



Continuent™ Tungsten Enterprise Solutions

Robust Solutions for
Continuous Data Availability
and Data Protection

Key Features White Paper

February 16, 2010

Overview

Data is the heart of the modern enterprise. Continuous Data Availability means having your data when you need it, where you need it, 24/7/365.

Continuent offers a wide range of data availability and performance scaling solutions for problems ranging from basic database availability to high performance database clustering.

Our solutions are designed to be both economical as well as effective -- combining the economy of open source with the advanced capabilities of commercial database offerings.

Continuent offers a wealth of solutions for MySQL[®], and PostgreSQL users proving the following key features on five (5) main categories:

Data Management

Tungsten allows users to manage a group of databases linked through replication as a single, integral unit, using simple commands and without any single point of failure in the system

Data Management functionality includes command line management, automatic datasource detection, manager command backgrounding, remote management API, service configuration/removal tool, and simple procedures for updates and maintenance,

Advanced Data Management features include business rule-driven management with automated failover and management for SaaS multi-tenant systems.

Monitoring features include remote monitoring API, event notifications, remote problem diagnosis.

High Availability

Tungsten provides simple, automated failover procedures to allow applications to switch to replicas quickly, seamlessly, and without risk of data loss.

High availability issues can be divided between **continuous operation during emergency failures** (unplanned outage) and **continuous operation during maintenance** operations (planned outage).

Tungsten features that cover the **unplanned database outage** include automatic master failover, automated failure detection and recovery, master recovery and failback, fail-safe data application on slave, and automatic routing to new master.

For the **planned database outages** Tungsten offers seamless, transparent switchover for maintenance and zero downtime upgrades.

High Performance

Faster hardware and increasing data volumes raise the bar for application performance. Tungsten offers capabilities that are well-suited to take advantage of multi-core architectures and cheap memory.

Tungsten High Performance functionality addresses **replication** issues and **load balancing** operations.

Replication functionality includes advanced 1+N database replication, fragmentation protocol for large transactions, fast unlogged replication, relay replication, parallel replication of shards, bi-directional replication with conflict resolution, direct TCP/IP proxying.

Load balancing features include intelligent load balancing, session consistency model load balancing, read-write splitting, and slave latency tracking.

Data Protection

Tungsten integrates advanced replication and backups to ensure complete protection from the DBA's worst nightmare: losing data.

Tungsten data protection functionality addresses **data loss and data consistency** issues, **backup**, and **disaster recovery** deployments.

Tungsten **data loss and data consistency** features include automated data consistency checks, zero data loss during planned failover, and recovery from admin data deletion

For **backup** solutions Tungsten offers integrated backup/restore, and scheduled backups (with cron jobs).

Tungsten **disaster recovery** functionality includes creating replicate databases on remote site, replicating data to remote site, and fast offload of DBMS log to alternate host.

Data Integration, Filtering and Archiving

Data integration functionality includes extending to new data sources and configurable replication pipelines.

Data filtering and archiving features include extensible data filters, data filter libraries, per-database and per-table filtering and shard identification filters.

Data Management

Tungsten allows users to manage a group of databases linked through replication as a single, integral unit, using simple commands and without any single point of failure in the system.

Command line management. Failover is a complex operation in master/slave clusters that may involve applications, proxies, replication tools, and of course databases. Tungsten reduces failover to a single command that takes databases safely on- and offline (without losing data) and reconfigures the rest of the system accordingly

Automatic data source detection. Tungsten group communication eliminates administration tasks by auto-configuring system components to adjust load balancing as servers become available, disappear, or exceed latency requirements of client applications.

Manager command backgrounding. Command job control allows users to move long-running commands like backup or restore operations into the background, where they can be managed using familiar Unix-like job control commands. All user commands run as separate threads so that they do not block other operations of the cluster.

Service configuration/removal tool. The Tungsten service configuration tool allows end-users of ISP and SaaS organizations to enable, disable, and manage Tungsten services like backup, point-in-time-recovery, and reporting database construction remotely without requiring help from the hosting organization.

Event notifications. Event notifications inform users of significant occurrences in the cluster, such as detecting failed hosts or marking successful recovery of a database server. Notifications can plug into a variety of messaging mechanisms.

High Availability

Tungsten provides simple, automated failover procedures to allow applications to switch to replicas quickly, seamlessly, and without risk of data loss. This helps address not just failures but also maintenance, which is the main source of downtime in modern IT environments.

Continuous Operation during Emergency Failure

Automated failure detection and recovery. Tungsten failover is a one-step procedure that automatically switches masters and reconfigures the rest of the cluster, including directing other slaves to the new master and updating load-balancing.

Master recovery and fallback. Tungsten can promote any slave to be master without data loss or requiring other slaves to be stopped or reprovisioned. Masters can be recovered with a single command. Likewise, a fallback to the recovered master is a single command.

Fail-safe data application on slave. Slaves remain consistent even following crashes thanks to transactional data replication, which ensures that SQL updates are neither dropped nor repeated.

Automatic Routing of SQL. Tungsten provides a variety of options to route SQL from application to an appropriate database server, including Java routing libraries, language-independent proxies, and floating IP addresses. Tungsten automatically shifts new connections as masters move and as administrators add and remove databases in the cluster.

Continuous Operation during Maintenance

Seamless, transparent switchover for maintenance. Tungsten allows users to switch master databases without data loss or halting applications, thereby facilitating daylight maintenance operations.

Zero-downtime upgrades. Run upgrade and maintenance processes on slaves, then switch masters after the upgrade has finished successfully. This eliminates the risk of data loss due to upgrade mistakes.

High Performance

Faster hardware and increasing data volumes raise the bar for application performance. Tungsten offers capabilities that are well-suited to take advantage of multi-core architectures and cheap memory.

Replication

Advanced 1+N database replication. Tungsten allows users to set up new databases easily and remove them when no longer needed. We also provide integrated management for the entire cluster.

Fragmentation protocol for large transactions. Tungsten allows users to replicate transactions that extend into millions of rows. The Tungsten Replicator breaks such transactions up into parts that can propagate independently without requiring excessive memory.

Parallel replication of shards Slave lag is a problem for any system that has high update (write) rates that grow worse as hardware advances raise master throughput. Tungsten includes parallel replication that allows slaves to apply events using multiple simultaneous channels. Multi-channel replication allows clusters to scale farther without requiring intrusive application changes. The Tungsten approach is particularly well suited for multi-tenant applications, an area of rapid growth for both open source as well as commercial databases.

Data partitioning/sharding. Partitioning data into separate databases or “shards” is a standard approach to scale database performance and keep data volumes at reasonable levels. Tungsten includes an advanced parallel replication that can shard data automatically out of a single master into multiple slaves and also take care of routing SQL transactions to the correct database for each shard.

Bi-directional replication with conflict resolution. Tungsten will support replication in both directions between databases with adjustable policies for handling conflicts.

Load balancing

Intelligent load balancing algorithms. Tungsten SQL Router allows Java applications to load balance and route SQL requests efficiently with no performance overhead across multiple data replicas in master/slave clusters. The load balancing is based on an innovative algorithm that allows users to trade availability versus data consistency requirements to increase performance.

Session consistency model load balancing. Session consistency load-balancing distributes user read transactions to slave databases that are up to date with the user's last write. It is a powerful technique to offer the benefits of multi-master clustering using simpler and more economical master/slave clusters.

Read-Write splitting. Read-write splitting enables applications to dispatch read traffic within transactions to slave servers. Operation is completely transparent to the applications that use it.

Slave latency tracking. Tungsten allows users to specify acceptable latency levels on slaves and then monitors them automatically. Connections are only load-balanced to up-to-date slaves.

Direct TCP/IP proxying. Tungsten includes lightweight proxies that route database requests between client and database without intermediate translations. Proxies make very fast routing decisions, then basically allow the client and server to converse with minimal additional overhead.

Data Protection

Tungsten integrates advanced replication and backups to ensure complete protection from the DBA's worst nightmare: losing data. Tungsten makes it easy to set up live copies that protect from failures as well as administrative data deletion. Built-in consistency checks ensure data integrity so users can be confident that replicas are ready for use and correct.

Data Loss and Data Consistency

Automated consistency checks. Perform incremental checks of consistency between master and slave databases, without stopping the cluster and without causing a performance hit on applications.

Zero-data loss for planned and unplanned failovers. Tungsten has multiple options to eliminate loss of data when failing over from an old master to a promoted slave. These include flush events that clear transactions from the master.

Recovery from admin data deletion. Tungsten can delay replication to offer a "window of opportunity" to recover data lost through administrative mistakes on the master, such as deleting application tables.

Backup

Integrated backup/restore. Tungsten integrates backups using file system snapshots and also through database dumps. Tungsten replication is tightly integrated with backup and restore. Users can dump or restore databases with a single command. In addition, Tungsten supports slave “auto-provisioning” in which newly configured replicators automatically load the most recent backup on start-up, a feature that is ideally suited for rapid slave provisioning. Slaves can even roll forward to a specific point in the log. Tungsten can integrate with a wide variety of backup software, so users can continue to use any backup with which they are familiar.

Scheduled Backups. Scheduled backups run at user-defined intervals to ensure that there are always archive copies available of user data. Tungsten uses innovative technology to ensure backups run once and only once each time they are scheduled regardless of the current composition of the cluster.

Disaster Recovery

Provision databases on remote site. Tungsten can provision slaves on a remote site using backups. This allows remote, live copies at a fixed point in time.

Replicate data to remote site. Tungsten can similarly replicate to slaves on a remote site. This allows users to maintain fully-up-to-date live copies on the remote site.

Fast offload of DBMS log to alternate host. The database log is the Achilles for both HA and performance. This feature moves the current database log to to a slave machine, which means that database transactions get off the master host faster. It also has the potential to speed replication by a factor of three or more for many users.

Data Filtering

Tungsten offers powerful filtering capabilities that can change or drop replicated SQL both on the master as well as slaves

Extensible data filters. Tungsten includes built-in filters, and users can easily develop new ones using either Java or Javascript. Filtering is so useful that Tungsten itself uses filters to implement time-delay replication, handle flush operations, and assign events to channels for parallel replication.

Data filter libraries. Tungsten integrated filtering technology can change SQL “on the fly” from the source dialect to the target. This enables a much wider range of SQL to be replicated across database platforms such as Oracle to MySQL.

Data Integration

Databases in large enterprises do not operate in isolation. Users need to share data between different database types as well as between

databases and applications. Tungsten provides features that enable this to occur easily.

Extend to new data sources. Tungsten provides interfaces that allow users to add additional data sources. Tungsten will provide an integration package to describe how to implement new database types.

Configurable replication pipelines. {Robert to add explanation}

Replication to/from applications. Tungsten offers applier and extractor implementations that allow users to replicate from popular applications. Users can of course write and contribute their own.

Replication to/from message bus software. Likewise, Tungsten is easy to extend to message bus software like IBM MQ Series. This can help with moving data between a message bus and the database.

Replication to any JDBC-enabled database. Tungsten will provide generic applier logic that works on any database that supports JDBC.

Summary And More Information

Tungsten commercial solutions help users harness innovative open source technologies to provide continuous data availability that meets the needs of enterprises from small to large. Tungsten offers features that are critical to ensuring that databases operate efficiently and economically under a wide range of conditions.

Tungsten Enterprise solutions and services are available from Continuent, a Continuous Data Availability Company.

Our website is www.continuent.com, and you may download an evaluation version of Tungsten Enterprise at www.continuent.com/software.

Tungsten Open Source components are available at the Continuent Community web site www.continuent.com/community.

Feature Matrix

Product Matrix	Tungsten Community Edition	Tungsten Enterprise Edition
Data Management		
Easy master/slave cluster management	√	√
Command line management	√	√
Automatic datasource detection	√	√
Scriptable commands	√	√
Manager command backgrounding	√	√
Event notifications		√
Advanced Data Management		
Business rule-driven management with automated failover		√
Complex Topologies		√
Multi-site database configurations		√
Management for multi-tenant systems		√
High Availability		
Database failure detection	√	√
Fail-safe data application on slave	√	√
Master recovery	√	√
Simple procedures for database maintenance	√	√
Unattended Availability		√
Automated failover		√
Handling of network failures		√
Floating IP Support		√
PostgreSQL Warm Standby		√
PostgreSQL Hot Standby/Log Streaming		√
High Performance		
Advanced 1+N database replication	√	√
Intelligent load balancing	√	√
Session consistency model load balancing	√	√
Slave latency tracking	√	√
Direct TCP/IP proxying	√	√
Fragmentation protocol for large transactions	√	√
Fast unlogged replication		√
Transparent session consistency for SaaS apps		√
Data Protection		
Automated data consistency checks	√	√
Integrated backup/restore	√	√
Zero-data loss planned failover	√	√
Protect from admin data deletion		√
Automated cron jobs for backups		√
Data Filtering and Archiving		
Extensible data filters	√	√
Data filter libraries	√	√

Per-database and per-table filtering	√	√
Data Integration		
Extend to new data sources	√	√
Configurable replication pipelines	√	√
Installation and Configuration		
1-step interactive installation	√	√
1-step batch installation	√	√
Documentation and Support		
Professional documentation	√	√
Community support	√	√
Certified builds and patches		√
Commercial 24x7 enterprise support		√
Custom bug fixes		√
Guaranteed response time		√
Database Platforms		
MySQL 5.0/5.1	√	√
MariaDB 5.0/5.1	√	√
PostgreSQL 8.2-8.4		√
Operating System Platforms		
Linux 32/64 Bit	√	√
Mac OS X	√	√
Solaris 10 Intel/Sparc	√	√