

Teradata[®] Vantage 1.3 Release Summary

Deployment Platform: Teradata IntelliFlex™

August 2020

Copyright and Trademarks

Copyright © 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: 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.

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.

Advanced SQL Engine (was NewSQL Engine) 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.

The following lists the fixed and known issues in this release. If you experience any of the following issues, open an incident with Teradata Customer Support and include the Reference ID in your description.

Compatibility Matrix

For component compatibility information:

1. Go to support.teradata.com
2. Login
3. Search for KB0033995

Key Features

The Vantage 1.3 release contains a major update to the Machine Learning Engine where new HyperSegment functions enable customers to produce millions of models in parallel at scale with high performance, as well as performance and latency improvements.

Key Features:

- Machine Learning Engine Backup and Restore Automation
- Machine Learning Performance Enhancements:
 - Multi-threaded connection to vWorks - No customer impact
 - Proper re-use of Postgres connections - No customer impact
- HyperSegment functions: GLMPerSegment, CLMPerSegmentPredict, DecisionForestPerSegment, DecisionForestPerSegmentPredict
- ApproximateCloseness
- Support Variable Inflation Factor (VIF) for multicollinearity detection in Correlation function
- Support regression in XGBoost
- QueryGrid Containerization Standards
- NaiveBayes TextClassifier2
- Data Mover dropped support for ARC and Aster
- Data Mover support for Surplus Hadoop Nodes and in-place DSC upgrades
- Data Mover SQLE 17.00 support - Foreign Table NOS
- Vantage Analyst UI enhancements for Path
- Vantage Analyst Guided Tour (Path, Rules, Workflow)

Fixed Issues

Teradata QueryGrid

Reference ID	Description
QUERYGRID-13447	Fabric: Initiator nodes with high CPU during link bandwidth test was observed
QUERYGRID-13434	QGM: Nodes falsely reported offline due to ElasticSearch "No subject alternative DNS name matching" error
QUERYGRID-13134	QGI: KeyedObjectPool can destroy connection used by another request
QUERYGRID-12824	QGM: import doesn't work when a node is using a non-default network interface as the proxy network
QUERYGRID-12599	QGL: queries are failing with SHM alloc failure due to incorrect overall SHM memory validation
QUERYGRID-12542	Swagger doc: filterByProxySupport should be under GET /config/systems instead of GET /config/bridges
QUERYGRID-12519	Node: Existing rotated logs files are not pushed on tdqg-node restart
QUERYGRID-12469	QGM: "completed queries" portlet is showing "success" for failed "SQLs"

Teradata AppCenter

Reference ID	Description
UDAPP-9276	Keycloak Event Listener fails due to missing Environment Variables
UDAPP-9251	Fixed issues with CoreDNS pattern

Known Issues

Machine Learning Engine Analytic Functions

Reference ID	Description
ANLY-10612	Description: In VectorDistance, ORDER BY was mandatory in the ON clause for the reference table if InputTablesSame argument was true. Workaround: A performance optimization has been done and ORDER BY is not mandatory anymore for the reference table for this specific case. Using it will work fine as before but it will slow down the query unnecessarily. Deployment: All
ANLY-10226	Issue: For XGBoost function, if sparse format is used for input dataset, the function may fail. Workaround: add UniqueID() argument when sparse format is used in XGBoost function

ANLY-9660	<p>Description: Output table column names constructed from an argument value should satisfy the following conditions: 1) use only lower case, 2) not use special characters other than underscore, 3) not start with a digit, 4) not use reserved keywords. The following functions and arguments are affected:</p> <ul style="list-style-type: none"> - PathStart: PartitionNames - JSONParser: Nodes, NodeIdOutputColumn, ParentNodeOutputColumn, ErrorHandler - Pack_MLE: OutputColumn - Unpack_MLE: OutputColumns - Unpivoting: AttributeColumn, ValueColumn - URIUnpack: Queries - XMLParser: Sibling, Nodes, OutputColumnNameId, OutputColumnNameParentName, OutputColumnNameGrandparentName, ErrorHandler - LARPredict: CoefPositions - GTree: Results, EdgeResults <p>When an output column name is taken directly from one of the above arguments, an error is thrown indicating that the column name is invalid so that the user is aware of the invalidity of the specified name. When an output column name is formed by concatenating one of the above arguments with some additional string, the function will convert the name to one that satisfies the above conditions: 1) replaces upper case by lower case, 2) replaces special characters by underscore, 3) replaces initial digit by underscore, 4) appends underscore to the reserved keyword.</p> <p>Workaround: do not use a string value in any of the affected function arguments that violates any of the conditions in the Description.</p> <p>Deployments: All</p>
ANLY-8244	<p>Issue: For KNN function, automatic tuning of PartitionBlockSize might not be optimal.</p> <p>Workaround: Manually tune value of PartitionBlockSize.</p>
ANLY-6958	<p>Issue: If an error message exceeds 256 characters, it is truncated to 256 characters.</p> <p>Workaround: None.</p>
IETTU-16966	<p>Observations for the 2 hypersegmentation functions GLMPerSegment and DecisionForestPerSegment:</p> <ol style="list-style-type: none"> 1. The function can handle a very large number of segments, up to millions. 2. The function massively parallelizes the processing of the segments. 3. Each segment individually, should fit in memory, otherwise the function will gracefully fail with an error message. 4. The partitioning key should not be in the target columns neither partially nor in its entirety. 5. For GLMPerSegment, when numeric columns have different ranges, it is recommended to set FeatureScale as true. As a side effect, this may help to accelerate the convergence.

Advanced SQL Engine Analytic Functions

Reference ID	Description
TDAF-2531	<p>Description: when the input data is large, consecutive query runs of the DecisionForest function can sometimes return different results. The reason being that after the first query execution, Java JVM needs to do GC (garbage collection) to release memory. If the GC hasn't finished and the second query run is launched, there may not be enough memory and only the portion of data loaded into memory will be used by the DecisionForest execution. As a consequence, the results from the two executions are different.</p> <p>Workaround: Wait a few minutes for GC to finish</p> <p>Deployment: All</p>
TDAF-2450	<p>Description: The issue occurs when one or more input table columns is named after MLE reserved keywords (see corr_keywords.txt). The user will get a syntax error which points to some column names. The issue occurs because an internal query misses enclosing the reserved column names within single quotes (which in turn require escaping).</p> <p>Workaround: rename the columns named after reserved keywords so that they do not use any of the 141 reserved keywords.</p> <p>Deployment: All</p>

DBSQ

Reference ID	Description
DBSQ-3762	<p>Description: Error messages show old, nonstandardized argument and table names.</p> <p>Workaround: For old names that appear in error messages and their corresponding new names, see Teradata Vantage™ Machine Learning Engine Analytic Function Reference, B700-4003.</p>

Teradata AppCenter

Reference ID	Description
UDAPP-8661	<p>Description: Customer will need to delete the malformed prometheus data to resolve errors in thanos compactor</p> <p>Workaround - Remove the corrupted blocks and restart thanos compactor</p>
UDAPP-8648	<p>Description: Ambassador needs to be restarted once new certificates are installed. This issue is intermittent.</p> <p>Workaround: Restart ambassador pods, only if the browser does not show updated certificates after install.</p>
UDAPP-8601	<p>Description: Apps with permissions revoked are visible to user, but if clicked it will throw permission error.</p> <p>Workaround: None.</p>
UDAPP-8552	<p>Description: Multibyte character app names do not work.</p> <p>Workaround: None.</p>
UDAPP-8270	<p>Description: Scheduled and Manual backups fail if Postgres data size is very large. If node does not have twice the space that Postgres has, backup fails with OOM or Pod Evicted.</p> <p>Workaround: Free up space in /var/lib/docker mount on machine where backup pods run. The space in this folder must be twice the size of the Postgres data.</p>

UDAPP-8206	<p>Description: Execution of OS commands is blocked from BTEQ apps. The . OS directive on BTEQ apps does not execute, but job shows status as successful.</p> <p>Workaround: Do not rely on job status when using BTEQ apps with . OS directive. Instead, see logs of apps, which display error messages related to failure in command execution.</p>
UDAPP-8119	<p>Description: Postgres fails to store large results.</p> <p>Workaround: Reduce size of query or split query into multiple parts.</p>
UDAPP-7789	<p>Description: Parsing fails for parameters with double hyphens.</p> <p>Workaround: None.</p>
UDAPP-7192	<p>Description: Service accounts in AppCenter are not backed up by Scheduled or Manual backup.</p> <p>Workaround: Manually recreate all service accounts in AppCenter after restore.</p>

Teradata Viewpoint

Reference ID	Description
VP-50858	<p>Description: Upgrade to Tomcat 9.0.31 to address the following high (CVSS >=7.0) security risks:</p> <ul style="list-style-type: none"> * CVE-2020-1938 * CVE-2020-1935 (not yet rated) <p>Workaround: N/A</p> <p>Ease of exploitation:</p> <ul style="list-style-type: none"> * CVE-2020-1938: This only affects the AJP protocol connector which we do not use and do not have enabled. It is a serious vulnerability, but not for Viewpoint. * CVE-2020-1935: Very difficult. From the CVE, "a possibility of HTTP Request Smuggling if Tomcat was located behind a reverse proxy that incorrectly handled the invalid Transfer-Encoding header in a particular manner. Such a reverse proxy is considered unlikely." <p>Deployments: All</p>