Application & Administration Guide

Schooner Appliance for Memcached™

Version 2.0

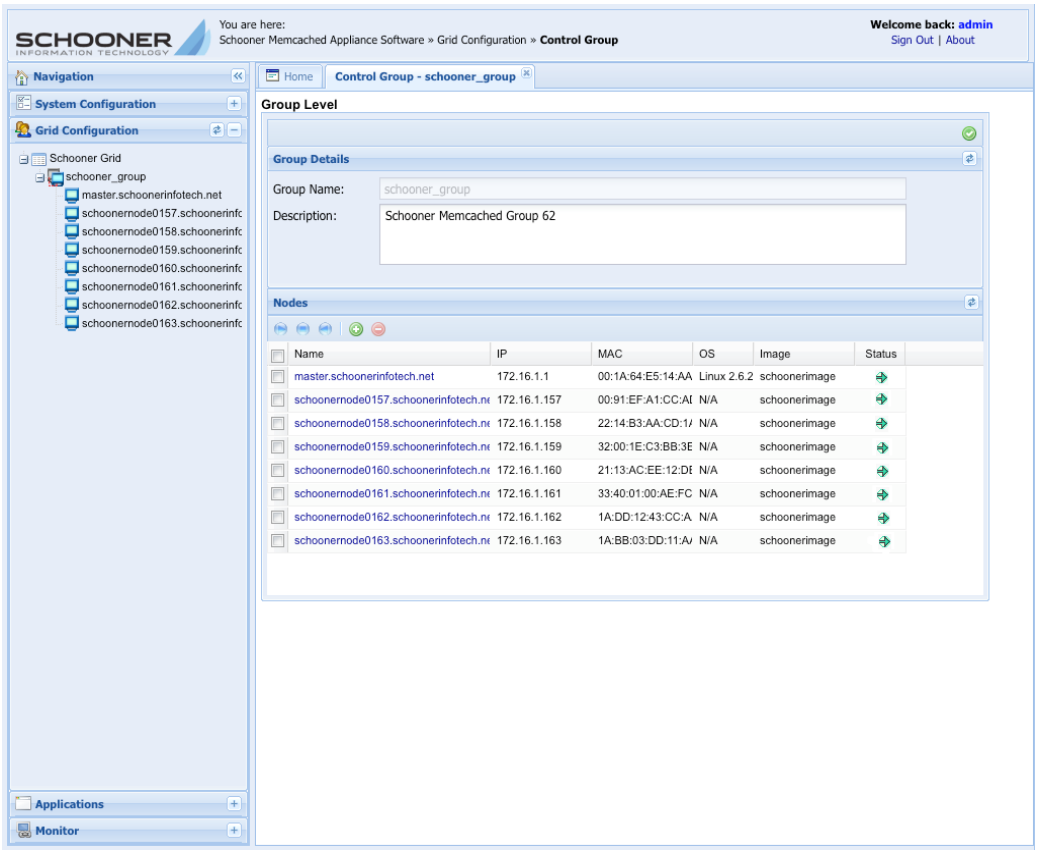

SCHOONER
INFORMATION TECHNOLOGY

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-M-v2.0-10120_619-AG-002

Schooner Appliance for Memcached™ - Application & Administration Guide

Issued November 2009.

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

Chapter 1. Introduction

Welcome

Thank you for purchasing the Schooner Appliance for Memcached™ and welcome to Schooner's new world of smart, fast, scalable, cost-effective data access.

Your Schooner appliance is compatible with your existing Memcached environment.

The information in this Guide is organized as follows:

- Chapter 2. "Configuring Your Memcached Environment" discusses the special considerations in optimizing Memcached for a Schooner system.

  Note: The remaining chapters describe the Schooner Administrator control console, which helps you manage and monitor your Schooner appliance.
- Chapter 3 "The Schooner Administrator Control Console" describes the Schooner Administrator Control Console user interface, showing its functions and how to get help and support.
- Chapter 4 "Common Management Tasks" describes how to manage and monitor your nodes, groups, and application instances.
- Chapter 5: "Common Administrative Tasks" describes how to upgrade the system, back up and restore configuration files, manage users, and change passwords.
- Chapter 6 "Troubleshooting" provides advice about resolving problems.
- Chapter 7 "Schooner Memcached Command Line Interface" describes how to use the Command Line Interface to configure and manage the system.
- Chapter 8 "Mirrored Group Network Setup Requirements" contains important information about the network environment that must be taken into consideration when configuring mirrored groups.

Prerequisites

System Requirements

Schooner Administrator can operate on a Windows or Linux operating system running either:

- Internet Explorer 7.0 or later; or
- FireFox 3.0 or later

Install, Connect, and Configure the Schooner Appliance

Before using Schooner Administrator, rack-mount and cable-connect the Schooner appliance and run the Schooner FirstTime Wizard. For instructions, see the *Schooner Appliance for Memcached - Installation & Setup Guide*.

External Network Management Tools

Schooner Appliance for Memcached is compatible with standard network management tools like Nagios, Cacti, and Ganglia. It also supports SNMP (Simple Network Management Protocol).

You will be able to monitor and administer your entire network from one management console. Schooner Administrator is needed only for configuring the Schooner system, managing the grid, and controlling application instances.

Integrating Schooner Memcached

Normally, integrating the Schooner appliance into your Memcached environment is simply a matter of adding the Schooner appliance's network interfaces to the configuration file of your Memcached clients. There should be no need to alter your application client code unless you require data persistence. In that case, see "Persistence" on page 4.

Schooner Administration Architecture

Master Node

Schooner Administrator operates through a "master node" that maintains the configuration of all Schooner nodes in the "grid" or collection of Schooner nodes. To perform maintenance on any node in the grid, log into the Schooner Administrator located on the master node and access the "client node" you wish to administer.

In the event of a master node failure, a new master node will automatically be selected from the remaining nodes. Operation will continue normally with the remaining nodes.

Client Node

A client node in the Schooner Administrator system is simply a non-master member of the Schooner grid. Client nodes belong to a "group" of which there may be several in a grid.

The image installed on client nodes and master node is the same.

All maintenance operations are performed via remote calls from the master node to the client nodes.

Alternative Administration

Schooner Appliance for Memcached nodes may also be administered via Command Line Interface (CLI), which is covered in **Chapter 7**. If this method is chosen, your administration system is responsible for maintaining all nodes without a Schooner master node.

Schooner Memcached High Availability

RAID

The Schooner Appliance for Memcached supports software RAID 5 SSD configuration that can tolerate a single drive failure.

Mirrored Groups

Schooner Appliance for Memcached supports replication of Memcached keys and data between a pair of mirrored nodes.

A “mirrored group” defines a high availability domain within a Schooner cluster:

- Replication of updates on one node to the other node
- Automatic fail-over of Memcached access via VIP (virtual IP)
- Automatic recovery of a failed node

When nodes are added to the system, they are assigned to a pool of non-replicating nodes, the Independent Group. These appliances do not support replication.

In order to start replication between two nodes, a mirrored group is created via Schooner Administrator by simply selecting the nodes and creating a new group. From this point on, all Memcached data will be replicated between the two nodes.

In the event of a failure of a mirrored node, the survivor node will take over the VIPs of the failed node. This allows Memcached clients to continue operation.

Upon recovery of the failed node, the failed node will sync all data from the survivor node. Once the recovery is complete, the failed node will take over its VIPs and operation should return to the normal state.

Backup

Schooner Appliance for Memcached supports backup and restore of persistent Memcached data. For more information, refer to Back Up and Restore Procedures in Chapter 5.

Chapter 2. Configuring Your Memcached Environment

Schooner Appliance for Memcached is compatible with existing Memcached environments. This chapter discusses special considerations in configuring Memcached for a Schooner system.

Configuring Your Memcached Clients

Configure your Memcached clients to use your Schooner Appliance for Memcached's *listening addresses*. These are the appliance's IP addresses on your application subnets, which you entered in Step 1 of the FirstTime Wizard (FTW).

In most cases, you can do this by simply adding the Schooner appliance's network interfaces to the configuration file of your Memcached clients.

Persistence

Because the Flash memory in your Schooner Appliance for Memcached is non-volatile, the data stored there is persistent. However, to optimize performance, DRAM memory is used as a cache front-end. DRAM memory is volatile and will lose data in the event of a power loss or node failure.

If you are using your appliance as a standard cache, you may not need your data to be persistent. In that case, set Persistence to **NO**, and the cache (first DRAM, then Flash) will start off empty on restart. However, it will take some "warm-up" time for the cache to fill. The "warm-up" time could be shortened if the cache data was persistent.

If you are using your appliance as a back-end storage medium, it is important that all data be persistent.

If you set Persistence to **YES**, beware that the persistent data in Flash memory may be "stale" (out of date) upon restart after failure. This is because, upon restart, a data object that might have been updated in DRAM might not yet have been updated in Flash. Therefore, if data in DRAM is lost, the persistent data retrieved from Flash may be stale.

Ensuring Current, Persistent Data

To help meeting data consistency requirements, Schooner provides the SYNC API (which includes the sync and sync_all commands) in Schooner Memcached for persistent containers. These two commands are Schooner-specific additions to the Memcached protocol and follow the standard Memcached conventions. Both respect data integrity in the sense that successful completion guarantees that complete, not partial, objects have been stored

Upon successful return, the `sync` command guarantees that the specified object has been stored on persistent storage, while `sync_all` guarantees that all

objects that have been successfully updated before the invocation of the `sync_all` command are committed to flash once the command returns successfully.

Keep the following points in mind when using these two commands:

- Any writes that are issued after (or concurrent with) the last `sync/sync_all` may or may not be visible after a recovery.
- If the container is running in eviction mode, it is possible that a committed object won't show up after a restart due to eviction. However, the system does guarantee that if an object is recovered, it will have a value no older than the last value before the most recent successful `sync/sync_all` command.

Note that the use of the SYNC API is not mandatory for persistent containers. Even if you do not use the `sync/sync_all` commands, you should still be able to recover most of your objects since Schooner Memcached automatically syncs in the background. Typically only a small number of the most recent updates will not be recovered.

sync()

Synchronizes a specific item between DRAM cache and Flash.

Attribute	Detail
Syntax	`sync <key>\r\n`
Parameter	`<key>` the object key for the item to synchronize
Returns if successful	`SYNCED\r\n`
Errors	`NOT_FOUND\r\n` if the item with this key was not found Other error conditions as found in Memcached protocol

sync_all()

Synchronizes all unsynchronized items between DRAM cache and Flash.

Attribute	Detail
Syntax	`sync_all\r\n`
Parameters	None
Returns if successful	`SYNCED\r\n`
Errors	Error conditions as found in Memcached protocol
Notes	An updated item is synchronized only if the update is completed before the server receives the `sync_all` request.

Recovery

On recovery, data is guaranteed to be consistent with the last sync call.

Multiple Memcached Containers

In the Schooner Appliance for Memcached, multiple storage domains may be created in order to provide fine-grain control over storage resources. A storage domain is supported by a “container” that defines the policies controlling eviction (the standard Memcached eviction protocol), storage capacity, and persistence.

Containers are identified by TCP/UDP port number. By connecting to a specific port number, you can direct traffic from a Memcached client to a specific container.

A single Schooner Appliance for Memcached can support up to 8 containers.

Chapter 3. The Schooner Administrator Control Console

Starting Schooner Administrator

1. Point your browser to **http://*master_node*/**, where *master_node* is the master node's IP address (or host and domain name) on your administration subnet (**eth0**), which you entered in Step 1 of the FirstTime Wizard.
2. Log in with the username and password assigned by your administrator.

   Once you have logged in, you can change your password if you want to. If you are the administrator, you can also start to add new users Screen Layout

Screen Layout

The main Schooner Administrator window is made up of two panels: the navigation panel and the main panel, as highlighted in Figure 1.

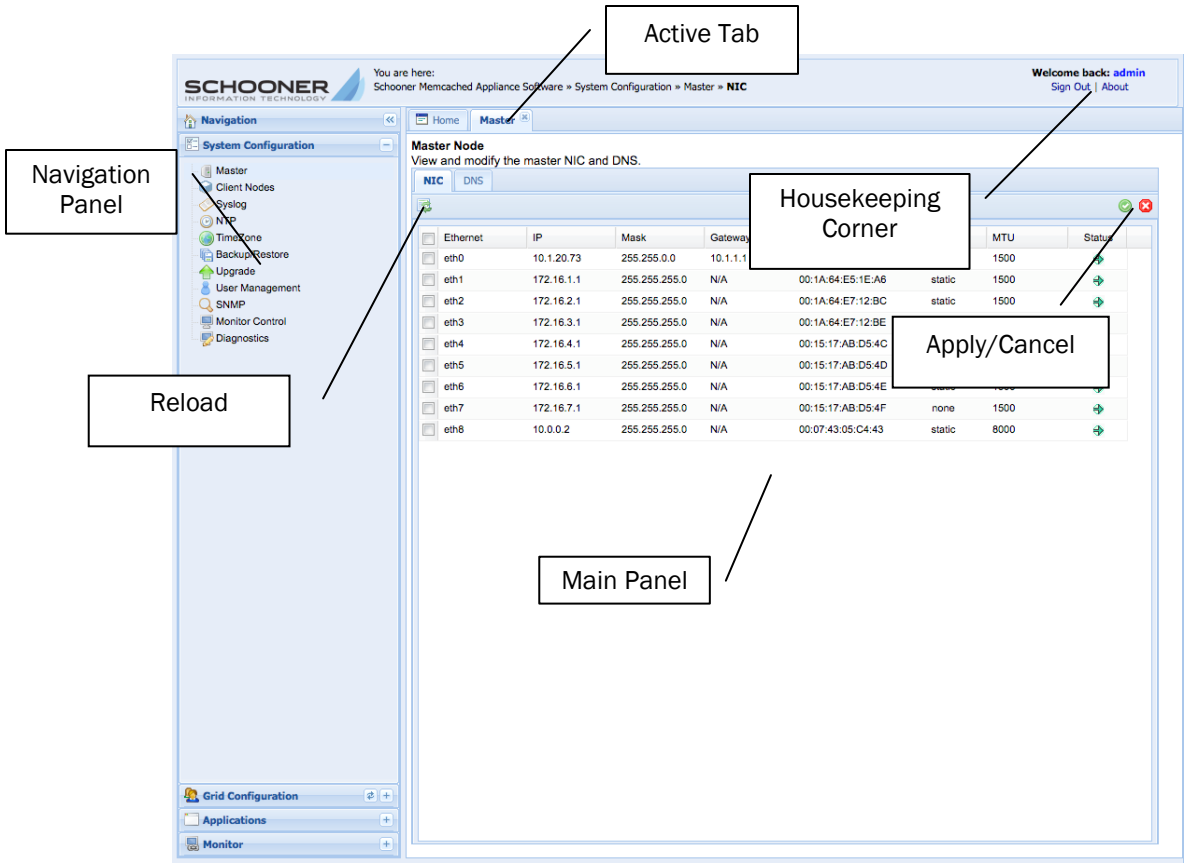

Figure 1: Schooner Administration interface

Navigation Panel

The navigation panel on the left lets you select the administration tasks you want to perform. The tasks are grouped into the following categories:

- System Configuration
- Grid Configuration
- Applications
- Monitor

You can click to expand any of these categories and select a task within that category.

Main Panel

When you first log into Schooner Administrator, you'll see the **Home** tab in the upper-left corner of the main panel, displaying the welcome screen. As you start to select entries in the navigation panel, their names will be automatically added to the top of the main panel as tabs. Each time you click an entry in the navigation panel or a tab across the top of the main panel, the main panel will refresh to display the content of the entry or tab being selected. The tabs help keep track of the screens you have visited. You can navigation from screen to screen simply by clicking the tabs or entries. The active tab which indicates the screen you are currently on is highlighted in WHITE.

Housekeeping Corner

The housekeeping corner in the upper-right corner of the Schooner Administrator GUI contains two buttons: **Sign Out** and **About.** The former allows you to log off Schooner Administrator whereas the latter provides version information about the Schooner product you are using.

Action Buttons

The following action buttons can be found on some or many of the screens:

- (**Apply**) button to commit changes made to the configuration.
- (**Cancel**) button to remove changes made to the configuration that has not yet been committed.
- (**Reload**) button causes Schooner Administrator to reload the configuration values displayed in the panel.
- (**Add**) button to add a new item.
- (**Delete**) button to remove a selected item.

Tour of Schooner Administrator

This brief tour of Schooner Administrator, organized around the activities on the navigation panel, will help you find the functionality you need.

System Configuration

From the System Configuration panel, you can view, set up, and configure your Schooner system. See Figure 2.

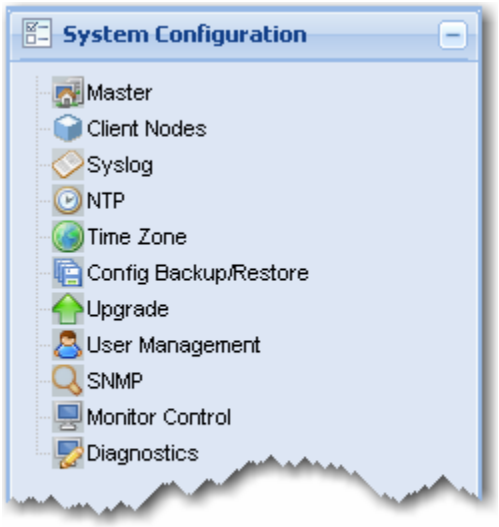

Figure 2: System Configuration menu options

- **Master** – The Schooner grid automatically chooses a master node from all of the member nodes. The Master panel shows your master node and its network interfaces, and you can also configure the DNS server from this panel.
- **Client Nodes** – This panel shows any physical nodes of your system that have not yet been added to a group. The first (master) node is added automatically, but if you add physical units to your system later, you can start, stop, identify and authenticate them here. (You would then go to Grid Configuration to add them to a group.)
- **Syslog** – This panel lets you review, search, and download the system event logs. For more on logs, see “Check System Logs” on page 31.
- **NTP** –Your Schooner system uses Network Time Protocol servers to ensure that all nodes operate on a consistent date and time. This panel shows the NTP servers in your network and allows you to add or remove them from Schooner’s synchronization list. Click to synchronize your Schooner system with the NTP server of your choice.
- **Time Zone** –This panel shows the current time zone and you can also configure it to another one as required.
- **Backup/Restore** – This panel shows your system configuration files and allows you to back up and restore them.

- **Upgrade** – This panel lets you upgrade the Schooner application software from an RPM file provided by Schooner.

  Note: You can also upgrade your Schooner appliance software using ISO files which contain updates for the Schooner disk image via a CD-ROM or PXE.
- **User Management** – This panel lets you change your password or (if you are an administrator) add or delete users from the system.
- **SNMP** – This panel lets you configure SNMP read and write access policies.
- **Monitor Control** – This panel lets you enable or disable the Web monitors hosted by the master node. These may include Ganglia, Cacti and Nagios. Alternatively, you may disable each Web monitor and connect to the Schooner appliances with your own installations of these Web monitors.
- **Diagnostics** – This panel lets you transfer system diagnostic information to a file for analysis and support.

Grid Configuration

Your *grid* comprises the set of groups that make up your system. In the navigation panel, Grid Configuration presents an expandable tree of your groups and nodes. The icons appear blue if the groups and nodes are operating normally, or gray if they are down. Click refresh to update the status. See Figure 3.

Figure 3: Grid Configuration menu options

Clicking at any level presents descriptive information about that level (if available), plus configuration information about the next level down.

After initial setup with the FirstTime Wizard, your Schooner grid will consist of one group comprising one node (the master), which is running a Memcached instance. If later you add new nodes, you would find the available nodes in the clients nodes list, and then from Grid Configuration to add them to groups.

After making changes to your grid configuration, click refresh to update the grid hierarchy in the navigation panel.

Applications

The Applications panel displays your available applications. Selecting Schooner Memcached lets you view and manage its instances. See Figure 4.

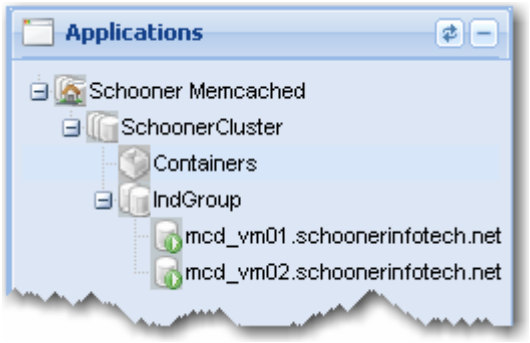

Figure 4: Grid Configuration menu options

From here you can:

- change parameters of existing instances
- start, stop, and restart existing instances
- add/delete/modify containers

Monitor

The Monitor page gives you access to the Ganglia, Cacti and Nagios open source monitoring tools. With these graphical monitors, you can watch the performance of the grid, individual Schooner Appliance for Memcached nodes and the Memcached applications. You may also connect with these interfaces using your own monitoring tools. See Figure 5.

Figure 5: Grid Configuration menu options

Getting Support

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

Chapter 4. Common Management Tasks

After initial setup with the FirstTime Wizard (FTW), each appliance provides an operational Memcached instance. You can use a Linux monitoring program or Schooner Administrator to:

- monitor and manage Memcached instances
- monitor and manage the Schooner system
- add new nodes to expand your Schooner group

Manage a Memcached Instance

You can perform basic management functions on application instances using the **Applications** menu.

1. From the navigation panel, expand the **Applications** menu.
2. Expand **Schooner Memcached** and then **SchoonerCluster**. A list of instances shows up in the main panel. See Figure 6.

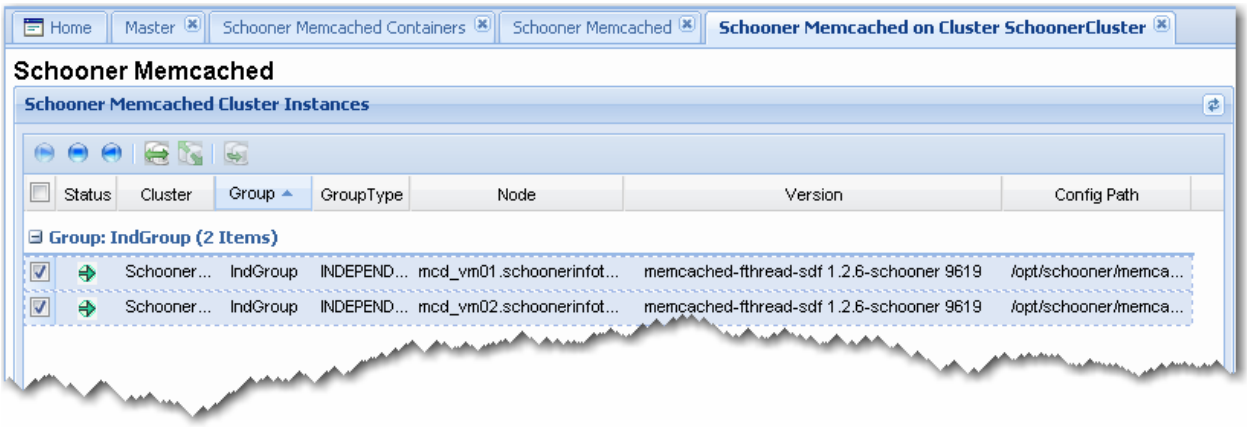

Figure 6: Managing Memcached instances

3. Select the instance or instances you want to manage.
4. Use the buttons to (start), (stop), or (restart) the checked (selected) instance or instances.

Create a Mirrored Group

Creating a mirrored group allows the two instances in the same group to share the same information. This means that operations, such as Add, Delete, Start, Stop, or Restart, performed on one container will automatically be applied to the other container in the mirrored group.

To create a mirrored group,

1 Select two nodes and then click the **(Create a Mirror Group)** button.

2 When the confirmation message pops up, click **Yes**. The Group Properties dialog box appears. See Figure 7.

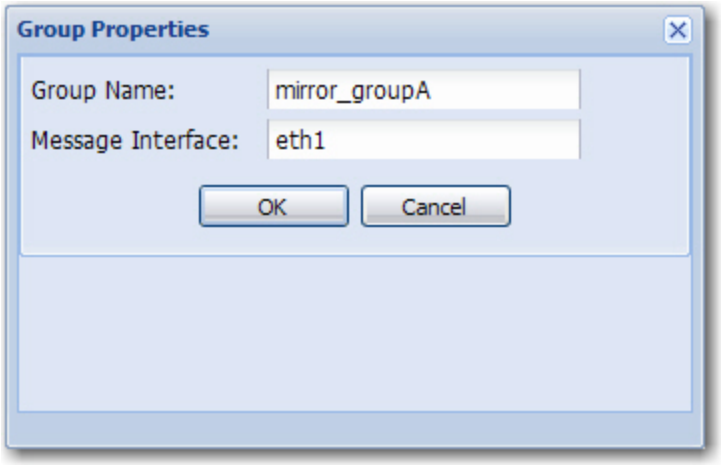

Figure 7: Specifying group properties

3 Make the desired entries as described in the following table:

Parameter	Description
Group Name	The name of the mirror group, which can only consist of alphanumeric characters, plus hyphen (-) and underscore (_). No special character or space is allowed.
Message Interface	One of the system interfaces on the Schooner appliance used for replication traffic. Any interface, including dynamic interfaces such as a bonded interface or a VLAN interface, can be used as a messaging interface. See “Messaging Interface” in Chapter 8.

4 Click **OK** when done.

5 When the Success message pops up, click **OK.** Figure 8 shows the mirror group that has just been created.

Figure 8: A mirror group

Add a Container

To add a container:

1 Click the (**Add a Container**) button. The Add a Schooner Container dialog box appears. See Figure 9.

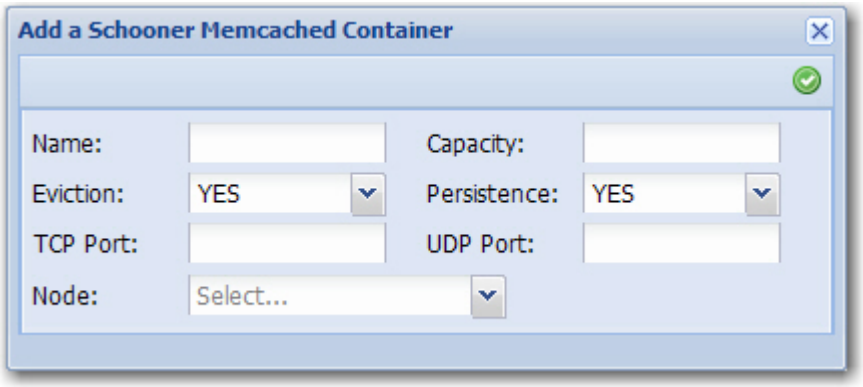

Figure 9: Creating a new container

2 Make the desired entries and/or selections as described in the table below.

Parameters	Description
Name	The name of the container, which can consist of any number of alphanumerical characters, plus hyphen (-) and underscore (_). No space or special character is allowed.
Capacity	The amount of the Flash memory (in GB) allocated to the container.
Eviction	Select either of the following: ▪ **YES** – If selected, the server will start discarding old data, when the container is full, to make room for new data. ▪ **NO** – If selected, the server will not discard any data.

	Instead, it will stop accepting new data once the container is full,
Persistence	Select either of the following: ▪ **YES** – If selected, the server will retain the data in the Flash when the system is shut down so that the same data can still be available upon system reboot. ▪ **NO** – If selected, the server will not retain any data when it shuts down. As a result, all historical data will be lost upon system reboot.
TCP Port	Enter the TCP port number of your Schooner appliance.
UDP Port	Enter the UDP port number of your Schooner appliance.
Node	Select a node for the container.

3 Click **(Apply)** when done.

Delete a Container

You can delete any container that is no longer needed.

To delete a container;

1 From the main panel, select the container of interest by clicking the corresponding radio button.

2 Click **(Stop Container)** to stop it.

3 Click the **(Remove Container)** button.

About Data Persistence

Setting Persistence to **YES** will cause data stored in Flash memory to be persistent across a reboot, power loss, or node failure. However, because your Schooner appliance uses DRAM memory as a cache front-end, some "persistent" data in Flash may actually be "stale." The Schooner-specific `sync/sync_all` commands allow applications to periodically force transient object data to flash (see page 4).

Manage Replication

In the event a mirrored node encounters a hardware failure (cannot be restarted) and must be replaced with a new node, you can create a mirrored peer using the "clone" function. The clone function copies the container configuration and data from the surviving mirrored node onto the replacement for the failed node. The following instructions show how to clone a mirrored node using the **(Clone the selected Memcached to a new node)** button.

Before cloning the node, you must first configure the replacement node's services (e.g., network, DNS, SMTP, etc.) to be the same as those on the failed node. This may be done using the configuration backup/restore facility or by running the FirstTime Wizard CLI on the replacement node. If you use the FirstTime Wizard CLI, make sure that the administration interface is configured properly for the cluster.

Note: Make sure that the hostname of the replacement node is different from the hostname of the failed node. Otherwise, the clone feature will not work.

Create Mirrored Groups

You need to create a replication (mirrored) group in order to take advantage of this feature. You can create a mirrored group using the Schooner Administrator's GUI.

To create a mirrored group:

1. From navigation panel, click the IndGroup level to display all nodes in the (independent) group.
2. From the main panel, checkmark the nodes that will form the mirrored group.
3. From the top of the main panel, click the  (**Create Mirror Group**) button. The Group Properties dialog box appears. See Figure 10.

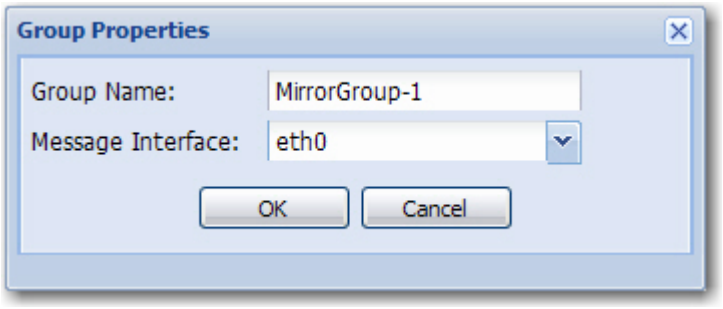

Figure 10: Creating a mirrored group

4. Specify a group name, select a messaging interface, and click **OK.** The newly created mirror group appears in the Schooner Memcached tree structure in the navigation panel. See Figure 11.

Figure 11: A mirrored group

When creating a mirrored group, make sure that you keep the following in mind:

- You must specify the two hostnames.
- You must restart Memcached on both hosts.
- The two hosts must have identical container configurations and consistent networking setup with the VIPs. Container consistency is automatically handled by the Administration GUI or CLI. When a container is created, that container will be added to both nodes with the same properties. Network consistency mandates that the corresponding interfaces on the two mirrored nodes must be

on the same subnet, i.e., eth1 on Node 1 and eth1 on Node 2 must be on the same subnet; if eth7 on Node 1 is used as the messaging interface on Node 1, then eth 7 must be used as the messaging interface on Node 2 as well.

Mirrored Group Backup Procedures

Refer to Backup and Restore Procedures in Chapter 5.

Mirrored Group Restore Procedures

Refer to Backup and Restore Procedures in Chapter 5.

Disaster Recovery (Clone a Mirrored Node)

Once the replacement node's configuration has been brought up to date you may perform the clone operation:

1. Log into the Schooner Administrator GUI on the survivor node and expand the **Applications** menu.
2. Click the mirror group containing the survivor and failed nodes.
3. From the main panel, select the failed node.
4. Click  (**Clone the selected Memcached to a new node**) button. A window will pop up, prompting you to select the recipient of the cloned data (the replacement node).
5. Locate the replacement node from the list and select it.
6. Confirm the selection when prompted. Now, the GUI will display the clone progress window.
7. When the clone operation is completed, click **OK** to finish the clone operation.

Data Recovery upon “Double Failure”

Double failure refers to a situation in which both nodes in a mirrored group go down. By design, the nodes are unable to recover persistent containers themselves even though they are able to automatically recover non-persistent containers soon after they come back to life. Human intervention is required to recover persistent containers on the nodes.

To recover persistent containers after a double-failure situation:

1 From the Schooner Administrator GUI, select either of the nodes and click the (**Start Container**) button. The Nodes List dialog appears. See Figure 12.

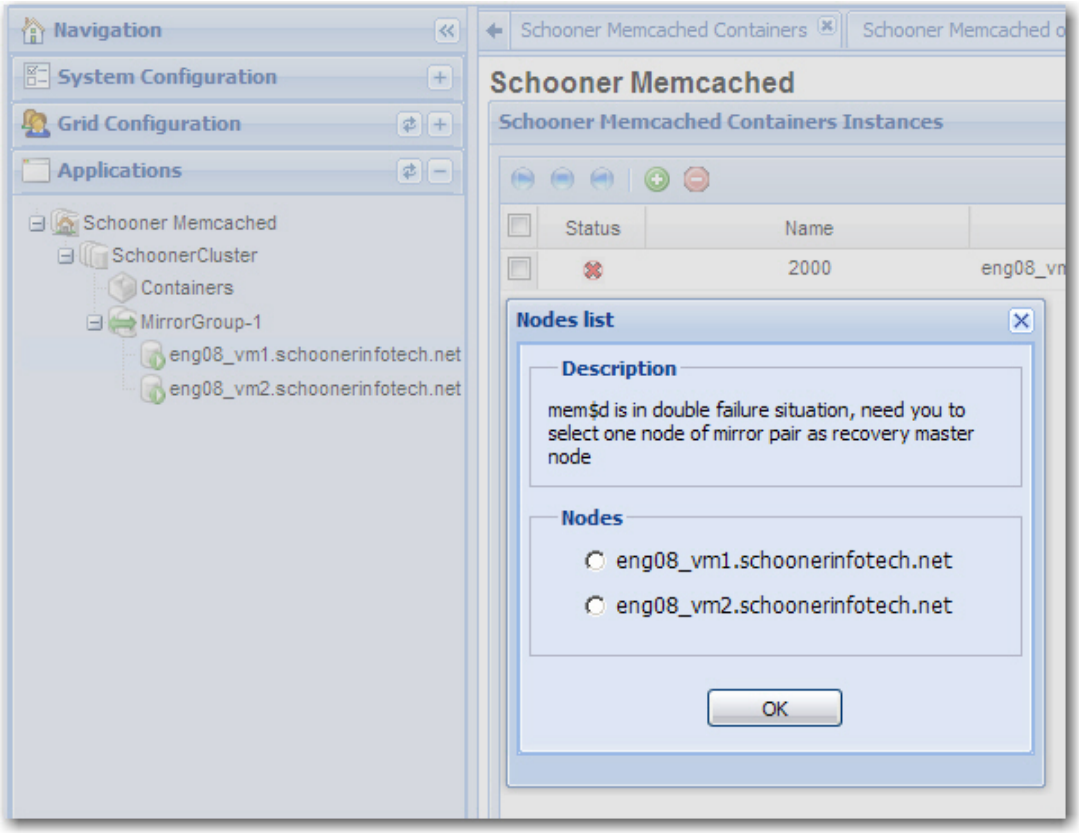

Figure 12: Recovering persistent containers after a double failure

2 Identify the node that contains the most recent data by using system mechanisms such as Syslog and make it the “recovery master node”.

3 Click **OK.**

Note: This will cause the other node go down and come back up again. That node, when it has come up, will automatically start to recover both persistent and non-persistent container data from the “recovery master node”.

4 When the recovery finishes, verify the stats on the mirrored nodes to make sure they match with each other.

Data Recovery Phases

During data recovery, the recovering node copies data from the surviving node. The length of time it takes to complete the data recovery depends on a number of variables, including the amount of data to be recovered.

During data recovery, no operation such as add, delete, start, or stop container is allowed. No such action is allowed either when a node is down.

Manage Nodes

Adding client nodes to your Schooner Memcached system can dramatically increase its capacity without adding overhead or management complexity.

The Schooner Appliance for Memcached *Installation & Setup Guide* explains how to install and connect the new nodes to your network. Once the nodes are connected, they will automatically connect with the master node. Here is how to add them to the Schooner system.

View Available Nodes

To view the available nodes in your Schooner system:

1. From the navigation panel, expand **System Configuration**.
2. Click **Client Nodes**. All physical nodes that are available will be listed in the main panel. If some new nodes were added after this menu or tab was opened, you can view them by clicking the (**Refresh**) button

View and Modify Node Properties

You can view and modify the properties of any node in a cluster.

To view and modify the properties of a node:

1. From the navigation panel, expand **Grid Configuration**.
2. Expand the Schooner Cluster of interest.
3. Double-click the node of interest. The screen refreshes, with a number of tabs shown across its top. See Figure 13.

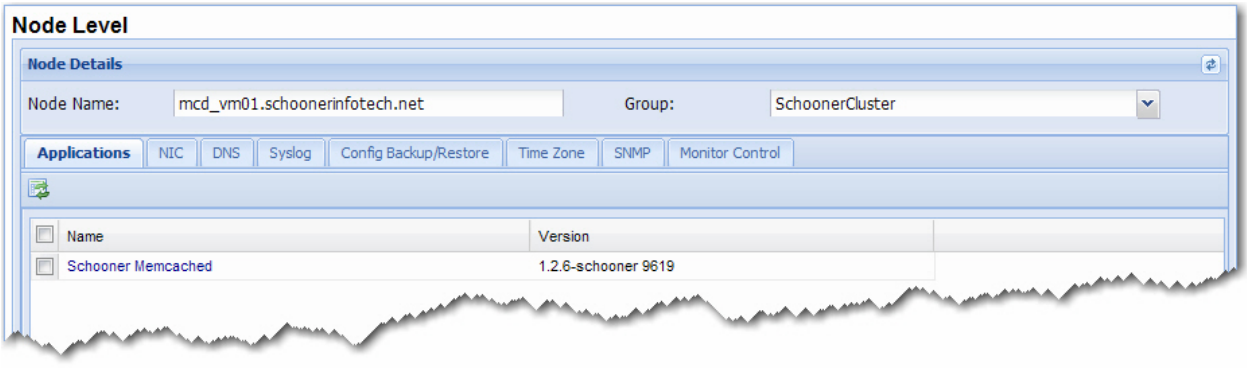

Figure 13: Viewing system parameters of a node

4. Click the tabs to review or modify the system parameters of the node.

Add New Nodes to a Group

To add nodes to a group:

1. From the navigation panel, expand **Grid Configuration.**
2. Select the group to which you want to add the nodes. The nodes of the group will appear.
3. Click  to indicate that you want to add nodes to the group. The Client Nodes List dialog will pop up.
4. Select the nodes of interest and click  to apply. This will add the selected nodes to the group.

Monitor System Status

Whenever a node, group, or instance is visible in the main panel, either of the following icons will appear in front of it, indicating its operational status, as explained below:

Status Icon	Significance
	Function properly
	Non-functioning

If your system is not functioning, you can start troubleshooting using the procedures outlined in Chapter 6: Troubleshooting.

Monitor System and Network Data

The Monitor menu allows the user to display statistics and graphs of CPU, memory, Flash, disk, and network usage at the node, group, or system level. Schooner Memcached is shipped with the following three monitoring software applications which can be used to monitor your system performance and alerts:

- **Ganglia** – for monitoring clusters
- **Cacti** – for monitoring MySQL data
- **Nagios** – for monitoring alerts

Note that you do not have to use any of these monitoring applications if you choose not to. However, if you do, you must enable these tools from inside Schooner Memcached.

Enable Monitoring Software Applications

To enable the monitoring software applications:

1 From the navigation panel, expand **System Configuration** and click **Monitor Control**. The screen refreshes, showing all the embedded monitoring software tools, as shown in Figure 14.

Figure 14: Enabling monitoring tools

2 Select the tool or tools you intend to use and click (**Apply**) in the upper-right corner of the screen.

Monitor System Data

Once a monitoring software tool has been enabled, it is very easy to use it to view system and network data on Schooner appliance.

To view data on a Schooner appliance:

1 From the navigation panel, expand **Monitor**. See Figure 15.

Figure 15: Embedded monitoring software tools

2 Click to launch a desired monitoring tool, e.g., **Ganglia.**

3 Log in using “admin” as the default user name and password. The **Ganglia** grid overview page appears. See Figure 16.

Figure 16: Ganglia grid overview screen

Since these monitoring tools differ in the type of data they monitor, it is good idea to use them all at the same time.

Use an External Monitoring Program

You can also use a Linux monitoring program like Nagios, Cacti, or Ganglia to monitor the hosts and services of your Schooner Appliance for Memcached system. Such a tool can provide centralized monitoring of your entire network.

You can integrate Schooner Memcached with a monitoring program of your choice by adding the Schooner Appliance for Memcached system’s listening address (the IP address on your application or the eth2 interface) to the monitoring program’s server list or configuration file.

Chapter 5. Common Administrative Tasks

Upgrade the System

Upgrade packages are periodically posted on the IBM support Web site at http://www.ibm.com/systems/support/. Be sure to register with IBM for notifications.

Upgrade packages come in two forms:

- RPM files - contain updates for Schooner applications and services.
- ISO files - contain updates for the Schooner disk image (via CD or PXE).

Upgrading the Memcached Appliance will require a reboot and interruption of service. Your configuration files will not be disturbed.

To upgrade the Schooner Appliance for Memcached applications and services:

1. Download the update files (.rpm) from the IBM support Web site (http://www.ibm.com/systems/support/).
2. From the navigation panel, select **System Configuration>Upgrade.**
3. From the main panel, check the **RPM Upgrade** radio button and click **Next.**
4. Click on the  button to load the RPM update file into the Schooner Administrators RPM repository.
5. Select the RPM update file and click the  button to apply the update.

Note: The system will automatically restart each time Memcached is upgraded.

Rolling Upgrade

For mirrored groups, upgrade of mirrored nodes can be done in the so-called “rolling upgrade” fashion, which means that the upgrade can be performed without an outage of Memcached.

When performing upgrade on a mirrored pair, the user first needs to shut down one node while leaving the other running. This will cause the surviving node (the one left running) to assume that the “down” node failed and start take over client traffic to the down node. the failover mechanism comes into play.

Back Up Configuration Files

As a best practice, back up your configuration files regularly.

To back up configuration files:

1 On the navigation panel, expand **System Configuration**.

2 Click **Config Backup/Restore**. All configuration files appear in the main panel. See Figure 17.

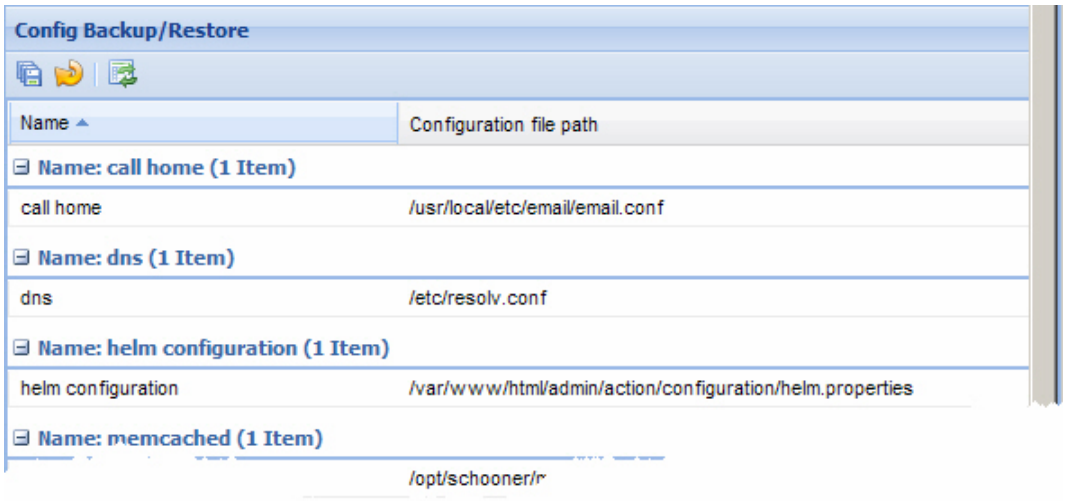

Figure 17: A list of configuration files

3 Click the (**Backup all configs**) button. The screen refreshes showing the following options:

- back up on the local node
- back up by download as a file
- back up on another node or destination

Note: There may be additional prompts to specify the exact backup location.

4 Click **Next**.

The files are backed up to the destination you specified. A message will pop up, showing where the files are stored.

Notes:

- On a Windows system, the download may be interrupted by a security warning. Simply ignore it and accept the file download. You may have to return to **Config Backup/Restore** and request the download again.
- To avoid the security warning entirely, you may want to change your Internet security options. In your browser, choose **Tools>Internet Options**. Select the **Security** tab, click **Custom** level, and enable both **File download** and **Automatic prompting for file downloads**.

Restore Configuration Files

If your Schooner configuration is corrupted or otherwise lost, you can restore your configuration files from backup copies.

1 On the navigation panel, expand **System Configuration**.

2 Click **Backup/Restore**.

3 Click the (**Restore all config**) button. The following options appear in the main panel:

- Restore from the local host
- Restore by uploading a file
- Restore by copying the file from another node

4 Select an option and click **Next.**

Note: The configuration files are restored from the destination you specified. Depending on the option you chose, you may have to go through some additional steps.

5 Click **Finish** when done.

Add or Delete a User

Schooner appliance allows the network administrator to add or delete Schooner Administrator users at any time.

Add a User

To add or delete a user:

1 From the navigation panel, expand **System Configuration**.

2 Click **User Management**.

3 From the main panel, click (**Add a user**) to add a user. The Add a user dialog box appears. See Figure 18.

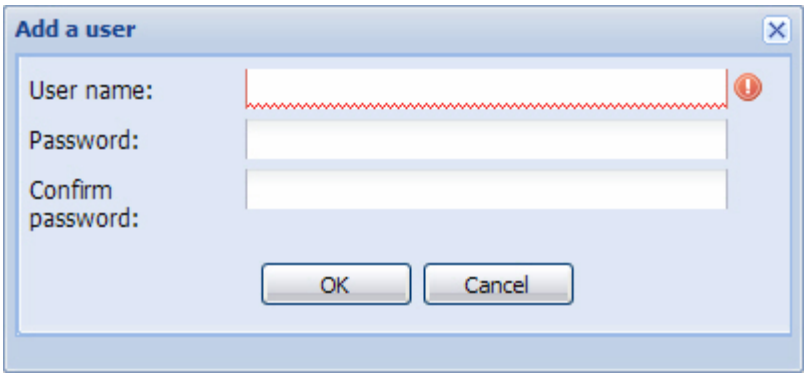

Figure 18: Create a new user account

4 Specify the user name and password and confirm the password.

Note: The user name can be 1 to 32 characters long, containing only alphanumeric characters and the underscore (_).

5 Click **OK.**

6 When a confirmation message appears, click **OK.**

Delete a User

To delete a user:

1. From the main panel, select a user account and then click (Delete a user).
2. When the confirmation message appears, click OK.

Change Your Password

1. From the navigation panel, expand **System Configuration** and click **User Management**.
2. From the main panel, select the desired entry and click  **(Change password).** The Change password dialog box appears.
3. Enter the old password and then the new password.
4. Confirm the new password by entering it for a second time.
5. When the confirmation message appears, click **OK**.

Container Backup/Restore

Overview

Schooner Appliance for Memcached supports the backup and restore of persistent container data:

- A container data backup/restore utility is provided to save/restore the contents of Schooner Memcached persistent containers.
- The container contents are retrieved from the SSD storage and placed on an HDD attached to the Schooner Appliance for Memcached for later transfer to a remote storage device; the backup can also transfer the contents of the container directly to a remote storage.

Backup

Schooner Appliance for Memcached supports two types of backup operations: full backup and incremental backup. A full backup is a logical copy of all objects in a container. An incremental backup is a logical copy of objects in a container that are new or have changed since the previous backup, as well as a logical representation of deleted objects. A full backup is taken to start a new backup "series", which contains the full backup, plus zero or more incremental backups. There is no limit on the number of incremental backups that may be taken in a backup series (though there are practical considerations).

Restore

Restoring a backup is the process of replaying backup streams to the server. A backup can be restored to any container. The target container must already exist at the time of the backup because the restore process does not create a container. The user must also make sure that there is enough space in the container to hold all the restored objects.

Users can restore to any arbitrary backup within a series, but all backups within the selected series are restored as part of a single restore operation.

Backup Commands

```
Usage:
  backup –hv | [-d] -l | -m -s [-d] [-t]

Options:
  -m, --mode=<full|incr>
  -s, --servers=<addr:port>
  -d, --dir=     default:/schooner/backup
  -t, --threads= default: 8 (min=1,max=8)
  -l, --list      lists backup files specified by --dir
  -h, --help      displays help
  -v, --version   displays version

Examples:
    backup -m full -s 127.0.0.1:11211
    backup -m full -s 127.0.0.1:11211 -d /schooner/backup
    backup -d /schooner/backup -l
    backup -l
```

To do a full backup:

```
backup --servers=10.1.20.82:11211 --mode=full
```

To do an incremental backup:

```
backup --servers=10.1.20.82:11211 --mode=incr
```

Restore Commands

```
Usage:
    restore -hv | [-d] -l | -s -d

Options:
    -s, --servers=<addr:port>
    -d, --dir=
    -l, --list       lists backup files in --dir, default dir:
/schooner/backup
    -h, --help       displays help
    -v, --version    displays version

Examples:
    restore -s 127.0.0.1:11211 -d
/schooner/backup/Schooner/127.0.0.1:11211:1/
    restore -d /schooner/backup -l
    restore -l
```

To restore the most recent backup series (full+incremental) for a container:

```
restore --servers=10.1.20.82:11211 dir=backups/Foo:11211:1
```

To restore a particular backup series (full+incremental) for a container:

```
restore --servers=10.1.20.87:44444 --dir=backups/Foo:11211:1/BACKUP-1
```

List Commands

The --list option on either the backup or restore commands lists the contents of the directory specified by the --dir option recursively.

Below are examples of how to use the directory structure above:

```
backup|restore --list --dir=backups

Cluster       Container               Backup  Type  Date
FooBaz        Foo:11211:1             1.0     FULL  2009-07-17 11:13
FooBaz        Foo:11211:1             1.1     INCR  2009-07-18 09:05
FooBaz        Foo:11211:1             2.0     FULL  2009-07-17 10:55
FooBaz        Foo:11211:1             2.1     INCR  2009-07-18 08:23
FooBaz        Foo:11211:1             2.2     INCR  2009-07-18 12:46
FooBaz        Bar:11212:1             1.0     FULL  2009-07-18 15:09
FooGoo        Boo:11222:1             1.0     FULL  2009-07-18 18:30

backup|restore --list --dir=backups/FooGoo

Cluster       Container               Backup  Type  Date
FooGoo        Boo:11222:1             1.0     FULL  2009-07-18 18:30

backup|restore --list --dir=backups/FooBaz/Bar:11212:1

Cluster       Container               Backup  Type  Date
FooBaz        Bar:11212:1             1.0     FULL  2009-07-18 15:09

backup|restore --list --dir=backups/FooBaz/Foo:11211:1/BACKUP-1

Cluster       Container               Backup  Type  Date
FooBaz        Foo:11211:1             1.0     FULL  2009-07-17 11:13
FooBaz        Foo:11211:1             1.1     INCR  2009-07-18 09:05

backup|restore    --list    --dir=backups/FooBaz/Foo:11211:1/BACKUP-1/FULL-0

Cluster       Container               Backup  Type  Date
FooBaz        Bar:11212:1             1.0     FULL  2009-07-17 11:13
```

The cluster column is a list of subdirectories under the top-level directory; the container column is a list of subdirectories in the cluster subdirectory; the backup column corresponds to directory names from the next two levels of the directory structure; and the date column is the backup time from any of the backup files within the incremental (bottom-level) directory, formatted in the user's locale.

Backup and Restore Procedures

Backup operations can only be performed on persistent containers only. Depending on your need, you can perform either a full backup or an incremental backup. If you plan to do a combination of full and incremental backups, you are advised to start with a full backup first, followed by one or more incremental backups, depending on your situation. You can start a backup operation even when the client machines are running.

Restore can only be performed to an empty, persistent container. It does not have to be the same container from which the data being restored was backed up. The following procedure restores the most recent full backup and all subsequent incremental backups.

Backup and/or restore operations can be performed on either a single node or a mirrored one. The following paragraphs provide the basic instructions for each scenario.

Back Up a Single Node

To back up a container on a single node:

1. Make sure to verify that your backup hard drives are mounted properly and that there is enough disk space available on them.
2. Start a full backup of the persistent container at port 11212 by executing:

```
opt/schooner/memcached/bin/backup --mode=full
--servers=127.0.0.1:11212
--dir=/schooner/backup
```

3. Start an incremental backup (if necessary) by executing:

```
/opt/schooner/memcached/bin/backup --mode=incr
--servers=127.0.0.1:11212
--dir=/schooner/backup
```

Restore a Single Node

To restore a container on a single node:

1. Make sure that no clients are accessing the container.
2. Stop the container (on both nodes) by executing:

```
# memcli stop container --tcp_port=11212
```

3. Format the container by executing:

```
# memcli format container --tcp_port=11212
```

4. Start the container by executing:

```
# memcli start container --tcp_port=11212
```

5. Restore the container by executing:

```
# /opt/schooner/memcached/bin/restore
--servers=127.0.0.1:11212
--dir=/schooner/backup/Schooner/Data:11212:10
```

   Note: The string “Data:11212:10” is the unique identifier for the container.

6. Verify the restore, e.g., use the memcached “stats” command to verify the number of objects in the container matches the number of restored objects.

Back Up a Mirrored Group

To back up a container on a mirrored group:

1. Make sure to verify that your backup hard drives are mounted properly and that there is enough disk space available on them.
2. Start a full backup of the persistent container at port 11212 by executing:

```
opt/schooner/memcached/bin/backup --mode=full
```

```
--servers=127.0.0.1:11212
--dir=/schooner/backup
```

3 Start an incremental backup (if necessary) by executing:

```
/opt/schooner/memcached/bin/backup --mode=incr
--servers=127.0.0.1:11212
--dir=/schooner/backup
```

Restore a Mirrored Group

To restore a container on a mirrored group:

1 Make sure that no clients are accessing the container.

2 Stop the container (on both nodes) by executing:

```
# memcli stop container --tcp_port=11212
```

3 Format the container by executing:

```
# memcli format container --tcp_port=11212
```

4 Start the container by executing:

```
# memcli start container --tcp_port=11212
```

5 Stop Node B by executing:

```
# memcli stop node --node_name=xyz.domain.net
```

6 Restore the container on Node A by executing:

```
# /opt/schooner/memcached/bin/restore
--servers=127.0.0.1:11212
--dir=/schooner/backup/Schooner/Data:11212:10
```

Note: The string “Data:11212:10” is the unique identifier for the container.

7 Verify the restore, e.g., use the memcached “stats” command to verify the number of objects in the container matches the number of restored objects.

8 Bring Node B up by executing:

```
# memcli start node --node_name=xyz.domain.net
```

9 When the recovery finishes, verify the stats on the mirrored nodes to make sure they match with each other.

Chapter 6. Troubleshooting

This chapter suggests what to do if you experience problems. It includes:

- Problem Indicators: How to tell if a problem exists
- General Troubleshooting Tips applicable to multiple situations

If none of these suggestions solves your problem, log into the IBM support Web site (http://www.ibm.com/systems/support/). There you can browse for additional support information and, if necessary, open a case. You will receive help on diagnosing the problem and finding a solution.

Problem Indicators

Status Red

A red status icon indicates that the entity is not operational or not recognized.

For a client node in a group configuration, it will be red if the node is down.

General Troubleshooting Tips

Refresh and Restart

If an application, node, or group is not working properly, or its status is red, first click to refresh its status. It may turn green without further intervention.

If that doesn't work, try the following:

- Click the start icon to start it (if it is stopped), or
- Click the restart icon to reboot it.

Check System Logs

You can often trace system problems by inspecting the system logs.

1. From the navigation panel, expand **System Configuration.**
2. Click **Syslog.** The Syslog screen appears in the main panel. See Figure 19.

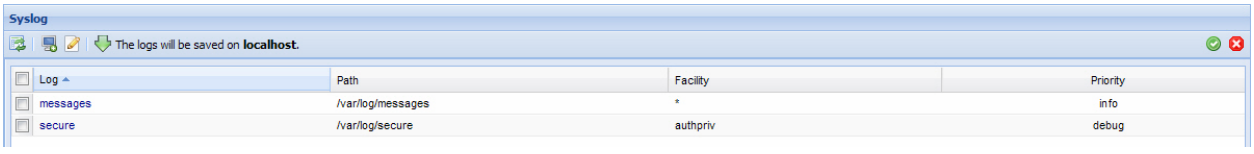

Figure 19: View Syslog messages

3. Click the log you wish to examine. You can browse the log or use your browser's search facility to look for a particular kind of events.

   Note: The system generates two kinds of logs: Messages and Secure. Messages (logs) contain information about general system events, whereas Secure (logs) contains information about system access and connection events.

4. To download logs onto your local workstation, checkmark the ones you want and click the download icon. On a Windows system, if you see a security warning, go ahead, and accept the file download.
5. To change the log file storage location to another node within the Schooner system, click, make the desired changes, and click to download the log files to that destination.
6. To change the log file storage location to another location outside the Schooner system, click, make the desired changes, and click to download the log files to that destination.

Restore Default Values

Some problems, especially slow performance, are caused by selecting NIC, node, or instance properties that are not optimal. Schooner recommends the following:

- Accept default values when possible.
- Keep a record of the default values that you've changed.

If you experience problems, try restoring the default values that you've changed. You will find these properties:

- For NICs, in **System Configuration>Masters>NIC.**
- For nodes, in **Grid Configuration** under each node.
- For Memcached instances, in **Applications>Schooner Memcached.**

Reboot

If all else fails, try rebooting using the following instructions:

1. Using a terminal program like PuTTY, connect to the appliance using its eth0 host name or IP address.
2. Log in as "admin".
3. Halt any Memcached traffic:
   monit stop memcached
4. Reboot the appliance:
   `reboot`

The system will reboot and resume operation in approximately 3 to 5 minutes, when you can again log into Schooner Administrator.

Diagnostics

If Schooner Support requests system information, use the **Diagnostics** interface to save logs and traces to be used for troubleshooting:

To perform system diagnostics:

1 From the navigation panel, expand **Diagnostics.** The Diagnostic screen appears. See Figure 20.

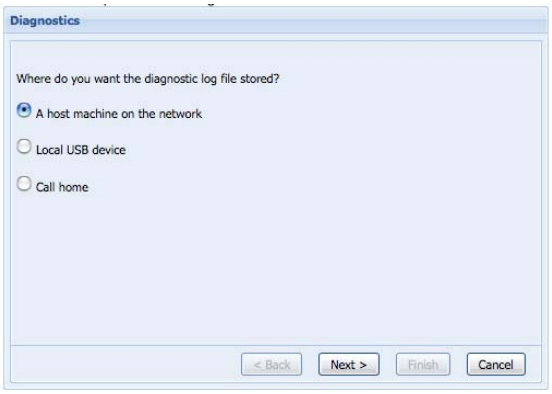

Figure 20: Specify Diagnostic log file storage location

2 To store log information on a network server, select “A host machine on the network” and click **Next.** See Figure 21.

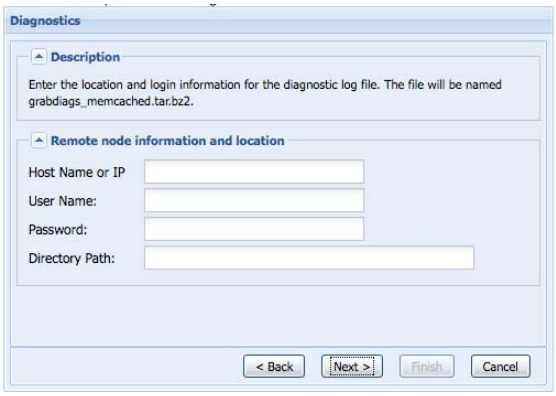

Figure 21: Create Diagnostic file directory

3 Make the required entries and/or (i.e., Host Name, User Name, Password, and Directory Path) and click **Next** and then **Finish.**

Note: The Directory Path should be /tmp.

4 To send log files to a USB device on a local machine, select **Local USB device** and click **Next** and then **Finish.**

Note: Make sure to have the USB device installed (inserted) in the USB drive before you select this option.

5 To send the log files directly to IBM Support (recommended), select **Call home** and click **Next** and then **Finish.**

Chapter 7. Schooner Memcached Command Line Interface

Schooner Memcached CLI

If you wish to use your existing administration system instead of the Schooner Administrator, you must initialize each Schooner Appliance for Memcached using the Schooner FirstTime Wizard CLI as documented in the *Schooner Appliance for Memcached - Installation & Setup Guide.*

Once the appliance has been configured, you must start the Schooner Memcached application using the Schooner Appliance for Memcached CLI. You may also change the configuration of the Schooner Memcached application though the default configuration is compatible with the default memcached configuration.

Basic Operational Procedures

To log into a node through CLI:

1. Using a terminal program like **PuTTY** (if you are using a Windows operating system) or SSH, connect to the appliance using its **eth0** host name or IP address.
2. Log into the appliance as user **admin** with default password **admin.**
3. Run the CLI:
   `/opt/schooner/memcached/bin/scripts/memcli`
4. You will see the list of commands available to configure and control the Schooner Memcached application:

```
Usage:  memcli <command> [<instance|group|container>
[option...]]
        memcli report [--core-file=<path>]
   Valid commands are:
      start          Start memcached
      stop           Stop memcached
      restart        Restart memcached
      status         Print memcached status
      list           List the instances, groups or containers
      add            Add a group or a container
      delete         Delete a group or a container
      report         Report the system status to Schooner
      init           Create independent node property file
      config         Configure messaging interface
      help           List all the commands
```

```
help <command> Get help on a property file editing
command
```

5 To get help on how to list existing instance, groups, or containers, execute:

```
memcli list

Usage:
   1. memcli list group
   2. memcli list container
         valid options are:
               -- node_name=<node name>
   3. memcli list instance
```

Then use the proper command to list the existing groups, container, or instances. For example, to list the existing containers, execute:

```
            memcli list container


                Existing Containers
Id  TCP   UDP  VIP  IF  TCP  UDP  Evict  Persist  Capacity
0   11211 11211 yes               yes     no       0
```

The example shows the default container configuration; a single container attached to the standard memcached TCP/UDP ports with eviction enabled, persistence disabled using the full capacity of the Schooner Appliance for Memcached storage system.

6 To get help on how to add a group or container, execute:

```
memcli add

The "add" command requires you to specify group or container.

Usage:
  1. memcli add group
      Valid options are:
          --cluster_name=<cluster name>
          --node_name=<node1(,node2,node3...)>
          --group_type=<group type>, MIRRORED or N+1
          --group_name=<group name>
  2. memcli add container
      Valid options are:
          --node_name=<node name>
          --tcp_port=<TCP port number>
          --udp_port=<UDP port number>
```

```
--capacity_gb=<gigabytes>
--eviction=<yes | no>
--persistent=<yes | no>
--container_name=<container name>
```

Then use the proper command to add a group or container.

7 To add a container, execute:

```
Memcli add container --node_name=engin1 --tcp_port=11212
--udp_port=11212 --capacity_gb=100 --eviction=no
--persistent=yes --container_name=Schooner1
```

All values must be specified. By default, both eviction and persistence are enabled. The capacity of the container to be created must not exceed the total capacity of the SSD drive on the appliance. If the container is created at the Container level in the Schooner Memcached tree structure (in the navigation panel), then you must also select a node as well.

The Schooner Memcached application must be restarted whenever a container has been added.

8 To restart the Schooner application, execute:

```
memcli restart
```

9 To delete a container, execute:

```
memcli delete container --node_name=engin1 --tcp_port=11212
```

The Schooner Memcached application must be restarted whenever a container has been deleted.

To restart the Schooner application, execute:

```
memcli restart
```

10 To get help on how to start a node or container, execute:

```
memcli start memcached
The "start" command requires to specify container or node.
Usage:
    1. memcli start container
       Valid options are:
            --node_name=<node name>
            --tcp_port=<TCP port number>

    2. memcli start node
        Valid options are:
          --node_name=<node name>
```

Then use the proper command to start a node or container.

11 To get help on how to stop a node or container, execute:

```
memcli stop memcached

The "stop" command requires to specify container or node.
Usage:
    1. memcli stop container
        Valid options are:
            --node_name=<node name>
            --tcp_port=<TCP port number>

    2. memcli stop node
        Valid options are:
            --node_name=<node name>
```

12 To get help on how to restart a node or container, execute:

```
memcli restart memcached

The "restart" command requires to specify node.
Usage:
    1. memcli restart node
        Valid options are:
            --node_name=<node name>
```

13 To print Memcached status, execute:

```
sudo memcli status
```

14 To create independent node property file, execute:

```
sudo memcli init
```

This command configures the current node as an independent node with one container. It also specifies the default values of the container, such as its capacity, eviction, persistence, and TCP or UDP port, etc.

15 To get help on how to configure the messaging interface, execute:

```
sudo memcli config

The "config" command requires to specify MSGIF.
Usage:
    1. memcli config MSGIF
```

```
Valid options are:
    --cluster_name=<cluster name>
    --node_name=<node name>
    --interface=<interface>
```

This command configures the messaging interface. You need to run this command whenever you change the messaging interface on a node.

16 To view the entire list of Help commands, execute:

```
memcli help
```

17 To get help on how to use a specific help command, execute:

```
memcli help <command>
```

For instance, to get help on how to use “add” command, you need to execute:

```
memcli help add
Usage:
    1. memcli add group
        Valid options are:
            --cluster_name=<cluster name>
            --node_name=<node1(,node2,node3...)>
            --group_type=<group type>, MIRRORED or N+1
            --group_name=<group name>

    2. memcli add container
        Valid options are:
            --node_name=<node name>
            --tcp_port=<TCP port number>
            --udp_port=<UDP port number>
            --capacity_gb=<gigabytes>
            --eviction=<yes | no>
            --persistent=<yes | no>
            --container_name=<container name>
```

Chapter 8. Mirrored Group Network Setup Requirements

Mirrored Group Failover

When a node fails, the Virtual IP of the failed node would automatically be configured on the surviving node's corresponding interfaces. eth0 VIP of the failed node would be configured on eth0 of the surviving node, eth1 VIP of the failed node would be configured on eth1 of the surviving node, bond0 VIP of the failed node would be configured on bond0 of the surviving node and so on

Network Interface IP/VIP/Subnet

- Each configured network interface of a node must have a unique IP address.
- The VIP must be a valid IP address on the associated interface subnet.
  Example: If eth0's subnet is 10.1.20.0/24, then the VIP must be between 10.1.20.1 to 10.1.20.254
- Similar interfaces on both the nodes should be on the same subnet.
  If a VIP is configured on eth0, the subnet of eth0 of Node 1 and Node 2 should be the same as shown in Figure 22 below:

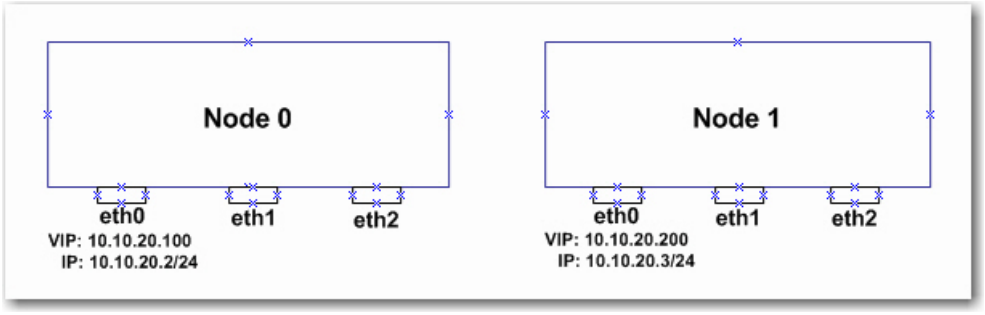

Figure 22

Note that the following configuration is NOT supported because eth0 of Node 0 and eth0 of Node 1 are configured on two different networks. See Fgiure 23.

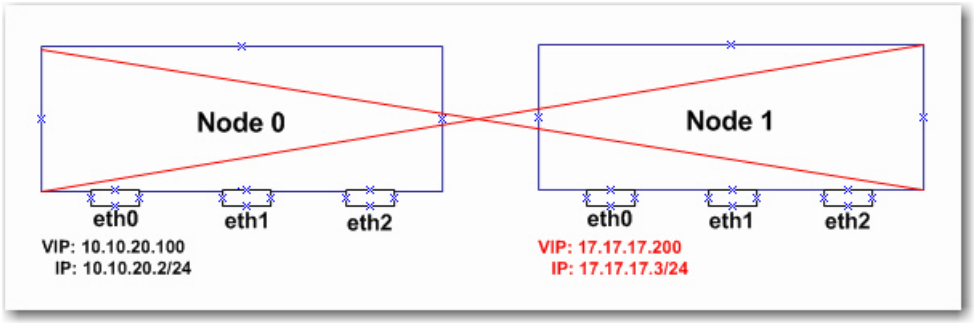

Figure 23

VLAN

If **VLAN tagging** is enabled on an interface on Node 0, then the corresponding interface on Node 1 should be configured with the same VLAN TAG as shown in Figure 24 below.

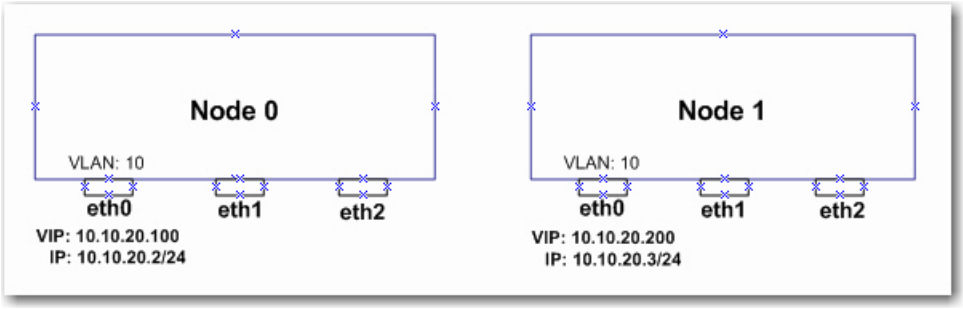

Figure 24

Note that the following configuration is **NOT** supported because eth0 on Node 0 and eth0 on Node 1 have different VLAN TAG IDs. See Figure 25.

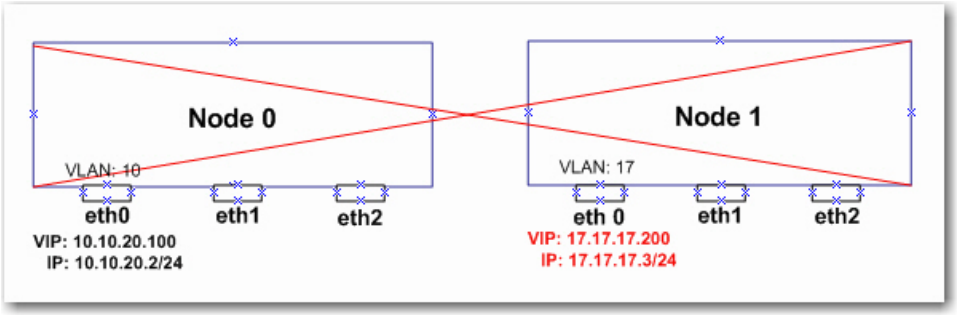

Figure 25

Bonding

If **bonding** is configured on Node 0, the bonding interface names should be the same on both nodes as shown in Figure 26 below.

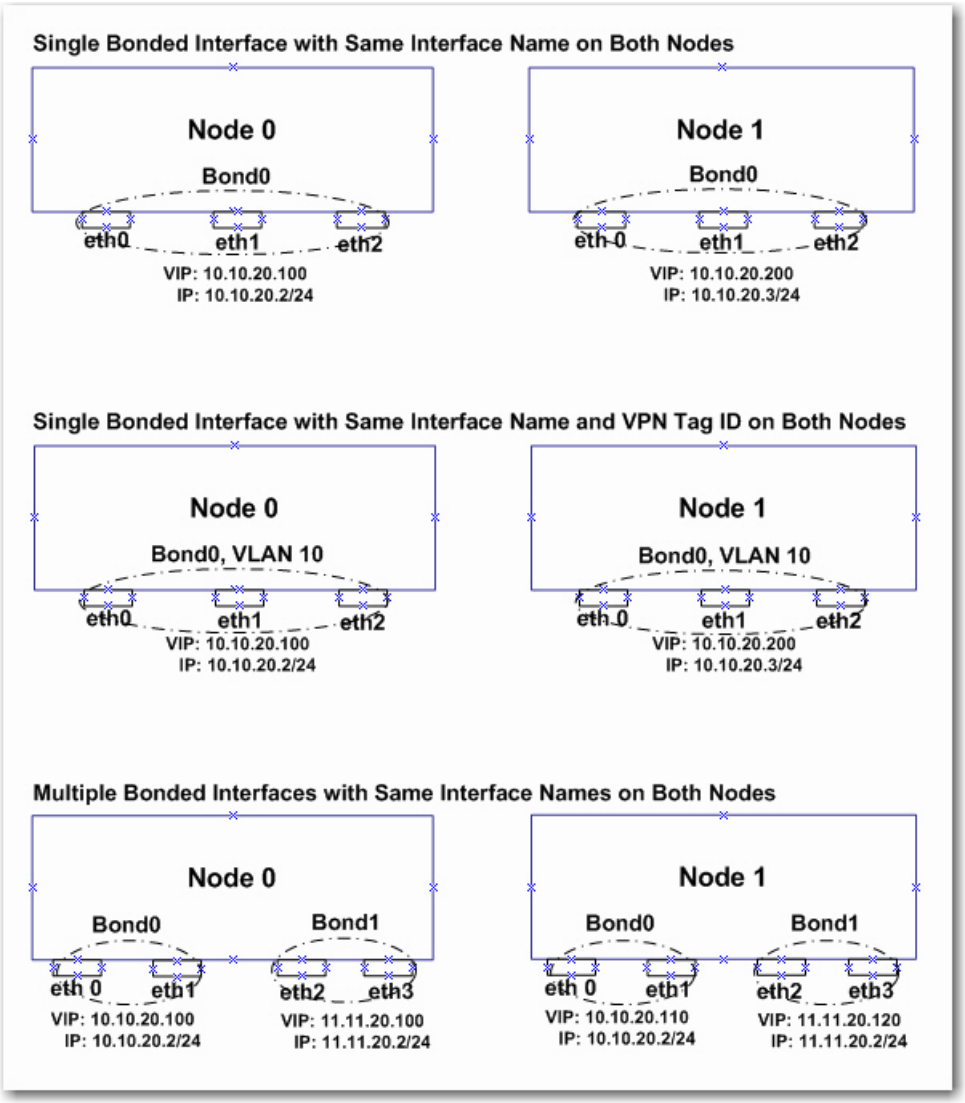

Figure 26

Note that the following configurations (as shown in Figure 27) are NOT supported:

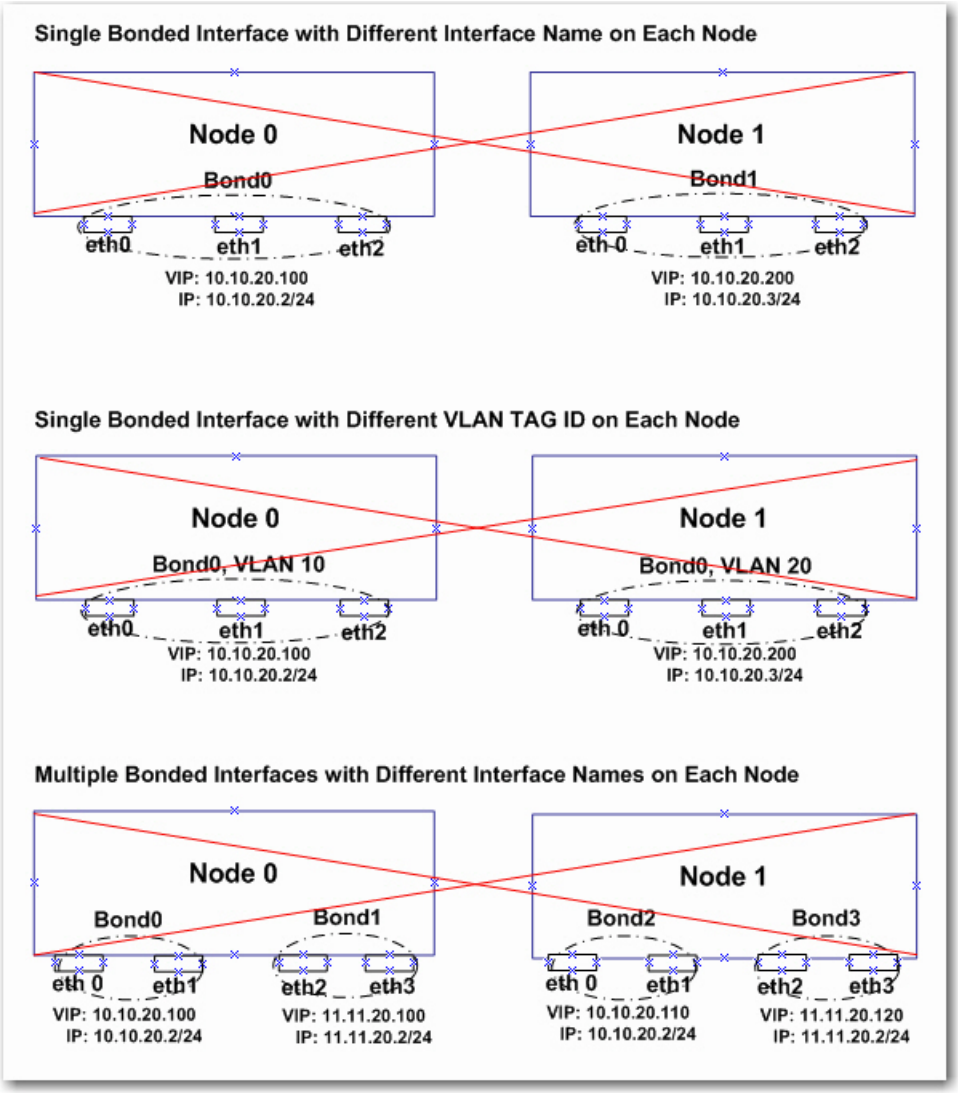

Figure 27

Messaging Interface

The messaging interface must be configured with a valid interface name after interface configuration is completed.

If a bonded interface is to be configured as the messaging interface, the corresponding bond interface name (e.g., Bond0) must be configured as the messaging interface.

If a VLAN-enabled interface (either ethx interface or bondx interface) is to be configured as the messaging interface, the corresponding VLAN interface name must be specified as the messaging interface on the memcli/gui.

The VLAN interface name is derived as vlan<tag>.

Example 1:

If eth8 is configured with VLAN 100 and eth8 needs to be the messaging interface, then vlan100 must be specified as the messing interface, not eth8.

Example 2

If bond0 is configured with the VLAN 10 and the bonded interface needs to be the messaging interface, then interface vlan10 must be specified as the messing interface, not bond0.

Split Brain Handler

Schooner Memcached may run into a split-brain situation when it is set in mirrored mode. It occurs when the messaging interface of a mirrored node is down while the client interfaces are still up and running. The situation, if left unattended, will cause unpredictable client connection behavior.

With the split-brain handler, Schooner Memcached can actively monitor for spit-brain conditions on the network and automatically stop Memcached on one node the moment a split-brain condition has been detected. Once stopped, the node will remain inactive until the system administrator manually restarts it, either from the Schooner Memcached GUI or CLI.

It is important to note that, in order for the split-brain handler to function, the two nodes in a mirrored group must be able to communicate with each other through client interfaces.