7 Key MySQL Clustering Technologies – Looking at Specifics

Peter Zaitsev
CEO, Percona

Percona Technical Webinar
March 26, 2014

In this part

Additional technical details

Our practical experience at Percona

Look at specific solutions

Why “Cluster” ?

Single Node Systems are easy to deal with for Developers and Ops

Because we need

High Availability
- Business critical
- Can't have

Scalability
- Larger systems

What Do we need to Scale ?

Reads

Writes

Data Size

Cluster is

Multiple Nodes + Network

- Many more moving parts
- More room for mistakes
- Frequent hardware bugs and failures
- Software Bugs

Cluster – Testing = Low Availability

Cluster – Automation = Extreme Operations Pain

… and usually mistakes causing downtime

Automation Needed

Efficiency

Consistency

Cluster Types

Tightly Coupled

- Seen as a single system
- Simple to work with
- Danger of global failures

Loosely Coupled

- Seen as multiple systems
- Harder on developers
- More resilient

Replication Types

Synchronous
- Performance

Asynchronous
- Consistency

Semi-Synchronous
- Asynchronous with better recovery

Virtually Synchronous
- Asynchronous with conflict prevention and better recovery

Failover Process

Taking a Decision

- Automatic
- Manual

Doing Actual Failover

- Automatic
- Manual

Solution from MySQL World

There is no Silver Bullet

MySQL Replication Based

Asynchronous or Semi-Synchronous

Application must be Replication-Aware for Scaling

Proxy Solutions may help but do not offer a full solution

Often used to create building blocks for Sharded Environment

Management Tools

MHA
- Minimizing data loss by choosing the right slave

PRM
- Pacemaker Based
- Handling complicated failure modes

MySQL Replication Utilities
- Using MySQL 5.6 GTIDs

Home Grown Solutions
- Targeting Specific Needs

External Replication

Continuent Tungsten
- Asynchronous
- Intelligent Proxy and Failover Process

SymmetricDS
- Trigger Based
- Many-to Many Synchronization

Home Grown Solutions
- Using Triggers
- Using Timestamps
- Using Binlog Parsing

Storage Level Replication

Technologies
- DRBD
- SAN
- NAS

Limitations
- Active-Passive (resource waste)
- Performance overhead
- Long failover time

MySQL Cluster

- High Availability and Scalability
- Replication and built-in Sharding
- Tightly coupled
- Complicated
- Needs a very high performance network
- Limits and differences vs. InnoDB

Management Tools

MySQL Cluster Manager

- By Oracle

Cluster Control

- By SeveralNines

Percona XtraDB Cluster
(and other Galera based clusters)

- Based on well understood Innodb tables
- Virtually Synchronous or Synchronous
- Loosely Coupled (Each node has all data)
- Scale Reads and some writes
- Network Communication on COMMIT only (Local Reads)
- Bad for transactions changing many rows
- Be aware of Cluster-wide Optimistic Locking

Management Tools

SeveralNines Cluster Control
- Market Leader
- Available with Percona Support Subscription

MariaDB Enterprise
- Rather Early Release
- Fully works with MariaDB Galera Cluster Only

MySQL Compatible “NewSQL”

Clustrix

MemSQL

GenieDB

ScaleDB

Sharding

Multi-Server Horizontal Partitioning

Can be built in into Database Engine

Spread data on Replication based “Clusters”

Spread data on PXC based “Clusters”

Solutions for Sharding

Built-In
- MySQL Cluster, Clustrix etc

Manual
- Many ad-hoc in-house solutions

Proxy Level
- ScaleArc, ScaleBase, Tesora, MySQL Proxy, MaxScale

MySQL Fabric
- New API developed by Oracle

Open Source Frameworks
- Hive, Vitess, JetPants

Getting Help

Multiple Vendors available for OpenSource products

Cover the most commonly used yechnologies in Percona Support, Consulting, RemoteDBA

MySQL Replication with PRM and MHA

Percona XtraDB Cluster and other Galera Solutions

Partnership with SeveralNines

Partnership with Continuent

PERCONA LIVE: MYSQL CONFERENCE AND EXPO 2014

Learn from leading MySQL experts. Santa Clara, CA. April 1 - 4, 2014.

Special Discount for Webinar Attendees:
Use Code **WebinarSC** to receive 15% off of standard rates (new registrations only)

http://bit.ly/MySQLUC2014

Now with Open Source Appreciation Day!

- March 31, 2014
  - ✓ http://bit.ly/1gysW6B
- Two Events To Chose From
  - ✓ CentOS Dojo
  - ✓ OpenStack Today
- Attendance is **FREE**

Thank You!

Peter Zaitsev
pz@percona.com