



Tungsten Release Notes

Tungsten Release Notes

Tungsten version 1.3

This document was published on 23 March 2010.

www.continuent.com [<http://www.continuent.com>]

Copyright © 2008-2010 Continuent

The trademarks, logos, and service marks in this Document are the property of Continuent or other third parties. You are not permitted to use these Marks without the prior written consent of Continuent or such appropriate third party. Continuent, Tungsten, uni/cluster, m/cluster, p/cluster, uc/connector, and the Continuent logo are trademarks or registered trademarks of Continuent in the United States, France, Finland and other countries.

All Materials on this Document are (and shall continue to be) owned exclusively by Continuent or other respective third party owners and are protected under applicable copyrights, patents, trademarks, trade dress and/or other proprietary rights. Under no circumstances will you acquire any ownership rights or other interest in any Materials by or through your access or use of the Materials. All right, title and interest not expressly granted is reserved to Continuent.

All rights reserved.

Table of Contents

1. Introduction to Tungsten Enterprise Release Notes	1
2. What's New with Tungsten?	2
3. Platform Support	5
3.1. Operating System Support	5
3.2. Database Platform Support	5
4. Release Summary	7
4.1. Tungsten Enterprise Release 1.3	7
4.1.1. Upgrade from Previous Releases	7
4.1.2. Notable Improvements in This Release	8
4.2. Tungsten Enterprise Release 1.2.3	8
4.2.1. Upgrade from Previous Releases	8
4.2.2. Notable Improvements in This Release	8
4.3. Tungsten Enterprise Release 1.2.2	9
4.3.1. Upgrade from Previous Releases	9
4.3.2. Notable Improvements in This Release	10
4.4. Tungsten Enterprise Release 1.2.1	10
4.4.1. Upgrade from Previous Releases	10
4.4.2. Notable Improvements in This Release	11
4.5. Tungsten Enterprise Release 1.2	11
4.5.1. Upgrade from Previous Releases	11
4.5.2. Notable Improvements in This Release	12
4.6. Tungsten Enterprise Release 1.1	13
4.6.1. Upgrade from Previous Releases	14
4.6.2. Notable Improvements in This Release	14
4.7. Tungsten Enterprise Release 1.0.9	16
4.7.1. Upgrade from Previous Releases	16
4.7.2. Notable Improvements in This Release	16
4.8. Tungsten Enterprise Release 1.0.7	17
4.8.1. Upgrade from Previous Releases	17
4.8.2. Notable Improvements in This Release	17
5. Known Issues in Tungsten Replicator Current Release	18
5.1. General Limitations with the Current Release	18
5.2. Specific Bugs and Missing Features	18
5.3. Installation of Patched MySQL JDBC Driver	19
A. Network DataBase Support	21
A.1. Know issues	21

Chapter 1. Introduction to Tungsten Enterprise Release Notes

This document contains release notes for Tungsten Enterprise. The release notes include a list of changes in each recent release, instructions for upgrade, and a list of known problems that users should be aware of before deploying Tungsten.

For additional questions about release contents, please contact our support group through the Continuent Support Portal (<http://www.continuent.com/support>). Meanwhile, we hope you enjoy the product!



Tip

Be sure to check Chapter 5, *Known Issues in Tungsten Replicator Current Release* for a summary of current limitations before planning a new deployment.

Chapter 2. What's New with Tungsten?

Tungsten Enterprise Version 1.3 is a major release that includes a number of features to improve the flexibility, speed, and reliability of Tungsten clusters. Version 1.3 also includes a large number of bug fixes and minor improvements.

- Significant architectural improvements to allow Tungsten Replicator to define custom replication flows (pipelines) for native Tungsten Replication as well as to manage non-Tungsten mechanisms such as PostgreSQL warm standby replication. Due to this change there are changes to the replicator.properties file format, as well as changes to JMX interfaces for the replicator. In addition, there is a change to the SQL tables used to track replication progress that requires users to clear Tungsten catalogs when upgrading.
- Extensions to MySQL replication capabilities including correct replication of session variables, processing of LOAD DATA INFILE statements, a fragmentation protocol for large transactions, ability to read binlog events from MySQL master servers using the MySQL client protocol, and finally commands to start replication at a user-selected location in the binlog. These greatly improve replicator reliability and reduce to a minimum data loss in the event of master failures.
- Improved support for clusters based on PostgreSQL warm standby. This release also has alpha support for clustering using PostgreSQL 9 log streaming and hot standby. These features eliminate standby lag and allow applications to read from standby databases.



Note

Tungsten 1.3 is in controlled release at this time. We recommend this version for PostgreSQL deployments. Please consult with Continuent before using Tungsten 1.3 in production MySQL deployments.

Tungsten Enterprise Version 1.2.3 release is a maintenance release that adds the following new features plus a number of key bug fixes.

- Transparent session consistency for multi-tenant applications. This feature allows SaaS applications with large numbers of tenants sharded by database or using separate database user IDs to distribute reads automatically to slave databases without application changes.
- Greatly improved purge history script for Enterprise users. This script works for all 1.2.x versions and should be applied immediately to all production deployments.
- Updated documentation for all parts of Tungsten. This includes greatly expanded descriptions of Tungsten cluster concepts, management procedures, and connectivity methods.
- Improvements to handling of session load balancing in the Tungsten Connector and SQL Router. These increase the reliability of SQL routing considerably.
- Fixes to MySQL binlog extraction to handle enum and set data types correctly when using MySQL row replication.

Tungsten Enterprise Version 1.2.2 release is a maintenance release that adds the following new features plus a number of important bug fixes.

- Managers now share the same group communications channel with monitoring. This greatly improves stability of the management framework.

What's New with Tungsten?

- The offers a `passThroughMode` property in `connector.properties` to direct the to proxy network buffers directly for queries rather than translating to JDBC. Pass-through mode can speed up query throughput by 200% or more.
- The offers extensions to session load balancing including a new setting that helps SaaS applications among others perform efficient read/write splitting between master and slave databases without application code changes.
- Tungsten includes optional rules to implement floating IP addresses for master data sources. These rules can be found in `cluster-manager/rules-ext`. A README explains how to install them.

Tungsten Enterprise Version 1.2.1 release is a maintenance release that adds two features as well as a number of stability improvements.

- Handling of binary and character set introducers in MySQL statement replication. You may now replicate embedded binary data when using statement replication on MySQL. Unicode character sets are correctly handled. Alternative character sets embedded using introducers are also handled.
- Data source recovery using a single command. You can now recover failed and shunned data sources by entering a single command: **`datasource dsname recover`**.

Tungsten Enterprise Version 1.2 release includes major enhancements to ease of use as well as stability of database clusters. These are of course in addition to smaller feature additions and bug fixes.

- Detection and correct handling of all major failure cases by automatic failover rules. This includes cases involving host failure, host reboot, and loss of networking connectivity.
- Addition of manager commands to support all major operations for replication, including backup/restore, heart-beat events, flush operations, and incremental consistency checks. The new commands virtually eliminate the need for users to connect to specific hosts to perform replicator operations.
- Ability to push long-running commands into the background and manage them using 'jobs' command similar to Linux shells. Also, long-running commands no longer block the cluster.
- Improvements to backups including the ability to run hot backups on slaves. Previously it was necessary to take the slave offline.
- Numerous improvements to management to reduce CPU overhead and use memory more efficiently during long-term operation.
- Support for managing clusters based on PostgreSQL warm standby. Support includes provisioning of clusters, simple failover procedures, and backup/restore.
- Preview support for installation and management of clusters based on Oracle 10g running on Linux.
- Improved MySQL replication support including ability to replicate temporary tables, correct handling of partial binlog writes, and better replication of numeric datatypes.

The Tungsten Enterprise 1.1 release provided a number of important enhancements in addition to many smaller feature additions and fixes.

- Enhanced rules for failover including support better fencing of failed resource and protections against split-brain scenarios
- Selectable modes of operation including automatic failover, manual failover, and maintenance mode

What's New with Tungsten?

- Across-the-board improvement of resource management to support long-term cluster operation with millions of transactions
- Built-in protections against accidental cluster and data source misconfiguration
- Support for cluster installation and operation using RightScale cloud administration tools
- Upgrades and expansions to system documentation

Tungsten Enterprise 1.0.x releases introduced a number of important features, including the following.

- Automated failover using configurable production rules
- Built-in support for backup and restore
- Features for cloud operation including efficient batch installations and non-multicast network support
- Support for Solaris platform
- Support for alternate cluster names
- Ability to control operating system services from `ccctl` management console
- Greatly improved documentation and up-front configuration
- Simple installation of all cluster components including service script deployment for Linux hosts
- More efficient administration through flexible and powerful cluster management client
- Reduced administrator effort through automated tracking of data source online/offline state as well as latency
- Improved transparency for applications through session-consistency model of load balancing

Chapter 3. Platform Support

This chapter contains describes platform support for the current version of Tungsten clustering.

3.1. Operating System Support

The following table shows current status of Tungsten platform support.

Table 3.1. Operating System Platform Certifications

Operation Systems	Status	Notes
Linux (RedHat)	Primary platform	RHEL 4 and 5 as well as CentOS 5.x versions are fully supported
Linux (Ubuntu)	Primary platform	Ubuntu 9.x/10.x versions are fully supported
Linux (Debian/Suse/ Other)	Secondary Platform	Other Linux platforms are supported but are not regularly tested. We will fix any bugs reported by customers.
Solaris	Primary platform.	Solaris 10 is fully supported. OpenSolaris is not supported at this time.
Mac OS/X	Secondary platform.	Mac OS/X Leopard and Snow Leopard are used for development at Continuent but not regularly certified. We will fix any bugs reported by customers.
Windows	Limited Support	Tungsten 1.3 and above will support Windows platforms for connectivity (Tungsten Connector and SQL Router) but may require manual configuration. Tungsten clusters do not run on Windows.
BSD	Limited Support	Tungsten 1.3 and above will support BSD for connectivity (Tungsten Connector and SQL Router) but may require manual configuration. Tungsten clusters do not run on BSD.

3.2. Database Platform Support

The following table shows current status of Tungsten DBMS support.

Table 3.2. DBMS Platform Certifications

DBMS	Status	Notes
MySQL	Primary platform	Statement replication is supported for MySQL 5.0 and 5.1. Row replication is supported for MySQL 5.1. MyISAM and InnoDB table types are fully supported. Provisional support for NDB is available as of Tungsten 1.2.3. We recommend customers use the most recent available versions of both 5.0 and 5.1.

Platform Support

DBMS	Status	Notes
PostgreSQL	Primary platform	Warm standby clustering is supported for PostgreSQL 8.2-8.4.
Oracle 9i/10g/11g	Secondary Platform	Replication from MySQL into Oracle is supported but not regularly tested. Special configuration is required for replication to function correctly.
Drizzle	Secondary Platform	Experimental support for Drizzle is available in in Tungsten 1.3. Drizzle replication is not tested.

Chapter 4. Release Summary

This chapter contains a summary of the changes and upgrade notes for each recent release of Tungsten.

4.1. Tungsten Enterprise Release 1.3

Tungsten Release 1.3 includes architectural changes to the Tungsten Replicator, substantial upgrades to MySQL binlog parsing, connectivity improvements, and numerous bug fixes.

Warning

Tungsten 1.3 is in controlled release at this time. It is recommended for PostgreSQL deployments. Please consult with Continuent before using Tungsten 1.3 in production MySQL deployments.

4.1.1. Upgrade from Previous Releases

Release 1.3 includes a change to the `trep_commit_seqno` table that is incompatible with previous releases of Tungsten. In addition, the management protocol for Tungsten SQL Router and Connector has changed. Upgrade to Tungsten will therefore require a brief application outage to complete production deployments.

We recommend for testing that you perform a fresh installation of Tungsten on a test system. You can also upgrade existing Tungsten 1.2.x test environments to use Tungsten 1.3 with the following procedure.

1. Shut down all user applications and services using the Tungsten 1.2.x cluster. Stop all Tungsten 1.2.x services. Do not shut down databases.
2. In MySQL installations, drop the Tungsten catalog database on all MySQL servers as shown in the following example:

```
mysql> drop database tungsten;
```

Note

As long as master and slave databases are consistent it is not necessary to re-provision slaves.

3. Install Release 1.3 software on each host that is in the cluster. Run **configure** and start Tungsten 1.3 cluster services.
4. Upgrade SQL Router libraries in any application that uses them to Tungsten 1.3.
5. Restart applications and user services. Check the cluster status using the **ctrl ls** command. Ensure that all nodes are ONLINE.

Tip

As with any new software we strongly recommend that you test the cluster carefully. Please report bugs to Continuent through the support portal at <http://www.continuent.com/support>.

4.1.2. Notable Improvements in This Release



Note

The list of improvements for Tungsten 1.3 is incomplete. The following list contains a selection of items that are addressed in current code.

- *TENT-49* - Need 'adduser' script that can be run by a mysql 'superuser' and that will set up database users with the correct permissions/grants
- *TUC-1* - Mysql extractor may skip event while displaying "Unexpected failure while extracting event" error message in log files
- *TUC-5* - Open Replicator script plugin does not collect and print stdout for scripts
- *TUC-6* - Add ability to specify replication start point
- *TUC-7* - Develop Tungsten Open Replicator framework with plugin support
- *TUC-10* - Implement capability to break up large transactions into fragments
- *TUC-12* - Support LOAD DATA INFILE command in MySQL statement replication
- *TUC-19* - Implement replication pipelines
- *TUC-23* - Provide a means to support application-transparent session-consistent reads
- *TUC-24* - Unable to recover replicator to a 'slave' after the node it is on is isolated from the cluster and there is no other cluster activity
- *TUC-44* - Implement session variables replication support for MySQL

4.2. Tungsten Enterprise Release 1.2.3

Tungsten Release 1.2.3 upgrades MySQL binlog support, improves management, and fixes several problems encountered by users in large deployments.

4.2.1. Upgrade from Previous Releases

You can upgrade from any Tungsten 1.2 installation to Tungsten 1.2.3 release by replacing software. There are no replicator protocol changes. See Section 4.3.1, "Upgrade from Previous Releases" for a description of the procedure.

4.2.2. Notable Improvements in This Release

- *TENT-49* - Need 'adduser' script that can be run by a mysql 'superuser' and that will set up database users with the correct permissions/grants
- *TENT-98* - Manager/monitor outage triggered by setting manager to maintenance mode
- *TENT-101* - Monitoring or policy engine stops executing under some failover circumstances, deadlock found in one case
- *TENT-105* - Database grants/permissions should be checked (or granted) during installation time.

- *TENT-106* - Confirm that MySQL replication is NOT running during the installation process.
- *TENT-121* - Configure asks for a port number but presets node names as defaults
- *TENT-127* - Documentation requires updates and polishing
- *TENT-138* - Fix PostgreSQL warm standby configuration to work easily for advanced configuration
- *TENT-139* - Open replicator goes online for PostgreSQL even when there are script failures
- *TENT-141* - Document character set best practices for replicator on MySQL
- *TENT-143* - Cctl does not report accurate latency numbers for slaves
- *TENT-146* - Purge_history cron job causes replicator failures when history table size is large
- *TUC-13* - When a rotate event is extracted, the extractor might try to open the new binlog file before it really exists on disk
- *TUC-26* - Add windows version of wrapper binaries and libs
- *TUC-28* - Provide a means to support application-transparent session-consistent reads
- *TUC-32* - Replicator does not correctly handle skipping multiple bad events
- *TUC-33* - Add provisional support for MySQL NDB tables
- *TUC-40* - Replicator may try to apply events which were previously skipped
- *TUC-41* - MySQL enum datatype not extracted correctly.

4.3. Tungsten Enterprise Release 1.2.2

Tungsten Release 1.2.2 merges monitoring and management group communications, introduces pass-through mode for SQL proxying, adds new session load balancing features, and adds management of master floating IP addresses. It also fixes a number of outstanding bugs for both MySQL as well as PostgreSQL replication.

4.3.1. Upgrade from Previous Releases

You can upgrade from any 1.2 Tungsten release by replacing software. There are no replicator protocol changes.

1. Unpack Tungsten 1.2.2 software on each host. Copy the existing `tungsten.cfg` to the new release.
2. Stop Tungsten services from previous releases
3. Run **configure** in batch mode, that is, **configure -b** on the Tungsten 1.2.2 release.

Tip

The `tungsten.cfg` file is normally compatible between all 1.2.x versions of Tungsten. If you do encounter a problem, just rerun the **configure** manually, which should fill in any missing values.

To upgrade from versions prior to Tungsten 1.2 see release notes for those versions.

4.3.2. Notable Improvements in This Release

- *MYO-111* - Implement pass-through mode proxying
- *TENT-102* - Introduce TCPING option for group communications
- *TENT-107* - Should confirm that binlog is enabled in the my.cnf file at installation time
- *TENT-112* - Manager does not reconnect with SQL Router if router restarted more than a certain period of time after last exit
- *TENT-122* - Add support for floating IP address management
- *TENT-129* - Mysql installation fails to recognize the binlog pattern if my.cnf entry contains spaces or tabs
- *TENT-130* - Add cron script to purge Tungsten history table on a regular basis
- *TMAN-98* - Manager locks up at start-up time with missing transition message
- *TMAN-100* - Manager 'times out' while executing an internal distributed action - sync, directory connect, merge etc.
- *TMAN-102* - Rework monitor to share the same group as the manager
- *TPSQL-43* - Detect whether pg_standby is installed during configuration
- *TPSQL-44* - WAL file progress does not increment when slave is lost
- *TPSQL-45* - Slave is out of sync if master DB is restarted during slave outage
- *TREP-358* - MySQL extraction may fail because of a timestamp value
- *TREP-395* - Mysql extractor may skip event while displaying "Unexpected failure while extracting event" error message in log files
- *TREP-397* - When using statement replication, applier fails if the database name contains special characters (e.g. a minus symbol)
- *TSR-26* - Session consistency is broken

4.4. Tungsten Enterprise Release 1.2.1

Tungsten Release 1.2.1 adds binary and character set processing capabilities to the replicator, adds a recover command, and fixes a number of bugs encountered in operation as well as in-house testing.

4.4.1. Upgrade from Previous Releases

Release 1.2.1 can be upgraded by replacing software. There are no replicator protocol changes. Here are the steps for upgrade.



Note

It is not necessary to replace SQL Router libraries embedded in applications.

1. Unpack Tungsten 1.2.1 software on each host. Copy the existing Tungsten 1.2 `tungsten.cfg` to the new release.
2. Stop Tungsten 1.2 services.
3. Run **configure** in batch mode, that is, **configure -b** on the Tungsten 1.2.1 release.

To upgrade from versions prior to Tungsten 1.2 see release notes for those versions.

4.4.2. Notable Improvements in This Release

- *TENT-112* - Manager does not reconnect with SQL Router if router restarted more than a certain period of time after last exit
- *TENT-113* - Add ability to handle binary data in statement replication
- *TENT-114* - Fix problems with `mysqsession` filter used to support replication of temp tables on MySQL
- *TENT-115* - Setting a datasource to FAILED or SHUNNED does not wait for/cancel current connections for that datasource
- *TENT-116* - Script `deployall` incorrectly installs boot script for monitor
- *TENT-117* - The 'thl info' command fails on MySQL with a syntax error
- *TENT-118* - Turn off manager data source lookup on local SQL Router
- *TENT-120* - Network partitions or other node-leave/join issues can corrupt cluster datasource state. Also includes implementation of data source recover.

4.5. Tungsten Enterprise Release 1.2

Tungsten Release 1.2 adds new commands to the manager, improves memory management and CPU usage, adds new failover cases, fixes a number of problems with MySQL replication, and adds provisional support for PostgreSQL. There are of course numerous bug fixes as well.

4.5.1. Upgrade from Previous Releases

Release 1.2 includes an incompatible change to the replication protocol that was required to fix a severe hang problem encountered in earlier versions of MySQL. The replicator protocol change necessitates a maintenance window for operational systems. Here are the steps for upgrade that minimize down-time.



Note

It is not necessary to replace SQL Router libraries embedded in applications.

1. Install Release 1.2 software on each host that is scheduled for upgrade. Run **configure** but *do not* start services.
2. Quiesce applications that write to the master database. Applications that read data only do not need to be quiesced.
3. Stop Tungsten 1.1 services.
4. Drop and recreate 'tungsten' database on each DBMS server as shown below.

```
DROP DATABASE tungsten;  
CREATE DATABASE tungsten;
```

5. Start Tungsten 1.2 services.
6. Start applications.

Once the upgrade is finished you should immediately back up a slave database to ensure that you have a valid backup to restore the cluster in the event of a failure.

To upgrade from releases prior to 1.1 you should reprovision replicas as there have been changes in the replication storage protocol that may prevent proper operation if you upgrade directly.

4.5.2. Notable Improvements in This Release

- *TENT-77* - Need to be able to run commands in the background and, also, to avoid having a single command hang the managers
- *TENT-78* - Amazon: shutting down a master machine does not introduce a failover
- *TENT-83* - Correct hang condition that is sometimes encountered in switch command
- *TENT-85* - Database port modification in configure does not update all required settings
- *TENT-90* - Automatic start-up of connectors causes problems for data source management
- *TENT-95* - Configurator needs to ask about the replicator binlog filename pattern. Many customers have the hostname instead of mysql-bin
- *[TENT-12* - Update Oracle dslisten builds for 32- and 64-bit Linux
- *TENT-71* - Add cctrl commands to perform replicator flush, check, and heartbeat commands
- *TENT-94* - At installation time, confirm java is in \$PATH and is of an appropriate version. Also confirm JAVA_HOME is correctly set.
- *[TENT-29* - Handle host and network failures where nodes and processes simply stop responding
- *TENT-39* - Add support for managing PostgreSQL replication with warm standby
- *TENT-42* - Eliminate split brain conditions due to manager failing and not informing SQL Router of data source changes
- *TENT-59* - Cctrl command should provide convenient backup and restore commands
- *[TMAN-22* - ClusterManagementHandler does not support JMX methods with non-String arguments
- *TMAN-28* - "exit" does not work in cctrl
- *TMAN-31* - Cctrl manual failover command needs cleanup
- *TMAN-32* - Add diag command to manager to print diagnostic output

- *TMAN-34* - Race condition in `ctrl` switch command may lead to master failing to log some transactions immediately after promoting a master
- *TMAN-35* - Add `ctrl` command to output last lines of component's log
- *TMAN-38* - Multiple hostnames for the same IP in `/etc/hosts` may introduce errors
- *TMAN-48* - A hung operation on one of the managers blocks the entire cluster
- *TMAN-64* - Memory exhaustion, high-CPU usage and hanging issues reported by customers
- *TMAN-75* - Failover rules do not take reliably take failed data sources offline
- *TMAN-76* - Manually failing a datasource causes it to disappear
- *TMAN-77* - Leaving `ctrl` idle for a long period results in a reconnect failure loop
- *TMAN-78* - Memory leaks and deadlocks encountered during intensive concurrent load
- *TMAN-84* - Remove help that is automatically printed when entering a bad command in `ctrl`
- *[TREP-298]* - Calls to get max/min sequence numbers slow replicator and generate floods of messages in log
- *TREP-337* - Slave to master transition may hang
- *TREP-339* - Tungsten replicator does not replicate temporary tables correctly in MySQL statement replication
- *TREP-345* - MySQLExtractor does not handle `DECIMAL(precision, scale)` type properly
- *TREP-347* - `database/OracleDatabase.createTable()` does not handle non-unique indexes. These are critical for the `tungsten.history` table.
- *TREP-351* - Implement PostgreSQL backuper
- *TREP-352* - Default `mysqldump` backup configuration fails
- *TREP-353* - Intermittent issue with master side hang while waiting for heartbeat
- *TREP-358* - MySQL extraction may fail because of a timestamp value
- *TREP-361* - MySQLExtractor does not properly handle `MySQLExtractException` which is generated by `Log_event.read_log_event` method
- *TREP-363* - `Log_event.read_log_event` sometimes fails with "EOFException while reading 0 bytes from binlog"
- *TREP-366* - Backup fails with error when trying to delete expired backups dumped by previous replicator version
- *TREP-367* - Implement slave hot backup support
- *TREP-376* - `thl purge -y -high $seqno` should never completely empty the history table.

4.6. Tungsten Enterprise Release 1.1

Tungsten Release 1.1 provides a host of improvements including improved rules for automated failover, better memory management for long-term operation, fixes for a number of common configuration problems such as

handling hosts that have alternative names in `/etc/hosts`, improved help and documentation, and numerous bug fixes. Finally, there are a number of improvements to support operation with RightScale management tools on Amazon.

4.6.1. Upgrade from Previous Releases

To upgrade from version 1.0.9, perform the following steps:

1. Prior to upgrade you should add an index on the `tungsten.history` table using the following command. This index is created automatically on new tables as part of the fix for TREP-316.

ALTER TABLE history ADD KEY eventid (eventid);



Note

The ALTER TABLE command can take a while to complete if the history table is very large.

2. Follow the instructions for upgrade in *Tungsten Installation Guide*.

To upgrade from releases prior to 1.0.9 you must reprovision replicas as there have been changes in the replication storage protocol that may prevent proper operation if you upgrade directly.

4.6.2. Notable Improvements in This Release

- *TENT-5* - Create simple utility to dump diagnostic information for support purposes
- *TENT-9* - Implement RightScale installation scripts for Tungsten Enterprise
- *TENT-20* - Cluster upgrade procedure
- *TENT-28* - A 'maintenance' mode is needed in `cctrl`
- *TENT-30* - Manager sometimes 'hangs' in `cctrl` when changing the state of a datasource - looks like it's notification related
- *TENT-34* - Managers intermittently fail to recognize new data sources
- *TENT-35* - Clean-up wrapper binaries
- *TENT-36* - Typo in replicator applier `url_options`
- *TENT-37* - Add recommended options to default mysql datasource URLs
- *TENT-38* - Do not remove datasources when failures are detected
- *TENT-43* - Non-default MySQL port is not supported nor by configurator nor by Replicator
- *TENT-44* - Multiple hostnames for the same IP in `/etc/hosts` may introduce errors
- *TENT-45* - Policy manager should not re-enable a shunned datasource if this creates more than one master
- *TENT-46* - Ability to change master/slave role from `cctrl`
- *TENT-47* - New datasource was being created even though DataServer was shown as stopped, and the removed because DataServer shows as stopped

Release Summary

- *TENT-48* - Install issue: need to create the Tungsten database if it does not already exist.
- *TENT-50* - Failover attempts to use slave on disconnected node
- *TENT-51* - Documentation improvements for Tungsten-Installation-Guide-2009 (version 1.1)
- *TENT-52* - Host property is not getting set, during configuration, in checker.mysqlserver.properties
- *TENT-53* - Session management uses an offline replicator
- *TENT-54* - cctrl allows users to do a number of potentially very damaging operations in logical mode - need to warn appropriately
- *TENT-56* - Documentation improvements for Tungsten-Enterprise-SQL-Router-Guide
- *TENT-57* - Connector should reflect default datasource option zeroDateTimeBehavior=convertToNull
- *TENT-61* - Failure of a replicator does not cause corresponding fencing/failure of datasource
- *TENT-62* - Configure script fails with an unhelpful error if required Ruby libraries are not installed
- *TENT-64* - Root installation does not correct set file ownerships; causes strange data source behavior
- *TENT-65* - Serialized ReplicatorNotification is not backward compatible
- *TENT-68* - Create functional tests to validate basic manager functionality, including switch and failover.
- *TENT-69* - Configuration script does not change monitor channel name if replicator is not running
- *TENT-70* - Documentation improvements for Tungsten-Enterprise-Replicator-Guide-2009
- *TENT-72* - Alfresco requires mysql connector java version 5.1.6 to run properly
- *TENT-73* - Documentation improvements for Tungsten-Enterprise-Connector-Guide-2009
- *TMON-10* - Every MySQLServerChecker.check() call consumes at least 10% of CPU
- *TMON-13* - MySQL server checker leaks memory
- *TREP-290* - MySQL extractor can miss events inserted into the database immediately after Replicator reaches ONLINE:MASTER state
- *TREP-291* - Logs fill disk up when encountering binlog read error
- *TREP-316* - Replicator hangs on restart when history table is very large
- *TREP-321* - Tungsten replicator reacts poorly to a request for a non-existent sequence number from a slave
- *TREP-323* - Replicator mysqlbackup agent generates corrupt output files
- *TREP-324* - Master replicator misses events to the MySQL binlog if master goes offline immediately after master control event
- *TREP-325* - Master dumps cannot be reloaded into a slave after master has written transactions
- *TREP-327* - Removing events on master side using a filter creates sequence number gaps that break replication

- *TREP-329* - Mysqldumps are occasionally corrupted
- *TREP-330* - There is no way to change default MySQL port
- *TREP-332* - Flush implementation may return a wrong sequence number
- *TREP-336* - MySQLExtractor may fail with an IOException while extracting an event from the binlog, showing a "binlog read error" message in the logs

4.7. Tungsten Enterprise Release 1.0.9

Tungsten Release 1.0.9 introduces automated backup/restore as well as configuration support for manual vs. automated configuration policies. There are documentation improvements and underlying management improvements including the ability to manage operating system services such as MySQL database servers directly from the management framework. A number of open source bug fixes are also included.

4.7.1. Upgrade from Previous Releases

We recommend that users upgrade all earlier versions to this release. To upgrade from version 1.0.7, perform the following steps:

1. Install version 1.0.9 but do not start services. Do not reuse the `tungsten.cfg` file from release 1.0.7 as it is not compatible.
2. Stop all 1.0.7 services on all hosts in the cluster. This includes SQL Router and Connector processes. Use the `cluster-home/bin/stopall` script.
3. Start all 1.0.9 services on all hosts in the cluster. This includes SQL Router and Connector processes. Use the `cluster-home/bin/startall` script.

To upgrade from previous releases, reprovision replicas as there have been changes in the replication protocol that may prevent proper operation if you upgrade directly.

4.7.2. Notable Improvements in This Release

- *TENT-4* - Add support for configuring JGroups gossip router for group communications
- *TENT-8* - Provide cluster support for database backup and restore
- *TENT-21* - Disabling Connector read/write splitting inserts a r/w configuration line
- *TENT-22* - Add debian startup script support in configuration
- *TENT-25* - Unable to use other port than 12001 with gossiprouter
- *TENT-26* - Configure script fails on hosts that have ruby 1.9.1 installed
- *TENT-27* - Upgrade Installation and Configuration Guide to cover planning, extra checks, and additional configuration information
- *TENT-31* - Installation and Configuration manual has a number of small inaccuracies
- *TENT-32* - Prefix monitoring channel by cluster name

4.8. Tungsten Enterprise Release 1.0.7

Tungsten Enterprise release 1.0.7 introduces automated failover using configurable production rules. It also includes bug fixes to underlying open source components as well as documentation and configuration improvements.

4.8.1. Upgrade from Previous Releases

Not applicable; upgrade to version 1.0.9.

4.8.2. Notable Improvements in This Release

- *TENT-16* - Implement automatic failure detection and failover invocation
- *TENT-18* - Support other than "default" for cluster name

Chapter 5. Known Issues in Tungsten Replicator Current Release

Tungsten is undergoing steady development and each new release contains numerous bug fixes and improvements. This chapter provides guidance on significant known issues in the current release.

5.1. General Limitations with the Current Release

Users of Tungsten should be aware of the following general limitations. These may prevent Tungsten from being used on particular types of applications or in particular operating environments. We are actively working on improvements in all of these areas.

- *Autonomic Management* - Network partitions in which both partitions elect masters may require operator intervention to reconcile replicator transaction histories. All other known cases of failover are handled correctly.
- *MySQL Replication* - Tungsten does not replicator session variables such as SQL_MODE settings and does not support **LOAD DATA INFILE** commands. In addition, UNSIGNED numeric types are not correctly replicated under row replication when the high order bit is set.



Note

The Tungsten 1.3 release adds support for session variables and **LOAD DATA INFILE**.

- *MySQL JDBC Driver Limitations* - Due to bugs in the MySQL JDBC driver handling of JDBC escape sequences, Tungsten may not correctly replicate SQL statements that contains HTML or usual text patterns. We recommend use of a patched MySQL driver until these bugs can be fixed. See Section 5.3, "Installation of Patched MySQL JDBC Driver" for instructions.
- *MySQL Versions* - Tungsten cannot replicate from MySQL 4.1 and earlier versions. Replication into these versions is supported though there may be limitations due to syntax changes in MySQL itself.
- *Windows Support* - Tungsten is not certified for Microsoft Windows and does not support cluster operation at this time. Individual components like replication do, however, work on Windows but are unsupported.
- *SQL Routing* - Only a single SQL Router library may operate on a single host due to port conflicts. This includes Connectors with embedded SQL Router libraries.
- *Online Upgrade* - Tungsten does not support a full online upgrade. It is not necessary to stop databases or reprovision replicas, but users must accept a short outage to replace application libraries and Tungsten components. Full online upgrade will be supported in a later release.
- *PostgreSQL support* - PostgreSQL replication and clustering supports warm standby clusters only.

5.2. Specific Bugs and Missing Features

The current Tungsten release has the following specific bugs or missing features. All of these are scheduled for resolution in future releases.

- *TENT-1* - Need to set up Tungsten boot scripts on Solaris and Mac OS X. Users can solve this problem for now by setting up the scripts by hand.

Known Issues in Tungsten Replicator Current Release

- *TENT-41* - Need to permit multiple Tungsten components of the same time to run on the same host. For example only one or SQL Router process may run on echo host.
- *TENT-42* - Split brain conditions may occur in the SQL Router if a local manager fails, as changes to data sources will not be passed to the SQL Router. Workaround is to ensure managers are restarted promptly.
- *TENT-51* - The `cctrl` utility must provide a convenient method to recover failed data sources either from another data source or a backup. The current `recover` command reconfigures the data source but does not handle data re-provisioning.
- *TMAN-40* - Manager should support JMX server with configurable registry for all cluster components. This is required to support multiple SQL Routers on a single host.
- *TREP-70* - The replicator does not replicate session variables that affect behavior of MySQL statements on the the slave. This results in slave failures when attempting to load `mysqldump` files and example schemas like the MySQL Sakila sample code on the master. This problem can be avoided by using 5.1 row replication.
- *TREP-197* - MySQL 5.1 row replication does not correctly populate column names in tables. This is a problem when filtering row data or replicating across database types because only the column numbers are recorded in the protocol. Check with Continuent Support for workarounds.
- *TREP-281* - Tungsten replicator should provide options for parallel replication to slaves. This is necessary to get acceptable performance for many clustering and replication scenarios.
- *TUC-3* - Unsigned numeric values with the high order bit set are not correctly replicated when using MySQL row replication. Values that do not have high order bits set replicate correctly.

5.3. Installation of Patched MySQL JDBC Driver

There is a serious problem with Tungsten replication due to bugs in the MySQL JDBC driver. The MySQL JDBC driver incorrectly treats certain strings of characters in SQL statements as escape sequences, which leads to immediate failures when slave replicators try to apply these statements to the database. This problem has only been seen when using statement replication so far. We have logged this problem with MySQL as bug 51313 (<http://bugs.mysql.com/bug.php?id=51313>) and provided suggestions for fixes, so we hope a future version of the MySQL JDBC driver will not have this problem.

In the meantime, we have a patched version of the MySQL 5.1.6 Connector/J driver that we are making available to interested users. We recommend that all MySQL sites update their Tungsten releases to use this driver. Here is the procedure for putting in the new driver. You must do this on each node where Tungsten is installed. It takes about 5 minutes to complete and does not require an application outage.

1. Download the driver from the following Amazon S3 link into a local directory like `/tmp`. You can use the `wget` command on Linux systems to do this from the command line.

<https://releases.continuent.com.s3.amazonaws.com/mysql-connector-java-5.1.6WithEscapeProcessingFix-bin.jar>

This driver is released under the GPL. You can download the source code from the following link:

<https://releases.continuent.com.s3.amazonaws.com/mysql-connector-java-5.1.6WithEscapeProcessingFix-src.tar.gz>

2. Copy the driver into the `tungsten-replicator/lib` directory and rename the existing driver as shown below:

Known Issues in Tungsten Replicator Current Release

```
cd tungsten-replicator/lib
mv mysql-connector-java-5.1.6-bin.jar mysql-connector-java-5.1.6-bin.jar.old
cp /tmp/mysql-connector-java-5.1.6WithEscapeProcessingFix-bin.jar .
```

3. Restart the replicator from the cctrl utility using the following commands.

```
(start cctrl)
set policy maintenance
service node_name/replicator restart
replicator node_name online
ls (Ensure replicator shows up as online before next command)
set policy {automatic | manual} (Use whichever is your preferred mode)
```

That is all that is required. Please follow all steps very carefully to avoid accidentally triggering a failover. You may wish to apply the procedure to a slave first and do the master later.



Note

The JDBC fix is required only for Tungsten. Java applications that run against MySQL are not affected and do not need an upgraded driver.

Appendix A. Network DataBase Support

Tungsten has basic support for replicating tables that use the Network DataBase (NDB) storage engine. NDB is the multi master clustering solution from MySQL. MySQL defines it as follows:

MySQL Cluster integrates the standard MySQL server with an in-memory clustered storage engine called NDB (which stands for "Network DataBase"). In our documentation, the term NDB refers to the part of the setup that is specific to the storage engine, whereas "MySQL Cluster" refers to the combination of one or more MySQL servers with the NDB storage engine.

For more information, see <http://dev.mysql.com/doc/refman/5.1/en/mysql-cluster-overview.html>

NDB support has been tested against `mysql-cluster-gpl-7.0.5-linux-x86_64-glibc23`.



Note

The use of a filter is still required to drop insert statement, such as `table ndb_apply_status`. The filter is not delivered in the release.



Note

Tungsten tables can now be stored into a different storage engine. Change the following property into `replicator.properties`:

```
replicator.table.engine=
```

A.1. Know issues

- *Mixing `ndbcluster` and other storage engines*

Mixing `ndbcluster` and other storage engines can lead to out of order logs. In this case, a write statement against an `innodb` table can be logged before another statement on a `ndbcluster` table, even though it was executed after it. This seems to happen when write statements against `ndbcluster` tables are done in `autocommit`.

- *Logging latency*

It seems that NDB can slightly delay logging. As a consequence, `flush` or `waitForAppliedSeqno` operations can occur against an old event id (**show master status** returns an outdated event id).

- *Logging several transactions as part of the same event*

NDB can log several transactions as part of the same event. This has a side effect on consistency checks, which are expected to be in their own event, but this is not the case. Consistency might not work for now, at least under load.