

Maximum Availability for MySQL/InnoDB with Synchronous Replication, Automated Failover, Full Data Consistency, Simplified Management, and Industry-Leading Performance

Dr. John R. Busch

Schooner Information Technology

Founder and Chief Technology Officer

<http://www.schoonerinfotech.com/>

John.Busch@SchoonerInfoTech.com

MySQL Conference

Tuesday, April 12, 2011, 2:00PM, Ballroom H

Maximum Availability for MySQL/InnoDB with Synchronous Replication, Automated Failover, Full Data Consistency, Simplified Management, and Industry-Leading Performance

Abstract

This session discusses innovative technology developed to achieve both industry-leading availability AND industry-leading performance for MySQL/InnoDB. Enabled by fully-synchronous replication which is tightly integrated into MySQL and InnoDB, downtime is reduced by over 90% through: automatic fail-over of masters and slaves transparent to client applications; full data consistency and durability across the cluster with zero slave lag; on-line hardware and software upgrades; and consistent on-line back-ups. This also provides fully-certified MySQL/InnoDB compatibility and greatly simplified data center management. Coupled with maximum MySQL/InnoDB availability and ease of administration, this new technology provides the industry's highest MySQL/InnoDB performance through highly-optimized multi-core thread, lock concurrency, DRAM, and flash management algorithms. This is now available for software download and trial on a broad range of commodity X86 servers. In this session, we present the technology, demonstrate the product simplicity in operation, and present benchmark results of its industry-leading availability and performance.

Dr John R Busch

John is the founder, Chairman and CTO of Schooner Information Technology, Inc. focusing on Optimal Scale-Up and Scale-Out for MySQL. Prior to Schooner, John was director of computer system architecture at Sun Microsystems Laboratories from 1999 through 2006. In this role, John led research in multi-core processors, multi-tier scale-out architectures, and advanced high-performance computer systems. Prior to Sun, John was VP Engineering and Business Partnerships with Diba, Inc, co-founder and VP Engineering of Clarity Software, and, from 1976 to 1993, John led many successful R&D programs at Hewlett Packard in Computer Systems Research and Development. John earned a Ph.D. in computer systems architecture from UCLA, an M.S. in mathematics from UCLA, an M.S. in computer science from Stanford University, and attended the Sloan Program at Stanford



Agenda

- Mission Critical MySQL Requirements
- Today's Challenges and Limitations
- New Architectural Approach
- Industry Leading
 - Availability
 - Performance
 - Administration
 - TCO/ROI
- Deployment Experiences

Mission Critical MySQL



High
Performance



Consistent
Data



Simple and
Powerful
Administration



Instant
Failover with
no data loss



Cost Effective

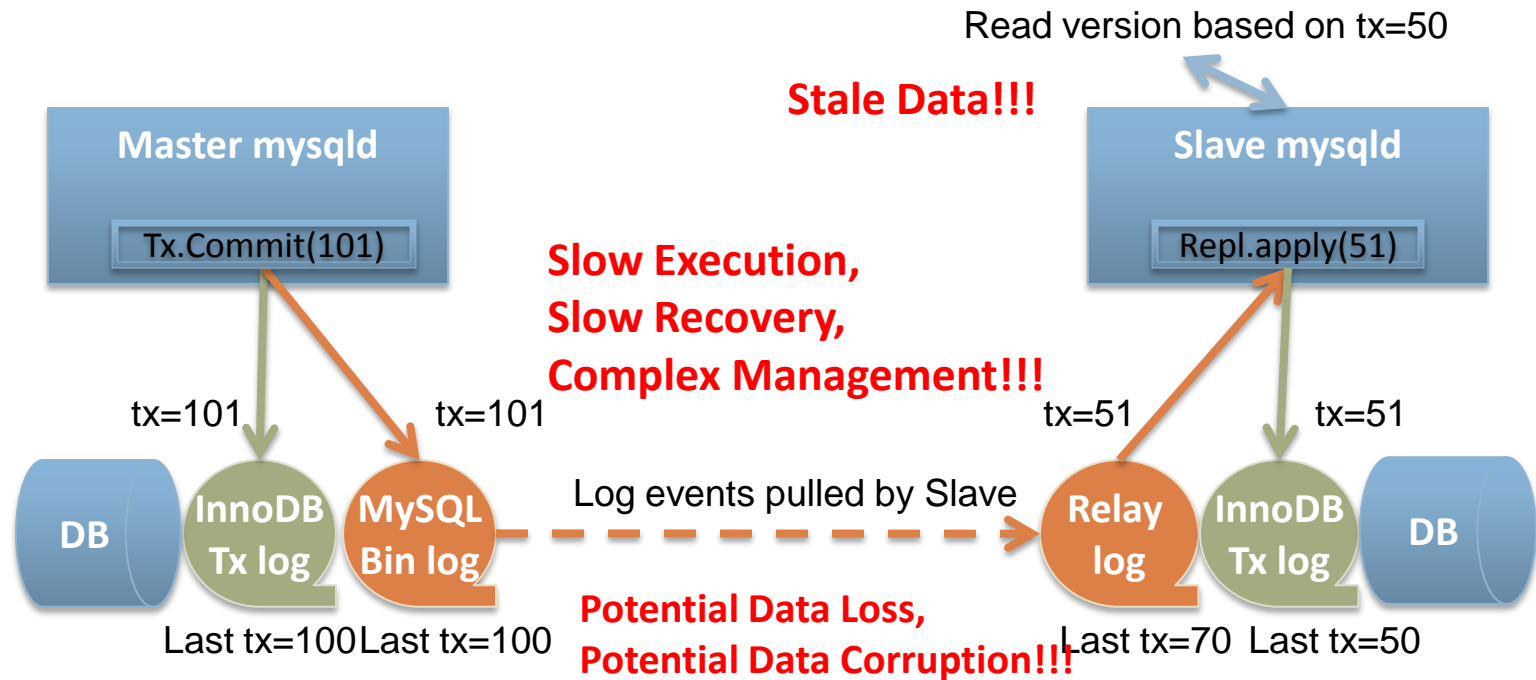


Truly
Compatible

Mission Critical

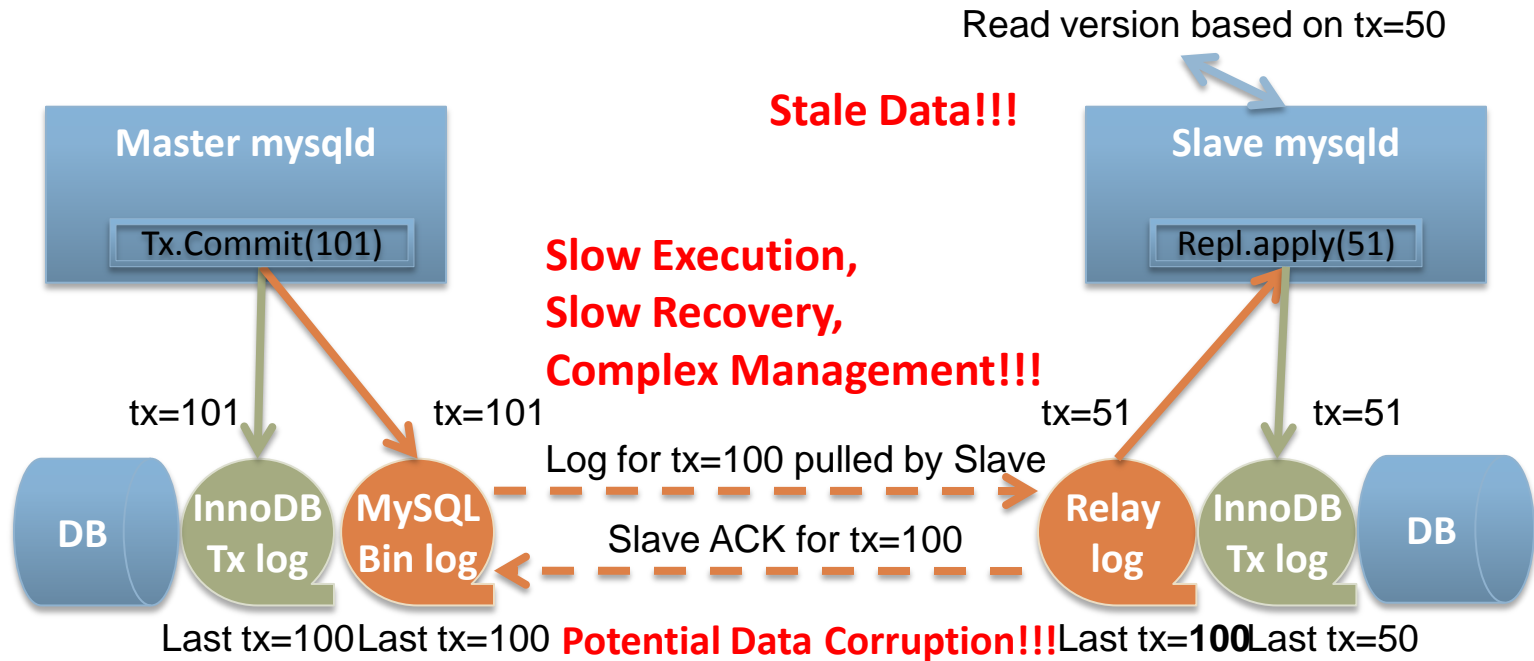
Some Current Challenges and Limitations

MySQL Asynchronous Replication



- Loosely coupled master/slave relationship
 - Master does not wait for Slave
 - Slave determines how much to read and from which point in the binary log
 - Slave can be arbitrarily behind master in reading and applying changes
- Read on slave can give old data
- No checksums in binary or relay log stored on disk, data corruption possible
- Upon a Master's failure
 - Slave may not have latest committed data resulting in data loss
 - Fail-over to a slave is stalled until all transactions in relay log have been committed – not instantaneous

MySQL Semi-synchronous Replication

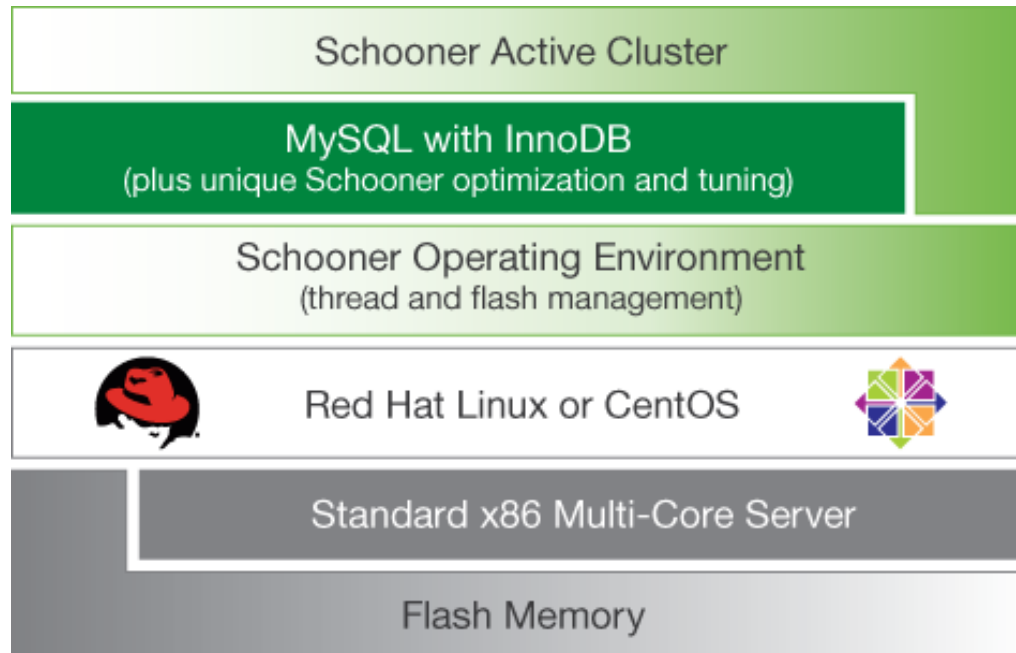


- Semi-coupled master/slave relationship
 - On commit, Master waits for an ACK from Slave
 - Slave logs the transaction event in relay log and ACKs (may not apply yet)
 - Slave can be arbitrarily behind master in applying changes
- Read on slave can give old data
- No checksums in binary or relay log stored on disk, data corruption possible
- Upon a Master's failure
 - Fail-over to a slave is stalled until all transactions in relay log have been committed – not instantaneous

A New Approach

Schooner MySQL Scale Smart Architecture

Schooner smart scaling is done with innovative software that fully leverages commodity processor, flash memory, and networking technologies for breakaway availability, performance, and capex and opex savings

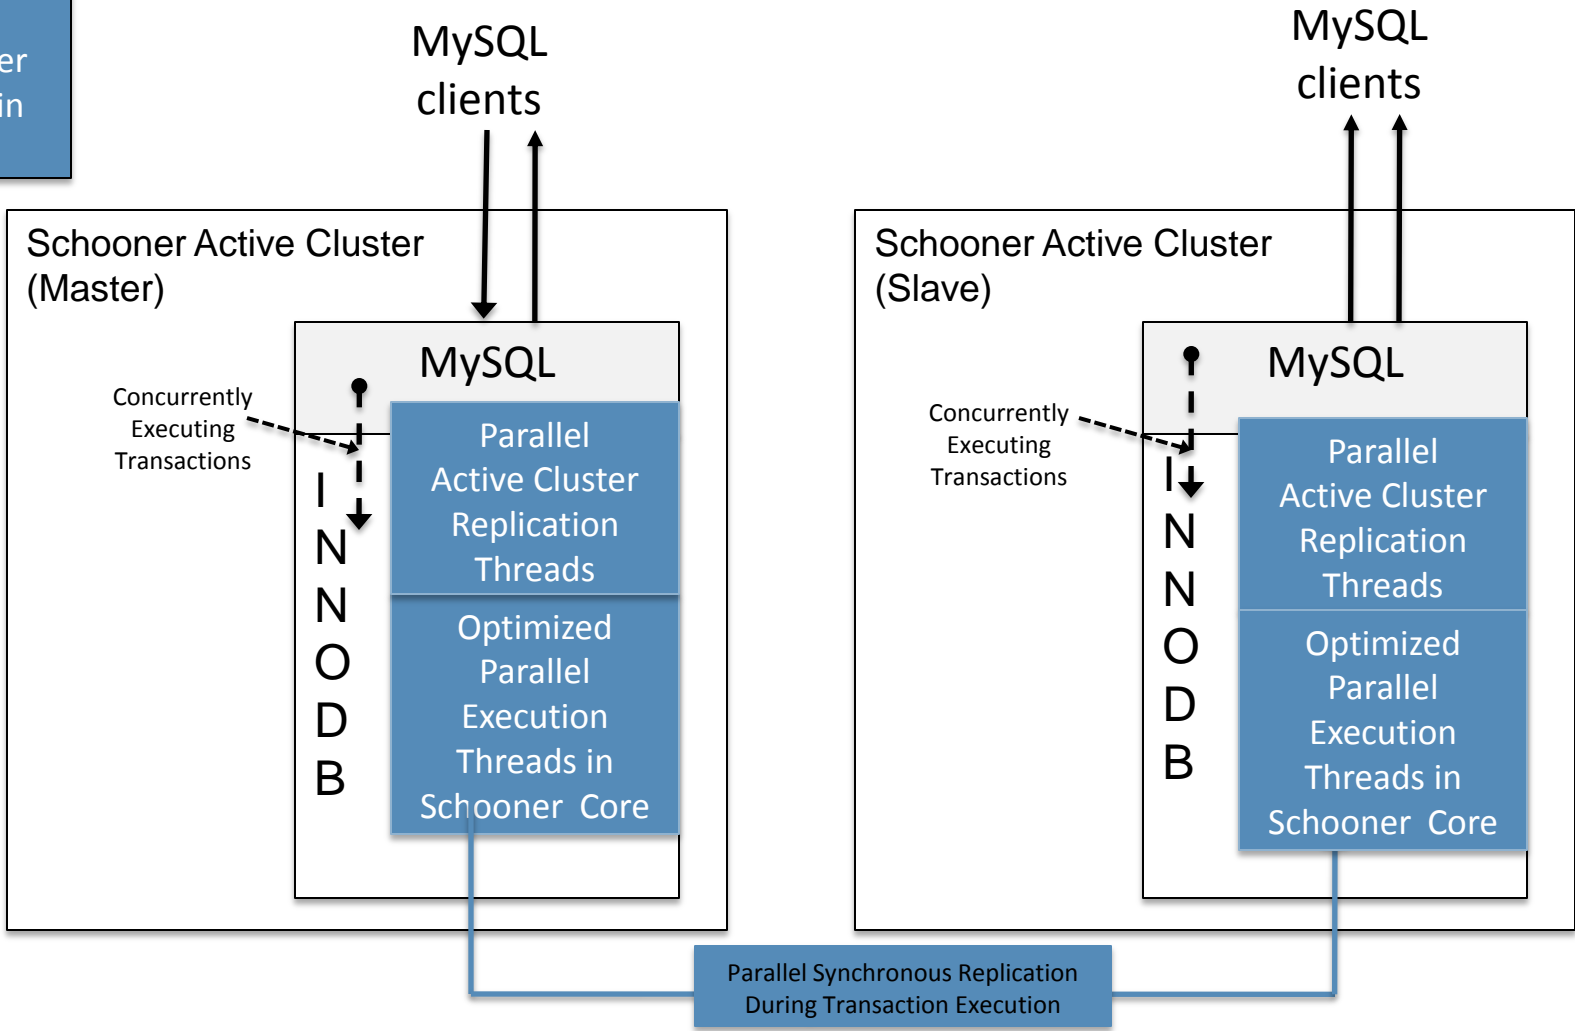


- 90% Less Downtime
- 4 - 20x Better Performance
 - Cut TCO by 60%
 - 100% Compatible

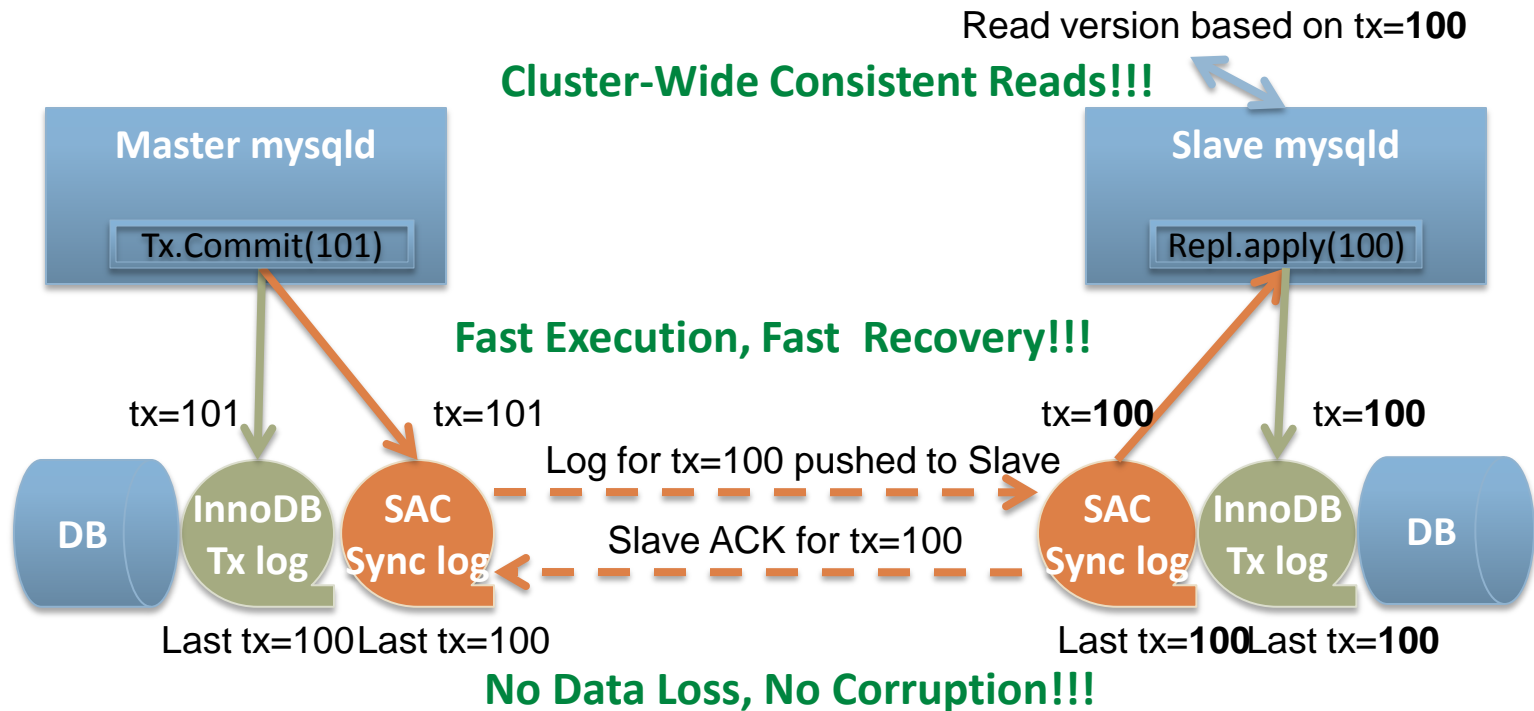
Schooner Active Cluster

Deeply Integrated Parallel Synchronous Replication

Cluster Admin



Schooner MySQL Active Cluster (SAC): An integrated HA and replication solution for MySQL/InnoDB



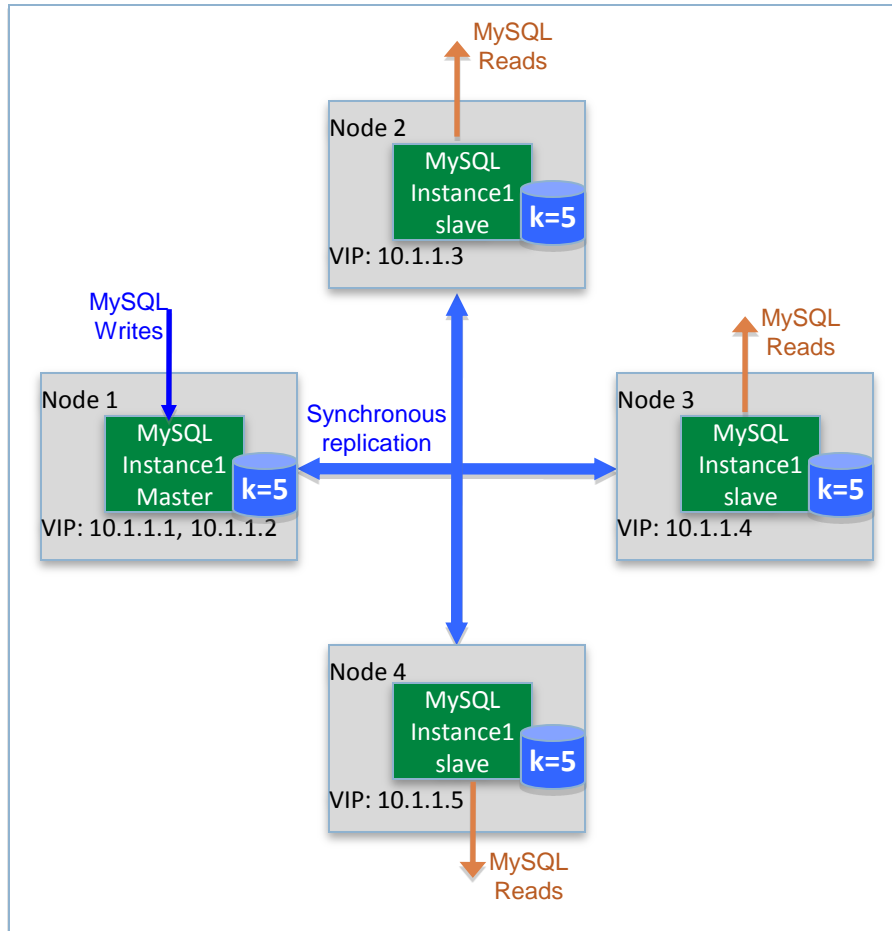
- Tightly-coupled master/slave relationship

- After commit, all Slaves guaranteed to receive and commit the change
- Slave in lock-step with Master

- Read on slave gives latest committed data
- Checksums in log stored on disk
- Upon a Master's failure
 - Fail-over to a slave is fully integrated and automatic
 - Application writes continue on new master instantaneously

Schooner Active Cluster

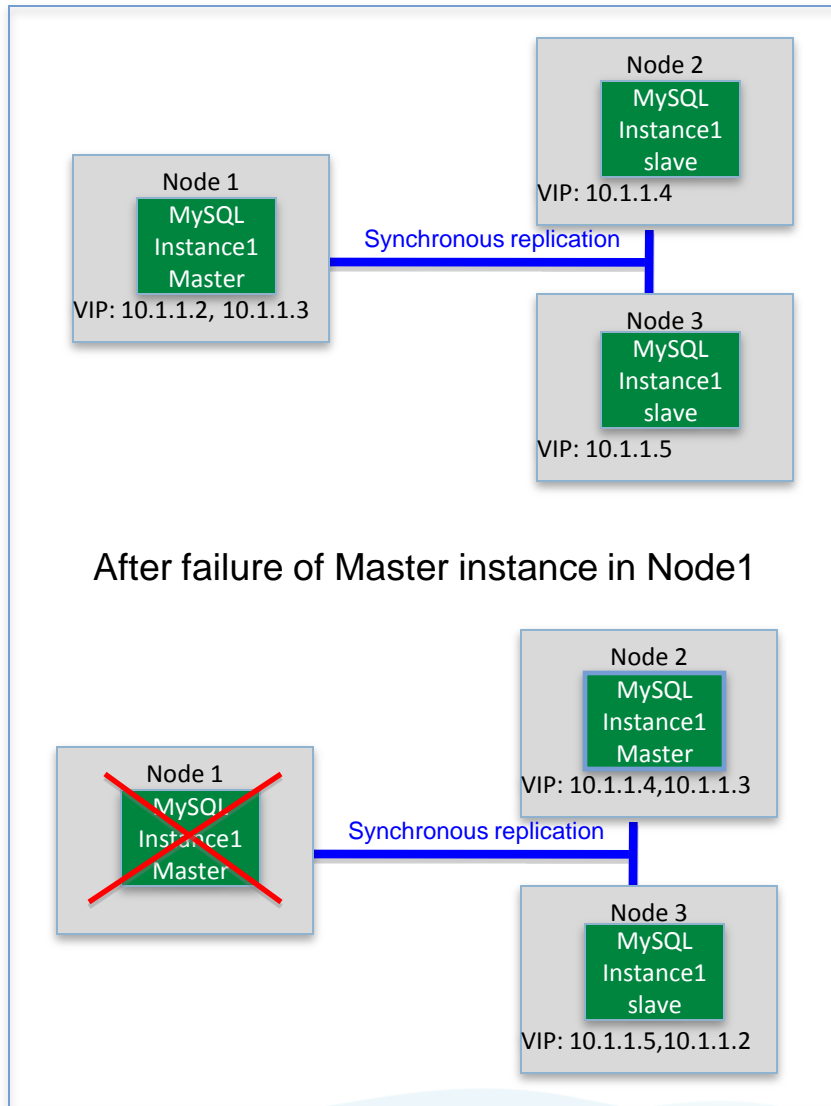
Fully Consistent Cluster-Wide Data



- Zero Slave lag
- Zero Data loss
- Apps do not have to read from Master to get consistent data
- Linear read scaling
- No data corruption

Schooner Active Cluster

Instant Failover with No Data Loss



- Instant Failure detection and automatic VIP migration
- Recovering Instance syncs with master and VIPs are re-balanced after the sync automatically
- Foundation of SAC on-line upgrades and on-line migration
- Simple user interface for failure management and recovery tracking

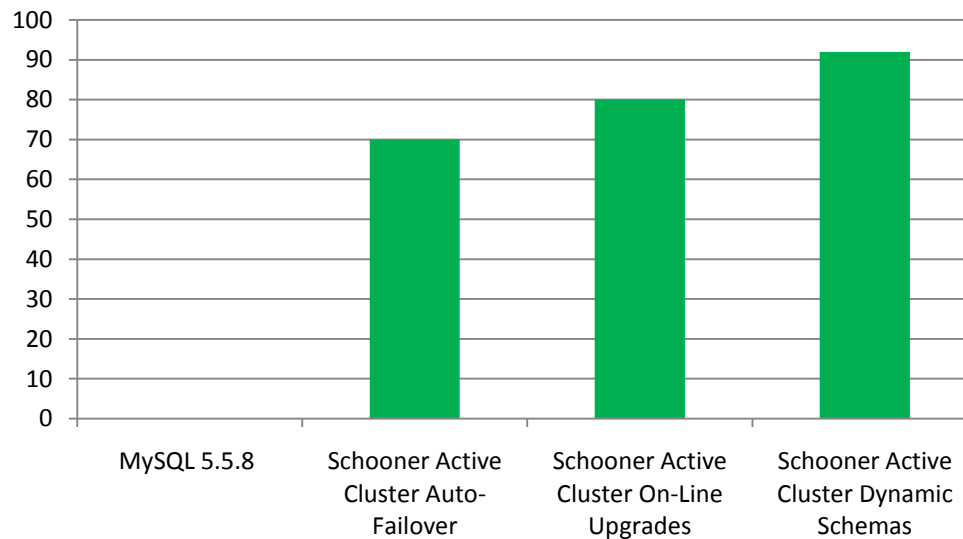
Breakaway Availability, Performance, TCO, Administration

Delivers Highest MySQL Availability

Schooner Active Cluster has
90% lower downtime than MySQL 5.5

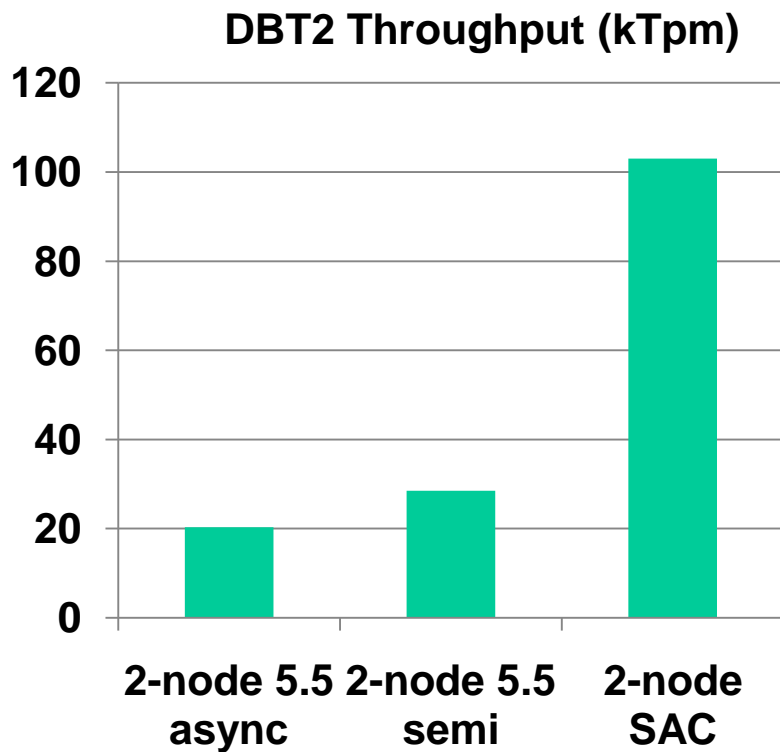


Availability Improvement
(% Cumulative Down Time Reduction)



Delivers Best MySQL Performance

Schooner Active Cluster gives
4x the throughput of MySQL 5.5



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

Fusion-io flash memory with HDDs

Schooner SAC High Performance with Stability

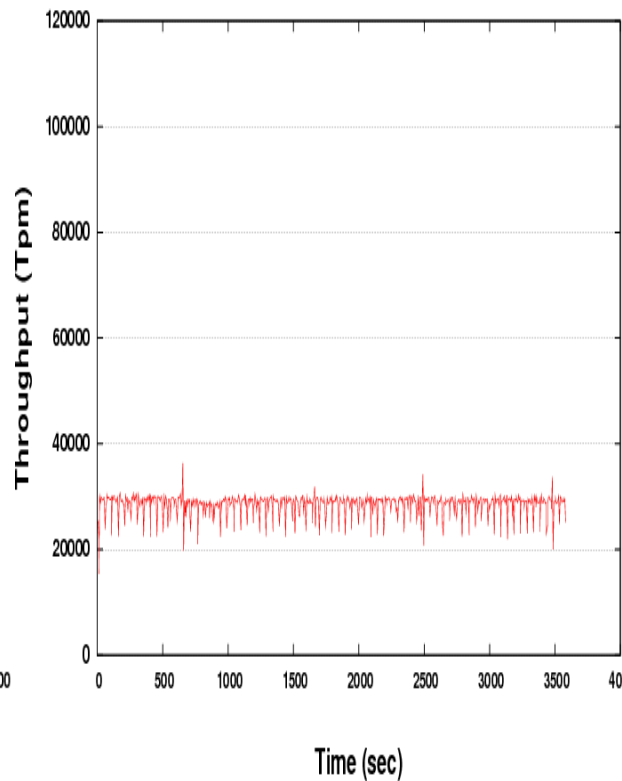
5.5 Async

Master Throughput vs. Time




5.5 Semi-sync

Master Throughput vs. Time




SAC

Master Throughput vs. Time




Schooner SAC Balances Systems, Drives Consolidation


CPU Utilization (%)



IOPS Utilization (%)




Network Utilization (%)



Delivers Lowest MySQL TCO

Schooner Active Cluster cuts TCO by 60% over MySQL 5.5



Delivers Easiest MySQL Administration

- Radical simplicity
 - Cluster view for administration and monitoring
 - Powerful 1-click actions for
 - On-line provisioning of servers and MySQL instances
 - On-line migration, upgrades
 - Monitoring and optimization
- Automatic failover & failback
- Integrated hot back-up
- GUI and CLI



Schooner MySQL Cluster-Wide Administration

View and manage all nodes and instances in the cluster from a single console

The screenshot displays the Schooner MySQL administration interface. The top navigation bar includes the Schooner logo, 'SCALE SMART', and user information 'Welcome back: admin' with links for 'Setting', 'About', and 'Sign Out'. The left sidebar shows a tree view of the 'Schooner Grid' containing a 'demo' group with two MySQL instances: 'mysqld1:lab137.schoonerinfotech.net' and 'mysqld1:lab136.schoonerinfotech.net'. The main content area is titled 'Overview' and features 'Attach Instance' and 'Setting' buttons. Below this is a 'Group Metric' table and an 'Instance Members' table. The 'Instance Members' table lists two instances: a Master instance on lab137.schoonerinfotech.net and a Slave instance on lab136.schoonerinfotech.net. Both are in a 'MySQL_READY' state. A 'Tasks' section at the bottom shows a successful 'Add Backup' task for the Master instance on 4/8/2011. The interface is updated as of 16:46:31 on 4/8/2011.

Group Metric

Type	Synchronous	VIP Policy	Balanced
User	admin	Read VIPs	10.1.137.3,10.1.136.3
Interface	eth4	Write VIPs	10.1.137.2
Async Slave	0	Schooner Data Format	Disabled

Instance Members

Name	Host	Version	Role	Progress	State	Commit/s	Select/s	Status
mysqld1	lab137.schoonerinfotech.net	5.1.52-3.1.547.393	Master	N/A	MYSQL_READY	0.00	0.20	up
mysqld1	lab136.schoonerinfotech.net	5.1.52-3.1.547.393	Slave	N/A	MYSQL_READY	0.00	0.00	up

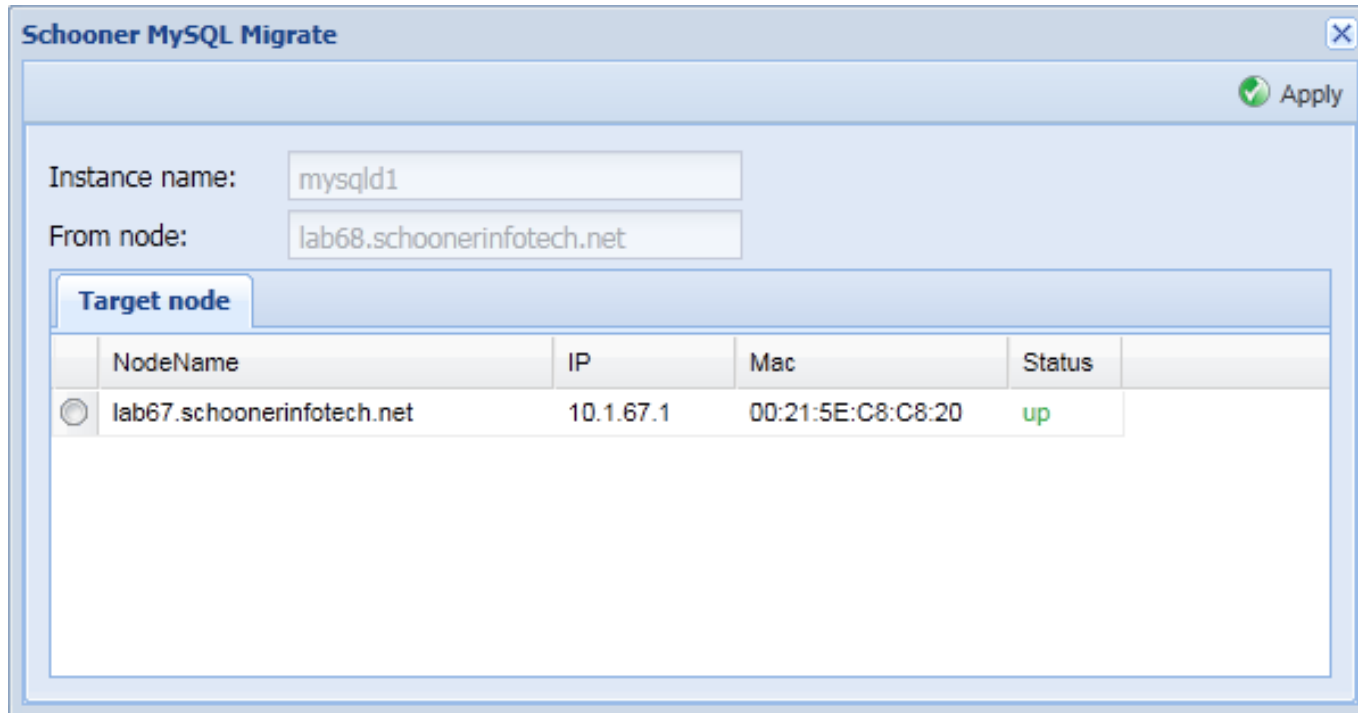
Tasks

Status	Name	Node	Instance	Group	Time(start)	Time(end)	Description
✓	Add Backup	lab137.schoonerinfotech.net	mysqld1	N/A	4:46:21 PM Apr/08/2011	4:46:22 PM Apr/08/2011	Add backup task successful.

Schooner MySQL Live Migration

Single click instance migration

Moves a live instance (including data) to a new node without service interruption



Schooner MySQL Migrate

Apply

Instance name:

From node:

Target node

	NodeName	IP	Mac	Status
<input checked="" type="radio"/>	lab67.schoonerinfotech.net	10.1.67.1	00:21:5E:C8:C8:20	up

Schooner MySQL Dashboard

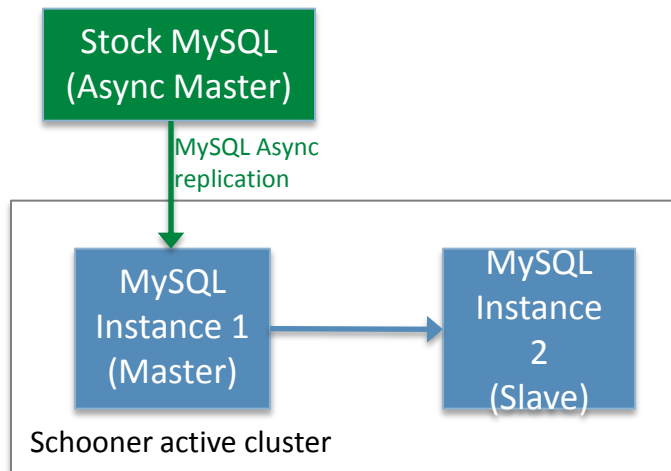
Cluster and Instance Monitoring , Troubleshooting , Optimization



Schooner MySQL : 100% MySQL Compatibility and Interoperability

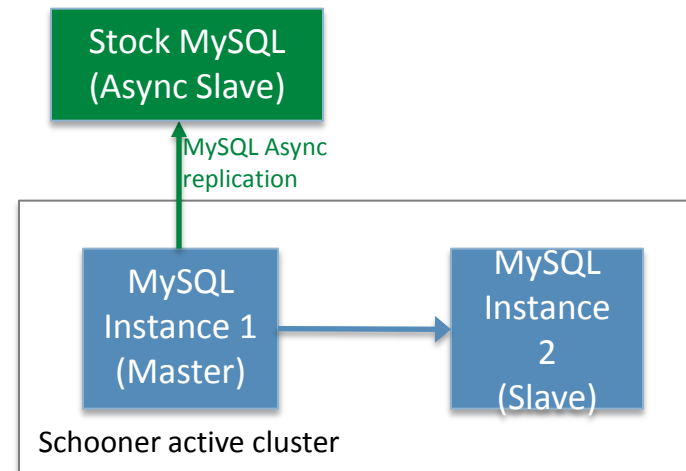


- Schooner Active Cluster is licensed, highly optimized, fully certified Oracle MySQL
- Works with all MySQL applications and databases – no migration required
- Interoperates with all other MySQL versions through asynchronous replication



Schooner Active Cluster as MySQL async slave

- Master functions as async slave and replicates synchronously to active cluster slaves
- Async link is automatically re-established from Instance 2 when master Instance1 fails



Schooner Active Cluster as MySQL async Master


- Master instance 1 functions as async master to stock slave and sync master to instance 2

Some Schooner MySQL Case Studies



"We were very pleased with the benefits from the Schooner NoSQL solutions. We are now deploying schooner Active Cluster MySQL 3.0 database solutions across our organization. The availability, performance, and management are industry leading, they solve our pain points completely..

– Ethan Erchinger, Director of Operations, Plaxo



"Schooner is the perfect solution for any MySQL enterprise whose business success requires great performance, exceptional reliability, and the ability to smoothly scale the datacenter as demand increases. Schooner helps us create a new wave of social networks, bringing technology that helps us create and sustain social communities like never before – efficiently, effectively, and effortlessly."

– Rayes Lemmens, CEO at MyLivePage



"We are very, very impressed with the performance and availability of Schooner Active Cluster 3.0. We plan to roll this out world-wide."

– Peter Leong, R&D Director, British Telecom



"In our business, Website performance and efficiency is key to the success of our Web properties. The Schooner MySQL Appliances have significantly helped GuteFrage improve their overall Website response time while at the same time allowing them to consolidate their database slaves onto a single Schooner appliance, dramatically reducing the time necessary for database administration."

– Frank Penning, CTO of Holzbrinck Digital



"We explored a variety of options from commodity SSD drives to PCI-express based flash memory cards. We decided to purchase Schooner MySQL. The performance has been great..

-- Mark Imbriaco DBA 37signals



"Our ad hoc MySQL queries run at least five times faster after installing Schooner. They deliver a huge performance benefit and are a breeze to install and manage."

– Darryl Weatherspoon, VP of Eng at Xoom

Schooner MySQL

MISSION CRITICAL



Highest Availability

- 90% downtime reduction vs. MySQL 5.5
- Fast, on-line fail-over, migration, h/w + s/w upgrades
- No data loss



Best Performance

- 4 - 20x throughput increase vs. MySQL 5.5
- Dramatic server consolidation



Fully Consistent Data

- Zero slave lag
- No data corruption



True Compatibility

- Licensed, optimized Oracle MySQL
- No schema, DB, or app changes needed



Lowest TCO

- Cuts TCO by 60% vs. MySQL 5.5



Easiest Admin

- Dramatically simpler cluster management
- Easy-to-use, powerful GUI and CLI

Schooner Software Products

Download and try for free!

Let us prove our breakaway availability, performance, consolidation, and savings in your shop with your data

- ❑ Schooner and our partners are always here to help you get there
- ❑ Just download our software for a free trial

Please stop by Schooner's Booth 308
see live demos, discuss with us,!

Thank You!