

Teradata[®] Vantage Release Summary

Release 1.1.0.3

October 2019

Copyright and Trademarks

Copyright © 2000 – 2019 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.

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.

Fixed Issues

Teradata Machine Learning Engine

Key	Release Note Description	Release Definition Description
ANLY-9076	GLML1L2 was taking a long time to complete in AWS	

GSC Tools

Key	Release Note Description	Release Definition Description
GSCT-1595	Enhanced error messages to better qualify no logs found during request	

Teradata AppCenter

Key	Release Note Description	Release Definition Description
UDAPP-8511	App-service throwing requested route [/config] has not been mapped	This was due to a race condition in the code. Fixed as part of 1.8.7
UDAPP-8504	Deleted the corrupt blocks from the thanos-compact. This issue was resolved after we applied the workaround in the cluster.	
UDAPP-8481	Thanos compactor will not encounter out of memory issues.	
UDAPP-8476		Elasticsearch maintains 2 copies of data spread over 3 data nodes. When 1 elasticsearch data node goes down, ES master starts replicating that data to other data nodes. The default timeout is 1 min, which means if the data node does not come back up in 1 min, the data replication will begin. If we increase the delay time to 60 mins, it will give enough time for the node to come back up and elasticsearch can skip the data replication.
UDAPP-8475		We found that sometimes the logs don't get pushed from filebeat and filebeat has a lot of warning messages for large log lines. For. Eg: Some MLE printf's are 16452 characters long. We are not certain that this log is choking filebeat and preventing the log to be pushed, but applying the following workaround helped filebeat to recover and push logs to elasticsearch. Although this issue can happen anytime when there are large loglines, but we have noticed that this issue gets worse with multiple node reboot.
UDAPP-8416		Sometimes if filebeat starts before the kubernetes APIs are available, filebeat is unable to fetch kubernetes metadata for logs. Hence, it ships these logs in default (system-*) index. A fix for this issue was applied.
UDAPP-8374	AppCenter pods go into crashloop if consul registration fails	Added a monitoring job to check if the pods are crashing and then restarting the containers to fix registration issues.
UDAPP-8249	We fixed a bug that led to group permissions that weren't properly checked.	

UDAPP-8248	RI: Group users who are added as managers can now edit and run those apps and scripts.	We fixed a bug that led to group permissions that weren't properly checked.
UDAPP-8247	RI: When the date parameter for a job is changed at script run time, the job no	We fixed a bug that led to a failure when changing the dates for a job run.

Known Issues

Teradata Machine Learning Engine

Key	Release Note Description	Release Definition Description
ANLY-9322	The Modularity function examples use an optional output table, CommunityEdgeTable, that should only be used when the information is required, as it can cause the function to take a very long time to complete.	
ANLY-8929	Classification functions do not support class names that differ only by case.	Description: Classification functions do not support class names that differ only by case. For example, functions may treat classes named "DeKalb" and "Dekalb" as one class. Workaround: Distinguish class names by more than case.
ANLY-6958	If an error message exceeds 256 characters, it is truncated to 256 characters. There is no workaround.	
ANLY-6855	DecisionTree and AdaBoost functions sometimes generate no rows for model table.	Description: DecisionTree and AdaBoost functions sometimes generate no rows for model table. Several situations can cause DecisionTree and AdaBoost functions to generate no rows for model table. For example: <ul style="list-style-type: none"> • Data set has large skew. • Every row in data set has at least one NULL value. • MinNodeSize value is too large. Workaround: None
MLE-3838	Time zone may be different on ML Engine analytic nodes and NewSQL Engine.	Description: Analytic applications, such as ML Engine, use UTC time zone. NewSQL Engine may have been set to non-UTC time zone. Time zone must match between NewSQL Engine and the analytic applications; otherwise, data transferred between nodes is impacted. Workaround: Contact Teradata Customer Support for assistance.
MLE-3491	The MLE connector stats handler thread continues to use old connector password after it is changed. Workaround: Contact Teradata Customer Support for assistance.	
MLE-3405	ML Engine does not support QueryGrid link names that contain whitespace.	Description: ML Engine does not support QueryGrid link names that contain whitespace. Workaround: The administrator should not use whitespace in names when they create links between the ML engine and other components. Doing so causes an error in the Failure Detection and Restart capability of the ML engine.
MLE-3031	Some rows in MLE stats table could contain empty stats. No workaround available. The rows with empty stats are just noise but they could be confusing.	
MLE-2363	Occasionally there may be utilization stats which seem to be "missing". This is a limitation of the infrastructure providing the stats and the very occasional dropped data point is expected.	Description: Some data points for query-level utilization stats are missing from collected data. Workaround: This happens too rarely to have material impact on collected statistics.
MLE-2244	HELP FOREIGN FUNCTION calls for nPath@ and NTree are empty in some fields.	Description: For "HELP FOREIGN FUNCTION" calls for nPath@ and NTree, the following fields are empty: <ul style="list-style-type: none"> • Input Columns • Output Columns • Function Owner • Creation Time • Function Version • Interfaces Implemented Workaround: As db_superuser, use act to provide these grants: <pre>beehive> grant execute on function nc_system.ntree to public; beehive> grant execute on function nc_system.npath to public;</pre>

MLE-2240	CPUIOWaitCentiseconds collects stats by node, not by pod. CPUIOWaitCentiseconds node is displayed for each pod on which it is deployed.	Description: CPUIOWaitCentiseconds collects stats by node, not by pod. CPUIOWaitCentiseconds node is displayed for each pod on which it is deployed. Workaround: None
MLE-2220	PERM space size of Query Level Monitoring (QLM) database is created with 10 Mb for each AMP. PERM space is full if QLM queries fail with [ERROR]: No more room in database td_mle_db.	
MLE-1392	Stored procedures in pm database, such as pm.install_afile, fail if first master node in analytic cluster is unavailable.	Description: Stored procedures in pm database, such as pm.install_afile, fail if first master node in analytic cluster is unavailable. Workaround: Administrator must log into each TPA node using ssh, then change the file /home/tdatuser/.ssh/config to point to another node in analytic cluster.

Teradata QueryGrid

Key	Release Note Description	Release Definition Description
QUERYGRID-11020	In some cases, error message returned to end user lacks information about cause of error. No workaround. If you need additional information, contact Teradata Customer support to retrieve support bundle for failed query.	
QUERYGRID-10421	If metadata request fails, Viewpoint reports incomplete metrics for failed query until fabric cleans up query state. There is no workaround.	
QUERYGRID-10024	System health monitoring for Vantage systems (NewSQL Engine and ML Engine) may no longer work after upgrade if prior Fabric or Driver restart issues exist. Workaround: After upgrade, delete restart issues.	

Teradata Advanced SQL Engine

Key	Release Note Description	Release Definition Description
TDAF-287	Function arguments that specify multiple columns accept only lists of column names, not column ranges or a combination of column names and column ranges. There is no workaround.	
DBSQ-3762	Error messages show old, nonstandardized argument and table names. 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.	

Teradata AppCenter

Key	Release Note Description	Release Definition Description
UDAPP-8661	customer will need to delete the malformed prometheus data to resolve errors in thanos compactor	
UDAPP-8270	Scheduled and Manual backups fail if Postgres data size is very large.	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. 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.

UDAPP-8206	Execution of OS commands is blocked from BTEQ apps.	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. 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.
UDAPP-8119	Postgres fails to store large results.	Description: Postgres fails to store large results. Workaround: Reduce size of query or split query into multiple parts.
UDAPP-7192	Service accounts in AppCenter are not backed up by Scheduled or Manual backup.	Description: Service accounts in AppCenter are not backed up by Scheduled or Manual backup. Workaround: Manually recreate all service accounts in AppCenter after restore.

Teradata Vantage

Key	Release Note Description	Release Definition Description
VAN-31	If table being transferred from NewSQL Engine to ML Engine has VARBYTE column and type of corresponding ML Engine column is incompatible with VARBYTE, an error occurs.	Description: If table being transferred from NewSQL Engine to ML Engine has VARBYTE column and type of corresponding ML Engine column is incompatible with VARBYTE, error message says Found: bytea instead of Found: varbyte ; for example: NAIVEBAYESPREDICT: The column 'c_varbyte' specified in CategoricalInputs must be a member of one of the following SQL type groups: [INTEGER, STRING], Found: bytea Workaround: On ML Engine, change column type from VARBYTE to BYTEA by calling procedure TD_SYSFNLIB.QGExecuteForeignQuery .