



MySQL Replication and Clustering for Web Technologies

Brian Aker

November 2008

MySQL AB

Terms

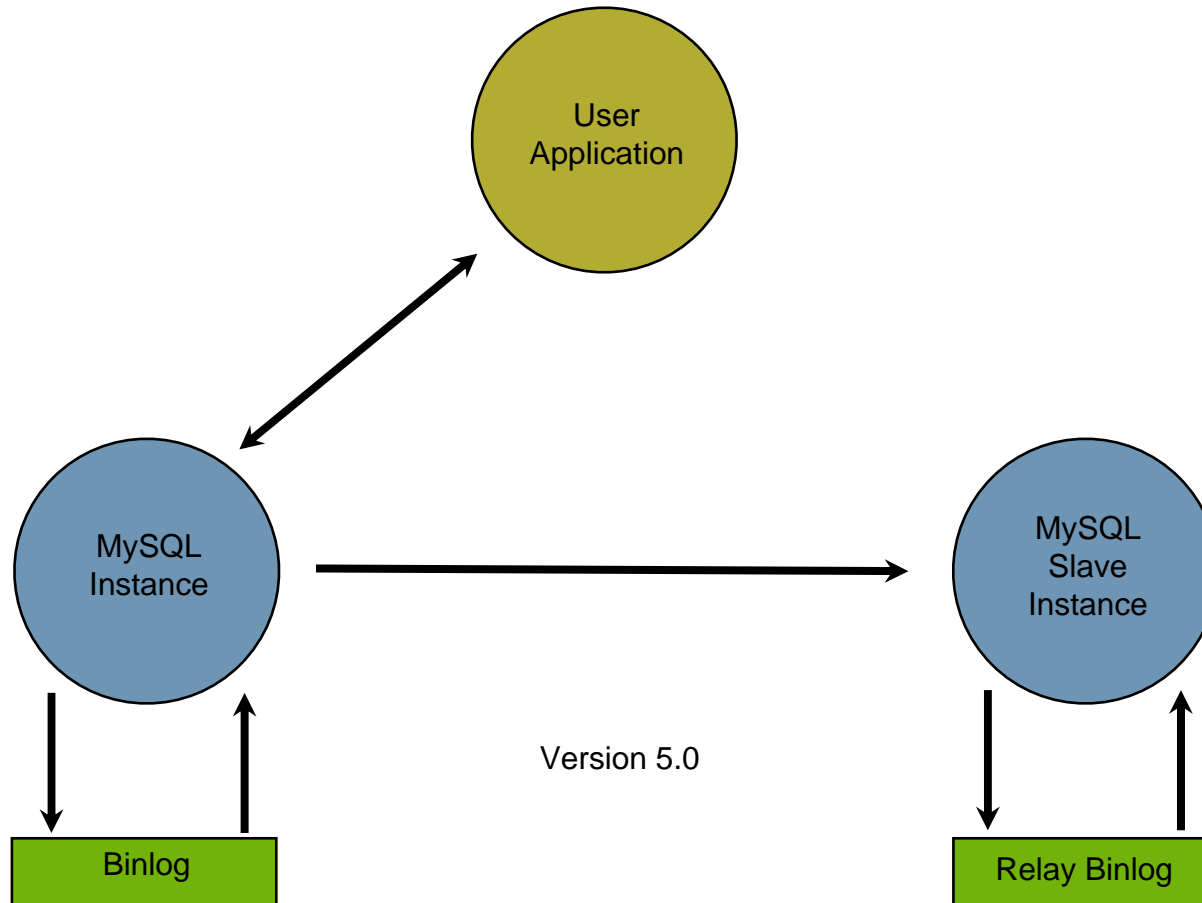
- Instance - one running process of MySQL
- Schema - the container of your objects (the result of a CREATE DATABASE/SCHEMA command)
- Master - the instance where data is first written to by a client application.
- Slave - the copy of the master
- binlog -file containing replication data



What is Replication?

- Copying data from one instance of MySQL to another where all results are guaranteed to be the same on both end.
- Atomic
- Non-blocking

Today's Replication



5.0 Architecture

- Master
 - Connecting Slaves
 - Connecting Masters
- Slaves
 - IO Threads
 - SQL Threads



Clustering != Replication

- Replication is not real time
- Replication by itself is not fault tolerant
- No two-phase commits

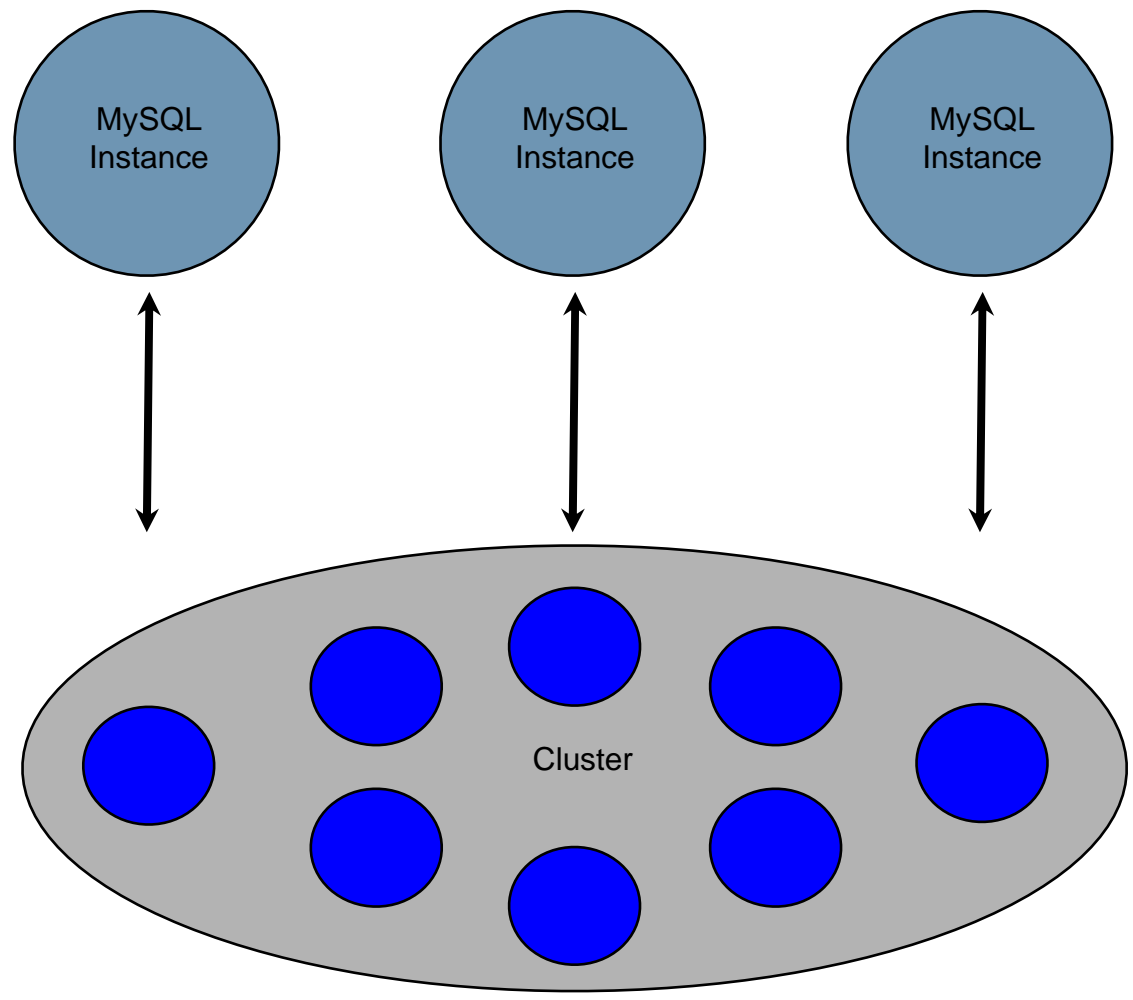


Replication == Clustering

- Allows you to spread your reads out over multiple hosts thus reducing your single points of failure.
- Create databases for specific tasks



MySQL Cluster





MySQL Cluster

- Shared Nothing
- Transactional
- In Memory (though not in 5.1!)
- Fast Primary Key Lookup
- Integrated with MySQL
 - Data Marts are as easy as
“CREATE TABLE SELECT”

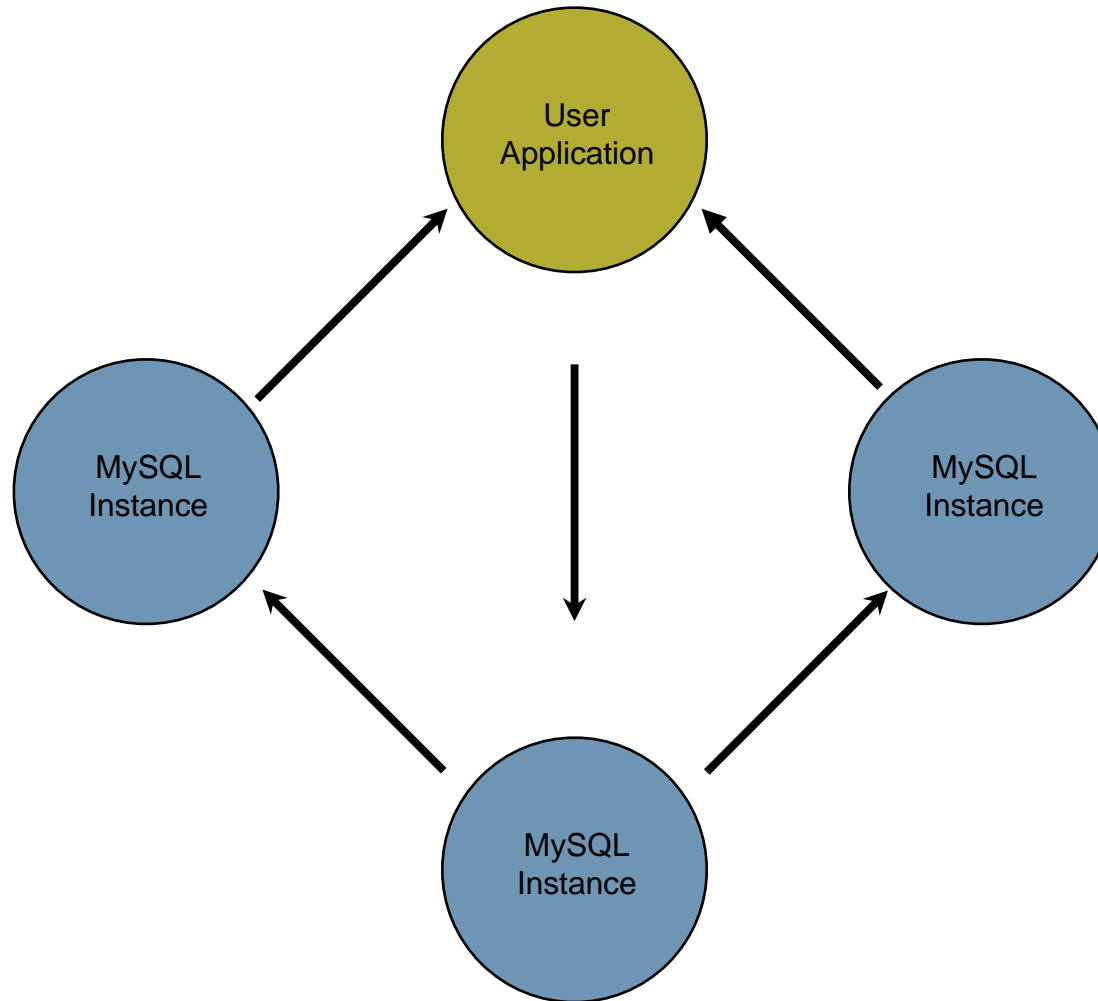


Common Paradigms

- Read Replication
- Split Loads
- Star Clusters



Read Replication Cluster

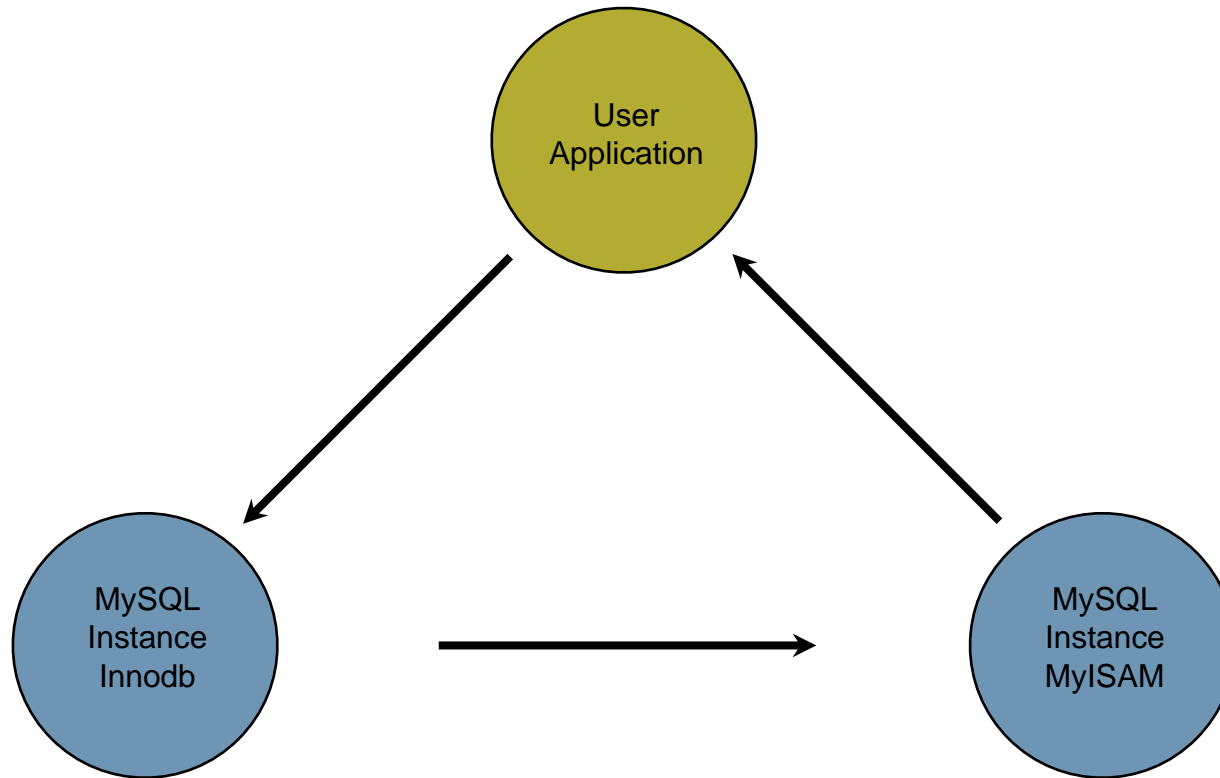




Read Replication Cluster

- Disadvantages
 - Not synchronous
 - Central point of failure in writing
- Advantages
 - Fail over to read only state made easy
 - Distributed reads that are easy to scale

Application Cluster

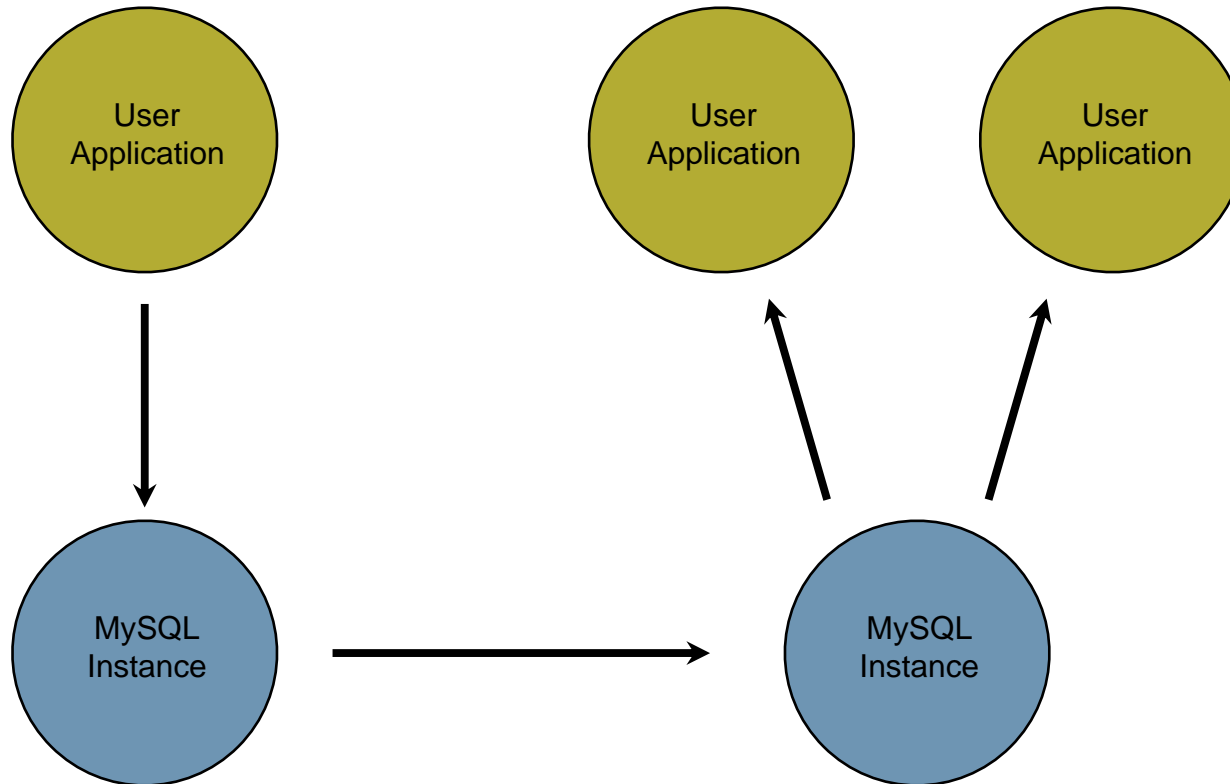




Application Cluster

- Full Text and Innodb
 - Search for Slashdot
 - DOS waiting to happen
 - High Cost
 - Full Text vs Innodb
- Replicate Data to separate host!

Logging Example



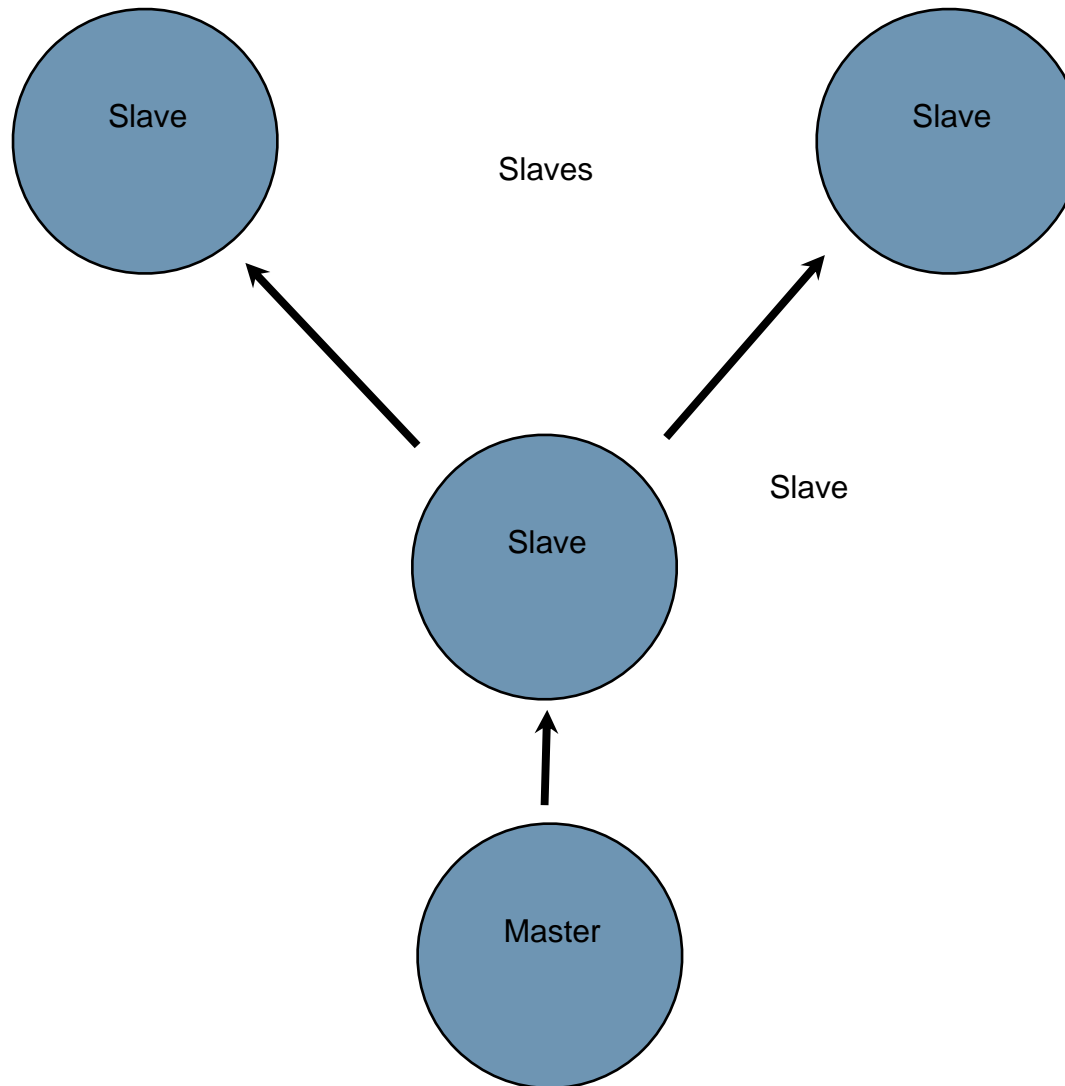


Logging Example

- Insert intensive / IO Critical
- Replicate to slave where expensive queries can be run



Master to Slave to Slaves

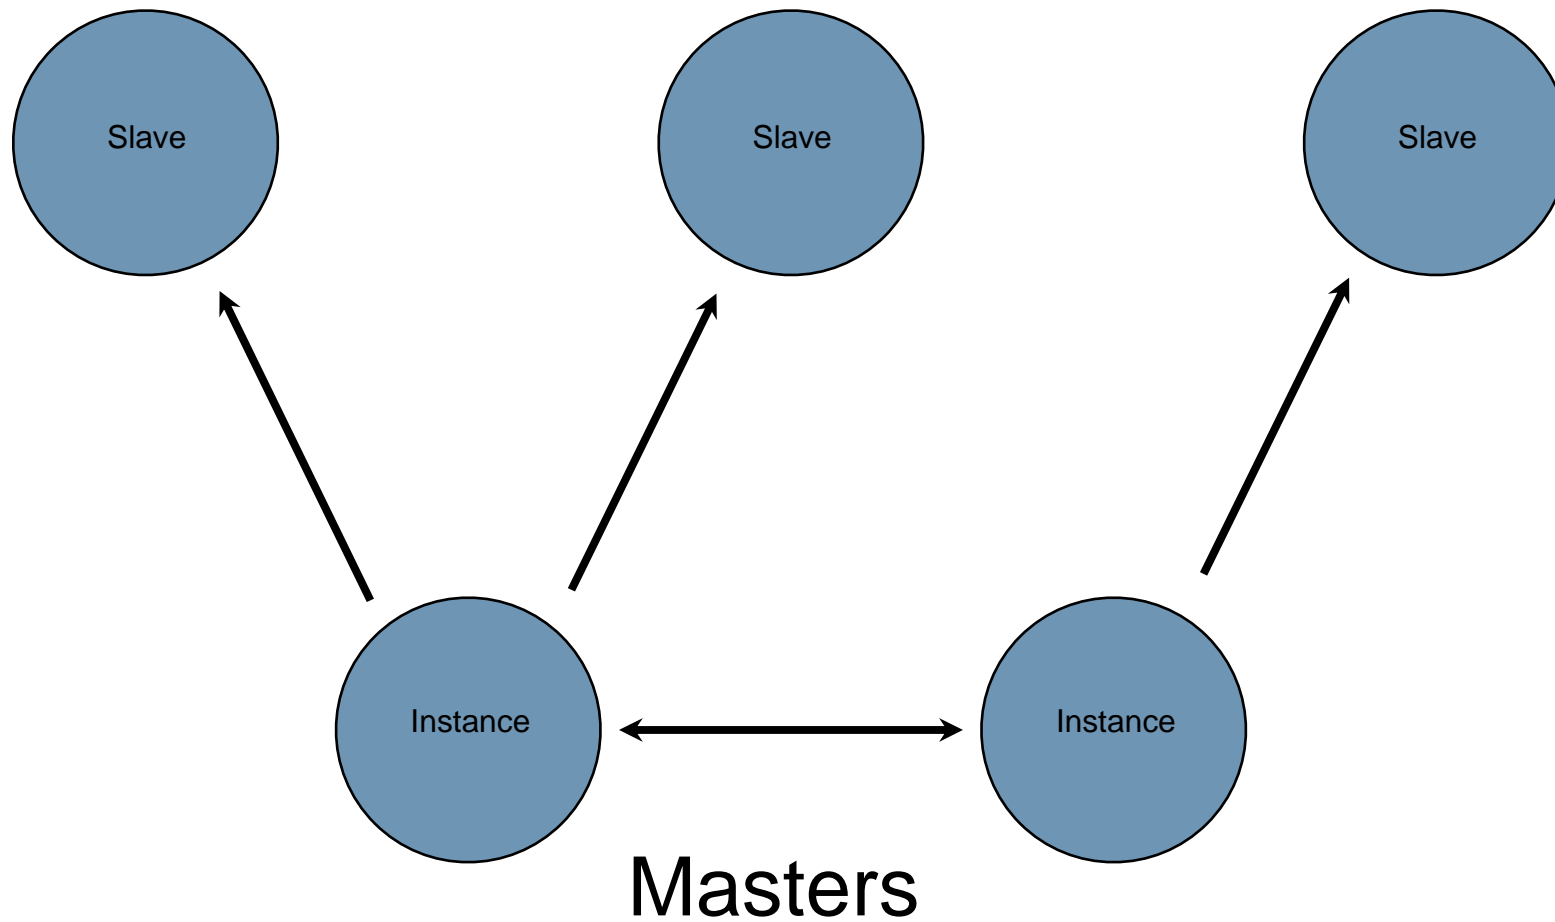




Master to Slave to Slaves

- Lowers IO on Master
- Speedier access to logs with dedicated host


Writing Cluster



Writing Cluster

- Advantage
 - No single point of failure
- Disadvantages
 - Must rethink primary keys (GUIDs)
 - Relational weaknesses

Split Master

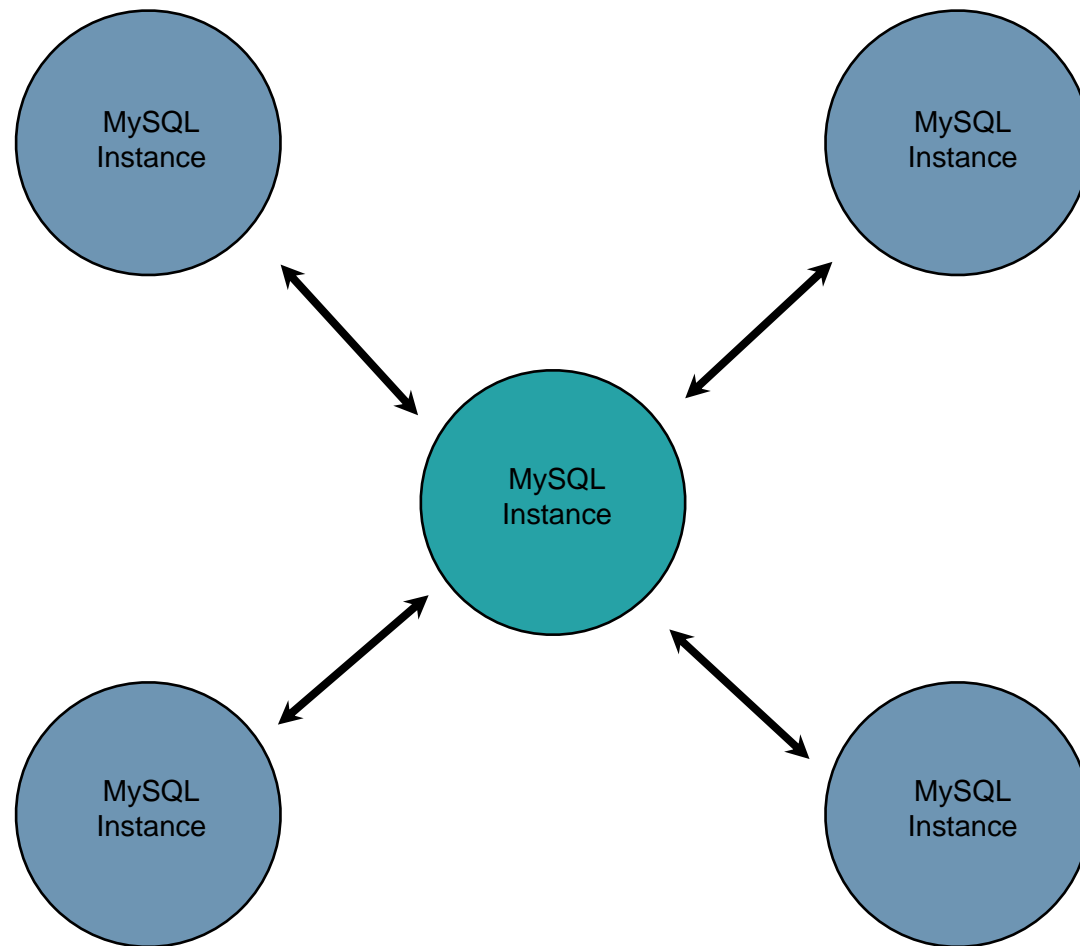




Why Split Masters

- Easier to failover
- Slow Queries
- Easier Backups


Star Cluster



Star Cluster

- End master nodes feed master node that in returns feeds end nodes.
- System can come and go as needed.
- Good for situations where masters can not always be assumed to exist.
- Does require a script to be used to push replication logs from edges to center.
- Requires primary key conflict resolution

Circular







Circular Cluster

- Easy to add new servers
- Removing servers is one command
- All hosts can be masters

Adding a Node to Circular



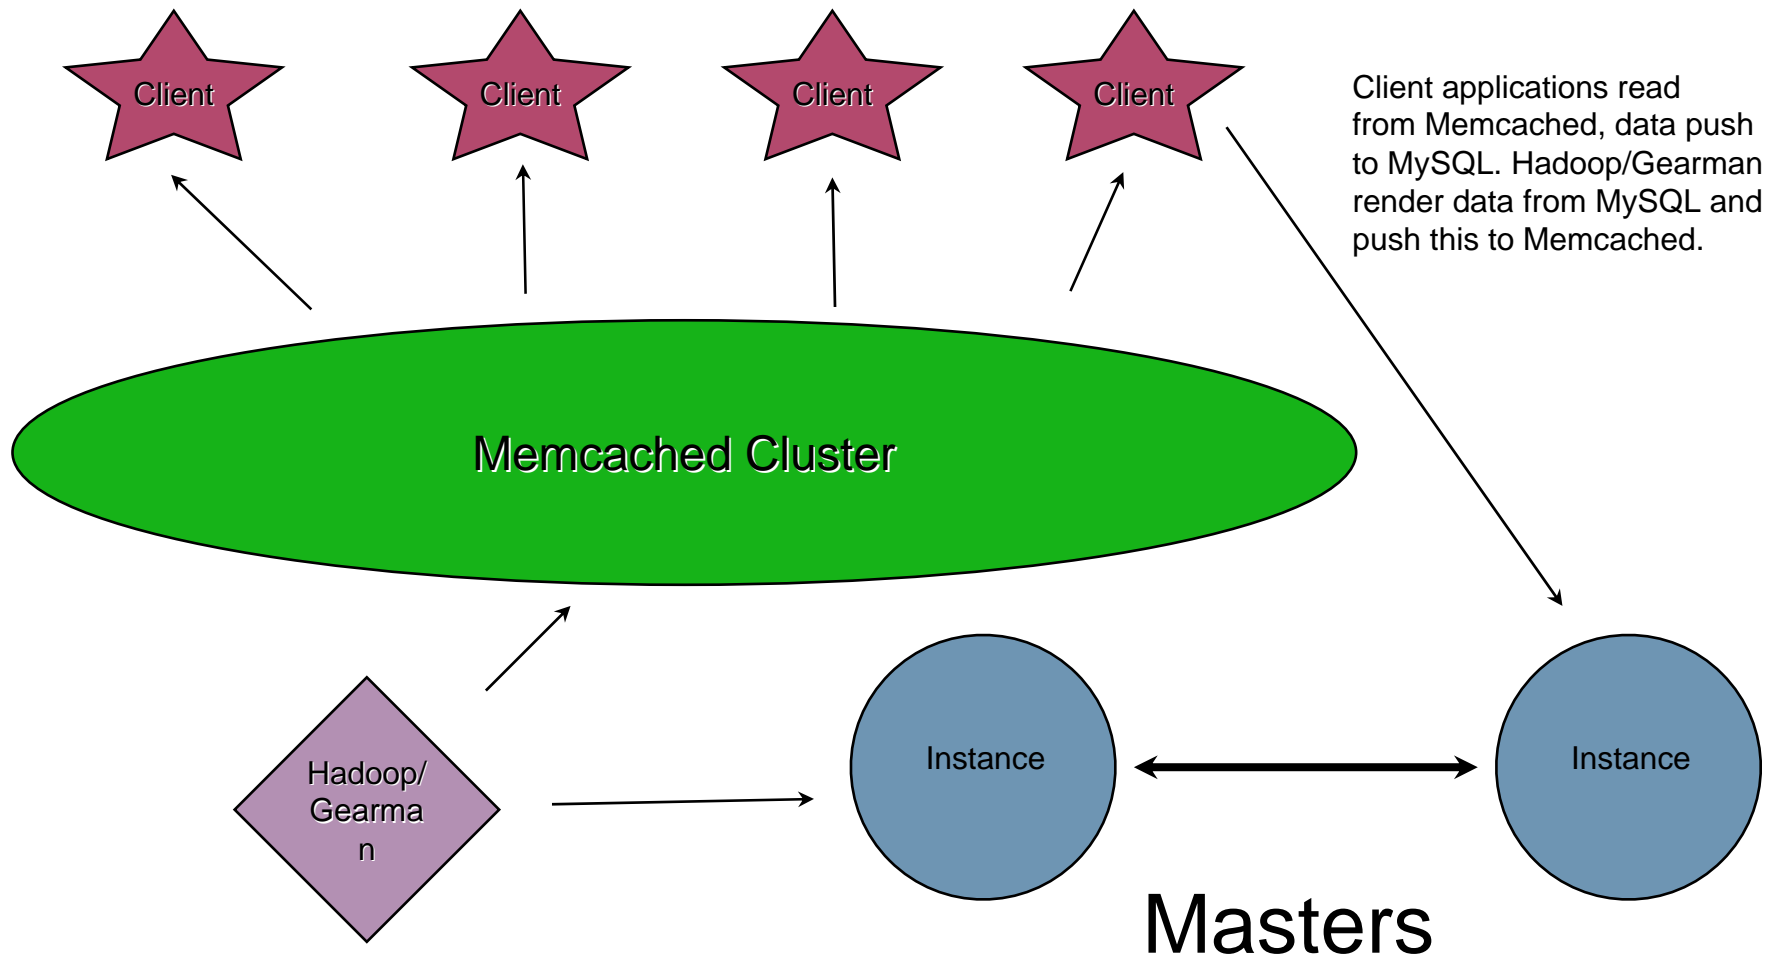
Adding a Node to Circular



Federated

- Remote Tables look like local tables (full join support)
- Can be used with Triggers

Memcached Clustering





More Information

- Look at Documents on mysql.com
- lists.mysql.com
- MySQL Support Contracts