

MySQL CLUSTER

WEB SCALABILITY WITH CARRIER-GRADE AVAILABILITY

HIGHLIGHTS

- Auto-sharding for high read and write scalability
- SQL & Key-Value NoSQL interfaces
- 99.999% availability
- Sub-second failover and self-healing recovery
- On-demand, elastic scaling with virtualization support
- On-line schema changes, maintenance & upgrades
- Low latency, real-time responsiveness
- Schema and schemaless data storage
- Low TCO: open source solution running on commodity hardware
- 24 x 7 global support and management tooling

KEY FEATURES OF MYSQL CLUSTER

- ACID compliant, transactional database
- Distributed shared-nothing design
- Multi-master database with parallel architecture
- Automatic and user-defined sharding
- Cross-shard JOINS with Adaptive Query Localization
- Synchronous & asynchronous replication with integrated heartbeating, failover and recovery
- Active / active geographic replication
- In-memory and disk-based data storage
- Online backup
- Data and index cache
- Configurable checkpoints
- SQL access
- NoSQL C++, Memcached, HTTP, Java and JPA interfaces

Data volumes are exploding – driven by increasing internet penetration rates, social networking, high-speed mobile broadband connecting ever smarter devices, and new Machine to Machine (M2M) interactions.

The databases needed to support this massive growth in data have to meet new challenges, including:

- **Scaling write operations**, not just reads, across commodity hardware;
- **Low latency** for a real-time user experience;
- **24 x 7 availability** for continuous service uptime;
- **Reducing barriers to entry**, enabling developers to quickly launch new, innovative services.

Many new applications need the back-end database to meet the above challenges, while still:

- Preserving transactional integrity with ACID compliance;
- Enabling deep insight by running complex, ad-hoc queries against the data;
- Leveraging the proven benefits of industry standards and skillsets to reduce cost, risk and complexity.

If your workloads have these demands, it is time to consider MySQL Cluster.

MySQL Cluster is a write-scalable, real-time, ACID-compliant transactional database, combining 99.999% availability with the low TCO of open source. Designed around a distributed, multi-master architecture with no single point of failure, MySQL Cluster scales horizontally on commodity hardware with auto-sharding to serve read and write intensive workloads, accessed via SQL and NoSQL interfaces.

MySQL Cluster's real-time design delivers predictable, millisecond response times with the ability to service millions of operations per second. Support for in-memory and disk-based data, automatic data partitioning (sharding) with load balancing and the ability to add nodes to a running cluster with zero downtime allows linear database scalability to handle the most unpredictable web, enterprise and telecoms workloads.

Auto-Sharding for Scaling Read and Write Operations

MySQL Cluster is implemented as an active/active, multi-master database ensuring updates can be made to any node and are instantly available to the rest of the cluster.

Tables are automatically sharded across a pool of low cost commodity data nodes, enabling the database to scale horizontally to serve read and write-intensive workloads, accessed both from SQL and directly via NoSQL APIs.

By automatically sharding tables at the database layer, MySQL Cluster eliminates the need to shard at the application layer, greatly simplifying application development and maintenance. Sharding is entirely transparent to the application which is able to connect to any node in the cluster and have queries automatically access the correct shards needed to satisfy a query or

TARGET APPLICATIONS

- High volume OLTP
- Real-time analytics
- Ecommerce, financial trading & payment gateways
- Mobile and micro-payments
- Session management and caching
- Feed streaming, analysis and recommendations
- Content management and delivery
- Massively multiplayer online games
- Presence and location-based services
- User profile and entitlement management
- Subscriber databases (HLR, HSS, VLR, etc.)
- Domain Name System (DNS) / Dynamic Host Configuration Protocol (DHCP) for broadband access
- IP Multimedia Subsystem (IMS) services
- Intelligent network nodes
- Service delivery platforms
- Softswitches
- VoIP, IPTV, and video on demand

commit a transaction.

Unlike other distributed databases, users do not lose the ability to perform JOIN operations or sacrifice ACID-guarantees when performing queries and transactions across shards.

Adaptive Query Localization pushes JOIN operations down to the data nodes where they are executed locally and in parallel, significantly reducing network hops and delivering high throughput and low latency. As a result, users can perform complex queries against their databases, enabling MySQL Cluster to serve those use-cases that have the need to run real-time analytics across live data sets, alongside high throughput OLTP operations.

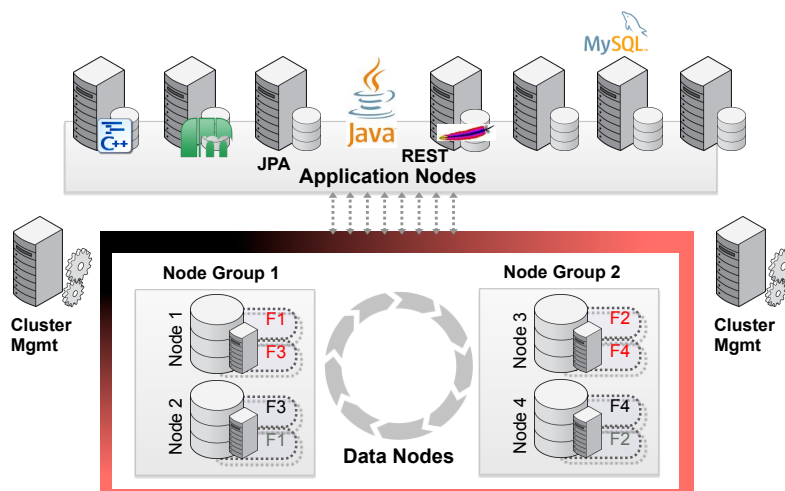


Figure 1. The MySQL Cluster architecture is designed for high write scalability and 99.999% availability with SQL and NoSQL APIs

With its real-time design, MySQL Cluster provides the response time needed to meet the needs of demanding applications. MySQL Cluster limits I/O bottlenecks by asynchronously writing transaction logs to disk, thereby achieving predictable response times in just a few milliseconds.

Five 9s Availability

MySQL Cluster is designed for five 9s availability, eliminating both planned and unplanned downtime. It achieves this via a distributed, shared-nothing architecture and synchronous replication of data that automatically propagates transaction information to all appropriate database nodes and ensures consistency of reads and writes.

Sub-second Failover and Recovery

Any failures are detected instantly and control is automatically failed over to other nodes in the cluster, without interrupting service to the clients. Plus, MySQL Cluster database nodes can automatically restart, recover, and dynamically reconfigure themselves in case of failures. The MySQL Cluster self-healing features are completely transparent to all applications.

On-Line Operations

To further support continuous operation, MySQL Cluster enables the dynamic addition of nodes to running clusters, re-sharding, back-ups and upgrades / maintenance of the cluster's underlying hardware and software infrastructure. MySQL Cluster is designed to scale on-demand, allowing services to start small and grow rapidly as demand takes-off.

In addition, MySQL Cluster supports on-line updates to live database schema, enabling users to rapidly evolve new applications by adding new columns and tables and adding or removing indexes – all while continuing to serve read and write requests, and without affecting response

“MySQL Cluster Carrier Grade Edition is a product of high quality, extremely robust and meets our demands in terms of performance and high availability. We evaluated shared-disk clustered databases but the cost would have been at least 10 times more.”

—Alain Chastagner,
Systems Manager,
Alcatel-Lucent

“Since deploying MySQL Cluster as our e-commerce database, we have had continuous uptime with linear scalability, enabling us to exceed our most stringent SLAs.”

—Sean Collier,
CIO and COO,
Shopatron

times.

Cross Data Center Deployment: Scaling and Disaster Recovery

Today’s services are global and so developers will want to ensure their databases can scale-out across regions, with added resilience to local data center failures. MySQL Cluster offers Geographic Replication that distributes clusters to remote data centers, serving to reduce the affects of geographic latency by pushing data closer to regional users, as well as providing disaster recovery. Replication between clusters is asynchronous, and each local cluster is fully active, eliminating the overhead of maintaining hardware that is largely idle.

MySQL Cluster also supports the splitting of a cluster’s data nodes across data centers. Users can synchronously replicate updates between sites with automatic failover in the event of a node failure.

SQL and NoSQL Interfaces

MySQL Cluster offers multiple APIs for data access, each of which can be used simultaneously, across the same data set, to provide the ultimate in application flexibility:

- Relational queries using the SQL API;
- Key/Value-based web services using the Memcached and REST/HTTP APIs;
- Enterprise applications using the ORM ClusterJ and JPA APIs;
- Real-time services using the C++ NDB API.

Schemaless Data Structures with the Memcached API

To allow rapid innovation in new web-based services, developers do not have to define a database schema.

When using the Memcached API for MySQL Cluster, every Key-Value is written to the same table with each Key-Value pair stored in a single row – thus allowing schema-less data storage. Alternatively, the developer can define a key-prefix so that each key and value are linked to pre-defined columns in a specific table.

If the application needs to access the same data through SQL then developers can map key prefixes to existing table columns, enabling Memcached access to schema-structured data already stored in MySQL Cluster.

Low Total Cost of Ownership

MySQL Cluster requires no additional infrastructure, such as shared storage, and runs on the latest commodity hardware and operating system (OS) platforms. An open source solution, MySQL Cluster is an extremely cost-effective database for services demanding web scalability with carrier-grade availability.

Proven Deployments

Alcatel-Lucent, BT Plusnet, Cisco, Docudesk, Neckermann, Shopatron, Telenor, Zillow.com and many more deploy MySQL Cluster in highly demanding web, broadband and mobile communications environments.

Managing and Monitoring MySQL Cluster

The commercial MySQL Cluster Carrier Grade Edition (CGE) includes a comprehensive set of tools to manage and monitor your MySQL Cluster environment, supporting you at every stage of your project – from provisioning to scaling to optimizing and upgrading.

MySQL Cluster Manager simplifies the creation and management of the MySQL Cluster

SYSTEM REQUIREMENTS PER MYSQL CLUSTER DATA NODE (RECOMMENDED)

- **OS:** Linux, Oracle Solaris, Microsoft Windows.

Note: Mac OS X for development only.

- **CPU:** Intel and AMD x86, UltraSPARC.
- **Memory:** 16 GB+ RAM (1 GB minimum).
- **HDD:** 18 GB+ (3 GB minimum).
- **Network:** Recommended minimum: Gigabit Ethernet

database by automating common management tasks. As a result, DBAs and system administrators are more productive and able to focus on strategic IT initiatives. At the same time, risks of database downtime, which often resulted from manual configuration errors, are significantly reduced.




Figure 2. MySQL Cluster expert advisors recommend best practices and reduce the risks of downtime

The MySQL Enterprise Monitor provides at-a-glance views of the health of your cluster. It continuously monitors your MySQL servers and data nodes, alerting you to potential problems before they impact your system using a series of Expert Advisors to recommend best practices developed by the database professionals who build the MySQL database.

Oracle Premier Support

Oracle offers 24x7, global support for MySQL Cluster. The MySQL Support team is composed of seasoned MySQL developers, who are database experts and understand the issues and challenges you face. With Oracle Premier Support, you can more rapidly innovate in the development of new services, lower cost and complexity and optimize the value of your database-driven solutions.

Oracle Premier Support for MySQL includes the following features:

- 24 x 7 global production support in 29 languages
- Direct access to MySQL support engineers, backed by the MySQL developers
- Unlimited support incidents
- Knowledge Base
- Maintenance releases, bug fixes, patches and updates
- MySQL consultative support

For end users, annual subscriptions include Oracle Premier Support for MySQL. For ISVs and OEMs, Premier Support must be purchased separately from commercial licenses.

Learn More

For additional resources, including white papers, on-demand Webinars, and customer case studies, visit <http://mysql.com/products/cluster>

To contact MySQL online, visit <http://mysql.com/about/contact/>

About MySQL

MySQL is the world's most popular open source database software. Many of the world's largest and fastest-growing organizations use MySQL to save time and money powering their high-volume Websites, critical business systems, communication networks, and commercial software.



Copyright © 2012, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0410

SOFTWARE. HARDWARE. COMPLETE.