How to use next generation Mission Critical MySQL and NoSQL to dramatically improve Availability, Performance, and TCO

Dr. John R. Busch
Founder and CTO
Schooner Information Technology
John.Busch@SchoonerInfoTech.com

How to use next generation Mission Critical MySQL and NoSQL to dramatically improve your Availability, Performance, and TCO

Abstract

The founder of Schooner Information Technology, Dr John R Busch, discusses the architecture, benefits, and industry deployments of next generation Mission Critical MySQL. Mission Critical MySQL completely eliminates data loss, automates failover, lowers response times, provides unlimited query and update transaction scaling, and provides multi-datacenter high availability and high performance. Designed to fully exploit standard commodity servers, storage, and networking, Mission Critical MySQL dramatically reduces datacenter footprint and power requirements, typically reducing Total Cost of Ownership by more than 50%. Building on Mission Critical MySQL allows much higher availability, data integrity, performance, scalability, and ease of administration than can be achieved from legacy MySQL or application/middleware layer sharding and replication. In addition to Mission Critical MySQL architecture, benefits, and measurements, Dr. Busch will discuss industry deployments in high volume websites, including eCommerce, Social Media, Telco, and Financial Services. program.

Bio:

Dr. John R. Busch is Founder and CTO of Schooner Information technology, focusing on high performance and highly available databases for high volume web facing, cloud, and enterprise services. Prior to Schooner, John was director of computer system research at Sun Microsystems Laboratories, VP of Engineering with Diba, founder of Clarity Software, and R&D director of computer systems at Hewlett Packard. John holds a Ph.D. in Computer Science from UCLA, M.S. in Mathematics from UCLA, M.S. in Computer Science from Stanford University, and attended the Stanford Sloan program.

The Mission-Critical Imperative

- Providing high data availability, excellent response time is critical for key classes of businesses
  - Web 2.0
  - eCommerce
  - High-volume websites
  - Telecommunications

- They require a mission critical databases and data stores

Mission-Critical Database Requirements

High Availability

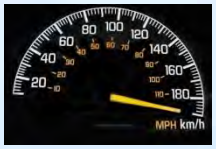

High Performance and Scalability

Simple and Powerful Administration

Data Integrity

Cost Effective

Standards and Compatibility

Mission Critical

Mission-Critical Database and DataStore Goals and Metrics

Goals	Metrics
High Availability	Service unavailability (minutes/year) from failures, disaster recovery, or during planned administration
High Data Integrity	Probability of data loss or corruption; data consistency levels
High Performance and Scalability	Transaction throughput, response time; performance scalability; performance stability
Simple and powerful administration	Ease of cluster administration; fail-over automation; monitoring and optimization tools
Cost effective	Total cost of ownership (TCO); return on investment (ROI)
Standards and Compatibility	Level of standards compliance and certification

Legacy MySQL : Loose Coupling and Asynchronous and Semi-Synchronous Replication

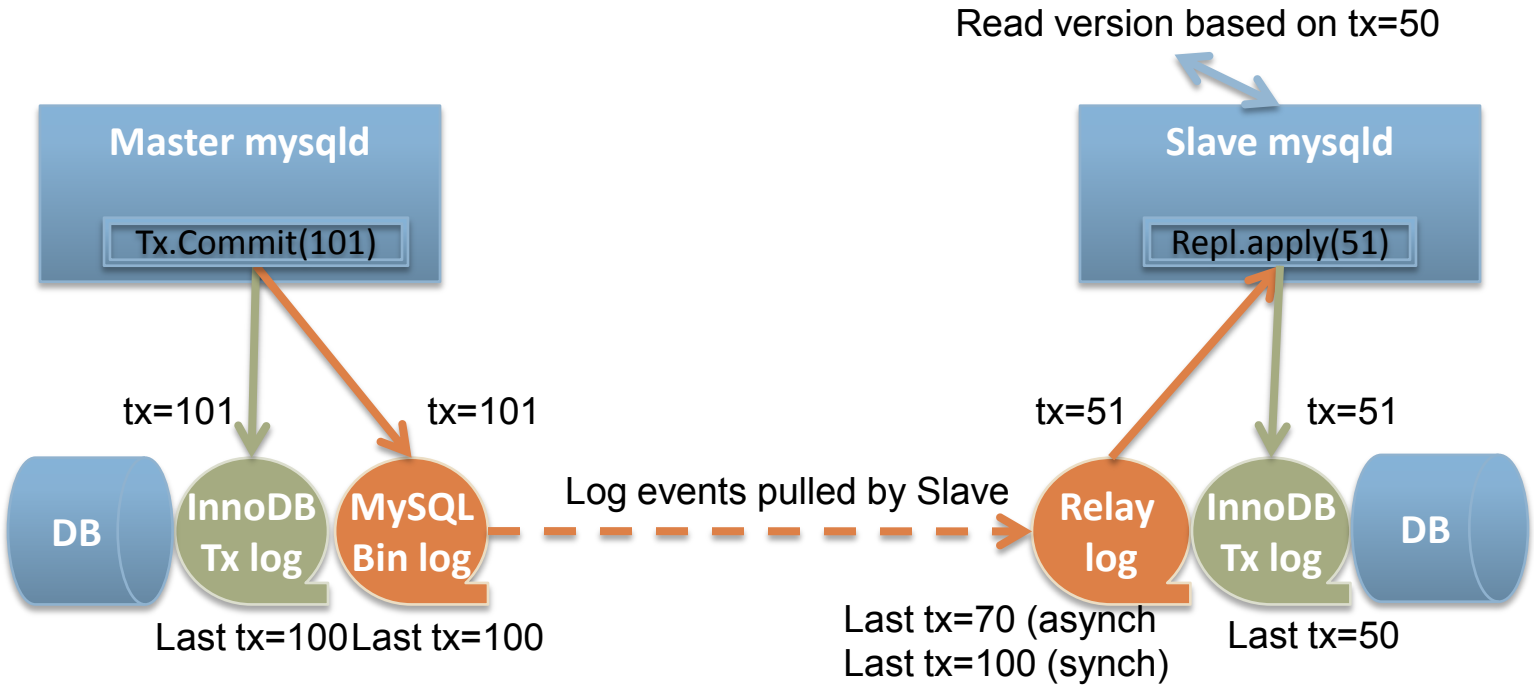

Example Products : MySQL Enterprise 5.1 Asynchronous and 5.5/5.6 Semi-Synchronous Replication

Loosely-Coupled Asynchronous and Semi-Synchronous Replication

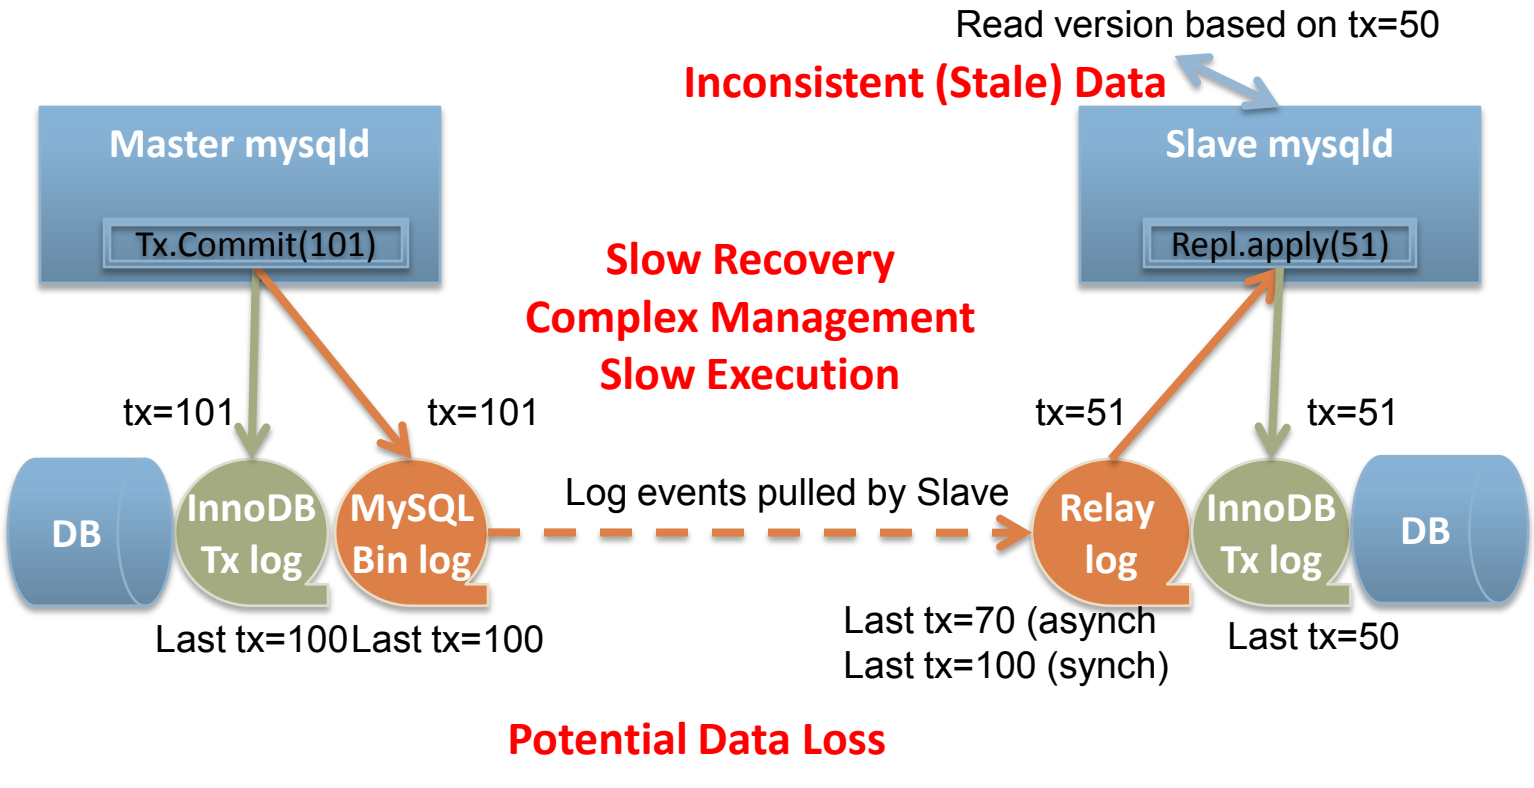

Example Products : MySQL Enterprise 5.1 Asynchronous and 5.5/5.6 Semi-Synchronous Replication

Loosely-Coupled Asynchronous and Semi-Synchronous Replication

Limited Service Availability
- Master fail-over, re-synch of slaves

Limited Data Integrity
- Lost data; inconsistent Data

Limited Performance and Utilization
- Low throughput and low utilization

Complex Administration
- Manual processes, slave re-synch

High Cost of Ownership
- High capital expense from server sprawl
- Increased operating expense from power, space, admin
- Reduced revenue and customer satisfaction from service downtime

SchoonerSQL Database Architecture for Commodity Servers
Tight Coupling and Fully Synchronous Replication

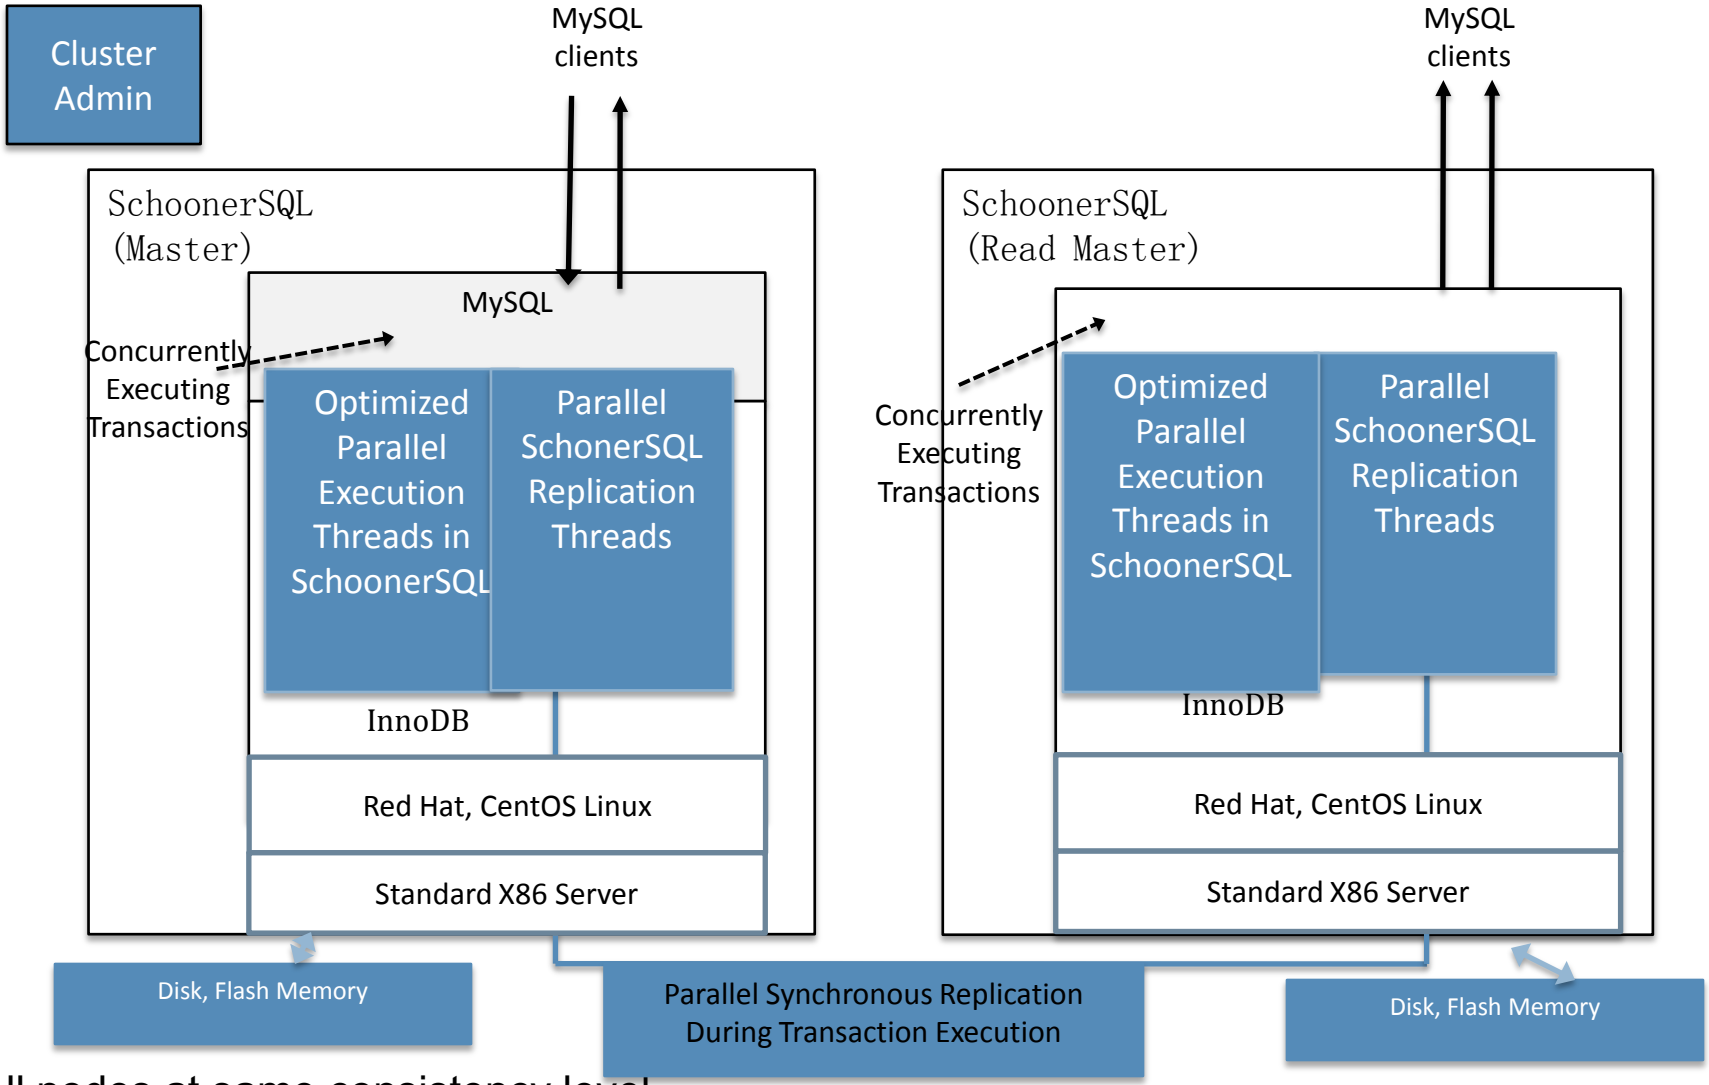

- All nodes at same consistency level
- At master transaction commit, all nodes guaranteed to have received and committed the changes

SchoonerSQL: Tight-Coupling and Synchronous Replication

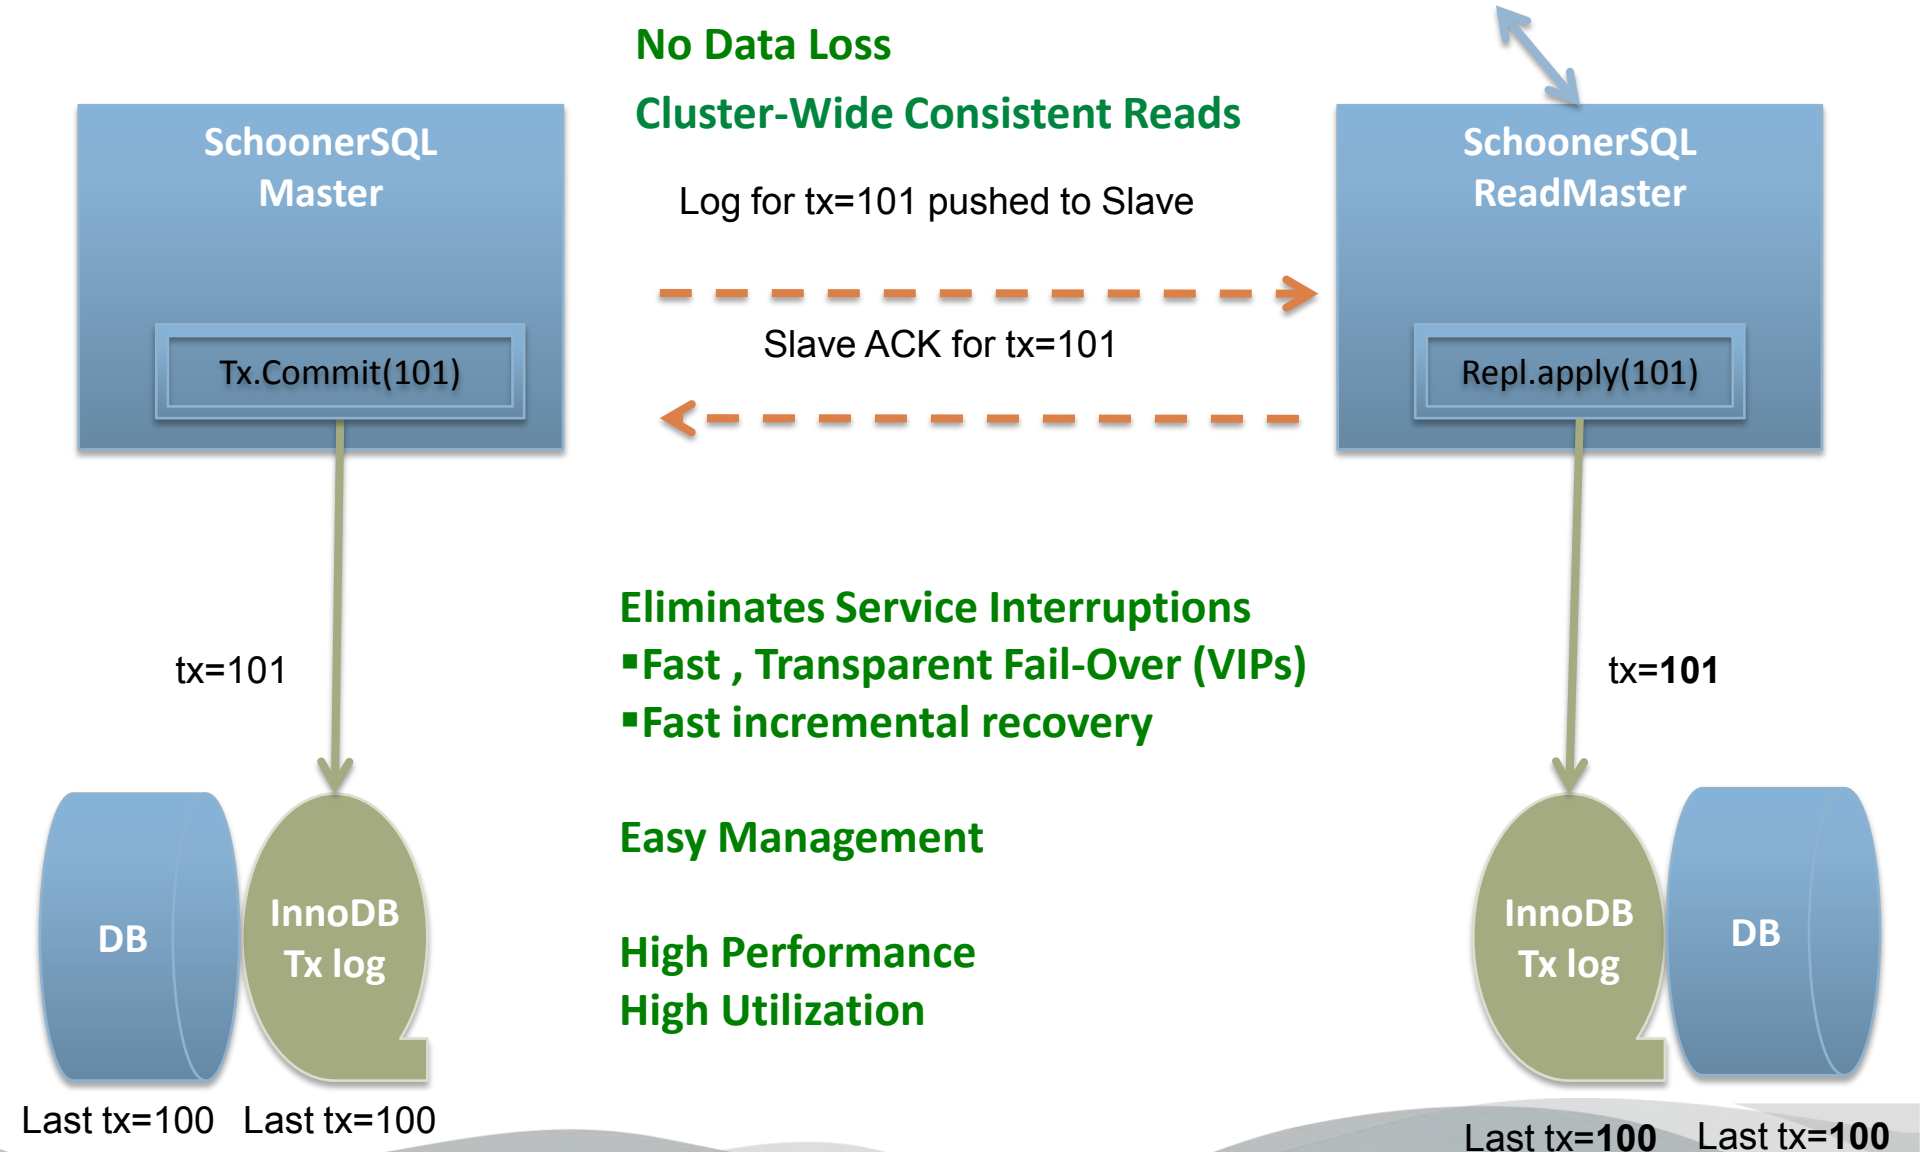

SchoonerSQL Improves Service Availability by 90%

Tightly-coupled MySQL synchronous replication with fast, automated application transparent fail-over and can provide much higher service availability than that achievable with asynchronous or semi-synchronous replication

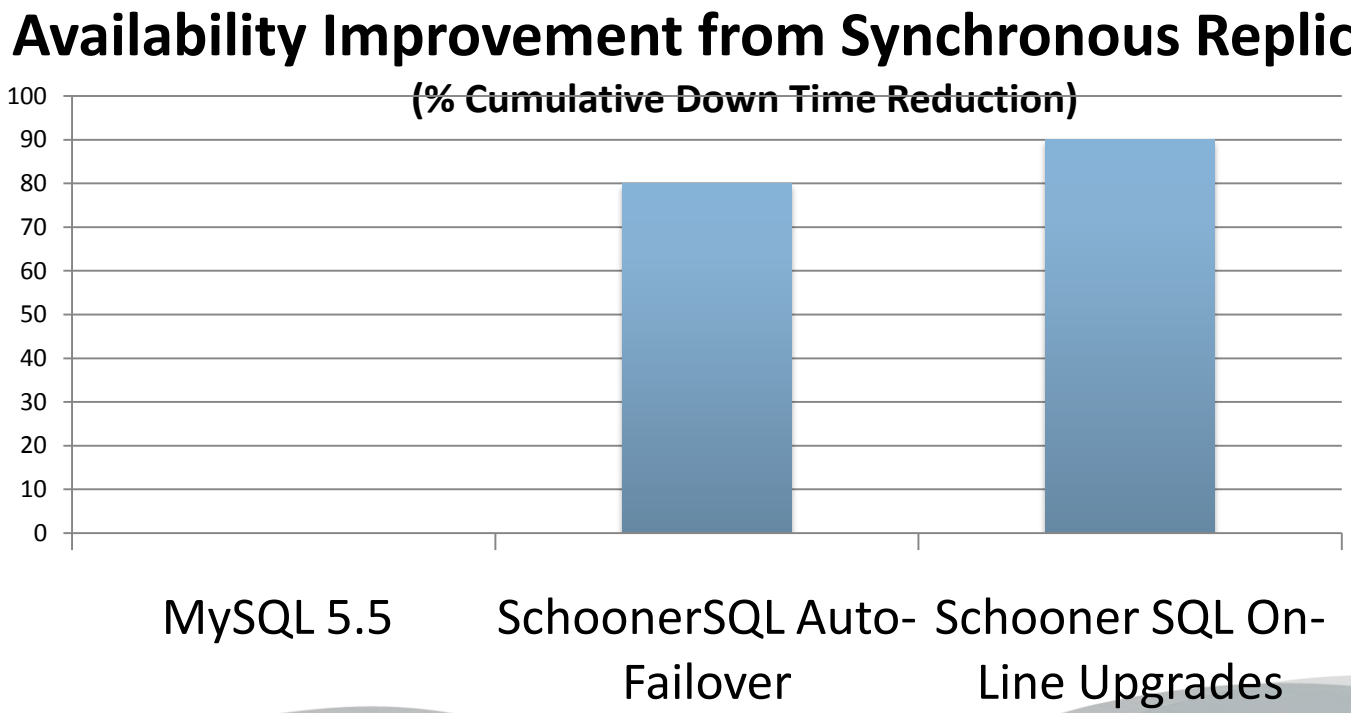

SchoonerSQL Provides Much Higher Performance Throughput per Server

Synchronous Transaction Throughput per Server can be Much greater Than Asynchronous or Semi-Synchronous (with hard disc drives (HDDs))

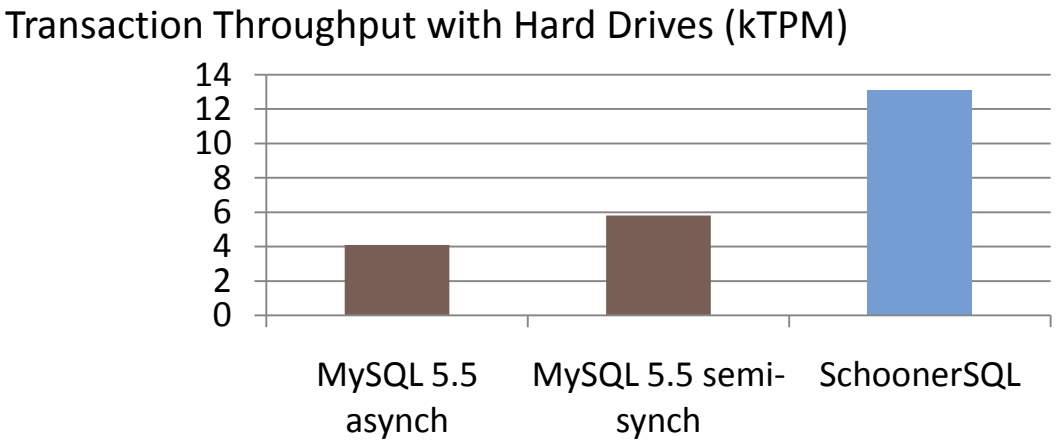

Measurement Configuration
- 2 node Master-Slave configuration
- 2 socket Westmere
- 72GB DRAM

DBT2 open-source OLTP version of TPC-C
- 1000 warehouses, 32 connections
- 0 think-time
- Result metric: TPM (new order)

SchoonerSQL Enables Vertical Scaling with Commodity Flash Memory, Cores

DBT2 open-source OLTP version of TPC-C
- 1000 warehouses, 32 connections
- 0 think-time
- Result metric: TPM (new order)

Measurement Configuration
- 2 node Master-Slave configuration
- 2 socket Westmere
- 72GB DRAM

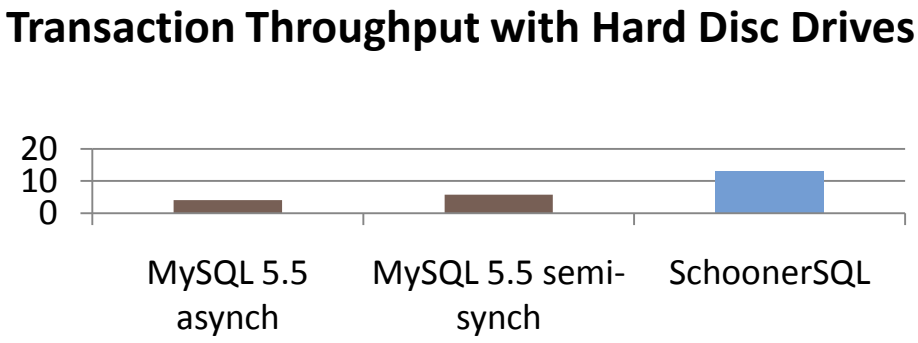

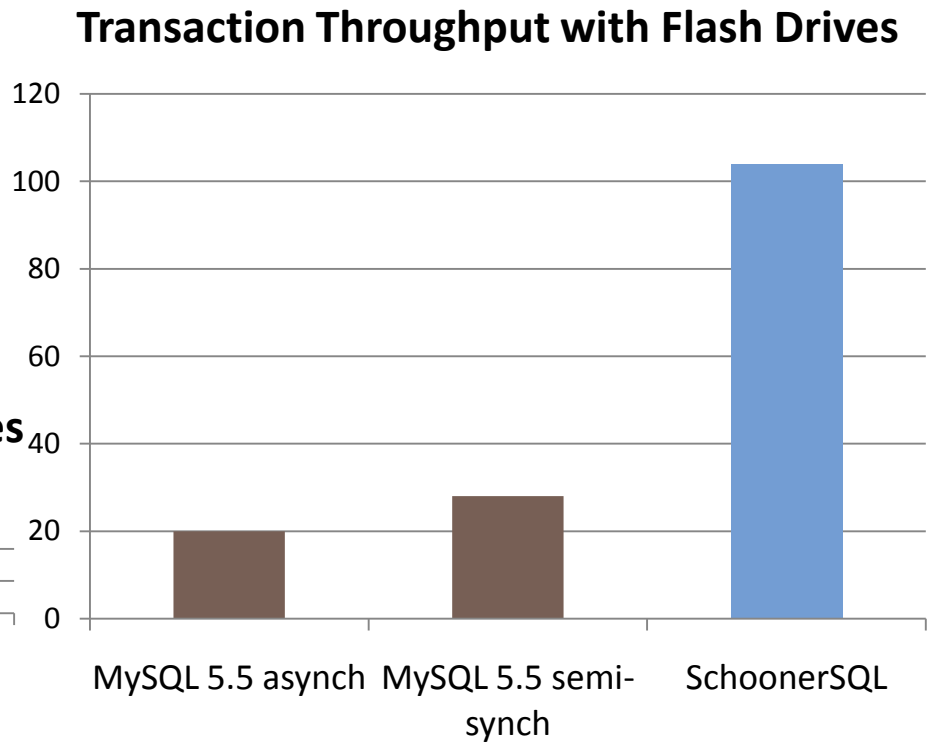

SchoonerSQL Lowers Response Times

Response Time (ms)

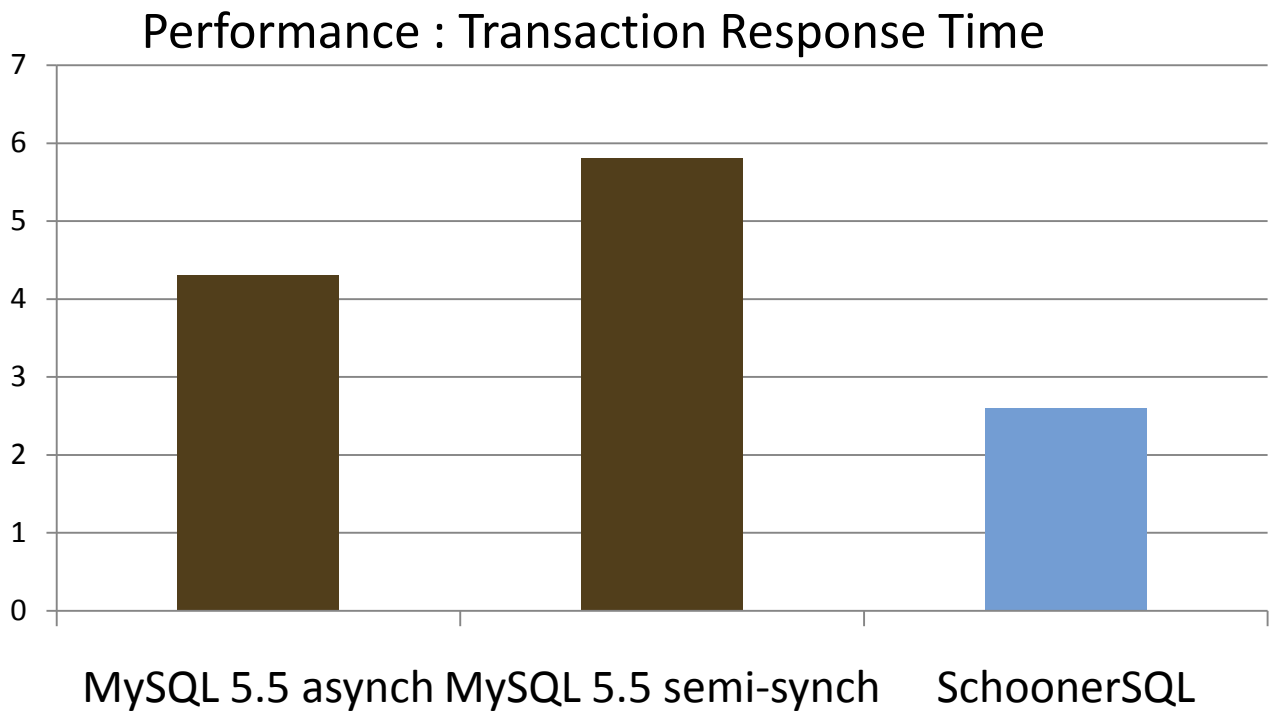

SchoonerSQL Provides High Performance Stability

MySQL 5.5 Asynchronous

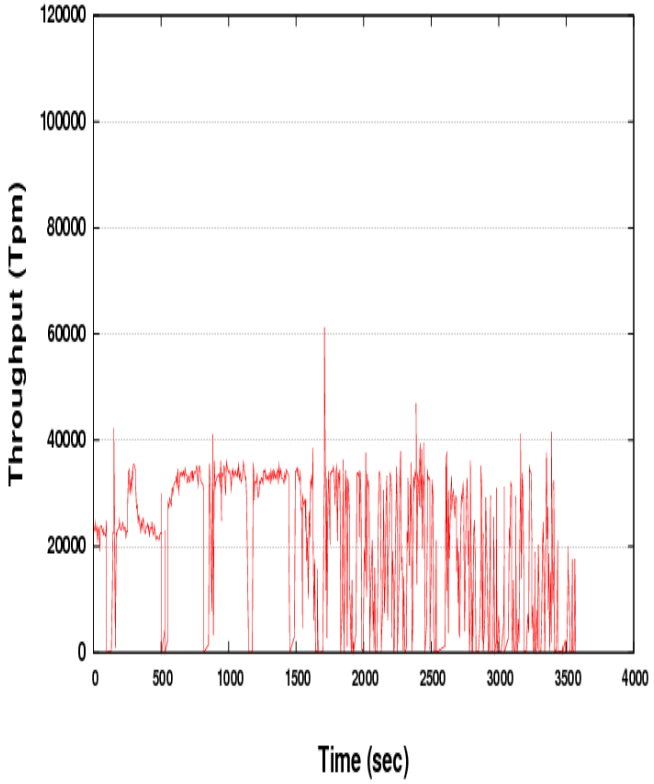

SchoonerSQL

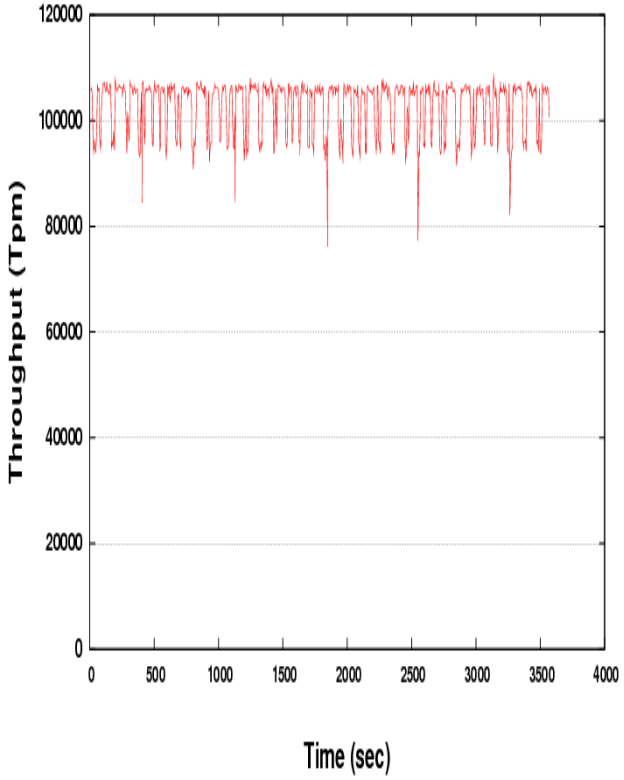

SchoonerSQL Lowers Total Cost of Ownership

Lower Cost

- Reduced capital and operating costs through reduction in servers, power, space, admin
- Savings from increased service availability and associated revenue and customer retention

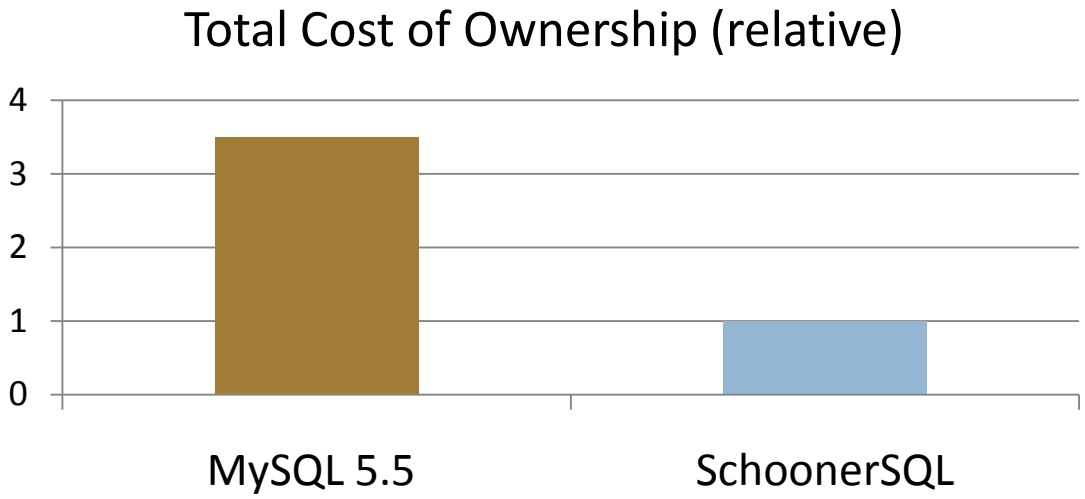

- TCO and ROI models are customer and workload specific
- Function (throughput/server; server, rack, and network costs, software license and support costs, admin costs; space and power costs; cost of downtime)

SchoonerSQL Greatly Simplifies Administration

- **Fail-over can be completely automatic and instant**
  - requiring no administrator intervention or service interruption
- **Cluster Administrator GUI and CLI can provide a single point for cluster-wide management**
  - single click slave creation and database migration

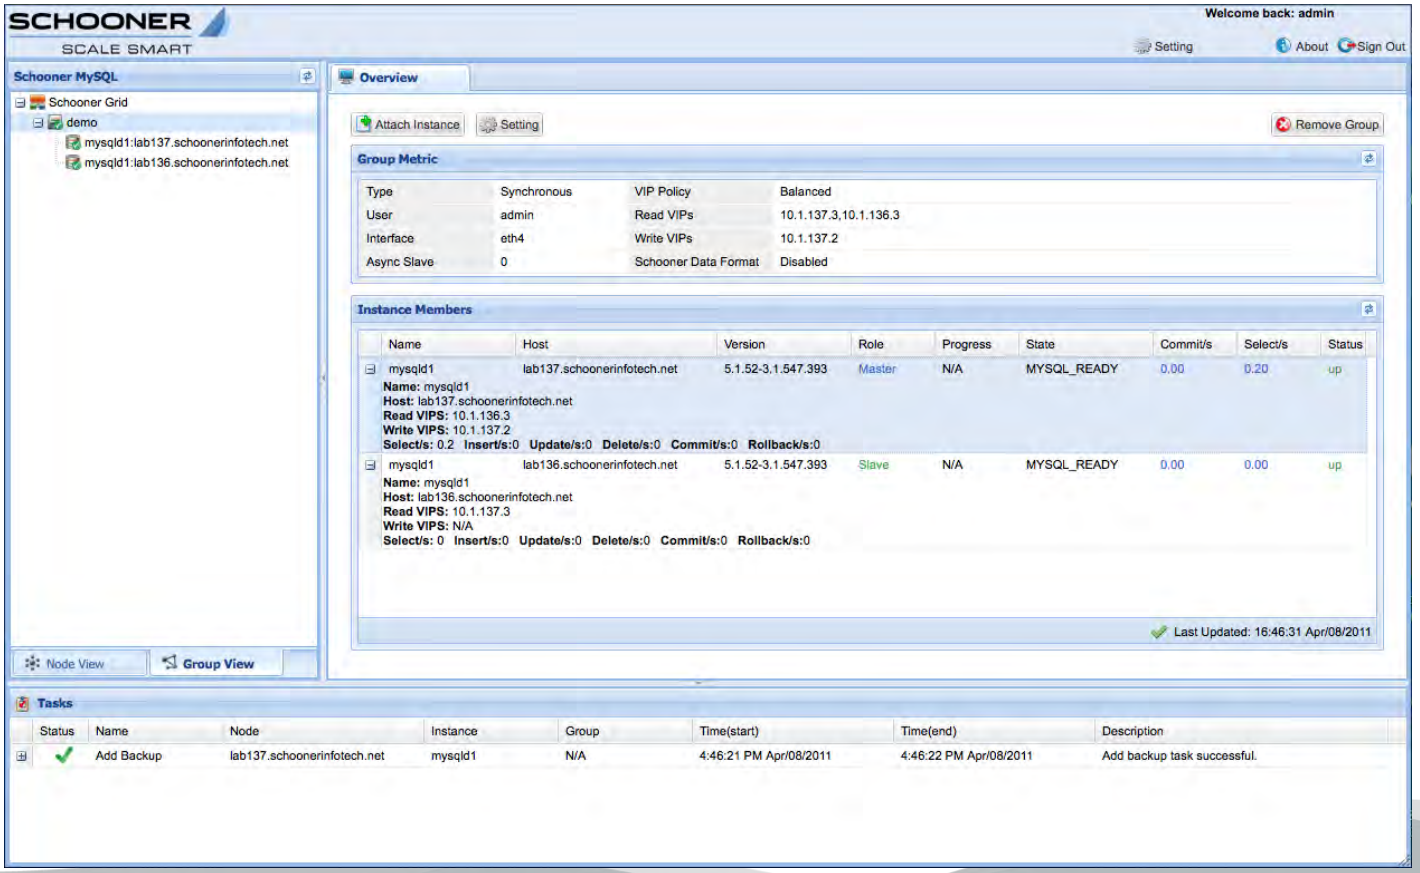

SchoonerSQL :Powerful Administration : Monitoring, Trouble-shooting, Tuning, Alerts

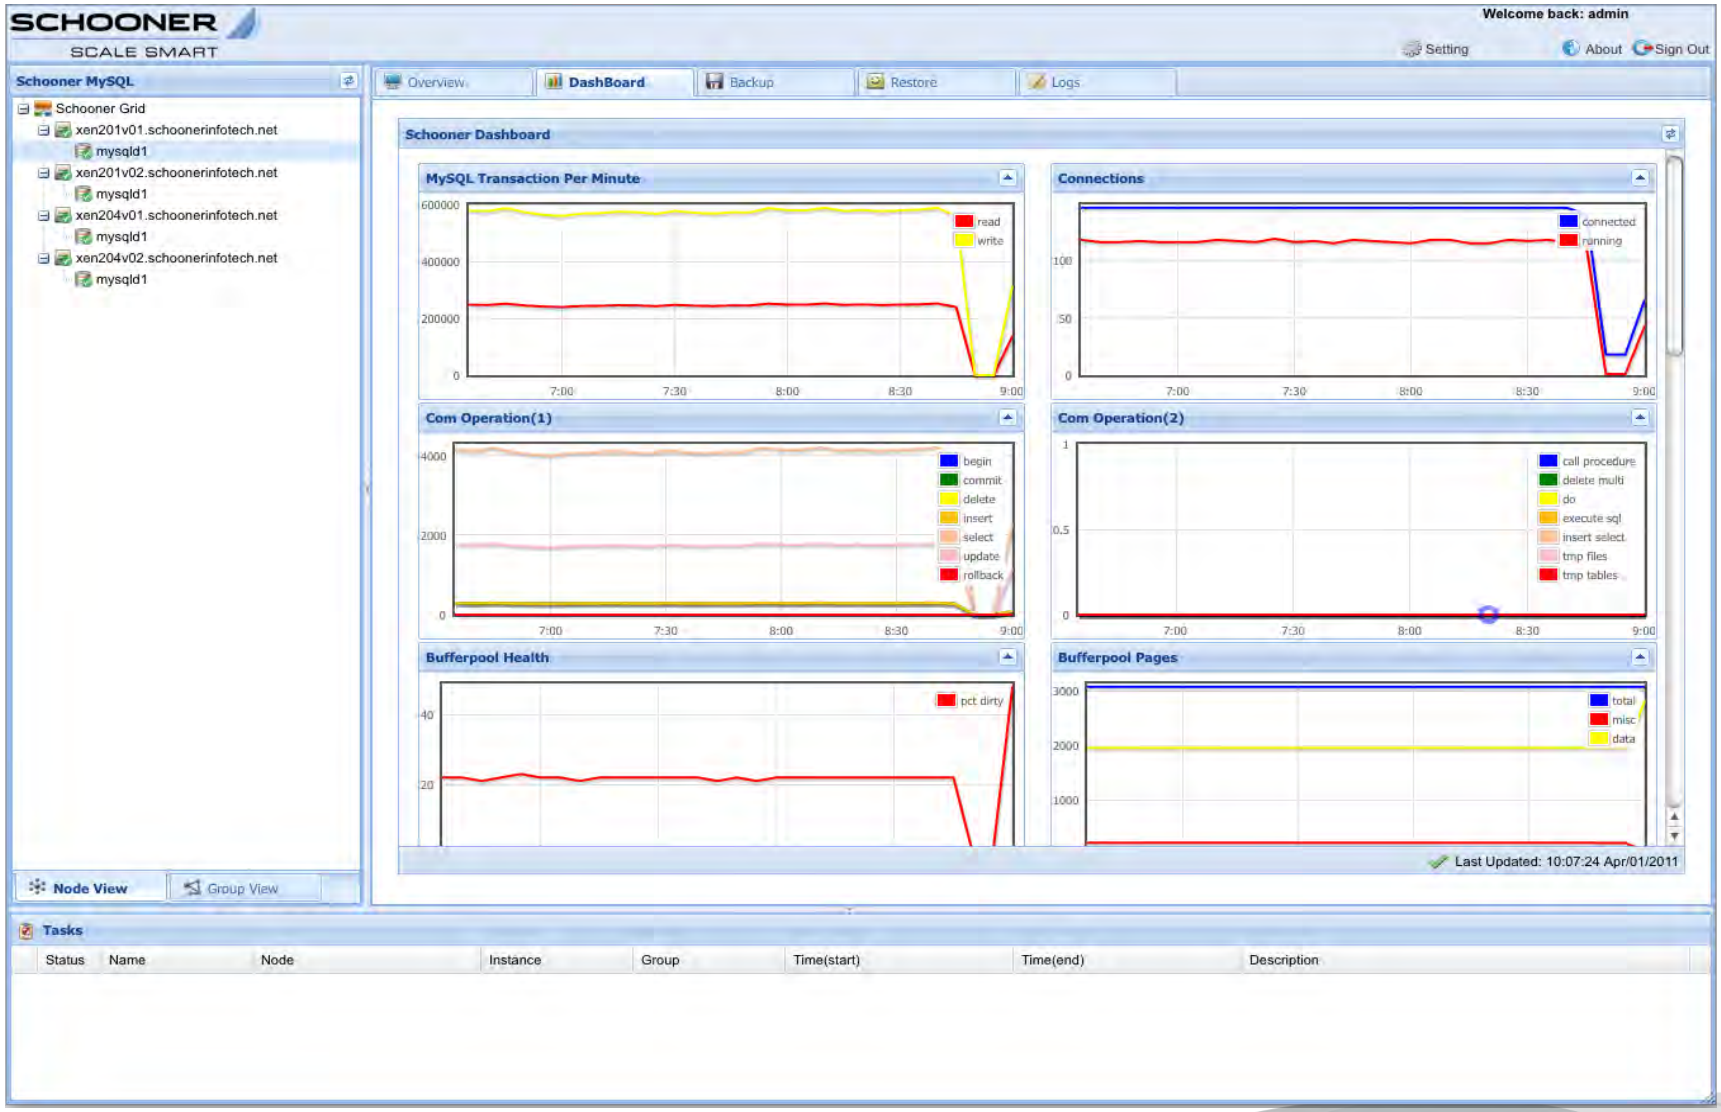

SchoonerSQL Extends Synchronous Replication with Parallel Asynch and Auto-Failover in WAN Geographic Distribution/DR

- WAN/geographically dispersed data centers
  - Requires Asynchronous replication
    - Can't add additional ~100ms with high potential variance to query response time for synchronous replication
- Data Integrity Requirement : Remote consistency lag and recovery time should be ~ WAN latency
    - Maximize WAN data consistency
    - Minimize disaster recovery time
  - Requires high performance asynchronous replication
    - Must have multi-threaded asynchronous parallelizing updates

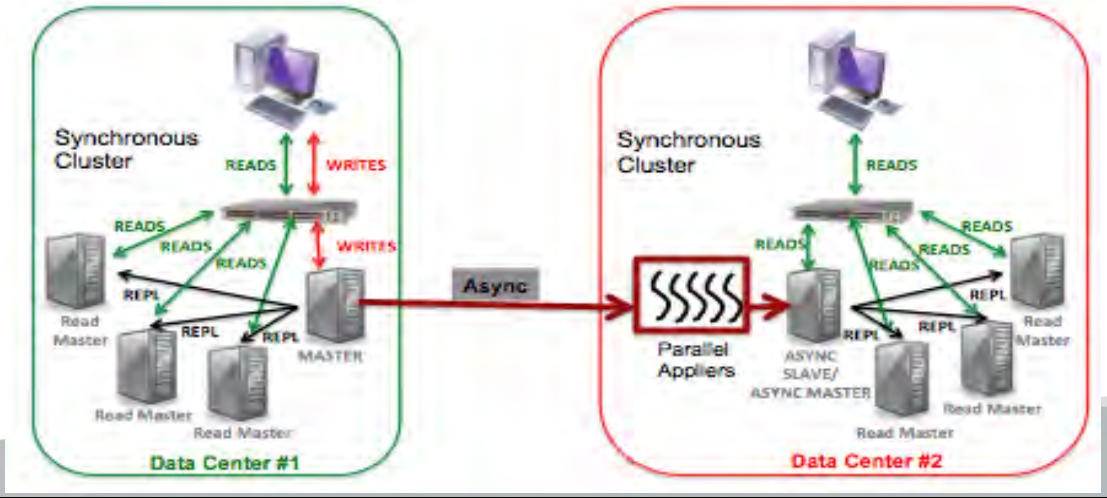

Legacy MySQL : Must Process Bin Log in a Single Thread -> Slave lag, low master and slave performance

Master

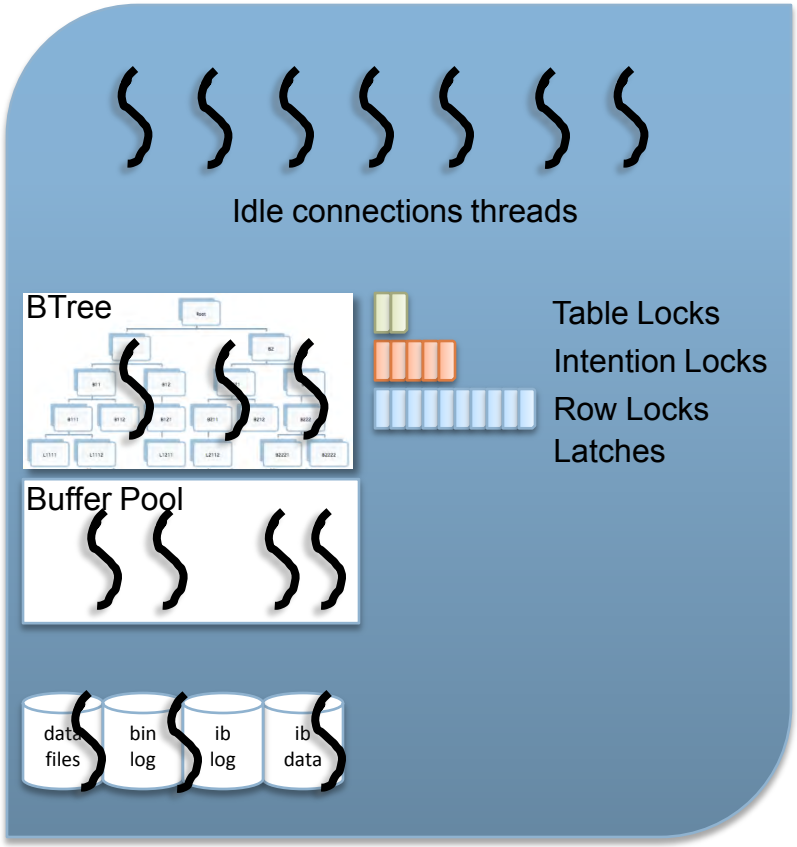

Slave – single thread

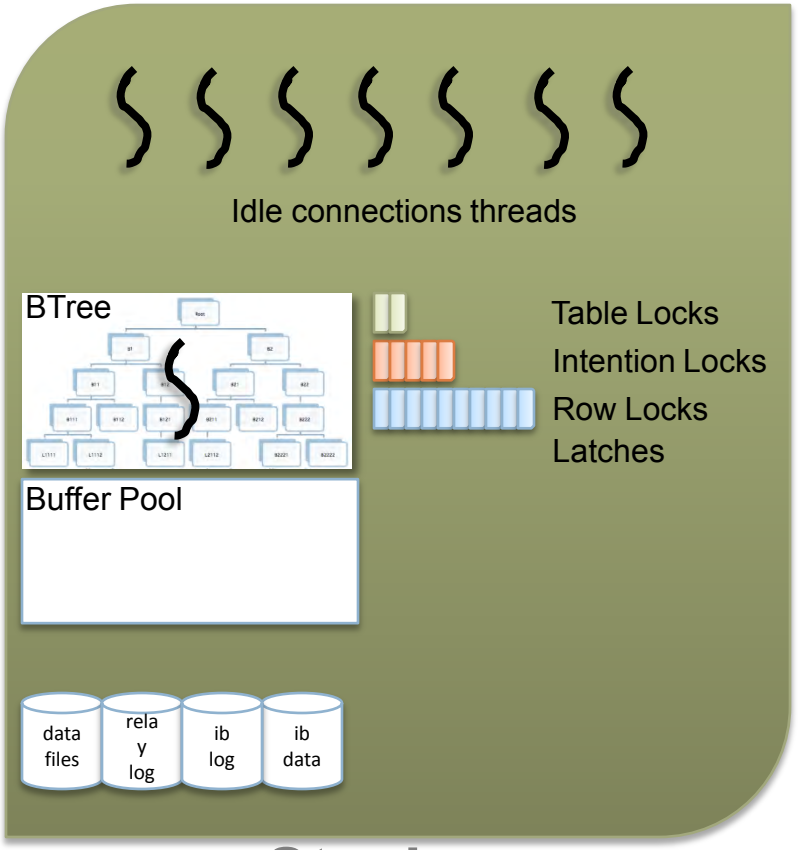

Stock MySQL

SchoonerSQL Parallel Asynchronous Replication

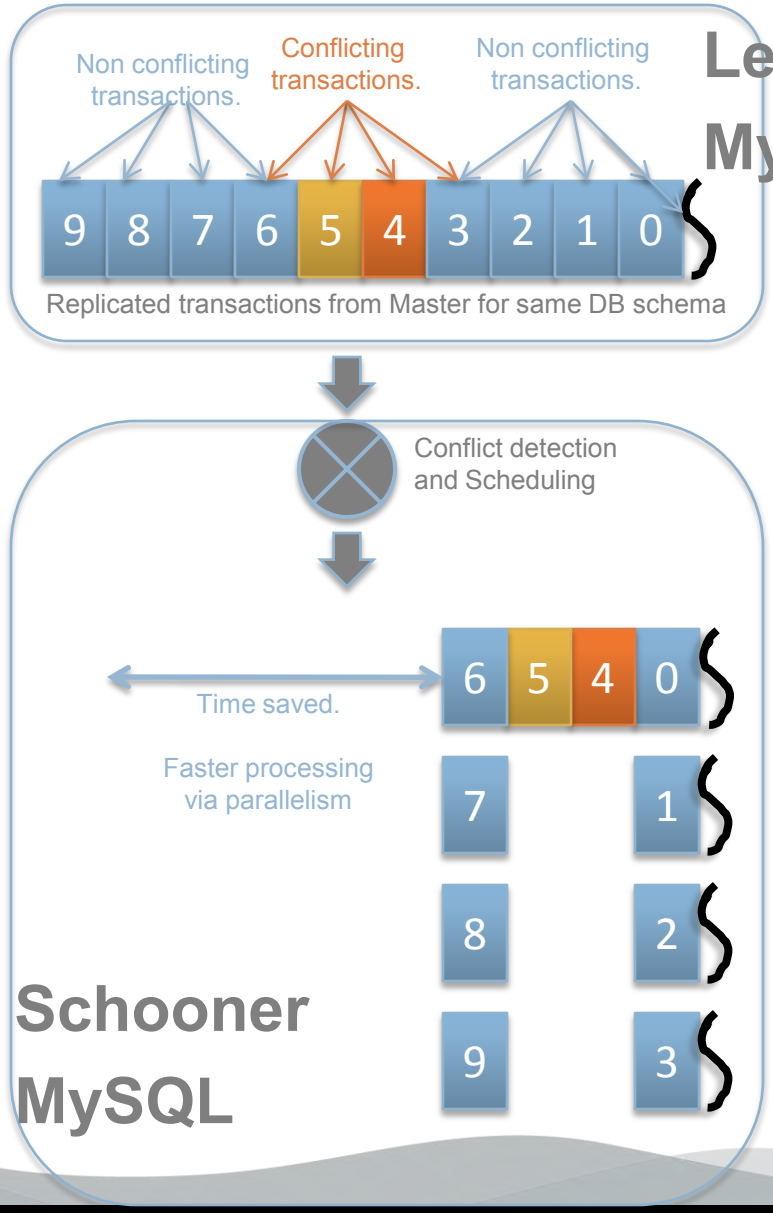

Schooner MySQL Slave (Multi-threaded)

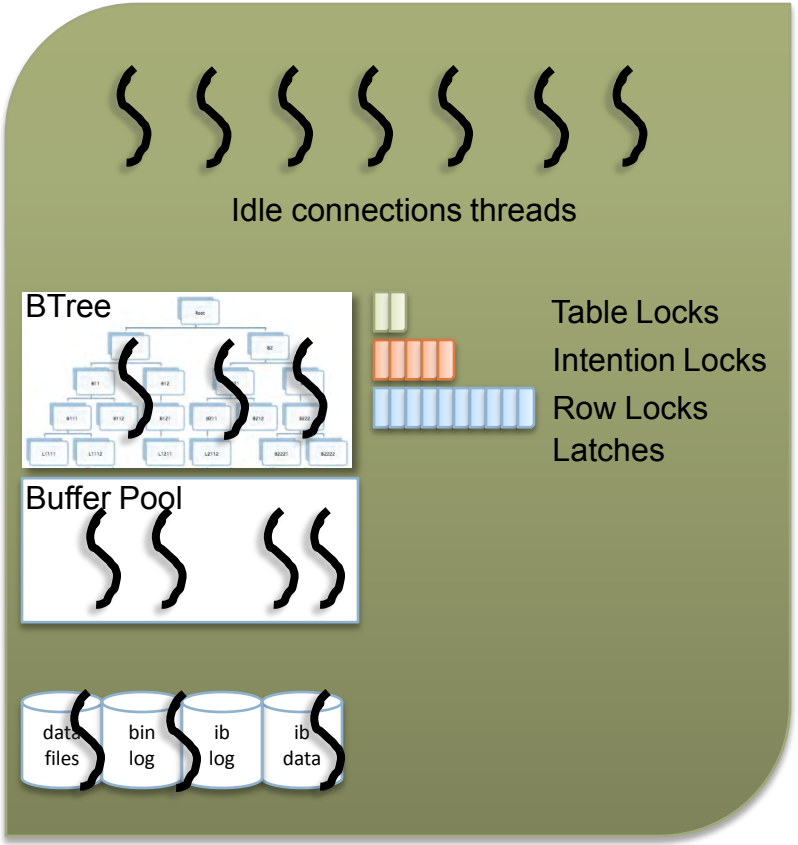

Reduces Slave lag by 95% with DBT2

SCHOONER SCALE SMART

SchoonerSQL Extends Synchronous Replication with Parallel Asynch and Auto-Failover in WAN Geographic Distribution/DR

- HA Requirement: Must automatically fail-over when synchronous master fails-over or asynch slave fail-over occurs, requires WAN asynchronous replication tight coupling with synch replication group

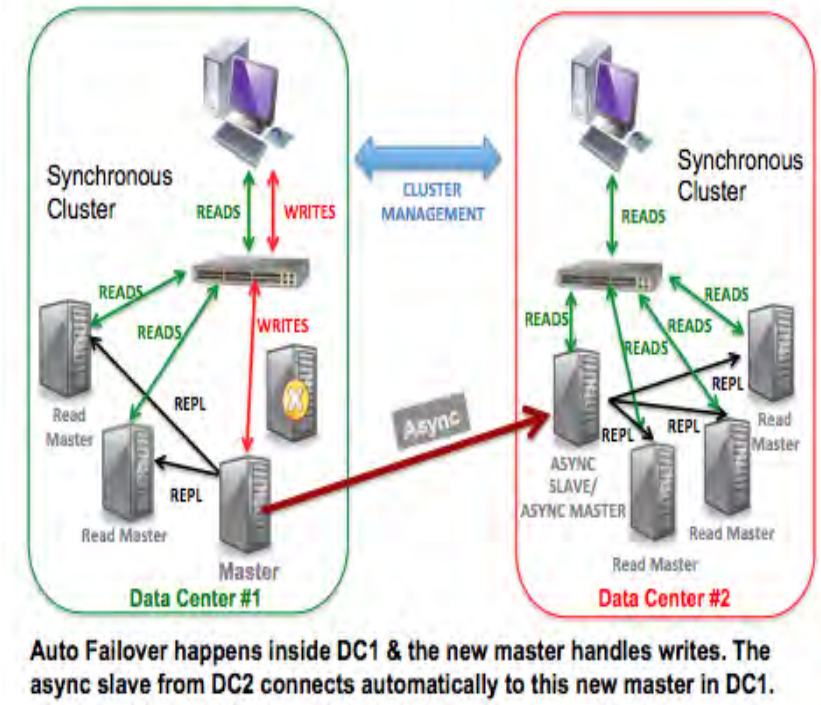

Auto Failover happens inside DC1 & the new master handles writes. The async slave from DC2 connects automatically to this new master in DC1.

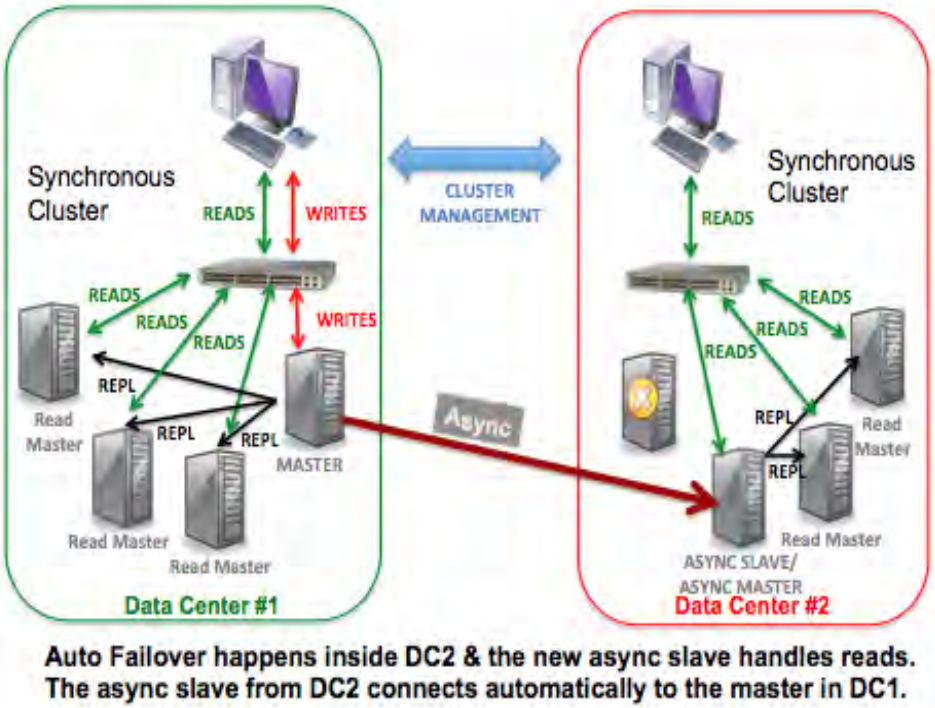

Auto Failover happens inside DC2 & the new async slave handles reads. The async slave from DC2 connects automatically to the master in DC1.

SchoonerSQL
Fully Synchronous within Data Center

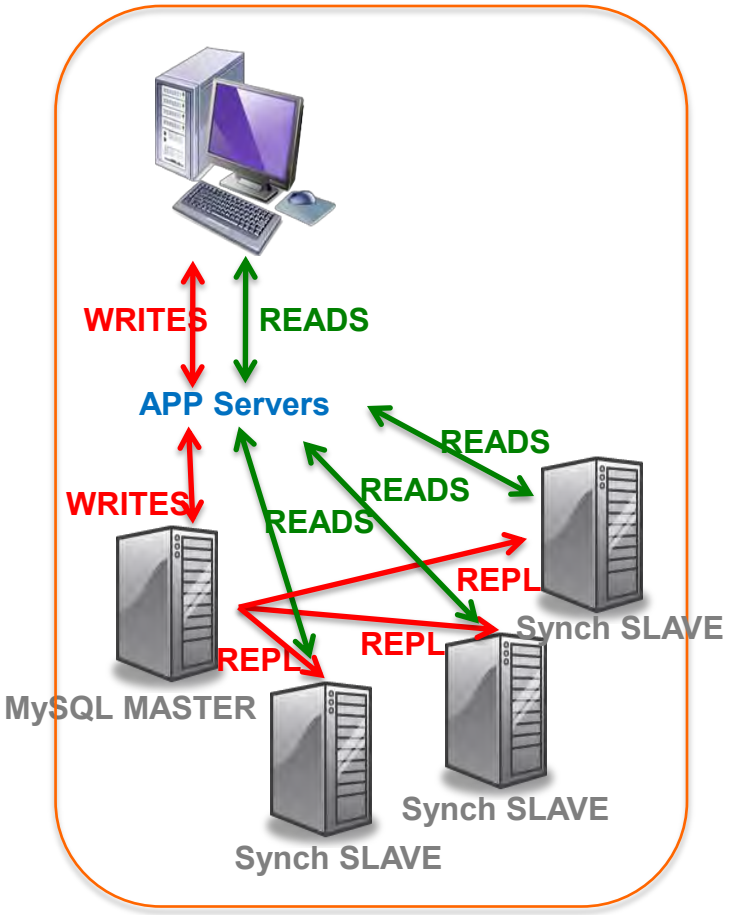

No Data Loss
Read Consistency
Auto Fail-Over
High Performance

Availability: up to 8 nodes
Scalability: up to 8 nodes

Schooner SQL with Parallel Synchronous Replication: Unlimited Scaling of LAN/MAN/WAN Clusters with Automated Failover

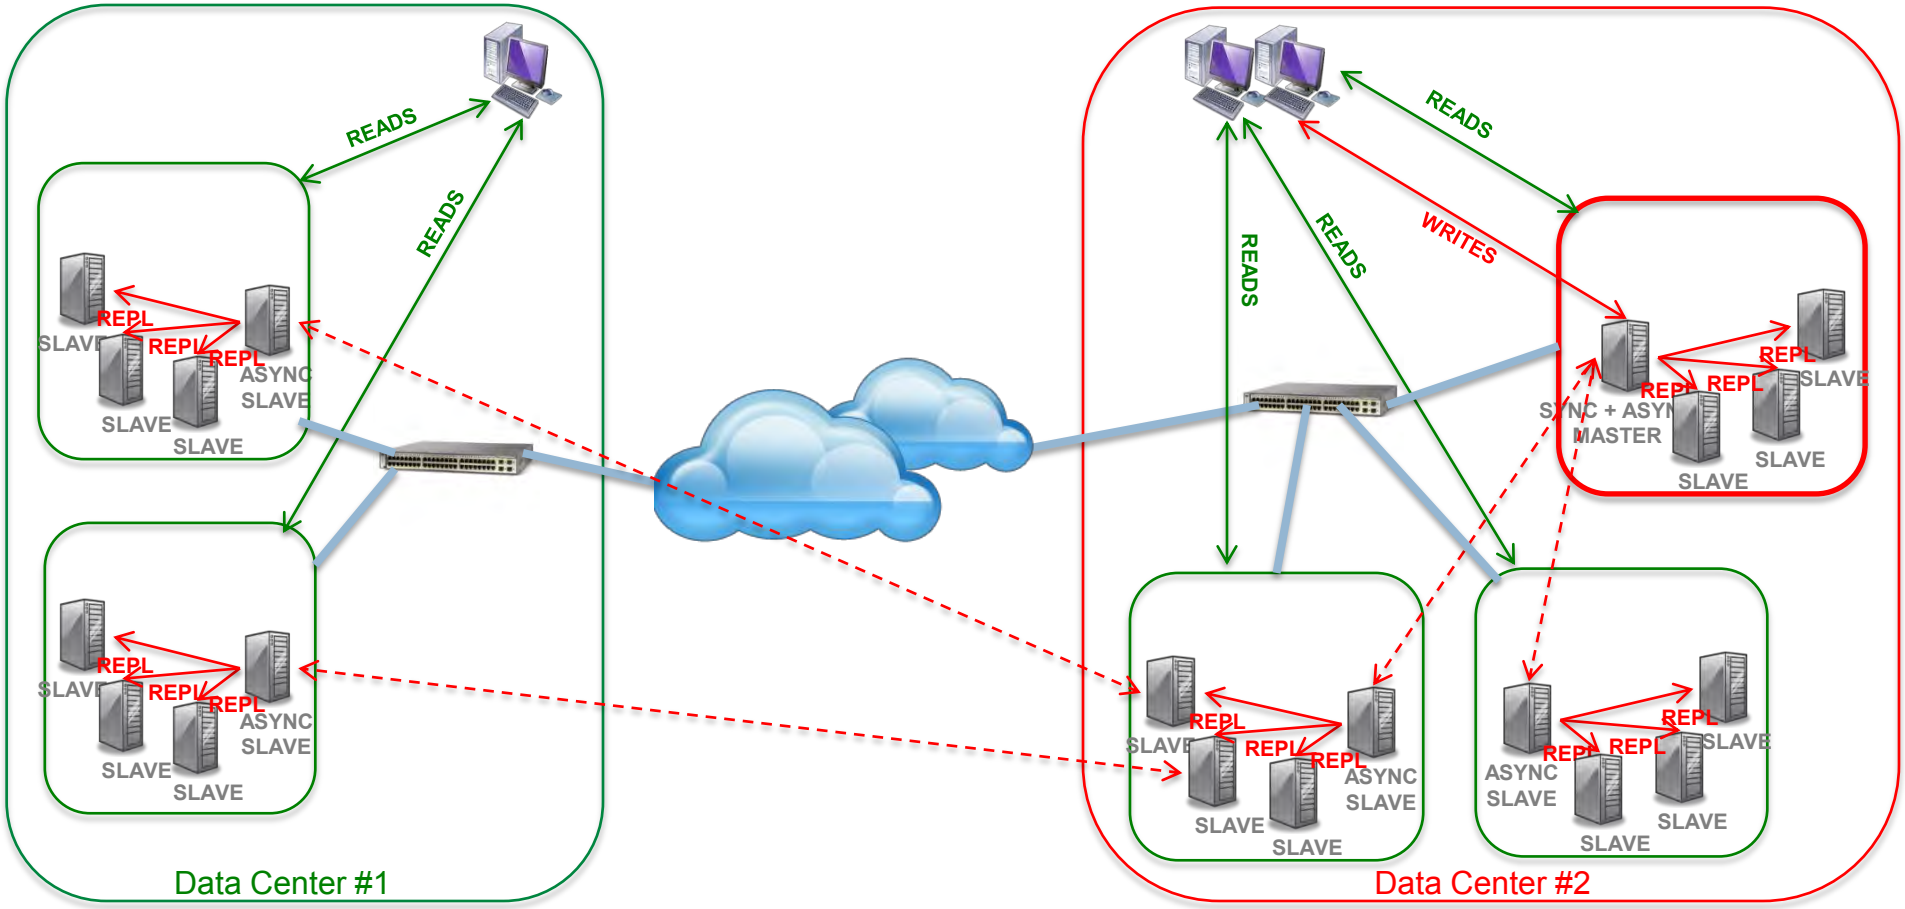

Parallel Slave Appliers with Sync and Async replication ensure a geographically distributed MySQL database cluster runs at high throughput (no slave lag), the same as within a data-center, and with high availability (auto-fail-over)

SchoonerSQL Provides Unlimited Query Scaling

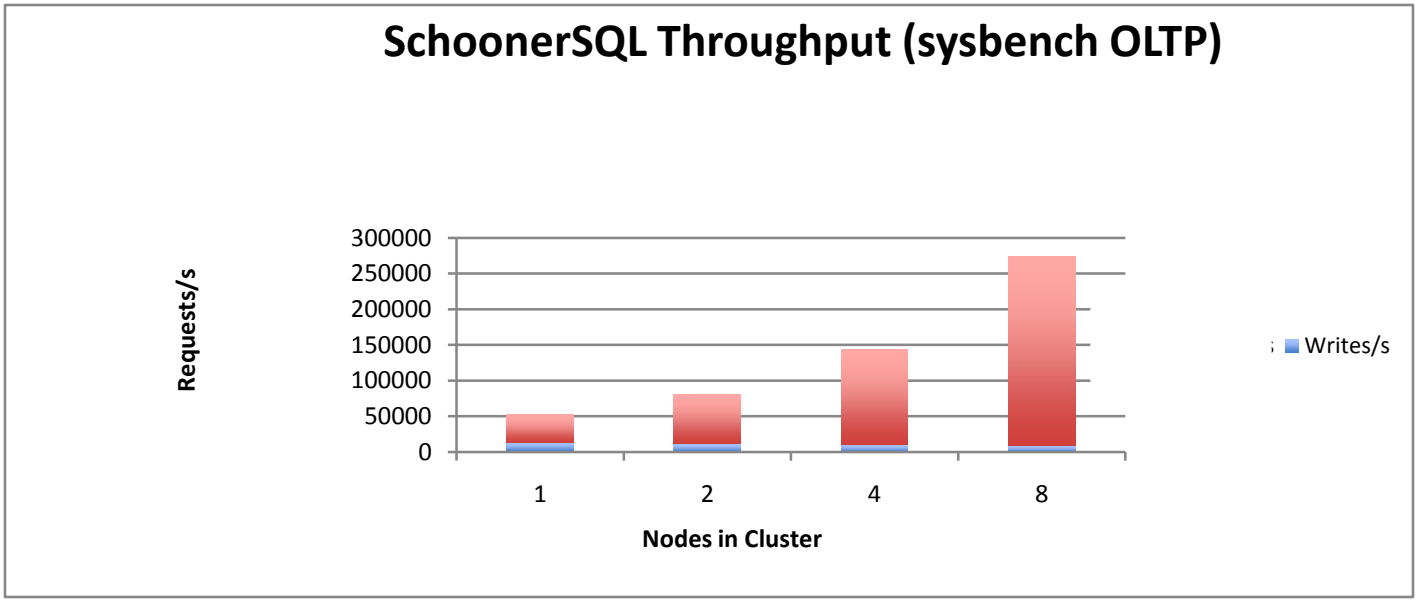

Scaling Updates

- Database Update Scalability
  - SchoonerSQL Vertically scale with commodity : flash memory, more cores, higher frequency

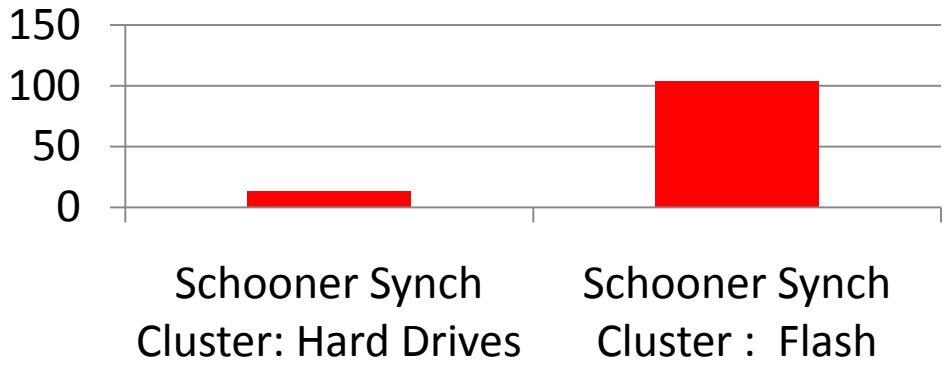

  - Compelling option exploiting low cost, high performance commodity technology

SchoonerSQL ProvidesUnlimited Update Scaling Through Optimized Transparent Multi-Node Sharding

- Database Update Scalability

…After Optimal Vertical Scaling:

Unlimited Query Scaling : Horizontally Scale Through Transparent Workload Aware Database Partitioning and Replication

- Database workload aware
  - Administrator analysis and configuration tools
  - Allows optimal layout, including replicated structures, for query data access optimization
- Application Transparent
  - High performance, highly available dynamic query execution across shards

SchoonerSQL Provides Industry Leading Availability, Data Integrity, Performance, Scalability, Ease of Management, TC

FEATURES & BENEFITS	MYSQL 5.5	DRBD	ScaleDB	MYSQL NDB CLUSTER	CONTINUENT (TUNGSTEN)	CLUSTRIX	SCHOONER SQL
Synchronous Replication for InnoDB (Guaranteed Data Consistency)	No	Limited	No	No	No	No	Yes
# Node Failures before Service Downtime (Failure Resistance)	Two	Two	Three	Four	Two	Two	Eight
Eliminates Slave Lag (100% Data Consistency and Zero Data Loss)	No	No	N/A	N/A	No	N/A	Yes
Automated Fail-Over (LAN/MAN/WAN)	No	No	No	No	No	No	Yes
Performance Across WAN	Low	Low	Low	Low	Low	Low	High
Full & Incremental Online Backup Integrated with GUI (Zero Downtime)	Limited	No	No	No	No	No	Yes
Online Software & Hardware Upgrades (Zero Downtime)	No	No	No	Low	No	Low	High
Elastic Cluster (add or remove nodes with ease - Zero Downtime)	No	No	Medium	Medium	Low	Medium	High
Performance with Flash Memory	Low	Low	Low	Low	Low	Medium	High
Cost (TCO)	Medium	High	High	High	High	High	Low

Schooner Membrain NoSQL : Flash/Multi-Core/Network Optimized = Breakaway Performance, TCO, and Availability

Highly optimized multi-core, DRAM and flash management with Synchronous Replication and Transparent Failover

- 1TB Fusion-io based key value store per server at DRAM-like speeds

Both a pure transient cache and persistent key-value data store

Standards based : memcapable

Schooner Membrain vs. CouchBase Key-Value Store Performance on 2 socket Westmere server with 1/2TB of Fusion-io avg 2kB objects, (90% get: 10% put)

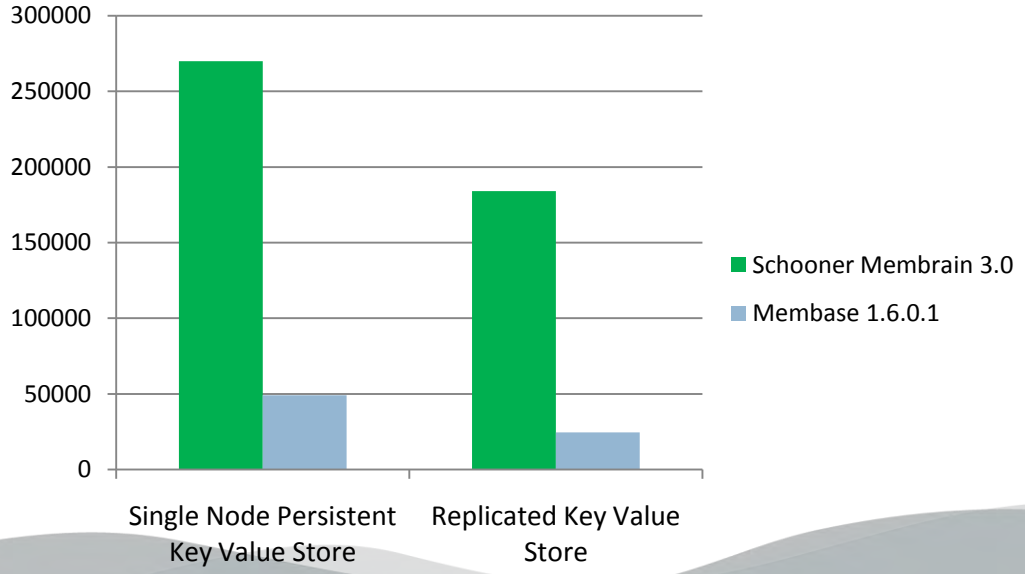

Schooner Membrain NoSQL : Flash/Multi-Core/Network Optimized = Breakaway Performance, TCO, and Availability

- 8:1 server consolidation + 10:1 power reduction
- 95% Availability Improvement
  - With transparent synchronous replication and automated failover

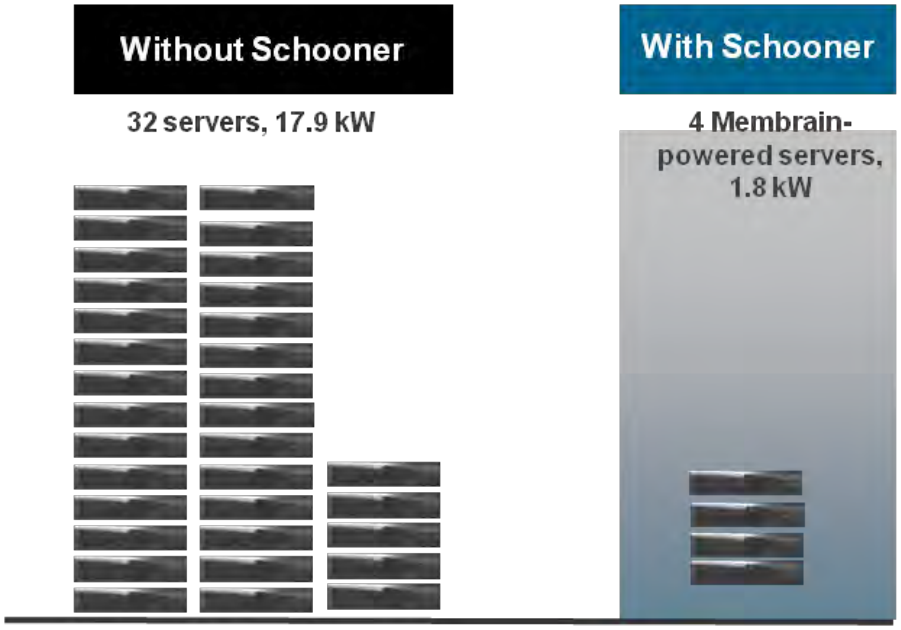

Schooner Membrain's Consolidation and Power Savings

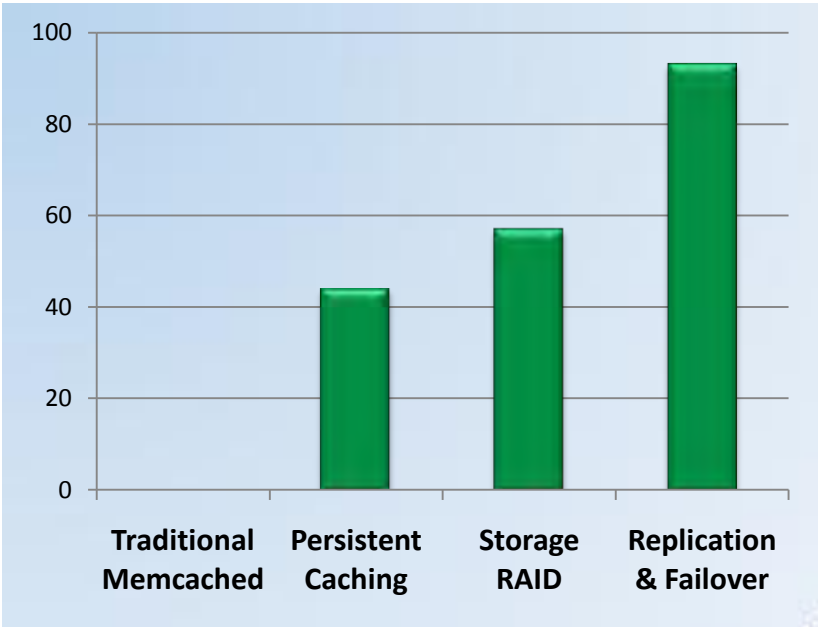

Cumulative % reduction in downtime

SCHOONER SCALE SMART

Mission Critical SQL and NoSQL : Flash/Multi-Core/Network Optimized

TPS/Node, Random Queries per node	In DRAM	In Flash
Cassandra	10,500	1,790
MongoDB	49,000	4,000
SchoonerSQL	115,000	101,000
Schooner Membrain	310,000	160,000

NoSQL benchmark is a key-value random query of 32M and 64M 1kByte items, on the same dual quad-core Intel Nehalem processors with 64 GB of DRAM and ½ TB of flash

Schooner Advantage

Highest Availability

- No service interruption for planned or unplanned database downtime
- Instant automatic fail-over
- On-line upgrade and migration
- 90% less downtime
- WAN/DR auto-failover

Highest Performance and Scalability

- 4-20x more throughput/server
- Highest performance synchronous and asynch replication clusters

Compelling Economics

- TCO 70% cheaper

Broad Industry Deployment

- eCommerce, Social Media, Telco, Financial Services, Education
- High volume web sites
- Geographically distributed websites

Highest Data Integrity

- No lost data
- Cluster-wide data consistency

Easiest Administration

- No error-prone manual processes
- Monitoring and Optimization

Out-of-the-box Product

- Full Products: not toolkits
- Free your staff to build your business

Standards Based 100% Compatible

SCHOONER SCALE SMART

About Schooner

- Founded in May 2007
- Headquarters in Sunnyvale, CA
- Deep and unique IP (19 patents filed)
- Exceptional team
- Products in production use from mid-2009
- Products Launched in China November 2011
  - HengTien Insigma Providing China-Wide Support
- Provides high availability, high performance database software products for demanding mission-critical use:
  - SchoonerSQL™: 99.999% availability OLTP database, 100% compatible with the widely-used MySQL and InnoDB
  - Membrain™: high-performance flash-optimized software cache

SCHOONER
SCALE SMART

Take Aways

Apps + Middleware on Legacy SQL, Legacy NoSQL
= Inefficient, High TCO, limited performance and availability

SchoonerSQL and Membrain = Highly Cost Effective, High Availability, High Performance Building Blocks

- 10x Improvement in Performance, Consolidation, Up-Time
- Dramatically reduces cost (servers, power, space, complexity)
- Eliminate Data Loss, Automate Fail-over, Ease Admin
  - Synchronous Replication
- Unlimited Scaling and Geo-Redundancy
  - Parallel Asynchronous Replication, Transparent Sharding
- Based on standards reduces risk, preserves application investment
- Installs in minutes on your existing servers and storage

Evaluating the Options and Trade-offs for Your Data Center? Let Schooner Help!

Try SchoonerSQL! We guarantee SchoonerSQL will increase your Availability, Performance, and Scalability

Contact us:

501 Macara Avenue, Suite 101
Sunnyvale, CA 94085 USA
Tel: +1 408-773-7500
www.schoonerinfotech.com
Email: info@schoonerinfotech.com

Schooner中国
地址：杭州市西湖区教工路23号百脑汇大厦18楼
传真：057189731509　电话：057189731653
销售电话：13867476875
Email: salescn@schoonerinfotech.com

SCHOONER SCALE SMART

Thank You!