Teradata® Vantage 1.1.6 Release Summary

Deployment Platform: Teradata Vantage on Azure

June 2020

B700-4000-070C
DOCS.TERADATA.COM

Copyright and Trademarks

Product Safety

Safety type	Description
NOTICE	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. Log in.
3. Search for KB0033995.

Key Features

- Support multiple VM types in same AKS cluster reducing costs of "Essential components":

  Existing Machine Learning Engine and Graph engine deployments only support one type of VM to be deployed in AKS cluster for "Essential components". With support for multiple types of VM, Managed App customers will be able to reduce infrastructure costs for their systems with Machine Learning and Graph engines.

- Support Machine Learning Engine in the same VNET as the Advanced SQL Engine to optimize AKS cluster configuration:

  Today, SQL Engine and ML Engine are deployed in different VNET peering. This incurs ingress AND egress charges for data that moves between the engines. We are simplifying the architecture by supporting ML Engine in the same VNET as the SQL Engine to optimize AKS cluster configuration. This will reduce infrastructure costs for customers.

- Enable JVM pooling:

  JVM pooling is a feature that enhances performance of the Machine Learning Engine. JVM pooling will be turned on by default from this release.

- Introducing Vantage Analyst

  Vantage Analyst provides a set of capabilities specifically designed for Business Analysts.

Fixed Issues

Machine Learning Engine

Reference ID	Description
MLE-5823	**Description**: CVE-2019-0211 In Apache HTTP Server 2.4 releases 2.4.17 to 2.4.38, with MPM event, worker or prefork, code executing in less-privileged child processes or threads (including scripts executed by an in-process scripting interpreter) could execute arbitrary code with the privileges of the parent process (usually root) by manipulating the scoreboard. Non-Unix systems are not affected. **Workaround**: None. Upgrade httpd to 2.4.41. **Deployment**: All
MLE-5822	This issue is outstanding for Vantage 1.1.1.1 and Vantage 1.1.2 **Description of the issue**: Threat record: CVE-2019-17571: https://nvd.nist.gov/vuln/detail/CVE-2019-17571 " Included in Log4j 1.2 is a SocketServer class that is vulnerable to deserialization of untrusted data which can be exploited to remotely execute arbitrary code when combined with a deserialization gadget when listening to untrusted network traffic for log data. This affects Log4j versions up to 1.2 up to 1.2.17. " **Workaround**: If we are unable to push this fix out, here's workaround provided by Apache: https://nsfocusglobal.com/apache-log4j-deserialization-remote-code-execution-cve-2019-17571-vulnerability-threat-alert/ " If users cannot upgrade to V2.8.2 or later for the time being, they can also prevent the socket port enabled by the SocketServer class in Log4j from being opened to the public network. " Since we do not expose socket port from this container to any public network (cluster is generally locked down), we should be fine with this issue for now. **Deployment** (affected platforms) All
MLE-5800	**Description**: During Black Duck scans, GCC and Python-devel package was found to be a security vulnerability. Hence it was needed to remove those from upgrade images. **Deployment (platform impacted)** Due to removal of those packages, Azure platform will be impacted. There is no impact on IFX/AWS. **Workaround** As a workaround, Azure upgrades would be to manually collect all MLE UDFs from old MLE and install them in the new MLE. **Deployment**: All

Server Management Software

Reference ID	Description
SM-29725	Update of Netty to resolve CVE-2019-3119 and CVE-2019-9512 vulnerabilities
SM-29354	Updating Spring Framework to resolve CVE-2020-5398 **Workaround**: CMIC is a closed system and does not allow outside software to run on it. External access to the host OS is very controlled to specific IP addresses/users. **Deployments**: All
SM-28655	Update Apache log4j for CVE-2019-17571 **Workaround**: CMIC is a closed system and does not allow outside software to run on it. External access to the host OS is very controlled to specific IP addresses/users. Apache workaround: https://nsfocusglobal.com/apache-log4j-deserialization-remote-code-execution-cve-2019-17571-vulnerability-threat-alert/ **Deployments**: All

Advanced SQL Engine

Reference ID	Description
SQLE-550	**Description**: ********** ********** 16.20.53.06 created 04/20/20 ********** ********** PDE: DR 191227 Priority 2 CIBT: PDE SLG-PSF - COD/HL test got hung on nb129 DR 191687 Priority 2 Info seg/fsys/mon is not priniting filename and line number as part of allocator information like puma -v option do. DR 193076 Priority 2 dbgcrnr is throwing SEGV when info seg is run under gdb on few dumps. DR 193091 Priority 2 System debugger is getting SEGV after issuing a kill command when the program is in initial stage(before unwind information is not constructed for the thread). DR 193447 Priority 2 Reset stuck in STOP/KILLTASKSDR TDBMS: 190061 Priority 2 Privileges on SYSLIB functions unnecessarily revoked on upgrade as the result of DROP FUNCTION statements in DIPDEM. DR 192350 Priority 2 Inconsistancy in restoring the data for CP with LOB tables when using online enabled backup DR 193061 Priority 2 3610 returning NULL SLOBs DR 193472 Priority 2 Deadlock involving an internal session that cannot be resolved automatically DR 193643 Priority 2 Aborting the TriggeredSP request while executing at parser side may lead to the restart with errorcode 3200 .

DR 193699 Priority 2
9124 AMP segment violation error occurred when XML memory limit crossed while processing the SP.
DR 193708 Priority 1
Mutliple input table operator with Cogroups returns incorrect number of rows
DR 193736 Priority 2
For sql commands that use function mapping/alias, TASM rules do not apply.
DR 193768 Priority 2
It is rare for a customer to run Interpretive EVL. However, users could experience the WARNING 3705 (when Compiled EVL runs out of memory in GNX) which would then store and run Interpretive EVL. It's up to the user to then follow the 'remedies' in the messages reference manual to eliminate the warning and store compiled EVL.
DR 193792 Priority 2
Aborting the DSA offline Job may cause the DBS hung if the parallelly any other DSA job is on going.
DR 193892 Priority 2
Errors are returned to secure zone related operations. Below are a couple of examples:

9869 Zone '' does not exist.

6704 Internal error: The ISF subsystem was requested to export a string to the client which is too large.
DR 193896 Priority 2
If we run DSA restore/copy job containing many DBS which uses the space more than the allocated, then it leads to physical disk space exhaustion.
DR 193918 Priority 1
DBC only restore job fails with 2693 error
DR 193929 Priority 2
During the DSA restore job, if user dropped some PPI tables after BUILD phase and before POSTSCRIPT phase,then job completed with complete_errors.
DR 193948 Priority 2
Snapshot dump for 3610 error is generated during the logon processing of a DSA restore job. Too many concurrent 3610 errors may escalate to a DBS restart.

DR 193956 Priority 2
7455 failure (Invalid Time Zone specified) may happen for query with AT timezone string.
DR 193959 Priority 2
Nullability may be set wrongly for fields in the outer table of an outer join using 2- step outer join.
DR 193991 Priority 1
Insert to JSON column less than 64k lead to corruption when we insert data more than 4096.
DR 194000 Priority 2
Parallel execution of RENAME TABLE SQL & DROP TABLE SQL on same subject table can lead to incorrect to 3610 for RENAME TABLE SQL
DR 194008 Priority 3
When "BEGIN ISOLATED LOADING" is initiated on tables with join indexes, 3610 is reported.

	DR 194014 Priority 2 SELECT sql can cause 3610 when >> SEL query with an UDF expression >> When reduced spools optimization kicks-in. DR 194015 Priority 1 Mutli-input tableOp stream return incorrect rows with single varchar in Part by field. DR 194016 Priority 2 When the table is dropped still trying to access that table header databaseid leads to SegV. DR 194017 Priority 2 During REVALIDATE, for BATCHRI on NOPI table can cause SEGV. DR 194023 Priority 2 Archived error table already had FLD4 on TD14 version when copied to TD16 version caused segV. DR 194047 Priority 2 The issue described in "Problem Descrip" comment is causing some issues with upgrades involving ?on-disk? interpretive EVL (in stored procedures, PPI tables etc..). DR 194053 Priority 1 Mutli-input tableOp stream return incorrect rows with single varchar in Part by field. **Workaround**: N/A **Deployments**: All

Teradata QueryGrid

Reference ID	Description
QUERYGRID-12986	Teradata Connector: datatype period(timestamp) is causing dbs restarts due to conversion of data
QUERYGRID-12976	Node: Unrestricted File Upload via Zip Slip
QUERYGRID-12909	QGM: Unrestricted File Upload via Zip Slip
QUERYGRID-12896	QGM: System hostname containing invalid domain name characters breaks ElasticSearch cross cluster search and causes nodes to be reported as offline
QUERYGRID-12822	**Description**: Upgrade jackson-databind to 2.9.10.2 * CVE-2019-20330 **Workarounds**: N/A **Deployments**: All
QUERYGRID-12624	QGM exception during migrate using backup file from newer QGM
QUERYGRID-12294	Node: QGLWatchDog crash or unresponsive after network change
QUERYGRID-12188	**Description**: Upgrade swagger to 3.24.2 * CVE-2019-17495 **Workarounds**: N/A **Deployments**: All

QUERYGRID-12187	**Description**: Upgrade jackson-databind to 2.9.10.1 * CVE-2019-16943, CVE-2019-16942, CVE-2019-17531 **Workarounds**: N/A **Deployments**: All
QUERYGRID-12147	QueryGrid queries are not reported in the Completed Queries and Query Monitor Viewpoint portlets
QUERYGRID-12121	Teradata Connector: Query_band cannot be parsed by MLE when profile includes a query_band
QUERYGRID-12017	**Description**: Upgrade commons-compress to 1.19 * CVE-2018-11771 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11909	**Description**: Upgrade jackson-databind to 2.9.10 * CVE-2019-14540, CVE-2019-17267, CVE-2019-16335 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11908	**Description**: Upgrade jackson-databind to 2.9.10 * CVE-2019-14540, CVE-2019-17267, CVE-2019-16335 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11906	Teradata Connector: Failure 7487 AMP step failure for select query when replace unsupported character used
QUERYGRID-11687	**Description**: Upgrade commons-beansutils to 1.9.4 * CVE-2019-10086 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11531	Node: Change of IP address for data node doesn't get reflected in QGM configuration
QUERYGRID-11528	**Description**: Upgrade jackson-databind version to 2.9.9.3 * CVE-2019-14379, CVE-2019-12384, CVE-2019-12814, CVE-2019-12086, CVE-2019-14439 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11527	**Description**: Upgrade jackson-databind version to 2.9.9.3 * CVE-2019-14379, CVE-2019-12384, CVE-2019-12814, CVE-2019-12086, CVE-2019-14439 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11416	**Description**: Upgrade cURL to 7.65.3 * CVE-2018-14618, CVE-2019-3822, CVE-2018-16839, CVE-2018-16840, CVE-2018-0500, CVE-2018-16842, CVE-2019-3823, CVE-2018-16890, CVE-2019-5436, CVE-2019-5443, CVE-2019-5481, CVE-2019-5482 **Workarounds**: N/A **Deployments**: All

QUERYGRID-11398	**Description**: Upgrade Bootstrap to 4.3.1 * CVE-2018-20676, CVE-2019-8331 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11397	**Description**: Upgrade Spring to 5.0.13.RELEASE * CVE-2018-15756, CVE-2019-3795 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11390	**Description**: Upgrade zlib to 1.2.11 * CVE-2016-9840, CVE-2016-9842, CVE-2016-9843, CVE-2016-9841 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11389	**Description**: Upgrade OpenSSL to 1.1.1c * CVE-2018-0732, CVE-2018-0734, CVE-2018-0737, CVE-2019-1563, CVE-2019-1547, CVE-2019-1552, CVE-2018-5407 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11387	**Description**: Upgrade protobuf to 3.5.1 * CVE-2015-5237 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11383	**Description**: Upgrade jackson-databind to 2.9.9.1 * CVE-2019-12384, CVE-2019-12814 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11317	**Description**: Upgrade ElasticSearch to 6.8.1 and remove Kibana * CVE-2018-17246, CVE-2018-3830, CVE-2018-3830, CVE-2019-7616 **Workarounds**: N/A **Deployments**: All
QUERYGRID-11088	Teradata Connector: For T2X, HELP foreign table fails with reading data from indic buffer error
QUERYGRID-11053	QGM: Boot did not recover after ElasticSearch was OOM killed by kernel
QUERYGRID-11020	**Description**: 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-11014	QGM: system health check does not complete
QUERYGRID-10891	Fabric: Error response during low shared memory conditions does not reflect true error condition
QUERYGRID-9920	Teradata Connector: importing CLOB data with invalid unicode characters does not return unsupported unicode character error

Teradata AppCenter

Reference ID	Description
UDAPP-8961	COREDNS rewrite rule does not work consistently on all IFX hardware machines. **Workaround**: Run - kubeclt edit cm -n kube-system coredns add below after ready section in Corefile rewrite stop { \n name regex <APPCTL_DOMAIN>.<APP-NAMESPACE>.svc.cluster.local.$ ambassador.td-platform.svc.cluster.local \n answer name ambassador.td-platform.svc.cluster.local.$ <APPCTL_DOMAIN>.<APP-NAMESPACE>.svc.cluster.local \n }

Known Issues

Data Stream Architecture (DSA)

Reference ID	Description
DSA-22112	**Description**: The Spring Framework is an application framework and inversion of control container for the Java platform. The Spring libraries 3.2.2 and 4.3.13.RELEASE contained the following security vulnerabilities: **Workaround**: Preventative Controls: The scale of impact is reduced because there are several mitigating controls in place, such as: 1. The user needs access to internal network. 2. The client device/server needs to have BARCmdline package installed or the barportlets package installed on Viewpoint to gain access to the DSC service. 3. In addition, Viewpoint User Authentication requires Viewpoint credentials to execute BARCmdline commands. 4. https protocol is available for DSA REST service. A valid CA certificate is required to invoke different endpoints of this service. 5. SSL connection type is also supported for ActiveMQ, which is the Message Queue the different DSA components use to communicate with one another. Detective Controls: 1. Unsuccessful attempts to exploit the vulnerabilities can cause the system to crash, prompting an alert when the particular system is down. Corrective Controls: 1. Disaster Recovery of DSC is available. If there is persistent code, etc. in the DSA projects, the user has an ability to wipe out their entire DSA environments and perform a fresh install using different passwords/credentials. Afterwards, the user can restore back the Repository data back to the state that they wish to revert back to. Compensatory Controls: 1. The attacker will need access to the Teradata Database in order to access/read data. Through DSA, the user can potentially see the database object names and types, but not the row data. 2. The credentials, etc. within the Job Plan sent over to the Teradata Database are encrypted. The data sent over to the different storage devices is encrypted as well. **Deployments**: All

DSA-21414	**Description**: Apache log4j is a Java-based logging utility. Apache log4j 1.2.14 and 1.2.17 contained the following security vulnerabilities: **Workaround**: Preventative Controls: The scale of impact is reduced because there are several mitigating controls in place, such as: 1. The user needs access to internal network. 2. The client device/server needs to have BARCmdline package installed or the barportlets package installed on Viewpoint to gain access to the DSC service. 3. In addition, Viewpoint User Authentication requires Viewpoint credentials to execute BARCmdline commands. 4. https protocol is available for DSA REST service. A valid CA certificate is required to invoke different endpoints of this service. 5. SSL connection type is also supported for ActiveMQ, which is the Message Queue the different DSA components use to communicate with one another. Detective Controls: 1. Unsuccessful attempts to exploit the vulnerabilities can cause the system to crash, prompting an alert when the particular system is down. Corrective Controls: 1. Disaster Recovery of DSC is available. If there is persistent code, etc. in the DSA projects, the user has an ability to wipe out their entire DSA environments and perform a fresh install using different passwords/credentials. Afterwards, the user can restore back the Repository data back to the state that they wish to revert back to. Compensatory Controls: 1. The attacker will need access to the Teradata Database in order to access/read data. Through DSA, the user can potentially see the database object names and types, but not the row data. 2. The credentials, etc. within the Job Plan sent over to the Teradata Database are encrypted. The data sent over to the different storage devices is encrypted as well. **Deployments**: All

DSA-20788	**Description** (same as DSA-22112): The Spring Framework is an application framework and inversion of control container for the Java platform. The Spring libraries 3.2.2 and 4.3.13.RELEASE contained the following security vulnerabilities: **Workaround**: Preventative Controls: The scale of impact is reduced because there are several mitigating controls in place, such as: 1. The user needs access to internal network. 2. The client device/server needs to have BARCmdline package installed or the barportlets package installed on Viewpoint to gain access to the DSC service. 3. In addition, Viewpoint User Authentication requires Viewpoint credentials to execute BARCmdline commands. 4. https protocol is available for DSA REST service. A valid CA certificate is required to invoke different endpoints of this service. 5. SSL connection type is also supported for ActiveMQ, which is the Message Queue the different DSA components use to communicate with one another. Detective Controls: 1. Unsuccessful attempts to exploit the vulnerabilities can cause the system to crash, prompting an alert when the particular system is down. Corrective Controls: 1. Disaster Recovery of DSC is available. If there is persistent code, etc. in the DSA projects, the user has an ability to wipe out their entire DSA environments and perform a fresh install using different passwords/credentials. Afterwards, the user can restore back the Repository data back to the state that they wish to revert back to. Compensatory Controls: 1. The attacker will need access to the Teradata Database in order to access/read data. Through DSA, the user can potentially see the database object names and types, but not the row data. 2. The credentials, etc. within the Job Plan sent over to the Teradata Database are encrypted. The data sent over to the different storage devices is encrypted as well. **Deployments**: All

Machine Learning Engine

Reference ID	Description
MLE-6526	**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 .
MLE-6355	**Issue Description**: The machine learning engine (MLE) includes its own set of three pods called consul-0 consul-1 and consul-2. If two of these re-start at the same time, they might never recover. One symptom of this is Queen pod stuck in a state like Init:2/14 . **Workaround**: In this case, MLE must be re-started by doing an appctl uninstall and install of the MLE chart. **Platform**: all

MLE-5120	**Description**: CVE-2013-1900: PostgreSQL 9.2.x before 9.2.4, 9.1.x before 9.1.9, 9.0.x before 9.0.13, and 8.4.x before 8.4.17, when using OpenSSL, generates insufficiently random numbers, which might allow remote authenticated users to have an unspecified impact via vectors related to the "contrib/pgcrypto functions". Other medium and low security risks which are related to above issue are CVE-2014-0060, CVE-2014-0066, CVE-2015-3165, CVE-2016-5424, CVE-2014-0067,CVE-2014-0062, CVE-2012-3488, CVE-2015-3167, CVE-2016-0773, CVE-2016-0768, CVE-2017-7484, CVE-2017-7486, CVE-2015-5289, CVE-2015-5288, CVE-2018-1115,CVE-2014-0063, CVE-2014-0065, CVE-2014-0064, CVE-2014-0061, CVE-2016-5423, CVE-2013-0255 and CVE-2017-14798. **Workaround**: Exploitation of PostgreSQL security issue can be Mitigated by network segmentation and firewall rules.There is no direct customer access to PostgreSQL rendering security issues unexploitable. **Deployment**: All
MLE-3997	**Description**: CVE-2013-0252: boost::locale::utf::utf_traits in the Boost.Locale library in Boost 1.48 through 1.52 does not properly detect certain invalid UTF-8 sequences, which might allow remote attackers to bypass input validation protection mechanisms via crafted trailing bytes. **Workaround**: Existing security issues in boost library are not exploitable due to the requirement of local access required within MLE services. There is no direct customer access to the Operating System and/or Containers rendering security issues unexploited. **Deployment**: All
MLE-3649	**Description**: CVE-2016-9013: Django 1.8.x before 1.8.16, 1.9.x before 1.9.11, and 1.10.x before 1.10.3 use a hardcoded password for a temporary database user created when running tests with an Oracle database, which makes it easier for remote attackers to obtain access to the database server by leveraging failure to manually specify a password in the database settings TEST dictionary. In MLE, This CVE is not applicable as MLE does not use Oracle. **Workaround**: Not applicable. **Deployment**: None. This CVE is not applicable for MLE **Description**: CVE-2014-0474: The (1) FilePathField, (2) GenericIPAddressField, and (3) IPAddressField model field classes in Django before 1.4.11, 1.5.x before 1.5.6, 1.6.x before 1.6.3, and 1.7.x before 1.7 beta 2 do not properly perform type conversion, which allows remote attackers to have unspecified impact and vectors, related to "MySQL typecasting." **Workaround**: Not applicable. **Deployment**: None. This CVE is not applicable for MLE **Description**: CVE-2015-5143: The session backends in Django before 1.4.21, 1.5.x through 1.6.x, 1.7.x before 1.7.9, and 1.8.x before 1.8.3 allows remote attackers to cause a denial of service (session store consumption) via multiple requests with unique session keys. **Workaround**: None. Upgrade Django to latest version (1.11.x) **Deployment**: All

MLE-3540	**Description**: CVE-2016-1234: Stack-based buffer overflow in the glob implementation in GNU C Library (aka glibc) before 2.7, when GLOB_ALTDIRFUNC is used, allows context-dependent attackers to cause a denial of service (crash) via a long name. Other issues that are related to above and glibc are CVE-2014-9402, CVE-2015-8779, CVE-2014-9761, CVE-2018-6485, CVE-2018-11236, CVE-2019-9169, CVE-2017-15670, CVE-2015-1472, CVE-2010-0015, CVE-2014-4043, CVE-2012-4412, CVE-2015-8778, CVE-2017-15804, CVE-2014-9984, CVE-2010-3856, CVE-2018-1000001, CVE-2010-0296, CVE-2017-1000366, CVE-2015-5277 **Workaround**: Existing security issues in glibc are not exploitable due to the requirement of local access required within MLE services. There is no direct customer access to the Operating System and/or Containers rendering security issues unexploitable. **Platforms impacted**: All
MLE-3491	**Description**: The MLE connector stats handler thread continues to use old connector password after it is changed. **Workaround**: Contact Teradata Customer Support for assistance. **Deployments**: All
MLE-3405	**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. **Deployment**: All
MLE-3031	**Description**: Some rows in ML Engine stats table may contain empty stats. **Workaround**: None. The rows with empty stats may be confusing, but do not affect anything. **Deployments**: All
MLE-2220	**Description**: 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. **Workaround**: https://techsupport.teradata.com/kb_view.do?sysparm_article=KB0026925 **Deployments**: All
MLE-1392	**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 **Deployment**: All analytic cluster.

Advanced Analytics Functions

Reference ID	Description
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-10087	**Issue**: SVMSparse has a backward incompatibility caused by making AttributeValueColumn mandatory. Being optional as it was before made no sense as it was defaulted to value 1 for all the attributes. **Workaround**: None.
ANLY-8534	**Description**: NaiveBayes is a new function that wraps the previous NaiveBayesMap and NaiveBayesReduce functions. We advise to use this function as it has a simpler syntax and other improvements. However, the previous nested syntax is still supported.

ANLY-8328	**Description**: The StringSimilarity_MLE function has 8 additional metrics: -OSA: OptimalStringAlignment -DL: Damerau-Levenshtein Distance -JACCARD: Jaccard Similarity -COSINE: Cosine Similarity -HAMMING: Hamming Distance -LDWS: Levenshtein Distance without Substitution -LCS: LongestCommonSubstring -SOUNDEXCODE: Soundex Code based Similarity (only for English strings).
ANLY-8244	**Issue**: For KNN function, automatic tuning of PartitionBlockSize might not be optimal. **Workaround**: Manually tune value of PartitionBlockSize.
ANLY-6958	**Issue**: If an error message exceeds 256 characters, it is truncated to 256 characters. **Workaround**: None.

DBSQ

Reference ID	Description
DBSQ-3762	**Description**: 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

Reference ID	Description
UDAPP-8661	**Description**: Customer will need to delete the malformed prometheus data to resolve errors in thanos compactor **Workaround** - Remove the corrupted blocks and restart thanos compactor
UDAPP-8648	**Description**: Ambassador needs to be restarted once new certificates are installed. This issue is intermittent. **Workaround**: Restart ambassador pods, only If the the browser does not show updated certificates after install.
UDAPP-8601	**Description**: Apps with permissions revoked are visible to user, but if clicked it will throw permission error. **Workaround**: None.
UDAPP-8552	**Description**: Multibyte character app names do not work. **Workaround**: None.
UDAPP-8270	**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	**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	**Description**: Postgres fails to store large results. **Workaround**: Reduce size of query or split query into multiple parts.
UDAPP-7789	**Description**: Parsing fails for parameters with double hyphens. **Workaround**: None.
UDAPP-7192	**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 Viewpoint

Reference ID	Description
VP-50858	**Description**: Upgrade to Tomcat 9.0.31 to address the following high (CVSS >=7.0) security risks: * CVE-2020-1938 * CVE-2020-1935 (not yet rated) **Workaround**: N/A Ease of exploitation: * 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." **Deployments**: All