QuickStart Guide

Schooner Appliance for Memcached

Version 2.0

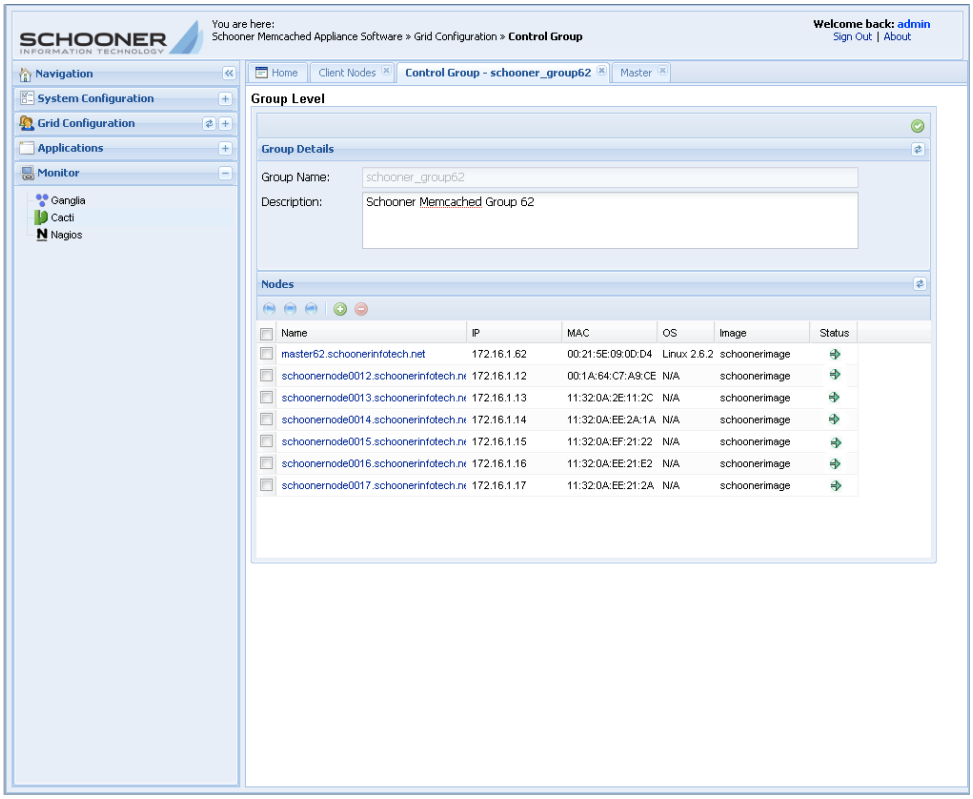

SCHOONER
INFORMATION TECHNOLOGY

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-QG-002
.

Schooner Appliance for Memcached™ - QuickStart 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

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
- 8MB cache (max)
- 64GB 1333MHz, fully buffered, DDR-3 RDIMM memory
- 12 hot-swappable drive bays
- 2 SAS hard disk drive
- 8 SATA solid state drives (512GB)
- 1 10-Gb Ethernet port
- 8 1-Gb Ethernet ports

Software Overview

Memcached

- Memcached 1.2.5
- Schooner support for:
    - Storage containers
        - Persistence
        - Replication
        - Auto Failover/Failback
        - Backup/Restore
    - Simple administration
        - CLI
        - Cluster-aware GUI
- Memcached web site: http://www.danga.com/memcached/

Supported Third-Party Monitoring Software

Schooner supports the following 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

Cacti (http://www.cacti.net/)

- Open source monitoring graph package
- SNMP, IOSTAT, MySQL templates

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 which offer:

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

Configuring the Appliance

To perform the following configuration tasks, you must open a terminal window with a Telnet/SSH client, like PuTTY. You 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 for at least 5 minutes for the system to boot.

Using a Terminal Window

- Log in 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
```

- Log in 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
```

- Log in to the IMM GUI.
- Click the Power/Restart link.
- Click the desired power command link:

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

Memcached Configuration and Control

The GUI provides buttons for these operations.

The CLI-equivalent commands are as follows:

- Log in as “admin”

Start, Stop and Restart Schooner Memcached

- To control the node you are logged into:

```
# sudo memcli start | stop | restart
```

- To control a remote node:

```
# sudo memcli start | stop | restart node
        --node name=<hostname>
```

Container Configuration and Control

Containers are storage entities supported by Schooner Memcached to give you fine-grained control over memory and flash resources. Memcached clients access a container via its TCP/UDP port. You may additionally assign a name to each container for ease of maintenance.

Once you have created a container, it will remain inactive until you start it. Similarly, you should stop the container before you delete it.

Add Container

```
# sudo memcli add container
        --node_name=<hostname>
        --container_name=<name>
        --tcp_port=<TCP port number>
        --udp_port=<UDP port number>
        --capacity_gb=<gigabytes>
        --eviction=<yes | no>
        --persistent=<yes | no>
```

Delete Container

```
# sudo memcli delete container
        --node_name=<hostname>
        --tcp_port=<TCP port number>
```

Start Container

```
# sudo memcli start container
        --node_name=<hostname>
        --tcp_port=<TCP port number>
```

Stop Container

```
# sudo memcli stop container
        --node name=<hostname>
        --tcp_port=<TCP port number>
```

List Containers

```
# sudo memcli list container
        --node_name=<hostname>
```

Create and Delete Mirrored Groups

Mirrored groups are the Schooner entities used to control container replication. Each Schooner Memcached node can belong to either the Independent Group (non-replicating) or the Mirrored Group (replicating).

Note that when you add/delete/start/stop a container belonging to a node in a Mirrored Group, the Admin interface will perform the operation automatically on both of the mirrored nodes.

Create Replication Group

```
# sudo memcli add group
        --cluster_name=<cluster name>
        --node_name=<hostname1(,hostname2, hostname3…)>
        --group_type=MIRRORED
        --group_name=<name>
```

Delete Replication Group

```
# sudo memcli delete group
        --cluster_name=<cluster name>
        --group_id=<id>
```

Container Backup/Restore

Overview

Schooner Appliance for Memcached supports the backup and restore of:

- Server configuration
  - This includes the service configuration files required to restore a fresh appliance to the backed up state of an existing server.
  - It also includes the Schooner Memcached container configuration.
  - Server configuration backup may be accessed via GUI or CLI menu.
- Container data
  - A container data backup/restore utility is provided to save/restore the contents of Schooner Memcached containers.
  - The container contents are retrieved from SSD storage and placed on an HDD attached to the Schooner Appliance for Memcached or on a remote storage device.

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; the restore process does not create a container. The user must ensure there is enough space in the container to hold all the restored objects.

Backups within a backup series must be restored in the same order in which they were taken. Any backup series can be restored, and a backup series can be restored to an arbitrary backup within the series. However, the restore process must begin by restoring a full backup, followed by zero or more incremental backups.

Note: Refer to the Schooner Appliance for Memcached™ - Application and Administration Guide for more information and instructions on backup and restore of server configuration and Memcached data.

Backup Commands

```
Usage:
    backup -hV | [-L] -l | -M -s [-L] [-T]

Options:
    -M, --mode=<full|incr>
    -s, --servers=<addr:port>
    -L, --dir=      default: /opt/schooner/memcached/data
    -T, --threads=  default: 8  (min=1, max=8)
    -l, --list      lists backup files in --dir
    -h, --help      displays help
    -V, --version   displays version
```

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 | [-L] -l | -s -L

Options:
    -s, --servers=<addr:port>
    -L, --dir=
    -l,  --list                lists  backup  files  in  --dir,  default  dir:
/opt/schooner/memcached/data
    -h, --help       displays help
    -V, --version    displays version
```

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:11211:1/1
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 list the contents of the directory specified by the –dir option recursively.

Examples using 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:11211:1/1

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:11211:1/1/0

Cluster       Container              Backup  Type  Date
FooBaz        Foo:11211:1            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 is made up of 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.

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.

- Log in 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”

- Log in 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, Ganglia and Cacti.

Net-SNMP 5.3.1

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

Nagios

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

Ganglia 3.0.6

- Includes standard Ganglia cluster monitors
- Includes standard Ganglia host monitors for:
    - Network traffic
    - CPU states and load
    - Disk I/O
    - NIC
    - RAID
    - IOSTAT for HDD and SSD
- Includes Schooner Memcached per container monitors for:
    - CPU
        - SIRQ
        - SYS
        - USER
    - Interface Bandwidth
        - Send

        - Receive
    - Flash
        - Space consumed
        - Number of objects
        - Number of objects created
        - Number of dead objects
        - Number of object overwrites
        - Reads/request
        - Gets/request
        - Writes/request
        - Puts/request
        - Deletes/request
        - Get exists/request
    - SDF Cache
        - Hit ratio
        - Hits/sec
        - Misses/sec
        - Evictions/sec
    - Memcached
        - Hit ratio
        - Hits/sec
        - Misses/sec
        - Evictions/sec

Cacti

- Includes standard Cacti monitors for:
    - IOSTAT for HDD and SSD
    - NIC
    - CPU load
    - Disk partitions
    - Memory
- Includes Schooner Memcached monitors for:
    - Gets/sec
    - Get hits/sec
    - Get misses/sec
    - Sets/sec
    - Transactions/sec (gets + sets)

Each Schooner Appliance for Memcached is able to host the data collection and display services for Nagios, Ganglia, and Cacti. 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.