HP Serviceguard for Linux Version A.11.18 Deployment Guide

Introduction

Purpose of this Document

The HP Serviceguard for Linux Deployment Guide provides administrators with step-by-step instructions for ordering, installing and configuring a Serviceguard cluster on two Linux hosts. Using this deployment guide, administrators can install and configure HP Serviceguard for Linux quickly and easily. This document pertains to specific hardware and software, but can easily be leveraged for use in deploying HP Serviceguard for Linux in different environments.

What Has Changed in this Version

The previous version used Device Mapper LUN identifiers of the form /dev/dm-N that were not persistent, that is, they could change after some failures and reboots. This version describes how to set up the cluster using /dev/mapper/mpathX format names that are persistent.

Scope

The HP Serviceguard for Linux Deployment Guide:

- Is focused on the steps required to install and configure HP Serviceguard for Linux on a two-node ProLiant DL cluster running Red Hat Enterprise Linux 5.1 with the HP StorageWorks 2000fc MSA.

- Is intended for customers who are new to clustering and HP Serviceguard. The goal is to simplify the deployment of a Serviceguard for Linux cluster.

- Describes the minimum hardware and software requirements for Serviceguard for Linux. It also includes instructions for setting up the hardware, installing the software, and configuring the cluster with a sample package.

- Assumes that the customer has access to the internet to download the required software.

Although this deployment guide is specific to Serviceguard for Linux Version A.11.18 on HP ProLiant DL servers (IA32/x86 architecture), the HP StorageWorks 2000fc MSA, and Red Hat Enterprise Linux 5.1, most steps apply to any supported hardware (servers, storage), Linux OS distribution, and Serviceguard for Linux version. Details of supported configurations can be found in the HP Serviceguard for Linux Certification Matrix (see the Related Documents section for where to find the latest version of this matrix).

Audience

This deployment guide is intended for Linux administrators who are new to Serviceguard and for those who want to use the HP Serviceguard for Linux installation and configuration scripts.

Installation and Configuration Overview

This overview shows the outline of the steps covered in this deployment guide.

1. What to Order
2. Getting Started
    a. Hardware Setup
    b. Network Planning
3. Gather Required Software
4. OS Installation

5. Storage Array Configuration
6. Preparing the Servers
   a. Fibre Channel Driver and Multipath configuration
   b. Shared storage configuration
   c. Network configuration
7. Serviceguard for Linux and Related Software Installation
8. Serviceguard for Linux Configuration
   a. Server Configuration
   b. Cluster Configuration
   c. Consolidated Log Package Configuration
   d. Sample Package Configuration
9. Verification

If you encounter errors in any of the steps, check the Troubleshooting section at the end of this document for possible remedies.

Terms and Definitions

Term	**Definition**
Consolidated Log (clog)	clog is the name of the Serviceguard package providing high availability for consolidated syslog and package log files across all nodes in the cluster. It is also the DSAU command for displaying a specified log file.
DSAU	Distributed System Administration Utilities provides several tools for simplifying the management of groups of systems and Serviceguard clusters. In Serviceguard, it is used for consolidating syslog and package logs from all nodes in the cluster to a central location.
Node	A host system or server that is configured to be a member of a Serviceguard cluster
Package	Application services (individual Linux processes) and resources grouped together and managed as a unit within Serviceguard in the event that a failover is required
Relocatable IP Address	A virtual IP address that is associated with a Serviceguard package; this type of IP address is not stationary to a specific node because it can move from one cluster node to another when a package is moved. Relocatable IP addresses are created and removed by Serviceguard.
Serviceguard Manager	A web application within HP System Management Homepage (HP SMH) used for monitoring, administering, and configuring Serviceguard clusters.
SGLX	Serviceguard for Linux

Table 1. Terms and Definitions

Related Documents

The following documents provide valuable information on the technology discussed in this deployment guide.

Tip:
Look at the date on each document to ensure you get the latest version.

- http://docs.hp.com → High Availability (under "Software Products & Solutions") → Serviceguard for Linux
  – Managing Serviceguard for Linux
  – HP Serviceguard for Linux Version A.11.18 Release Notes
  – Editing Security Files for Serviceguard

  - Installing and Configuring Apache Toolkit for Serviceguard for Linux
  - Securing Serviceguard

- http://docs.hp.com → Network and Systems Management (under "Internet and Networking") → System Administration → Installation and User's Guide
  - Distributed Systems Administration Utilities User's Guide

- http://www.hp.com/go/sglx/info (under "Downloadables")
  - HP Serviceguard for Linux Certification Matrix
  - Configuration Guide

- http://docs.hp.com → Storage Solutions (under "Hardware") → Storage Array Systems → Modular Array Systems → HP StorageWorks 2000fc Modular Smart Array → Manuals
  - HP StorageWorks 2000 Family Modular Smart Array CLI reference guide (under "General reference")
  - HP StorageWorks 2012fc Modular Smart Array user guide (under "User guide")

- http://docs.hp.com → Storage Solutions (under "Hardware") → Storage Software → Multi-path Software → Multipath Device Mapper for Linux Software → Manuals
  - Installation and Reference Guide Device Mapper Enablement Kit for HP StorageWorks Disk Arrays

- http://www.hp.com/go/ilo → iLO 2 Standard (under "Integrated Lights-Out 2") → Support and Documents (under "Support") → Manuals (under "Resources …")
  - HP Integrated Lights-Out 2 User Guide (under "User guide")

What to Order

This section provides a recommendation for what hardware and software to order for a two-node Serviceguard for Linux cluster. It specifies the minimum (and recommended) hardware requirements for high availability. Please refer to the current HP Quick Specs for part numbers and the latest ordering information.

1. 2 x ProLiant DL Servers:
   a. local hard drives for local storage and OS (1 at minimum, 2 recommended)
   b. 2 x Ethernet ports provided by the dual ported built-in Ethernet card (2 ports at minimum, 3 recommended)
2. 1 x HP StorageWorks 2000fc MSA (HP StorageWorks 2012fc Dual Controller Modular Smart Array), with at minimum:
   a. at least 2 hard drives
   b. 2 controllers
   c. 2 dual-port Fibre Channel host bus adapters (HBAs) (for example, FC1243 or FC1242SR, depending on the type of PCI slots available on the ProLiant servers)
3. 2 x Red Hat Enterprise Linux Advanced Platform 5 Update 1 Subscription Service
   a. 3 year 24x7 subscription with either 1 to 2 sockets or unlimited sockets depending upon the server chosen
4. 2 x Serviceguard for Linux for x86 Media, Server LTU & 1Yr 24x7

This list does not include some common items such as cables, power backup, and network switches.

Since this solution is configured on HP ProLiant servers, HP Serviceguard for Linux includes One Year of 24x7 Technical Software Support & Updates. This support can be extended to three years to coincide with support contracts on hardware if desired. This solution also includes support for Red Hat, 3-year Red Hat Enterprise Linux Advanced Platform 24x7 either 1 to 2 sockets or unlimited sockets depending upon the server chosen (visit http://www.hp.com/go/proliantrhel for more information). Please note that support is highly recommended for all high availability solutions.

For Linux management to complement the SGLX capabilities in multi-system environments, HP's ICE-Linux solution provides an integrated solution for system discovery, deployment & imaging, monitoring, and management. HP ICE-Linux auto-configures Nagios and other open source software to allow you to be quickly productive with Linux management, while still preserving flexibility and extensibility.

For more details on the various support options available from HP and ICE-Linux, please consult with your HP Sales Representative or Partner.

Getting Started

Hardware setup and networking planning must be completed before you start the installation.

Hardware Setup

The 2 servers and external storage device must be cabled for power, network, and storage connectivity. For initial set-up of the MSA2000 storage device, a serial console device will be needed. This can be done on a Windows laptop using the HyperTerminal software or with another system running similar terminal software along with the Mini DB9 RS232 Serial cable supplied with the MSA2000. A keyboard, mouse, and monitor will be required for use as a console to install the operating system on each of the servers. 5 Ethernet cables are required (two for each server and one for the storage device) along with access to 2 network switches (see Figure 1). Storage cabling is described in “Connecting the MSA2000 to the servers” and illustrated in Figure 2. Four Fibre Channel cables are required to connect the storage device to the servers.

Tip:
Do not connect the Fibre Channel cables to each server until after the OS has been installed on each server, otherwise the OS installation process may attempt to place the boot image on the external storage device.

Network Planning

The minimum network configuration requires 2 network connections configured as a bonded pair running a heartbeat and the public network connections from each server. Each network connection should be plugged into separate switches (see Figure 1). If three network connections are used, one should be configured in a private heartbeat network (shown in grey in Figure 1) and the other should be a bonded pair for the public and heartbeat networks. The manual “Managing HP Serviceguard for Linux” goes into more detail.

Figure 1. Network and Storage Topology

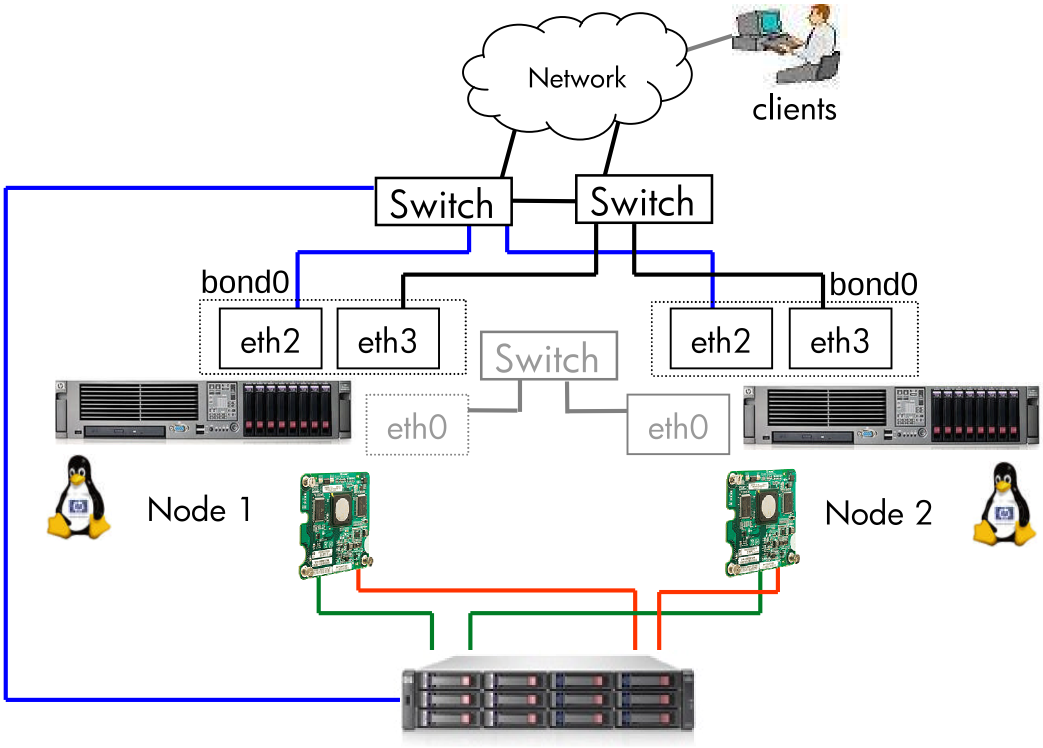

Required Networking Information:

1. Get static IP Addresses for the following:
   a. 2 IP addresses and hostnames with subnet mask, one for each server
   b. 2 relocatable IP addresses (1 for consolidated logging package, 1 for sample package), DNS names (optional) and subnet mask
   c. 1 IP address, hostname, and subnet mask for management access to the external storage device (MSA2000)
   d. Gateway address, DNS server address(es)
2. NTP Server (recommended but not required)
   a. name and address of NTP server that will be reachable by the servers and external storage device

Table 2 shows the list of IP addresses and hostnames needed for deploying a Serviceguard for Linux cluster. Column 2 shows the addresses used for demonstration purposes in the step-by-step instructions. You should fill in Column 3 with the addresses for your network environment.

Field	Example	Customer [fill-in]
Node 1 Static IP Address	16.89.84.245	
Node 1 Hostname	adam.cup.hp.com	
Node 1 Private heartbeat network IP Address		
Node 2 Static IP Address	16.89.84.247	
Node 2 Hostname	eve.cup.hp.com	
Node 2 Private heartbeat network IP Address		
MSA2000 IP Address	16.89.84.235	

Field	Example	Customer [fill-in]
Relocatable IP Address for clog Package	16.89.84.233	
DNS Name for clog package (optional)		
Relocatable IP Address for ws Package	16.89.84.218	
DNS Name for ws package (optional)		
Subnet (Network) and Subnet Mask	16.89.84.128, 255.255.255.128	
Gateway	16.89.84.129	
DNS Server Address	16.110.135.51, 16.110.135.52	
Subnet Mask (Private heartbeat network)		
NTP Server Address	15.36.88.4	

Table 2. Planning worksheet for Networking Parameters

Gather Required Software

To prepare for the software installation steps, gather the software packages in the following list. Store the downloaded files in a common directory or folder, for example on your PC in a folder such as c:\sglx_install. This step can also be done once the first server has been installed and connected to the Internet.

1. Linux OS Distribution

This deployment guide includes instructions for deploying Serviceguard on Red Hat Enterprise Linux 5.1. It assumes the administrator will install the operating system from CDs. Alternatively, the administrator may need to register their subscription service and download the appropriate images or use an update service such as Yellow Dog Updater (YUM), an open source rpm package management utility for Linux. For more information about YUM, refer to the Red Hat Enterprise Linux Deployment Guide at http://www.redhat.com/docs/manuals/enterprise/.

2. Kernel Errata

There may be a newer version of the kernel available online from Red Hat than the one provided with Red Hat Enterprise Linux 5.1 (RHEL5.1). The kernel version released with Red Hat Enterprise Linux 5.1 is 2.6.18-53.el5.

1. Check the HP Serviceguard for Linux Certification matrix for the latest RHEL5.1 kernel that has been certified for Serviceguard for Linux (see the Related Documents section of this paper for details on where to find the latest version of the certification matrix). At the time of publication the latest kernel version certified for Serviceguard for Linux A.11.18 was 2.6.18-53.1.14.EL5.
2. If you choose to use a newer kernel errata that has been certified by Serviceguard for Linux, download it from http://rhn.redhat.com using your subscription service login, or if you use some other update process such as YUM, update the servers after the operating systems have been installed.

3. Serviceguard for Linux CD

When you purchase Serviceguard for Linux from Hewlett-Packard, you should receive a CD (or image) containing the software. The instructions in this deployment guide are for Serviceguard for Linux Version A.11.18.

4. Serviceguard for Linux Patches

It is recommended that you download the most recent Serviceguard patches available for Serviceguard for Linux A.11.18.

At the time of publication, the following patches were available:

SGLX_00222 (Red Hat 5.0 IA32) Serviceguard A.11.18.03
SGLX_00204 (Red Hat 5.0 IA32) Serviceguard Manager B.01.01.03

For the specific installation recommended in this deployment guide, Serviceguard for Linux version A.11.18.02 or later must be used. This is required for support of the cluster lock LUN with HP Device Mapper Multipath on Red Hat 5.1 using the MSA2000.

Patches can be downloaded from HP at the following URLs:

http://itrc.hp.com (Americas and Asia Pacific)
http://europe.itrc.hp.com (Europe)

Instructions for downloading the patches from itrc.hp.com:

1. Go to http://itrc.hp.com.
2. Login.
3. Find the section (in center) titled ">> maintenance and support (hp products)".
4. Click ">> patch database".
5. Find the section titled ">>find individual patches".
6. Click ">>Linux".
7. On the "search for patches" page, choose the following options:
   Step 1: Select vendor and version: "redhat", version "5.0".
   Step 2: Select "Search by Keyword" (default), enter the keyword "Serviceguard".
   Step 3: Select Search Criteria: all words (default).
   Step 4: Select Results per page: 25 (default).
8. Click the "search>>" button.
9. On the "search results" page, select the most recent Serviceguard and Serviceguard Manager patches.
10. Click the "add to selected patch list>>" button.
11. On the "selected patch list" page, click the "download selected >>" button.
12. On "download patches" page, go to the "download items individually" section.
13. Click the "FTP>>" button to download the Serviceguard patch.
14. Select "Save" to save file to the local system (in the download directory, for example, c:\sglx_install). The downloaded file name is, for example, sglx_00222.tar.
15. Click the "FTP>>" button to download the Serviceguard Manager patch.
16. Select "Save" to save file to the local system (in the download directory, for example, c:\sglx_install). The downloaded file name is, for example, sglx_00204.tar.

NOTE: The instructions in this deployment guide assume each patch is downloaded individually.

5. Java JDK

Serviceguard Manager, the browser-based management interface for Serviceguard, requires Java 1.4.2.16 JDK or greater. The latest version of the Java 5 JDK is recommended. As of the publication date of this deployment guide, Java 6 has not been tested by HP for Serviceguard Manager.

The Java JDK is available from Sun at the following URL:

http://java.sun.com/j2se/1.5.0/download.html

1. Go to this website to download the latest Java 5 JDK.
2. Follow the instructions for downloading the Linux RPM in a self-extracting file.

The instructions in this document assume the reader downloads the "Linux RPM in self-extracting file".

At the time of publication, JDK 5.0 Update 15 was the latest release available from Sun. The download file for this version is named jdk-1_5_0_15-linux-i586-rpm.bin. When the .bin file is executed, the .rpm file is extracted.

6. HP Distributed Systems Administrator Utilities (DSAU)

DSAU is recommended for use with Serviceguard to facilitate troubleshooting by consolidating syslog and package log files from all nodes in the cluster to a central shared location. DSAU is required if you plan to configure the clog package. The installation scripts require that you install DSAU if you are installing Serviceguard Manager.

This software is normally provided on the Serviceguard for Linux CD, but the version for Red Hat 5, was unavailable when the CD was released.

To download DSAU:

1. Go to http://www.hp.com/go/softwaredepot/ha.
2. Search for "Linux Distributed Systems Administration Utilities".
3. Click "Receive for Free >>".
4. Select "Red Hat Enterprise Linux (AS and ES) 5 for x86" in the Software specification box.
5. Fill in required fields and agree to the terms.
6. Click the "Next" button.
7. On the "Software download confirmation" page, click "Download Directly >>".

At the time of publication, the version of HP DSAU available was 1.4-1. The downloaded file for this version is named hpdsau-1.4-1.rhel5.i386.rpm.

7. Serviceguard for Linux Installation and Configuration Scripts

Installation (sgEasyInstall) and configuration (sgEasyConfig) scripts are available to install Serviceguard for Linux, Serviceguard Manager, and related packages and to configure system settings to work with Serviceguard for Linux. The installation script walks the user through the installation of the rpm packages on the Serviceguard for Linux CD. The configuration script configures the system settings, such as services, /etc/hosts, PATH, and the firewall settings for each node planned for the Serviceguard cluster.

To download the Serviceguard for Linux Installation and Configuration Scripts:

1. Go to http://www.hp.com/go/softwaredepot/ha.
2. Search for "Serviceguard for Linux Installation and Configuration Scripts".
3. Click "Receive for Free >>".
4. Select "Red Hat Enterprise Linux (AS and ES) 5" in the Software specification box.
5. Fill in required fields and agree to the terms.
6. Click the "Next" button.
7. On the "Software download confirmation" page, click "Download Directly >>".

8. Serviceguard for Linux Free Toolkit Suite

The optional Serviceguard for Linux package included in these instructions requires the Serviceguard for Linux Apache Toolkit, which is part of the HP Serviceguard for Linux Free Toolkit Suite, available at no charge from HP.

To download this package:

1. Go to http://www.hp.com/go/softwaredepot/ha.
2. Search for "Serviceguard for Linux Free Toolkit Suite".
3. Click "Receive for Free >>".
4. Select "Red Hat Enterprise Linux (AS and ES) 5" in the Software specification box.
5. Fill in required fields and agree to the terms.

6. Click the “Next” button.
7. On the “Software download confirmation” page, click “Download Directly >>”.

At the time of publication, the version of the toolkit available was A.03.02-0. The downloaded file for this version is named sglxtools-A.03.02-0.product.redhat.tar.

9. Linux Driver Kit for the Fibre Channel HBAs

Download the latest driver from Hewlett-Packard for the Fibre Channel host bus adapters ordered with the MSA2000. The steps here are for the FC1243 card, which requires the latest “Driver Kit for QLogic HBAs and QLogic mezzanine-based HBAs”. For the FC1242SR, look for the Driver Kit for Emulex HBAs.

To download the kit:

1. Go to: http://www.hp.com.
2. Click “Software & Driver Downloads”.
3. Select “Download drivers and software (and firmware)”.
4. Enter product name, for example “FC1243”, in “for product:” box.

   NOTE: Select the HBA Product ordered with your MSA2000, for example, FC1243 (StorageWorks PCI-X 4 GBit Host Bus Adapter).
5. Click “>>”.
6. Select Operating System: Red Hat Enterprise Linux 5 Server (x86).

   NOTE: Select Red Hat Enterprise Linux 5 Server (x86-64), if appropriate.
7. From the “Driver – Storage Controllers – FC HBA” section, locate the most recent “Linux Driver Kit for HP Qlogic HBAs and mezzanine HBAs”.
8. Click “Download”.

At the time of publication, the most recent version was 8.01.07.25 (1 Nov 2007). The download file for this version is named hp_qla2x00-2007-10-05.tar.gz.

10. HP Device Mapper Multipath Enablement Kit

Download the HP Device Mapper Multipath Enablement Kit. This kit is required to support multiple paths from each server to the external storage device, the MSA2000.

To download the kit:

1. Go to http://www.hp.com/go/devicemapper.
2. Select Operating System, for example: Red Hat Enterprise Linux 5 Server (x86).

   NOTE: Select Red Hat Enterprise Linux 5 Server (x86-64), if appropriate.
3. Click the “Download” button for the latest version of “Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays”.

At the time of publication, the most recent version was v4.0.0 28 Feb 2008. The download file for this version is named HPDMmultipath-4.0.0.tar.gz.

OS Installation

The OS installation must be performed on each server.

A monitor, keyboard, and mouse can be attached directly to the server as a console device or a remote console can be used over the network via the HP iLO2 remote management port. To set up the iLO2 browser-based management interface, refer to the HP Integrated Lights-Out 2 User Guide.

Linux OS Distribution

For more detailed instructions for installing Red Hat Enterprise Linux 5.1 refer to the Red Hat Enterprise Linux Installation Guide at http://www.redhat.com/docs/manuals/enterprise/.

NOTE: Make sure the Fibre Channel cables from the storage server are not connected to the servers at this time. Otherwise, the Linux install process will attempt to install the boot partition on the storage server.

Follow these steps to install from the Red Hat Enterprise Linux 5.1 Distribution CD set:

1. Hook up the monitor, keyboard and mouse to the server.
2. Insert RHEL 5.1 CD#1.
3. Power cycle the server.
4. Watch the screen on the console for the following messages:
   ```
   Slot 0 HP Smart Array 6i Controller …
   Press [F8] to run the Optional ROM Configuration Arrays Utility
   ```
5. Hit F8 to pull up ORCA (Optional ROM Configuration for Arrays) to create a logical drive on the local disk.
   a. Select Create logical drive.
   b. Select appropriate RAID level (depends on number of local hard drives available and preference).
   c. Save the configuration.
   d. Exit ORCA.
6. Let the boot process flow through to install from CD.
7. When prompted for Mode, hit <Enter> for Graphical Mode, or type "linux text" and <Enter> for text mode.
8. If desired, run the media check to verify the CDs.
   a. Each CD is ejected after each media check.
   b. When done, reinsert CD#1.
9. At the welcome message, click "Next" to continue.
10. Select Language, choose "English" and "U.S. English", for example. Click "OK".
11. Enter Installation number. The installation number is a 16-digit hexadecimal text string provided by Red Hat with a subscription to Red Hat Enterprise Linux 5.1. The installation number enables a user to install the full set of supported packages included with the subscription.
12. Select (a) the appropriate target drive for installing the operating system (this should be the logical drive on the local disk created in Step 5), (b) partition removal options, and (c) partition layout options. For example, select "Wipe out partitions on local disk(s) and create default layout". Follow the prompts to confirm your selection.
13. Set up the initial network connection (See Table 2 for addresses).
    a. In the Network Devices section, select the Ethernet port to have "Active on Boot" (for one of the interfaces connected to your network). For example, click the checkbox for eth2 and un-check eth0.
    b. Highlight eth2, then click the "Edit" button.
       i. Select "Enable IPv4", manually enter the IP address and prefix (subnet mask).
       ii. Select "Enable IPv6" if appropriate for your network environment, check "Automatic neighbor discovery".
       iii. Click "OK".
    c. In the Hostname section, enter the fully qualified hostname.
    d. In the Miscellaneous section, enter the gateway address and primary and secondary DNS addresses.
    e. Click the "Next" button.
14. Set Time zone: select "system clock uses UTC", and, for example, "America" and "Los Angeles" for the time zone.
15. Set Root password, enter twice.

16. Select the following in the check boxes:
    Software Development
    Webserver

    NOTE: Software Development is required for installing Serviceguard, the QLogic driver, and HP Device Mapper Multipath. The Webserver is required for the optional Serviceguard package. You can omit the Webserver if not configuring the optional Serviceguard package and a Webserver is not otherwise needed.

17. Choose "Customize Later" (default). Click "Next".
18. Installation to Begin, click "Next".
19. Have CD#1, CD#2, and CD#3, or the DVD, ready. Click "Continue".
20. Insert each disk as prompted.
21. When done, click "Reboot".
22. After reboot, go through the post-installation steps (from the console window).
    a. Welcome
    b. License Agreement – Select "Yes" to agree to the terms and click "Forward".
    c. Enable Firewall (default is enabled)
        i. Make sure to allow ssh, allowed by default, and www (http).

           NOTE: http (port 80) is required by the optional Serviceguard package. You can omit http if not configuring the optional Serviceguard package and http is not otherwise needed. ssh is required by the Serviceguard for Linux configuration script and for remote terminal access to the servers.

        ii. Click "Forward". Then "Yes", to confirm when prompted.

    d. Enforce SELinux (default). Click "Forward".
    e. Don't enable Kdump (default). Click "Forward".
    f. Set date and time (manually or enter NTP server address). Click "Forward".
    g. Set up for SW Updates – Choose "No, I prefer to register at a later time", click "Forward". In the pop-up window, select "No thanks, I'll connect later", click "Forward". Or set up SW updates, if appropriate.
    h. Create user – Leave blank, click "Forward", Continue.
    i. Sound card – Click "Forward".
    j. Additional CDs to install? No.
    k. Click "Finish".

Copy Downloaded Software to Servers

At this point, you should copy all downloaded software packages to each server.

1. Login to the server (from the console or an ssh client).
2. Create a directory on the server for the downloaded software, for example:
   ```
   mkdir /tmp/sglx_install
   ```
3. From the location of the downloaded software, for example c:\sglx_install on your PC, copy the files to each server, for example:
   ```
   scp * root@adam:/tmp/sglx_install/.
   ```
4. Repeat on the other server.

Kernel Errata

It is recommended that the Red Hat Linux 5.1 kernel be updated to the latest kernel that has been certified by Hewlett-Packard for use with Serviceguard for Linux.

These steps should be performed on both servers.

The following instructions are for updating the kernel if the rpms are downloaded from Red Hat as an rpm packages:

1. To get the version of the kernel currently installed on the system, execute the command “uname –a”:
   ```
   uname -a
   Linux eve.cup.hp.com 2.6.18-53.el5 #1 SMP Wed Oct 10 16:34:02
   EDT 2007 i686 i686 i386 GNU/Linux
   ```
2. Compare the current kernel version, for example 2.6.18-53.el5, to the downloaded kernel version to verify that the downloaded kernel is more recent.
3. Change to the directory containing the kernel rpms. For example:
   ```
   cd /tmp/sglx_install
   ```
4. Install the rpms. For example:
   ```
   rpm –ivh kernel-2.6.18-53.1.14.el5.i686.rpm
   rpm –ivh kernel-devel-2.6.18-53.1.14.el5.i686.rpm
   ```
5. Edit (or check) the grub.conf file so that it boots the appropriate kernel. Update the “title” line, the “kernel” line to point to the new vmlinuz file under /boot and the “initrd” line to the new initrd file. The rpm install may have done these edits for you. For example:
   ```
   vi /boot/grub/grub.conf
   default=0
   timeout=5
   splashimage=(hd0,0)/grub/splash.xpm.gz
   hiddenmenu
   title Red Hat Enterprise Linux Server (2.6.18-53.1.14.el5)
        root (hd0,0)
        kernel /vmlinuz-2.6.18-53.1.14.el5 ro \
               root=/dev/VolGroup00/LogVol00 rhgb quiet
        initrd /initrd-2.6.18-53.1.14.el5.img
   ...
   ```
6. Reboot the server:
   ```
   reboot
   ```

NOTE: This step must be done prior to installing and configuring the QLogic driver and HP Device Mapper Multipath. Otherwise, special steps must be taken to ensure that the servers can boot off of the new kernel. Refer to the Red Hat documentation at http://www.redhat.com/docs/manuals/enterprise for more detail.

Linux Post-OS Installation

Serviceguard for Linux and the software for multipath support require several other packages from the Red Hat Linux 5.1 Distribution CDs that are not installed by default. You can install them now, or will be required to install them later.

These packages need to be installed on each server. The installation order below is determined by the CD on which each package was found.

NOTE: It may be easier to copy all of the packages from the CDs to your PC, then copy (scp) them to the servers rather than mounting each CD on each server.

1. Login to the server (from the console or an ssh client).

Serviceguard for Linux depends on the xinetd service.

1. To check if xinetd is installed, run the following command:
   ```
   rpm –qa | fgrep xinetd
   ```

If xinetd is not installed, there will be no output when running this command. If it is already installed, you will see output. For example:

```
xinetd-2.3.14-10.el5
```

2. If not installed, locate the xinetd rpm on one of the Linux OS CDs (probably CD#2) under the Server directory, and install it. For example:

```
rpm –ivh xinetd-2.3.14-10.el5.i386.rpm
```

The Serviceguard for Linux SNMP subagent requires the lm_sensors and net-snmp packages. lm_sensors must be installed before net-snmp. Instructions for net-snmp follow later.

1. To check if lm_sensors is installed, run the following command:

```
rpm –qa | fgrep lm_sensors
```

2. If not installed, locate the lm_sensors rpm on one of the Linux OS CDs (probably CD#2) under the Server directory, and install it. For example:

```
rpm –ivh lm_sensors-2.10.0-3.1.i386.rpm
```

The libXp package is required for Serviceguard Manager.

1. To check if libXp is installed, run the following command:

```
rpm –qa | fgrep libXp-
```

2. If not installed, locate the libXp rpm on one of the Linux OS CDs (probably CD#2) under the Server directory, and install it. For example:

```
rpm -Uhv libXp-1.0.0-8.1.el5.i386.rpm
```

The HP Device Mapper Multipath requires several packages. The two that were not installed by default are identified here.

1. To check if the required libraries are installed, use the following command:

```
rpm –qa | fgrep –e libsysfs-devel –e libaio-devel
```

2. If the dependencies are missing, locate them on the Linux OS CDs (probably CD#2) under the Server directory, and install. For example:

```
rpm –ivh libaio-devel-0.3.106-3.2.i386.rpm
rpm –ivh libsysfs-devel-2.0.0-6.i386.rpm
```

NOTE: Not all required dependencies are listed here. Refer to the "Installation and Reference Guide Device Mapper Enablement Kit for HP StorageWorks Disk Arrays" for more detail.

The Serviceguard for Linux SNMP subagent requires the net-snmp and lm_sensors packages. lm_sensors must be installed before net-snmp. Instructions for net-snmp are here.

1. To check if net-snmp is installed, run the following command:

```
rpm –qa | fgrep net-snmp
```

NOTE: By default, net-snmp-libs is already installed. Inspect the output carefully to determine if net-snmp-5.3.1-19.el5, for example, is installed.

2. If not installed, locate the net-snmp rpm on one of the Linux OS CDs (probably CD#3) under the Server directory, and install it. For example:

```
rpm –ivh net-snmp-5.3.1-19.el5.i386.rpm
```

The Serviceguard CIM provider requires the tog-pegasus package.

1. To check if tog-pegasus is installed, run the following command:

```
rpm –qa | fgrep tog-pegasus
```

2. If not installed, locate the tog-pegasus rpm on one of the Linux OS CDs (probably CD#4) under the Server directory, and install it. For example:

```
rpm –ivh tog-pegasus-2.6.1-2.el5.i386.rpm
```

Storage Array Configuration

The following instructions are for the MSA2000. It has a built-in browser-based management interface called SMU (Storage Management Utility). The device can also be configured using the CLI. This document has instructions to use the CLI to configure the IP address, then the remainder of the configuration steps are shown using SMU.

In this section there are instructions for creating external shared storage to be used for the cluster lock and for package data.

Connecting the MSA2000 to the Servers

Once the OS has been installed on each server, plug the Fibre Channel cables from the MSA2000 into each server. Make sure the green lights on the cards and controllers are on. If not, double check that the cables are completely plugged in. See Figure 2 for the wiring specification.

To avoid a single point of failure (SPOF), a path from each server to each storage controller should be created, providing two paths from each server. The MSA2000 has a feature called "host port interconnects" that provides high availability using internal connections between the host ports, providing redundancy in the event that one controller fails. This feature will be configured in a later step.

Plug in the Fibre Channel cables as follows:

1. Host 1 – plug 2 Fibre Channel cables from the dual channel HBA to the MSA2000 to
    a. Controller A Port 0
    b. Controller B Port 0
2. Host 2– plug 2 Fibre Channel cables from the dual channel HBA to the MSA2000 to
    a. Controller A Port 1
    b. Controller B Port 1

Figure 2. Wiring the servers to the storage controllers for redundancy

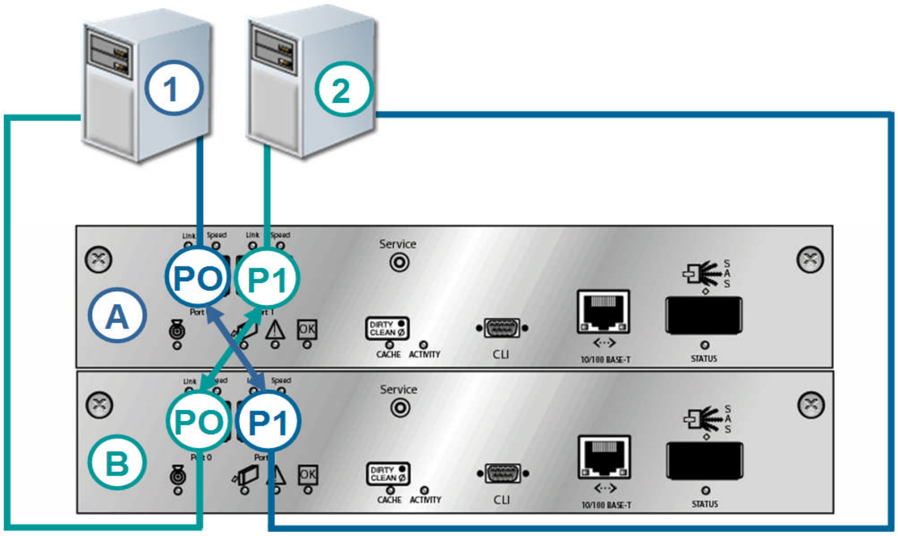

Setting Management Port IP Address using the CLI

The IP address can be configured using the CLI (command line interface) from a serial console or through DHCP.

To configure the IP address using the CLI, perform the following steps:

1. Plug the serial cable (Mini DB9 RS232 Serial cable) into the COM1 port of a PC to access the MSA2000 via a terminal emulator such as HyperTerminal (available on most Windows PCs) and into the serial port of one of the storage device controllers.
2. Run HyperTerminal from the PC. The settings are:

   Terminal Emulation Mode: VT-100 or ANSI
   Font: Terminal
   Translations: None
   Columns: None
   Connector: COM1 (typically)
   Baud rate: 115,200
   Data bits: 8
   Parity: None
   Stop bits: 1
   Flow control: None

   NOTE: For more detailed information, refer to the MSA2000 CLI Reference Guide.

3. Once connected, press <Enter> to display the prompt (#).
4. At the prompt, enter the "set network-parameters" command to configure the IP address for the controller connected to the network. For example, type the following command (as a single command):

   ```
   set network-parameters ip 16.89.84.235 netmask
   255.255.255.128 gateway 16.89.84.129 controller a
   ```

   NOTE: It is assumed that one of the controllers is already connected to a network switch using an Ethernet cable. In this example, the Ethernet cable is plugged into controller A.

5. To verify the setting, enter the following command:

   ```
   show network-parameters
   ```
6. Disconnect from the CLI and exit the emulator.

To verify, ping the address from another device on the network.

If there are problems after 3 minutes, type the following command at the CLI of the MSA2000 to restart the management controller on both controllers:

```
restart mc both
```

Creating a User Login from the CLI

By default, a user and password of manage/manage should already exist on the storage server to provide access to the browser-based Storage Management Utility (SMU).

If you want to change the defaults or add additional users, this can be done using the CLI through a terminal emulator. Follow the instructions above to establish a Hyperterminal connection to the controller. Here is an example of a command to create a new user, jsmith, with "manage" (modify) capabilities, and access to the command line interface and the web-browser interface. You will be prompted to input the password:

```
create user jsmith level manage interfaces cli, wbi
```

```
Enter Password for new user jsmith:*****
Re-enter Password:*****
Info: User Type not specified, defaulting to Advanced.
Success: New user created
```

To change the default manage password, use the following command:

```
set password manage
Enter new password:****
Verify new password:****
Info: Changing password for user: manage
Success: Password set
```

Configuring the MSA2000

This section provides instructions for the initial configuration of the storage device, configuration settings for the connected hosts, and the creation of a virtual disk and its volumes.

1. Go to SMU from a browser with access to the network, for example:

   http://16.89.84.235
   login: **manage/manage**

   NOTE: If the menu options in the left hand column do not appear, try the "Refresh" option in the browser.

2. Set the date and time.
   - a. Click "MANAGE" in the left-hand pane.
   - b. Click "GENERAL CONFIG" in the left-hand pane.
   - c. Click "set date/time" in the left-hand pane.
   - d. If NTP is available, locate the "Obtain Time with NTP" frame.
     - i. Select Enable NTP
     - ii. Enter the NTP Server Address, for example 15.36.88.4
   - e. Set current date, if incorrect, in "Set MSA Storage System Date" frame.
   - f. Set current time, if incorrect, in "Set MSA Storage System Time" frame.
   - g. To save changes, click the "Save Date/Time" button.
3. Configure the host ports
   - a. Click "MANAGE" in the left-hand pane.
   - b. Click "GENERAL CONFIG" in the left-hand pane.
   - c. Click "host port configuration" in the left-hand pane.
   - d. Set the Controller Link Speeds to match the link speed of the FC HBA cards on each host. The default is 4 GBit/sec. If the speed does not match, change it in each of the "Link Speed" pull-down menus in the top two frames.
   - e. Click "Update Host Port Configuration" to save changes.
   - f. Set the Advanced Options (in the "Advanced Options" frame).
     - i. Click "Change FC Loop ID".
     - ii. If necessary, change Loop ID to "Soft" (default) for each controller module.
     - iii. If changed, click "Save and Continue..", otherwise click "Return to Main Host Port Configuration Page".
     - iv. Click "Change FC Port Interconnect Settings".
     - v. In the "Host Port Configuration" frame, select "Interconnected".
     - vi. If changed, click "Save and Continue..", otherwise click "Return to Main Host Port Configuration Page".
     - vii. Click "Change Host Port Topology".
     - viii. If necessary, change Topology to "Loop" in the pull-down menu for each port. Loop topology is required with "Interconnected" host ports.
     - ix. If changed, click "Save and Continue..", otherwise click "Return to Main Host Port Configuration Page".

4. Create the virtual disk.
   a. Click “MANAGE” in the left-hand pane.
   b. Click “VIRTUAL DISK CONFIG” in the left-hand pane.
   c. Click “create a vdisk”.
   d. Select “Manual Virtual Disk Creation” for the Virtual Disk Creation Method.
   e. Enter Virtual Disk Name, for example “edenshare”.
   f. Select Virtual Disk RAID level, for example “RAID 1 – Disk Mirroring”.
   g. Click “Create New Virtual Disk”.
   h. Select drives to add to virtual disk by checking at least 2 green (available) disks. You may need to select more disks based on the RAID level configured in step f.
   i. Skip the “Calculate Virtual Disk Size” option.
   j. Would you like to add dedicated spare drives for this virtual disk? Select “No” (default).
   k. Click the “Continue” button to “Add Selected Drives to “edenshare” and Continue Creating Virtual Disk”.
   l. Configure Volumes for Virtual Disk “edenshare”.
      i. How Many Volumes? Select “3” from the pull-down menu.
      ii. Create Volumes of Equal Size? Click the “No” button.
      iii. Expose Volumes to All Hosts? Click the “Yes” button.
      iv. Automatically Assign LUNs? Click the “Yes” button.
      v. Would You Like to Name Your Volumes? Click the “Yes” button.
      vi. Advanced Virtual Disk Creation Options [skip]
      vii. Click the “Create Virtual Disk” button.
   m. To add volumes to virtual disk “edenshare”, enter the following:

Volume #	Volume Size – Mbytes	Volume Name
1	100	edenshare_lockLUN
2	200	edenshare_ws
3	300	edenshare_clog

   n. Click the “Add Volumes” button.

NOTE: It could take a few hours to initialize the virtual disk depending on the size and RAID level. By default the virtual disk will be initialized “Online”, so it can be used immediately. You do not need to wait for the vdisk initialization to complete before moving on to the next step.

5. Verify the volume mapping.
   a. Click “MANAGE” in the left-hand pane.
   b. Click “VOLUME MANAGEMENT” in the left-hand pane.
   c. Click “volume mapping” in the left-hand pane.
   d. Click “map hosts to volume” in the left-hand pane.
      i. Select each volume to view settings. From the “edenshare” Volume Menu, select each volume (one at a time) to view the “Current Host-Volume Relationships” to verify “All Hosts” and “rw” access on each port.
   e. Click “manage host list” in the left-hand pane.
      i. From the “Current Global Host Port List”, verify that the storage device sees the 4 host connections by the Host WWN and Manufacturer (for example, Qlogic).

NOTE: If all 4 host connections are not listed, check the cable connections and reboot the controllers. To reboot the controllers, take the following steps:
   a. Click “MANAGE” in the left-hand pane.
   b. Click “RESTART SYSTEM” in the left-hand pane.
   c. Select “Restart Both RAID Controllers” from the pull-down menu.
   d. Click the “Restart” button.

e. Click “OK” to the warning.

Preparing the Servers

In preparation for the Serviceguard for Linux installation and configuration steps, the two servers need to be configured to meet the minimum requirements for Serviceguard in terms of redundant paths to the external storage server, shared storage configuration and network redundancy.

NOTE: If you haven’t already done so, copy all downloaded software files to a directory on each server, /tmp/sglx_install, for example.

Install and Configure Driver for Fibre Channel HBA

New drivers will probably be required for the Fibre Channel HBA card used to connect each server to the shared storage device. These instructions are for the QLogic HBA cards recommended in the “What to Order” section. Be sure to get the driver appropriate for your systems’ HBA cards.

These steps need to be executed on each server.

The following examples are for the QLogic driver, version 8.01.07.25-1. If a newer driver is available, change the command arguments accordingly.

1. Check to see if a QLogic driver is already installed:
   ```
   rpm –qa | fgrep hp_qla2x00src
   ```
   If the QLogic driver is not installed, there will be no output from this command. If the QLogic driver is already installed, you will see command output showing the version, for example:
   ```
   hp_qla2x00src-8.01.07.25-1
   ```
2. If it is an earlier version than the one downloaded, remove it:
   ```
   rpm –e hp_qla2x00src
   ```
3. Change to the directory where the kit has been downloaded, for example:
   ```
   cd /tmp/sglx_install
   ```
4. Untar the driver kit using the following command:
   ```
   tar xfzv hp_qla2x00-2007-10-05.tar.gz
   ```
5. Change directory to the hp_qla2x00-yyyy-mm-dd directory. For example,
   ```
   cd hp_qla2x00-2007-10-05
   ```
6. Install the driver and fibreutils packages, for example:
   ```
   rpm -ivh hp_qla2x00src-8.01.07.25-1.linux.rpm
   rpm -ivh fibreutils-2.4-1.linux.i386.rpm
   ```

Configuring Multipath Support

Support for multiple Fibre Channel paths from each server to the MSA2000 is provided using HP Device Mapper Multipath.

The steps in this section need to be executed on each server.

1. Go to the directory where the HP Device Mapper Multipath Enablement Kit has been downloaded, for example:
   ```
   cd /tmp/sglx_install
   ```
2. Untar the driver kit using the following command:
   ```
   tar xfzv HPDMmultipath-4.0.0.tar.gz
   ```
3. Change to the HPDMmultipath-X.Y.Z directory, for example:
   ```
   cd HPDMmultipath-4.0.0
   ```

   NOTE: The documentation for HPDMmultipath is included in the kit and can be found under the docs directory, if needed.

4. Run the install script and answer the prompts as follows (user inputs in bold):

```
./INSTALL.sh
HP Device Mapper MultiPath Ver4.0.0 - Installation Menu

        1. Install HP Device Mapper MultiPath
        2. Uninstall Multipath Utilities
        3. Exit

Enter choice [1/2/3] :1

Checking for Build dependencies...


Building HPDMmultipath-tools-4.0.0 ....
Checking for previous installation. Please wait...

Do you wish to uninstall device-mapper-multipath-0.4.7-12.el5
?(y/n) : y

Checking for dependencies...


Do you still wish to uninstall device-mapper-multipath-0.4.7-
12.el5 ?(y/n):y

device-mapper-multipath-0.4.7-12.el5 is Uninstalled
successfully

Do you wish to install HPDMmultipath-tools-4.0.0? (y/n) : y


Warning: Restoring your previous configuration, you will have
      to manually edit the configuration file for HP
      recommended parameters. Please refer user documentation
      for more details.


Do you wish to restore previous configuration? (y/n) : y

Configuring multipath daemon to start at boot time....  OK

HPDMmultipath-tools-4.0.0 is installed successfully.
```

5. Configure the QLogic HBA parameters. If you have Emulex HBA cards, refer to the HP Device Mapper Installation and Reference Guide for specifics. For QLogic 2xxx family, complete the following steps:
   a. Edit the /etc/modprobe.conf file to add the following line (or if the line beginning with "options qla2xxx" exists, change it to the following):

      ```
      options qla2xxx qlport_down_retry=10 ql2xfailover=0
      ```

   b. Rebuild initrd by executing the following script:

      ```
      /opt/hp/src/hp_qla2x00src/make_initrd
      ```

   c. Reboot the host.

      ```
      reboot
      ```

6. Edit the /etc/multipath.conf file to add an entry for the MSA2000 family of storage devices.
   a. Create a copy of the original /etc/multipath.conf file, for example:

      ```
      cp /etc/multipath.conf /etc/multipath.conf.orig
      ```

b. Copy the sample multipath.conf file provided in the kit to the /etc directory (input as one line):

```
cp /usr/share/doc/HPDMmultipath-tools-
      4.0.0/multipath.conf.HPTemplate /etc/multipath.conf
```

c. Open the /etc/multipath.conf file to add a “device” subsection entry for the MSA2000 within the “devices { }” block as follows:

```
device
{
        vendor                  "HP”
        product                 "MSA2[02]*"
        path_grouping_policy    multibus
        getuid_callout          "/sbin/scsi_id -g -u -s /block/%n"
        path_selector           "round-robin 0"
        rr_weight               uniform
        prio_callout            "/bin/true"
        path_checker            tur
        hardware_handler        "0"
        failback                immediate
        no_path_retry           12
        rr_min_io               100
}
```

7. Restart the multipath daemon by executing the following commands:

```
/etc/init.d/multipathd restart
/sbin/multipath
```

8. To view status of multipath devices, execute the following command:

```
/sbin/multipath -ll
```

The sample output shows there are 3 mpaths to the 3 volumes created on the MSA, dm-4, dm-3, and dm-2 and the physical paths they are mapped to, for example, dm-4 maps to sdc and sdf.

```
mpath2 (3600c0ff000d50322ff8a3e4803000000) dm-4 HP,MSA2012fc
[size=287M][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][active]
 \_ 0:0:1:2 sdc 8:32  [active][ready]
\_ round-robin 0 [prio=0][enabled]
 \_ 1:0:0:2 sdf 8:80  [active][ready]
mpath1 (3600c0ff000d50322ff8a3e4802000000) dm-3 HP,MSA2012fc
[size=191M][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][active]
 \_ 0:0:1:1 sdb 8:16  [active][ready]
\_ round-robin 0 [prio=0][enabled]
 \_ 1:0:0:1 sde 8:64  [active][ready]
mpath0 (3600c0ff000d50322ff8a3e4801000000) dm-2 HP,MSA2012fc
[size=96M][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][active]
 \_ 0:0:1:0 sda 8:0   [active][ready]
\_ round-robin 0 [prio=0][enabled]
 \_ 1:0:0:0 sdd 8:48  [active][ready]
```

Configuring the Lock LUN and Shared Volumes

In this section, we will be configuring 3 shared volumes. The first is for the Serviceguard for Linux cluster lock LUN, a special area of the disk used for cluster arbitration during cluster re-formation. While this deployment guide provides steps to configure a lock LUN, a quorum server can also be used for cluster arbitration. The quorum server is a separate software package that comes with Serviceguard for Linux and requires a third server. The second volume will be used as a shared data volume for the sample Serviceguard for Linux package to be created in a later step. And the third volume is for the DSAU consolidated log (clog), which allows the user to view consolidated syslog and Serviceguard package logs for the entire cluster from a central location. The consolidated log will be configured as a Serviceguard for Linux package, also in a later step.

Most of the shared volume configuration is performed on the first node. The configuration can then be imported to the second node. A few of the steps must be performed on both nodes.

Table 3 shows the details for the 3 shared volumes used in this sample deployment.

Shared Volume	Volume Name (on the MSA)	Standard device paths	Device Mapper persistent name	Partition size	Mount point
Cluster lock LUN	edenshare_lockLUN	/dev/sda, /dev/sdd	/dev/mapper/mpath2	100 MB	n/a
ws package LUN	edenshare_ws	/dev/sdb, /dev/sde	/dev/mapper/mpath3	200 MB	/ws
clog package LUN	edenshare_clog	/dev/sdc, /dev/sdf	/dev/mapper/mpath4	300 MB	/clog

Table 3. Shared Volumes

Create the Partitions

When configuring the shared storage device earlier in this document, you created 3 shared volumes. They are for the cluster lock LUN, the ws package LUN, and the clog package LUN. The cluster lock LUN needs to be configured as a “Linux” partition (Hex code 83) and the package LUNs need to be configured as “Linux LVM” partitions (Hex code 8e). For each volume, you will create a new (n) partition (Partition #1), set the type (t), and save (w) the settings for the partition.

The steps in this section are to be executed on one server only.

1. To create a partition on the cluster lock LUN (/dev/sda), run the fdisk command and answer the prompts as follows (user inputs in bold):

   ```
   fdisk /dev/sda

   Command (m for help): n
   Command action
      e   extended
      p   primary partition (1-4)
   p
   Partition number (1-4): 1
   First cylinder (1-1024, default 1): 1
   Last cylinder or +size or +sizeM or +sizeK (1-1024, default 1024): 1024

   Command (m for help): t
   Selected partition 1
   Hex code (type L to list codes): 83

   Command (m for help): w
   The partition table has been altered!

   Calling ioctl() to re-read partition table.
   Syncing disks.
   ```

2. To create a partition on the ws package LUN (/dev/sdb), run the fdisk command and answer the prompts as follows (user inputs in bold):

   ```
   fdisk /dev/sdb

   Command (m for help): n
   Command action
   ```

```
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-24, default 1): 1
Using Last cylinder or +size or +sizeM or +sizeK (1-1016, default 1016):
1016

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

3. To create a partition on the clog package LUN (/dev/sdc), run the fdisk command and answer the prompts as follows (user inputs in bold):

```
fdisk /dev/sdc

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-1013, default 1): 1
Last cylinder or +size or +sizeM or +sizeK (1-1013, default 1013): 1013

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Create the Logical Volumes

In this section, you will create the shared volume groups for the 2 Serviceguard for Linux packages (ws and clog). This includes creating physical volumes, volume groups, and the logical volumes. You will also create the file system on each logical volume and create the mount points for each logical volume on the first server.

The steps in this section are to be executed on one server. This should be the same server where the partitions were created in the previous step.

1. Execute the following commands to create a logical volume for the ws package:

```
pvcreate -f /dev/mapper/mpath3
vgcreate /dev/vgws /dev/mapper/mpath3
lvcreate -L 180M -n lvol1 vgws
mke2fs -j /dev/vgws/lvol1
mkdir /ws
vgchange -a n vgws
```

2. Execute the following commands to create a logical volume for the clog package:

```
pvcreate -f /dev/mapper/mpath4
vgcreate /dev/vgclog /dev/mapper/mpath4
```

```
lvcreate -L 280M -n lvol1 vgclog
mke2fs -j /dev/vgclog/lvol1
mkdir /clog
vgchange -a n vgclog
```

NOTE: The size specified in the lvcreate command, specified by the "-L" option, is related to the size of the physical partition created on the storage device.

Backup the Volume Groups

The steps in this section should be performed on the same server where the partitions and volumes were created.

1. Execute the vgcfgbackup command to backup the volume groups, for example:

   ```
   vgcfgbackup /dev/vgws /dev/vgclog
   ```

Import and Configure the Volume Groups on the Second Server

At this point, you need to import the volume groups (e.g. vgws and vgclog) and create the mount points on the second server.

NOTE: The Serviceguard for Linux documentation suggests using vgexport and vgimport to achieve the same result, but the vgscan on the second server is sufficient.

The steps in this section should be performed on the second server.

1. First, run the "fdisk –l" command to verify that the external devices are visible on the second server. For example, look for /dev/sda through /dev/sdf, and /dev/mapper/mpath2, /dev/mapper/mpath3, and /dev/mapper/mpath4.

   ```
   fdisk –l
   ```

   NOTE: If you do not see these devices, test the connections and reboot the server.

2. Import the volume groups by running the vgscan command.

   ```
   vgscan
     Reading all physical volumes.  This may take a while...
     Found volume group "vgclog" using metadata type lvm2
     Found volume group "vgws" using metadata type lvm2
   ```

   NOTE: If the volume groups, vgws and vgclog are not found, reboot the server.

3. Create the mount points

   ```
   mkdir /ws
   mkdir /clog
   ```

4. Backup the volume group configurations

   ```
   vgcfgbackup /dev/vgws /dev/vgclog
   ```

5. Deactivate the volume groups on this server

   ```
   vgchange –a n vgws vgclog
   ```

Create alternate disk monitoring script

In this step, you will create a configuration script that will be used to monitor disk failures. This alternate script is required to overcome the limitation that the device mapper names /dev/dm-N, as described in the Serviceguard documentation, are not persistent. That is, if the LUN associated with /dev/dm-3 failed, then /dev/dm-4 would be renamed to /dev/dm-3 on the following reboot of the system.

Create the conversion script */usr/local/cmcluster/bin/cmresserviced_custom*. The contents must be exactly as shown:

```
#!/bin/sh

. /etc/cmcluster.conf

typeset -a realdm dm
typeset -i i

(( i = 0 ))
for arg in "$@"
do
dm[$i]=${arg##/dev/mapper/}
realdm[$i]=/dev/$(multipath -ll ${dm[$i]}| grep mpath | awk '{print
$3}' )
(( i = i + 1 ))
done

$SGLBIN/cmresserviced "${realdm[*]}"
```

Network Configuration

Configure Network Bonding

In this step, you will configure a bonded network interface. You will need to determine which interfaces are participating in the bond. In the following example, eth2 and eth3 are configured to create the bond0 logical interface to meet the minimum requirements for Serviceguard for Linux.

These steps should be performed on each server, using the appropriate IP and MAC addresses.

In this example, we show the content of the 3 configuration files in the /etc/sysconfig/network-scripts directory required to create the bonded interface. The ifcfg-eth# files may already exist, but the ifcfg-bond0 file probably does not exist.

1. Change directory:
   ```
   cd /etc/sysconfig/network-scripts
   ```
2. Create ifcfg-bond0 and add the IP address and other network addresses according to your environment as well as the other fields listed here, for example:
   ```
   DEVICE=bond0
   IPADDR=16.89.84.247
   NETMASK=255.255.255.128
   NETWORK=16.89.84.128
   BROADCAST=16.89.84.255
   ONBOOT=yes
   BOOTPROTO=none
   USERCTL=no
   BONDING_OPTS='miimon=100 mode=1'
   ```
3. Edit the ifcfg-ethN file to contain the following content. Remove existing entries, if any, except for "DEVICE" and "HWADDR". Make sure that the HWADDR field matches the specific MAC address for the card, do not use the value listed here. Here is an example for ifcfg-eth2:
   ```
   DEVICE=eth2
   USERCTL=no
   ```

```
ONBOOT=yes
MASTER=bond0
SLAVE=yes
BOOTPROTO=none
HWADDR=00:14:C2:C0:44:C5
```

4. Edit the ifcfg-ethN file to contain the following content. Remove existing entries, if any, except for "DEVICE" and "HWADDR". Make sure that the HWADDR field matches the specific MAC address for the card, do not use the value listed here. Here is an example for ifcfg-eth3:

   ```
   DEVICE=eth3
   USERCTL=no
   ONBOOT=yes
   MASTER=bond0
   SLAVE=yes
   BOOTPROTO=none
   HWADDR=00:14:C2:C0:44:C4
   ```

5. Edit /etc/modprobe.conf to add the following 2 lines (the second and third lines shown here should be input as one line):

   ```
   alias bond0 bonding
   install bond0 /sbin/modprobe tg3; /sbin/modprobe e1000;
         /sbin/modprobe --ignore-install bonding –o bond0
   ```

6. Restart the network from console:

   ```
   /etc/init.d/network restart
   ```

   NOTE: You may see an error while shutting down bond0, since it previously didn't exist. For example:

   ```
   /etc/sysconfig/network-scripts/ifdown-eth: line 101:
         /sys/class/net/bond0/bonding/slaves: No such file or
         directory
   ```

7. To verify the bond configuration, check the `/proc/net/bonding/bond0` file. This listing shows two slave interfaces with eth2 as the currently active slave.

   ```
   more /proc/net/bonding/bond0
   Ethernet Channel Bonding Driver: v3.1.2 (January 20, 2007)

   Bonding Mode: fault-tolerance (active-backup)
   Primary Slave: None
   Currently Active Slave: eth2
   MII Status: up
   MII Polling Interval (ms): 100
   Up Delay (ms): 0
   Down Delay (ms): 0

   Slave Interface: eth2
   MII Status: up
   Link Failure Count: 0
   Permanent HW addr: 00:14:c2:c0:ba:8f

   Slave Interface: eth3
   MII Status: up
   Link Failure Count: 0
   Permanent HW addr: 00:14:c2:c0:ba:8e
   ```

Identify Free Port for DSAU Consolidated Logging

Before configuring the consolidated log package, you should choose an available TCP port. Any port that is not in use can be selected, but HP recommends that you choose a port from the reserved range of ports below 1024. These steps are required only if configuring the DSAU consolidated logging package, clog.

The steps in this section should be performed on each server to ensure the port is free on each server.

1. Check the /etc/services file for well-known reserved ports to look for an unreserved port, for example:

   ```
   more /etc/services
   ```

2. Verify that an unreserved port from step 1 is not in use by running the "netstat –an" command and searching the output to ensure the selected port is not in use:

   ```
   netstat –an
   ```

Serviceguard for Linux and Related Software Installation

Java JDK

Serviceguard Manager requires the Java JDK 1.4.2 or greater. We recommend you install the latest version of Java 5 that is available from Sun.

This needs to be installed on each server.

To install JDK 5.0, (for example, jdk-1_5_0_15) from the .bin file, run the following commands:

1. Change to the directory where the Java JDK install file is located. For example:

   ```
   cd /tmp/sglx_install
   ```

2. Make sure that the .bin file is executable.

   ```
   chmod a+x jdk-1_5_0_15-linux-i586-rpm.bin
   ```

3. Run the .bin file to extract the rpm package. This step may also install the rpm.

   ```
   ./jdk-1_5_0_15-linux-i586-rpm.bin
   ```

   [Hit space bar until the end of the license terms]

   ```
   Do you agree to the above license terms? [yes or no] yes
   ```

4. Install the Java JDK, if necessary.

   ```
   rpm -ivh jdk-1_5_0_15-linux-i586.rpm
   ```

Record the path to the java binary from JDK. If unsure, run the "find" command to find all java files and directories. For example "find / -name java" will search the entire system under "/". The path from the install above is: /usr/java/jdk1.5.0_15/bin/java.

HP Distributed Systems Administration Utilities

HP Distributed Systems Administration Utilities (DSAU) provides consolidated syslog and package log files. While this is optional, it is recommended to facilitate troubleshooting when there are problems with the Serviceguard for Linux cluster or its packages.

This needs to be installed on each server.

To install DSAU, run the following commands:

1. Change to the directory where the DSAU package is located. For example:

   ```
   cd /tmp/sglx_install
   ```

2. Install DSAU, for example:

   ```
   rpm –ivh hpdsau-1.4-1.rhel5.i386.rpm
   ```

NOTE: The installation script will instruct the user to set up the PATH environment variable. This will be taken care of by the Serviceguard System Configuration Automation script later.

Serviceguard for Linux Installation Script

The Serviceguard for Linux Installation script will guide you through installing Serviceguard and Serviceguard Manager from the CD. It will also check for dependencies and prompt you to install the required dependencies, if not already installed. This script will streamline the installation of the several rpm packages that are included on the CD, reducing the installation from many commands to just one command to invoke the script.

This script needs to be run on each server.

The script will ask for the directory path where the Serviceguard for Linux CD is mounted. It will check for Linux packages that are required as prerequisites for Serviceguard and Serviceguard Manager. Instructions for installing the prerequisites are provided earlier in this deployment guide. If not already installed, you will be prompted to install these packages for the installation to complete successfully. This can be done in a separate terminal window.

You will also be prompted at least once to verify or install Java JDK 1.5. If you have already installed Java, just type “yes” to continue. It will also prompt you to provide the path to the JDK java binary, which is required by Serviceguard Manager and Tomcat. For the Java JDK installation described earlier, the path will be /usr/java/jdk1.5.0_15/bin/java. If you do not know where java is installed, you can try to search the server’s file system using the “find” command, for example “find / -name java”. Look for the java file in the bin directory of the jdk.

The script will also check whether the authd package is installed on the server. If installed, you will be prompted to remove it. It will also to attempt to ensure that authd is not installed in the future by a software update service such as YUM. Authd conflicts with another package required by Serviceguard for Linux.

To run the script, follow these steps:

1. Mount the Serviceguard for Linux CD.
2. Open a terminal window to the server from the console or an ssh client.
3. Change to the directory where the script resides. For example,
   ```
   cd /tmp/sglx_install
   ```
4. Untar the file if necessary,
   ```
   tar xf sglx_easy_install.tar
   ```
5. Make sure that the script is executable, for example,
   ```
   chmod +x sgEasyInstall
   ```
6. Invoke the script.
   ```
   ./sgEasyInstall
   ```

   NOTE: If you get a “not a valid executable” error from the tomcat_cfg step when entering the path to Java, hit [Return] to Proceed anyway. The wrong path to Java was provided. You will need to re-run the tomcat_cfg script (/opt/hp/hpsmh/tomcat/bin/tomcat_cfg) and provide the correct path the Java binary from the JDK.

The script comes with its own README. Refer to the README for additional information.

Serviceguard Patches

HP recommends that you install the latest patches available for Serviceguard for Linux A.11.18 and Serviceguard Manager B.01.01 on each server. The configuration recommended in this deployment guide, specifically for the use of HP Device Mapper multipath for the cluster lock LUN, requires Serviceguard for Linux version A.11.18.02 or later.

Serviceguard Patch

The following instructions assume that you are installing Serviceguard for Linux for the first time, and not upgrading from a previous version. If you are patching an existing Serviceguard for Linux cluster node, please see the patch instructions in the text file included with the patch.

These steps need to be performed on each server.

To install the patch:

1. Change to the directory containing the patch file, for example:
   ```
   cd /tmp/sglx_install
   ```
2. Untar the patch file, for example:
   ```
   tar xf SGLX_00222.tar
   ```
3. Change to the tools directory.
   ```
   cd tools
   ```
4. Run the sgupdate command.
   ```
   ./sgupdate
   ```

NOTE: If the script terminates with the following error:

```
An update for sgproviders exists - loading.
Error executing command "rpm -U /tmp/sglx_install/rpms/sgproviders-
A.02.00.01-0.rhel5.i386.rpm".
See file /tmp/rpminst.4550 for details.
```

a. Stop and restart the cimserver using the following commands
   ```
   cimserver –s
   cimserver
   ```
b. Rerun the sgupdate command to complete the patch installation:
   ```
   ./sgupdate
   ```

Serviceguard Manager Patch

These steps need to be performed on each server.

To install the patch:

1. Change to the directory containing the patch file, for example:
   ```
   cd /tmp/sglx_install
   ```
2. Untar the patch file, for example:
   ```
   tar xf SGLX_00204.tar
   ```
3. Install the rpm
   ```
   cd rpms
   rpm -i --force sgmgrpi-B.01.01.03-1.rhel5.i386.rpm
   ```
4. Restart HP System Management Homepage to make sure it recognizes the new version of the Serviceguard Manager:
   ```
   /etc/init.d/hpsmhd restart
   ```

Serviceguard for Linux Free Toolkit Suite

For the optional sample package, we will be using Serviceguard for Linux to monitor and protect an instance of the Apache webserver. The Apache Toolkit is used for the sample package.

This toolkit needs to be installed on only one of the servers.

To install the toolkit:

1. Change to the directory containing the toolkit, for example:
```
cd /tmp/sglx_install
```
2. Untar the file, for example:
```
tar xf sglxtools-A.03.02-0.product.redhat.tar
```
3. Install the Apache Toolkit, for example:
```
rpm -ivh apache-toolkit-A.03.01-0.product.redhat.noarch.rpm
```

Serviceguard for Linux Configuration

In this section, you will perform the server configuration steps that are required prior to configuring a Serviceguard for Linux cluster. Then you will configure a 2-node Serviceguard for Linux cluster, create the clog package, and configure a sample package.

Most steps will be performed using the browser-based Serviceguard Manager. Some of the sample package configuration steps will need to be performed on one of the nodes in the cluster.

Serviceguard System Configuration Automation script

The Serviceguard System Configuration Automation script will perform most of the server configuration tasks that need to be done prior to creating a Serviceguard for Linux. It will configure the following:

- Start and configure xinetd and identd services
- Set up the PATH and MANPATH environment variables
- Add cluster node entries to /etc/hosts
- Configure the firewall (if enabled) with Serviceguard-specific settings
- Configure the /etc/nsswitch.conf settings
- Configure lvm for exclusive volume group activation

This script should be run from one of the servers and will perform configuration updates for all nodes intended for the cluster. The user will be prompted for the list of servers. Public ssh key authorization should be configured to permit the script to run on each (remote) server without the need to prompt for passwords.

To set up ssh keys, replace <othernode> with the hostname (or IP address) of the remote server(s), for example “eve.cup.hp.com”. Execute the following command on the server where the script will be invoked to configure remote access to the other server(s) intended for the cluster:

1. Generate the keys
```
ssh-keygen –t rsa
```
   NOTE: Use the default file name (id_rsa) for the keys and leave the passphrase empty (if desired).

2. Copy the public key to the other node using the following command (as one line):
```
scp /root/.ssh/id_rsa.pub
      root@<othernode>:/root/.ssh/authorized_keys2
```
   NOTE: This step will fail if the “/root/.ssh” does not exist on the <othernode>. If it does not exist, create it.

3. Append the public key from the current node to the authorized_keys file on the other node by executing the following command (as one line):

   ```
   ssh root@<othernode> 'cat /root/.ssh/authorized_keys2
         >> /root/.ssh/authorized_keys'
   ```

4. Set the shell for the ssh authentication agent using the following command:

   ```
   ssh-agent $BASH
   ```

5. Add the rsa identity to the ssh authentication agent

   ```
   ssh-add
   ```

NOTE: To verify that the ssh keys have been set up properly, try to ssh to the other node. If not prompted for a password, the ssh keys have been set up properly.

To run the script, follow these steps:

1. Change to the directory containing the configuration script, for example:

   ```
   cd /tmp/sglx_install
   ```

2. Make sure the configuration script is executable:

   ```
   chmod +x sgEasyConfig
   ```

3. Run the script:

   ```
   ./sgEasyConfig
   ```

4. You will be put into an edit (vi) window to specify the inputs for the script. The content is shown here. User inputs are shown in bold. For the list servers (including the server where the script is being executed), make sure the hostname and IP Address are TAB separated. When done, save and quit ("[Esc]:wq!"). You will be returned to the script for further processing.

```
############################################################
# Enter yes or no to indicate whether Virtual Machines are
# participating as nodes in the cluster. Please provide the
# information on same line as the question.
############################################################

Will this cluster have any VM(s) participating as nodes? [yes|no] : no


############################################################
# Enter list of servers that will be part of the current cluster
# and their corresponding IP addresses in the format :
# HOSTNAME.DOMAINNAME<tab>IP-ADDRESS
# Example: xyz.abc.hp.com       111.222.333.444
############################################################
adam.cup.hp.com 16.89.84.245
eve.cup.hp.com 16.89.84.247


############################################################
# Enter the list of non-root users allowed to access this
# cluster and their corresponding roles in the format :
# USER<tab>ROLE
# Example: abcd monitor
# Please note that monitor, full_admin and package_admin are
# the allowed roles for non-root users
############################################################
```

NOTE: This last section is left blank indicating that only "root" users are enabled to manage the cluster.

The script comes with its own README. Refer to the README for additional information.

Test Volume Group Activation On Each Node

This should be done for each volume group, vgws and vgclog to verify volume group activation on each server. The steps shown here are for vgws.

1. On the first server (for example, adam)

   a. Activate volume group:

   ```
   vgchange --addtag $(uname –n) vgws
   vgchange -a y vgws
   mount /dev/vgws/lvol1 /ws
   ```

   b. Write test file to volume group:

   ```
   echo “Written by” `hostname` “on” `date` > /ws/test
   cat /ws/test
   ```

   NOTE: If you try to activate `vgws` on the other server, it should fail. For example:

   ```
   vgchange -a y vgws
     0 logical volume(s) in volume group "vgws" now active
   ```

   c. Deactivate the volume group:

   ```
   umount /dev/vgws/lvol1
   vgchange -a n vgws
   vgchange –-deltag $(uname –n) vgws
   ```

2. Repeat the steps on the second server.

Configure the cluster

In the next steps, you will create the cluster, define the node membership, configure the cluster heartbeat and cluster lock LUN device.

1. From an internet browser such as Internet Explorer, invoke HP System Management Homepage, https://[hostname]:2381. For example:
   https://eve.cup.hp.com:2381
2. Login (use root user and password set during installation of the operating system).
3. Go to the Tools tab.
4. Click the “Serviceguard Manager” link to launch the Serviceguard Manager.
5. Click the “Create Cluster” button on the right.

   NOTE: You may see the following message: “`There is no cluster configured.`”

6. In the Create Cluster window, enter the Cluster Name, for example “**`Test`**” and enter checkmarks in the boxes for both nodes, for example: adam, eve.
7. Go to the Network tab,
   Enter in the “Subnets” section, for example:
   Subnet: 16.89.84.128, Type: Heartbeat
   Enter in the “Select Subnet Configuration” section, for example:

Node	Network	Address
adam	bond0	16.89.84.245
eve	bond0	16.89.84.247

8. Go to the Lock tab
   For the Cluster Lock Type, Select “Lock Lun”
   Enter the Lock Lun Path for each node, for example:

Node	Lock Lun Path
adam	**`/dev/mapper/mpath0`**
eve	**`/dev/mapper/mpath0`**

   Select OK.

NOTE: When using Device Mapper Multipath, the path to the cluster Lock LUN, for example /dev/mapper/mpath0, must be the same on each node.

9. Select Check Configuration. Look for any errors.

   NOTE: You may get a warning about the default NODE_TIMEOUT value. This warning can be ignored here, but refer to the documentation when finalizing your cluster.

10. Select Apply Configuration. Select "OK" in the pop-up dialog box.
11. To verify the cluster configuration, run the following options from Administration menu of the HP Serviceguard Manager Summary page to test that each node can run the cluster in the event that the other node fails:
    a. Administration -> Run Cluster (on both nodes)
    b. Administration -> Halt Node (select adam)
    c. Administration -> Run Node (on adam)
    d. Administration -> Halt Node (select eve)

Configure the consolidated log package

The consolidated logging tool is used to consolidate syslog and package log files from all nodes in the cluster. Configuration is done through a wizard in the Serviceguard Manager to create a clog package in Serviceguard. This makes the consolidated logging tool highly available.

1. In your browser on the HP Serviceguard Manager Summary page, go to the "Configuration" menu and select "Configure Log Consolidation Tool…"
2. Enter the "Package Storage Parameters", for example:
```
Volume Group       /dev/vgclog
Logical Volume     /dev/vgclog/lvol1
Mount Point        /clog
Filesystem Type    ext3
Mount Options      -o rw
```
3. Enter the "Package Network Parameters" (Package Relocatable IP address and subnet mask), for example:
```
IP Address         16.89.84.233
Subnet Address     16.89.84.128
```
4. Select "Perform package log consolidation" by putting a check in the box.
5. Select "Use TCP?" by putting a check in the box. Then enter the free TCP port number identified earlier.
6. Click "OK".
7. Click "Yes" in the pop-up dialog box to apply this configuration.
8. Check the operations log to verify success. The operations log should appear in a separate pop-up window.

   NOTE: You may see the following errors in the operations log:
```
ERROR: Command /etc/init.d/syslog-ng start
failed on node adam.

ERROR: Command /etc/init.d/syslog-ng start
failed on node eve.
```
   To remedy this error, try starting the service manually on each node using the following command:
```
/etc/init.d/syslog-ng start
```
   Once syslog-ng has been started on each node, click "OK" on the "Configure Log Consolidation Tool" page to reapply the configuration from Serviceguard Manager.

To test, try to run the clog package on each node in the cluster. From the “Summary” page, select the clog package, then from the “Administration” menu, select “Run Package” to run the package on the first node, then “Move Package” to run the package on the second node.

Configuring a sample package

In this section, you will configure a Serviceguard for Linux package, called ws, for an Apache web server using scripts from the Apache Toolkit. The content for the web server will reside on the shared volume group, vgws, and will be accessed from a browser at a relocatable IP address. This way the user browsing the content will not need to know which node is hosting the web server. He or she can access the web browser from using a single URL, regardless of which node is hosting the web server.

Configure Package from Serviceguard Manager

1. From your browser on the “HP Serviceguard Manager Summary” page, go to the “Configuration” menu and select “Create a Single Package…”.
2. Under the “Parameters” tab.
   Enter the Package name, for example: **`ws`**
   Keep other defaults (for example, Type = Failover package).
3. Go to the “Monitored Resources” tab.
   a. Check the Subnet, for example: 16.89.84.128
   b. Under “Specify Services”, add the following service:
   ```
   http_monitor
   ```
   c. Click “<<Add”.
   d. Under “Specify Services”, add the following service:
   ```
   diskmon
   ```
   e. Check “Halt Timeout” and enter 300 seconds.
   f. Click “<<Add”.
4. Edit the control script by clicking the “Edit Control Script” button.
   a. For the “Run Script File Path”, use the pre-filled, default path.
   b. Select “Edit”.
   c. Select “Generate Default” (if the file does not exist).

   NOTE: You will be placed in an edit window containing the default control script. Make sure you save your control script changes within 20 minutes or you could lose them and have to start over.

   d. Locate the `VOLUME GROUPS` section, look for “`#VG[0]`”.
   e. Add the volume group for the package, for example:
   ```
   VG[0]=”vgws”
   ```
   f. Locate the `FILESYSTEMS` section, look for “`#LV[0]`”.
   g. Add the following details about the logical volume and file system, for example:
   ```
   LV[0]=/dev/vgws/lvol1;
   FS[0]=/ws
   FS_TYPE[0]="ext3";
   FS_MOUNT_OPT[0]="-o rw";
   FS_UMOUNT[0]="";
   FS_FSCK_OPT[0]="";
   ```
   h. Locate the `IP ADDRESSES` section, look for “`#IP[0]`”.
   i. Add the following details for the package relocatable IP address and subnet, for example:
   ```
   IP[0]="16.89.84.218"
   SUBNET[0]="16.89.84.128"
   ```
   1. Locate the SERVICE NAMES AND COMMANDS section, look for “#SERVICE_NAME[0]”.
   j. Add the following details:
   ```
   SERVICE_NAME[0]="http_monitor"
   ```

```
SERVICE_CMD[0]="/usr/local/cmcluster/conf/ws/toolkit.sh monitor"
SERVICE_RESTART[0]="-r 0"
SERVICE_NAME[1]="diskmon"
SERVICE_CMD[1]="/usr/local/cmcluster/bin/cmresserviced_custom
/dev/mapper/mpath3"
SERVICE_RESTART[1]=""
SERVICE_FAIL_FAST_ENABLED[1]="no"
SERVICE_HALT_TIMEOUT[0]="300"
```

NOTE: Make sure the SERVICE_NAME's match the names of the Services added earlier under "Monitored Resources" in Step #3.

k. Locate the string "`function customer_defined_run_cmds`".

l. Add the following customer defined run command, "/usr/local/cmcluster/conf/ws/toolkit.sh start", such that the function looks like the following:

```
function customer_defined_run_cmds
{
# ADD customer defined run commands.
    : # do nothing instruction, because ..
    /usr/local/cmcluster/conf/ws/toolkit.sh start
    test_return 51
}
```

m. Locate the string "`function customer_defined_halt_cmds`".

n. Add the following customer defined halt command, "/usr/local/cmcluster/conf/ws/toolkit.sh stop", such that the function looks like the following:

```
function customer_defined_halt_cmds
{
# ADD customer defined halt commands.
    : # do nothing instruction, because ..
    /usr/local/cmcluster/conf/ws/toolkit.sh stop
    test_return 52
}
```

5. Then click "Save and Distribute".
6. Check for any error messages in the Operations Log. If none, click "OK".
7. Click "Close" in the "Edit Control Script" window.
8. Click "OK" on the "Create a Single Package" page.
9. Click "Check Configuration".
10. Check for any error messages in the Operations Log. If none, click "OK".
11. Click "Apply Configuration".
12. Click "OK" in the pop-up dialog box.

Customize and Distribute Toolkit Files

To customize and distribute the Toolkit files and the httpd.conf file, perform the following steps:

1. Login to one of the cluster nodes, for example eve.
2. Change to the ws package directory.
   ```
   cd /usr/local/cmcluster/conf/ws
   ```
3. Copy the Apache Toolkit files into the package directory.
   ```
   cp /usr/local/cmcluster/apachetoolkit/* .
   ```
4. Edit hahttp.conf to change following line for the location of the httpd.conf file.
   ```
   HTTPD_CONFIG_FILE="/usr/local/cmcluster/conf/ws/httpd.conf"
   ```
5. Copy the default httpd.conf file to the package directory, for example:
   ```
   cp /etc/httpd/conf/httpd.conf .
   ```
6. Edit the package's httpd.conf file. These tell the Apache instance (httpd) in the "ws" package to use content from the shared volume group (/dev/vgws mounted at /ws) and specifies the listen address to be the package's relocatable IP address. Change the following directives, for example:

```
DocumentRoot "/ws"
Listen 16.89.84.218:80
```

7. Copy the Toolkit files and httpd.conf file to other node, for example adam:

```
scp ha* root@adam:/usr/local/cmcluster/conf/ws/.
scp toolkit.sh root@adam:/usr/local/cmcluster/conf/ws/.
scp httpd.conf root@adam:/usr/local/cmcluster/conf/ws/.
```

Start the Sample Package

You will need to run the package to complete the configuration:

1. From your browser on the "HP Serviceguard Manager Summary" page, select the "ws" package.
2. From the "Administration" menu and select "Run Package".

NOTE: You may need to start the cluster if it is not already running. To run the cluster, go the "Administration" menu and select "Run Cluster".

Create Sample html content for the Package

In the following steps, you will create an html file for the "ws" package.

1. Login to the cluster node that is currently running the "ws" package.
2. Create a sample content file (this command should be input on one line):

```
echo "<html><body> Sample Serviceguard package:
      ws </body></html>" > /ws/index.html
```

NOTE: The sample package must be running and /ws must be mounted in order to create the file.

Verification

Make sure the package is running on the first node, for example, eve.

Test that you can access the ws package content from your browser at the following URL, http://<ws_ip_address>/index.html, for example:
http://16.89.84.218:80/index.html

From another browser window, go to the "Serviceguard Manager Summary" page, move the package to the other node, for example, adam, using the "Move package" option from the "Administration" menu. Select the node you want to move the package to.

Test that you can access the sample web page when the package is running on the other node. In your browser, go to the following URL, http://<ws_ip_address>/index.html, for example:
http://16.89.84.218:80/index.html
Hit the "Refresh" button on your browser to make sure you are still able to access the page.

Troubleshooting

This section contains a list of common errors and potential remedies.

Problem: Cannot boot off local disk after OS installation
Possible Remedy: do not plug FC cables into servers until after the OS is installed. The standard install process may try to place the boot partition on an external storage device if it detect that external storage is connected.

Problem: Storage Management Utility browser interface does not display any "Manage" menu options in the left-hand column
Possible Remedy: Refresh the browser window.

Problem: When installing the Serviceguard patch, the sgupdate command encounters an error. For example:

```
      An update for sgproviders exists - loading.
      Error executing command "rpm -U /tmp/julie/rpms/sgproviders-
      A.02.00.01-0.rhel5.i386.rpm".
      See file /tmp/rpminst.4550 for details.
```

Possible Remedy:

a. Stop and restart the cimserver.

```
cimserver -s
cimserver
```

b. Rerun the update command from the patch "tools" directory:

```
./sgupdate
```

Problem: Error invoking the Serviceguard manager

```
Bad Gateway
The proxy server received an invalid response from an upstream
server.
```

Possible Remedy: The Java path for tomcat is wrong. To fix, make sure you have Java JDK 1.4.2 or greater (Java 5 JDK is recommended). If not, install it. If you do not know the path to Java, you can run the "find" command to find all java files and directories. For example "find / -name java" will search the entire system under "/". Look for the result under the JDK, not the JRE, for example, /usr/java/jdk1.5.0_15/bin/java. Then run /opt/hp/hpsmh/tomcat/bin/tomcat_cfg to input the correct path to the Java.

Problem: Error configuring consolidated log using the Wizard when attempting to apply the configuration.

```
ERROR: Command /etc/init.d/syslog-ng start
failed on node adam.

ERROR: Command /etc/init.d/syslog-ng start
failed on node eve.
```

Possible Remedy: Start the service manually on each node using the following command: "/etc/init.d/syslog-ng start", then re-apply the configuration from Serviceguard Manager.

Support of other distributions and architectures

HP Serviceguard for Linux, version A.11.18 supports Novell SLES 10 and Red Hat Enterprise Linux 4 as well as Red Hat 5. In addition to running on the ProLiant IA32/x86 architecture, it also runs on ProLiant x86_64 and Integrity architectures. The installation and configuration scripts are designed to support those distributions and architectures as well as Red Hat 5 on IA32. It is beyond the scope of this white paper to fully describe the installation steps for Red Hat 4 and SLES10 and the 64-bit architectures.

The key information needed for those distributions is the list of rpms that must be installed from the distribution CDs. These are:

For Red Hat 4: `net-snmp, kernel-devel, glibc-devel, glibc-headers, glibc-kernheaders, kernel_smp_devel, xinetd, and gcc`

For SLES 10: `pidentd, glibc, glibc-locale, glibc-devel, glibc-info, net-snmp, kernel-source, libmudflap, cpp, xinetd, gcc`

As with Red Hat 5, the "development" option should be selected during installation if available.

NOTE: For SLES 10, if the cluster does not autostart please rename /etc/init.d/cmcluster.init to /etc/init.d/cmcluster.

Additionally, there are several different requirements, listed below, for the installation of HP Serviceguard for Linux on the Integrity platforms. These steps should be performed prior to running the Serviceguard for Linux Installation script.

On Integrity, Serviceguard Manager requires the following Java JDK: BEA JRockit 5.0 R27.4 JDK for Linux (Intel Itanium- 64-bit) or later version of JRockit 5.0.

1. To download BEA JRockit:
   a. Go to http://commerce.bea.com/products/weblogicjrockit/jrockit_prod_fam-bea.jsp.

   NOTE: This link was valid at publication time. If the link no longer works, contact HP Support.

   b. Click on the link to download JRockit 5.0.
   c. Agree to the License Terms.
   d. Select the latest JRockit 5.0 available for Intel Itanium – 64-bit, for example:
      JRockit 5.0 R27.5 JDK Linux (Intel Itanium - 64-bit)
2. To install BEA JRockit:
   a. Go to the directory where you have downloaded the file. For example:
      ```
      cd /tmp/sglx_install
      ```
   b. Make the downloaded file executable, for example:
      ```
      chmod +x jrockit-R27.5.0-jdk1.5.0_14-linux-ipf.bin
      ```
   c. For Red Hat 5 Update 1, you must disable SELinux before installing JRockit, otherwise bypass this step.
      ```
      echo 0 > /selinux/enforce
      ```
   d. Execute the installation file, for example:
      ```
      ./jrockit-R27.5.0-jdk1.5.0_14-linux-ipf.bin
      ```
   e. Write down the path where the JDK gets installed, for example:
      ```
      /root/jrockit-R27.5.0-jdk1.5.0_14
      ```
   f. For Red Hat 5 Update 1, you must re-enable SELinux, otherwise bypass this step.
      ```
      echo 1 > /selinux/enforce
      ```

For Serviceguard Manager on Integrity, HP SMH and Tomcat software must be acquired and installed separately. They are not included on the Serviceguard for Linux CD.

1. To download HP SMH and Tomcat, get the latest HP Integrity Essentials Foundation Pack for Linux Support Pack from HP.
   a. Go to: http://www.hp.com.
   b. Click on "Software & Driver Downloads".
   c. Select "Download drivers and software (and firmware)".
   d. Enter your server model in the "for product" box.
   e. Click ">>".
   f. Select your distribution.
   g. Find "HP Integrity Essentials Foundation Pack for Linux Support Pack" on that page and download the latest version. At publication time, the latest version of the Support Pack was 4.27 (5 May 2008); the downloaded file is SupportPack-4.27-sles10.ia64.tar.

NOTE: This link was valid at publication time. If the link no longer works, contact HP Support.

2. To install HP SMH and Tomcat:
   a. Go to the directory where you have downloaded the file. For example:
      ```
      cd /tmp/sglx_install
      ```
   b. Use the tar command to extract the files. For example:
      ```
      tar xf SupportPack-4.27.sles10.ia64.tar
      ```
   c. Change to the new directory that is created. For example:
      ```
      cd SupportPack-4.27/distros/sles10
      ```
   d. Install HP SMH. For example:
      ```
      rpm –Uvh  hpsmh-2.1.10-08.linux.ia64.rpm
      ```
   e. Install Tomcat. For example:
      ```
      rpm –Uvh hpsmh-tomcat-1.0-23.linux.ia64.rpm
      ```

      NOTE: Be sure to get the latest patch for Serviceguard Manager. There is a problem in which HP SMH does not show the Serviceguard Manager link under the Tools tab.

   f. If necessary, run tomcat_cfg to complete the installation by providing the path to the JRockit JDK when prompted.
      ```
      /opt/hp/hpsmh/tomcat/bin/tomcat_cfg
      ```

For more information

Learn more about HP Serviceguard for Linux at http://www.hp.com/go/sglx, or consult with your HP Sales Representative or Partner.