

Installation & Setup Guide

Schooner Membrain

Version 3.1

Software that transforms standard x86 servers and flash memory into NoSQL and caching super-servers



Technical Support

Technical support for Schooner Information Technology products in North America is available from the following sources:

- Phone: (877) 888-5064; (408) 888-1619, Monday through Friday
- Fax: (408) 736-4212
- Email: support@schoonerinfotech.net
- Website: www.schoonerinfotech.com/support

Documentation ID: S-M-v3.1IG-01

© 2009 ~ 2011 Schooner Information Technology™, Inc. All rights reserved.

Schooner Membrain Installation & Setup Guide

Issued January 2011

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™ with InnoDB, Schooner Membrain™: the smart NoSQL data store, based on 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 |
| Chapter 1: Introduction..... | 1 |
| Hardware and Software Prerequisites | 1 |
| Hardware | 1 |
| Software..... | 1 |
| Typical Schooner Membrain Configuration | 2 |
| Small Schooner Membrain Configuration | 2 |
| Secondary Storage/DRAM Ratio | 2 |
| Hostname | 2 |
| Networking..... | 2 |
| Miscellaneous | 3 |
| Chapter 2: Installation and Licensing | 4 |
| Prerequisites for Installation..... | 4 |
| Software..... | 4 |
| 3 rd Party Software Installation..... | 4 |
| Schooner Membrain Software Installation..... | 5 |
| Platform Verification..... | 6 |
| Uninstall..... | 7 |
| Licensing..... | 7 |
| Chapter 3: Initializing Schooner Membrain | 8 |
| Schooner Membrain Administrator Account..... | 8 |
| Initializing Schooner Membrain..... | 8 |
| Chapter 4: Supported Hardware Platforms..... | 10 |
| Base Server | 10 |
| Flash Memory | 10 |
| Availability Options | 11 |
| Recommended Configurations..... | 11 |
| DELL | 12 |
| HP | 12 |
| IBM | 12 |
| System Tuning..... | 13 |
| SSD Controller | 13 |
| BIOS | 13 |

Chapter 1: Introduction

This guide explains how to install the Schooner Membrain software appliance and prepare it for execution. For information about managing and monitoring Schooner Membrain, see the *QuickStart Guide* and the *Schooner Membrain Application & Administration Guide*.

Hardware and Software Prerequisites

Hardware

Schooner Membrain supports an enhanced version of the Memcached cache protocol and provides advanced performance, availability and management features. Schooner has qualified specific hardware configurations that support high performance:

- x86_64 CPU
 - 2 sockets
 - 16 cores
 - hyper-threading enabled
- 64 GB DRAM
- Flash storage
 - 1 PCIe flash card, or
 - 4 x SAS/SATA SSD (2.5" format)
- Controllers
 - For SSD installations, separate controllers for HDD and SSD.
 - For PCIe flash card installations, a single HDD controller suffices.
- 10GE network interface or multiple bonded 1GE interfaces
 - Used for both application and replication traffic.
- BIOS
 - Power saving mode disabled.

Please see Chapter 4 for details on supported hardware platforms.

Software

Schooner Membrain is supported on the following Linux releases:

- Centos 5.4/5.5
- RHEL 5.4/5.5

Schooner Membrain supports a performance charting package that requires PHP 5.1.6 or later. The Schooner Membrain charting package is compatible with Internet Explorer 7.0 (or later), or Firefox 3.0 (or later).

Typical Schooner Membrain Configuration

During initialization, you will be asked to assign CPU, DRAM and secondary storage to Schooner Membrain. The amount of each system resource allocated depends on the performance and capacity needs of your application.

A typical, high-performance, high-capacity installation might use the following Membrain configuration:

- 6 of 8 CPU cores (75%)
- 64 GB DRAM
- 512 GB secondary storage
- 10GE network interface or bonded 1GE interfaces

Small Schooner Membrain Configuration

For evaluation systems, you may use a small Membrain configuration:

- 1 CPU core
- 4 GB DRAM
- 32 GB secondary storage
- 1GE network interface

Secondary Storage/DRAM Ratio

In general, a ratio of 8:1 secondary storage to DRAM is required for best performance. You may adjust DRAM and secondary storage allocations, but if this ratio is increased beyond the recommended value your performance may be reduced due to increased DRAM cache misses.

Hostname

Ensure that the node has a hostname that is unique in the cluster.

Networking

Ensure that the interface you will be using for the Schooner Cluster Manager ("SCM") is running and able to communicate with all other SCM interfaces in the Schooner cluster.

Ensure that the interfaces you will be using for the membrain messaging service (used in replication) are running and able to communicate with the other node in the mirrored pair. For example, if you want to use eth0 on node 0 and eth1 on node 1, you must set up the routing tables on the two nodes so that packets from eth0 on node0 are routed to eth1 on node 1, and vice-versa.

Membrain uses the following ports, which must be opened on each Membrain node:

- 44160: used by Membrain engine [hard-coded]

- 51350: admin port [configurable]
- 51360: split brain handler [hard-coded]
- 51503: SCM cluster guard [configurable]
- 51515: SCM messaging [configurable]
- 51512: JMX remote monitor [hard-coded]
- plus any user-configured ports for membrain containers

Miscellaneous

SELinux must be disabled on nodes running Membrain.

Chapter 2: Installation and Licensing

This section will show you how to install the Schooner Membrain server and required packages.

Prerequisites for Installation

Software

The Schooner Membrain server and CLI have been qualified to execute under RHEL5 and Centos 5 operating systems and require the following packages:

```
* php-5.2.9
* php-cli-5.2.9
* php-common-5.2.9
* php-gd-5.2.9
* php-ldap-5.2.9
* php-pdo-5.2.9
* httpd-2.2.3-31
* java-1.6
* perl-DBI-1.52-2.e15
* perl-Net-Telnet-3.03-1.2.e15
* expect-5.43.0-5.1
* sysstat-7.0.2-3.e15
* gmp-4.1.4-10.e15
```

If you do not have these packages installed, you may either download them or install them using the Schooner Membrain 3rd-party software installer.

3rd Party Software Installation

You may use the Schooner 3rd party software installer to bring your system up to date. Here are the steps to install the 3rd party packages:

- Copy the 3rd party installer package to your server.
- To view the installer options:

```
sh schooner-memcached-3rdparties3.1.0_15048_1303.sh -h
Usage:
schooner-memcached-3rdparties-3.1.0_15048_1303.sh -h
schooner-memcached-3rdparties-3.1.0_15048_1303.sh -r
schooner-memcached-3rdparties-3.1.0_15048_1303.sh -v
schooner-memcached-3rdparties-3.1.0_15048_1303.sh -i

-h : displays this help message
-r : displays README
-v : displays version
-i : installs
```

- Install the 3rd party packages:

```
sh schooner-memcached-3rdparties3.1.0_15048_1303.sh -i
Starting install process for Schooner 3rd-party Package, this may take several
minutes...

Saving copy of package to '/var/schooner/packages'...
Extracting temporary files into /tmp
```

```

Running deploy script...
Performing package install...
Preparing... #####
Successfully applied package schooner-memcached-3rdparties-3.1.0_15048_1303
Cleaning up temporary package files from /tmp

Copy of package can be found in the following directory:

    /var/schooner/packages

Patch install log file can be found here:

    /var/schooner/packages/log/schooner-memcached-3rdparties-3.1.0_15048_1303-
2010_11_16-17_59_46.log

```

Schooner Membrain Software Installation

Once you have ensured that the required 3rd party packages have been installed, you may install the Schooner Membrain server:

- Copy the Schooner Membrain installer package to your server.
- To view the installer options:

```

sh schooner-memcached-3.1.0_15048_1303.sh -h

Usage:

    schooner-memcached-3.1.0_15048_1303.sh -h
    schooner-memcached-3.1.0_15048_1303.sh -r
    schooner-memcached-3.1.0_15048_1303.sh -v
    schooner-memcached-3.1.0_15048_1303.sh -i
    schooner-memcached-3.1.0_15048_1303.sh -u

    -h : displays this help message
    -r : displays README
    -v : displays version
    -i : installs
    -u : rollback

```

- Install Schooner Membrain:

```

# sh schooner-membrain-3.1.0_16083_1407.sh -i
Date: 12/01/2010
Schooner Application: Membrain
Package: schooner-memcached-3.1.0_16083_1407.sh

=====
README - Schooner Membrain 3.1.0
=====

This document provides information and/or instructions on the following:

    * POINT PACKAGE VERSION
    * PREREQUISITE
    * IMPORTANT NOTES
    * PACKAGE CONTENTS
...

```

- After the README is displayed, answer “y” to the “install package” question.

```

Starting install process for Schooner Membrain/NoSQL, this may take several
minutes...
Do you want to install this package? [y/N]: y
Saving copy of package to '/var/schooner/packages'...
Extracting temporary files into /tmp
Running deploy script...

```

- The minimum hardware configuration will be checked. Answer “y” to continue if the hardware does not meet requirements:

```

| Required Hardware Configuration Not Met
| -----
| CPU Cores: 2 is less than the required 16
| Hyperthreading: no is less than the required yes
| Total DRAM: 3.00G is less than the required 64G
| Ethernet 10GE: 1GE is less than the required 10GE
| -----
| End
| Do you want to continue? [y/n]: y

```

- The new version will be installed.

```
warning: /tmp/schooner-membrain-3.1.0_16083_1407/schooner_helm-3.1.0-
16083.x86_64.rpm: Header V3 DSA signature: NOKEY, key ID ffa79912
Preparing... #####
schooner_helm #####
warning: /tmp/schooner-membrain-3.1.0_16083_1407/schooner_scm-3.1.0-
16083.x86_64.rpm: Header V3 DSA signature: NOKEY, key ID ffa79912
Preparing... #####
schooner_scm #####
warning: /tmp/schooner-membrain-3.1.0_16083_1407/schooner_utils-3.1.0-
16083.x86_64.rpm: Header V3 DSA signature: NOKEY, key ID ffa79912
Preparing... #####
schooner_utils #####
warning: /tmp/schooner-membrain-3.1.0_16083_1407/schooner_membrain-3.1.0-
16083.x86_64.rpm: Header V3 DSA signature: NOKEY, key ID ffa79912
Preparing... #####
schooner_membrain #####
warning: /tmp/schooner-membrain-3.1.0_16083_1407/schooner_sendincident-3.1.0-
16083.x86_64.rpm: Header V3 DSA signature: NOKEY, key ID ffa79912
Preparing... #####
schooner_sendincident #####
/dev/mem: mmap: Bad address
Preparing... #####
emt-membrain #####
Preparing... #####
emt-charts-membrain #####
Stopping httpd: [ OK ]
Starting httpd: [ OK ]

Installation completed successfully.
User "madmin" is created to run Schooner CLI.
Please grant permission to user "madmin" to run all commands without a password
After that, you can run Schooner CLI as user "madmin".

User "membrain" is created to run Schooner Membrain.
Please grant permission to user "membrain" to run these commands without a
password :
/opt/schooner, /sbin/ifconfig, /sbin/arping, /sbin/ip, /sbin/iptables,
/sbin/route ,/usr/bin/killall

Successfully applied package schooner-membrain-3.1.0_16083_1407
Cleaning up temporary package files from /tmp/schooner-membrain-3.1.0_16083_1407
```

The Schooner Membrain installer creates a “membrain” account for use by the Schooner Membrain server. This account is set “nologin”.

Also, an “madmin” account is created which invokes the CLI. The password is the same as the login id. You may choose to create your own login for CLI administration. The path to the Schooner Membrain CLI is:

```
/opt/schooner/admin/bin/schcli
```

Platform Verification

The Schooner Membrain installer will check for required hardware resources:

- CPU cores and hyper-threading
- DRAM capacity
- 10GE network interface

If any of the requirements are not met an error message will be displayed:

```
| Required Hardware Configuration Not Met
| -----
| CPU Cores: 2 is less than the required 16
| Hyperthreading: no is less than the required yes
| Total DRAM: 3.00G is less than the required 64G
| Ethernet 10GE: 1GE is less than the required 10GE
| -----
| End
| Do you want to continue? [y/n]: y
```

Please note that if the hardware resource check fails and you decide to proceed, your performance may not reach the levels qualified by Schooner Information Technology.

Uninstall

- Execute the install script using the “-u” option. You have the option to keep the default admin user account and Membrain data.

```
sh schooner-membrain-3.1.0_16083_1407.sh -u
Starting uninstall process for Schooner Membrain/NoSQL, this may take several
minutes...
Do you want to uninstall this package? [y/N]: y
Extracting temporary files into /tmp
Running deploy script...
stopping service...done
removing schooner rpms...done
Do you want to keep the user and data of this application? [y/n]: y
Successfully rolled back package schooner-membrain-3.1.0_16083_1407
```

Licensing

Schooner Membrain comes with a 96-hour trial license. After the license expires Schooner Membrain will cease to function.

The remaining time on your license is displayed whenever you login to the CLI:

```
License Status : (You have 95 hours to get a production license)
xen201v01.schoonerinfotech.net>
```

Please contact Schooner Sales to purchase a product license.

Once you have purchased a license, install it by copying the file to the license directory:

```
# cp schooner_license /opt/schooner/license
```

You do not have to restart Schooner Membrain after installing the license file.

Chapter 3: Initializing Schooner Membrain

To initialize Schooner Membrain, you must run the First Time Wizard CLI.

Schooner Membrain Administrator Account

The Schooner Membrain installer creates the “madmin” user account for use as the default administration user. The default password is also “madmin”.

The administration account requires privileged access to certain files and commands. Please grant “madmin” these rights using the “sudoers” facility. The easiest way to perform this is to execute the command “visudo” and add the following line:

```
madmin ALL=(ALL) NOPASSWD: ALL
membrain ALL=(ALL) NOPASSWD: ALL
```

Initializing Schooner Membrain

Once the Schooner Membrain is installed, you need to initialize the server using the FTW CLI:

1. Log into the appliance using the ‘madmin’ user name and password.
Note: If the login is successful, you’ll see the welcome message and the Schooner Membrain End User License will appear.
2. Read the End User License by tapping the Enter key on your keyboard until you have reached the end of it.

```
Do you agree? (yes/no) [no]:
```

3. Type ‘yes’ to accept the End User License. The Schooner First Time Wizard starts. See Step 4.

Note: You must accept the End User License in order to start the Schooner First Time Wizard.

4. Configure the number of cores to assign to Schooner Membrain.

Schooner Membrain requires that a set of CPU cores be assigned to it. No other applications will be able to use the CPU cores assigned to Schooner Membrain.

```
Your system core number are 8
Schooner recommends setting the number of dedicated cores to 75% of all
available cores.

CPU Cores          [6]:6
```

5. Configure the amount of DRAM to use for the Schooner Membrain cache.

```
Your system total memory is 15.7 GB.
Schooner recommends setting 90% of system memory for Memcached for best
performance.

Memory(GB) allocation for memcached: (1GB at least) [14]:1
```

6. Configure the Schooner Membrain storage.

Storage can be either a file path or a device path.

```
Primary storage type, physical device or file system partition(d/f) [d]:f
File System Storage [/opt/schooner]:
Storage Size[G] [25]:20
```

7. Apply the configuration.

Schooner Membrain will be initialized, including a default container.

```
Do you want to apply changes?(y/n) [y]:
We are dding file /opt/schooner/schooner0 for memcached, please wait for
few minutes.
This Configuration Type : Small Server
10240+0 records in
10240+0 records out
10737418240 bytes (11 GB) copied, 66.2423 seconds, 162 MB/s
Warning: all configs are reset! And a demo container is generated.
```

8. After initialization, the CLI prompt will be displayed.

If your license is about to expire, a warning will be printed.

```
License Status : (You have 94 hours to get a production license)
```

At this point, you are now ready to start using your Schooner Membrain server. Please see the Schooner Membrain Quickstart Guide for an explanation of concepts and configuration tasks.

Chapter 4: Supported Hardware Platforms

Schooner Membrain will run on many hardware and software platforms, but is optimized to run with flash storage. We don't list required configurations but do list recommended configurations on which we've done extensive testing.

Here are the main items in your selection of Base Server, Flash Memory and Availability Options.

Base Server

You can use any HP, Dell, or IBM multi-core x86 server or blade (abbreviated simply as servers below), running Red Hat Linux 5.4 / 5.5 or CentOS 5.4 / 5.5.

We optimized our software for maximum concurrency to fully exploit multiple cores. Schooner-powered servers deliver great performance when they have at least two 4-, 6-, or 8-core processors, and performance scales linearly with number of cores. Our software runs on less powerful machines, but performance will suffer.

We recommend 64 GB of DRAM, since our software uses DRAM as a cache for the flash memory to give you DRAM-type speeds from all of the flash memory. More DRAM never hurts, but you usually won't need it since we fully exploit the flash IOPS to deliver performance equivalent to what you'd get if you used much more DRAM.

Pay attention to your server power supply as you add flash drives and controllers. PCIe-based flash drives tend to consume much more power than SAS drives. The bottom line is that you may have a fine server with, say, a standard 460W power supply, but find that you use 465W or more when it's fully outfitted.

One "gotcha" in many servers (as of January 2011) is the default disk controller. SAS and SATA flash drives (see next section) can be seen as super-fast energy-efficient disk drives, and are inserted into HDD slots. But many servers ship with a default disk controller that is ineffective with SAS and SATA drives. We've found this quite common with HP and Dell, for example. Be careful about the controller if you use SAS or SATA drives; you can always upgrade to something better such as the LSI 9211. Or use PCIe flash drives.

Many modern servers have a power-saving mode, settable in the BIOS. You'll want to turn this off to get the maximum benefit from Schooner software.

Flash Memory

We support flash memory in the SATA, SAS, and PCIe form factors. This includes flash memory qualified and sold by HP, Dell, and IBM, and from Fusion-io, LSI, Unigen, Intel, OCZ, Smart Modular, and Pliant.

Because our software is tuned for maximum concurrency and optimal use of multiple cores, it can exploit the high I/O rates of multiple flash drives in parallel. There are tradeoffs in the economics, performance, and capacity of using (a) PCIe drives with very high IOPS and (b) SATA / SAS drives with lower IOPS but

which are probably cheaper for the same capacity. Schooner MySQL will work either way, but the most compelling economics and best performance usually come from using a larger number of cheaper drives than from using a smaller number of expensive but faster drives.

The number of available slots on your server can limit your flash drive choices: for example, you can't plug in eight low-cost SAS or SATA SSDs if you only have two slots.

What about SLC, MLC, or eMLC flash? There's a definite tradeoff between the cost and wear – which equates to lifetime – of flash memory. SLC (Single Level Cell) tends to be very expensive. MLC (Multi-Level Cell) is cheaper but less durable than SLC. eMLC (Enterprise MLC) provides a nice balance of low cost and high durability. While Schooner supports all forms of flash memory, we recommend eMLC.

Availability Options

To RAID or not to RAID? Although flash memory is very reliable, flash drives can still fail. Schooner MySQL supports the option of RAID5 among SATA and SAS drives in a server, with hot-swap of SSDs if one fails. But you can't hot swap PCIe drives. One option is to use the RAID provided in some drives, such as Fusion-io, as recommended by the manufacturer. Schooner MySQL works with all of these choices, but your choice does affect the user experience.

Database backup can be done using on-server HDDs, the capacity of which should match your flash memory capacity. But this is optional and you can always back up across a LAN or SAN.

Recommended Configurations

These servers have been qualified by Schooner:

- DELL R710
- HP DL 380 G6/G7
- IBM 3650 M3

While the supported server platforms include widely deployed models using standard configurations, price and performance considerations may dictate an upgrade of installed servers or affect the bill of materials for new server orders.

In particular, Schooner has qualified two sets of platform and flash configurations, one that maximizes performance and another that maximizes price/performance.

Flash storage devices are extremely fast compared with hard disks. They are best used for database table storage while logging is more suited to hard disks. In order to support the highest level of transaction throughput, the I/O path to each type of drive must be configured for optimal performance.

The following sections describe in detail the hardware platform configurations and settings supported by Schooner Membrain.

DELL*Price/Performance Configuration*

- Controller: 2 x PERC H700 (one for HDD, one for SSD)
- Storage: Unigen eMLC SSDs (200GB)

Maximum Performance Configuration

- Controller: 1 x PERC H700 or PERC 6i
- Storage: Fusion-io ioDrive Duo 320GB or Duo 640GB

Additional Notes for DELL:

- As of this writing, Unigen SSDs are not supported by DELL. You may be asked to remove the SSDs before DELL support will support your environment.
- Schooner recommends 256MB or more of non-volatile (NV) cache for the H700 controllers.

HP*Price/Performance Configuration*

- Controller: 2 x HP Smart Array P410i (one for HDD, one for SSD)
- Storage: Unigen eMLC SSDs

Maximum Performance Configuration

- Controller: 1 x HP Smart Array P410i
- Storage: Fusion-io ioDrive Duo 320GB or Duo 640GB

Additional Notes for HP:

- HP does not support Unigen SSDs. You may be asked to remove the SSDs before HP support will support your environment.
- Schooner recommends 256MB or more of non-volatile (NV) cache for the P410i controllers.

IBM*Price/Performance Configuration*

- Controllers:
 - 1 x IBM ServeRAID MR10i attached to HDD devices.
 - 2 x LSI 9211-8i attached to SSD data devices.
- Storage: Unigen eMLC SSDs

Maximum Performance Configuration

- Controller: 1 x IBM ServeRAID MR10i
- Storage: Fusion-io ioDrive Duo 320GB or Duo 640GB

Additional Notes for IBM:

- IBM does not support Unigen SSDs. You may be asked to remove the SSDs before IBM support will support your environment.
- Schooner recommends 256MB or more of non-volatile (NV) cache for the two controllers.

System Tuning

While the supported server platforms include widely deployed models using standard configurations, price and performance considerations may dictate an upgrade of installed hardware and reconfiguration of controller and BIOS settings.

SSD Controller

For SSD installations, a dedicated controller will yield the best performance. Installations that connect both HDD and SSD on a single controller will be performance limited.

- Controller cache
 - Minimum of 256MB
- File system
 - XFS

BIOS

BIOS settings that control CPU clock rates may have an effect on performance of Schooner Membrain. In particular, power saving modes can decrease the clock rate.

- Power saving mode off.
- Processor hyper-threading on.