Continuent Clustering

Benefits And Configurations

Table of Contents

What is Continuent Clustering?

Continuent Clustering™ allows enterprises running business-critical database applications to affordably achieve commercial-grade high availability (HA), globally redundant disaster recovery (DR) and performance scaling.

Continuent Clustering makes it simple to create database clusters in the cloud or in your private data center, to keep the data available even when systems fail, freeing you to focus on your business and applications.

In addition, Tungsten Replicator™ provides data replication from relational databases to high-performance NoSQL and data analytics engines to derive insight from big data for better business decisions.

The key benefits include:

- **Deployment Flexibility** – Build database clusters in any bare-metal, private cloud (VMware, Open Stack) or public cloud (Amazon AWS, Google Cloud Platform, Microsoft Azure) environments.
- **Zero Downtime Maintenance** – Perform database maintenance, such as schema changes, upgrades to a new version of MySQL or adding new hardware, without application or service interruptions.
- **High Availability** – Replace a failed master DBMS server within seconds while maintaining continuous availability and without losing any transactions.
- **Disaster Recovery** – Switch to a hot-standby disaster recovery site running a Continuent Clustering cluster with a single command without losing application connectivity.
- **Multisite Operations with Geo-Scaling** – Implement multi-master, multi-site database clusters that span sites and are ready for immediate failover.
- **Supports off-the-shelf MySQL, MariaDB and Percona Server** – No software upgrade, data migration, application or DDL changes are required.

Continuent Products

Continuent Clustering

Continuent Clustering allows enterprises running business-critical database applications to affordably achieve commercial-grade high availability (HA), globally redundant disaster recovery (DR) and performance scaling.

Continuent Clustering makes it simple to create database clusters in the cloud or in your private data center, to keep the data available even when systems fail, freeing you to focus on your business and applications.

Tungsten Replicator

Tungsten Replicator provides core replication between MySQL, MariaDB and Percona servers and replication to and from Oracle up to Oracle 12c. Additionally, Tungsten Replicator can extract from MySQL and Oracle sources and apply to various Hadoop distributions, Amazon Aurora, RDS, Redshift, Vertica by HPE, Elasticsearch, Kafka, Cassandra and MongoDB.

Feature	Continuent Clustering	Tungsten Replicator
Flexible MySQL clustering	Yes	
Zero downtime maintenance	Yes	
High availability	Yes	
Disaster recovery	Yes	
Multimaster, multisite	Yes	Yes
Improved performance	Yes	Yes
Real-time data loading to analytics and big data		Yes
Oracle replication		Yes
Supports MySQL, MariaDB and Percona Server	Yes	Yes

continuent

Key Benefits of Continuent Clustering

High availability (HA), disaster recovery (DR)

- Provide redundancy within and across data centers
- Create database clusters ready for immediate failover for maximum availability and data protection of business-critical applications
- Reduce database recovery time from hours or days to seconds

True multimaster distributed solution

- Mix-and-match on-premises, private and public clouds (incl. Amazon AWS, Google Cloud Platform, Microsoft Azure)
- Withstand node, data center, availability zone or region failures
- Bring data close to your application users for faster response times and better user experience

Zero downtime

- Site-level and cross-site failover ensures application availability
- Upgrade hardware, software and data without taking applications offline
- MySQL compatibility means easy migration of your data and applications

Scaling and performance

- Load balance read operations across multiple slaves
- Rapidly add more read slaves as needed for infinite scaling
- Increase replication performance and transaction volumes with parallel replication

No migration, no application changes

- Setup in minutes
- Run your off-the-shelf MySQL, MariaDB and Percona Server
- MySQL compatibility means easy migration of your data and applications

Flexibility and choice

- Deploy within cloud, virtual and bare metal environments
- Avoid vendor lock-in
- Select your environment(s) and providers at better price points

Cost savings

- Optimize costs by selecting your environment(s)
- Eliminate downtime risks, reduce DBA time spent on these, lowering your costs while increasing reliability
- Re-use your code and applications with little or no change

Customer service

- 24/7 customer support
- 1-hour SLA, with average response of less than 5 minutes
- Uptime measured in months or years

Advanced Features

Multimaster replication

- Enable global real-time transaction processing with flexible multimaster replication configurations between clusters
- Make real-time data aggregation of reporting data simple by replicating from multiple locations into a single database server without the need for data transformation

Parallel replication

- Increase replication performance over native MySQL and increases transaction volumes using parallel slave apply

Oracle replication

- Tungsten Replicator can replicate MySQL transactions to and from Oracle databases
- Tungsten Replicator can also replicate from Oracle to Oracle

Real-time analytics and big data

- Tungsten Replicator provides real-time analytics replicating data quickly and efficiently to high-performance analytics engines, such as Vertica
- Tungsten Replicator can publish data in real-time from SQL to NoSQL implementations, such as MongoDB

Continuent Clustering Components

Continuent Clustering includes three main components:

Tungsten Replicator

Tungsten Replicator™ is an open source replication engine supporting a variety of different extractor and applier modules. During replication, Tungsten Replicator assigns data a unique global transaction ID, and enables flexible statement and/or row-based replication of data. This enables data to be exchanged between different databases and different database versions. Information can be filtered and modified in-flight, and deployment can be between on-premise or cloud-based databases.

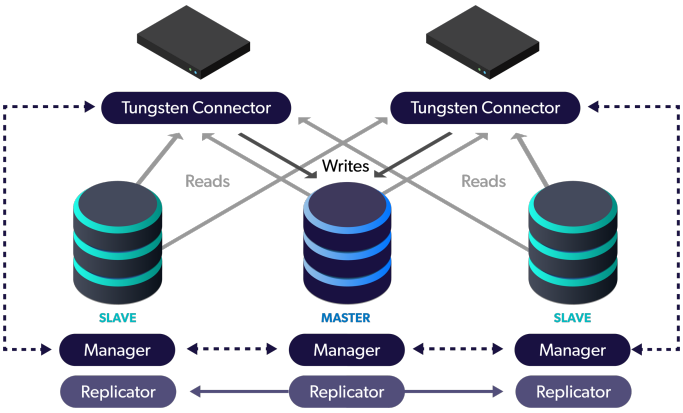

Tungsten Manager

The Tungsten Manager™ is responsible for monitoring and managing a Continuent Tungsten data service. The manager has a number of control and supervisory roles for the operation of the cluster, and acts both as a control and a central information source for the status and health of the data service as a whole.

Tungsten Connector

The Tungsten Connector™ is a service that sits between your application server and your MySQL database. The connector routes connections from your application servers to the data sources within the cluster, automatically distributing and redirecting queries to each data source according to load balancing and availability requirements.

Tungsten Dashboard

In addition to the command-line user interfaces, we provide our Tungsten Dashboard™ for easy monitoring and management of your clusters.

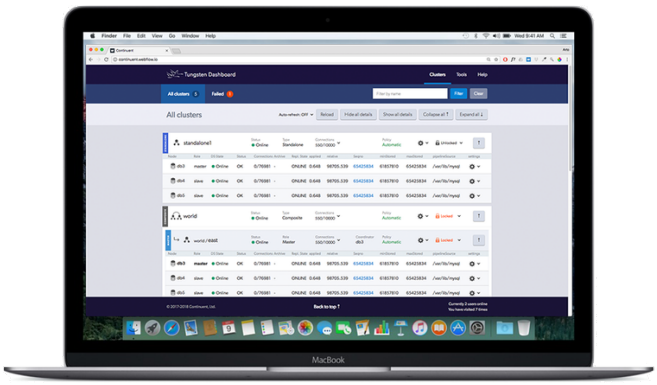

Multisite Solutions

Primary/DR Clustering

1 master + 2 failover/read slaves and 1 relay/read node + 2 failover/read slaves

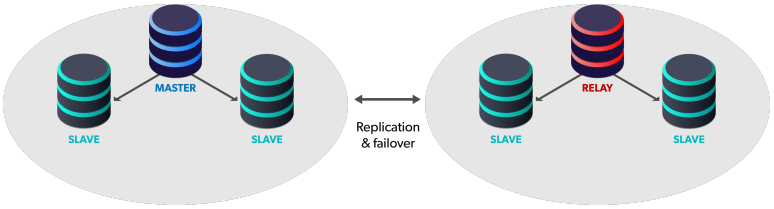

In the Primary/DR Composite Cluster topology, there is a single writeable master node across all sites, and all writes are directed to that master. The Connectors are multi-site aware and will automatically react to a site failover. Replication flows in one direction only, from the current Primary cluster to the DR cluster(s).

This topology provides a highly available local operation and an active failover site that is also available for reads. It is well-suited for customers who seek the highest level of HA, with the ability to operate applications on both sites, without application changes.

Multimaster Clustering

Two or more clusters each consisting of 1 master + 2 failover/read slaves

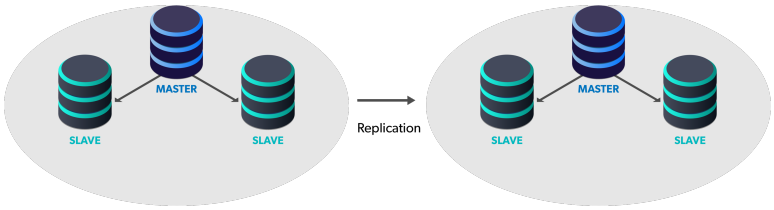

In the Multimaster Composite Cluster topology, there is one writeable master node per cluster, and all writes are directed to that master by the local Connectors. The Connectors are able to use any other site in the event of a local outage; all sites are write-able at all times, and replication goes in both directions.

This topology links highly available clusters across sites to enable constant availability for updates in two or more locations separated by high-latency networks. It is recommended for SaaS applications, credit card payment gateways, or online services that must always be available for business.

Single Site Solutions

Standalone HA Clusters

1 master + 2 failover/read slaves

The Continuent Clustering HA configuration is designed for sites that strive to ensure constant availability. This topology ensures the ability to perform maintenance on any database server and still maintain an environment capable of HA failover at all times.

This is the preferred basic configuration for all applications.

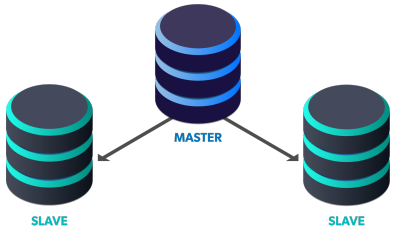

HA with Read Scaling

1 master + 4 or more failover/read slaves

Continuent Clustering HA with Read Scaling is ideal for sites that require HA and read-scaling for enhanced performance under peak loads. Continuent Clustering enables the use of additional read slaves with minimal performance impact to the HA cluster. Media and consumer web sites should consider this option.

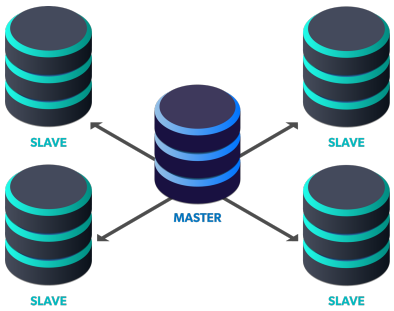

Additional Continuent Clustering Information

It is possible to have a 2-node cluster with one master and only one failover/read slave. Clusters with only two database servers must have either a passive or an active witness server to ensure proper quorum during failovers. **We do not recommend a 2-node cluster topology; always consider adding a third node to your cluster.** In order to be able to avoid split brain, a cluster needs an odd number of members such that if there is a network partition, there's always a chance that a majority of the members are in one of the network partitions. If there is not a majority, it's not possible to establish a quorum and the partition with the master, and no majority, will end up with a shunned master until such time a quorum is established. To operate with an even number of database nodes, a witness node is required, preferably an active witness, since the dynamics of establishing a quorum are more likely to succeed with an active witness than with a passive witness.

About Continuent

Continuent Ltd. is a leading provider of database clustering and replication, enabling enterprises to run business-critical applications on cost-effective open source software. Our customers represent many innovative and successful organizations throughout the world, handling billions of transactions daily across a wide range of industries.

For more information on our products and services, please visit www.continuent.com, email us at sales@continuent.com or call us at (800) 270-9035, and follow us on Twitter @Continuent.