QuickStart Guide

Schooner Appliance for MySQL Enterprise[™] Version 2.0

You are here: Schooner Memcached Applance Software » Grid Configuration » Control Group Welcome back: admin Sign Out | About Client Nodes 🛞 Control Group - schooner_group62 🛞 Master 🗵 Avigation System Configuration + Group Level 🧏 Grid Configuration 🛛 🛊 + \bigcirc Applications + Group Details Monitor = Group Name: schooner_group62 😤 Gandia Schooner Memcached Group 62 Description: Cacti Nodes \$ 🗎 🔿 🔿 Name IP MAC OS Image Status master62.schoonerinfotech.net 00:21:5E:09:0D:D4 Linux 2.6.2 schoonerimage 172.16.1.62 ٩
 master62.schoonerinfotech.net
 1/2.1b.1.b2
 UU.21.b2.U3.U0.U4
 LINAL 2012
 Option 12.1b.1.b2

 schoonerinde0012.schoonerinfotech.net
 172.1b.1.12
 00.1A.64.C7.A9.CE
 N/A
 schooneringe
 ٩
 schoonernode0013.schoonerinfotech.rx
 172.16.1.13
 11:32.04:2E:11:2C
 N/A

 schoonernode0014.schoonerinfotech.rx
 172.16.1.14
 11:32.04:EE:2A:1A
 N/A
 • schoonerimage schoonerimage ٩ 11:32:0A:EF:21:22 N/A 11:32:0A:EE:21:E2 N/A schoonernode0015.schoonerinfotech.nt 172.16.1.15 schoonerimage • schoonernode0016.schoonerinfotech.nt 172.16.1.16 schoonerimage • schoonernode0017.schoonerinfotech.nr 172.16.1.17 11:32:0A:EE:21:2A N/A schoonerimage •



Technical Support

Go to the IBM support Web site at http://www.ibm.com/systems/support/to check for technical information, hints, and tips.

Documentation ID: I-I-v2.0-10120_619-QG-003

© 2009 Schooner Information Technology™, Inc. All rights reserved.

Schooner Appliance for MySQL Enterprise™ – QuickStart Guide

Issued December 2009.

Duplication or distribution without written permission is prohibited. Schooner reserves the right to revise this manual without notice.

Schooner Information Technology, Schooner Appliance for MySQL Enterprise, Schooner Appliance for Memcached, and the Schooner Logo are trademarks or registered trademarks of Schooner Information Technology in the USA and other countries.

InnoDB is the trademark of Innobase Oy. MySQL is a registered trademark of MySQL AB in the United States and other countries. Other products mentioned herein may be trademarks or registered trademarks of their respective owners.

Schooner Information Technology

501 Macara Ave., Suite 101 Sunnyvale, CA 94085, USA Tel: (408) 773-7500 (Main) (877) 888-5064 (Sales and Support) Fax: (408) 736-4212

Contents

Contents	iii
Hardware Overview	1
Software Overview	1
MySQL Enterprise	1
Supported Third-party Monitoring Software	2
IBM Diagnostics Interface	2
Schooner Capabilities/Features	2
Configuring the Appliance	2
Start, Stop, and Reboot the Appliance	3
Using the Front Panel	3
Using a Terminal Window	3
Using IMM Console	3
Start, Stop and Restart MySQL	3
Start, Stop, and Restart Schooner MySQL	3
Start, Stop, and Restart Admin MySQL	3
Initialize the Database and Perform the Initial Setup	4
Replication/Failover	4
Percona XtraBackup	4
Cache Replacement Policies	4
Network Bonding	5
VLAN Tagging	6
Monitoring	6
Overview	6
FAQs	7
How do I change a MySQL configuration?	7
How is a Solid State Drive (SSD) removed and replaced?	8
Why can't I stop MySQL after I've changed the domain name or hostname?	8
Why can't I log in when multiple appliances are used in a grid?	9

This QuickStart Guide provides a brief summary of the Schooner Appliance's capabilities and features, along with pointers to the specific information needed to get the appliance up and running in a minimum amount of time. This Guide is not meant to replace the FirstTime Wizards; it is provided as a quick reference for administrators who are already familiar with the configuration process.

Hardware Overview

The Schooner appliance leverages IBM's X-Architecture with the System x3650 M2 server.

- 2 Quad-core Intel Xeon 5560 "Nehalem" Series Processors
- 8 MB cache (max)
- 64GB 1333MHz, fully buffered, DDR-3 RDIMM memory
- 12 hot-swappable drive bays
- 4 SAS hard disk drives
- 8 or 16 SATA solid state drives (512GB/1TB)*
- 4 1-Gb Ethernet ports
 - * Capacity expansion is possible with an external drive enclosure EX3000.

Software Overview

MySQL Enterprise

- MySQL Enterprise 5.1
- Partitioning for very large database environments
- Row-based/Hybrid Replication for improved replication security
- Event Scheduler to create and schedule database tasks
- XPath Support
- Dynamic General/Slow Query Log
- Performance/Load Testing Utility (mysqlslap)
- Full Text Search (faster, new dev templates)
- Archive engine (better compression, more features)
- User session and problem SQL identification
- MySQL embedded library (libmysqld)
- Additional INFORMATION_SCHEMA objects
- Faster data import operations (parallel file load)
- MySQL web site: http://www.mysql.com/
- InnoDB web site: http://www.innodb.com/

Supported Third-party Monitoring Software

Schooner supports these system-monitoring tools:

Net-SNMP (http://net-snmp.sourceforge.net)

• Net-SNMP and IBM MIB

Nagios (http://www.nagios.org/)

- Open Source script-based monitoring
- Standard NRPE and Schooner-specific plug-ins

Ganglia (http://ganglia.info/)

- Lightweight cluster monitor
- Host and Schooner application metrics

IBM Diagnostics Interface

Schooner supports the IBM IMM service processor and its fully featured system diagnostics capabilities:

- Access to critical server settings
- Remote console
- Automatic notification and alerts
- Advanced failure analysis

Schooner Capabilities/Features

- Advanced performance upgrades tuned for SSD
- CAR and LRE replacement policies to match workloads
- Detailed monitoring of MySQL, server and cluster metrics
- Simple setup, backup/restore and master/slave configuration
- Locations of important files and directories:

MySQL configuration file: /opt/schooner/mysql/config/my.cnf MySQL data files: /schooner/data/db0 InnoDB table-space and transaction file: /schooner/txlog/db0 MySQL binary replication log: /schooner/binlog MySQL database backup: /schooner/backup MySQL root directory: /opt/schooner/mysql MySQL admin: /opt/schooner/mysql/bin/mysqladmin MySQL client: /opt/schooner/mysql/bin/mysql

Configuring the Appliance

To perform the following configuration tasks, you must open a terminal window with a Telnet/SSH client, like PuTTY. You may also use a Web browser to access the IMM service processor console.

Start, Stop, and Reboot the Appliance

Using the Front Panel

- Press the Power button.
- Wait at least 5 minutes for the system to boot.

Using a Terminal Window

- Login as "admin".
- Execute the desired command:

sudo reboot
sudo shutdown -h now

Using IMM Console

Using the IMM CLI from a terminal window:

• Use the telnet command to access the IMM IP address. For example:

#telnet 192.168.254.100

- Login to the IMM CLI.
- Execute the desired power command:

```
# power on
# power off
# power cycle
```

Using the IMM GUI from a Web browser:

- Enter the IMM IP address in the Web browser. For example:
- #http://192.168.254.100
- Login to the IMM GUI.
- Click on the Power/Restart link.
- Click on the desired power command link:

```
Power On Server Immediately
Power Off Server Immediately
Restart the Server Immediately
```

Start, Stop and Restart MySQL

The GUI provides buttons for these operations. The CLI-equivalent commands follow (log in as "admin"):

Start, Stop, and Restart Schooner MySQL

service mysql-sch start | stop | restart

Start, Stop, and Restart Admin MySQL

service mysqld start | stop | restart

Initialize the Database and Perform the Initial Setup

- Login as "admin"
- # sudo init_database_setup.php
- Enter "1" to select Initial setup.
- Enter "1" to select the Standard Configuration of MySQL.
- Enter "e" as many times as required, to return to the main menu.
- Enter "2" to select **Setting access and permissions**. Follow the prompts to create users and assign permissions to them.

Replication/Failover

Replication and failover management is provided by the *Multi-Master Replication Manager for MySQL (MMM)*. MMM (http://mysql-mmm.org/) is a set of scripts that performs the following functions:

- Configures a MySQL instance as a "master".
- Configures a MySQL instance as a "slave".
- Configures the MySQL Master-to-Master Management tool.
- Manages replication and failover, using the Master-to-Master Management tool.

To configure the replication option:

- Login as "admin"
- # sudo replication.php
- Follow the prompts to configure the system's replication and failover policies.

Percona XtraBackup

XtraBackup is an open-source back-up solution for the InnoDB storage engine. See http://www.percona.com/docs/wiki/percona-xtrabackup:xtrabackup_manual for more information.

To configure the backup and restore options:

- Login as "admin"
- # sudo backup.php
- Follow the prompts to configure the system backups.

Note: When entering the name of the root back-up directory, use the name of the directory specified in the "backup" option, in the "mysqld" section of the "my.cnf" file.

Cache Replacement Policies

Here is a brief overview of the two supported cache replacement policies.

Least Recently Used (LRU) – This policy tracks page usage over a relatively short time period, and discards the Least Recently Used page in the cache. The "recency" of each cache line is tracked. A "hit" to a cache page moves that page to the Most Recently Used (MRU) position, with a subsequent change in the status of all of the other cache lines. The MRU position is protected by a lock. This lock can be a source of contention,

because all cache hits are serialized behind it. Schooner has optimized the LRU policy, and made multi-threaded access more concurrent, by using fine-granularity locking.

Clock with Adaptive Replacement (CAR) – This policy combines the features of the "CLOCK" buffer pool replacement algorithm with those of ARC (Adaptive Replacement Cache), to create a cache replacement policy that considers frequency, as well as recency, when determining which page to discard from the cache.

CAR maintains two pairs of lists: each pair consists of a list of the pages in the cache, and a corresponding "history" list, containing the pages recently discarded from the cache. The first pair of lists represents the most *recently* used pages. The second pair represents those pages most *frequently* used.

Queries that incorporate scans can severely effect concurrently-executing transactions, by bringing "cold blocks" into the buffer pool. LRU is particularly susceptible to scans, because of the serial nature of its MRU lock. CAR, with no MRU and its ability to track frequency, provides improved scan resistance and better concurrent-transaction performance.

Upon selecting the cache replacement policy, MySQL must be restarted using the "service mysql-sch restart" command in order for the changes to take effect.

To select a cache replacement policy:

- Login as "admin".
- # cd /opt/schooner/mysql/config
- # sudo vi my.cnf
- Scroll down until you find this line, innodb_page_replacement_algorithm=lru The LRU policy is the default. To select the CAR policy, replace "Iru" with "car".

Network Bonding

Network Bonding (or Link Aggregation) is the grouping of multiple physical interfaces into a single "virtual" interface. Bonding interfaces together improves their combined link speed, and offers higher availability, due to the greater number of redundant interfaces.

- Login as "admin"
- # sudo /run_cli_ftw
- From the NIC menu, select option 3 to enter the bonded interface menu.
- Select option 1 to create a bonded interface then select the interfaces to bond:

```
Please select one option from the list above (type e/q to escape): 1
Enter the bond device name (Default "bond0"): bond
+-----+
| ID Select interfaces to bond |
+-----+
| 1 eth0 |
```

| 2 eth1

Please select options from the list above (comma-seperated, type e/q to escape): 1,2 Enter the bond device ipaddress: 172.16.1.28 Enter the bond device netmask (Default "255.255.255.0"): Enter the bond device mtu (Default "1500"): Enter the bond miimon (Default "100"): Enter the bond mode (Default "0"): SUCCESS: Create or modify bonded device

VLAN Tagging

VLAN Tagging (or IEEE 802.1Q) allows multiple networks to share a single physical network link transparently, and without any data leakage between networks.

To configure an interface's VLAN tag"

- Login as "admin"
 - # sudo /run_cli_ftw
- From the NIC menu, select option 2 to edit the configuration.
- Select an interface id number to edit the interface:

```
Please select one option from the list above (type e/q to escape): 2
Enter the ip address for interface eth1 (Default "172.16.1.76"):
Enter the virtual ip address for interface eth1 (Default "N/A"):
Enter the VLAN number (Default "N/A"): 1
```

Monitoring

Overview

Schooner Appliance for Memcached supports standard and open source event and performance monitoring interfaces and tools including SNMP, Nagios, and Ganglia.

Net-SNMP 5.3.1

0

- Includes standard Net-SNMP MIBs
- Includes IBM IMM MIB (service processor monitor)
 - Hardware event traps
 - Temperature
 - Voltage
 - Fan
 - o Power
 - o System health
 - o Machine ID
 - o Event log
 - o System clock
 - o Network interface

Nagios

- Nagios 3.0.6
- NRPE 2.5.1
- Includes standard Nagios plugins for:
 - o Disk partitions
 - o CPU, Memory, Swap, NIC
 - o RAID, HDD, SSD
 - o MySQL
- IPMI hardware monitor (sdr)

Ganglia 3.0.6

- Includes standard Ganglia cluster monitors
- Includes standard Ganglia host monitors for:
 - o Network traffic
 - o CPU states and load
 - o Disk I/O
 - o NIC
 - o RAID
 - o IOSTAT for HDD and SSD
 - o MySQL
- Includes Percona MySQL monitors for:
 - o Innodb
 - o MyISAM
 - o MySQL queries
 - o MySQL commands
 - o MySQL query cache
 - o MySQL locks

Each Schooner for MySQL Appliance is able to host the data collection and display services for Nagios and Ganglia. These are provided for reference and may also be used in small deployments. It is expected that most customers will host data collection servers on their own servers and would disable the appliance-based data collection services.

FAQs

How do I change a MySQL configuration?

The MySQL configuration file is /opt/schooner/mysql/config/my.cnf.

You may edit it using either of the following methods:

• Login as "admin" to a terminal window and edit the file using the command

sudo vi/opt/schooner/mysql/config/my.cnf

- Login as "admin" to the GUI.
 - Click on the "Applications" link to display the MySQL instances:
 - Click on the icon to pop up the MySQL edit dialog:

nooner HysQL			~
) 🛛			
	Key 🔺	Value	
∃ Group: client (2 Items)			
	port	3307	-
	socket	/schooner/data/mysql.sock	
3	Group: isamchk (4 Items))	
	key_buffer	2014	_
	read_buffer	2M	
	sort_buffer_size	20M	
	write_buffer	2M	
E	Group: myisamchk (4 Ite	ms)	
	key_buffer	20M	
	read_buffer	214	100

Edit the desired fields and click the log button to apply.

How is a Solid State Drive (SSD) removed and replaced?

The SSDs are hot swappable. Simply remove and replace them as follows:

- 1 Lift up on the release latch.
- 2 Rotate the handle down and pull the drive out of the chassis.
- 3 Installation is the reverse of this procedure.



Note: To maintain proper cooling, do not operate the system for more than 10 minutes, without either a drive, or a filler panel, installed in each drive bay.

Why can't I stop MySQL after I've changed the domain name or hostname?

After you have updated the domain name or hostname while running MySQL, you may not be able to stop MySQL. This is because the command such as "service mysql-sch stop" looks for the process ID of MySQL that is stored in the /schooner/data/db0/{hostname}.pid file. When the hostname or domain name is updated while MySQL is running, the name of the .pid file still corresponds to the old settings. The discrepancy makes it impossible to stop MySQL.

To solve this problem, do the following:

- 1. Issue the command: "/opt/schooner/mysql/bin/mysqladmin defaults- file=/opt/schooner/mysql/config/my.cnf
 shutdown"
- 2. Manually update the name of the {hostname}.pid file.

Why can't I log in when multiple appliances are used in a grid?

When multiple Schooner MySQL Appliances are used in a grid fashion, it may be difficult, under certain situation, to log into the master node from the Schooner Administrator GUI because the node does not become the master node.

To solve this problem, do the following:

- 1. Log into the shell.
- 2. Make the nodes independent rather than being part of a grid by executing:

/opt/schooner/mesh/bin/configure_mesh.pl -c <key>
key could be any unique name such as the host name.