

# **Teradata Vantage™ - Advanced SQL Engine Release Summary**

---

Release 17.10

July 2021

# Copyright and Trademarks

Copyright © 2000 - 2021 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
	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.
	Indicates a hazardous situation which, if not avoided, could result in minor or moderate personal injury.
	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 law, the information available from the Teradata Documentation website or contained in Teradata information products is provided on an "as-is" basis, without warranty of any kind, either express or implied, including the implied warranties of merchantability, fitness for a particular purpose, or noninfringement.**

The information 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 information 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 this information at any time without notice.

## 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](mailto: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.

# Contents

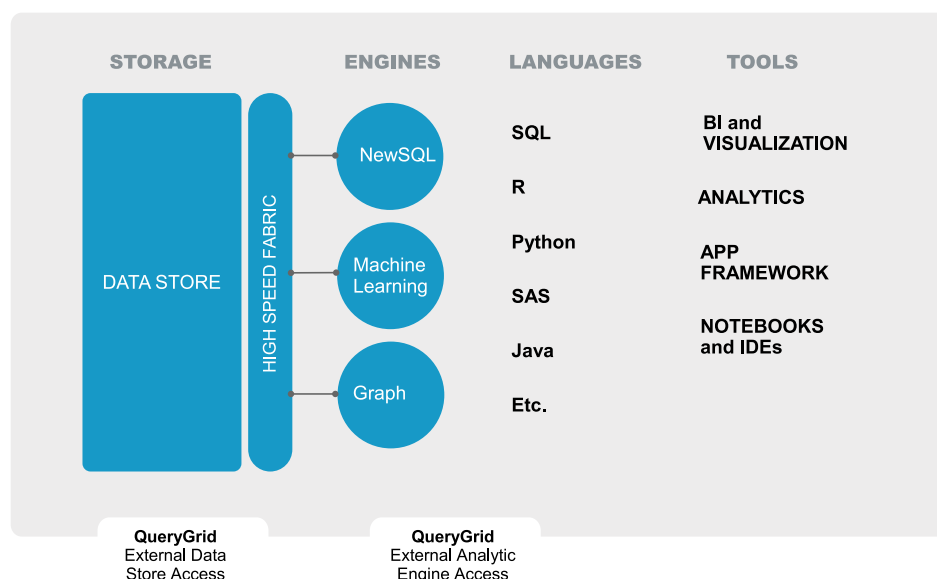
<b>Chapter 1: Introduction</b>	<b>4</b>
What is Teradata Vantage?	4
Release Summary Overview	5
Analytics	6
Industry Compatibility	6
Quality and Supportability	6
Simplicity and Ease of Use	7
<b>Chapter 2: Release 17.10 Features</b>	<b>8</b>
Arrival Rate Meters and Throttles	8
CheckTable Enhancements	8
Diagnostic Tool tdgssgetinfo	9
JSON Web Token Enhancements	9
Modify TDGSS Without a TPA Reset	10
Native Object Store Enhancements	11
Performance Enhancement	12
TLSv1.2 Support for Client Connections to SQL Engine and Cipher Suite Enhancements	12
Table Rebuild Enhancements	13
<b>Chapter 3: Release 17.05 Features</b>	<b>15</b>
Native Object Store Enhancements	15
WRITE_NOS for Native Object Store	15
<b>Chapter 4: Release 17.00 Features</b>	<b>17</b>
Customer Data Space Tools	17
DBQL Enhancements	18
Ferret SHOWBLOCKS Per AMP Information	19
Improved Resource Usage Monitoring for UDFs	19
Native Object Store	19
Security Enhancements in Release 17.00	21
TD_JSONSHRED Table Operator	21
UDF Search Path	22
<b>Appendix A: Teradata Vantage Restricted Words</b>	<b>23</b>
<b>Appendix B: Additional Information</b>	<b>25</b>

# Introduction

## What is Teradata Vantage?

Teradata Vantage™ is our flagship analytic platform offering, which evolved from our industry-leading Teradata® Database. Until references in content are updated to reflect this change, the term Teradata Database is synonymous with Teradata Vantage.

## Teradata Vantage



© 2019 Teradata

teradata.

The diagram illustrates the initial release of Vantage, which enabled our customers to leverage all of their data, all of the time, so they can analyze anything, deploy anywhere, and deliver analytics that matter most to their business. At a high-level, Vantage consists of the following components.

### Storage

**Datastore:** The file system that helps isolate Vantage from hardware platform dependencies and supports the creation and maintenance of database tables under the direction of Vantage. The file system stores data in physical format.

## High-speed Fabric

The high-speed fabric ties the processing technologies together. On Teradata® IntelliFlex®, it is the BYNET on InfiniBand, running the BYNET protocol. On a public cloud deployment such as Amazon Web Services (AWS) and Microsoft Azure (Azure), it is the network switching fabric used by that environment.

## Engines

- **Advanced SQL Engine (was NewSQL Engine):** It is a core capability of Teradata Vantage, based on our best-in-class Teradata Database. Advanced SQL refers to the ability to run advanced analytic functions beyond that of standard SQL. At a high-level, the Advanced SQL Engine:
  - Performs SQL-based analytics as was possible with Teradata Database
  - Provides additional advanced analytics capability such as 4D Analytics, path and pattern analysis, scoring, and more
  - Invokes Machine Learning Engine and Teradata Graph Engine functions
  - Executes additional query optimization functions as existed in Teradata Database, such as Optimizer and Workload Management
- **Machine Learning and Graph Engines:** These optional Engines provide the ability to run over 180 built-in functions, previously available from Teradata Aster®. The functions include math and statistics, text analysis, sentiment analysis, clustering, path and patterns, decision trees, graph, and so on. They are invoked through the Advanced SQL Engine, using standard SQL.

## Languages

As with the language support provided by Teradata Database, Vantage provides support for various languages including SQL, R, Python, SAS, and Java.

## Tools

Vantage provides access to:

- BI visualization tools including Tableau, Qlik, and IBM Cognos
- Analytic tools including Dataiku and SAS
- Notebooks and IDEs including RStudio, Jupyter, and Teradata Studio™
- Application framework such as Teradata® AppCenter

## Release Summary Overview

This document provides an overview of the features introduced in this Teradata Vantage Advanced SQL Engine release, including benefits and considerations for using each feature. It includes references to other documents for more details on each feature.

This material can help you with the following tasks:

- Determine whether to upgrade to a new release.

- Plan for upgrading to the new release.
- Prepare for and introduce new features to end users.

For details on Advanced SQL Engine hardware and software requirements, see *Teradata Vantage™ - Advanced SQL Engine Release Definition*, B035-1725.

This section lists major features according to their strategic category.

## Analytics

This strategic category includes features that:

- Address capabilities to support new analytical applications.
- Improve performance of analytics.
- Support new data types and data sources.

### Release 17.00

[TD\\_JSONSHRED Table Operator](#)

## Industry Compatibility

This strategic category includes features that:

- Improve migration from other databases.
- Improve areas where data and applications can be moved without impact.
- Make Teradata Vantage easier to use.
- Provide security features and functionality that comply with users' security policies and standards.

### Release 17.00

[Security Enhancements in Release 17.00](#)

### Release 17.10

[JSON Web Token Enhancements](#)

[Modify TDGSS Without a TPA Reset](#)

[Native Object Store Enhancements](#)

[TLSv1.2 Support for Client Connections to SQL Engine and Cipher Suite Enhancements](#)

[Table Rebuild Enhancements](#)

## Quality and Supportability

This strategic category includes features that:

- Reduce unplanned downtime.

- Reduce restarts.
- Reduce failures.
- Reduce issues using Teradata product.
- Use CLCAs to target certain areas for implement.

**Release 17.00**[DBQL Enhancements](#)[Ferret SHOWBLOCKS Per AMP Information](#)[Improved Resource Usage Monitoring for UDFs](#)[UDF Search Path](#)**Release 17.10**[Arrival Rate Meters and Throttles](#)[CheckTable Enhancements](#)[Diagnostic Tool tdgssgetinfo](#)

## Simplicity and Ease of Use

This strategic category includes features that:

- Leverage new tooling to attract new development users and address new markets and new applications.
- Deliver products that are easy to use and simple to implement.
- Provide new tunables to provide better control over features.

**Release 17.00**[Customer Data Space Tools](#)

## Release 17.10 Features

### Arrival Rate Meters and Throttles

Arrival Rate Meters (ARMs) are TASM rules that provide a way to regulate the flow of requests admitted by TASM into the system. An ARM rule specifies the maximum number of requests per time unit that TASM will admit into the system. Qualification criteria can be used to target specific types of requests.

TASM's filter and throttle processing sequence for a SQL request is enhanced to check for ARM. New Workload Management ruleset APIs allow you to automate the creation of a system throttle or arrival rate meter without using Viewpoint Workload Designer.

#### Benefits

- Prevents a predictable request surge from flooding the TASM delay queue and impacting service time of subsequent queries.
- Limits resource-intensive requests to a specific rate.
- Prevents you from submitting many extremely complex queries frequently.

#### Considerations

- Use arrival rate meters carefully after an evaluation period in warning mode.

#### Additional Information

For more information, see *Teradata Vantage™ - Application Programming Reference*, B035-1090 and *Teradata Vantage™ - Workload Management User Guide*, B035-1197.

### CheckTable Enhancements

CheckTable is enhanced to:

- Report dictionary errors in batch mode.
- Optionally skip the dictionary check.
- Optionally skip the user table check.
- Use more parallel sessions for user table checking.

#### Benefits

- Provides the capability of achieving high concurrency of table checks when executed on large systems.

#### SQL Changes

New keywords are added to the CheckTable syntax:



- MAXSESSIONS
- SCOPE { DICTIONARY | USER | ALL }

### Additional Information

For more information, see the CHECK command in *Teradata Vantage™ - Database Utilities*, B035-1102.

## Diagnostic Tool tdgssgetinfo

tdgssgetinfo is a new diagnostic tool in 17.10 that collects and displays information useful in determining the health of the TeraGSS or TDGSS installed on the system.

### Benefits

tdgssgetinfo provides better diagnostics capabilities when errors are encountered. It displays:

- Details about the system from which the tool is run.
- Information returned in the fourth parameter from a call to tdgss\_configure().
- The currently active version, as well as all of the installed versions on the system.
- The available security mechanisms and the locations of the configuration files.

### Additional Information

For more information, see *Teradata Vantage™ - Advanced SQL Engine Security Administration*.

## JSON Web Token Enhancements

The JSON Web Token (JWT) mechanism is enhanced in Release 17.10:

- The JWT mechanism now dynamically updates JSON Web Keys (JWKs). When logging on to Teradata Vantage using JWT, an Identity Provider (IdP) signs the token using its private key and the Teradata server verifies the token's signature using the corresponding public key. If the IdP rotates the private keys the Teradata Gateway now automatically updates the public keys.
- The JWT mechanism now accepts JWT logons from third-party applications. For example, a user logs into a web app from a browser. The web app federates the logon to the customer's IdP. If the user then connects to Teradata Vantage, the web app provides the JWT token to the database to successfully complete the logon.

### Benefits

- Automatic JWK update during logon.
- Allows logons from third-party applications.
- JWKs are cached for improved performance.

## Considerations

- JWT key rotation and third-party logons are disabled by default. To enable these features, update the configuration file and execute the `run_tdgssconfig` command. No database reset is needed.

## Additional Information

For more information about security, see *Teradata Vantage™ - Advanced SQL Engine Security Administration*, B035-1100.

## Modify TDGSS Without a TPA Reset

Previously, when the TDGSS configuration changed a TPA reset was required for the new values in the TDGSSCONFIG GDO to take effect. Now, the following can be modified without a TPA reset:

- Any attribute or property whose name begins with "Ldap" for KRB5 and LDAP
- MechanismEnabled property for KRB5, LDAP, JWT, and PROXY
- AuthorizationSupported property for KRB5 and LDAP
- LDAP Service ID and password with no impact to user LDAP logons
- The following properties in the PROXY mechanism:
  - CertificateFile
  - PrivateKeyFile
  - PrivateKeyPassword
  - PrivateKeypasswordProtected
  - CACertFile
  - CACertDir
  - SigningHashAlgorithm
- Any JWT mechanism property whose name begins with "JWT"
- All canonicalizations including the lightweight authorization structures

Additionally, `tdgsstestcfg` is a new tool to test configuration changes before making them permanent with `run_tdgssconfig`.

## Benefits

- Decreases downtime previously caused by mechanism property reconfiguration.
- Simplifies steps when modifying mechanism properties.
- The `run_tdgssconfig` tool informs you when a `tpareset` is required.

## Considerations

The following configuration changes still require a `tpareset`:

- Changes to any mechanism property not mentioned above require a `tpareset`
- QoP configuration

- Local or global policy configuration, including service name changes
- TDNEGO and SPNEGO

### Additional Information

For more information about security, see *Teradata Vantage™ - Advanced SQL Engine Security Administration*, B035-1100.

## Native Object Store Enhancements

NOS enhancements for 17.10 are:

- SQL Syntax
  - Introduced new system-level authorization object security model.
  - Reduced minimal DDL syntax used for defining the NOS foreign table schemas.
  - Reduced minimal DQL statements used for performing queries.
- Schema Discovery
  - Removed manual and error prone schema discovery activities.
  - Automated the discovery of file data formats, field delimiter, column names, column data types and sizes, and object store path definition.
- Error Diagnostics
  - Expanded definitions providing more insight into error and warning messages.
  - Expanded the Error Table framework to NOS foreign tables.
- Performance
  - Simplified CSV processing aligned with Parquet level performance.
  - Reduced spool used for processing queries.
  - Improved path filtering performance.
  - Optimized TOP N query processing.
- READ\_NOS now supports Parquet data in addition to CSV and JSON.

### Benefits

- Provides simpler SQL syntax for reading and writing of data living on object storage.
- Reduces the level of familiarity with the data living on object storage and accelerates data exploration.
- Provides better diagnostics capabilities when errors are encountered.
- Provides overall performance improvements.

### SQL Changes

- RETURNTYPE for READ\_NOS has an additional argument, NOSREAD\_SCHEMA. NOSREAD\_SCHEMA returns the schema from sampling the first 16MB of data from the latest timestamp file from the LOCATION URL.

- The CREATE AUTHORIZATION syntax is updated with a USING clause to allow Azure Key Vault authentication.

## Additional Information

For practical guidance and code samples, see [Teradata Vantage™ Native Object Store - Advanced SQL Engine 17.0 \(Orange Book\)](#). You must log in to view the Orange Book.

For more information, see:

- *Teradata Vantage™ - Native Object Store Getting Started Guide*, B035-1214
- *Teradata Vantage™ - SQL Data Definition Language Syntax and Examples*, B035-1144
- *Teradata Vantage™ - SQL Data Manipulation Language*, B035-1146
- *Teradata Vantage™ - SQL Operators and User-Defined Functions*, B035-1210

## Performance Enhancement

Bloom Filter optimization is added to classical hash join algorithms for Release 17.10. A Bloom Filter is a compact metadata structure on the join column of one relation of a hash join step. The Bloom Filter is used while reading the other relation of the join as an existence check or test. Specifically, the build side of a classical hash join produces a bitmap spool based on the join column while building the hash table. This bitmap spool is passed down to the probe side of the join to be able to filter rows while the probe side is being read.

### Benefits

This feature can improve the performance of some joins.

### Considerations

Bloom Filter is enabled by default.

## TLSv1.2 Support for Client Connections to SQL Engine and Cipher Suite Enhancements

TLSv1.2 is supported between clients and the database server in Release 17.10. This includes new tools and utilities:

- **gtwcontrol** - New options to enable or disable TLSv1.2 and trace the protocol.
- **gtwglobal** - New options to update the TLS configuration from the gateway TLS configuration.
- **nodenames** - Displays the network interfaces that a node is known by. The node names are discovered by querying DNS. This information is helpful when you generate a Certificate Signing Request (CSR) because nodenames provides the common name (CN) and subject alternative names (SANs) that are used in the CSR.
- **tlsutil** - A tool used to obtain and install signed certificates and private keys that are required for TLS. tlsutil uses the nodenames command internally.

The cipher suite has been updated.

- The following ciphers are included by default:
  - TLS\_AES\_256\_GCM\_SHA384
  - TLS\_CHACHA20\_POLY1305\_SHA256
  - TLS\_AES\_128\_GCM\_SHA256
  - ECDHE-ECDSA-AES256-GCM-SHA384
  - ECDHE-RSA-AES256-GCM-SHA384
  - DHE-RSA-AES256-GCM-SHA384
  - ECDHE-ECDSA-CHACHA20-POLY1305
  - ECDHE-RSA-CHACHA20-POLY1305
  - DHE-RSA-CHACHA20-POLY1305
  - ECDHE-ECDSA-AES128-GCM-SHA256
  - ECDHE-RSA-AES128-GCM-SHA256
  - DHE-RSA-AES128-GCM-SHA256
  - AES256-GCM-SHA384
  - AES128-GCM-SHA256
- Customers may change the ciphers as needed.

## Benefits

- TLSv1.2 is an industry standard protocol to secure network traffic between client and server.
- SQL Engine Release 17.10 system can be configured to prevent client applications from logging on if they are not using TLSv1.2. As a default configuration, SQL Engine Release 17.10 is capable of supporting TLSv1.2. Certificates must be configured appropriately for SQL Engine to take advantage of the TLSv1.2 feature.

## Considerations

- Client applications must be updated to Teradata Tools and Utilities (TTU) Release 17.10 to use TLSv1.2.
- DNS setup is required to use the `tlsutil` and `nodenames` tools.

## Additional Information

- *Teradata Vantage™ - Advanced SQL Engine Security Administration*, B035-1100
- *Teradata Vantage™ - Database Utilities*

## Table Rebuild Enhancements

Tables from different databases can now be rebuilt in parallel. The number of tables that can be rebuilt in parallel increased from 6 to a maximum of 255. The default number of parallel streams remains 6.

## Benefits

- The Rebuild Utility processes tables from different databases in parallel, increasing the overall rebuild performance.

## Additional Information

For more information, see the IN PARALLEL option of REBUILD AMP and REBUILD AMP FALLBACK TABLES in *Teradata Vantage™ - Database Utilities*, B035-1102.

# Release 17.05 Features

## Native Object Store Enhancements

Native Object Store (NOS) is available in this release on the following platforms:

- Vantage delivered as-a-Service platforms, such as Amazon Web Services, Microsoft Azure (including Azure Data Lake Storage Gen2 in Blob Interop Mode), and Google Cloud Services.
- Vantage DIY platforms, such as Amazon Web Services, Microsoft Azure (including Azure Data Lake Storage Gen2 in Blob Interop Mode), and Google Cloud Services.
- Vantage on IntelliFlex (IFX) on-premises object stores including Hitachi HCP, Dell EMC/ECS, MinIO, and NetApp StorageGrid.

### Additional Information

For more information about Native Object Store, see *Teradata Vantage™ - Native Object Store Getting Started Guide*, B035-1214.

## WRITE\_NOS for Native Object Store

Native Object Store (NOS) allows external data to be accessed and integrated into Teradata Vantage. For this release, NOS has been enhanced with the WRITE\_NOS feature.

WRITE\_NOS is a table operator that allows you to write data from database tables to external object storage and store it in Parquet data format.

### Benefits

- Allows you to free up database space by offloading old or cold data to cheaper external object storage
- Offloading data to object storage may reduce the overall cost of storage
- Data written to object store using WRITE\_NOS is optimized to be read back by foreign table
- Offloaded Parquet data can be queried (with foreign table). This enables you to meet compliance standards for retaining data
- Write data to various external sources, such as Amazon Web Services, Microsoft Azure (including Azure Data Lake Storage Gen2 in Blob Interop Mode), and Google Cloud Services.

### Considerations

- The maximum size of the output object is 16,776,192 bytes (including LOB and LOB-based UDT columns), which is the same as the read size of 16,776,192 bytes
- GZIP and SNAPPY are the supported compression algorithms for Parquet data
- Redshift is the supported format for the manifest files

- Parquet is the supported data format
- It is the responsibility of the user to clean up object store files on write failures interrupted write operations. Write operations can be interrupted on transaction aborts or system resets, among other reasons.
- Teradata recommends using WRITE\_NOS to offload data during a batch window
- WRITE\_NOS uses a large amount of memory.

## SQL Changes

- New WRITE\_NOS table operator

## Additional Information

For more information about creating foreign tables and accessing external data, see *Native Object Store: Teradata Vantage™ Advanced SQL Engine*, TDN0009800, *Teradata Vantage™ - SQL Data Definition Language Syntax and Examples*, B035-1144, *Teradata Vantage™ - SQL Data Manipulation Language*, B035-1146 and *Teradata Vantage™ - Native Object Store Getting Started Guide*, B035-1214.



# Release 17.00 Features

## Customer Data Space Tools

Customer Data Space (CDS) Tools let you examine data space consumption at the object, database, and system level. The tools report the logical size of your data in its uncompressed form. Customer Data Space tools consist of a set of SystemFE macros and a new Data Dictionary view.

### Benefits

- Determining current data space consumption can help determine trends and future space requirements.
- Eliminates the need for laborious manual space accounting calculations to find object sizes and estimated compression ratios for block-level compressed data.
- You can provide a list of tables to exclude from the CDS calculations, and use the tools to manage and view the list.
- Most of the macros are available in two versions, one for Teradata session mode and one for ANSI mode. Some have a single version that can be used for both session modes.
- CDS macros are installed automatically.

### Considerations

- Reported space consumption does not include fallback data.
- You must have the EXECUTE privilege on the macros.
- You must grant the SystemFE user STATISTICS and SELECT ACCESS privileges in order to use the *CollectSummaryStats* option of CDS\_Database or CDS\_System macros. This option samples the largest tables in the database or system, and can give more accurate space estimates.

### SQL Changes

New SystemFE macros:

- SystemFE.CDS\_Object  
SystemFE.CDS\_Object\_ansi
- SystemFE.CDS\_Database  
SystemFE.CDS\_Database\_ansi
- SystemFE.CDS\_System  
SystemFE.CDS\_System\_ansi
- SystemFE.CDS\_ViewPercentData - for both Teradata and ANSI session modes.
- SystemFE.CDS\_ModifyPercentData - for both Teradata and ANSI session modes.

- SystemFE.CDS\_ViewExclusions  
SystemFE.CDS\_ViewExclusions\_ansi
- SystemFE.CDS\_AddExclusion  
SystemFE.CDS\_AddExclusion\_ansi
- SystemFE.CDS\_RemoveExclusion  
SystemFE.CDS\_RemoveExclusion\_ansi
- SystemFE.CDS\_LogicalTableSizeEstimate  
SystemFE.CDS\_LogicalTableSizeEstimate\_ansi
- SystemFE.CDS\_PhysicalTableSizeEstimate  
SystemFE.CDS\_PhysicalTableSizeEstimate\_ansi

New Data Dictionary view:

- DBC.CDSTableSizeV

### Additional Information

For more information, see *Teradata Vantage™ - SystemFE Macros*, B035-1103 and *Teradata Vantage™ - Data Dictionary*, B035-1092.

## DBQL Enhancements

The Database Query Log (DBQL) is enhanced for Advanced SQL Engine 17.00 with more complete and more accurate logging.

### Benefits

- DBQL now uses Algorithm 3 by default, which includes collecting statistics on aborted and parallel steps, and results in more accurate resource usage statistics.
- By default, Vantage now logs the number of SUM and JOIN steps at the request level to gauge query complexity.
- The DBQL default behavior now is to include Feature Use Logging (FUL).

### Considerations

- Upgraded systems that were not using Algorithm 3 will not use Algorithm 3 until explicitly set to do so. This requires that you explicitly change DBS Control General field 64, DBQL CPU/IO Collection. Teradata recommends you switch to using Algorithm 3.

### Additional Information

- For more information about DBQL, see *Teradata Vantage™ - Database Administration*, B035-1093.
- For more information about the DBS Control utility, see *Teradata Vantage™ - Database Utilities*, B035-1102.

## Ferret SHOWBLOCKS Per AMP Information

The SHOWBLOCKS command of the Ferret utility displays information about data blocks for tables and sub-tables in the Advanced SQL Engine file system. This new feature adds a PERVPROC option to SHOWBLOCK that displays data block information per AMP vprocs within the current scope.

### Benefits

- Displays data block information for specified AMPs.
- Allows scoping to AMPs and maps, to display information about only the specified AMPs that are within the specified map.

### Considerations

- The PERVPROC option is not supported by the CreateFsysInfoTable and PopulateFsysInfoTable macros.

### Additional Information

- For more information about the Ferret utility, see *Teradata Vantage™ - Database Utilities*, B035-1102.
- For more information about the CreateFsysInfoTable and PopulateFsysInfoTable macros, see *Teradata Vantage™ - SQL Operators and User-Defined Functions*, B035-1210.

## Improved Resource Usage Monitoring for UDFs

Resource usage monitoring of user defined functions and of system provided UDFs, such as table operators has been improved in Teradata Vantage Advanced SQL Engine 17.00.

### Benefits

- Extends ability to collect resource usage statistics to all UDFs, including system UDFs like the SCRIPT and ExecR table operators.
- More accurate resource usage monitoring, that now includes all types of UDFs.
- Improved workload management is possible due to the more extensive and inclusive UDF monitoring.

### Additional Information

- For more information about workload management, see *Teradata Vantage™ - Workload Management User Guide*, B035-1197.
- For more information about resource usage statistics, see *Teradata Vantage™ - Resource Usage Macros and Tables*, B035-1099.

## Native Object Store

The Native Object Store (NOS) feature allows external data to be integrated into Teradata Vantage easily. You can create new *foreign tables* or use the new READ\_NOS table operator to access this data as if it were native Vantage relational data.

### Benefits

- Allows Vantage to read and process external data using standard SQL. This can vastly speed extracting valuable business insights from external data.
- Use data from external storage without the need to manually copy the data to Vantage permanent storage from where it natively resides.
- Read data from various external sources, such as Amazon S3, Azure Blob storage (including Azure Data Lake Storage Gen2 in Blob Interop Mode), and Google Cloud Storage.
- External data can be unstructured or semi-structured, stored in several formats, such as JSON, CSV, and Parquet.
- Join external data to existing relational data in Vantage for further processing by Teradata analytic functions.
- New CREATE FOREIGN TABLE statement and READ\_NOS table operator can access external data. READ\_NOS does not require the user to have CREATE TABLE database privileges.
- CREATE FOREIGN TABLE and READ\_NOS can include authorization information for access to the external data. (Access to the authorization information can be limited by storing it in a function mapping that has limited access.)

### Considerations

- NOS is available only on Vantage delivered as-a-service platforms (such as Teradata Vantage on AWS, Teradata Vantage on Azure, and Teradata Vantage on Google Cloud).
- Currently supports JSON, CSV, and Parquet data formats.
- You cannot modify (INSERT, UPDATE, or DELETE) data in foreign tables.
- You cannot use ALTER TABLE on foreign tables.
- BAR and DSA only archive the table definitions for foreign tables. Foreign data is not archived.

### SQL Changes

- New CREATE FOREIGN TABLE statement
- New FOREIGN optional keyword for DROP TABLE
- READ\_NOS table operator
- New EXTERNAL SECURITY clause for CREATE AUTHORIZATION
- New EXTERNAL SECURITY clause for CREATE FUNCTION MAPPING

### Additional Information

For more information about creating foreign tables and accessing external data, see *Native Object Store: Teradata Vantage™ Advanced SQL Engine*, TDN0009800, *Teradata Vantage™ - SQL Data*

*Definition Language Syntax and Examples*, B035-1144, and *Teradata Vantage™ - SQL Data Manipulation Language*, B035-1146.

## Security Enhancements in Release 17.00

- TDNEGO now supports JWT (JSON web token) authentication.
- New LdapServicePasswordFile property.
- The DIGEST-MD5 authentication protocol used by LDAP is deprecated.

### Benefits

- With the addition of JWT, Teradata Vantage now supports a greater variety of authentication protocols and mechanisms.
- The LdapServicePasswordFile property provides a way for you to change the LDAP service password without requiring a system restart.

### Considerations

- If you use the LdapServicePasswordFile property, the LdapServicePassword and LdapServicePasswordProtected properties are ignored, and passwords are read exclusively from the password file.
- Passwords listed in the LDAP service password file must be encrypted using the tdspasswd command-line utility.
- Teradata strongly recommends you stop using DIGEST-MD5, and instead use simple binding with TLS protection. A future release of Vantage will remove support for DIGEST-MD5.

### Additional Information

For more information about security, see *Teradata Vantage™ - Advanced SQL Engine Security Administration*, B035-1100.

## TD\_JSONSHRED Table Operator

The TD\_JSONSHRED table operator allows you to shred data contents of an array in a JSON document into columns of a Teradata Vantage database table.

### Benefits

- TD\_JSONSHRED can be significantly faster at shredding data than the JSON\_TABLE table operator.
- The JSON document is input to TD\_JSONSHRED as the contents of a table column. The column can be of data type JSON, CLOB, or VARCHAR.
- You can specify the data type for the shredded output data columns, which can include CLOB data.
- The output table can include additional columns (not JSON data) from the input table. These columns are passed through to the output table unchanged.

- String matching for identifying JSON objects to shred in the JSON document can optionally be case-insensitive.

### Considerations

- You must use nested calls to TD\_JSONSHRED to shred the contents of nested arrays in the original JSON document.
- TD\_JSONSHRED does not support JSONPath expressions.
- TD\_JSONSHRED does not support as many output data types as the JSON\_TABLE table operator.

### SQL Changes

- New TD\_JSONSHRED table operator.

### Additional Information

For more information about TD\_JSONSHRED, see *Teradata Vantage™ - JSON Data Type*, B035-1150.

## UDF Search Path

The new SET SESSION UDFSearchPath statement lets you set a custom path for Teradata Vantage to use when searching for invoked UDFs.

### Benefits

- Allows you to specify where Vantage should look first for UDFs. This ensures that the UDF you want will be executed in cases where there might be more than one UDF in the system with the same name.
- The UDF search path you set is applied to all forms of UDFs, including user-defined functions, user-defined types, and embedded services system functions.

### Considerations

- Best practice is to use the fully qualified UDF path and name when invoking a UDF, but if you cannot or might not, you should use the new SET SESSION UDFSearchPath statement to control the location order that Vantage uses to search for the UDF.

### SQL Changes

- New SET SESSION UDFSearchPath statement

### Additional Information

For more information about how Teradata Vantage searches for UDFs, see "UDF Locations" in *Teradata Vantage™ - SQL External Routine Programming*, B035-1147.

# Teradata Vantage Restricted Words

This section lists new Teradata reserved and nonreserved words, as well as ANSI SQL:2011 reserved and nonreserved words introduced in Teradata Vantage Advanced SQL Engine Releases 17.xx.

For information on how to generate a complete list of restricted words, see *Teradata Vantage™ - SQL Fundamentals*, B035-1141.

## Teradata Vantage Reserved Words

Teradata Vantage reserved words cannot be used as identifiers to name host variables, correlations, local variables in SQL stored procedures, objects (such as databases, tables, columns, or stored procedures), or parameters, such as macro or stored procedure parameters, because Teradata Vantage already uses the word and might misinterpret it.

## Teradata Vantage Nonreserved Words

Teradata Vantage nonreserved keywords are permitted as identifiers but are discouraged because they may in the future be used as reserved keywords.

## Teradata Vantage Future Reserved Words

Teradata Vantage future reserved words are words reserved for future Vantage use, and cannot be used as identifiers.

## Teradata Parallel Transporter Restricted Words

For a list of Teradata Parallel Transporter (TPT) restricted words, see *Teradata® Parallel Transporter User Guide*, B035-2445.

## ANSI SQL:2011 Reserved Words

The danger in using an ANSI reserved word that is not currently a Teradata reserved word is that it may in the future become a Teradata reserved word if Teradata implements the ANSI feature that uses it.

## ANSI SQL:2011 Nonreserved Words

The danger in using an ANSI nonreserved word is that it may in the future become a Teradata reserved word if Teradata implements the ANSI feature that uses it.

## New Restricted Words

Word	Teradata Vantage Status			ANSI SQL-2011 Status	
	Reserved	Future	Nonreserved	Reserved	Nonreserved
FUNCTIONDESCRIPTOR	17.10				

Word	Teradata Vantage Status			ANSI SQL-2011 Status	
	Reserved	Future	Nonreserved	Reserved	Nonreserved
TD_ROWSIZE	17.00				
TD_ROWREVISION	17.00				

## Getting the Restricted Words for the Current Release

This query returns the restricted words for the current database release. Because it returns all columns from the view, it includes the Teradata and ANSI categories for the words.

A portion of the output is shown below.

```
SELECT * FROM SYSLIB.SQLRestrictedWords;
```

restricted_word	category	ANSI_restricted
-----	-----	-----
ABORT	R	T
ABORTSESSION	R	T
ABS	R	R
ACCESS_LOCK	R	T
...	...	...

For more information, see *Teradata Vantage™ - SQL Fundamentals*.



## Additional Information

### Teradata Links

Link	Description
<a href="https://docs.teradata.com/">https://docs.teradata.com/</a>	Search Teradata Documentation, customize content to your needs, and download PDFs. Customers: Log in to access Orange Books.
<a href="https://support.teradata.com">https://support.teradata.com</a>	One-stop source for Teradata community support, software downloads, and product information. Log in for customer access to: <ul style="list-style-type: none"><li>• Community support</li><li>• Software updates</li><li>• Knowledge articles</li></ul>
<a href="https://www.teradata.com/University/Overview">https://www.teradata.com/University/Overview</a>	Teradata education network
<a href="https://support.teradata.com/community">https://support.teradata.com/community</a>	Link to Teradata community