

Esgyn Corporation

# Release Notes for EsgynDB R2.6.2



Published: October 2019  
Edition: EsgynDB Release R2.6.2

EsgynDB is based on the Apache Trafodion project – a webscale SQL-on-Hadoop solution enabling transactional or operational workloads on Hadoop. Trafodion builds on the scalability, elasticity, and flexibility of Hadoop, and further extends Hadoop to provide guaranteed transactional integrity.

The EsgynDB Enterprise and Enterprise Advanced Editions extend the functionality available in Apache Trafodion by including additional components such as EsgynDB Manager for managing the database, ADO.NET driver for Windows-based client applications, and features such as support for ORC and Parquet storage formats, Backup/Restore, multi-tenancy and better performance and stability.

This EsgynDB release is an incremental update based on release R2.6.1. It primarily includes defect repairs. Fixes shipped up to EsgynDB release R2.4.10, R2.5.2, R2.6.0 and R2.6.1 are included.

## Features

### R2.6.2

Category	Features
DB Manager	<ul style="list-style-type: none"> <li>Elevated user can change the DB Admin password using DB Manager settings page. A restart is not needed.</li> </ul>

### R2.6.1

Category	Features
Database	<ul style="list-style-type: none"> <li>Support for Hive table property skip.header.line.count</li> </ul>

### R2.6.0

Category	Features
Infrastructure	<ul style="list-style-type: none"> <li>Support for CDH 5.14 – 5.15</li> </ul>
Usability	<ul style="list-style-type: none"> <li>Support CASCADE/RESTRICT option for DROP TENANT operation</li> <li>In a Multitenant configuration, allow tenants to be dropped without removing dependencies</li> </ul>
Database	<ul style="list-style-type: none"> <li>Support for ROWNUM function</li> <li>Support Oracle CONNECT BY function</li> <li>Support Hive partitioned tables</li> <li>Support for Hive subdirectories</li> <li>Implicit data type conversion in DEFAULT clause in DDL</li> </ul>
Advanced	<ul style="list-style-type: none"> <li>Support for composite data type in Parquet</li> <li>Support for LATERAL VIEW OUTER UNNEST</li> </ul>
Performance	<ul style="list-style-type: none"> <li>Improved performance for EsgynDB metadata GET commands</li> </ul>

## Migration Considerations

Users currently on pre-EsgynDB R2.4 versions will need to migrate in two steps: first to R2.5, then to R2.6. The 2-step process is needed because there has been a change in EsgynDB SQL metadata in R2.6 which makes it incompatible with prior releases. This means once upgraded to R2.6, the system cannot be reverted back to a pre-EsgynDB R2.6 release without a re-initialize of EsgynDB database.

Users currently on EsgynDB R2.4 or R2.5 versions can migrate directly to EsgynDB R2.6 version.

### Migration from R2.4 or R2.5

The EsgynDB installers (whether command-line, or Cloudera Manager or Ambari-based) will automatically upgrade the SQL metadata. In case of a failure during the upgrade, the change will be rolled back to the existing release version. Prior to starting the upgrade, a backup should be taken of the EsgynDB SQL metadata.

Use the following command to backup the EsgynDB SQL metadata before upgrade (as `trafodion` user)

```
hbase shell createUpgradeSnapshots > createSnapshots.log
```

If an error occurs during upgrade, determine where upgrade fails and execute the appropriate commands.

`restoreUpgradeStep1` - restores backup snapshots

```
hbase shell restoreUpgradeStep1 > restoreSnapshots.log
```

`restoreUpgradeStep2` - verifies that the restored tables are good

```
hbase shell restoreUpgradeStep2 > verifyRestore.log
```

`restoreUpgradeStep3` - removes temporary files created during upgrade

```
hbase shell restoreUpgradeStep3 > removeTempTablesFromHbase.log
```

`restoreUpgradeStep4` - removes temporary tables from EsgynDB

```
sqlci < restoreUpgradeStep4 > removeTempTablesFromSQL.log
```

Once a successful upgrade occurs, you can remove the snapshots

```
hbase shell deleteUpgradeSnapshots > deleteSnapshots.log
```

Users upgrading from R2.4 may be affected by the following changes

## Application

1. SQL ALTER TABLE syntax has been enhanced to conform with ANSI SQL syntax.

### Previous syntax

```
ALTER TABLE <table> ALTER COLUMN <column> <datatype>
```

### New syntax

```
ALTER TABLE <table> ALTER [COLUMN] <column> [SET DATA TYPE] <datatype>
```

## Migration from R2.3 and earlier

Refer to Release Notes for EsgynDB release R2.5.

## Notes

### R2.6

- Existing **get** commands have been enhanced to
  - return information from the HIVE repository in addition to the EsgynDB repository with the necessary privileges. Privileges required:
    - DB\_\_ROOT or user granted DB\_\_ROOTROLE or granted SHOW component privilege
    - cqd TRAF\_GET\_HIVE\_OBJECTS is ON (the default)
  - change get schemas for user/role commands to return all schemas where the requested user has been granted at least one privilege. Prior to this release, it returned schemas owned by the requested user/role.
  - return names as delimited, only if they are delimited for:
    - get schemas for user <user>
    - get schemas for role <role>
    - get tables for user <user>
    - get tables for role <role>
    - get views for user <user>
    - get views for role <role>

The following command

```
get objects for user <user>
```

returns all securable objects that the specified user or one of the user's roles has been granted at least one privilege. Securable objects include:

- BT - base tables
- LB - libraries
- PA - packages
- SG - sequence generators
- UR - routines (procedures, functions, and table\_mapping functions)
- VI - views

If DB\_\_ROOT (or user granted DB\_\_ROOTROLE) and the cqd TRAF\_GET\_HIVE\_OBJECTS is ON, then the list of objects in the Trafodion catalog plus the list of objects from HIVE is returned. If a regular user, then only the list of objects from the Trafodion catalog are returned. This list includes any registered Hive objects.

### Pre-R2.6

- Generate statistics concurrent with Bulk Load operation  
Statistics can now be generated during the Bulk Load operation eliminating the need for a separate UPDATE STATISTICS step at the end.  
The syntax is

```
LOAD WITH UPDATE STATISTICS INTO table-name SELECT * FROM source-table-name
```

This will bulkload data into *table-name*, and concurrently update the histograms for the table.

## 2. Recovery after a hardware failure

The default TCP timeout to detect a broken connection in Linux is 2+ hours. Client applications connected to EsgynDB over TCP/IP may appear to hang after a failure (node failure, network failure, etc). In order to reduce the failure detection time for such broken connections, use the following keepalive settings across the cluster –

```
pdsh $MY_NODES 'sysctl -w net.ipv4.tcp_keepalive_time=240'  
pdsh $MY_NODES 'sysctl net.ipv4.tcp_keepalive_intvl=15'  
pdsh $MY_NODES 'sysctl net.ipv4.tcp_keepalive_probes=4'
```

## 3. HBase lease timeout patch

HBase uses a lease mechanism to protect against memory leaks in Region Servers caused by potential client instabilities that would open scanners, but die before having the opportunity to close cleanly and release resources. This mechanism relies on a server side timer, configured by the **hbase.client.scanner.timeout.period** parameter in *hbase-site.xml*. If a client fails to call next() within the timeout period, the server will assume the client died, and will force close the server side scanner and release resources. However, in EsgynDB, there are legitimate use cases where client is busy doing heavy processing, and needs more time than specified in the default scanner timeout value. Increasing the **hbase.client.scanner.timeout.period** value has the side effect of weakening the safety mechanism previously described.

The HBase community agrees that the correct behavior of this safety feature should be to have the client “reset” the scanner and resume where it left off instead of giving up and throwing an exception. The change will be implemented in a future release of HBase. In the meantime, this EsgynDB release includes a mechanism to invoke the correct behavior via a custom setting. You can enable the behavior by adding this parameter in *hbase-site.xml*

```
<property> <name>hbase.trafodion.patchclientscanner.enabled</name><value>true</value> <description>Enable an EsgynDB feature to allow a client to reset the HBase scanner and resume where it left off instead of throwing an exception upon expiry of the HBase hbase.client.scanner.timeout.period timer</description> </property>
```

The default value of the parameter is **false**.

## Fixed Issues

### R2.6.2

Mantis	Component	Description
M-11319	SQL	Compilation times are long for normal users when the number of views in the database is very large (90,000+)
M-12284	Security	Internal error returned while trying to retrieve schema privileges
M-12611	Installer	Cloudera Manager non-fatal error trying to set permissions on a non-existent file
M-12699 M-12705	Monitor	EsgynDB instance failure
M-12700	Tools	HBase shell errors and EsgynDB script failures when /tmp has NOEXEC attribute set
-	DTM	Miscellany fixes (a) Abort active transaction if the beginner dies on receiving a monitor message for connection close (b) Configurable abort-timer like behavior in HBase Region Server via HBase property <code>hbase.transaction.lease.timeout</code> (c) Reduce commit conflicts
-	Tools	Additional logging during the phase when an Esgyn Node role is started.

### R2.6.1

Mantis	Component	Description
M-10623	DTM	A NullPointerException error causes the HBase RegionServer to abort
M-10640	DTM	putMultiple interface continues trying to put into a region despite receiving a fatal exception
M-10572	Security	DB Manager cannot login successfully with db__root user
M-10687	DB Manager	Normal user who owns MANAGE_WMS privilege cannot get nodes list in Add/Alter Profile sub-window
-	DB Manager	Fix process stack error handling
M-10756	Security	User granted DB__SERVICESROLE cannot connect
M-10866	SQL	orc table which contains decimal (18,12) returns incorrect result
M-10749	DB Manager	On RH7 systems, UNREGISTER TENANT operation sometimes reports a cgroup delete warning and the memory cgroup is not removed for tenant
M-10731	DB Manager	Enable submit button for RESTORE operation
M-10653	DB Manager	"Client IP Address" configuration is not saved in Workloads Configuration page
M-10740	Manageability	Optimize tcollector scripts, drop collection of redundant Hadoop datanode/namenode metrics
M-9929	Connectivity	The "upsert using load into ...values(...),(...)..." causes the MXOSRVR to crash if the statement is too long

<b>M-10359</b>	SQL	The <> not equal condition caused wrong query plan and bad performance
-	SQL	"createUpgradeSnapshots" script for capturing EsgynDB metadata before upgrade
<b>M-10876</b>	DB Manager	DB Manager should not read traf_authentication_config if authorization is not enabled

## R2.6.0

<b>Mantis</b>	<b>Component</b>	<b>Description</b>
<b>M-10006</b>	BR	Data incorrect after restore failure
<b>M-10113</b>	Tools	Make USE_NEW_EDB_PDSH the default.
<b>M-10144</b>	Foundation	Fix to default the sqwatchdog timer to 60 seconds
<b>M-10145</b>	SQL	Avoid mxosrvr cores due to Hive DDL
<b>M-10172</b>	Installer	DBmgr config overwritten across nodes
<b>M-10179</b>	SQL	Select query returns error 4099 on a is null predicate
<b>M-10189</b>	SQL	Turn off salting on UPSHIFT columns
<b>M-10190</b>	Tools	rc_client abends when running some edb_pdsh commands
<b>M-10213</b>	SQL	HIVE: is null predicates return compiler internal error on joined tables with no join predicate
<b>M-10223</b>	SQL	HIVE: select count(*) returns 8001 internal error on a sample predicate with composite data types
<b>M-10321</b>	SQL	Hive text format table does not return last row, when this row does not have a column or row delimiter after last column
<b>M-5603</b>	Installer	Adding nodes while the first node down, and adding nodes on the first node
<b>M-7144</b>	DB Manager	"Grant/Revoke Privileges" window is broken at times
<b>M-7151</b>	DB Manager	Grant/Revoke SQL Privileges window is broken on all objects except for Schema.
<b>M-7580</b>	SQL	Create index cores tdm_arkcmp on a salted table with numeric column
<b>M-7781</b>	SQL	Large number of tdm_arkcmp,tdm_arkesp,mxosrvr cores in NAHeap
<b>M-7785</b>	SQL	When the column type is VARCHAR2(200000) character set utf8, the insert fails when the inserted data length exceeds 3393
<b>M-8286</b>	SQL	SELECT list with multiple CASE WHEN expression in USING clause of MERGER INTO with WHEN MATCHED clause cause ERROR[2006]
<b>M-8364</b>	SQL	Insert the maximum +1 supported by the nchar type, cause ERROR[2006] Internal error
<b>M-8498</b>	SQL	Sum (1) over(partition by pkey) in Olap function, cannot
<b>M-8538</b>	SQL	Lower-case objects cannot be dumped
<b>M-8539</b>	SQL	hbase table cannot be dumped through "_MAP_" / "_CELL_" / "_ROW_" entry
<b>M-8746</b>	SQL	Chinese Characters Error in Case When Expression columns
<b>M-8839</b>	SQL	Keyword case maxing as table or column name may migrate

<b>M-8877</b>	SQL	Aggr pushdown on orc/parquet tables with numerous files
<b>M-8880</b>	SQL	Support added columns to Parquet tables, first try
<b>M-8894</b>	SQL	Load returns a 8448 error and fails to populate the index
<b>M-8895</b>	SQL	Update stats ... like ... returns 9253 error when copying
<b>M-8898</b>	SQL	Fixes for Mantis-8898, Mantis-8988, Mantis-8925, Mantis-9406
<b>M-8909</b>	SQL	[first N], [any N], limit N in a predicate subquery
<b>M-8926</b>	SQL	Create table as with 'load if exists' wipes out the existing data when error occurs
<b>M-8940</b>	SQL	support expression in order by list and the expression is
<b>M-8959</b>	SQL	support the oracle syntax for retriving next/current value of sequence
<b>M-8988</b>	SQL	trafodion.sql.java.log rotation creates many small files in the same time period, rather than sharing logs
<b>M-9122</b>	SQL	Internal error on query to get list of procedures.
<b>M-9123</b>	SQL	support to_date('17-07-2018 09:44:57.196000', 'dd-mm-yyyy
<b>M-9169</b>	DB Manager	Add url auto login function
<b>M-9182</b>	SQL	insert chinese string into column whose character set is
<b>M-9199</b>	SQL	syntax error when update set clause has table alias
<b>M-9200</b>	SQL	The first operand of function TO_CHAR must be a datetime
<b>M-9201</b>	SQL	[Oracle Compatibility] syntax error when order by non-used column in a group by query
<b>M-9203</b>	SQL	[Oracle Compatibility] support ORDER options in sequence creation
<b>M-9208</b>	SQL	EsgynDB NUMERIC data type cannot completely match equivalent Oracle data type
<b>M-9231</b>	SQL	A difference of two DATE data types returns an INTERVAL data type, should convert to INTEGER internally
<b>M-9291</b>	SQL	long parameter for instr
<b>M-9294</b>	SQL	update the value itself with unique index fail
<b>M-9368</b>	JDBC T4	set param equals to null for a not null column would give an error.
<b>M-9385</b>	SQL	Create index fails with 8448 bulkload error
<b>M-9406</b>	SQL	lobtostring API reports an error when LOB column is NULL or the LOB is empty
<b>M-9476</b>	SQL	Trunc() function input (select count(*) from t)/(select
<b>M-9513</b>	SQL	The select statement contains an empty string, causing more compilation time
<b>M-9518</b>	Installer	Simplify cm_settings to avoid over-writing configs
<b>M-9519</b>	BR	Support accessing tables while BR operation is in progress
<b>M-9525</b>	SQL	Disable SQL_NOT_ATOMIC insert query using undo tree.
<b>M-9533</b>	SQL	Poor plan for groupby with CASE expression
<b>M-9582</b>	SQL	CONNECT BY condition using a compound condition will cause server to abort
<b>M-9589</b>	SQL	CONNECT BY does not work when table has primary key
<b>M-9613</b>	SQL	SQL with rownum takes 40 ms to compile every time
<b>M-9655</b>	Tools	DCS component cannot be stopped by ckillall cmd

<b>M-9664</b>	Tools	edb_pdsh miss last line when doing edb_pdsh -a cat /etc/hosts
<b>M-9750</b>	BR	Failed to cleanup backup after restore failure due to OOM
<b>M-9830</b>	Installer	Delete of some master nodes from EslynDB configured with DcsHA MultiMaster fails
<b>M-9841</b>	SQL	CQD parallel_num_esps does not honor tenant configuration
<b>M-9857</b>	Foundation	Fixed propagation of TRAF_LOG environment variable to child
<b>M-9860</b>	Connectivity	JDBC fetch will cause MXOSRVR crash when fetch size is very large
<b>M-9866</b>	Foundation	Local node fails because of znode expired when remote node triggers "sync thread timeout"
<b>M-9886</b>	Installer	Parameter logfile-name is not used when -l option is specified for the command line installer
<b>M-9887</b>	Foundation	Monitor Sync Thread Timer should log CRITICAL status event messages when counting down to a timeout
<b>M-9899</b>	BR	User who has select & insert privilege on schema can't restore the schema
<b>M-9903</b>	DB Manager	Disable weak SSL ciphers for DB Manager
<b>M-9919</b>	SQL	Fix UPDATE STATS hang on Avro tables
<b>M-9944</b>	SQL	Query cache causes OLAP query to return wrong result
<b>M-9946</b>	SQL	Incremental update stats returns "No record delimiter found" 8446 error on partitioned hive text table with large varchar cols
<b>M-9953</b>	Foundation	Bypassed Allgather() reconnect/ping logic when in a virtual environment
<b>M-9972</b>	SQL	core in datetimeType::getRepresentableValue.

The following Trafodion JIRAs have been addressed.

<b>JIRA</b>	<b>Component</b>	<b>Description</b>
<b>T-2626</b>	Configuration	Make logs directory location configurable
<b>T-3218</b>	Security	User still has privilege after user's role has been revoked or privilege has been revoked from the user's role
<b>T-3220</b>	SQL	Extend the existing compiler GUI to display query execution as well
<b>T-3223</b>	SQL	Row count estimation code works poorly on time-ordered aged-out data
<b>T-3234</b>	SQL	Refactor hive meta calls to be less resource intensive
<b>T-3237</b>	SQL	Fix incorrect PCode optimization
<b>T-3246</b>	JDBCT4	Add slf4j in jdbc
<b>T-3250</b>	Connectivity	optimize get/set schema
<b>T-3260</b>	SQL	SSMP may wait 3 seconds before handling requests
<b>T-3266</b>	SQL	Change ustats base table accesses to SKIP CONFLICT
<b>T-3270</b>	SQL	Handle incompatible types in key expressions

## Known Issues

ID	Component	Description
M-7552	BR	Concurrent backup/restore operations or concurrent DDL operation with backup/restore operation return errors.
M-7324	BR	Export/Import of backups in Hortonworks HDP environment hangs
M-7162	Foundation	When hostname is either all upper-case or mixed case, the EsgynDB instance fails to start
M-6909	SQL	Compiler crashes for certain ORC insert select statement during codegen
M-6796	SQL	Cached Hive NATable cause MXOSRVRs to fail with a SIGSEGV

## Platform Support

<b>Operating Systems</b>	CentOS 6.5 – 6.9, 7.2 – 7.4 Red Hat 6.5 – 6.9, 7.2 – 7.4
<b>Hadoop Distributions</b>	Cloudera distributions CDH 5.9 – 5.16 Hortonworks distributions HDP 2.4 – 2.6.3
<b>Java</b>	Minimum Recommended JDK 1.8

## Documentation

<b>EsgynDB</b>	EsgynDB Installation Guide R2.6.0 EsgynDB SQL Reference Manual R2.6.0 EsgynDB CQD Reference Manual R2.6.0
<b>Clients</b>	Trafodion Client Installation Guide R2.3.0 EsgynDB ODBC Driver Installation Guide R2.6.0
<b>EsgynDB Manager</b>	EsgynDB Manager User Guide R2.6.0

© 2019 Esgyn Corporation. All rights reserved.