SchoonerSQL

Meng Wang

Feb 27th 2012

Roadmap

- Introduction
- Architecture
  - Cluster Architecture
  - Synchronous replication
- Failover and Recovery
- Scaling cluster sets
  - Asynchronous Replication and failover

- Built in 2007 at Sunnyvale, CA
- Founded by John Busch
- Use souce code of MySQL and InnoDB
- The predominant elements of SchoonerSQL's availability and performance enhancements include:
  - synchronous replication.
  - Scaling by adding additional clusters synchronized with the primary cluster via asynchronous replication.
  - Highly parallelized synchronous replication within a cluster across LANs and MANs(metropolitan area network using fiber optic connections).
  - Highly parallelized asynchronous replication between clusters across LANs, MANs, and WANs.
  - Automated failover within seconds over LANs, MANs, and WANs.

Cluster Architecture

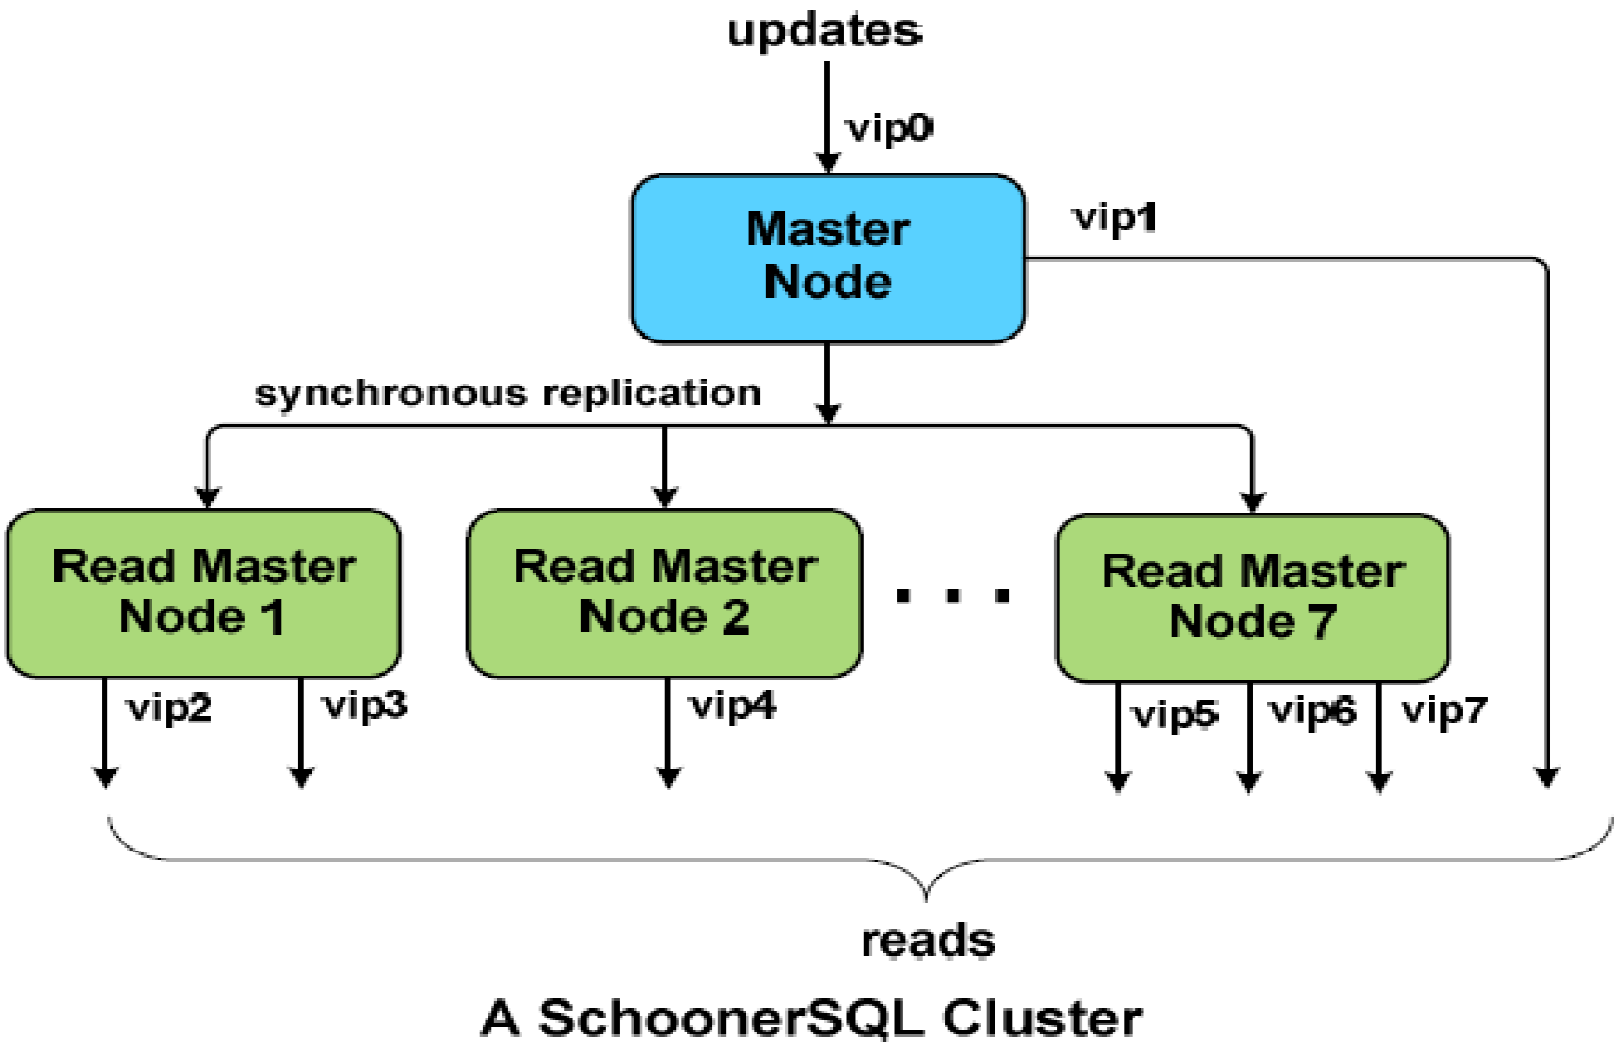

Source:http://www.availabilitydigest.com/public_articles/0701/schoonersql.pdf

Replication

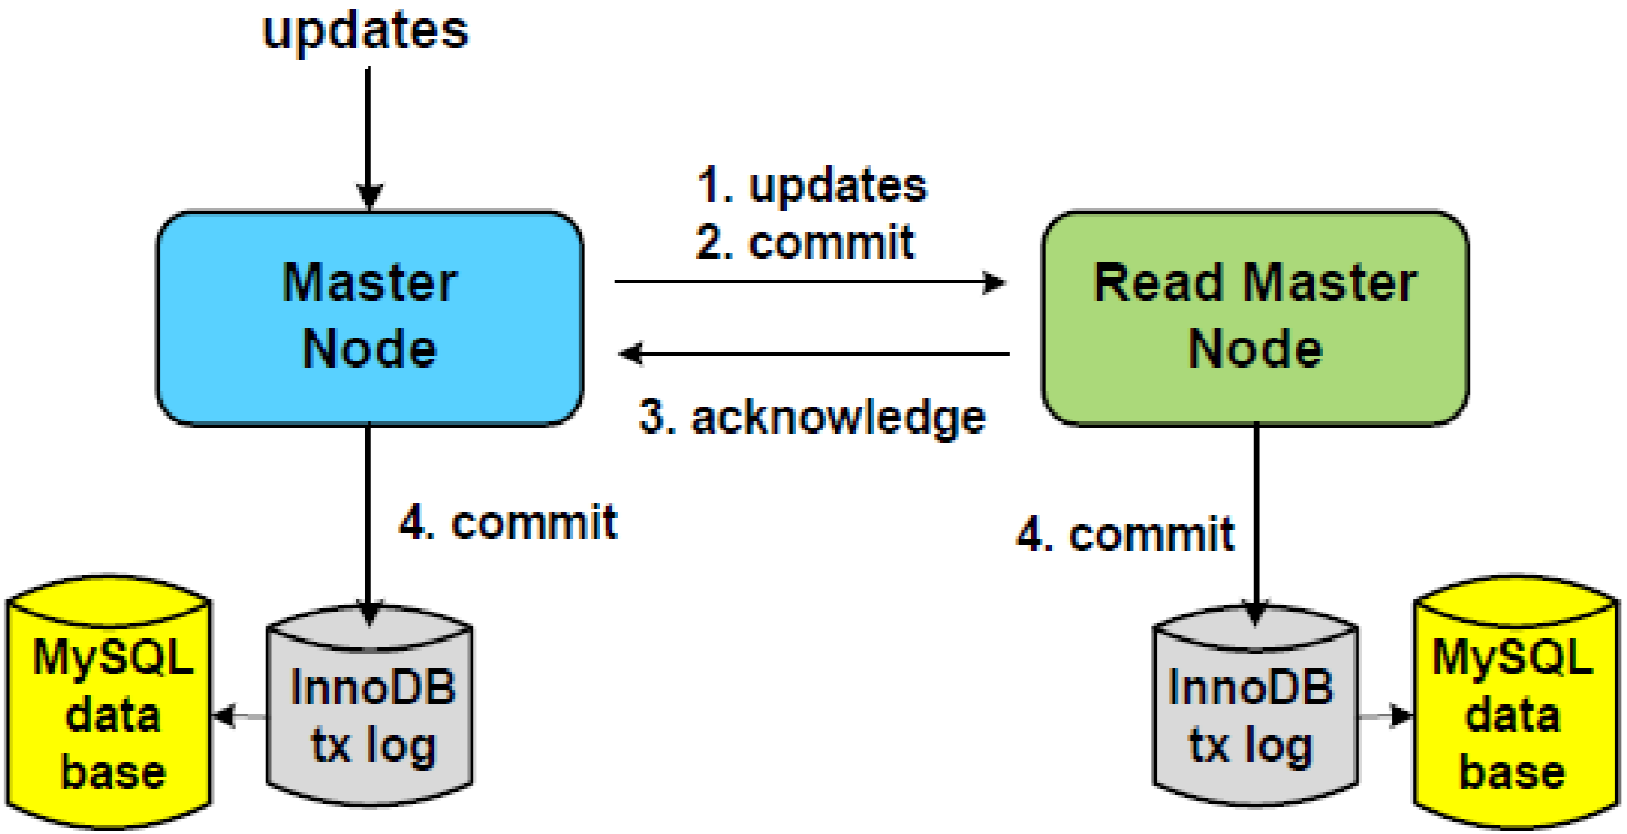

SchoonerSQL Synchronous Replication

Source: http://www.schoonerinfotech.com/downloads/white-papers/Achieving_Five_9s_with_SchoonerSQL.pdf

Replication

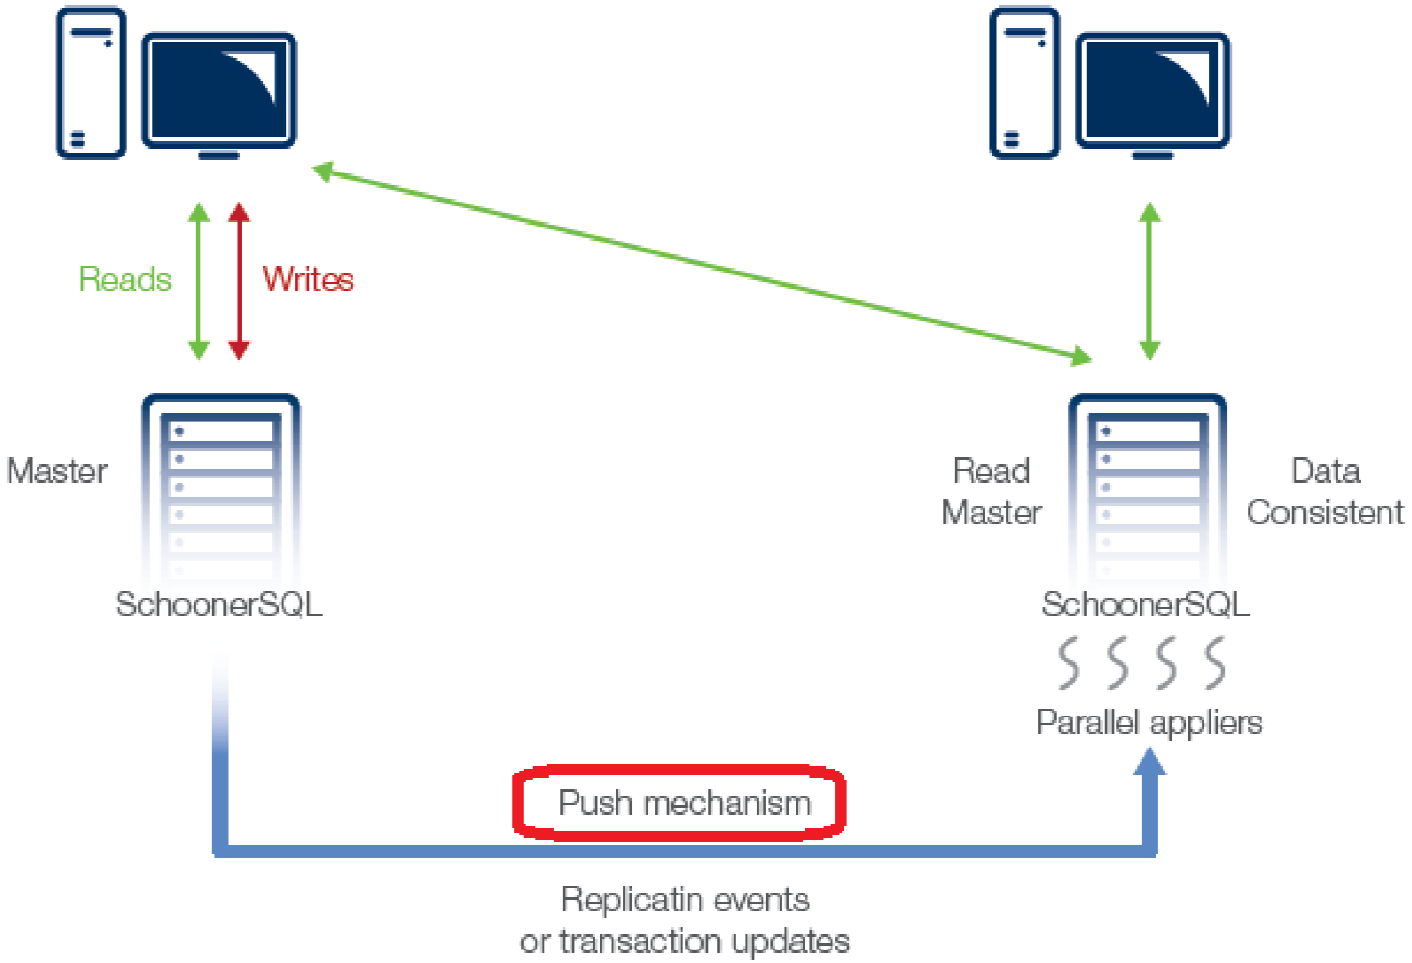

Source: http://www.schoonerinfotech.com/downloads/white-papers/Achieving_Five_9s_with_SchoonerSQL.pdf

VIP Auto Failover

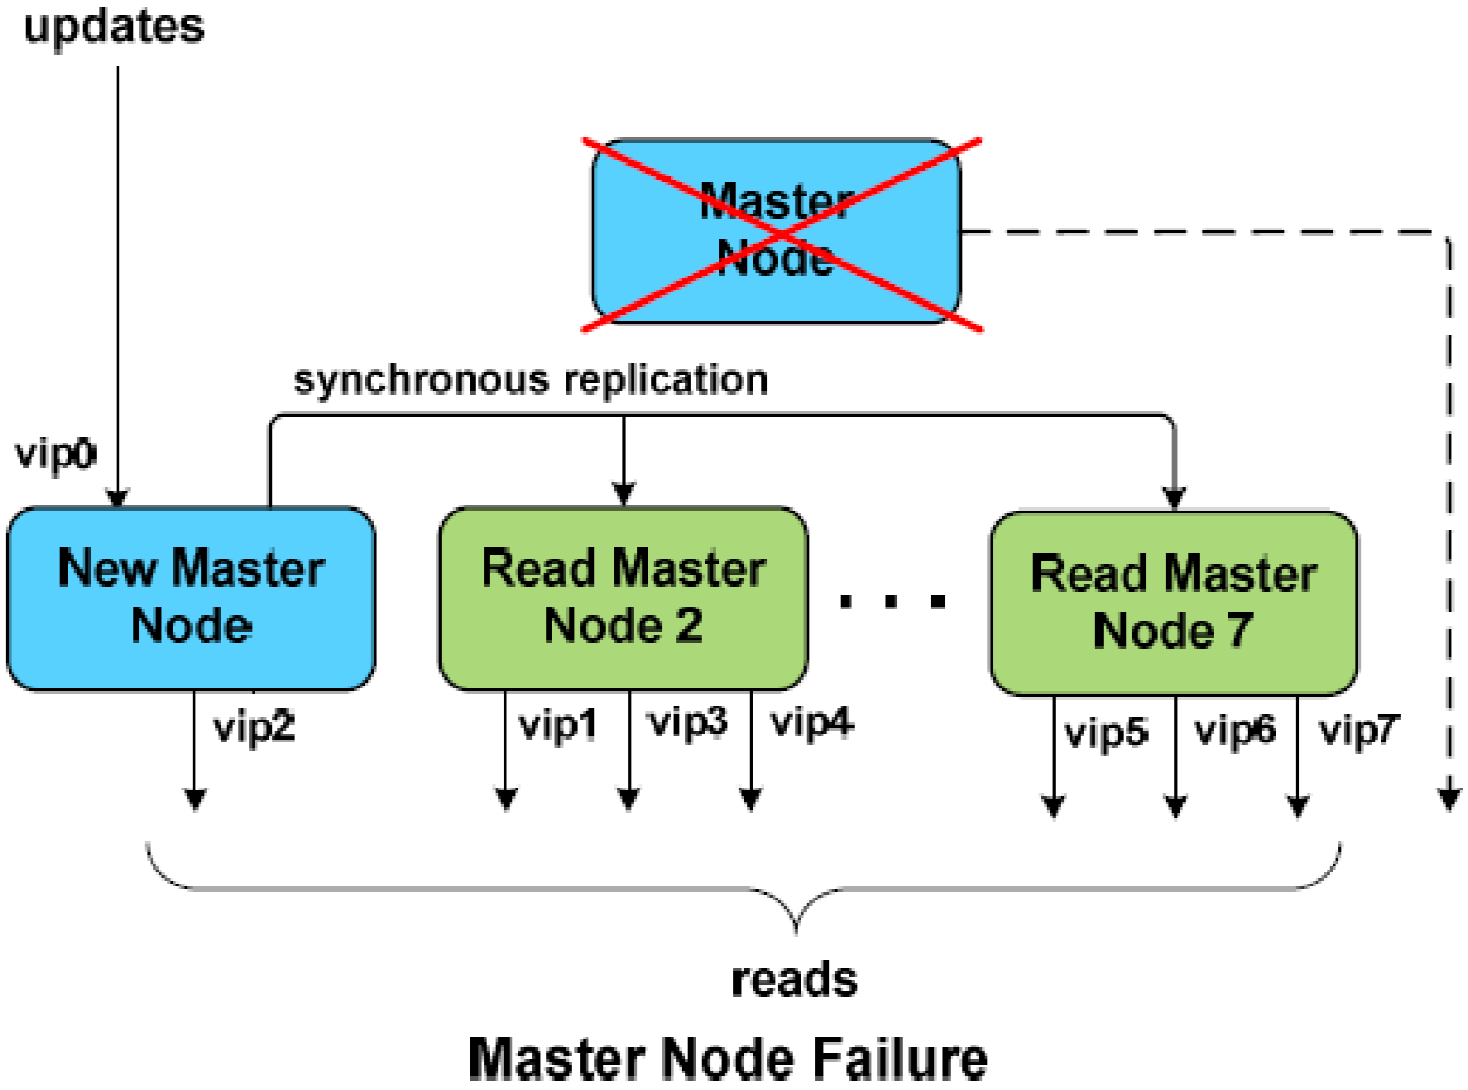

Former read - master node1 with read VIP2 and VIP3

Source:http://www.availabilitydigest.com/public_articles/0701/schoonersql.pdf

Scaling with Asynchronous Clusters

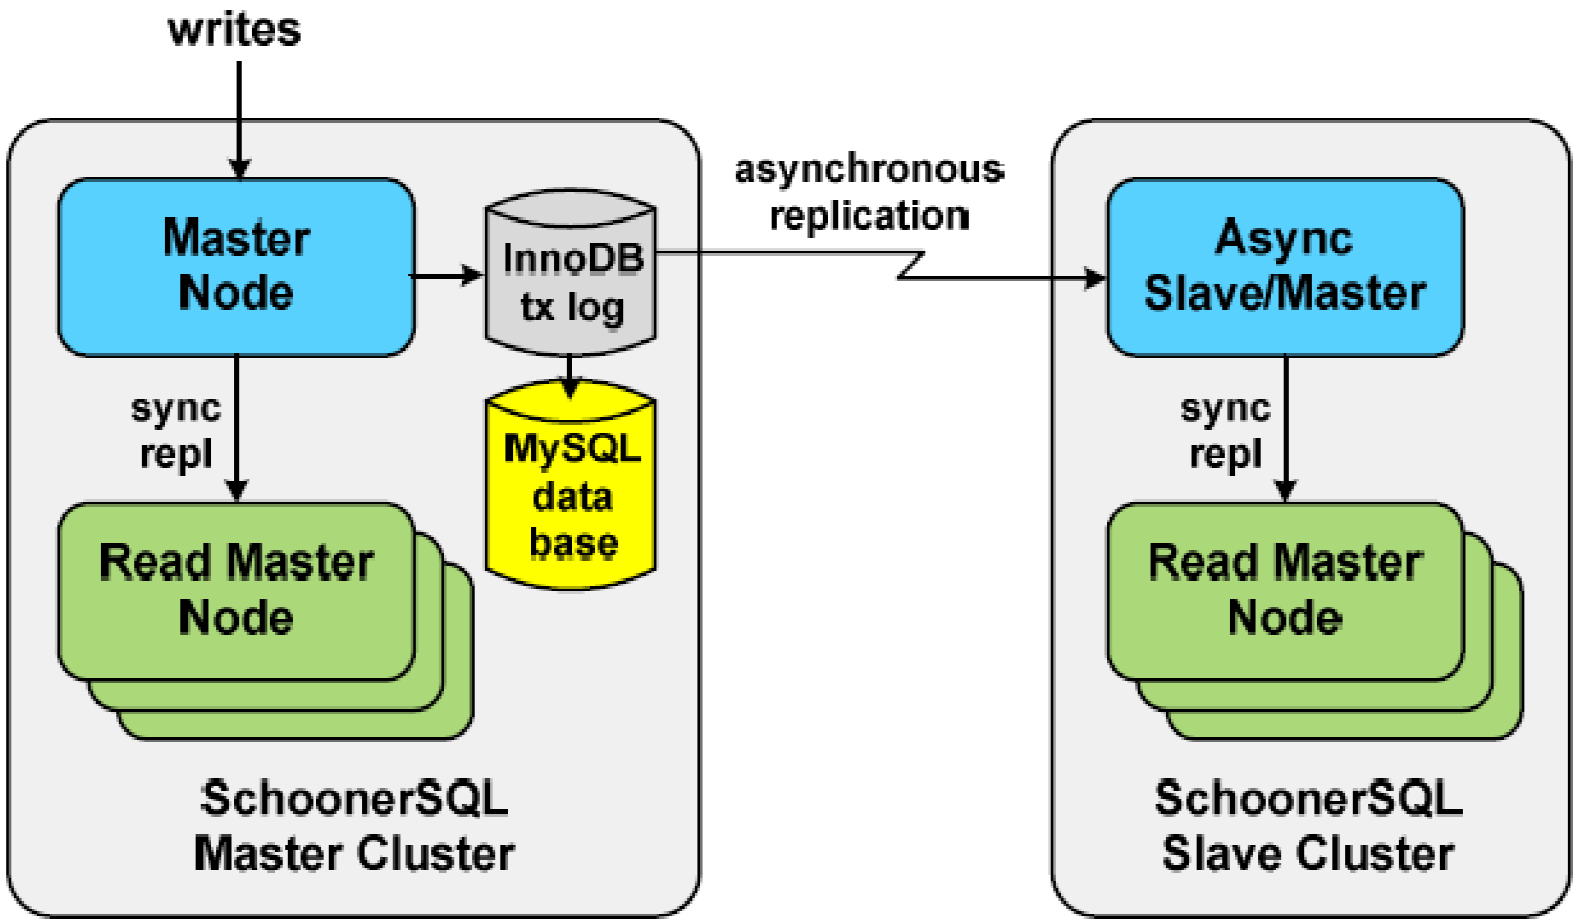

Multi-Cluster SchoonerSQL Configuration

Source:http://www.availabilitydigest.com/public_articles/0701/schoonersql.pdf

Performance

- The benchmarks are based upon a 1,000 warehouse configuration with 32 connections.
  - show a 3x performance advantage over MySQL asynchronous replication and a 2x performance advantage over MySQL semisynchronous replication when using hard disk drives.
  - For flash memory, SchoonerSQL shows a 5x performance advantage over MySQL asynchronous replication and a 4x performance advantage over MySQL semisynchronous replication.
- from hard disk drives to flash memory. Performance will also improved.

Summary

- SchoonerSQL's performance improvement is achieved by highly parallelized multithreaded replication threads for use with multi-core processors.

- SchoonerSQL's high availability is achieved through the use of synchronously replicated multi-node clusters that provide fast failover (within seconds) of a node failure.

- A SchoonerSQL cluster is scalable by configuring it with up to eight nodes. Cluster-sets are scalable by using asynchronous replication

Thank you