

Teradata Vantage™ NewSQL Engine Release Definition

Release 16.20

December 2020

Copyright and Trademarks

Copyright © 2000 – 2020 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
<i>NOTICE</i>	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.

Warranty Disclaimer

Except as may be provided in a separate written agreement with Teradata or required by applicable law, the information contained in this document 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 contained in this document 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 contained in this document 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.

Feedback

To maintain the quality of our products and services, email your comments on the accuracy, clarity, organization, and value of this document to: teradata-books@lists.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.

Table of Contents

Introduction	5
Understanding This Release	5
Additional Information	5
Software and Hardware Requirements	6
Supported Software and Hardware	6
Software Maintenance Schedule	6
Teradata Database Maintenance Roadmap	7
Platform Coexistence and Coresidence Requirements	7
Node Memory Recommendations	7
Memory-Consuming Features	8
Reserved Words	10
Software and Hardware Restrictions	11
Obsolete and Unsupported Tools, Utilities, Options, Record Types, and Other Software	11
TDGSS Single Mechanism to Log-On (TDNEGO) Teradata Unity Support	13
1 MB Perm and Response Rows	14
Teradata Archive/Recovery Utility	14
Deprecated Kanji1 Character Set	14
Kanji1 Restrictions	14
Compatibility Views	14
Security-Related Restrictions	15
Custom Authentication Mechanisms	15
System-Level Software Restrictions	15
Analytic Functions	16
Teradata Unity	16
Changes in System Behavior	17
Default Feature Status	17
1 MB Response Rows	19
1 MB Maximum Request Bytes	19
Data Dictionary Tables and Views Column Data Type Changed	19
DBQL Cache Sizes	20
DBS Control Changes	20
Execute SCRIPT Table Operator with a Specific Authorization/OS User	20
Fallback	21
Fallback and Clustering	21
Function Mapping for Native Functions	21
Function Mapping Variable Substitution	21
Legacy versus Dedicated Fallback	22
Heatmap Table Function Integration and Enhancements	22

IPE Enhancements: Cache Dynamic Plans (CDP) for IPE Statistics Feedback - Nonparameterized Requests	22
Java Runtime Environment (JRE 1.8)	22
JSON Web Token (JWT).....	22
Link-local IP Address Restrictions	23
Minimum Bounding Box (MBB) and Minimum Bounding Rectangle (MBR) Output	23
Packageless TeraGSS.....	23
Replace CFS Hard Limits with PSF Hard Limits	23
Resource Usage Macros.....	24
SHOW QUERY LOGGING FeatureUsage Field.....	24
Teradata Database MAPS Architecture	25
TDGSS Single Mechanism to Log-On Client Support.....	25
TDWM UDF GDO Size	25
Trigger Enhancement for Row Level Security Tables	26
Trusted Session Support in X Views	26
Unicode Pass Through	26
Cleansing and Loading Data with Unicode Pass Through	27
Recompiling Stored Procedures.....	27
Installation, Upgrade, Migration, and Backdown (IUMB)	28
Supported IUMB Operations	28
IUMB Planning.....	28
Upgrading Teradata Temporal Tables	28
Parallel Upgrade Utility (PUT).....	29
Unicode Pass Through (UPT) IUMB	29
Replacing Unsupported Operating Systems	30
Upgrading from Older Releases.....	30
About Returning to an Older Release	30
System Performance	31
Performance Regressions.....	31
Running Teradata Database with Other Applications	31

Introduction

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 NewSQL Engine is the core of Teradata Vantage, based on our best-of-breed Teradata Database processing capability. NewSQL refers to the ability to run advanced analytic functions beyond that of standard SQL.

Understanding This Release

This *Release Definition* applies to Teradata Database Release 16.20 and the two updates to this release: 16.20 Feature Update 1 and 16.20 Feature Update 2. This *Release Definition* uses the term Release 16.20.xx when the text applies to Release 16.20, 16.20 Feature Update 1, and 16.20 Feature Update 2

Before you install or upgrade to this release of Teradata® Database, read the following publications:

- *Release Definition*, B035-1725, which provides information about basic system requirements. If you are upgrading from an older Teradata Database release, read the *Release Definition* for every intervening release.
- *Release Summary*, B035-1098, which describes the new features in a release. If you are upgrading from an older release of Teradata Database, you should also read the *Release Summary* for every intervening release to understand how the features in the latest release differ from your current version.

Additional Information

Link	Description
https://docs.teradata.com	Teradata documentation (HTML)
https://www.info.teradata.com	Teradata documentation (PDF)
https://access.teradata.com/	Customer portal (one stop source for Teradata services and products)
http://www.teradata.com/products-and-services/TEN	Teradata Education Network
http://developer.teradata.com/	Public downloads (also available from the customer portal)
https://community.teradata.com/	Link to Teradata community (also available from the customer portal)

Software and Hardware Requirements

Supported Software and Hardware

Supported	Description
Operating Systems	<ul style="list-style-type: none"> SUSE Linux Enterprise Server (SLES) 12 SP3: Teradata Database 16.20 Feature Update 2. SLES 11 SP3: Teradata Database Release 16.20, 16.20 Feature Update 1, and 16.20 Feature Update 2.
Hardware Platforms	See <i>Teradata Platform and Minimum Database Version Compatibility Matrix</i> in KCS008400, available at https://access.teradata.com .
Cloud Platforms	<ul style="list-style-type: none"> Public Cloud: Amazon AWS, Microsoft Azure Private Cloud: Teradata Database on VMware
Teradata Vantage	<p>Teradata Vantage includes analytic functions and engines, preferred tools and languages, and support for multiple data types:</p> <ul style="list-style-type: none"> Analytic engines include SQL, machine learning, and graph engines. Languages include SQL, R, and Python. Workbenches and tools include Teradata Studio, Teradata AppCenter, Jupyter, and RStudio. Data support includes relational, spatial, temporal, XML, JSON, Avro, and time-series formats.
Compilers	Installation of a C++ compiler is required on at least one database node configured with a PE vproc. C++ compilers are included with each Teradata Database release as part of the operating system software disks.
Free Disk Space	For information on the amount of free disk space required on each Teradata Database node to upgrade to this release, see IDA00108C82, available at https://access.teradata.com . (You must log into the website to view this article.)
Additional Disk Space for Trace Files	The Write Ahead Logging (WAL) feature requires 5 MB per AMP of disk space for File System trace files. For example, if there are 10 AMPs per node, then trace files would require 50 MB per node of additional disk space, located in /var/opt/teradata/tdtemp.
Supported External Disk Arrays	See Product and Site Preparation Guide for your platform, available at https://www.info.teradata.com .
Backup, Archive, and Restore (BAR)	<p>Supported BAR Software and Compatibility Matrix</p> <p>For information on backup and restore software versions and version compatibilities, see the DSA Ecosystem Compatibility Matrix in KCS000003 available at https://access.teradata.com.</p>
Teradata Tools and Utilities (TTU)	For detailed information on the full range of tools and utilities and the individual Teradata client product versions compatible with Teradata Database 16.20, see <i>Teradata Tools and Utilities Supported Platforms and Product Versions</i> . Search for Document/Product ID: B035-3119 at: https://www.info.teradata.com .
UDFs	<p>A library of UDFs is available at Teradata Downloads.</p> <p>The terms under which these UDFs are available (“as is” with no support) are described on the website.</p>

Software Maintenance Schedule

Purchasers of Teradata Database software are entitled to a period of continuing support after initial installation or upgrade.

Teradata Database Maintenance Roadmap

The Teradata Maintenance Release Roadmap shows the detailed code-level remedy and support for each Teradata Database version.

- Customers with active Service Agreements can find the Teradata Maintenance Release Roadmap at <https://access.teradata.com>.
- Customers without support agreements should contact their sales or support team.

Mixed Operating Systems

A mixed operating system includes both Trusted Parallel Application (TPA) nodes and non-TPA nodes.

The following are the basic system requirements for mixed operating systems:

- All TPA nodes must run the same operating system, but you can mix non-TPA nodes in any combination of the allowed operating systems.
- TPA nodes need not run on the same operating system as non-TPA nodes.
- Mixed OS allows two consecutive generations of nodes/storage across TPA and/or non-TPA nodes.
- The same version of BYNET must be deployed across the system.

Allowed platform combinations are the same as those allowed for coexistence and coresidence systems.

Platform Coexistence and Coresidence Requirements

You can upgrade Teradata system nodes or add new nodes to a system to enhance system capacity or performance. Teradata offers the following system enhancements, which are subject to some limitations by platform type:

- **Coexistence** allows you to combine existing system nodes with certain newer, more powerful nodes. You can assign an unbalanced number of AMPs per node to optimize both current and new platform resources.
- **Coresidence** allows you to combine existing system nodes with certain newer, more powerful nodes. You can assign only a fixed number of AMPs per node, which may limit the use of newer platform resources.
- **Upgrade** allows you to replace processors in an existing node to make it functionally equivalent to a more powerful node model.

Determining whether a Teradata node can coexist or coreside with other node models depends on the platform type, performance class, operating system, and the version of BYNET installed on each node.

Coexistence and coresidence requirements are not part of the standard user documentation set. For more information, contact the Teradata Support Center.

Node Memory Recommendations

For best performance, Teradata recommends that each node has at least the minimum recommended RAM. See *Teradata Platform and Minimum Database Version Compatibility Matrix* in KCS008400,

(available at <https://access.teradata.com>) for node memory recommendations, including the default installed RAM per node, recommended minimum RAM per node, and recommended maximum RAM per node.

When upgrading to Release 16.20.xx, several factors can cause some systems, especially large ones, to require additional memory.

The general guidelines for memory follow; however, memory requirements are workload-dependent, so your system's actual memory requirements may differ.

General Guidelines:

- Teradata recommends a minimum of 4 GB memory per vproc to achieve the most value and performance from Teradata Database 16.20.xx. The absolute minimum requirement is 2 GB memory per vproc (including AMP, PE, TVS, and GTW).
- Some Teradata Database 16.20.xx features require 512 GB per node; for example, Teradata In-Memory Optimizations.
- Additional memory consumed is based on the size of system, AMPs per node, AWT, and feature use.

Note: The maximum amount of memory allowed per node may be increased between database releases. Always double-check what the maximum memory amount is for your platform.

Release 16.20.xx can run on a system with the recommended minimum RAM, but performance may not be optimal, depending on the system configuration and the Teradata Database features you use. You should also factor in the following to determine the optimal memory configuration:

- Workload
- Memory-consuming features
- Performance requirements
- Cost of memory

Memory-Consuming Features

These features may require more memory for optimum system performance:

Release First Available	Memory Consuming Features
16.20, 16.20 Feature Update 1, 16.20 Feature Update 2	<ul style="list-style-type: none">• None
16.10	<ul style="list-style-type: none">• None
16.00	<ul style="list-style-type: none">• 1 MB Perm and Response Rows• In-Memory Enhancements• AVRO DATASET• Multiple Count Distinct Performance• Queryable Column Information on Views

Release First Available	Memory Consuming Features
15.10	<ul style="list-style-type: none"> • XSLT_SHRED_BATCH • In-Memory Optimizations • BSON and UBSON • Columnar Primary AMP/Primary Index • SQL Interface for Ferret SHOWBLOCKS • Parameterized Query Logging
15.00	<ul style="list-style-type: none"> • JSON Data Type • 3D Geospatial • Scripting and Language Support • DBQL – Show Parameters • 1 MB Phase 2 • Script Table Operators • QueryGrid: Teradata Database-to-Hadoop
14.10	<ul style="list-style-type: none"> • Auto Stats Enhancements • Data Stream Architecture • Extended Object Naming • Geospatial Indexing • Incremental Planning and Execution • Teradata Intelligent Memory • Teradata XML • 1 MB Data Block • 1 MB Spool Row • 128K Parser Tree Segments
14.0	<ul style="list-style-type: none"> • Teradata Columnar • Partial Online Reconfiguration • SLES 11 (first available in Release 14.0.2)
Pre-14.0	<ul style="list-style-type: none"> • PPI and Multivalue Compression • Join Index, Hash-Join, Stored Procedures, and 128K Data Blocks • Cylinder Read • LOBs and UDFs • 1 MB Response Buffer • Larger than 1 MB Plan Cache • External Stored Procedures • Table Functions • Array INSERT • Java Stored Procedures • Online Archive Memory Enhancements • More than 80 AWTs per AMP • Expanded Table Header • Geospatial Data Type • Increased Join/Subquery Limits • Teradata Virtual Storage • Tunable UDF Memory Limit • Algorithmic Compression and Block Level Compression • XML DBQL Logging

Release First Available	Memory Consuming Features
	<ul style="list-style-type: none"> • Global and Persistent Data (GLOP) • Large Cylinder with Cylinder Read • More Than 20 AMPs/Vprocs per Node (All Releases) • Temporal DBS Support

Reserved Words

Teradata Database reserved words cannot be used as identifiers to name host variables, correlations, local variables in stored procedures, objects (such as databases, tables, columns, or stored procedures), or parameters (such as macro or stored procedure parameters).

New reserved words for this release are listed in the “Restricted Words” appendix in *Release Summary*, B035-1098, available at <https://www.info.teradata.com> or <https://docs.teradata.com>.

Software and Hardware Restrictions

Obsolete and Unsupported Tools, Utilities, Options, Record Types, and Other Software

Utility	Final Release that Contains this Feature	Replacement Feature, if any	Additional Information
DBS Check tool (dbschk)	TTU 15.0 TD 15.10	Mailbox Check (mboxchk) tool; see the man page for information	
dbscontrol	This utility is still supported.		Obsolete DBS Control fields include: <ul style="list-style-type: none"> • DisplacementOnOverlap • MDS Is Enabled
dbscsp	TD 12.00		The dbscsp tool, used only on MP-RAS systems, is no longer supported. The executable /usr/ntos/bin/dbgcsp now links to fdclsp instead of dbscsp.
DULTAPE	TD 16.00	You can continue to use DUL.	
gdoviewer	Removed prior to TD 12.00		
gtwcontrol	This utility is still supported.		Removed -b option. Deprecated logons are no longer allowed.
HP-UX Itanium and IBM Mainframe z/Linux (RedHat and SUSE)	TTU 16.10		
ITEQ, HUT CNS	TD 15.00		
Meta Data Services (MDS)	TD 15.00		
OLE DB Provider for Teradata	TTU 15.0 TD 15.10	Use Microsoft's OLE DB Provider for ODBC and Teradata ODBC Driver products together	
PMON	TD 13.10	Teradata Viewpoint	
Priority Scheduler	TD 12.00		Priority scheduler functions must be controlled through Teradata Viewpoint, Workload Designer portlet
rcvmanager			F7 help is not available.
Replication Services	TD 14.10		Teradata Replication Services (Teradata to Teradata replication) was discontinued for new sales as of August 2011. Aligned with that discontinuation, no further enhancements have been made since the

Utility	Final Release that Contains this Feature	Replacement Feature, if any	Additional Information
			<p>TRS 13.10 release. TRS 13.10 has been certified with Teradata DB 14.00 and 14.10 versions but with no new feature support and for existing customers only. Teradata Unity™ is the replacement for TRS.</p> <p>Note: Replication from third-party solutions to Teradata is still supported by Oracle GoldenGate.</p>
rssmon	TD 13.00		The rssmon utility (Resource Sampling System Monitor) was only usable on MP-RAS systems and is obsolete now that Teradata Database is no longer supported on MP-RAS.
SLES 10	TD 15.10	Later versions of SLES, depending on your release.	
SQL Assistant	TTU 16.20 TD 16.00	Teradata Studio and Teradata Studio Express	
tdgsspkgm	TD 15.10		
tdgssversion	TD 16.00		
tdssearch	TD 13.00		Due to limited functionality, tdssearch has been replaced by ldapsearch, which is included with Teradata Database 13.10 and later.
Teradata Administrator	TTU 15.10 TD 15.10	Replaced by Teradata Studio	
Teradata Archive/Recovery Utility (ARC)	TTU 16.20	Teradata Data Stream Architecture (DSA) or Teradata Data Stream Utility (DSU)	
Teradata Data Mart Edition	TTU 15.0 TD 15.10	Teradata Database is no longer natively supported on 3 rd party SMP computers running specific versions of SUSE Linux. The replacement product is Teradata Virtual Machine Edition (TVME). Please refer to the separately published <i>Order and Configuration Information</i> document for TVME.	
Teradata Dynamic Workload Manager	TD 13.00	TASM, controlled by Teradata Viewpoint	
Teradata IDE-Plugin for Eclipse	TTU 16.00		

Utility	Final Release that Contains this Feature	Replacement Feature, if any	Additional Information
Teradata Index Wizard	TTU 16.20		
Teradata Manager	TD 13.00	Teradata Viewpoint	
Teradata Method 1 (TD1), NTLM, NTLMC, and KRB5C	TD 16.00		These authentication mechanisms must be manually enabled to use them in Release 16.00.
Teradata Monitor	TTU 15.0 TD 15.10	A set of Teradata table functions embedded in the Teradata Database	
Teradata Multitool	TD 13.10	Teradata Database command line utilities, such as Database Window	
Teradata Preprocessor2 (PP2) for C and COBOL <ul style="list-style-type: none"> Network Platforms (Windows, Linux, Unix) 	TTU 16.20	ODBC	
Teradata Query Director	TD 13.10	Teradata Unity	
Teradata Query Scheduler (TQS)	TD 16.00		
Teradata Statistics Wizard (TSWIZ)	TTU 14.10 TD 15.00	Teradata Viewpoint Statistics Manager	
Teradata Visual Explain (VEComp)	TTU 16.20 TD 16.00	Visual Explain App in Teradata App Center	
Transparency Series/Application Programming Interface (TS/API) (mainframe)	TTU 15.0 TD 15.10	An SQL query tool, such as Teradata Studio	Capped at supporting Query Management Facility (QMF) 9.1.
vpacd	TD 12.00		
Windows Vista	TTU 16.10		

TDGSS Single Mechanism to Log-On (TDNEGO) Teradata Unity Support

Unity for Teradata Database 15.10 and later does not support TDNEGO. Teradata recommends disabling TDNEGO on Unity servers when the Unity version for Release 15.10 or later is installed. For more information, see *Teradata Vantage™ NewSQL Engine Security Administration*, B035-1100.

1 MB Perm and Response Rows

This feature is not supported on Small Cylinder systems (such as systems using a maximum cylinder size of 3872 sectors or approximately 1.9 MB).

Teradata Archive/Recovery Utility

When Teradata Database MAPS Architecture (MAPS) is enabled, the Teradata Archive/Recovery Utility (ARC) can only archive/restore objects contained in the all-AMPs contiguous map. For more information, see “Running Teradata ARC in a Multiple Hash Map (MAPS) Environment” in *Teradata Archive/Recovery Utility Reference*, B035-2412.

If MAPS is enabled, you can use DSA for archive/recovery operations. For more information about DSA, see *Teradata Data Stream Architecture (DSA) User Guide*, B035-3150.

Deprecated Kanji1 Character Set

As of Release 14.0 and later, Kanji1 support is deprecated and planned for discontinuation. Although many Kanji1 queries and applications may continue to operate, you should prepare to convert Kanji1 data to another character set as soon as possible.

During an upgrade to Teradata Database from a pre-14.0 release, the system automatically replaces DEFAULT CHARACTER SET KANJI1 with DEFAULT CHARACTER SET UNICODE in existing user definitions.

Kanji1 Restrictions

As part of the plans for discontinuing Kanji1 support, creation of new Kanji1 objects is highly restricted. For example, inclusion of the phrase CHARACTER SET KANJI1 in the following statements returns a syntax error:

- CREATE USER/MODIFY USER
- CREATE TABLE/ALTER TABLE
- CREATE FUNCTION/REPLACE FUNCTION
- CREATE TYPE/ALTER TYPE
- CREATE PROCEDURE/REPLACE PROCEDURE
- CREATE MACRO/REPLACE MACRO
- CREATE VIEW/REPLACE VIEW
- CAST function

Use the TRANSLATE function to convert existing Kanji1 data to Unicode or another supported server character set. For more information, see “TRANSLATE” in *Teradata Vantage™ SQL Functions, Expressions, and Predicates*, B035-1145.

Compatibility Views

Compatibility views convert the native variable-length Unicode object names into 30 bytes of either Latin or Kanji1. This can cause loss of information by truncation or inability to convert object names

longer than 30 characters into Latin or Kanji1. Characters that cannot be converted are replaced by the substitution character, which is 0x1A for both Latin and Kanji1.

As with all character data, when object names are returned to the user they are converted to the session character set. This conversion can produce loss of information if the characters in the object name cannot be converted to the session character set or exceed the export width for the character data. Teradata recommends using Unicode views. For more information, see *Teradata Vantage™ Data Dictionary*, B035-1092.

Security-Related Restrictions

Custom Authentication Mechanisms

Development and deployment of a custom authentication mechanism for use with this release can only be achieved by purchasing the Teradata Security Software Developer Kit.

Note that development and deployment of a custom authentication mechanism results in the creation of a custom encryption software product. Distribution of a custom encryption product outside the U.S. and Canada is regulated by the U.S. Department of Commerce (DOC). An export authorization must be filed for and obtained from DOC in order to export such custom products. For information on obtaining an export authorization, please see the DOC website at [BIS Website](#).

If additional information is required, customers should contact their Teradata representative. Teradata representatives requiring help should contact Teradata Corporate Export Compliance, Law Department.

Note: A custom authentication mechanism is a user authentication mechanism that is above and beyond the mechanisms that are provided with Teradata Database:

1. Teradata Method 2
2. KRB5 (Kerberos authentication)
3. LDAP
4. SPNEGO (used for Kerberos authentication for logons from Windows .NET clients)
5. TDNEGO

System-Level Software Restrictions

- JRE 1.8 must be installed on the database server prior to installing Teradata Database 16.20.
- Only one instance of Teradata Database is supported on a system.
- Teradata Tools and Utilities, including utilities on mainframes, must be installed at or upgraded to Teradata Tools and Utilities 16.20 to all Teradata Database 16.20.xx features and functions.
- For additional restrictions, dependencies, and performance considerations when running the Teradata Database applications, see [Running Teradata Database with Other Applications](#).
- Backup and restore management utilities are not provided as part of Teradata Database but are available as separate products. For more information, see [Supported BAR Software](#).
- A maximum of 1,200 concurrent LAN-connected sessions are allowed per node.

-
- The maximum number of sessions for mainframe clients is 120 x the number of configured Parsing Engines (PEs) for each TDP (Logical Host ID).

Analytic Functions

Unicode is not supported for the nPath® and Attribution functions on NewSQL Engine.

Note: nPath and Attribution_MLE support Unicode on ML Engine.

Teradata Unity

Teradata Unity Release 16.20.33 provides Unity Passive routing support with ML Engine. Additionally, this release of Unity continues to support Passive and Managed routing with NewSQL Engine. Unity Managed routing with ML Engine is a roadmap item and scheduled for a future Unity release.

Changes in System Behavior

Default Feature Status

The following features are disabled by default in this release:

- COUNT DISTINCT Performance Improvement
- In-Memory Outer Hash Join Optimization
- UNION ALL Optimization
- Queryable Column Information on Views

In previous releases, some features were enabled by default and others were manually enabled, depending on whether the Teradata Database system had a fresh installation (sysinit) or an upgrade.

Release 16.00 and later features do not require a sysinit to be enabled. Features in earlier releases that required a sysinit still require a sysinit (for details, see Release 15.10 *Release Definition*, B035-1725).

Feature	License Tiers	Upgrade Sysinit	Effects and Comments
OVERWRITE ON ERROR	n/a	Upgrade: Enabled Sysinit: Enabled	OVERWRITE ON ERROR has become the default when creating tables. If a user names an invalid map when creating a table, the system uses the default map for the user, role, or profile instead. If there is no default map for the user, role, or profile, the system default map is used.
Always Fallback	n/a	Upgrade: Disabled Sysinit: Enabled	Fallback behavior and defaults have become platform-specific. Newer platforms now always use fallback, even if you specify NO FALLBACK. Older platforms that upgrade to this release still allow the NO FALLBACK option, and the default is NO FALLBACK for CREATE TABLE, ALTER TABLE, CREATE JOIN INDEX, CREATE HASH INDEX, CREATE DATABASE, MODIFY DATABASE, CREATE USER, and MODIFY USER requests.
In-Memory Optimization	Developer: Not Applicable Base: Not Applicable Advanced: Not offered Enterprise: Included	Upgrade: Disabled Sysinit: Disabled	In-Memory Optimization is enabled when the appropriate license and memory are purchased. To enable, contact Teradata Support Center.
Temporal	Developer: Included Base: Included Advanced: Included Enterprise: Included	Upgrade: Disabled Sysinit: Disabled	To enable, contact Teradata Support Center. Note: After Temporal is enabled, it cannot be disabled.
Teradata Secure Zones	Developer: Included Base: Included Advanced: Included	Upgrade: Disabled	Teradata Secure Zones must be enabled to get full functionality.

Feature	License Tiers	Upgrade Sysinit	Effects and Comments
	Enterprise: Included	Sysinit: Disabled	To enable, contact Teradata Support Center.
Teradata Database MAPS Architecture (MAPS)	Developer: Not Applicable Base: Included Advanced: Included Enterprise: Included	Upgrade: Disabled Sysinit: Enabled	This feature is: <ul style="list-style-type: none"> • Enabled for new installations • Disabled by default on upgraded systems To enable, contact Teradata Support Center. After this feature is enabled backdown is not supported. Note: When MAPS is enabled, Teradata Archive/Recovery Utility (ARC) can only archive/restore objects in the all-AMPs contiguous map. For more information, see “Running Teradata ARC in a Multiple Hash Map (MAPS) Environment” in the <i>Teradata Archive/Recovery Utility Reference</i> , B035-2412.
Adaptive Optimization	Developer: Includes IPE Base: Includes IPE Advanced: Includes IPE Enterprise: Includes Enhanced IPE	Upgrade: Enabled for Enterprise Sysinit: Enabled for Enterprise	Enhanced IPE provides: <ul style="list-style-type: none"> • Noncorrelated subqueries with small results sets • Single-row query blocks • Derived table or view produces zero rows or a single row • Spooled subqueries with smaller result sets • Advanced rewrites based on results feedback
Block loads between Row Level Security (RLS) tables and non-RLS tables	Developer: Included Base: Included Advanced: Included Enterprise: Included	Upgrade: Disabled Sysinit: Disabled	To enable, contact Teradata Support Center.
IN-list Rewrite	Developer: Included Base: Included Advanced: Included Enterprise: Included	Upgrade: Enabled Sysinit: Enabled	To disable this feature, use DBS Control fields. For more information, see <i>Utilities</i> , B035-1102.
TASM I/O Usage Event	Developer: Not Applicable Base: Not Applicable Advanced: Not Applicable Enterprise: Included	Upgrade: Disabled Sysinit: Disabled	To enable, contact Teradata Support Center. For more information, see <i>Workload Management User Guide</i> , B035-1197, or <i>Teradata Viewpoint User Guide</i> , B035-2206.
Workload Management	Developer: Not Applicable Base: Not Applicable Advanced: TIWM Enterprise: TASM	Upgrade: Disabled Sysinit: Disabled	To enable, contact Teradata Support Center. The appropriate license must be purchased first.
Intelligent Memory (TIM)	Developer: Not Applicable Base: Not Applicable Advanced: Included Enterprise: Included	Upgrade: Disabled Sysinit: Disabled	TIM is enabled when the appropriate license and memory are purchased. To enable, contact Teradata Support Center.

Feature	License Tiers	Upgrade Sysinit	Effects and Comments
TVS (Hybrid Storage)	Developer: Not Applicable Base: Not Applicable Advanced: Not Applicable Enterprise: Included	Upgrade: Enabled Sysinit: Enabled	TVS (Hybrid Storage) is enabled by default when using the hybrid storage system.
Concurrent Query Limit	Developer: 2 Base: 15 Advanced: None Enterprise: None	Upgrade: Enabled Sysinit: Enabled	Appropriate license must be purchased first.
Function Mapping	Developer: Included Base: Included Advanced: Included Enterprise: Included	Upgrade: Enabled Sysinit: Enabled	Teradata Database 16.20 Feature Update 1 and later supports function mapping to simplify executing functions on foreign servers.
Teradata Analytic Functions	Developer: Included Base: Included Advanced: Included Enterprise: Included	Upgrade: Enabled Sysinit: Enabled	Teradata Database 16.20 Feature Update 1 and later now natively executes some scoring and predictive analytical functions. See Teradata Analytic Functions for Release 16.20.xx in <i>Teradata Vantage™ NewSQL Engine Release Summary</i> , B035-1098. Teradata Database 16.20 Feature Update 2 additionally supports these functions: Antiselect, MovingAverage, NGramSplitter, Pack, StringSimilarity, and Unpack. See Teradata Analytic Functions for Release 16.20 Update 2 in <i>Teradata Vantage™ NewSQL Engine Release Summary</i> , B035-1098.

Default enabling may cause changes in system behavior compared with previous releases. Even after enabling, some features may require additional configuration.

1 MB Response Rows

In prior releases, client-server response messages were limited to 1 MB. In Release 16.00 and later, Teradata Database supports client-server response messages of up to 16 MB. Applications designed for Teradata Database Release 16.00 and later can take advantage of this larger message size.

1 MB Maximum Request Bytes

In prior releases, client-server requests were limited to 1 MB. In Release 16.00 and later, Teradata Database supports client-server SQL requests of up to 7 MB.

Data Dictionary Tables and Views Column Data Type Changed

In Release 16.00 and later, the following tables changed column data type from FLOAT to BIGINT:

-
- DBC.DatabaseSpace
 - DBC.DBQLStepTbl
 - DBC.DBQLLogTbl
 - DBC.Dbase
 - DBC.Profiles

The following views are affected by this change:

- DiskSpaceV[X][_SZ]
- TableSizeV[X][_SZ]
- DBC.AllSpaceV[X][_SZ]

Customers with pre-existing queries listing these tables and views must explicitly CAST BIGINT columns to FLOAT.

DBQL Cache Sizes

DBQL cache sizes are enhanced in Teradata Database 16.20 Feature Update 1 and later. The cache size increases from 64 KB to 2 MB. By default, the new cache size will be enabled on all systems for upgrades and migrations. If you prefer to use the old default (64 KB), then set the `DBQLDefCacheSize` performance parameter to 0 in DBS Control utility.

Users can now flush the object use count cache whenever their applications require it instead of waiting for the 10-minute system default period to expire or the cache to become full.

DBS Control Changes

The default value of the `MaxParseTreeSegs` field is increased to 3,000 in Release 16.0 and later. The value is automatically increased on upgrade if the current value is less than the new default.

The following DBS Control fields are obsolete in Release 16.00 and later:

- `DisplacementOnOverlap`
- `MDSIsEnabled`

In previous releases, the `ROUND` and `TRUNC` functions returned a `DATE` type if they were passed a `TIMESTAMP` type. In Release 16.10 and later, a new DBS Control General field `TruncRoundReturnTimestamp` can be set to `TRUE` to have these functions return `TIMESTAMP` values in this situation.

Execute SCRIPT Table Operator with a Specific Authorization/OS User

You can grant the `EXECUTE` privilege on an authorization object. Granting `EXECUTE` on an authorization object allows the grantee to execute the `SCRIPT` table operator using the credentials

provided by the authorization object. This mechanism essentially binds the execution of a script by a database user to a valid operating system user. You can create your own authorization objects or use the default authorization, SYSUIF.DEFAULT_AUTH. If a database user has the EXECUTE privilege WITH GRANT OPTION on an authorization object, they can grant EXECUTE on that authorization object to another user.

Fallback

Beginning in Teradata Database 16.20 Feature Update 1:

- **Newer platforms now always use fallback enforced by software.** All newly created databases, users, indexes, and tables are set to FALLBACK, even if you specify NO FALLBACK in the CREATE request. All databases or users are set to FALLBACK even if you specify NO FALLBACK in the ALTER TABLE or MODIFY DATABASE/USER request.
- **Older platforms that upgrade to this release still allow the NO FALLBACK option** and the default is NO FALLBACK for the CREATE TABLE, ALTER TABLE, CREATE JOIN INDEX, CREATE HASH INDEX, CREATE DATABASE, MODIFY DATABASE, CREATE USER, and MODIFY USER requests.

Fallback and Clustering

When a table has FALLBACK protection, Teradata Database stores two copies of every row in the table, a primary copy and a fallback copy. The fallback copy is stored on a different AMP than the primary copy. This protects against failure of the AMP storing the primary copy. AMPs are assigned to clusters, in which the AMPs provide fallback protection for each other. A cluster must consist of at least two AMPs and is limited to at most eight AMPs.

Function Mapping for Native Functions

You can use a function mapping to specify a simple name for executing a function or table operator within a database or user, in addition to executing a function on an external server. For CREATE FUNCTION MAPPING and REPLACE FUNCTION MAPPING:

- The SERVER clause is now optional, that is, not specified for a function within a database or user.
- You can now specify multiple ANY IN TABLE clauses.

Function Mapping Variable Substitution

You can now include variables in function mapping definitions, such as:

- User variables.
- System variables.
- Scalar subquery expressions (SSQs)
- Concatenated variable expressions

Legacy versus Dedicated Fallback

Teradata Database MAPS Architecture introduces a new *dedicated fallback* scheme. In this scheme, one AMP in the cluster stores all fallback rows for one other AMP in the cluster. For example, consider a four-AMP cluster (AMPs 0, 1, 2, and 3). If AMP 0 has nine rows, AMP 1 has all nine fallback rows for AMP 0; AMPs 2 and 3 do not have any fallback rows for AMP 0. Similarly, AMP 2 provides fallback protection for AMP 1, AMP 3 for AMP 2, and AMP 0 for AMP 3.

In prior versions of Teradata Database, *legacy fallback* was used. In this scheme, each AMP in a cluster contains fallback copies of primary data from all the other AMPs in the cluster. For example, consider a four-AMP cluster (AMPs 0, 1, 2, and 3). If AMP 0 contains nine rows of primary data, the fallback copies of those rows are spread among AMPs 1, 2, and 3.

New maps created on systems upgraded to Teradata Database 16.10 or later use dedicated fallback; however, tables that still use existing maps from releases prior to 16.10 use legacy fallback. For systems that use two-AMP clusters, dedicated and legacy fallback operate identically.

Heatmap Table Function Integration and Enhancements

If you installed the `td1410_heatmap` function, this older version will be deleted from the system when you upgrade to Teradata Database 16.00 or later. If you have scripts that reference `td1410_heatmap`, you must rewrite them to use the new function name, `tdheatmap`.

Note: Users cannot install the heatmap function from Developer Zone after upgrading to Release 16.00 or later.

IPE Enhancements: Cache Dynamic Plans (CDP) for IPE Statistics Feedback - Nonparameterized Requests

- Behavior prior to Release 16.20 Feature Update 2: IPE plan is not cached.
- Behavior change: IPE plan with only statistics feedback is cached and reused.
- IPE plan with results feedback is still not cached.

Java Runtime Environment (JRE 1.8)

Java 8.0 is required before installing Teradata Database 15.10.01 or later. Users can now run Java UDFs and external stored procedures compiled with JDK 8.0 on Teradata JRE 1.8

JSON Web Token (JWT)

TDNEGO is not enabled for JWT in this release.

Link-local IP Address Restrictions

In Teradata Database 16.00 or later, IPv4 and IPv6 link-local IP addresses (169.254.0.0/16 for IPv4 and fe80::/10 for IPv6) are blocked from connecting to the database. Once the link-local restrictions are configured (by upgrading or installing Release 16.00 and later), backing down will not remove the restrictions. If link-local IP addresses are needed, they must be manually allowed.

If the upgrade or installation detects that the Teradata Database system is currently using ipfilters, the link-local restriction will not be imposed, and a warning message will advise adding the link-local restriction manually.

For instructions on modifying the link-local IP address configuration, see *Teradata Vantage™ NewSQL Engine Security Administration*, B035-1100.

Minimum Bounding Box (MBB) and Minimum Bounding Rectangle (MBR) Output

Previously, MBR and MBB output had an extra terminating null character '\0' as part of the output. The extra null character is no longer output.

Packageless TeraGSS

TeraGSS is now embedded in client drivers, eliminating the need for TeraGSS installation and configuration. If there is a need to configure TeraGSS on a client, see the appendix regarding the Teradata GSS Administrative Package in *Teradata Vantage™ NewSQL Engine Security Administration*, B035-1100. Note: Teradata recommends not configuring TeraGSS.

The tdgssversion tool has been removed.

Teradata Unity uses TDGSS instead of TeraGSS. For information about Unity, see *Teradata Unity Installation, Configuration, and Upgrade Guide for Customers*, B035-2523, and *Teradata Unity User Guide*, B035-2520.

Replace CFS Hard Limits with PSF Hard Limits

The Teradata Priority Scheduler (PSF) was built on top of the SLES 11 Completely Fair Scheduler (CFS). In previous releases, there was no way to make a task immune to CPU throttling if quotas were met. The Priority Scheduler had to make a task real time to prevent it from being throttled.

As of Release 16.20 Feature Update 2, Priority Scheduler implements its own CPU throttling scheme. Tasks running in the internal system workload, such as PDE daemon tasks, are never CPU throttled. Tasks holding critical internal resources are not throttled. To prevent total PDE CPU usage from

exceeding Workload Management Capacity on Demand (WM COD), the CPU used by the system workload reduces the amount of CPU available to other workloads.

The benefits are:

- Better control: Task-level CPU throttling gives Priority Scheduler more control of when and how long a task is throttled. Priority Scheduler will only throttle at specific preemption points.
- Better overall system performance: System daemons can be made immune to CPU throttling.
- Visibility: Tasks are throttled at the task level, so you can examine statistics to determine the impact of hard limits or COD on a task, request, session, workload, and virtual partition.
- Process address lock contention eliminated: No CPU throttling of a task in the middle of a kernel call.

PSF can only CPU throttle those tasks that are visible to PDE. Similar to past behavior, CPU usage reported by the Resource Sampling Subsystem (RSS) SPMA table is different than CPU usage reported by RSS SPS and SVPR tables. This is because SPMA reports all OS CPU usage information from /proc, while tables like SPS/SVPR report only PDE-related usage.

For more information, see Carrie Ballinger, *Workload Management Capacity on Demand and Other Hard Limits*, Teradata Database 16.20 Feature Update 2, Teradata Database Orange Book TDN0009761.

For information about the RSS tables, see *Resource Usage Macros and Tables*, B035-1099.

Resource Usage Macros

Some ResUsage macros use Node ID as an input parameter to select the data for the desired node(s) and display the Node ID in the macro execution. For these macros, the Node ID parameter type is extended from CHAR(6) to CHAR(9) (from 'CCC-MM' to 'CCCC-MMMM') to accommodate the new node id format of a 4-digit cabinet number and a 4-digit module number. The affected macros are:

ResCPUByAMP	ResLdvByNode	ResPdiskOneNode
ResCPUByAMPOneNode	ResLdvOneNode	ResVdiskByNode
ResCPUByNode	ResMemMgmtByNode	ResVdiskOneNode
ResCPUByPE	ResMemMgmtOneNode	ResPsByGroup
ResCPUByPEOneNode	ResNetByNode	ResPsByNode
ResCPUOneNode	ResNetOneNode	ResAWTByNode
ResHostByLink	ResOneNode	
ResHostOneNode	ResPdiskByNode	

SHOW QUERY LOGGING FeatureUsage Field

The SHOW QUERY LOGGING statement output now includes the FeatureUsage field to indicate whether feature use logging is in effect. Feature use logging is specified using the WITH FEATUREINFO option of the BEGIN QUERY LOGGING and REPLACE QUERY LOGGING statements. For more information, see BEGIN QUERY LOGGING, REPLACE QUERY LOGGING, and SHOW QUERY LOGGING in *SQL Data Definition Language – Syntax and Examples*, B035-1144.

Teradata Database MAPS Architecture

Teradata Database MAPS Architecture, released in 16.10, changes what customers see in Teradata Database compared with prior releases, whether or not the feature is enabled.

The MAPS feature impacts the following:

- New dictionary tables and changes to existing dictionary tables to support maps.
- User tables are defined to use TD_Map1, which includes all the AMPs in the system. It is the system-default map.
- Data Dictionary tables are defined to use TD_DataDictionaryMap, except for nonhashed tables in database DBC, which are defined to use TD_GlobalMap.
- When MAPS is enabled, the Teradata Archive/Recovery utility (ARC) can archive/restore only objects contained in the all-AMPs contiguous map. For more information, see “Running Teradata ARC in a Multiple Hash Map (MAPS) Environment” in *Teradata Archive/Recovery Utility Reference*, B035-2412. If MAPS is enabled, use Teradata Data Stream Architecture (DSA) for archive/recovery operations. For more information about DSA, see *Teradata Data Stream Architecture (DSA) User Guide*, B035-3150.

If the MAPS feature is enabled, you cannot back down to a release prior to 16.10. For information on using MAPS, see *Database Administration*, B035-1093.

TDGSS Single Mechanism to Log-On Client Support

TDNEGO now supports the following clients, in addition to the CLI and ODBC clients:

- JDBC
- Windows .NET

For more information, see *Teradata Vantage™ NewSQL Engine Security Administration*, B035-1100.

TDWM UDF GDO Size

The GDO for TDWM.TDWMActiveWDS exceeds the size of the protected buffer, so this function is now unprotected.

Trigger Enhancement for Row Level Security Tables

This enhancement allows triggers to reference RLS-protected tables. Triggers on non-protected tables can update a protected table, and triggers on protected tables can update a non-protected table. Referenced tables in the triggers must have identical constraints.

Trusted Session Support in X Views

X views are enhanced to return results based on the current authorized user. Previously, results were not returned for proxy users granted access through query band. Now X views return results for the proxy user if the current authorized user is set to a trusted sessions proxy user. If the user is not a proxy user, the results returned are based on the access rights of the current user. For more information, see *Teradata Vantage™ Data Dictionary*, B035-1092.

Unicode Pass Through

Customers can now store and retrieve unsupported Unicode characters, including emoji and other ideographs. If enabled for a session, Unicode Pass Through lets users pass through Teradata unsupported Unicode characters and unassigned Unicode characters to the Teradata Database. Pass through characters include:

- BMP from Unicode versions 6.1.0 to 9.0.0, which Teradata does not support, including Emoji
- SMP from Unicode versions, which Teradata does not support
- Unassigned characters
- Private use characters

Note: Although pass through characters can be stored in Teradata Database, they are not fully supported. For example, collation, case sensitivity, and object name support are not included with this feature.

Noncharacter exceptions and invalid character encoding form (ill-formed code unit sequence) exceptions will not occur for a pass through session when importing from UTF-8 or UTF-16 sessions to the UNICODE server character set. They are changed to one or more replacement character(s), (U+FFFD), without consuming valid successor bytes as described in the Unicode Standard (see <http://www.unicode.org/versions/Unicode8.0.0>).

Unicode Pass Through is not supported with FastLoad, MultiLoad, Teradata Parallel Data Pump (TPump), FastExport, and OLE DB Provider for Teradata. Use Teradata Parallel Transporter (TPT), BTEQ, or the Teradata JDBC Diver to load and unload pass through data. If you do not specify Unicode Pass Through for a session, the default is OFF.

Because the internal representation of pass through characters is UTF-16, a supplementary character is represented in the UNICODE server character set as two 16-bit UTF-16 code units, or 4 bytes. So VARCHAR(2) or CHAR(2) are the minimum sizes required in order to store pass through characters.

Cleansing and Loading Data with Unicode Pass Through

You may have existing processes in place to cleanse Unicode data before loading it into Teradata Database, such as access modules and UDFs. To take full advantage of Unicode Pass Through, change or eliminate prior methods of cleansing Unicode data before loading. See the TPT documentation for details on configuring how TPT uses access modules. For details about configuring access modules to allow pass through characters, see *Teradata Tools and Utilities Access Module Reference*, B035-2425, particularly the sections regarding automatic character conversions.

For more information about Unicode Pass Through, see:

- *Teradata Vantage™ Database Administration*, B035-1093
- SET SESSION in *Teradata Vantage™ SQL Data Definition Language – Syntax and Examples*, B035-1144
- *Teradata Vantage™ NewSQL Engine Release Summary*, B035-1098
- *Teradata Vantage™ NewSQL Engine International Character Set Support*, B035-1125

Recompiling Stored Procedures

When you upgrade to or across a major Teradata Database release, for example, when upgrading or migrating to Release 16.20.xx from any previous release, you must recompile stored procedures. If PUT is used, it automatically recompiles stored procedures if the source is available on the system. It also generates a report of stored procedures without source (such as stored procedures originally compiled with the NOSPL option) that must be recreated. For more information, see *Teradata Database Node Software Upgrade Guide: Major/Minor*, B035-5943.

Installation, Upgrade, Migration, and Backdown (IUMB)

Teradata supports customer-performed maintenance and patch upgrades. Contact your sales or customer support representative for questions.

For changes in behavior that impact upgrade or migration, see [Changes in System Behavior](#).

Supported IUMB Operations

This release supports the following IUMB operations:

- Installation of this Teradata Database release on all supported platforms and operating systems.
- Migration from Teradata Database 15.00 and later. Migrations from previous releases require an intermediate migration.
- Upgrade to this release from the releases shown in Knowledge Article IDA00108C82, available from Teradata @ Your Service (<https://access.teradata.com>). To see the article, you must be logged into Teradata @ Your Service.

If your current Teradata Database version is not listed as an approved upgrade starting version, you must first upgrade to an approved starting version before upgrading to this release. Contact the Teradata Support Center for details. For information on upgrades from older releases, see [Upgrading from Older Releases](#).

IUMB Planning

- Upgrade scripts and the upgrade estimator tool are available in the PUTTools package. Always get the latest version. For all IUMB change controls obtain PUTTools from <https://access.teradata.com>. Go to Update Your Software and then Database and Applications.
- Teradata @ Your Service (<https://access.teradata.com>) provides access to copies of other items required for IUMB procedures such as:
 - The certified list of software packages for each supported Teradata Database version, including recently updated versions of software packages.
 - Required application and operating system software patches, firmware, drivers, service packs and hotfixes.
- You must upgrade your Teradata client software to at least the minimum supported release before or at the same time as you upgrade to this release.
- Some features are enabled by default during a sysinit when upgrading or migrating and may affect system behavior. To see if the current release is affected, see [Default Feature Status](#).

Upgrading Teradata Temporal Tables

Teradata originally introduced support for creating and manipulating temporal tables before an ANSI/ISO standard had been developed. Consequently, the original Teradata Temporal Tables and SQL syntax do not conform to the ANSI standard. If you upgrade from a Teradata Database release prior to

15.00 and you were using Teradata Temporal Tables, you can choose either to continue using them or convert to using ANSI standard temporal tables and syntax.

For more information on the differences and ramifications:

- See the “DBS Control” coverage in *Teradata Vantage™ Database Utilities*, B035-1102, and review the description of the Temporal Behavior DBS Control field.
- See the “ANSI Temporal Tables” coverage in *Teradata Vantage™ Temporal Table Support*, B035-1182.

Parallel Upgrade Utility (PUT)

Use PUT to install or upgrade Teradata Database and other software, as well as install and configure the Teradata Database.

- PUT is provided with each copy of this release, but you should download the latest version of PUT from <https://access.teradata.com>. Go to Update Your Software and then Database and Applications.

You can download *Parallel Upgrade Tool (PUT) Reference*, B035-5716, from <https://docs.teradata.com>.

Unicode Pass Through (UPT) IUMB

Upward Compatibility

When using this feature, users who rely on Teradata Database to screen out unsupported characters or the REPLACEMENT CHARACTER (U+FFFD) can no longer rely on Teradata to do so.

The hash function can handle pass through characters, including the hashing of surrogate code points (as of Unicode version 8.0). **Note:** In a future release when pass through characters are fully supported, the code points may be hashed differently.

Backward Compatibility

Because this feature stores characters that were previously unsupported, pass through characters on a Teradata system cannot be transported to an earlier release without replacing the pass through characters.

Prior to Unicode Pass Through, internal Unicode to internal Unicode translations of pass through characters stored a U+FFFD in the destination string. With Unicode Pass Through and for all sessions, including non-Pass Through Sessions, the translation of pass through characters stores, without loss, those same pass through characters in the destination string.

Prior to Unicode Pass Through, internal Unicode to external Unicode, such as, UTF-16/UTF-8 translations of pass through characters, stored a U+FFFD or EFBFBD, respectively, in the destination string. With Unicode Pass Through and for a pass through session, this translation of pass through characters will instead store, without loss, those same pass through characters in the destination string.

Upgrade and Migration

There are no issues.

Backdown

Backdown is not supported for Unicode Pass Through.

Replacing Unsupported Operating Systems

If your system runs on an unsupported OS (MP-RAS, SLES 9, SLES 10, or Windows), you must replace it with a supported version of SLES before upgrading or migrating to Teradata Database 14.10 or higher. For more information, see [Supported Software and Hardware](#).

Note: Installation of SLES 11 changes the workload management options available on the system.

Upgrading from Older Releases

Teradata Database 16.20 Feature Update 2 is a patch to Teradata Database 16.20, so Release 16.20 must be installed first before upgrading to Release 16.20 Feature Update 2.

You can upgrade from Release 15.x to Release 16.20 in one step.

To upgrade from Release 14.10, perform a multi-stage upgrade process. This requires two upgrades:

- 14.10 to 15.10
- 15.10 to 16.20

All customers upgrading to Release 16.20 from Release 14.10 automatically receive a copy of the necessary intermediate versions of Teradata Database to use during the upgrade process. Customers are licensed to use the intermediate software CD only as part of the upgrade process.

After the upgrade is complete, customers should dispose of the intermediate software CD, while retaining the Release 16.20 CD.

Contact the Teradata Support Center if you are upgrading from a version that is more than two major releases back.

About Returning to an Older Release

Although moving to a new Teradata Database release is automated, there is no automated way to reverse the process and the required conversions to move to a previous release. Backing down across a major release (xx.0), such as Release 16.00, is not supported. You can back down from Release 16.20.xx to Release 16.10 if the DBS Control NoDot0Backdown flag was not set to True to enable Teradata Database MAPS Architecture.

System Performance

Performance Regressions

Any regressions that Teradata identifies, either by further testing or in field-deployed systems, are fixed as soon as possible. To find out the latest information about performance regressions that have been identified for the new release and how they may affect your system, see: <https://access.teradata.com>.

Running Teradata Database with Other Applications

Other applications (including Teradata applications) may execute concurrently with Teradata Database on approved system platforms. However, this is not encouraged as it may negatively impact the database:

- Throughput and response time performance
- Availability

It is strongly recommended that you do not run applications that are large consumers of system resources (such as other databases) concurrently without understanding the performance and availability impact to both the applications and Teradata Database.

If you do run applications on the same system or node as Teradata Database:

- The system or node may need additional hardware (for example, memory) to support the applications.
- Monitoring and tuning the system may be more complex.
- Compromises (for example, in the settings of tuning parameters) may be required to provide satisfactory and consistent performance for both Teradata Database and applications.
- If a problem does occur, it may be necessary to determine whether the problem also occurs in isolation.

If an application does not run properly or interferes with Teradata Database, it may be necessary to move it to another node or system. For instance, such applications may:

- Require a different version of the operating system
- Require a different set of operating system-level patches
- Require different settings of tuning parameters
- Be unable to obtain adequate system resources or obtain too many system resources due to the UNIX scheduler or other OS-specific resource limitations
- Adversely affect performance

- Require significant use of BYNET bandwidth
- Cause a UNIX failure or a Teradata Database restart

By default, the Teradata memory allocation algorithms are based on Teradata Database using 100% of the memory on a node. If other applications use a significant amount of node memory, you may need to add memory, and you should adjust the option controlling this percentage. For these reasons, avoid running non-Teradata applications on nodes running Teradata Database, if possible, as shown in the table that follows.

Software Type	Runs on Nodes That Run Teradata Database?	Considerations
Non-Teradata applications	No	Run on nodes that do not run Teradata Database, so that: <ul style="list-style-type: none"> • Expected throughput, expected response time, and parallel efficiency are not impacted. • Detrimental impact on the system is reduced. • Problems can be more easily isolated. • Fixes needed by one application can be made without having to apply them to nodes that do not need those fixes (or for which those fixes are detrimental).
Teradata applications and Teradata client software	Yes	Run on nodes that run Teradata Database if the software: <ul style="list-style-type: none"> • Puts a very small load on a system. • Evenly distributes the workload across the nodes. • Is used periodically for system maintenance. • Has an impact that is well-understood and acceptable.

Note: Even if applications are run on separate nodes (non-TPA nodes), they may still share the BYNET and thereby potentially interfere with Teradata Database.