

Esgyn Corporation

Release Notes for EsgynDB R2.4



Published: April 2018
Edition: EsgynDB Release R2.4

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 better performance and stability.

This EsgynDB release includes a number of new features, support for the Parquet storage format, and numerous improvements in usability and enhanced stability. Fixes and enhancements shipped up to EsgynDB release R2.3.4 are included.

Features

Category	Features
Infrastructure	<ul style="list-style-type: none">• Support for HDP 2.5• Support for Linux pids greater than 64K• Integration with Cloudera Manager for installation Advanced <ul style="list-style-type: none">• Support for the Apache Parquet columnar storage format
High Availability	<ul style="list-style-type: none">• Support for Linux keepalived for DCS HA
Usability	<ul style="list-style-type: none">• Notification in EsgynDB Manager for license expiry• Windows ODBC driver now supports storing user-id and password (encrypted) in the ODBC DSN
Manageability	Advanced <ul style="list-style-type: none">• Support incremental backup and restore of Trafodion tables• EsgynDB Manager support for Hive object browsing
Database	<ul style="list-style-type: none">• Support for Region transactions
Security	Advanced <ul style="list-style-type: none">• Support for multi-domain Active Directory / LDAP

Migration Considerations

Users upgrading from EsgynDB R2.3 release may be affected by the following changes

System

1. Backup / Restore

The backup and restore utility now supports incremental backup capability. An arbitrary set of schemas and tables can be identified for incremental backup, so that only changes between two backup operations will be captured. This is an online operation.

Backups taken in prior releases (EsgynDB R2.2 and EsgynDB R2.3) are incompatible with EsgynDB R2.4. After an upgrade to EsgynDB R2.4, the administrator should take a full database backup.

Application

1. Keywords UUID and VALUE are no longer treated as reserved
UUID and VALUE are now treated as non-reserved ANSI words.
2. Region Transactions
Transactions that involve single row / single table (ie single region) are now handled by the DTM subsystem. Conflicts in such a case will result in SQL error 8448.

Users upgrading from pre-R2.2 or R2.2 EsgynDB releases should first upgrade to EsgynDB R2.3, then to EsgynDB R2.4. As part of the upgrade to R2.3, they may need to execute additional steps after the installation of the release and once EsgynDB is started.

System

1. Hbase namespace
EsgynDB metadata in R2.3 is now hosted in separate namespaces instead of the default HBase namespace in pre-R2.3 releases.

Migration of existing metadata from the default HBase namespace to EsgynDB reserved namespaces is done automatically by the installer when migrating from pre-R2.3 to R2.3 releases. Reserved namespaces are TRAF_RSRVD_1 through TRAF_RSRVD_7.

User data will continue to be hosted in HBase default namespace, unless it is migrated explicitly. Refer to the EsgynDB Product Guide for additional details.

When multi-tenancy is enabled and a tenant is created, a “tenant” namespace is automatically created to hold data. Generally there is one (maybe two) schema associated with a tenant. Any schema or objects created when tenants are enabled are created in the tenant namespace. There is also a default namespace if a tenant is not specified at connection time.

2. TRAF_HOME
A new environment variable named TRAF_HOME has been introduced to replace the variable SQ_HOME. This is a name change only. Just like SQ_HOME, TRAF_HOME represents the root of the EsgynDB installation tree.
3. Log files
EsgynDB log files for individual product modules have been simplified greatly. Each product module now has a single log file. Log directories remain unchanged from their locations in pre-R2.3 releases.

Connectivity subsystem	\$TRAF_HOME/dcs-2.3.0/logs
REST subsystem	\$TRAF_HOME/rest-2.3.0/logs
EsgynDB Manager	\$TRAF_HOME/dbmgr-2.3.0/logs
Other subsystems	\$TRAF_HOME/logs

4. Elasticity
Support for elastically adding nodes has been simplified in this release. New nodes can be dynamically added into the EsgynDB instance on demand. Prior to R2.3, virtual nodes had to be preconfigured in order to avoid an EsgynDB instance restart when adding new nodes. This is no longer needed.

5. Backup / Restore

Enhancements in the backup and restore utility now enable online backups and online restores of an arbitrary set of schemas and/or tables. This is in addition to support for an online full database backup and offline full database restore.

Prior to R2.3, separate backup (`run_full_trafodion_backup.sh`) and restore (`run_full_trafodion_restore.sh`) utilities were used for online full database backup and offline full database restore. Both utilities are no longer supported starting with EsgynDB R2.3 release.

Note: Support for online backup and online restore of schemas or tables is only available in the Advanced Edition of EsgynDB.

6. Installer

A new python-based installer has been introduced in EsgynDB R2.3, subsuming the old bash installer which is now obsolete.

7. Licensing

The EsgynDB product license has been extended to support new features. Existing production licenses for R2.2 will continue to work with R2.3 through to their expiry date. However, it is recommended to update to the new license key as part of the upgrade to R2.3.

Application

1. SQL Error 8606 is obsolete

Starting in EsgynDB R2.2, SQL applications should interpret error 8616 as a retryable error. This error replaces error 8606.

Old error:

```
8606: Transaction subsystem <name> returned error <number> on a commit transaction.
```

New error:

```
8616: A conflict was detected during commit processing. Transaction has been aborted.
```

Notes

1. 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'
```

2. 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

ID	Component	Description
M-7245	DB Manager	In Cloudera parcels installed cluster, DB Manager transactions and canary response graphs show no data
M-7009	Installer	Python installer will over-write HBase Advanced settings in a CDH environment at times
M-6998	Tools	sqcheck incorrectly shows TSD and EsgynDB Manager processes as not running
M-6969	SQL	An INSERT operation containing a DECODE function causes the MXOSRVR to abort
M-6966	Foundation	EsgynDB node down when the network is stopped for 30 sec under high load
M-6908	SQL	Java error during ORC scan with predicate or aggregate pushdown
M-6741	SQL	Comparison of string and numeric datatypes returns an error if the string operand contains a non-numeric value
M-6544	Security	Error “No valid credentials provided” when running queries
M-6380	SQL	Parquet filter stack trace when operand is null does not convey the column name
M-6315	Foundation	Monitor fails with a segmentation violation error due to use of a node ID value that is not configured
M-6274	SQL	Invalid node map for a query when CQD PARALLEL_NUM_ESPS exceeds number of nodes
M-6030	Security	Group members cannot get roles with command 'get roles for group'

M-5975	Security	Tenant administrator does not have privileges on the tenant default schema if the schema is set by 'ALTER TENANT'
M-5662	BR	Cannot drop EsgynDB backup tables if Trafodion is not initialized
M-3712	SQL	SQL compiler crashes in NATable::updateExtTableAttrs

Issues fixed since Apache Trafodion R2.1 are included.

Known Issues

ID	Component	Description
M-7214	SQL	ESP core during UDR stress testing
M-7162	Foundation	When hostname is either all upper-case or mixed case, the EsgynDB instance fails to start
M-7160	Foundation	When hostnames in /etc/sysconfig/network and /etc/hosts do not match, 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
M-6766	Foundation	Unable to start ESPs on a node after a reboot, causing queries to fail with OS Error 4022, TPC Error 53, error detail 2

Platform Support

Operating Systems	CentOS 6.5 – 6.9, 7.1 – 7.2 Red Hat 6.5 – 6.7, 7.1 – 7.2
Hadoop Distributions	Cloudera distributions CDH 5.4, 5.5, 5.6, 5.7, 5.8, 5.9 Hortonworks distributions HDP 2.4, 2.5 Apache Hadoop with Apache HBase 1.0
Java	Minimum Recommended JDK 1.8

Documentation

EsgynDB	EsgynDB Installation Guide R2.4.0 EsgynDB Backup Restore Guide R2.4.0 EsgynDB SQL Reference Manual R2.4.0 EsgynDB CQD Reference Manual R2.4.0.1
Clients	Trafodion Client Installation Guide EsgynDB ODBC Driver Installation Guide R2.4.0
EsgynDB Manager	EsgynDB Manager User Guide R2.4.0

© 2018 Esgyn Corporation. All rights reserved.