

Teradata Vantage Database Engine 20: New Features, Design, & Tuning Tips

TERADATA





Table of Contents

Overview	3
Introduction to Some New Features: (Teradata Vantage v20.00)	3
Native Open Table Format Support	3
Modernized Operating System Foundation	4
Analytics Database 17.20 End-of-Life (June 2027)	4
Enterprise Cluster Architecture	4
QueryGrid & Cross-Engine Connectivity	5
Enhanced AI / ML (Bring Your Own Model)	6
Some Design and Tuning Tips	6
Tuning Native Iceberg and Delta Lake Workloads	6
The Primary Index (PI) Choice Still Matters	7
Statistics on OTF and Native Tables	7
Query Redesign Across Hybrid Workloads	7
Planning the 17.20 to 20.xx Upgrade	8
XTIVIA's Proprietary Tools & Frameworks	8
MVC Tool	8
Table Skew Analyzer	9
Object DDL Comparison Tool	9
Tuning Environment Creator	10
OTF Readiness Assessor	10
Summary	10

Overview

Teradata Vantage v20.00 marks a shift with native support for Open Table Formats like Iceberg and Delta Lake, bypassing the need for data duplication. Built on a modernized SUSE Linux Foundation, this version enhances AI/ML and enterprise cluster architectures. As the 17.20 version reaches EOL in June 2027, organizations should plan upgrades, carefully tuning OTF workloads and maintaining precise statistics. XTIVIA's suite of proprietary tools, including an OTF Readiness Assessor and skew analyzers, helps teams navigate these architectural transitions and optimize database performance in hybrid environments.

Introduction to Some New Features: (Teradata Vantage v20.00)

Native Open Table Format Support

One of the most significant enhancements in Teradata Vantage 20.x is native support for Open Table Formats (OTF). Database Engine 20 enables Vantage to read and write Apache Iceberg and Delta Lake tables stored in cloud object storage such as Amazon S3, Azure Data Lake Storage (ADLS), and Google Cloud Storage (GCS)—without copying the data into Teradata-managed storage.

This closes a long-standing gap between Teradata and lakehouse-native platforms and enables hybrid data architectures without duplicating data.

Modernized Operating System Foundation

Teradata Vantage 20.x is built on SUSE Linux Enterprise Server (SLES) 15, replacing the legacy SLES 12 SP3 platform used by Analytics Database 17.20. This strengthens the platform's security posture and enables faster operating system patching, which matters more as OTF workloads pull in a wider set of open-source client libraries.

Analytics Database 17.20 End-of-Life (June 2027)

Teradata has set an end-of-life date of June 30, 2027, for Analytics Database 17.20. Only critical maintenance fixes will be provided until the EOL date. Customers still on 17.20 are expected to have a documented upgrade plan reviewed with Teradata's Global Change Control team, or an approved exception on file. New deployments of 17.20 are no longer available through cloud marketplaces starting in Q3 2026, which effectively makes v20.xx the default target for any new system.

Enterprise Cluster Architecture

Native Open Table Format functionality is delivered through the new Enterprise Cluster Architecture, the deployment model designed to unlock direct engine access to open table formats outside Teradata's traditional storage layer. This is a meaningful architectural shift from prior versions, where connecting external data required staging it into Teradata storage or bridging it through QueryGrid foreign server connections.

QueryGrid & Cross-Engine Connectivity

QueryGrid continues to serve as the connectivity layer for engine-to-engine access, handling logging, heartbeat monitoring, and configuration while the Vantage Database Engine runs the actual computation. With native OTF support, QueryGrid is no longer required when querying Iceberg or Delta Lake datasets directly. Instead, QueryGrid continues to provide value for cross-platform federation where native OTF access isn't available.



Enhanced AI / ML (Bring Your Own Model)

Database Engine 20 continues expanding in-database AI capabilities.

New capabilities include:

- Improved ONNX model support
- Binary input support for ONNX models (BLOB/VARBYTE)
- ONNX model property inspection
- Additional NLP functions such as ChunkText and NerReplace
- Expanded Spark MLeap model support

Some Design and Tuning Tips

Tuning Native Iceberg and Delta Lake Workloads

As of v20.00, native Open Table Format access utilizes the Apache Iceberg and Delta Lake Java APIs. Compared to traditional Teradata-native tables, these workloads may introduce additional CPU overhead, so size OTF-heavy systems with that overhead in mind. Beyond engine-level sizing, each OTF project has its own performance levers, indexing, data skipping, and file-level statistics, and those should be tuned according to the OTF vendor's own documentation rather than assumed to behave like a native Teradata table.

The Primary Index (PI) Choice Still Matters

While Vantage 20.x introduces new storage capabilities, Primary Index design remains fundamental for Teradata-managed tables, join performance, and data distribution. Skewed or poorly chosen PIs still produce the same hot-AMP symptoms they always have; OTF adoption changes where some data lives, but it does not change the fundamentals of PI design for data that stays inside Vantage-managed storage.

Statistics on OTF and Native Tables

Native tables continue to depend on accurate, current statistics for the optimizer to pick a good plan, and that discipline doesn't relax with OTF in the mix; if anything, hybrid queries that join a native table against an Iceberg or Delta table need both sides of the join well-characterized so the optimizer can choose a sensible join strategy across two very different storage models.

Well-maintained statistics improve join selection, partition elimination, and query execution efficiency.

Query Redesign Across Hybrid Workloads

Queries spanning Teradata-managed tables and Open Table Format datasets should be carefully designed to minimize unnecessary data movement.

Queries that span native Vantage tables and OTF tables benefit from the same query-rewriting discipline as any cross-system join: push filtering as close to the OTF source as possible, avoid pulling wide OTF scans into Vantage only to filter afterward, and validate the plan shows partition pruning or file-level skipping happening on the OTF side rather than a full table scan.

Planning the 17.20 to 20.xx Upgrade

Given the June 2027 end-of-life date for 17.20, upgrade planning should start now, not close to the deadline. That means inventorying which workloads would benefit from native OTF access, validating that existing PPI/MLPPI and secure-zone configurations carry forward cleanly, and testing the OS-level change (SLES 12 SP3 to SLES 15) against any custom scripts or tooling that assume the old base image.

XTIVIA's Proprietary Tools & Frameworks

MVC Tool

An automated DBA tool for performing Multi-Value Compression (MVC) to save space and improve performance. Teradata doesn't provide any tool out of the box to define or create a list of values to be compressed for tables. Collecting candidate column values for compression requires manual effort.

Not all data types can be compressed, and there is a limit to the number of values that can be compressed. XTIVIA's tool analyzes given tables and columns for the possibility of MVC and performs the MVC upon confirmation.

Table Skew Analyzer

An interactive command-line tool to verify the skew and create modified DDL at runtime without changing the Primary Index physically on the table. This tool identifies table skew, lets the DBA choose new columns at runtime, and generates the DDL for deployment. Skew can be analyzed for all tables in a database in one go, for a particular table, or at a given skew percentage. The tool does not change the PI on the analyzed table, but it emails the new DDL (with the updated PI).

Object DDL Comparison Tool

This automated tool identifies DDL changes between different environments. Often, the DBA identifies performance issues related to PI/PPI/Statistics being out of sync between environments. ETL developers also find quality issues when objects are out of sync. This tool identifies differences in DDL, Primary Index, Partition Primary Index, Secondary Index, and Statistics. The tool generates new DDL and a detailed report, helping to sync up the environments.

Tuning Environment Creator

When approached for SQL performance tuning, DBAs usually spend a lot of time setting up the objects, statistics, and data without impacting existing data and ongoing development and testing. This tool automates this entire process and saves DBA time by automatically creating the required objects related to offending queries and copying data in a separate environment. With this tool, a DBA can focus on tuning exercises instead of spending time copying objects and setting up environments.

OTF Readiness Assessor

A newer addition to the toolkit, built for the v20.00 upgrade cycle: inventories candidate tables for native Iceberg or Delta Lake conversion, flags workloads likely to see the Java-API CPU overhead discussed in Section 2, and produces a prioritized migration backlog for OTF adoption.

Summary

Teradata Vantage v20.00 introduces native Open Table Format (OTF) support for Iceberg and Delta Lake, a modernized SUSE Linux Foundation, and improved AI/ML capabilities. This update necessitates proactive upgrade planning ahead of the 17.20 EOL on June 30, 2027. Key design tips emphasize tuning OTF workloads, maintaining rigorous statistics for hybrid queries, and respecting fundamental Primary Index design. XTIVIA offers tools to streamline these transitions, including an OTF Readiness Assessor, MVC automation, skew analysis, and environment synchronization utilities to optimize database performance.

*"Your brand experience rests on the technology that delivers it.
Enhance your competitive advantage through the right technology.
The future of your business depends on it."*

Dennis Robinson
XTIVIA Chief Executive Officer

<https://www.xtivia.com>

