present.

Skew Information	Description
Highest	AMP with the highest CPU utilization or I/O count
2nd Highest	AMP with the second highest CPU utilization or I/O count
3rd Highest	AMP with the third highest CPU utilization or I/O count
Average	Average CPU utilization or I/O count across all AMPS
3rd Lowest	AMP with the third lowest CPU utilization or I/O count
2nd Lowest	AMP with the second lowest CPU utilization or I/O count
Lowest	AMP with the lowest CPU utilization or I/O count
Session Skew	Difference between the highest and the average values
Participating AMPs	Total number of AMPs participating for this session during the last session collection interval

Managing Queries and Sessions

In Teradata systems, you can manage queries and sessions to improve workload performance and balance system resources.

Abort

Abort the selected query or session

Change Priority

Change the priority of a query or session

Change Workload

Change the workload of the selected query or session

Release Query

Release the selected query from a delay queue

You must log in with a user ID that has permissions to abort, change workloads, or release queries. If you log out, close a portlet, or open a new portlet instance, you must log in again.

Aborting a Query or Session

For Teradata systems, you can abort a query or session that is blocking other queries or consuming too many resources.

1. Click the row of the query you want to abort.
2. Click ▼ to the right of the session number and select **Abort**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Select one of the following:
 - **Abort Query** to abort the selected query.
 - **Abort Session** to abort the selected query and log out of the session.
6. Click **Next**.
7. Click **Next** to confirm your selection.
8. Click **OK**.

Changing the Priority of a Query or Session

For Teradata systems, you can change the priority of a query or session to allow higher priority queries to run or balance session resources.

This option is only available when workloads are not enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Priority**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Do one of the following to change the priority of sessions:
 - Select an account string from the list of accounts that have been assigned to the user.
 - Type an account string.
6. [Optional] Select the check box to use the account string as the default for the selected session. The priority will be changed for the selected query and all subsequent queries in the current session.
7. Click **Next**.
8. Click **Next** to confirm your selection.

Changing the Workload of a Query or Session

For Teradata systems, you can change the workload of a query or session to allow higher priority workloads to run or to balance workload resources.

This option is only available when workloads are enabled and the system being monitored supports this feature.

1. Click the row of the query you want to change.
2. Click ▼ to the right of the session number and select **Change Workload**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Select a different workload from the list and click **Next**.
5. Click **Next** to confirm your selection.
6. Click **OK**.

Releasing a Query

For Teradata systems, you can release a query from the delay queue for immediate processing.

1. Click the row of the query you want to change.
2. Click ▾ to the right of the session number and select **Release Query**.

Note:

If JWT integration with Vantage is enabled (versions 16.50.05.00 and later), you will not be prompted to log in. A Disconnect option is available, which will allow the user to login using different credentials.

Note that once disconnected, the user will continue to be prompted to enter credentials. This setting is saved per portlet instance; a new portlet instance will bypass the login dialog.

3. Log in to Analytics Database, if prompted.
4. Click **Next**.
5. Click **Next** to confirm your selection.

Settings View

The **Settings** view allows you to customize the portlet.

The following tabs allow you to customize your settings.

Systems

Choose a system to monitor.

Metrics

In SLES 10, select metrics for workloads in the Tactical, Priority, Normal, or Background enforcement priorities.

In SLES 11, select metrics for workloads in the Tactical, SLG Tier, or Timeshare management methods.

Trend Interval


Set the time frame displayed for the metric graphs.

Past Averages

Choose the time frame for which a historical average is calculated. This average is shown by the skyline plots in the details views.

Managing Default Settings

In the **Settings** view, you can set or clear user-defined default settings for the portlet.


1. Click  in the portlet frame and select **Settings**.
2. Set or clear the user-defined default settings for the portlet:

Option	Description
Save as Default	Click to use the current user-defined settings as the default configuration each time you add this portlet to a portal page.
Clear Defaults	Click to clear all user-defined default settings created for this portlet.

3. Click **Close**.
4. Click **OK**.

Setting Metrics for Enforcement Priority or Tier/Access Level


For each enforcement priority group in SLES 10 or tier/access level in SLES 11, you can select and order metrics for display in the workload details view. Choose from a defined list of metrics.

1. Click  in the portlet frame and select **Settings**.
2. Click one of the metrics tabs for an enforcement priority group or tier/access level.

The display pane shows how metric rows appear in the portlet.

3. Do any of the following to change the way metric rows display:


Option	Description
Add a metric row	Select the check box of the metric in the Select metrics for display pane. That metric is added to the display pane beneath the rows currently displayed.

Option	Description
Remove a metric row	Clear the check box of the metric in the Select metrics for display pane.
Change metric row order	Use  to drag a metric already in the display pane to a new location in the pane (up or down).

- Click **OK**.

Setting the Workload Trend Interval


Use the **Trend Interval** tab in the **Settings** view to set the length of time represented in sparkline graphs.

- Click  in the portlet frame and select **Settings**.
- Click the **Trend Interval** tab.
- Select an interval value from the list.
The **Sample view** pane shows an example of how data would appear when viewed at the selected interval.
- Click **OK**.

Setting Past Averages

Specify the time frame used to calculate and display the shaded skyline graph in the details view.

Each segment of the skyline graph is the average of data samples taken each week at the same time and day of the week.

- Click  in the portlet frame and select **Settings**.
- Click the **Past Averages** tab.
- Enter the number of weeks up to 99.
- Click **OK**.

Metric Sources

Metric Sources Overview

Metric sources expand on the portlet metric descriptions in this guide and are for database administrators and Viewpoint administrators to understand how Viewpoint metrics are derived.

Metric sources are database data used by Viewpoint to perform calculations that produce the metrics shown in the Viewpoint portlets. Typically, the **Source** column indicates the **Field/Column Name** in the *Teradata® Database Application Programming Reference*.

Metric sources from Viewpoint, such as data collectors, are not included.

Query Portlets

The metric sources are for the following Viewpoint Query portlets:

- **My Queries**
- **Query Monitor**
- **Query Spotlight**

Metric descriptions are listed in the documentation for the Queries portlets.

Teradata System Metrics

Session Metrics

Metric	Source	Calculation	Type
CPU Use	Δ CPU: derived SessionRate: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 1 > Group I Data Fields > SessionRate CPUCount: Teradata PM/API > System PMPC APIs > Monitor Physical Config > Statement 2 > CPUCount	$(\Delta \text{ CPU} / (\text{SessionRate} * \text{total CPUCount on nodes with amps})) * 100$	Percent
Δ CPU	AMPCPUSec: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > AMPCPUSec	Current session's AMPCPUSec - previous sampled AMPCPUSec	Number
Δ I/O	AMPIO: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > AMPIO	Current session's AMPIO - previous sampled AMPIO	Number

Metric	Source	Calculation	Type
Hot AMP Spool	ReqHotAmpSpool: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group VIII Data Fields > ReqHotAmpSpool	(Teradata 16.0 and later) ReqHotAmpSpool	Number
Impact CPU	HotAmp1CPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > HotAmp1CPU UpAMPCount: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > UpAMPCount ReqHotAmpCPU: Teradata PM/API (15.10) > System PMPC APIs > Monitor Session > Statement 2 > Group VII Data Fields > ReqHotAmpCPU ReqInvolvedAMPCnt: Teradata PM/API (15.10) > System PMPC APIs > Monitor Session > Statement 2 > Group VII Data Fields > ReqInvolvedAMPCnt	(Teradata 15.0 and earlier) (HotAmp1CPU * UpAMPCount) (Teradata 15.10 and later) (ReqHotAmpCPU * ReqInvolvedAMPCnt)	Number
Partition	PartName: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > PartName	PartName	String
PJI	RequestAmpCPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpCPU RequestAmpl/O: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpl/O	$1000 * (\text{RequestAmpCPU} / \text{RequestAmpl/O})$	Number
Proxy Username	ProxyUser: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group V Data Fields > ProxyUser	ProxyUser	String
Request Count	ReqCount: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 1 > Group I Data Fields > ReqCount	ReqCount	Number
Request CPU	RequestAmpCPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpCPU	RequestAmpCPU	Number
Request CPU Skew	RequestAmpCPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpCPU ReqInvolvedAMPCnt: Teradata PM/API (15.10) > System PMPC APIs > Monitor Session	(Teradata 15.10 and later) $100 * (1 - (\text{RequestAmpCPU} / \text{ReqInvolvedAmpCnt}) / \text{ReqHotAmpCPU})$	Percent

Metric	Source	Calculation	Type
	> Statement 2 > Group VII Data Fields > ReqInvolvedAMPCnt ReqHotAmpCPU: Teradata PM/API (15.10) > System PMPC APIs > Monitor Session > Statement 2 > Group VII Data Fields > ReqHotAmpCPU		
Request I/O	RequestAmpI/O: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpI/O	RequestAmpI/O	Number
Request I/O Skew	RequestAmpI/O: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpI/O ReqInvolvedAMPCnt: Teradata PM/API (15.10) > System PMPC APIs > Monitor Session > Statement 2 > Group VII Data Fields > ReqInvolvedAMPCnt ReqHotAmpIO: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group VII Data Fields > ReqHotAmpIO	(Teradata 15.10 and later) $100 * (1 - (\text{RequestAmpI/O} / \text{ReqInvolvedAmpCnt}) / \text{ReqHotAmpIO})$	Percent
Snapshot CPU Skew	AvgAmpCPUsec: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > AvgAmpCPUsec HotAmp1CPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > HotAmp1CPU	$100 * (1 - (\text{AvgAmpCPUsec} / \text{HotAmp1CPU}))$	Percent
Snapshot I/O Skew	AvgAmpIOCnt: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > AvgAmpIOCnt HotAmpIO: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > HotAmpIO	$100 * (1 - (\text{AvgAmpIOCnt} / \text{HotAmpIO}))$	Percent
Spool	Request_AmpSpool: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > Request_AmpSpool	Request_AmpSpool	Number
Spool Skew	Request_AmpSpool: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > Request_AmpSpool ReqHotAmpSpool: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group VIII Data Fields > ReqHotAmpSpool ReqInvolvedAMPCnt: Teradata PM/API > System PMPC APIs > Monitor Session	(Teradata 16.0 and later) $100 * (1 - (\text{Request_AmpSpool} / (\text{ReqHotAmpSpool} * \text{ReqInvolvedAMPCnt})))$	Percent

Metric	Source	Calculation	Type
	> Statement 2 > Group VII Data Fields > ReqInvolvedAMPCnt		
Start	RequestStartTime: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestStartTime	RequestStartTime	Number
State	PEState: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > PESTate AMPState: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > AMPState	See MONITOR SESSION Response Combinations in <i>Teradata® Database Application Programming Reference</i> .	String
Temp Space	TempSpace: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > TempSpace	TempSpace	Number
Unnecessary I/O	RequestAmpl/O: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpl/O RequestAmpCPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpCPU	RequestAmpl/O / (RequestAmpCPU * 1000)	Number
Workload	TDWM.RuleDefs	rulename	String

Account String and User Metrics

Metric	Source	Calculation	Type
Δ Request Count	ReqCount: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 1 > Group I Data Fields > ReqCount	Difference between request counts between samples	Number

Blocker Metrics

Metric	Source	Calculation	Type
Blocking Count	Blk_x_SessNo: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > Blk_x_SessNo	Count of blocking session	Number
Blocking Type	Blk_1_SessNo: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > Blk_1_SessNo	If blocker session and there is no Blk_1_SessNo, then root cause. Otherwise, if any Blk_Status has a GRANTED status,	Number

Metric	Source	Calculation	Type
	Blk_x_Status: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > Blk_x_Status	then type is GRANTED. Otherwise, WAITING.	

Utility Metrics

Metric	Source	Calculation	Type
LSN	LSN: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > LSN	LSN	Number
Utility	LSN: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > LSN	LSN > 0	String
AMP CPU	RequestAmpCPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpCPU	RequestAmpCPU	Number
AMP I/O	RequestAmpI/O: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group III Data Fields > RequestAmpI/O	RequestAmpI/O	Number

Vproc Metrics

Metric	Source	Calculation	Type
PE	LogonPENo: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > LogonPENo	LogonPENo	Number
PE CPU	PECPUsec: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group I Data Fields > PECPUsec	PECPUsec	Number
Skewed CPU	AvgAmpCPUsec: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > AvgAmpCPUsec HotAmp1CPU: Teradata PM/API > System PMPC APIs > Monitor Session > Statement 2 > Group II Data Fields > HotAmp1CPU	$100 * (1 - (\text{AvgAmpCPUsec} / \text{HotAmp1CPU}))$	Percent

Additional Information

Audience

This guide is intended for use by:

- Database administrators
- System administrators
- Application developers and users

Changes and Additions

Date	Release	Description
April 2023	23.04.00.00	<ul style="list-style-type: none"> • Included metric analysis for one hour interval. • Included lockout details. • Included deleted users audit trial details. • Included CPU I/O Wait metric for Node and Vproc Alert Types.
October 2022	22.10.00.00	<ul style="list-style-type: none"> • Introduced the option to send TASM alerts to database (business) users directly, and added an example. • Added the notification alert configuration to notify the upcoming expiry dates of certificates in the Teradata Viewpoint key store. • Added the capability of User Details View to show the user details audit and roles audit information.
May 2022	17.10.02.00	<ul style="list-style-type: none"> • Added the time unit of Move Duration in the exported Workflow Actions metrics • Added information about the disk space alert notification on the Viewpoint dashboard • Added a note about the bridge information for the QueryGrid Manager 2.18 • Described the detailed view of any SQL log • Updated Status metric description for node and Vproc alert types
February 2022	17.10.01.00	<ul style="list-style-type: none"> • Added Defer Time metric under session alert type • Added ENABLED CPU metric, and added tooltip content for other metrics in Viewpoint dashboard • Added two new options under the Session Alert to cancel blocked and blocking sessions • Updated alert configuration under Backup portlet • Updated utility classification type to allow meters • Added a check-box "Actual Client IP Address" under Request Source Classification

Date	Release	Description
		<ul style="list-style-type: none"> Updated the supported web browsers version for Viewpoint portal
October 2021	17.10.00.00	<ul style="list-style-type: none"> Updated Daily Jobs in the Configuring Jobs section Updated Installing a Certificate from a Certificate Signing Request under the HTTPS Connections section
July 2021	16.50.06.00	Administration Portlets updated with information on VP_User_Manager role
April 2021	16.50.05.00	<ul style="list-style-type: none"> Added Alerts information in Stats Manager. Arrival Rate Monitor updates made in Workload Designer, Workload Monitor Portlet, My Queries, Query Groups, Query Monitor, and Viewpoint Dashboard along with updated images. Bridge information updated in Query Monitor and Completed Queries. Updates made in Administration Portlets under Metrics for Alert Types.
January 2021	16.50.04.00	<ul style="list-style-type: none"> Updated the Disk usage option in the Monitored Systems portlet to specify Viewpoint DB Disk Usage, and changed the functionality to reflect total disk usage. For TMSMonitor, HTTPS is enabled by default and HTTP connection is no longer supported.
September 2020	16.50.03.00	<ul style="list-style-type: none"> Added Cumulative Error Count to Stats Manager metrics. Added instructions to configure cleanup of space usage collector tables.
June 2020	16.50.02.00	<ul style="list-style-type: none"> Removed references to Presto as it is no longer supported. Added support for disabling HTTP for TMSMonitor on a managed system. Added support to generate Elastic Performance On Demand interval reporting at ten-minute intervals.
January 2020	16.50.00.00	<ul style="list-style-type: none"> Added table of Data Collector Sample Rates for Teradata Machine Learning Engine. Added Supported Web Browsers for Viewpoint Portal and Supported Web Browsers for Viewpoint Mobile sections.
November 2019	16.20.24.00	Added support for JWT integration with Vantage.
August 2019	16.20.23.08	Added TD_Webservices. Removed Presto from Monitored Systems portlet.
April 2019	16.20.23.06	Added new user alert selection in the Monitored Systems portlet. Changed viewable information in MAPS Manager portlet.
February 2019	16.20.23.04	Added new selection tool to MAPS Manager portlet.
January 2019	16.20.23.03	<ul style="list-style-type: none"> Added Machine Learning Engine metrics. Added Session view to Machine Learning Engine.

Date	Release	Description
December 2018	16.20.23.02	Updated alert metrics.
October 2018	16.20.23.01	<ul style="list-style-type: none"> Added information about reclaiming disk space. Added information about converting pre-expansion workflows.
September 2018	16.20.23.00	<ul style="list-style-type: none"> Added information about elastic TCore configuration. Added information about pre-expansion workflows.
July 2018	16.20.12.02	Added information about Teradata Machine Learning Engine and system selection. Updated portlet controls.
June 2018	16.20.12.01	Added overview of Viewpoint RESTful APIs and support for certificates using Subject Alternative Names.
May 2018	16.20.12.00	Added group alerting user sessions into a single email.
April 2018	16.20.00.05	Added Teradata Machine Learning Engine. Enhanced System Workload Report.
February 2018	16.20.00.03	Updated HTTPS connections.
November 2017	16.20.00.01	Removed External Content portlet.
October 2017	16.20	Initial release.

Teradata Links

Link	Description
https://docs.teradata.com/	Search Teradata Documentation, customize content to your needs, and download PDFs. Customers: Log in to access Orange Books.
https://support.teradata.com	Helpful resources in one place: <ul style="list-style-type: none"> Support requests Account management and software downloads Knowledge base, community, and support policies Product documentation Learning resources, including Teradata University
https://www.teradata.com/University/Overview	Teradata education network
https://support.teradata.com/community	Link to Teradata community

Related Documentation

Title	Publication ID
<i>Teradata® Viewpoint Installation, Configuration, and Upgrade Guide for Customers</i> Describes how to install Viewpoint software, configure settings, and upgrade a Teradata Viewpoint server.	B035-2207
<i>Teradata® Data Lab User Guide</i> Describes how to use Teradata Data Lab portlets.	B035-2212
<i>Teradata® Viewpoint and Teradata® Data Lab API Reference Guide</i> Describes how to perform various Viewpoint and Data Lab tasks programmatically as an alternative to using the Viewpoint portal.	B035-2215
<i>Teradata® Database Administration</i> Describes how to administer the Teradata Database.	B035-1093
<i>Teradata® Database Application Programming Reference</i> Describes various types of APIs and how they can be used to interface with your custom applications.	B035-1090