

HP Serviceguard Version A.11.18 Release Notes

HP Part Number: B3935-90119
Published: April 2008



Legal Notices

© Copyright 1998-2007 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

Table of Contents

Printing History	11
1 Serviceguard Version A.11.18 Release Notes.....	13
Announcements.....	13
Platform Dependencies.....	13
April 2008 Patch.....	13
Serviceguard Bundled Components - New Product Structure.....	13
Advantage of the New Product Structure.....	13
Serviceguard Optional Products Not Bundled.....	13
SMS A.01.01 or Later Required with June 2007 or Later MCOE.....	13
Support for Veritas 5.0.....	14
New Method for Configuring Packages.....	14
New Online Cluster-Management Capabilities.....	14
Serviceguard Manager Available from the System Management Homepage (SMH).....	14
Version 3.0 or 3.5 of HPVM Required.....	14
RS232 Heartbeat Obsolete.....	14
Token Ring and FDDI Obsolete.....	15
Parallel SCSI Dual Cluster Lock Obsolete.....	15
Parallel SCSI Not Supported for Lock LUN.....	15
Cluster Name Restrictions.....	15
Optimizing Performance when Activating LVM Volume Groups.....	15
High Availability Consulting Services.....	16
Announcements for HP-UX 11i v2.....	16
Possible Problem when Upgrading from HP-UX 11i v1 to 11i v2.....	16
Use -T option to identd.....	16
New Features for April 2008 Patch.....	17
Announcements for HP-UX 11i v3.....	17
New Bundles.....	17
Support for Veritas Volume Manager (VxVM).....	17
Native Multipathing, Veritas DMP, and Related Features.....	17
PV Links.....	18
PCI Error Recovery.....	18
Online Replacement of LAN Cards Requires Patch.....	18
New Features for April 2008 Patch.....	18
What's in this Release.....	18
New Serviceguard Features Common to HP-UX 11i v2 and 11i v3.....	19
New Serviceguard Features Available Only on HP-UX 11i v3.....	21
Changes First Introduced in Earlier Releases.....	21
What's Not in this Release.....	22
Support for Veritas 5.0 on HP-UX 11i v2 and 11i v3.....	22
About Modular Packages.....	22
Package Modules.....	23
Base Package Modules.....	23
Optional Package Modules.....	24
New Parameters.....	26
<i>module_name</i>	26
<i>module_version</i>	26
<i>successor_halt_timeout</i>	26
<i>operation_sequence</i>	27
<i>log_level</i>	27

<i>priority</i>	27
<i>monitored_subnet_access</i>	27
<i>ip_subnet_node</i>	27
<i>cluster_interconnect_subnet</i>	27
<i>enable_threaded_vgchange</i>	28
<i>vxvm_dg_retry</i>	28
<i>pev</i>	28
<i>external_pre_script</i>	28
<i>external_script</i>	29
Legacy and Modular Package-Configuration Parameters.....	29
Parameters whose Formal Requirements Have Changed.....	31
Migrating Legacy to Modular Packages.....	31
About Package Dependencies.....	31
About Multi-Node Packages.....	32
Differences between Failover and Multi-Node packages.....	32
Other Package Changes.....	32
About Cross-Subnet Configurations.....	33
Configuration Tasks.....	33
Restrictions.....	34
For More Information.....	34
New Online Networking-Configuration Capabilities.....	35
What You Can Do.....	35
What You Must Keep in Mind.....	35
New -c Option in cmquerycl.....	36
New Online Timing-Configuration Capabilities.....	37
About Lock LUNs.....	37
About Device Special Files (DSFs).....	37
About LVM 2.0.....	38
About olrad.....	39
About HPVM and Cluster Re-formation Time.....	39
About the Volume Monitor.....	40
Monitor Usage.....	40
Syntax.....	40
Examples.....	41
Scope of Monitoring.....	41
About vgchange -T.....	41
Alternate Quorum Server Subnet.....	42
Access changes as of A.11.16.....	42
Considerations when Upgrading Serviceguard.....	42
Considerations when Installing Serviceguard.....	43
About the Serviceguard Manager SMH Plug-In.....	43
DSAU Integration.....	43
Native Language Support.....	43
What You Can Do.....	43
New Features.....	43
Differences from A.05.01.....	44
Limitations.....	44
Help Subsystem.....	45
Before Using HP Serviceguard Manager: Setting Up.....	45
Launching Serviceguard Manager.....	45
Scenario 1 - Single cluster management.....	46
Scenario 2 -Multi-cluster management.....	48
Scenario 3 - Web-based multi-cluster management.....	49
Patches and Fixes.....	51
What Documents are Available for This Version.....	51

Further Information.....	52
Release Notes Revisions.....	53
Compatibility Information and Installation Requirements.....	53
Compatibility.....	53
Compatibility with Storage Devices.....	53
Bastille Compatibility.....	53
Mixed Clusters.....	53
Before Installing Serviceguard A.11.18	54
Memory Requirements.....	54
Port Requirements.....	54
System Firewalls.....	55
Installing Serviceguard on HP-UX.....	55
Dependencies.....	55
Installing Serviceguard.....	55
If You Need To Disable identd.....	56
Upgrading from an Earlier Serviceguard Release.....	56
Veritas Storage Management Products.....	57
Rolling Upgrade.....	57
Requirements for Rolling Upgrade.....	57
Rolling Upgrade Exceptions.....	57
HP-UX Cold Install.....	57
HP Serviceguard Storage Management Suite and standalone CVM product.....	57
Migrating to Agile Addressing if Using Cluster Lock Disk.....	57
Upgrading from an Earlier Release if You Are Not Using Rolling Upgrade (Non-Rolling Upgrade).....	58
Uninstalling Serviceguard.....	58
Patches for this Version.....	58
Fixed in This Version.....	59
JAGaf46654 (SR8606386500): cmclnd: software causes connection abort if standby LAN card fails.....	59
JAGaf67601 (SR8606407696): cmclnd does not consider newer drivers e.g. c8xx and mpt for cluster lock.....	59
JAGaf75382 (SR8606415528): lan2 has an IPv6 address (fec0:0:0:f08::36) but not in config. file.....	59
JAGag09818 (SR8606453031): Cannot modify SERVICE_HALT_TIMEOUT to default value (0 seconds).....	60
JAGag11398 (SR8606454772): cmclnd aborts with Failed to set keep alive: Invalid argument.....	60
JAGag12644 (SR8606456223): cmsrvassistd doesn't check script permissions.....	60
JAGag20225 (SR8606464542): Serviceguard probe to CD, DVD devices.....	60
JAGag21443 (SR8606465899): cmclnd aborts when select() is interrupted.....	60
JAGag25508 (SR8606470417): Log file permissions set to 666.....	61
JAGag35297 (SR8606481305): cfsmount failed since the package was already started by cmclnd....	61
JAGag41937 (SR8606489376): Hanging node problem causing 2 out of 4 nodes to die.....	61
JAGag48756 (SR8606496838): Serviceguard minimum lan requirements not updated for VM detection information.....	61
JAGag46475 (SR8606494542): cmquerycl -f line output may display incorrect VG disk device files.....	61
JAGag47195 (SR8606495273): Improper log message from cmmodnet on a system on which IPv6 is not enabled.....	62
JAGag48606 (SR8606496688): cmmigratepkg missing PEV and external script permission.....	62
QXCR1000594093: cmcheckconf failed to detect cluster lock error.....	62
QXCR1000739586: Serviceguard does not log .cm_start_time message correctly in syslog after abort.....	62
QXCR1000740150: cmclnd abort when deleting node online.....	63
QXCR1000740395: Should display warnings about subnets only for related packages.....	63
QXCR1000741571: Hanging node problem causing cluster to go down.....	63

QXCR1000746305: Site controller pkg failed to start after cmhaltnode -f.....	63
QXCR1000747823: cmviewconf does not display managed packages for SC package.....	63
QXCR1000749527: cmmakepkg -l does not list non-SG modules.....	63
QXCR1000749924: Dependent packages started on non-primary nodes.....	64
QXCR1000766057: cmviewcl package state "blocked" is not documented.....	64
QXCR1000766101: cmhaltpkg reports that auto_run is disabled for multi-node packages but it is not.....	64
Problems Fixed in this Version of the Serviceguard Manager Plug-in.....	64
JAGag05829: Operations log window should point to the operation just performed.....	64
JAGag05831: A few operations do not immediately show in the Operations Log.....	65
JAGag07820: Handling invalid cluster name with * / \ space characters is broken.....	65
JAGag07875: Edit cluster summary page: warning about pkg config should be restricted.....	65
JAGaf97851: Serviceguard Manager hangs when launching System Log Viewer	65
JAGag11832 - Launching consolidated syslog first time causes SMH logout.....	65
QXCR1000581069: Deleting multiple packages failed if a selected package was not down.....	65
QXCR1000591254: Tooltip value incorrect.....	65
QXCR1000741159: Incorrect error message.....	66
QXCR1000752875: Using Serviceguard Manager to edit a CFS package with environment variables corrupts the package configuration file.....	66
QXCR1000765143 SGMgr failed to discover Oracle DB and instance on 10G and 11G.....	66
Known Problems and Workarounds.....	66
JAGad39695 (SR8606170431):User error can result in "ghost" services:	66
JAGad54887 (SR8606185685): cmquerycl command may hang when probing disabled disks....	67
JAGaf32484 (SR8606372066): Problems doing a rolling upgrade when moving from CVM 3.2 to 3.5.....	67
JAGaf79389 (SR8606419559): Package not started on adoptive node after cmhaltnode -f.....	68
JAGag27186 (SR8606472298): OLR of a LAN Card in SG cluster fails on HP-UX 11i v3.....	68
JAGag31490 (SR8606477006): Modular package fails with the error: 'Unable to retrieve package attribute'.....	68
JAGag43145 (SR8606490767): cmcheckconf fails for non-existent deferred filesystem EMS resources.....	69
JAGag44319 (SR8606492148): cmcluster boot script does not print cmrunnode errors	69
QXCR1000768063: Serviceguard A.11.18 network validation errors missing.....	69
QXCR1000770172: cmgetconf fails if some PV Links are missing	69
Known Problems and Workarounds for the Serviceguard Manager Plug-in.....	69
JAGag35535: Adding a node to a cluster with Lock LUN fails.....	70
JAGag34285: Deleting multiple packages will fail if a selected package is not down.....	70
JAGag04773: Clicking on the browser refresh button takes you back to main page.....	70
JAGag32353: With Mozilla or Firefox browser, configuring a SGeRAC package is confusing.....	70
JAGag22514: Operations Log focus returns to previous selection after operation is over.....	70
JAGag05524: When editing roles and monitored resources, some buttons may be disabled....	70
JAGag29458: In matrix views, empty columns are not displayed.....	71
JAGag32176: With Mozilla 1.5, Auto Refresh button appears after first auto-refresh.....	71
JAGag31375: SGeSAP missing from cluster extensions.....	71
JAGag28015: All dropdown menu options become enabled when clog is re-launched.....	71
About Serviceguard Releases.....	71
Types of Releases and Patches.....	71
Platform Release.....	71
Feature Release.....	71
Patch.....	72
Supported Releases.....	72
Version Numbering.....	72
Software Availability in Native Languages	72

List of Figures

1-1	System Management Homepage Login.....	46
1-2	Serviceguard Cluster.....	47
1-3	System Management Homepage with Serviceguard Manager.....	48
1-4	Systems Insight Manager Sign-In Page.....	49
1-5	Expanded Left Panel.....	50
1-6	Serviceguard Manager Page.....	51

List of Tables

1	Printing History.....	11
1-1	Base Modules.....	24
1-2	Optional Modules.....	24
1-3	Legacy Configuration-File Parameters and Modular Equivalents.....	29
1-4	Legacy Control-Script Parameters and Modular Equivalents.....	30
1-5	Accessing Serviceguard Manager.....	45
1-6	Serviceguard Manager Installation Requirements.....	46
1-7	59

Printing History

Table 1 Printing History

Printing Date	Part Number	Edition
June 2007	B3935-90108	First Edition
October 2007	B3935-90108	Second Edition
December 2007	B3935-90114	Third Edition
April 2008	B3935-90119	Fourth Edition

The fourth edition includes information on new features included in patches PHSS_37601 and PHSS_37602. See “New Features for April 2008 Patch” (page 17) (for HP-UX 11i v2) and on (page 18) (for HP-UX 11i v3).

1 Serviceguard Version A.11.18 Release Notes

Announcements

This section announces the most important features and limitations of Serviceguard A.11.18. For more information, see “What’s in this Release” (page 18).

Platform Dependencies

This new version of Serviceguard is supported on HP-UX 11i v2 and 11i v3. There are some platform dependencies; see the following subsections for details:

- “Announcements for HP-UX 11i v2” (page 16)
- “Announcements for HP-UX 11i v3” (page 17)
- “New Serviceguard Features Available Only on HP-UX 11i v3” (page 21)

April 2008 Patch

For information about new features introduced in the March 2008 patches, see “New Features for April 2008 Patch” (page 17) (for 11i v2: patch PHSS_37601) and (page 18) (for 11i v3: patch PHSS_37602).

Serviceguard Bundled Components - New Product Structure

The following Serviceguard components are now available for HP-UX HP-UX 11i v2 and 11i v3:

- Product T1905CA — A.11.18 — software and license
- Product T1906CA — A.11.18 — documentation

The following products are bundled with Serviceguard, as **components** of the Serviceguard product, T1905CA:

- The new Serviceguard Manager web-based graphical user interface (GUI), version B.01.01. This version is also known as the **SMH Plug-in**, or simply **plug-in**, version. See “About the Serviceguard Manager SMH Plug-In” (page 43).
- The Cluster Object Manager (COM), version B.05.00.
- The SGWEBEM Providers product, version A.02.00.

Advantage of the New Product Structure

This new product structure allows Serviceguard customers who subscribe to Update services to download Serviceguard from the web (along with the three sub-products classed as Serviceguard components and listed above) via Software Update Manager (SUM).

Serviceguard Optional Products Not Bundled

The following optional product is not bundled with Serviceguard, but is delivered free with Serviceguard on the Serviceguard Distributed Components CD, and can also be downloaded from <http://www.hp.com/go/softwaredepot/ha>:

- Quorum Server, and the *Quorum Server Version A.03.00 Release Notes*.

SMS A.01.01 or Later Required with June 2007 or Later MCOE

If you are running one of the HP Serviceguard Storage Management Suite (SMS) version A.01.00 MCOE bundles (T2794BA, T2795BA, T2796BA, T2797BA) do not attempt an upgrade to an HP-UX Mission Critical Operating Environment (MCOE) version dated June 2007 or later, without also upgrading to a compatible version of SMS.

The upgrade will complete without errors, but it will not be possible to start the cluster. This is because the SMS version A.01.00 MCOE bundles support only Serviceguard version A.11.17, but the MCOE includes Serviceguard version A.11.18 after the June 2007 release.

If you need to update to the June 2007 or later version of the MCOE, you must upgrade to a compatible version of SMS (A.01.01 or A.02.00) as well. For instructions, see the *HP Serviceguard Storage Management Suite Release Notes* for your version of SMS at docs.hp.com -> High Availability -> HP Serviceguard Storage Management Suite.

Support for Veritas 5.0

Serviceguard A.11.18 now supports version 5.0 of Veritas VxVM, CVM and CFS from Symantec, on both HP-UX 11i v2 and 11i v3. See “Support for Veritas 5.0 on HP-UX 11i v2 and 11i v3” (page 22).

New Method for Configuring Packages

Serviceguard A.11.18 introduces a new process for configuring packages. Packages created by this method are known as **modular** packages, because you can now build a package using only the sets of functions (modules) that this package needs; and you no longer need to create and distribute a separate package control script.

Packages created using previous versions of Serviceguard are referred to as **legacy** packages and are still supported; you can continue to maintain as well as run them.

See “About Modular Packages” (page 22).

New Online Cluster-Management Capabilities

Serviceguard A.11.18 introduces important new online management capabilities, particularly as regards the cluster timing parameters and networking configuration. These new capabilities significantly reduce the number of maintenance tasks that require cluster downtime.

See “New Online Networking-Configuration Capabilities” (page 35) and “New Online Timing-Configuration Capabilities” (page 37).

Serviceguard Manager Available from the System Management Homepage (SMH)

For details, see “What’s in this Release” (page 18) and “About the Serviceguard Manager SMH Plug-In” (page 43).



NOTE: The earlier, station-management version of Serviceguard Manager is obsolete as of Serviceguard A.11.18. The last version of Serviceguard that supports that product is A.11.17 (A.11.17.01 on HP-UX 11i v3).

Version 3.0 or 3.5 of HPVM Required

If you intend to use **HP Integrity Virtual Machines (HPVM)** with Serviceguard A.11.18, you must install or upgrade to HPVM 3.0 or 3.5; these are the only versions currently supported with Serviceguard A.11.18. See “Changes First Introduced in Earlier Releases” (page 21) for more information about HPVM.

RS232 Heartbeat Obsolete

The current version of Serviceguard does not support RS232 for the cluster heartbeat (Serviceguard A.11.17 on HP-UX 11i v2 and A.11.16 on HP-UX 11i v1 were the last versions that did).

A minimum configuration of Serviceguard on HP-UX 11i v2 or 11i v3 needs two network interface cards for the heartbeat in all cases, using one of the following configurations:

- two heartbeat subnets; *or*
- one heartbeat subnet with a standby; *or*
- one heartbeat subnet using APA with two physical ports in hot standby mode or LAN monitor mode.

To provide the required redundancy for both networking and mass storage connectivity on servers with fewer than three I/O slots, you may need to use multifunction I/O cards that contain both networking and mass storage ports.

Token Ring and FDDI Obsolete

Serviceguard A.11.18 does not support Token Ring and FDDI technologies for the cluster heartbeat and data networks (Serviceguard A.11.17 on HP-UX 11i v2 and A.11.16 on HP-UX 11i v1 were the last versions that did). HP-UX 11i v3 does not support these two technologies.

The unsupported configurations include physical Token Ring and FDDI interfaces, Virtual LAN (VLAN) interfaces over FDDI or Token Ring, and failover groups of Token Ring and FDDI interfaces in the LAN Monitor Mode of the APA product.

Parallel SCSI Dual Cluster Lock Obsolete

As of Serviceguard A.11.18, you must use Fibre Channel connections for a dual cluster lock; you can no longer implement it in a parallel SCSI configuration. See “Dual Lock Disk” in Chapter 3 of *Managing Serviceguard* for more information about dual cluster locks.

Parallel SCSI Not Supported for Lock LUN

The new lock LUN functionality does not support parallel SCSI; you must use Fibre Channel for a lock LUN. If you need to use parallel SCSI, use an LVM cluster lock disk, or a Quorum Server.

See “About Lock LUNs” (page 37) for more information about lock LUNs. For more information about the cluster lock, see “Cluster Lock” in Chapter 3 of *Managing Serviceguard*.

Cluster Name Restrictions

The following characters must not be used in the cluster name if you are using the Quorum Server: at-sign (@), equal-sign (=), or-sign (|), semicolon (;).

These characters are deprecated, meaning that you should not use them, even if you are not using the Quorum Server, because they will be illegal in a future Serviceguard release. Future releases will require the cluster name to:

- Begin and end with an alphanumeric character
- Otherwise use only alphanumeric characters, or dot (.), hyphen (-), or underscore (_)

Optimizing Performance when Activating LVM Volume Groups

If a package activates a large number of volume groups, you can improve the package’s start-up and shutdown performance by carefully tuning the *concurrent_vgchange_operations* parameter in the package configuration file (or control script, for legacy packages).

Tune performance by increasing this parameter a little at a time and monitoring the effect on performance at each step; stop increasing it, or reset it to a lower level, as soon as performance starts to level off or decline.

Factors you need to take into account include the number of CPUs, the amount of available memory, the HP-UX kernel settings for *nfile* and *nproc*, and the number and characteristics of other packages that will be running on the node.



NOTE: Remember to do this exercise not only on the node on which the package will normally run, but also on the node with the least resources in the cluster, as a failover or other unexpected circumstances could result in that node running the package.

For more information, see the section “Optimizing for Large Numbers of Storage Units” in Chapter 6 of the latest edition of *Managing Serviceguard* (in the High Availability collection on docs.hp.com) and the comments in the package configuration file.

High Availability Consulting Services

Because Serviceguard configurations can be complex to configure and maintain, HP strongly recommends that you use its high availability consulting services to ensure a smooth installation and rollout; contact your HP representative for more information. You should also work with your HP representative to ensure that you have the latest firmware revisions for disk drives, disk controllers, LAN controllers, and other hardware.

Announcements for HP-UX 11i v2

Possible Problem when Upgrading from HP-UX 11i v1 to 11i v2

Under certain rare conditions the Serviceguard daemon `/usr/sbin/cmclconfd` could abort with a core dump during a system update from HP-UX 11i v1 to 11i v2 on a system on which EMS is installed.

This can happen if the Serviceguard node being updated receives a Serviceguard probe request from another Serviceguard node on the subnet while the EMS product is being updated.

During the EMS product update, the EMS libraries on which `cmclconfd` depends are removed and re-installed, leaving a very short period during which the system does not have an EMS library.

If the node receives a Serviceguard probe request during this time, the system will fail to load the EMS library and will send a `SIGABORT` signal to `cmclconfd`. The core dump created as a result of this is harmless and does not affect the product or system functionality in any way.

Once the EMS library is installed, `cmclconfd` will provide normal functionality.

Use `-T` option to `identd`

A new option, `-T`, is available for the HP-UX daemon `identd`. You can obtain this version of `identd` for an HP-UX 11i v2 system by downloading Sendmail release 8.13.3 (version B.11.23.1.005 or later): go to <http://www.software.hp.com> -> Internet ready and networking and choose Sendmail 8.13.3 Special Release Upgrade.

The download includes Release Notes; follow directions in that document to install the upgrade.

Once you have downloaded the software, set the `identd -T` option to 120 (`-T120`); this ensures that a connection inadvertently left open will be closed after two minutes. `identd` is usually started by `inetd` from `/etc/inetd.conf`, so you would normally edit that file to set the `identd -T` option.

The `-T` option may also be available in the form of a patch to earlier Sendmail versions; or for later HP-UX versions. (And if you are still running some clusters under Serviceguard A.11.16 on HP-UX 11i v1, note that the `-T` option *is* available for that configuration.)

Check the man page for `identd` to determine whether the `-T` option is supported for the version of `identd` currently on your system.

For more information, see:

- The section “Ensuring that the Root User on Another Node Is Recognized” in Chapter 5 of the 15th edition of *Managing Serviceguard* at <http://docs.hp.com> -> High Availability -> Serviceguard
- The section “What’s In This Version” in the *Sendmail 8.13.3 Release Notes*.
- The `identd` (1M) manpage.

New Features for April 2008 Patch

Patch PHSS_37601 for Serviceguard A.11.18 on HP-UX 11i v2 is needed to enable some of the features discussed in these Release Notes. See the following for more information about the features in question:

- “About Cross-Subnet Configurations” (page 33)
- “Migrating Legacy to Modular Packages” (page 31)
- “Other Package Changes” (page 32)
- “About the Volume Monitor” (page 40)
- “Alternate Quorum Server Subnet” (page 42)

Announcements for HP-UX 11i v3

This release of Serviceguard A.11.18 supports HP-UX 11i v3, which introduces important improvements, particularly in regard to the I/O subsystem. See the subsections that follow, and “What’s in this Release” (page 18).

New Bundles

Serviceguard A.11.18 is now available as a recommended product in the HP-UX 11i v3 HA-OE and DC-OE bundles.

Support for Veritas Volume Manager (VxVM)

Veritas (VxVM) 4.1 and 5.0 from Symantec are supported on HP-UX 11i v3, but VxVM 3.5 is not. If you are running VxVM 3.5 (as part of the Base Product included with HP-UX), you can do a **rolling upgrade** to Serviceguard A.11.18 on HP-UX 11i v3 with VxVM 4.1 or later. See “Upgrading from an Earlier Serviceguard Release” (page 56) for more information. See also “Support for Veritas 5.0 on HP-UX 11i v2 and 11i v3” (page 22).



NOTE: A file created with VxFS 5.0 cannot be opened by the 4.1 version.

Native Multipathing, Veritas DMP, and Related Features

The HP-UX 11i v3 I/O subsystem provides multipathing and load balancing by default. This is often referred to as **native multipathing**.

Veritas Volume Manager (VxVM) version 4.1 and Dynamic Multipathing (DMP) from Symantec are supported on HP-UX 11i v3, but *do not* provide multipathing and load balancing; DMP acts as a pass-through driver, allowing multipathing and load balancing to be controlled by the HP-UX I/O subsystem instead.

When you upgrade a system to HP-UX 11i v3, the I/O subsystem by default will start performing load balancing and multipathing for all multipath devices (whether or not they are managed by VxVM/DMP, and whether or not you decide to migrate the system to **agile addressing**); you do not have to take any additional steps to make this happen.

For more information about multipathing in HP-UX 11i v3, see the white paper *HP-UX 11i v3 Native Multipathing for Mass Storage*, and the *Logical Volume Management* volume of the HP-UX *System Administrator’s Guide* in the HP-UX 11i v3 Operating Environments collection at <http://docs.hp.com>. See also “About Device Special Files (DSFs)” (page 37).

PV Links

Previous editions of *Managing Serviceguard* recommended creating LVM PV links for LUNs defined in disk arrays. On HP-UX 11i v3, these PV Links are redundant; they are supported, but will be inactive unless you turn off native multipathing (in that case they will function as they did in previous releases of HP-UX).

PCI Error Recovery

PCI Error Recovery enables an HP-UX system to detect, isolate, and automatically recover from a PCI error. PCI Error Recovery is enabled on HP-UX 11i v3 systems by default, but HP recommends that it remain enabled in a Serviceguard cluster only if your storage devices are configured with multiple paths and you have not disabled HP-UX native multipathing.



IMPORTANT: If your storage devices are configured with only a single path, or you have disabled multipathing, you should disable PCI Error Recovery; otherwise Serviceguard may not detect when connectivity is lost and cause a failover.

For instructions on using the `pci_eh_enable` parameter to disable PCI Error Recovery, see the “Tunable Kernel Parameters” section of the latest edition of the *PCI Error Recovery Product Note* in the High Availability collection on docs.hp.com.

Online Replacement of LAN Cards Requires Patch

Before you can replace a LAN card online (without bringing down the cluster) as described under “Replacing LAN or Fibre Channel Cards” in Chapter 8 of *Managing Serviceguard*, you must apply patch PHNE_35894. See “Patches for this Version” (page 58) and “JAGag27186 (SR8606472298): OLR of a LAN Card in SG cluster fails on HP-UX 11i v3” (page 68).

For more information about replacing interface cards online, see the *Interface Card OL* Support Guide* for HP-UX 11i v3, at <http://docs.hp.com> -> High Availability -> System Administration.

New Features for April 2008 Patch

Patch PHSS_37602 for Serviceguard A.11.18 on HP-UX 11i v3 is needed to enable some of the features discussed in these Release Notes. See the following for more information about the features in question:

- “About Cross-Subnet Configurations” (page 33)
- “About vgchange -T” (page 41)
- “Migrating Legacy to Modular Packages” (page 31)
- “Other Package Changes” (page 32)
- “About LVM 2.0” (page 38)
- “About the Volume Monitor” (page 40)
- “Alternate Quorum Server Subnet” (page 42)

What’s in this Release

This release of Serviceguard A.11.18 runs on both HP-UX 11i v2 and 11i v3, and adds new functionality. There are some platform-dependent differences.

See the subsections that follow for details; see also the “Announcements” (page 13). For information about documentation, see “What Documents are Available for This Version” (page 51).

New Serviceguard Features Common to HP-UX 11i v2 and 11i v3

- Serviceguard A.11.18 on HP-UX 11i v2 and 11i v3 supports version 5.0 of the Veritas Volume Manager (VxVM), the Veritas Cluster Volume Manager (CVM), and the Veritas Cluster File System (CFS) from Symantec.

The 4.1 versions of these products are also supported with Serviceguard A.11.18.

See “Support for Veritas 5.0 on HP-UX 11i v2 and 11i v3” (page 22).

- There is a new, **modular** method of configuring packages; see “About Modular Packages” (page 22).
 - In addition, in the patches listed in the sections headed “New Features for April 2008 Patch” (page 17) and (page 18), Serviceguard provides a means of automatically migrating existing (**legacy**) packages to modular packages. See “Migrating Legacy to Modular Packages” (page 31).

(Additional changes specifically support migrating HP Toolkits to the modular method. The converted versions will be provided in support of a future Serviceguard release; do not try to convert HP Toolkit packages yourself.)
- You can now configure a package to be dependent on another package. See “About Package Dependencies” (page 31).
- You can now create **multi-node** packages. See “About Multi-Node Packages” (page 32).
- Other new package capabilities are provided in the patches listed in the sections headed “New Features for April 2008 Patch” (page 17) and (page 18); see “Other Package Changes” (page 32).
- A new Serviceguard Manager graphical user interface (GUI) is being supplied as a plug-in to the System Management Homepage (SMH). The standalone version of Serviceguard Manager (also known as the **station-management** version) is obsolete.

See “About the Serviceguard Manager SMH Plug-In” (page 43).
- There are new networking capabilities:
 - Serviceguard now supports **cross-subnet** configurations.

The patches listed on (page 17) and (page 18) allow you to configure multiple subnets, joined by a router, both for the cluster heartbeat and for data, with some nodes using one subnet and some another. See “About Cross-Subnet Configurations” (page 33).

 - Additional changes (including the addition of the parameters *SITE_NAME* and *SITE* to the cluster configuration file) support **site-aware disaster-tolerant clusters**, which require additional software (see “For More Information” (page 34)).
 - You can now make significant changes to the cluster’s networking configuration while the cluster is running; see “New Online Networking-Configuration Capabilities” (page 35).
 - Serviceguard now supports the *LACP_AUTO* (IEEE 802.3ad Link Aggregation Control Protocol) mode of Auto Port Aggregation (APA), with up to 32 ports per link aggregate; and supports up to eight ports per link aggregate for the *MANUAL* and *FEC_AUTO* (Cisco proprietary Fast EtherChannel) modes. (The former maximum for *MANUAL* and *FEC_AUTO* was 4 ports per link aggregate, and *LACP_AUTO* mode was not supported.)

For information about implementing APA with Serviceguard, see the latest version of the *HP Auto Port Aggregation (APA) Support Guide* and other APA documents posted at docs.hp.com in the IO Cards and Networking Software collection.

- Serviceguard now supports Virtual LANs (VLANs) over APA. (VLANs over physical interfaces were already supported.)

For more information, see the sections “Automatic Port Aggregation” and “Configuring VLANs” in Chapter 3 of the latest edition of *Managing Serviceguard* in the High Availability collection on docs.hp.com.

- You can now use `cmmodnet` to enable or disable a `LAN_name` currently configured in a cluster. (Note that if you use `cmmodnet` to disable a `LAN_name` on a node that subsequently leaves the cluster, the `LAN_name` will be re-enabled if the node then rejoins the cluster.) Extreme caution should be exercised when executing this command outside the context of a package; see the entry for `cmmodnet` in Appendix A of *Managing Serviceguard*, and the `cmmodnet (1m)` manpage.
- You can now change cluster timing-configuration parameters while the cluster is running; see “New Online Timing-Configuration Capabilities” (page 37).
- You can now configure a LUN as the cluster lock. See “About Lock LUNs” (page 37).
- Hostnames for Serviceguard cluster nodes can be up to 39 characters (bytes) long. (The former limit was 31 characters.)



NOTE: On HP-UX 11i v2 systems, the `vgdisplay` command will display only the first 31 characters.

- A new volume monitor for VxVM and CVM is provided in the patches listed in the sections headed “New Features for April 2008 Patch” (page 17) and (page 18). See “About the Volume Monitor” (page 40).
- You can configure an alternate subnet for communication between the cluster nodes and a Quorum Server. See “Alternate Quorum Server Subnet” (page 42).
- A quorum server can now support up to 150 clusters comprising up to 300 nodes. The previous limits were 50 clusters and 100 nodes. The new limits are supported in Quorum Server versions A.02.04 and later.
- There are new Serviceguard demons:
 - Lock LUN Daemon: `cmdisklockd`
If a lock LUN is being used, `cmdisklockd` runs on each node in the cluster and is started by `cmcl` when the node joins the cluster.
 - Network Manager Daemon: `cmnetd`
Runs on every node on which `cmcl` is running.
 - Utility Daemon: `cmlockd`
Runs on every node on which `cmcl` is running (though currently not actually used by Serviceguard on HP-UX systems).
- Serviceguard A.11.18 uses a new cluster lock algorithm, with a standard delay time irrespective of the type of bus or Host Bus Adapter (HBA) being used by the lock device. This is possible because many older, slower buses and HBAs are no longer supported by any version of HP-UX that supports Serviceguard A.11.18. In some cases, the new algorithm results in a substantial reduction in the overall cluster re-formation time.

New Serviceguard Features Available Only on HP-UX 11i v3

- Serviceguard A.11.18 on HP-UX 11i v3 supports HP-UX **agile addressing**, sometimes also called **persistent LUN binding**, for device special files (DSFs).
See “About Device Special Files (DSFs)” (page 37).
- Serviceguard A.11.18 on HP-UX 11i v3 supports HP-UX native multipathing and load balancing.
See “Native Multipathing, Veritas DMP, and Related Features” (page 17).
- There are new networking capabilities:
 - The HP-UX `olrad -C` command now identifies network interface cards (NICs) that are part of the Serviceguard cluster configuration.
You can remove a NIC from the cluster configuration, and then from the system, without bringing down the cluster. See “About olrad” (page 39).
 - Serviceguard now also supports the LAN Monitor mode of APA.
- Serviceguard supports Process IDs (PIDs) of any size, up to the maximum value supported by HP-UX and the node’s underlying hardware architecture.
Previous versions of HP-UX imposed a limit of 30,000; this limit has been removed as of HP-UX 11i v3. For more information, see the white paper *Number of Processes and Process ID Values on HP-UX* on docs.hp.com.
- Serviceguard now supports cell OL* (online addition and deletion of cells) on HP Integrity servers that support them).
For more information about using Serviceguard with partitioned systems, see the white paper *HP Serviceguard Cluster Configuration for HP-UX 11i or Linux Partitioned Systems* at <http://docs.hp.com> -> High Availability -> Serviceguard ->White Papers.
- You can now use the FSWeb utility to configure LVM volumes in a Serviceguard cluster (and SLVM volumes if the add-on product Serviceguard Extension for Real Application Cluster (SGeRAC) is installed).
For more information about FSWeb, see the `fsweb (1m)` manpage.
- Serviceguard A.11.18, with the patch listed under “New Features for April 2008 Patch” (page 18), supports LVM 2.0 volume groups. See “About LVM 2.0” (page 38).

Changes First Introduced in Earlier Releases

The following important changes have either been introduced since the previous edition of the Serviceguard Release Notes was published, or may affect the upgrade from an earlier release.

- As of A.11.16, Serviceguard supports **HP Integrity Virtual Machines (HPVM)**.



IMPORTANT: Serviceguard A.11.18 supports *only* **version 3.0 or 3.5** of HPVM.

See the *Serviceguard/SGeRAC/SMS/Serviceguard Mgr Plug-in Compatibility and Feature Matrix*, on docs.hp.com -> High Availability, under the heading Serviceguard -> Support Matrixes, for the most up-to-date compatibility information.

Serviceguard A.11.18 supports an HPVM either as a package or as a cluster node. If any Serviceguard cluster node is a virtual machine, the amount of time Serviceguard needs to wait for a failed node’s I/O to complete increases.

See “About HPVM and Cluster Re-formation Time” (page 39).

- As of A.11.16, Serviceguard uses **Access Control Policies** to control access to, and administrative control over, the cluster.

See “Access changes as of A.11.16” (page 42).

What’s Not in this Release

- RS232 is no longer supported for the cluster heartbeat.
- Token Ring and FDDI are no longer supported for the cluster heartbeat and data networks.
- Version 3.5 of Veritas Volume Manager (VxVM) from Symantec is no longer supported on HP-UX 11i v3, but VxVM 4.1 and 5.0 are supported.
- The `cmviewcl` command no longer supports the `-r 11.09` option.
`-r 11.12` and `-r 11.16` are still supported. See the `cmviewcl (1m)` manpage for details

See “Announcements” (page 13) and “Rolling Upgrade Exceptions” (page 57) for more information.

Support for Veritas 5.0 on HP-UX 11i v2 and 11i v3

Serviceguard A.11.18 supports Veritas VxVM, CVM and CFS 5.0 from Symantec, on both HP-UX 11i v2 and 11i v3, with the following exceptions:

- Serviceguard supports a maximum of eight nodes for CVM and CFS.
- Serviceguard does not support CFS nested mounts.

For more information, see the *HP Serviceguard Storage Management Suite A.02.00 Release Notes* at docs.hp.com -> High Availability -> HP Serviceguard Storage Management Suite.



NOTE: VxVM, CVM and CFS 4.1 are also supported with Serviceguard A.11.18. See the *HP Serviceguard Storage Management Suite A.01.01 Release Notes*.



IMPORTANT: For information about required patches, see the *HP Serviceguard Storage Management Suite Release Notes* for your version of the Storage Management Suite.

Serviceguard Extension for Faster Failover is supported with CVM and CFS 4.1 and 5.0.

Serviceguard NFS Toolkit supports CFS 5.0. See the white paper *Serviceguard NFS Toolkit Support for CFS* available from docs.hp.com -> High Availability under Highly Available NFS -> White Papers.

About Modular Packages

Serviceguard A.11.18 allows you to create **modular** packages, using building blocks that comprise only the functions that a particular package needs.



IMPORTANT: Serviceguard does not force you to use this new methodology. Existing packages, known as **legacy** packages, continue to run without modification, and can be maintained by means of the older methodology; see “Configuring a Legacy Package” in Chapter 7 of *Managing Serviceguard*. You can also still create a new legacy package. If you are using a Serviceguard Toolkit such as Serviceguard NFS Toolkit, consult the documentation for that product.

The new modular method of creating packages provides the following benefits:

- All package parameters are now configured in one place, the package configuration file. You no longer need to create and distribute a separate package control file.
- External scripts offer an improved means of application integration. These replace the Customer Defined Functions in the legacy package control script. You *do* need to distribute these scripts to all the nodes that can run the package.
- The modular approach allows you to build a package from building blocks containing only the functions needed by this package.

The subsections that follow provide more information. See also:

- *Managing Serviceguard*: “Package Configuration Planning” in Chapter 4; and Chapter 6 “Configuring Packages and Their Services”
- The comments in the editable package configuration file output by the `cmmakepkg` command, and the `cmmakepkg (1m)` manpage

Package Modules

The tables that follows show the new package modules and the configuration parameters each module includes.

You can use `cmmakepkg -l` (letter l) to see a list of all available modules, including non-Serviceguard modules such as those supplied in the HP Toolkits.

Base Package Modules

At least one base module (or `default` or `all`, which include the base module) must be specified on the `cmmakepkg` command line. Parameters in **bold** are new for A.11.18 and are discussed under “New Parameters” (page 26).

Table 1-1 Base Modules

Module Name	Parameters	Comments
failover	<i>package_name</i> <i>module_name</i> <i>module_version</i> <i>package_type</i> <i>node_name</i> <i>auto_run</i> <i>node_fail_fast_enabled</i> <i>run_script_timeout</i> <i>halt_script_timeout</i> <i>successor_halt_timeout</i> <i>script_log_file</i> <i>operation_sequence</i> <i>log_level</i> <i>failover_policy</i> <i>failback_policy</i> <i>priority</i>	Base module. Use as primary building block for failover packages.
multi_node	<i>package_name</i> <i>module_name</i> <i>module_version</i> <i>package_type</i> <i>node_name</i> <i>auto_run</i> <i>node_fail_fast_enabled</i> <i>run_script_timeout</i> <i>halt_script_timeout</i> <i>successor_halt_timeout</i> <i>script_log_file</i> <i>operation_sequence</i> <i>log_level</i> <i>priority</i>	Base module. Use as primary building block for multi-node packages.
system_multi_node	<i>package_name</i> <i>module_name</i> <i>module_version</i> <i>package_type</i> <i>node_name</i> <i>auto_run</i> <i>node_fail_fast_enabled</i> <i>run_script_timeout</i> <i>halt_script_timeout</i> <i>successor_halt_timeout</i> <i>script_log_file</i> <i>operation_sequence</i> <i>log_level</i> <i>priority</i>	Base module. Primary building block for system multi-node packages. System multi-node packages are supported only for applications supplied by HP.

Optional Package Modules

Add optional modules to a base module if you need to configure the functions in question.

Table 1-2 Optional Modules

Module Name	Parameters	Comments
dependency	<i>dependency_name</i> <i>dependency_condition</i> <i>dependency_location</i>	Add to a base module to create a package that depends on one or more other packages.
monitor_subnet	<i>local_lan_failover_allowed</i> <i>monitored_subnet</i> <i>monitored_subnet_access</i> <i>cluster_interconnect_subnet</i>	Add to a base module to configure subnet monitoring for the package.

Table 1-2 Optional Modules (continued)

Module Name	Parameters	Comments
package_ip	<i>ip_subnet</i> <i>ip_subnet_node</i> <i>ip_address</i>	Add to failover module to assign relocatable IP addresses to a failover package.
service	<i>service_name</i> <i>service_cmd</i> <i>service_restart</i> <i>service_fail_fast_enabled</i> <i>service_halt_timeout</i>	Add to a base module to create a package that runs an application or service.
resource	<i>resource_name</i> <i>resource_polling_interval</i> <i>resource_start</i> <i>resource_up_value</i>	Add to a base module to create a package that monitors critical resources via the Event Monitoring Service (EMS).
volume_group	<i>concurrent_vgchange_operations</i> <i>enable_threaded_vgchange</i> <i>vgchange_cmd</i> <i>cvm_activation_cmd</i> <i>vxvol_cmd</i> <i>vg</i> <i>cvm_dg</i> <i>vxvm_dg</i> <i>vxvm_dg_retry</i> <i>deactivation_retry_count</i> <i>kill_processes_accessing_raw_devices</i>	Add to a base module if the package needs to mount file systems on LVM or VxVM volumes, or uses CVM volumes but not CFS.
filesystem	<i>concurrent_fsck_operations</i> <i>concurrent_mount_and_umount_operations</i> <i>fs_mount_retry_count</i> <i>fs_umount_retry_count</i> <i>fs_name</i> <i>fs_directory</i> <i>fs_type</i> <i>fs_mount_opt</i> <i>fs_umount_opt</i> <i>fs_fsck_opt</i>	Add to a base module to configure filesystem options for the package.
pev	<i>pev_</i>	Add to a base module to configure environment variables to be passed to an external script.
external_pre	<i>external_pre_script</i>	Add to a base module to specify additional programs to be run before volume groups and disk groups are activated while the package is starting and after they are deactivated while the package is halting.
external	<i>external_script</i>	Add to a base module to specify additional programs to be run during package start and stopped at package halt time.
acp	<i>user_name</i> <i>user_host</i> <i>user_role</i>	Add to a base module to configure Access Control Policies for the package.

Table 1-2 Optional Modules *(continued)*

Module Name	Parameters	Comments
all	all parameters	Use if you are creating a complex package that requires most or all of the optional parameters; or if you want to see the specifications and comments for all available parameters.
default	(all parameters)	A symbolic link to the all module; used if a base module is not specified on the cmmakepkg command line.



NOTE: The default form for parameter names in the modular package configuration file is lower case; for legacy packages the default is upper case. There are no compatibility issues; Serviceguard is case-insensitive as far as the parameter names are concerned. These release notes use lower case, unless the parameter in question is used only in legacy packages, or the context refers exclusively to such a package.

New Parameters

The following parameters are new as of A.11.18. They can all be used for modular packages, and some (explicitly noted below) can also be used for legacy packages. With the exception of “*vxvm_dg_retry*”, none have direct equivalents in the configuration file or package control file for previous Serviceguard versions.



NOTE: For more information, see the comments in the editable configuration file output by the cmmakepkg command, and the cmmakepkg manpage. See also *Managing Serviceguard: “Package Configuration Planning”* in Chapter 4, and Chapter 6 “Configuring Packages and Their Services”

module_name

The module name (for example, *failover*, *service*, etc.) Do not change it. Used in the form of a relative path (for example *sg/failover*) as a parameter to cmmakepkg to specify modules to be used in configuring the package. (The files reside in the `$SGCONF/modules` directory; see “Understanding Where Files Are Located” in Chapter 5 of *Managing Serviceguard* for an explanation of Serviceguard directories.)

module_version

The module version. Do not change it.

successor_halt_timeout

Specifies how long, in seconds, Serviceguard will wait for packages that depend on this package to halt, before halting this package. Can be 0 through 4294, or `no_timeout`. The default is `no_timeout`. Can be used in both modular and legacy packages.

- `no_timeout` means that Serviceguard will wait indefinitely for the dependent packages to halt.
- 0 means Serviceguard will not wait for the dependent packages to halt before halting this package.

See “About Package Dependencies” in Chapter 4 of *Managing Serviceguard* for more information.

operation_sequence

Defines the order in which the scripts defined by the package's component modules will start up. See the package configuration file for details.

This parameter is not configurable; do not change the entries in the configuration file.

log_level

Determines the amount of information printed to `stdout` when the package is validated, and to the *script_log_file* when the package is started and halted. Valid values are 0 through 5 :

- 0 - informative messages
- 1 - informative messages with slightly more detail
- 2 - messages showing logic flow
- 3 - messages showing detailed data structure information
- 4 - detailed debugging information
- 5 - function call flow

priority

Assigns a priority to a failover package whose *failover_policy* is `configured_node`. Valid values are 1 through 3000 or `no_priority`. The default is `no_priority`. Can be used in both modular and legacy packages.

priority can be used to satisfy dependencies when a package starts, or needs to fail over or fail back: a package with a higher priority than the packages it depends on can force those packages to start or restart on the node it chooses, so that its dependencies are met.

If you assign a priority, it must be unique in this cluster. A lower number indicates a higher priority, and a numerical priority is higher than `no_priority`. HP recommends assigning values in increments of 20 so as to leave gaps in the sequence; otherwise you may have to shuffle all the existing priorities when assigning priority to a new package.

For more information, see "About Package Dependencies" in Chapter 4 of *Managing Serviceguard*.

monitored_subnet_access

In **cross-subnet** configurations, specifies whether each *monitored_subnet* is accessible on all nodes in the package's node list, or only some. Valid values are `PARTIAL`, meaning that at least one of the nodes has access to the subnet, but not all; and `FULL`, meaning that all nodes have access to the subnet. The default is `FULL`, and it is in effect if *monitored_subnet_access* is not specified.

See also "*ip_subnet_node*" and "About Cross-Subnet Configurations" (page 33).

New for modular packages. For legacy packages, see "Configuring Cross-Subnet Failover" in Chapter 7 of *Managing Serviceguard*.

ip_subnet_node

In a **cross-subnet** configuration, specifies which nodes an *ip_subnet* is configured on. If no *ip_subnet_nodes* are listed under an *ip_subnet*, it is assumed to be configured on all nodes in this package's *node_name* list.

See also "*monitored_subnet_access*" and "About Cross-Subnet Configurations" (page 33).

New for modular packages. For legacy packages, see "Configuring Cross-Subnet Failover" in Chapter 7 of *Managing Serviceguard*.

cluster_interconnect_subnet

Specifies an IPv4 address. Can be configured only for a multi-node package in a Serviceguard Extension for Real Application Cluster (SGeRAC) installation. Can be used in both modular and legacy packages.

(See the latest version of *Using Serviceguard Extension for RAC* at <http://www.docs.hp.com> -> High Availability -> Serviceguard Extension for Real Application Cluster (Serviceguard OPS Edition) for more information.)

enable_threaded_vgchange

Indicates whether multi-threaded activation of volume groups (`vgchange -T`) is enabled. New for modular packages. Available on HP-UX 11i v3 only.

Legal values are zero (disabled) or 1 (enabled). The default is zero.

Set `enable_threaded_vgchange` to 1 to enable `vgchange -T` for all of a package's volume groups. This means that when each volume group is activated, physical volumes (disks or LUNs) are attached to the volume group in parallel, and mirror copies of logical volumes are synchronized in parallel, rather than serially. That can improve a package's startup performance if its volume groups contain a large number of physical volumes.

Note that, in the context of a Serviceguard package, this affects the way physical volumes are activated *within* a volume group; `concurrent_vgchange_operations` controls how many volume groups the package can activate simultaneously.



IMPORTANT: Make sure you read the configuration file comments for both `concurrent_vgchange_operations` and `enable_threaded_vgchange` before configuring these options, as well as the `vgchange (1m)` manpage.

vxvm_dg_retry

Specifies whether to retry the import of a VxVM disk group, using `vxdisk scandisks` to check for any missing disks that might have caused the import to fail.

Legal values are `yes` and `no`. `yes` means `vxdisk scandisks` will be run in the event of an import failure. The default is `no`.

HP recommends setting this parameter to `yes` in Metrocluster installation using EMC SRDF.



IMPORTANT: `vxdisk scandisks` can take a long time in the case of a large IO subsystem.

New for modular packages as of the April 2008 Patch; see (page 13). Equivalent to `VXVM_DG_RETRY` in the legacy package control script.

pev_

Specifies a package environment variable that can be passed to an `external_pre_script`, `external_script`, or both, by means of the `cmgetpkgenv (1m)` command.

The variable name must be in the form `pev_<variable_name>` and contain only alphanumeric characters and underscores. The letters `pev` (upper-case or lower-case) followed by the underscore (`_`) are required.

The variable name and value can each consist of a maximum of `MAXPATHLEN` characters (1024 on HP-UX systems).

You can define more than one variable. For more information, see "About External Scripts" in Chapter 4 of *Managing Serviceguard*, and the comments in the package configuration file.

external_pre_script

The full pathname of an external script to be executed before volume groups and disk groups are activated during package startup, and after they have been deactivated during package shutdown (that is, effectively the first step in package startup and last step in package shutdown).

If more than one `external_pre_script` is specified, the scripts will be executed on package startup in the order they are entered into this file, and in the reverse order during package shutdown.

HP supplies a template for external scripts:
`$SGCONF/examples/external_script.template`.

For more information, see “About External Scripts” in Chapter 4 of *Managing Serviceguard*, and the comments in the package configuration file.

external_script

The full pathname of an external script. This script is often the means of launching and halting the application that constitutes the main function of the package.

The script is executed on package startup after volume groups and file systems are activated and IP addresses are assigned, but before services are started; and during package shutdown after services are halted but before IP addresses are removed and volume groups and file systems deactivated.

If more than one *external_script* is specified, the scripts will be executed on package startup in the order they are entered into this file, and in the reverse order during package shutdown.

HP supplies a template for external scripts:
`$SGCONF/examples/external_script.template`.

For more information, see “About External Scripts” in Chapter 4 of *Managing Serviceguard*, and the comments in the package configuration file.

Legacy and Modular Package-Configuration Parameters

There are some differences between the parameters used by modular packages and the legacy parameters. Legacy parameters are still supported, but, as the following tables show, some appear under different names in modular packages, and in a few cases a legacy parameter has no modular equivalent (these are shown in **bold**).

An asterisk in the table indicates that the formal requirements for the parameter value have changed (in terms of what characters are permitted, etc.). See “Parameters whose Formal Requirements Have Changed” (page 31).

Table 1-3 Legacy Configuration-File Parameters and Modular Equivalents

Legacy Parameter	Modular Parameter
<i>PACKAGE_NAME</i> *	<i>package_name</i> *
<i>PACKAGE_TYPE</i>	<i>package_type</i>
<i>NODE_NAME</i>	<i>node_name</i>
<i>AUTO_RUN</i>	<i>auto_run</i>
<i>NODE_FAIL_FAST_ENABLED</i>	<i>node_fail_fast_enabled</i>
<i>RUN_SCRIPT</i>	
<i>HALT_SCRIPT</i>	
<i>RUN_SCRIPT_TIMEOUT</i>	<i>run_script_timeout</i>
<i>HALT_SCRIPT_TIMEOUT</i>	<i>halt_script_timeout</i>
<i>SCRIPT_LOG_FILE</i>	<i>script_log_file</i>
<i>FAILOVER_POLICY</i>	<i>failover_policy</i>
<i>FAILBACK_POLICY</i>	<i>failback_policy</i>
<i>DEPENDENCY_NAME</i> *	<i>dependency_name</i> *
<i>DEPENDENCY_CONDITION</i>	<i>dependency_condition</i>
<i>DEPENDENCY_LOCATION</i>	<i>dependency_location</i>

Table 1-3 Legacy Configuration-File Parameters and Modular Equivalents *(continued)*

Legacy Parameter	Modular Parameter
<i>LOCAL_LAN_FAILOVER_ALLOWED</i>	<i>local_lan_failover_allowed</i>
<i>SUBNET</i>	<i>monitored_subnet</i>
<i>SERVICE_NAME*</i>	<i>service_name*</i>
<i>SERVICE_FAIL_FAST_ENABLED</i>	<i>service_fail_fast_enabled</i>
<i>SERVICE_HALT_TIMEOUT</i>	<i>service_halt_timeout</i>
<i>RESOURCE_NAME</i>	<i>resource_name</i>
<i>RESOURCE_POLLING_INTERVAL</i>	<i>resource_polling_interval</i>
<i>RESOURCE_START</i>	<i>resource_start</i>
<i>RESOURCE_UP_VALUE</i>	<i>resource_up_value</i>
<i>VXVM_DG_RETRY</i>	<i>vxvm_dg_retry</i>
<i>STORAGE_GROUP</i>	
<i>USER_NAME</i>	<i>user_name</i>
<i>USER_HOST</i>	<i>user_host</i>
<i>USER_ROLE</i>	<i>user_role</i>

Table 1-4 Legacy Control-Script Parameters and Modular Equivalents

Legacy Parameter	Modular Parameter
<i>SUBNET[]</i>	<i>ip_subnet</i>
<i>IP</i>	<i>ip_address</i>
<i>SERVICE_NAME[]</i>	<i>service_name</i>
<i>SERVICE_CMD[]</i>	<i>service_cmd</i>
<i>SERVICE_RESTART[]</i>	<i>service_restart</i>
<i>DEFERRED_RESOURCE_NAME</i>	<i>resource_name</i>
<i>CONCURRENT_FSCK_OPERATIONS</i>	<i>concurrent_fsck_operations</i>
<i>CONCURRENT_MOUNT_AND_UMOUNT_OPERATIONS</i>	<i>concurrent_mount_and_umount_operations</i>
<i>FS_UMOUNT_COUNT</i>	<i>fs_umount_retry_count</i>
<i>FS_MOUNT_RETRY_COUNT</i>	<i>fs_mount_retry_count</i>
<i>LV[]</i>	<i>fs_name</i>
<i>FS[]</i>	<i>fs_directory</i>
<i>FS_TYPE[]</i>	<i>fs_type</i>
<i>FS_MOUNT_OPT[]</i>	<i>fs_mount_opt</i>
<i>FS_UMOUNT_OPT[]</i>	<i>fs_umount_opt</i>
<i>FS_FSCK_OPT[]</i>	<i>fs_fsck_opt</i>
<i>CONCURRENT_VGCHANGE_OPERATIONS</i>	<i>concurrent_vgchange_operations</i>
<i>VGCHANGE</i>	<i>vgchange_cmd</i>
<i>CVM_ACTIVATION_CMD</i>	<i>cvm_activation_cmd</i>
<i>VX_VOL</i>	<i>vxvol_cmd</i>

Table 1-4 Legacy Control-Script Parameters and Modular Equivalents *(continued)*

Legacy Parameter	Modular Parameter
VG[]	<i>vg</i>
CVM_DG[]	<i>cvm_dg</i>
VXVM_DG[]	<i>vxvm_dg</i>
DEACTIVATION_RETRY_COUNT	<i>deactivation_retry_count</i>
KILL_PROCESSES_ACCESSING_RAW_DEVICES	<i>kill_processes_accessing_raw_devices</i>

Parameters whose Formal Requirements Have Changed

As of Serviceguard A.11.18, the rules for the values of the following parameters have changed:

- *package_name*
- *dependency_name*
- *service_name*

The value of each of these parameters must now be: any name, up to a maximum of 39 characters, that:

- starts and ends with an alphanumeric character
- otherwise contains only alphanumeric characters or dot (.), dash (-), or underscore (_)
- is unique among package names in this cluster



IMPORTANT: These rules apply to both legacy and modular packages. Restrictions on the values of these parameters in previous Serviceguard releases were less stringent. Packages that use values for these parameters that do not conform to the new rules will continue to run, but if you reconfigure them, you will need to change the values; `cmcheckconf` and `cmapplyconf` will enforce the new rules.

Migrating Legacy to Modular Packages

A new Serviceguard command, `cmmigratepkg`, provided with the patches listed under “New Features for April 2008 Patch” (page 17) and (page 18), automates the process of migrating legacy packages to modular packages as far as possible. Many, but not all, packages can be migrated in this way; for details, see the white paper *Migrating Packages from Legacy to Modular Style* at <http://docs.hp.com ->High Availability -> Serviceguard -> White papers>.



NOTE: The `cmmigratepkg` command requires Perl version 5.8.3 or higher on the system on which you run the command. It should already be on the system as part of the HP-UX base product.

About Package Dependencies

Starting in Serviceguard A.11.17, a package can have **dependencies** on other packages, meaning the package will not start on a node unless the packages it depends on are running on that node.

In Serviceguard A.11.17, package dependencies are supported only for use with certain applications specified by HP, such as the multi-node and system multi-node packages that HP supplies for use with Veritas Cluster File System (CFS) on systems that support it.

As of Serviceguard A.11.18, package dependency is no longer restricted; you can make a package dependent on any other package or packages running on the same cluster node, subject to the restrictions spelled out in Chapter 6 of *Managing Serviceguard*.

Make a package dependent on another package if the first package cannot (or should not) function without the services provided by the second. For example, `pkg1` might run a real-time web interface to a database managed by `pkg2`. In this case it might make sense to make `pkg1` dependent on `pkg2`.

For more information, see “About Package Dependencies” in Chapter 4 of *Managing Serviceguard*.

About Multi-Node Packages

As of Serviceguard A.11.18, multi-node packages are no longer restricted to HP-supplied applications; you can create a multi-node package for any purpose.

Multi-node packages run simultaneously on more than one node in the cluster. Failures of package components such as applications, services, EMS resources, or subnets, will cause the package to be halted only on the node on which the failure occurred.

Relocatable IP addresses cannot be assigned to multi-node packages.



IMPORTANT: If the package uses volume groups, they must be activated in shared mode: `vgchange -a s`, which is available only if the SGeRAC add-on product is installed.

For more information, see Chapter 6 of *Managing Serviceguard*.

Differences between Failover and Multi-Node packages

Note the following important differences in behavior between multi-node and failover packages:

- If a multi-node package has `auto_run` disabled (set to `no` in the package configuration file) it will not start when the cluster is started. You can use `cmmmodpkg` to enable package switching and start the package for the first time. But if you then halt the multi-node package via `cmhaltpkg`, it can be re-started only by means of `cmrunpkg`, not `cmmmodpkg`.
- If a multi-node package is halted via `cmhaltpkg`, package switching is not disabled. This means that the halted package will start to run on a rebooted node, if it is configured to run on that node and its dependencies are met.
- When a multi-node package is started the first time (either at cluster startup, or subsequently if `auto_run` is set to `no`, and package switching is then enabled) any dependent package will start on its primary node. But if a multi-node package is halted along with its dependent packages, and the multi-node package is then restarted, dependent packages which have had package switching re-enabled will start on the first eligible node on which an instance of the multi-node package comes up; this may not be the dependent packages' primary node.

To ensure that dependent failover packages restart on their primary node if the multi-node packages they depend on need to be restarted, make sure the dependent packages' package switching is not re-enabled before the multi-node packages are restarted. You can then either restart the dependent failover packages with `cmrunpkg`, specifying the node you want them to start on, or enable package switching for these packages after the multi-node package startup is complete.

Other Package Changes

The patches listed in the sections headed “New Features for April 2008 Patch” (page 17) and (page 18) provide the following new capabilities:

- Serviceguard now supplies information explaining why a package has shut down. Serviceguard will set the new environment variable `SG_HALT_REASON` in the package control script to one of the following values when the package halts:

- `failure` - set if the package halts because of the failure of a subnet, resource, or service it depends on
- `user_halt` - set if the package is halted by a `cmhaltpkg` or `cmhaltnode` command, or by corresponding actions in Serviceguard Manager
- `automatic_halt` - set if the package is failed over automatically because of the failure of a package it depends on, or is failed back to its primary node automatically (`failback_policy = automatic`)

You can add custom code to the package to interrogate this variable, determine why the package halted, and take appropriate action. For legacy packages, put the code in the `customer_defined_halt_cmds()` function in the CUSTOMER DEFINED FUNCTIONS area of the package control script; for modular packages, put the code in the package's external script (see "*external_script*" (page 29)).

For example, if a database package is being halted by an administrator (`SG_HALT_REASON` set to `user_halt`) you would probably want the custom code to perform an orderly shutdown of the database; on the other hand, a forced shutdown might be needed if `SG_HALT_REASON` is set to `failure`, indicating that the package is halting abnormally (for example because of the failure of a service it depends on).

- `cmviewcl -v -f line` has a new field, `last_halt_failed`, that shows whether the last invocation of the halt script of a package on a node succeeded or failed. The value is `no` if the halt script ran successfully, or was not run since the node joined the cluster, or was not run since the package was configured to run on the node; otherwise it is `yes`.
- A new parameter in the package configuration file, `vxvm_dg_retry`, allows you specify that a failed VxVM import should be retried; see "*vxvm_dg_retry*" (page 28).

About Cross-Subnet Configurations

As of Serviceguard A.11.18, with the patches listed on (page 17) and (page 18), Serviceguard allows you to configure multiple subnets, joined by a router, both for the cluster heartbeat and for data, with some nodes using one subnet and some another.

A cross-subnet configuration allows:

- Automatic package failover from a node on one subnet to a node on another
- A cluster heartbeat that spans subnets.

Configuration Tasks

Cluster and package configuration tasks are affected as follows:

- You must use the `-w full` option to `cmquerycl` to discover actual or potential nodes and subnets across routers.
- For modular packages, you must configure two new parameters in the package configuration file to allow packages to fail over across subnets:
 - `ip_subnet_node` - to indicate which nodes the subnet is configured on
 - `monitored_subnet_access` - to indicate whether the subnet is configured on all nodes (FULL) or only some (PARTIAL)

(For legacy packages, see "Configuring Cross-Subnet Failover" in Chapter 7 of the 15th edition of *Managing Serviceguard*.)

- You should not use the wildcard (*) for `node_name` in the package configuration file, as this could allow the package to fail over across subnets when a node on the same subnet is eligible. Instead, list the nodes in order of preference.

Restrictions

The following restrictions apply:

- All nodes in the cluster must belong to the same network domain (that is, the **domain** portion of the **fully-qualified domain name** must be the same).
- The nodes must be fully connected at the IP level.
- A minimum of two heartbeat paths must be configured for each cluster node.
- There must be less than 200 milliseconds of latency in the heartbeat network.
- Each heartbeat subnet on each node must be physically routed separately to the heartbeat subnet on another node; that is, each heartbeat path must be physically separate:
 - The heartbeats must be statically routed; static route entries must be configured on each node to route the heartbeats through different paths.
 - Failure of a single router must not affect both heartbeats at the same time.
- Because Veritas Cluster File System from Symantec (**CFS**) requires link-level traffic communication (LLT) among the nodes, Serviceguard cannot be configured in cross-subnet configurations with CFS alone.

But CFS *is* supported in specific cross-subnet configurations with Serviceguard and HP add-on products such as Serviceguard Extension for Oracle RAC (**SGeRAC**); see the documentation listed below.

- Each package subnet must be configured with a standby interface on the local bridged net. The standby interface can be shared between subnets.
- Deploying applications in this environment requires careful consideration; see “Implications for Application Deployment” in Chapter 4 of the 15th edition of *Managing Serviceguard*.
- `cmrunnode` will fail if the “hostname LAN” is down on the node in question. (“Hostname LAN” refers to the public LAN on which the IP address that the node’s hostname resolves to is configured).
- If a *monitored_subnet* is configured for `PARTIAL monitored_subnet_access` in a package’s configuration file, it must be configured on at least one of the nodes on the *node_name* list for that package. Conversely, if all of the subnets that are being monitored for this package are configured for `PARTIAL` access, each node on the *node_name* list must have at least one of these subnets configured.
 - As in other configurations, a package will not start on a node unless the subnets configured on that node, and specified in the package configuration file as monitored subnets, are up.
- You cannot use Serviceguard Manager to configure a cross-subnet cluster or package.

For More Information

For more information on the details of configuring the cluster and packages in a cross-subnet context, see the following sections in the 15th edition of *Managing Serviceguard*, at <http://docs.hp.com> -> High Availability -> Serviceguard: “Obtaining Cross-Subnet Information” in Chapter 5, “About Cross-Subnet Failover” in Chapter 4, and (for legacy packages only) “Configuring Cross-Subnet Failover” in Chapter 7.



IMPORTANT: Although cross-subnet topology can be implemented on a single site, it is most commonly used by extended-distance clusters, and specifically **site-aware disaster-tolerant clusters**, which require HP add-on software.

Design and configuration of such clusters are covered in the disaster-tolerant documentation delivered with Serviceguard. For more information, see the following documents at <http://www.docs.hp.com> -> High Availability:

- *Understanding and Designing Serviceguard Disaster Tolerant Architectures*
 - *Designing Disaster Tolerant HA Clusters Using Metrocluster and Continentalclusters*
 - *Using Serviceguard Extension for RAC*
 - The white paper *Configuration and Administration of Oracle 10g R2 RAC Database in HP Metrocluster*
 - The white paper *Technical Considerations for Creating a Serviceguard Cluster that Spans Multiple IP Subnets*
-

New Online Networking-Configuration Capabilities

You can now make significant changes to the cluster's networking configuration while the cluster is running.

What You Can Do

Online operations you can perform include:

- Add a network interface with its *HEARTBEAT_IP* or *STATIONARY_IP*.
- Add a standby interface.
- Delete a network interface with its *HEARTBEAT_IP* or *STATIONARY_IP*.
- Delete a standby interface.
- Change the designation of an existing interface from *HEARTBEAT_IP* to *STATIONARY_IP*, or *vice versa*.
- Change the *NETWORK_POLLING_INTERVAL* and *NETWORK_FAILURE_DETECTION* parameters.
- A combination of any of these in one transaction (`cmapplyconf`), given the restrictions below.

What You Must Keep in Mind

The following restrictions apply:

- You must not change the configuration of all heartbeats at one time, or change or delete the only configured heartbeat.
At least one working heartbeat, preferably with a standby, must remain unchanged.
- In a CVM configuration, you can add and delete only data LANs and IP addresses.
You cannot change the heartbeat configuration while a cluster that uses CVM is running.
- You cannot add interfaces or modify their characteristics unless those interfaces, and all other interfaces in the cluster configuration, are healthy.
There must be no bad NICs or non-functional or locally switched subnets in the configuration, unless you are deleting those components in the same operation.
- You cannot change the designation of an existing interface from *HEARTBEAT_IP* to *STATIONARY_IP*, or *vice versa*, without also making the same change to all peer network interfaces on the same subnet on all other nodes in the cluster.

- You cannot change the designation of an interface from *STATIONARY_IP* to *HEARTBEAT_IP* unless the subnet is common to all nodes.

Remember that the *HEARTBEAT_IP* must be on the same subnet on all nodes, and must be an IPv4 address.

- You cannot delete a primary interface without also deleting any standby interfaces, unless the standby is being used by another primary interface that is not being deleted.
- You cannot delete a subnet or IP address from a node while a package that uses it (as a *monitored_subnet*, *ip_subnet*, or *ip_address*) is configured to run on that node.
- You cannot change the IP configuration of an interface used by the cluster in a single transaction (`cmapplyconf`).

You must first delete the NIC from the cluster configuration, then reconfigure the NIC (using `ifconfig (1m)`, for example), then add the NIC back into the cluster.

Examples of when you must do this include:

- moving a NIC from one subnet to another
- adding an IP address to a NIC
- removing an IP address from a NIC

For more information and examples, see “Changing the Cluster Networking Configuration while the Cluster Is Running” in Chapter 7 of *Managing Serviceguard*.

New -c Option in `cmquerycl`

You can use `cmquerycl -c <cluster_name> -C <filename>` to produce a text file (*filename*) showing the current configuration of the cluster, with additional, commented-out, entries for network subnets that are not currently part of the cluster configuration, but can be configured into it.

For example, suppose that a subnet `15.13.170.0` is shared by nodes `ftsys9` and `ftsys10` in a two-node cluster `cluster1`, and you want to add it to the cluster configuration as a heartbeat subnet. You can run `cmquerycl -c` to get a cluster configuration template file that includes commented out entries for interfaces on that subnet; for example:

```
cmquerycl -c cluster1 -C clconfig.ascii
```

The networking portion of the resulting `clconfig.ascii` file looks something like this:

```
NODE_NAME          ftsys9
NETWORK_INTERFACE  lan1
HEARTBEAT_IP       192.3.17.18
#NETWORK_INTERFACE lan0
#STATIONARY_IP     15.13.170.18
NETWORK_INTERFACE  lan3
# Possible standby Network Interfaces for lan1, lan0: lan2.
```

```
NODE_NAME          ftsys10
NETWORK_INTERFACE  lan1
HEARTBEAT_IP       192.3.17.19
#NETWORK_INTERFACE lan0
# STATIONARY_IP     15.13.170.19
NETWORK_INTERFACE  lan3
# Possible standby Network Interfaces for lan0, lan1: lan2
```

You can now uncomment the entries and apply the new configuration; see “Changing the Cluster Networking Configuration while the Cluster Is Running” in Chapter 7 of *Managing Serviceguard* for more information.

New Online Timing-Configuration Capabilities

You can now change the following cluster-configuration parameters while the cluster is running:

- `NETWORK_POLLING_INTERVAL`
- `HEARTBEAT_INTERVAL`
- `NODE_TIMEOUT`
- `AUTO_START_TIMEOUT`



NOTE: If you are using CVM or CFS, you cannot change `HEARTBEAT_INTERVAL`, `NODE_TIMEOUT`, or `AUTO_START_TIMEOUT` while the cluster is running. This is because they affect the aggregate failover time, which is only reported to CVM on cluster startup.

See “Reconfiguring a Cluster” in Chapter 7 of *Managing Serviceguard* for more information.

About Lock LUNs

The **cluster lock** is a tie-breaker that ensures that a cluster does not re-form into two equal halves; see “Cluster Quorum to Prevent Split Brain Syndrome” in Chapter 3 of *Managing Serviceguard* for details.

LUN stands for Logical Unit Number. The term can refer to a single physical disk, but these days is more often used in a **SAN** (Storage Area Network) or **NAS** (Network-Attached Storage) context to denote a virtual entity derived from one or more physical disks.

As of Serviceguard A.11.18 you can configure a LUN as the cluster lock.

Keep the following points in mind when choosing a device for a lock LUN:

- All the cluster nodes must be physically connected to the lock LUN.
- Connections between the nodes and the lock LUN must be Fibre Channel only; parallel SCSI is not supported.
- All existing data on the LUN will be destroyed when you configure it as a lock LUN.
- A lock LUN cannot also be used in an LVM physical volume or VxVM or CVM disk group.
- A lock LUN cannot be shared by more than one cluster.
- A lock LUN cannot be used in a dual-lock configuration.
- You do not need to back up the lock LUN data, and in fact there is no way to do so.
- A lock LUN needs only a small amount of storage, about 100 KB.

For more information, see “Setting Up a Lock LUN” in Chapter 5 of *Managing Serviceguard* and the `cmdisklock (1m)` manpage.

About Device Special Files (DSFs)

HP-UX releases up to and including 11i v2 use a naming convention for device files that encodes their hardware path. For example, a device file named `/dev/dsk/c3t15d0` would indicate SCSI controller instance 3, SCSI target 15, and SCSI LUN 0. HP-UX 11i v3 introduces a new nomenclature for device files, known as **agile addressing** (sometimes also called **persistent LUN binding**).

Under the agile addressing convention, the hardware path name is no longer encoded in a storage device’s name; instead, each device file name reflects a unique instance number, for example `/dev/[r]disk/disk3`, that does not need to change when the hardware path does.

Agile addressing is the default on new 11i v3 installations, but the I/O subsystem still recognizes pre-11i v3 device files, which as of 11i v3 are referred to as **legacy device files**. Device files using the new nomenclature are called **persistent device files**,

When you upgrade to HP-UX 11i v3, a set of new, persistent device files is created, but the existing, legacy device files are left intact and by default will continue to be used by HP-UX and Serviceguard.

This means that you are not required to migrate to agile addressing when you upgrade to 11i v3, though you should seriously consider its advantages (see the white paper *The Next Generation Mass Storage Stack under Network and Systems Management -> Storage Area Management* at [docs.hp.com](http://www.docs.hp.com)). Migration involves modifying system and application configuration files and scripts to use persistent device files and in some cases new commands and options; the process is described in the white papers *Migrating from HP-UX 11i v2 to HP-UX 11i v3* and *LVM Migration from Legacy to Agile Naming Model HP-UX 11i v3* at <http://www.docs.hp.com>.

If you cold-install HP-UX 11i v3, sets of both legacy and persistent device files are automatically created. In this case, by default the installation process will configure system devices such as the boot, root, swap, and dump devices to use persistent device files. This means that system configuration files such as `/etc/fstab` and `/etc/lvmtab` will contain references to persistent device files, but Serviceguard's functioning will not be affected by this.



CAUTION: You cannot migrate to the agile addressing scheme during a rolling upgrade if you are using cluster lock disks as a tie-breaker, because that involves changing the cluster configuration. But under certain conditions, you *can* migrate the cluster lock device file names to the new scheme without bringing the cluster down. For the requirements and a procedure, see "Updating the Cluster Lock Configuration" in Chapter 7 of *Managing Serviceguard*.



NOTE: It is possible, though not a best practice, to use legacy DSFs on some nodes after migrating to agile addressing on others; this allows you to migrate different nodes at different times, if necessary.

For more information about agile addressing, see following documents at <http://www.docs.hp.com>:

- the *Logical Volume Management* volume of the *HP-UX System Administrator's Guide* (in the 11i v3 -> System Administration collection)
- the *HP-UX 11i v3 Installation and Update Guide* (in the 11i v3 -> Installing and Updating collection)
- the following white papers:
 - *The Next Generation Mass Storage Stack* (under Network and Systems Management -> Storage Area Management)
 - *Migrating from HP-UX 11i v2 to HP-UX 11i v3*
 - *HP-UX 11i v3 Native Multi-Pathing for Mass Storage*
 - *LVM Migration from Legacy to Agile Naming Model HP-UX 11i v3*

See also the HP-UX 11i v3 intro (7) manpage.

About LVM 2.0

Logical Volume Manager (LVM) 2.0 volume groups, which remove some of the limitations imposed by LVM 1.0 volume groups, can be used on systems running HP-UX 11i v3 0803 Fusion with Serviceguard A.11.18 and the Serviceguard patch listed under "New Features for April 2008 Patch" (page 18).



IMPORTANT: LVM2 volume groups are not supported for cluster lock disks.



NOTE: You are not required to move to LVM 2.0 volume groups and everything will work as before if you do nothing.

If you *do* use LVM 2.0 volume groups, you can still manage them with the same commands as before, although you may have to make minor changes to any scripts you use that parse the output of `lvdisplay`, `vgdisplay`, `pvdisplay` or `vgscan`, as the output of these commands has changed slightly. In addition, new options are available for some commands.

For more information, see the white paper *LVM 2.0 Volume Groups in HP-UX 11i v3* at docs.hp.com -> 11i v3 -> LVM Volume Manager. For information about all other aspects of LVM on HP-UX 11i v3, see the *Logical Volume Management* volume (volume 3) of the *HP-UX System Administrator's Guide*, at the same address.

About olrad

You must remove a LAN or VLAN interface from the cluster configuration before removing it from the system. You can do this without bringing down the cluster.

HP-UX 11i v3 provides a new option for the `olrad` command, `olrad -C`, to help you determine whether or not an interface is part of the cluster configuration: run `olrad -C` with the affected I/O slot ID as argument.

If the NIC is part of the cluster configuration, you'll see a warning message telling you to remove it from the configuration before you proceed. See the `olrad (1M)` manpage for more information about `olrad`.

After removing the NIC from the cluster configuration, you can remove it from an HP-UX 11i v3 cluster node without shutting down the system by running `olrad -d`.

See "Removing a LAN or VLAN Interface from a Node" in Chapter 7 of *Managing Serviceguard* for more information.

About HPVM and Cluster Re-formation Time

When a node fails and the cluster re-forms, Serviceguard must wait a certain amount of time to allow I/O from the failed node to be written out to the target storage device. Only after that time has elapsed can Serviceguard allow an adoptive node access to that device; otherwise data corruption could occur. The amount of time Serviceguard waits is calculated by Serviceguard and is not user-configurable.

The above is true whether or not the cluster includes virtual machines (VMs), but using VMs as Serviceguard nodes increases the amount of time Serviceguard needs to wait before it is safe to allow another node access to the same storage. This additional wait can increase cluster re-formation time by as much as 70 seconds.

The additional time Serviceguard needs to wait depends in part on whether or not a **VM guest depot** is installed on the VM node. (See *HP Integrity Virtual Machines Installation, Configuration, and Administration*, at the address given below, for information on installing a guest depot.) Serviceguard uses information it derives from the VM guest depot to set the timeout to the optimal value. If any VM node does not have a VM guest depot, Serviceguard may not be able to obtain the information it needs to set the optimal timeout, and in that case it sets the additional timeout to the maximum value, 70 seconds.



IMPORTANT: This additional timeout extension represents a net addition to the time it takes for the cluster to re-form. For example, if the cluster typically took 40 seconds to re-form before any VM nodes were added, it will take about 80 seconds when one or more VM nodes are members of the cluster, if all those nodes have a VM guest depot. If any VM node without a VM guest depot is a member of the cluster, it will take about 110 seconds. This is true whenever VM nodes are cluster members, whether or not the re-formation is caused by the failure of a VM node.

For more information about HP Integrity Virtual Machines, see *HP Integrity Virtual Machines Installation, Configuration, and Administration* at <http://docs.hp.com> -> HP Virtual Server Environment (VSE) -> HP Integrity Virtual Machines.

About the Volume Monitor

Simply monitoring each physical disk in a Serviceguard cluster does not provide adequate monitoring for volumes managed by Veritas Volume Manager from Symantec (VxVM), because a physical volume failure is not always a critical failure that triggers failover (for example, the failure of a mirrored volume is not considered critical).

For this reason, and because VxVM allows you to resize a volume while it is online, it can be very difficult to determine which physical disks must be monitored to ensure that a logical volume is functioning properly. The HP Serviceguard VxVM Volume Monitor, supplied with the patches listed in the sections headed “New Features for April 2008 Patch” (page 17) and (page 18), provides a means for effective and persistent monitoring of VxVM volumes. The VxVM Volume Monitor supports Veritas Volume Manager versions 3.5, 4.1, and 5.0, as well as Veritas Cluster Volume Manager (CVM) versions 3.5, 4.1, and 5.0.

Monitor Usage

You configure the VxVM Volume Monitor as a service within a package which requires access to a VxVM volume.

When a monitored volume fails or becomes inaccessible, the monitor service will exit, causing the package to fail on the current node. The package’s failover behavior depends on its configured settings, as with any other failover package.

Syntax:

```
cmvx serviced [-h, --help] [-v, --version]
[-O, --log-file <log_file>]
[-D, --log-level <1-7>]
[-t, --poll-interval <seconds>]
<volume_path> [<volume_path>...]
```

A brief description of each parameter follows:

-h or --help

Displays the usage, as listed above, and exits.



NOTE: Do not include the help or version parameters in your service command; this will result in immediate package failure at runtime.

-v or --version

Displays the monitor version and exits.

-O or --log-file

Specifies a file for logging (log messages are printed to the console by default).

`-D` or `--log-level`

Specifies the log level. The level of detail logged is directly proportional to the numerical value of the log level. That is, a log level of 7 will provide the greatest amount of log information. The default log level is 0.

`-t` or `--poll-interval`

Specifies the interval between volume probes.

You can specify a polling interval of as little as 1 (one second), but bear in mind that a short polling interval (less than 10 seconds) may impair system performance if you are monitoring a large number of volumes. HP recommends a polling interval of at least 10 seconds if 50 or more volumes are being monitored by a single service command.

The default polling interval is 60 seconds.

`volume_path`

The full path to at least one VxVM volume device file for monitoring (required). The pathname must identify a block device file.

Examples

```
/usr/sbin/cmvserviced -O /log/monlog.log -D 3 /dev/vx/dsk/cvm_dg0/lvol2
```

This command monitors a single volume, `/dev/vx/dsk/cvm_dg0/lvol2`, at log level 3, with a polling interval of 60 seconds, and prints all log messages to `/log/monlog.log`.

```
/usr/sbin/cmvserviced /dev/vx/dsk/cvm_dg0/lvol1 /dev/vx/dsk/cvm_dg0/lvol2
```

This command monitors two volumes at the default log level of 0, with a polling interval of 60 seconds, and prints all log messages to the console.

```
/usr/sbin/cmvserviced -t 10 /dev/vx/dsk/cvm_dg2/lvol3
```

This command monitors a single volume at log level 0, with a polling interval of 10 seconds, and prints all log messages to the console.

Scope of Monitoring

The VxVM Volume Monitor detects the following failures:

- Failure of the last link to a storage device or set of devices critical to volume operation
- Failure of a storage device or set of devices critical to volume operation
- An unexpected detachment or disablement of a volume

The VxVM Volume Monitor *does not* detect the following failures:

- Failure of a redundant link to a storage device or set of devices where a working link remains
- Failure of a mirrored plex within a volume (assuming at least one plex is functional)
- Corruption of data on a volume which Veritas Volume Manager regards as enabled and active

About `vgchange -T`

With the patch listed under “[New Features for April 2008 Patch](#)” (page 18), Serviceguard supports `vgchange -T`, which allows multi-threaded activation of volume groups on HP-UX 11i v3 systems.

This means that when the volume group is activated, physical volumes (disks or LUNs) are attached to the volume group in parallel, and mirror copies of logical volumes are synchronized in parallel, rather than serially. That can improve a package’s startup performance if its volume groups contain a large number of physical volumes.

To enable `vgchange -T` for all of a package’s volume groups, set `enable_threaded_vgchange` to 1 in the package configuration file (the default is 0, meaning that multi-threaded activation is disabled).

Note that, in the context of a Serviceguard package, this affects the way physical volumes are activated *within* a volume group; another package parameter, *concurrent_vgchange_operations*, controls how many volume groups the package can activate simultaneously.



IMPORTANT: Make sure you read the configuration file comments for both *concurrent_vgchange_operations* and *enable_threaded_vgchange* before configuring these options, as well as the `vgchange (1m)` manpage.

Alternate Quorum Server Subnet

Serviceguard 11.18, with the patches mentioned under “New Features for April 2008 Patch” on (page 17) and (page 18), allows you to configure an alternate subnet for communication between the cluster nodes and the Quorum Server. You can do this from the command line or in Serviceguard Manager.



IMPORTANT: This capability requires Quorum Server Version A.03.00. For details and instructions, including patch information, see the *HP Serviceguard Quorum Server Version A.03.00 Release Notes* at <http://docs.hp.com> -> High Availability -> Quorum Server.

Access changes as of A.11.16

Serviceguard version A.11.16 introduced a new access method. As of A.11.16, Serviceguard uses **Access Control Policies**, also known as **Role-Based Access**, rather than `cmclnodelist` or `.rhosts`, to authenticate users.

For more information about Access Control Policies, see Chapter 5 of the *Managing Serviceguard* manual (docs.hp.com -> High Availability), the Serviceguard Manager help, and the cluster and package configuration files themselves.

Considerations when Upgrading Serviceguard

- `.rhosts`

If you relied on `.rhosts` for access in the previous version of the cluster, you must now configure Access Control Policies for the cluster users. For instructions on how to proceed, see the subsection “Allowing Root Access to an Unconfigured Node” under “Configuring Root-Level Access” in Chapter 5 of the *Managing Serviceguard* manual.

- `cmclnodelist`

When you upgrade from an earlier version, Serviceguard converts `cmclnodelist` entries into new entries written into the cluster configuration file during the upgrade, as follows:

```
USER_NAME <user_name>
```

```
USER_HOST <host_node>
```

```
USER_ROLE Monitor
```

A wild-card + (plus) is converted as follows:

```
USER_NAME ANY_USER
```

```
USER_HOST ANY_SERVICEGUARD_NODE
```

```
USER_ROLE Monitor
```

After you complete the upgrade, use `cmgetconf` to create and save a copy of the new configuration. If you do a `cmapplyconf`, you want to be sure it applies the newly migrated Access Control Policies.

Considerations when Installing Serviceguard

When you install Serviceguard for the first time on a node, the node is not yet part of a cluster, and so there is no Access Control Policy. For instructions on how to proceed, see the subsection “Allowing Root Access to an Unconfigured Node” under “Configuring Root-Level Access” in Chapter 5 of the *Managing Serviceguard* manual.

About the Serviceguard Manager SMH Plug-In

HP Serviceguard Manager B.01.01 is a web-based, HP System Management Homepage (HP SMH) plug-in application, that replaces the functionality of the earlier Serviceguard management tools. HP Serviceguard Manager allows you to monitor, administer and configure a Serviceguard A.11.18 cluster from any system with a web browser.

Unlike previous Serviceguard management tools, Serviceguard Manager does not require additional software installation. Instead, using your browser, you log into an HP Systems Management Homepage (SMH) and access HP Serviceguard Manager tool, as well as other system management tools.

HP Serviceguard Manager Main Page provides you with a summary of the health of the cluster including the status of each node and each package.

DSAU Integration

HP Serviceguard Manager uses Distributed Systems Administration Utilities (DSAU) to display consolidated cluster log (`syslog`) and consolidated package logs.

You can find more information on DSAU in the *Distributed Systems Administration Utilities User's Guide*, available from <http://docs.hp.com> -> Network and Systems Management -> System Administration.

Native Language Support

HP Serviceguard Manager Version B.01.01 is available in English, Japanese, Korean, Simplified Chinese and Traditional Chinese, French, German, Spanish, and Italian.

What You Can Do

Depending on your SMH and Serviceguard security privileges, you can do the following:

- Monitor, create, and modify, and run and halt a cluster.
- Monitor and run and halt nodes.
- Create and modify failover and multi-node packages, including configuring package dependencies. You can also modify `Auto Run` and `Node Switching` settings.
- Monitor, run, halt, and move failover, multi-node, and system multi-node packages.

New Features

HP Serviceguard Manager version B.01.01 supports Serviceguard A.11.18 on HP-UX 11i v2 and HP-UX 11i v3. The following are new capabilities in B.01.01:

- Monitoring enhancements:
 - Identify HPVM guest and host, nPar, and vPar
 - Identify the following Serviceguard extensions:
 - HP Serviceguard Extension for RAC
 - HP Serviceguard Extension for Faster Failover

- HP Metrocluster
- HP Continentalclusters
- Administration and configuration enhancements:
 - Support multi-node packages
 - Supports same node dependencies for failover packages
 - Support for Lock LUN
- Usability enhancements:
 - Mouse-over tooltips for read-only property pages
 - General consolidation of information displayed on the main page and read-only property pages
 - New icons
 - Many read-only property pages include multiple view modes, including the default tabular view and an optional consolidated matrix view.

Differences from A.05.01

The following are differences between B.01.01 and A.05.01:

- Single cluster view only – A.05.01 manages multiple clusters, but B.01.01 manages only the local cluster and depends on HP SIM (Systems Insight Manager) to provide multi-cluster view.




NOTE: For more information about SIM, go to <http://www.hp.com/go/hpsim> and follow the Information Library link for access to white papers and other documents.

- Cluster topology map is replaced by HP Serviceguard Manager main page.
- The SNMP event viewer is replaced by consolidated `syslog` and package logs.
- Drag and drop of a package is replaced by the “Move Package” operation.
- Openview integration is done via HP SIM.
- No file save – Serviceguard Manager A.05.01 had the ability to save a snapshot of the current clusters monitored, and display it later. This feature is mainly used for comparing configuration changes, which you can now do with DSAU.

Limitations

- Using Serviceguard Manager, you can create only **legacy** packages.
But you can monitor and administer both legacy and **modular** packages. See “About Modular Packages” (page 22) for more information about modular and legacy packages.
- You cannot use Serviceguard Manager to configure **cross-subnet** clusters or packages (see “About Cross-Subnet Configurations” (page 33)).
 - The Serviceguard Manager Configuration pull-down menu is disabled in a cross-subnet configuration.
Although you can view the cluster, node, and package property pages, you cannot create or modify a cluster, package, or package control script.
 - You cannot configure DSAU log consolidation.
If you currently have DSAU configured to consolidate the logs locally, you will need to set up a remote log consolidation server. See Chapter 3 of the *Distributed Systems Administration User's Guide* for more information.
If you use messages that are sent to `syslog.log` to trigger warning and alert alarms, you will need to configure the alarm server to point to the log server.

Help Subsystem

Use this section to help you become familiar with Serviceguard Manager. Once Serviceguard Manager is running, use the tooltips by moving your mouse over a field from the read-only property pages for a brief definition for each field. You can also access the online help by clicking the  button located in the upper-right hand corner of the screen to view overview and procedure information. Start with the help topic Understanding the HP Serviceguard Manager Main Page. You should read the help topic About Security, as it explains HP Serviceguard Manager Access Control Policies, as well as root privileges.

Before Using HP Serviceguard Manager: Setting Up

You must have, or have done, the following before you can start using HP Serviceguard Manager:

- At least one cluster member node with Serviceguard A.11.18 and Serviceguard Manager B.01.01 installed.

- The hpuxswTOMCAT product.

hpuxswTOMCAT is installed by default with HP-UX. To check that it is on your system, use a command such as:

```
swlist -l fileset | grep TOMCAT
```

- Version 2.2.5 or higher of SMH (System Management Homepage)
 - A web browser (Internet Explorer 6.0 or higher, Mozilla 1.5 or higher) with access to SMH.
 - Have launched SMH (settings -> Security -> User Groups) to configure user roles for SMH.
 - A user with HP SMH Administrator access has full cluster management capabilities.
 - A user with HP SMH Operator access can monitor the cluster and has restricted cluster management capabilities as defined by the user's Serviceguard role-based access configuration.
 - A user with HP SMH User access does not have any cluster management capabilities.
- See the online help topic About Security for more information.
- Have created the security "bootstrap" file `cmclnodelist`. See "Configuring Root-Level Access" in Chapter 5 of the *Managing Serviceguard* manual for instructions.

Launching Serviceguard Manager

There are three scenarios for accessing Serviceguard Manager based on the number of clusters you want to manage, and which Serviceguard version you are using.

The following table details the options you have to access Serviceguard Manager:

Table 1-5 Accessing Serviceguard Manager

Scenario	Scenario Use Case	Number of Clusters	Serviceguard Version	Serviceguard Manager Version	See for more info
1	Single cluster management	1	A.11.18	B.01.01	(page 46)
			A.11.17.01	B.01.00	
2	Multi-cluster management	1 or more cluster(s)	A.11.16 - A.11.17.01	A.05.01	(page 48)
3	Web-based multi-cluster management	more than 1 cluster	A.11.16 - A.11.18	A.05.01 with HP SIM 5.10	(page 49)

Table 1-6 details the Serviceguard Manager workstation installation requirements.

Table 1-6 Serviceguard Manager Installation Requirements

Scenario	Scenario Use Case	Serviceguard Manager Version	Installation Requirements
1	Single cluster management	B.01.01, B.01.00	For HP-UX - No installation is required.
2	Multi-cluster management	A.05.01	Requires Serviceguard Manager installation on each system that will be used as a management workstation.
3	Web-based multi-cluster management	A.05.01 with HP SIM 5.10	Requires the following software installations on a server: <ul style="list-style-type: none"> • HP SIM 5.10 <i>and</i> • Serviceguard Manager A.05.01

Scenario 1 - Single cluster management

Scenario 1 applies if the following is true:

- Manage a single cluster
- Have installed Serviceguard version A.11.18.

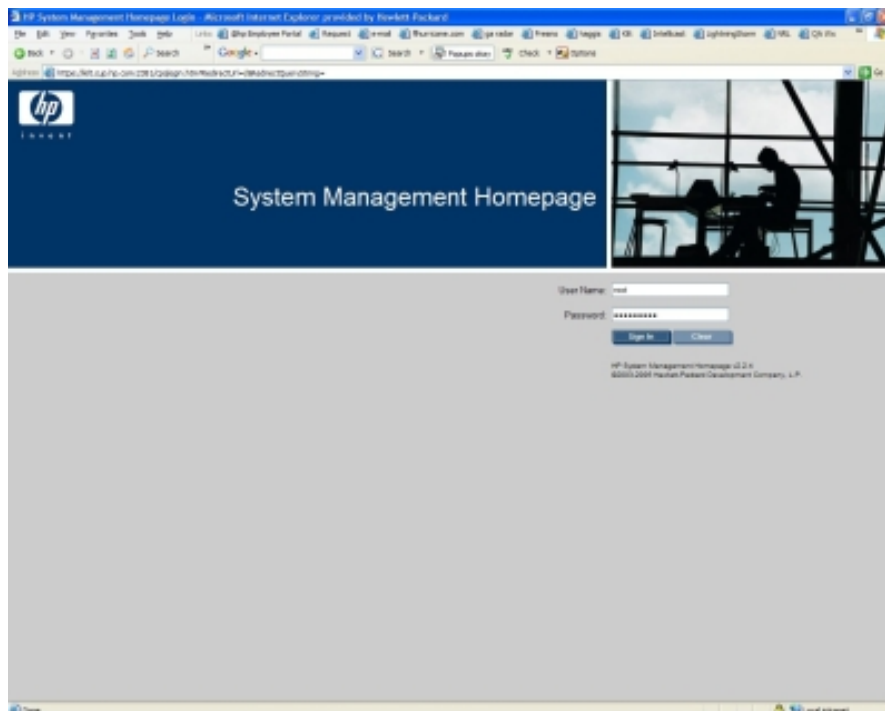
In order to administer Serviceguard version A.11.18 through SMH, a user must login from a browser using the following URL: `http://<full hostname of a cluster node>:2301`.

- A user with HP SMH Administrator access has full cluster management capabilities.
- A user with HP SMH Operator access can monitor the cluster and has restricted cluster management capabilities as defined by the user's Serviceguard role-based access configuration.
- A user with HP SMH User access does not have any cluster management capabilities.

1. Enter the standard URL "`http://<full hostname of server>:2301/`"

For example `http://clusternode1.cup.hp.com:2301/`

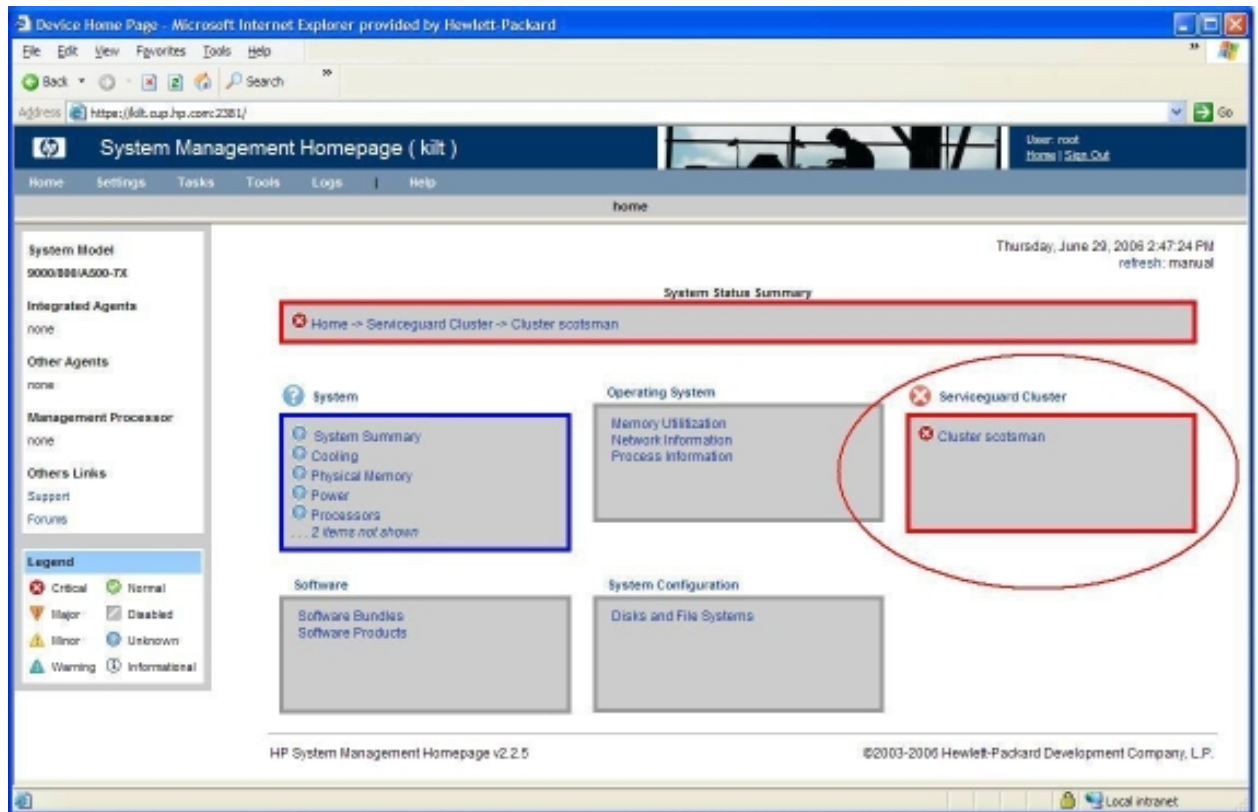
Figure 1-1 System Management Homepage Login



2. When the System Management Homepage login screen appears, enter your login credentials and click Sign In.

The System Management Homepage for the selected server appears. From the Serviceguard Cluster box, click the name of the cluster.

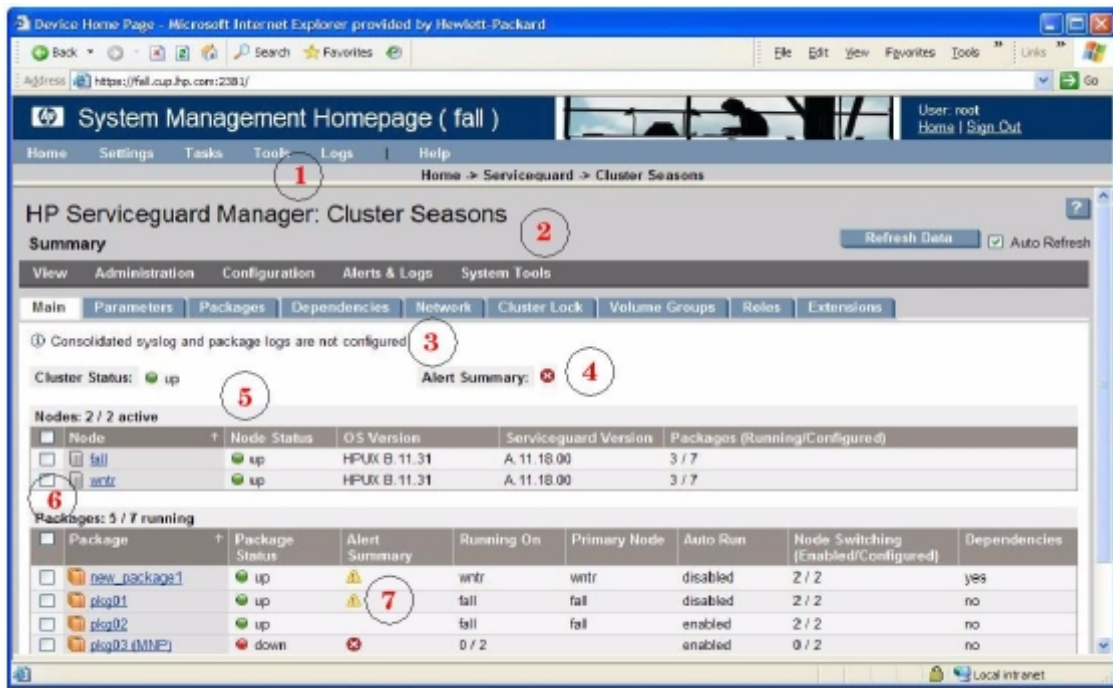
Figure 1-2 Serviceguard Cluster



NOTE: If a cluster is not yet configured, then you will not see the Serviceguard Cluster section on this screen. To create a cluster, from the SMH Tools menu, you must click Serviceguard Manager link in the Serviceguard box first, then click Create Cluster.

The figure below shows a browser session at the HP Serviceguard Manager Main Page.

Figure 1-3 System Management Homepage with Serviceguard Manager



Number	Description
1	System Management Homepage (SMH) menus.
2	HP Serviceguard Manager pull-down menus.
3	HP Serviceguard Manger tabs.
4	Overall cluster status indicator. Click this icon to troubleshoot issues.
5	Easy to understand status icons.
6	Use the check boxes to multi-select nodes or packages from which to take action. <ul style="list-style-type: none"> To select all, select the check box located in the left corner of the table header. To select some or specific nodes or packages, select the check box corresponding to the node or package from which you want to take an action. <p>NOTE: You can only select one package at a time to modify its configuration.</p>
7	Clickable icons allow you to troubleshoot issues.

Scenario 2 -Multi-cluster management

Scenario 2 applies if the following is true:

- One or more clusters with Serviceguard version A.11.16 through A.11.17.01
- Serviceguard Manager version A.05.01 installed

Launching Serviceguard Manager from Windows

Here are 2 ways to launch Serviceguard Manager in Windows:

- By default, a shortcut icon is placed on your desktop (labelled SG Manager). Click it to launch the program. When installed, it has the simple command without options. You will be prompted to enter the options once the program opens.
- From the *Start* menu, select Programs -> Serviceguard Manager. This launches the simple command, with no options.



NOTE: You can also launch Serviceguard Manager from an HP-UX or Linux system.

Scenario 3 - Web-based multi-cluster management

Scenario 3 applies if the following is true:

- One or more clusters with Serviceguard version A.11.15.01 through A.11.18
- Serviceguard Manager version A.05.01 with HP SIM 5.10 installed on a server



NOTE: Serviceguard Manager can be launched by HP Systems Insight Manager version 5.10 if Serviceguard Manager is installed on an HP Systems Insight Manager Central Management Server.

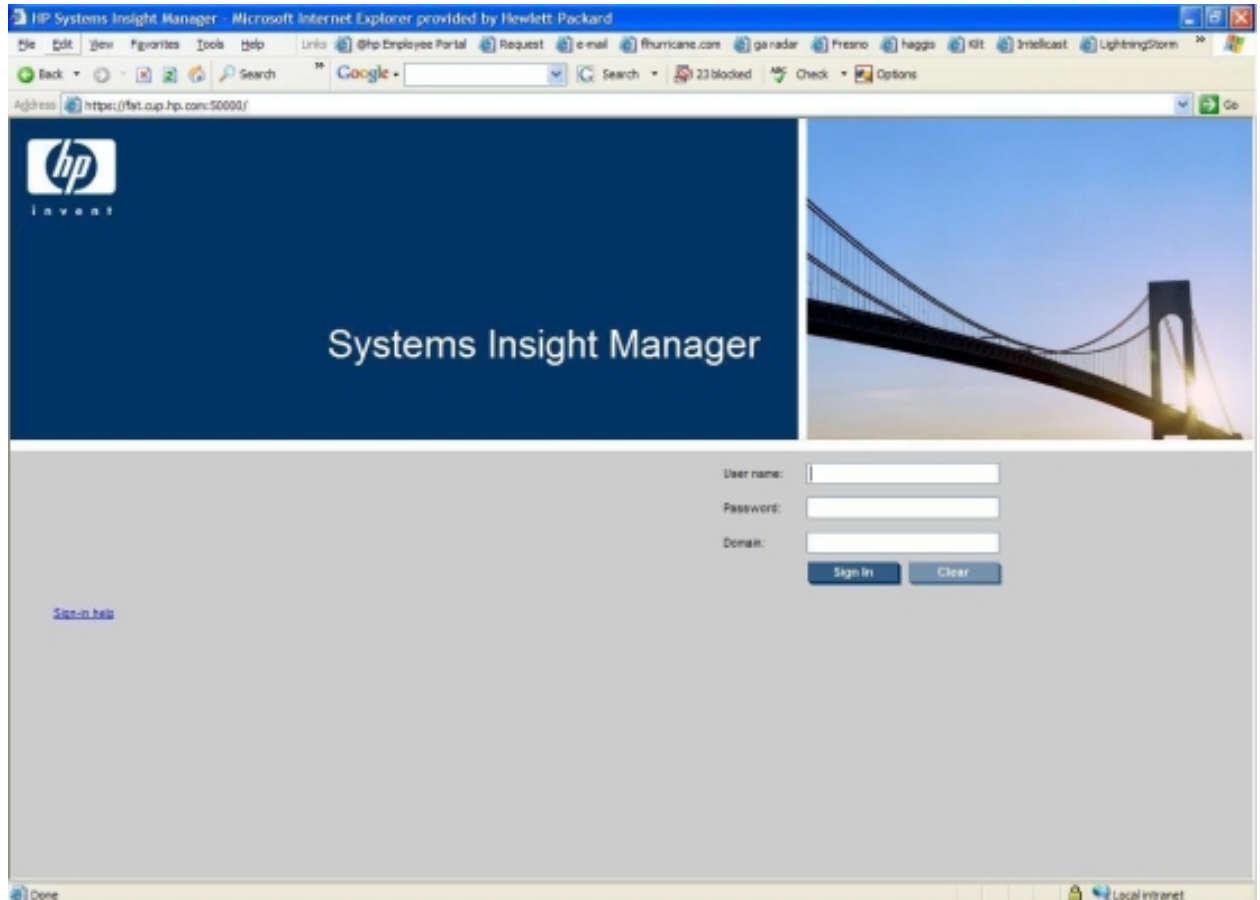
For a Serviceguard A.11.18 cluster, Systems Insight Manager will attempt to launch Serviceguard Manager B.01.01 from one of the nodes in the cluster; for a Serviceguard A.11.17.01 cluster, Systems Insight Manager will attempt to launch Serviceguard Manager B.01.00 from one of the nodes in the cluster.

For a Serviceguard A.11.17 cluster or earlier, Serviceguard Manager A.05.01 will be launched via Java Web Start. You must ensure that the hostname for each Serviceguard node can be resolved by DNS. For more information about this older version of Serviceguard Manager, see the Serviceguard Manager Version A.05.01 Release Notes at <http://docs.hp.com> -> High Availability ->Serviceguard Manager.

1. Enter the standard URL, `https://[full hostname of SIM server]:50000/`

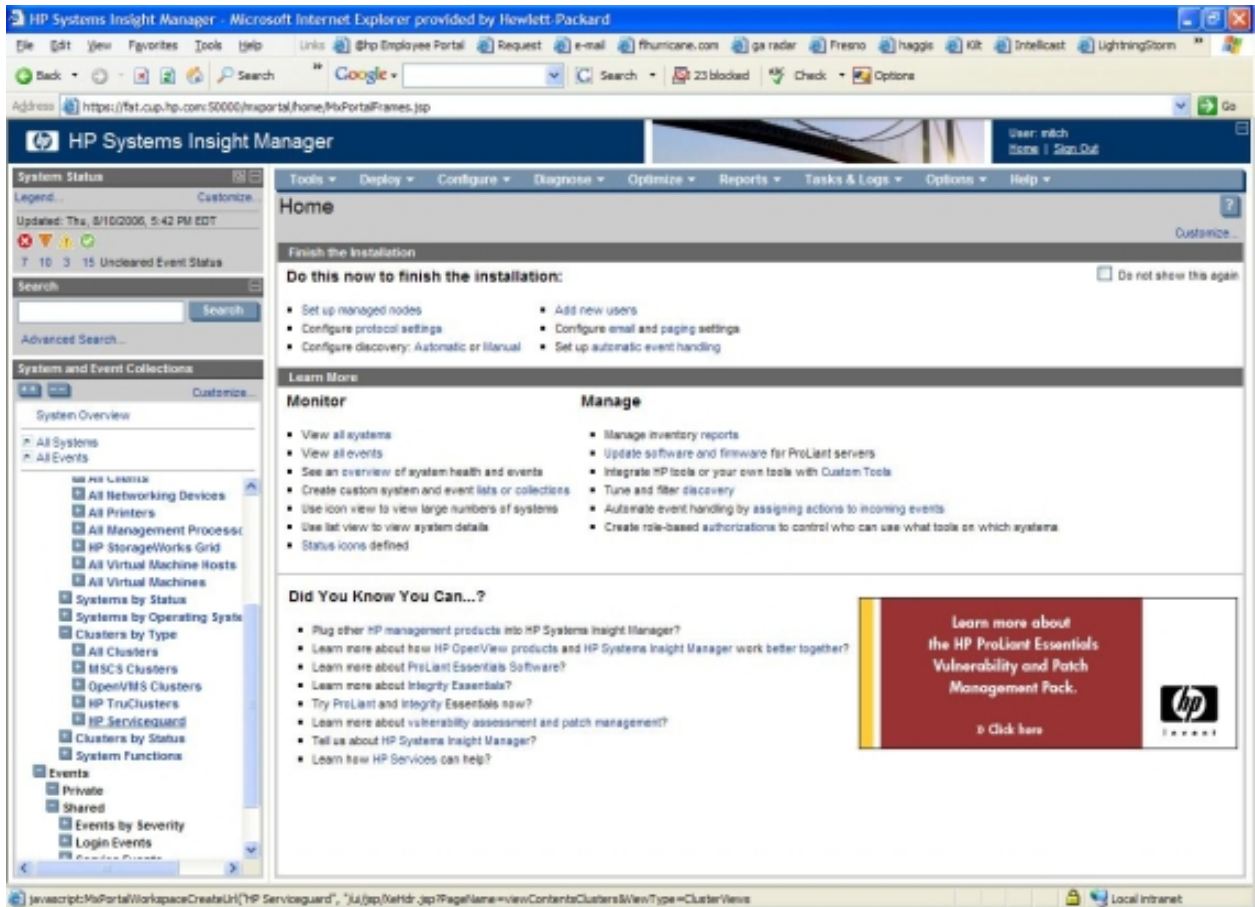
For example `https://SIMserver.cup.hp.com:50000/`

Figure 1-4 Systems Insight Manager Sign-In Page



2. When the Systems Insight Manager login screen appears, enter your login credentials and click **Sign In**.
3. From the left-hand panel, expand Cluster by Type.

Figure 1-5 Expanded Left Panel



4. Expand HP Serviceguard, and click on a Serviceguard cluster. The following screen appears if you click on a Serviceguard cluster running A.11.18:

Figure 1-6 Serviceguard Manager Page

The screenshot displays the HP Serviceguard Manager web interface. The browser window shows the URL <https://guist8.cup.hp.com:50000/mxportal/home/MxPortalFrames.jsp>. The page title is "HP Systems Insight Manager". The user is logged in as "administrator". The main content area is titled "HP Serviceguard Manager" and shows a cluster named "scotsman" with 2/2 active nodes. A table lists 31 running packages with columns for Package, Status, Alert Summary, Running On, Primary Node, Auto Run, Node Switching, and Dependencies.

Package	Package Status	Alert Summary	Running On	Primary Node	Auto Run	Node Switching (Enabled/Configured)	Dependencies
alexander1	up	ⓘ	kilt	kilt	enabled	1 / 1	no
clb	up		haggis	haggis	enabled	2 / 2	no
msch2	up		haggis	haggis	enabled	2 / 2	no
msProg-1	up		haggis	haggis	enabled	2 / 2	no
new_msc	up		kilt	kilt	enabled	2 / 2	no
new_package1	up		haggis	haggis	enabled	2 / 2	no
new_package2	up		haggis	haggis	enabled	2 / 2	no
new_package3	up		haggis	haggis	enabled	2 / 2	no
new_package4	up		haggis	haggis	enabled	2 / 2	no
pkg1	up		haggis	haggis	enabled	2 / 2	yes
pkg2	up		haggis (adoptive)	kilt	enabled	2 / 2	yes
pkg3	up		haggis	kilt	enabled	2 / 2	yes



NOTE: If you click on a cluster running an earlier Serviceguard release, the page will display a link that will launch Serviceguard Manager A.05.01 (if installed) via Java Webstart.

Patches and Fixes

No patches are required for Serviceguard Manager B.01.01. For information about known problems and workarounds, see “Known Problems and Workarounds for the Serviceguard Manager Plug-in” (page 69).

What Documents are Available for This Version

For information about the current version of Serviceguard, and about older versions, see the Serviceguard documents posted at <http://docs.hp.com/> -> High Availability -> Serviceguard. Documents for related products, including the older, management-station version of Serviceguard Manager, are also on the High Availability page.

The following documents relate to Serviceguard A.11.18 and related high availability products. The latest versions can be found on the web at <http://docs.hp.com/>->High Availability.

- *Managing Serviceguard*, 15th Edition. This manual has been revised for the A.11.18 release and the additional new features announced in these Release Notes under “New Features for April 2008 Patch” (page 17) and (page 18).
- *HP Serviceguard Quorum Server Version A.03.00 Release Notes*
- *Serviceguard Extension for RAC Version A.11.18 Release Notes*
- *Using Serviceguard Extension for RAC*

- *Understanding and Designing Serviceguard Disaster Tolerant Architectures*
- *Designing Disaster Tolerant HA Clusters Using Metrocluster and Continentalclusters*
- *Enterprise Cluster Master Toolkit Version Release Notes*
- *Clusters for High Availability: A Primer of HP Solutions*, second edition (HP Press: Prentice Hall, ISBN 0-13-089355-2). This guide describes basic cluster concepts.
- *Serviceguard/SGeRAC/SMS/Serviceguard Mgr Plug-in Compatibility and Feature Matrix*, at <http://docs.hp.com> -> High Availability, under the heading Serviceguard -> Support Matrixes.
- *Securing Serviceguard* and other Serviceguard white papers at <http://docs.hp.com> -> High Availability, under the heading Serviceguard -> White Papers.

For information on the Distributed Systems Administration Utilities (DSAU), see the latest version of the *Distributed Systems Administration Utilities Release Notes* and the *Distributed Systems Administration Utilities User's Guide* at <http://docs.hp.com>: go to the Network and Systems Management collection and choose System Administration.

For information about the Event Monitoring Service, see the following documents at <http://www.docs.hp.com> -> High Availability -> Event Monitoring Service and HA Monitors -> Installation and User's Guide::

- *Using the Event Monitoring Service* (B7609-90049)
- *Using High Availability Monitors* (B5736-90074)

The Event Monitoring Service (EMS) and the Event Monitoring Service (EMS) Developer's Kit are available free from <http://www.hp.com/go/softwaredepot> -> High Availability.

Other relevant HP-UX documentation posted at <http://docs.hp.com> includes:

- *HP-UX System Administrator's Guide*.

This multi-volume manual replaces *Managing Systems and Workgroups* as of HP-UX 11i v3. For information about the organization of the set, see the Preface to the *Overview* volume (5991-7436). All volumes are under System Administration in the HP-UX 11i v3 Operating Environments section of docs.hp.com.



NOTE: For HP-UX 11i v2 systems, continue to use *Managing Systems and Workgroups* in the HP-UX 11i v2 Operating Environments section of docs.hp.com.

- The latest *HP Auto Port Aggregation Release Notes* and other APA documentation under Auto Port Aggregation (APA) in the I/O Cards and Networking Software section of docs.hp.com.
- *The latest version of Using HP-UX VLANs* and other documentation under Virtual LAN in the I/O Cards and Networking Software section of docs.hp.com.

Further Information

Additional information about Serviceguard and related high availability topics can be found at:

<http://www.hp.com/go/softwaredepot> -> High Availability

Online versions of user's guides and white papers are available on Hewlett-Packard's HP-UX Documentation web page:

<http://docs.hp.com> -> High Availability

Support tools and information are available from the Hewlett-Packard IT Resource Centers:

<http://us-support.external.hp.com> (Americas and Asia Pacific)

<http://europe-support.external.hp.com> (Europe)

Release Notes Revisions

Occasionally, important new information warrants revising the Release Notes after they have gone to press. In such cases HP updates the Release Notes on docs.hp.com (<http://www.docs.hp.com> -> High Availability -> Serviceguard -> Release Notes). Versions with the same part number are differentiated by the publication date.

Compatibility Information and Installation Requirements

Read this entire document (and any other Release Notes or READMEs for related products you may have) before you begin an installation.

Compatibility

For complete compatibility information see the *Serviceguard Compatibility and Feature Matrix* posted at <http://www.docs.hp.com> -> High Availability -> Serviceguard -> Support Matrixes.

Compatibility with Storage Devices

For the matrix of currently supported storage devices and volume managers, see <http://h71028.www7.hp.com/enterprise/downloads/External-SG-Storage6.pdf>.

Bastille Compatibility

To ensure compatibility between Serviceguard (and Serviceguard Manager) and Bastille, do the following, depending on your environment: If Bastille is started using Sec10Host (`host.config`) level lock down, change `SecureInetd.deactivate_ident=Y` to `SecureInetd.deactivate_ident="N"`

If Bastille is started using Sec20MngDMZ (`mandmz.config`) level lock down, change `SecureInetd.deactivate_ident=Y` to `SecureInetd.deactivate_ident=N`

If Bastille is started using SIM.config, change `SecureInetd.deactivate_ident=Y` to `SecureInetd.deactivate_ident=N`

If Bastille is started using Sec30DMZ (`dmz.config`) level lock down, change `SecureInetd.deactivate_ident=Y` to `SecureInetd.deactivate_ident=N` and add the following rules to `ipf.customrules`.

```
pass in quick proto tcp from any to any port = 2301
pass in quick proto tcp from any to any port = 2381
pass in quick from <clusternodes> to any
pass out quick from any to <clusternodes>
```

In the above rules, `<clusternodes>` are all nodes in the cluster, including the local node. The `ipf.customrules` file is located under the Bastille directory itself.

IPFilter-Serviceguard rules are documented in the latest *HP-UX IPFilter Version Administrator's Guide*, posted at <http://docs.hp.com/> -> Internet and Security Solutions -> IPFilter

For information on how to configure HP-UX Bastille Sec10Host to allow the `identd` daemon to run, see the latest *HP-UX 11i v3 Installation and Update Guide* posted at <http://docs.hp.com> under Core HP-UX/Operating Environments -> 11i v3.

Mixed Clusters

As of HP-UX 11i v2 Update 2 (0409) and Serviceguard A.11.16, Serviceguard supports "mixed" clusters, consisting of HP 9000 and Integrity servers. Mixed clusters support the same volume managers, at the same version level, as homogeneous Serviceguard clusters. The following restrictions apply (specifics for the current release in parentheses):



NOTE: **HP-UX version** in this context means major release, such as 11i v2. It is acceptable to have a mix of different HP-UX **fusion** releases for the same major revision (for example, 11i v2 September 2004 and 11iv2 September 2006), although it is generally best to have all nodes running the same fusion release.

- All nodes must be running:
 - The same version of HP-UX (11i v3 or 11i v2)
 - The same version of Serviceguard (A.11.18)
 - The same version of any volume manager or file system that is independent of HP-UX
 - The same patch level for LVM and SLVM
 - The same patch level for HP-UX, Serviceguard, and volume managers and related subsystems (for example Veritas VxVM, and VxFS)
- All applications running in the cluster must adhere to the vendor's requirements for mixed Integrity and HP 9000 environments.

Before Installing Serviceguard A.11.18

Before you install Serviceguard A.11.18, you need to make sure that your cluster has the correct hardware upgrades. If you are upgrading older systems, make sure your HP representative reviews the firmware levels of SCSI controller cards and installs the latest versions.

Memory Requirements

Serviceguard needs approximately 15.5 MB of lockable memory on each cluster node.



NOTE: Remember to tune the swap space and the HP-UX kernel parameters *nfile*, *maxfiles* and *maxfiles_lim* to ensure that they are set high enough for the number of packages you are configuring.

Port Requirements

Serviceguard uses the ports listed below. Before installing, check `/etc/services` and be sure no other program has reserved these ports.

- `discard 9/udp`
- `snmp 161/udp`
- `snmp 162/udp`
- `hacl-qs 1238/tcp`
- `clvm-cfg 1476/tcp`
- `hacl-hb 5300/tcp`
- `hacl-hb 5300/udp`
- `hacl-gs 5301/tcp`
- `hacl-gs 5301/udp`
- `hacl-cfg 5302/tcp`
- `hacl-cfg 5302/udp`
- `hacl-probe 5303/tcp`
- `hacl-probe 5303/udp`
- `hacl-local 5304/tcp`
- `hacl-test 5305/tcp`

Serviceguard also uses port `9/udp discard` during network probing setup when running configuration commands such as `cmcheckconf` or `cmapplyconf` and `cmquerycl`. If the port

is disabled (in `inetd.conf`), the network probing may be slower and under some conditions error messages may be written to `syslog`.

Serviceguard also uses dynamic ports (typically in the range of 49152 - 65535) for some cluster services. If you have adjusted the dynamic port range using kernel tunable parameters, alter your rules accordingly.

System Firewalls

When using a system firewall such as HP-UX IPFilter with Serviceguard, you must leave open the open the ports listed above and follow specific IPFilter rules required by Serviceguard; these are documented in the HP-UX IPFilter Administration Guide, available from <http://www.docs.hp.com> -> Internet and Security Solutions.

Installing Serviceguard on HP-UX

Dependencies

The following are required by Serviceguard. They are part of the HP-UX Base Operating Environment:

- Open SSL, which includes the `OPENSSL-RUN` and `OPENSSL-LIB` filesets.
- The “EventMonitoring” bundle, which contains the `EMS-CORE` and `EMS-CORE-COM` filesets.

Installing Serviceguard

Serviceguard will automatically be installed when you install the HP-UX Mission Critical Operating Environment (MCOE).

To install Serviceguard independently, follow these broad steps:



NOTE: This release of Serviceguard requires HP-UX 11i v2 or 11i v3. If you are already running earlier versions of Serviceguard and HP-UX, see “Upgrading from an Earlier Serviceguard Release” (page 56) for more information.

1. Install or upgrade to HP-UX 11i v2 or 11i v3 before loading Serviceguard Version A.11.18.
For information and instructions, see the *HP-UX Installation and Update Guide* for the target release at <http://docs.hp.com>: choose the target version of HP-UX under Operating Environments and then choose Installing and Updating.
2. Use the `swinstall` command to install Serviceguard, product number T1905CA.
For more information about `swinstall`, see the `swinstall (1M)` manpage and the *Software Distributor Administration Guide* for HP-UX 11i v2 or 11i v3.
3. Verify the installation. Use the following command to display a list of all installed components: `swlist -R T1905CA`

The filesets that make up the Serviceguard product are:

- `Serviceguard.CM-SG`
- `SGManagerPI.SGMGRPI`
- `SGWBEMProviders.SGPROV-CORE`
- `SGWBEMProviders.SGPROV-DOC`
- `SGWBEMProviders.SGPROV-MOF`
- `CM-Provider-MOF.CM-MOF`
- `CM-Provider-MOF.CM-PROVIDER`
- `Cluster-OM.CM-DEN-MOF`
- `Cluster-OM.CM-DEN-PROV`
- `Cluster-OM.CM-OM`

- Cluster-OM.CM-OM-AUTH
- Cluster-OM.CM-OM-AUTH-COM
- Cluster-OM.CM-OM-COM
- Cluster-OM.CM-OM-MAN
- Package-CVM-CFS.CM-CVM-CFS
- Package-CVM-CFS.CM-CVM-CFS-COM
- Package-Manager.CM-PKG
- Package-Manager.CM-PKG-MAN
- Cluster-Monitor.CM-CORE
- Cluster-Monitor.CM-CORE-COM
- Cluster-Monitor.CM-CORE-MAN



NOTE: There are files in CM-CORE that are reserved for HP support. Do not change these files. Do not move, alter, or delete the following:

- /usr/contrib/bin/cmcorefr
- /usr/contrib/bin/cmdumpfr
- /usr/contrib/bin/cmfmtfr
- /usr/contrib/Q4/lib/q4lib/cmfr.pl
- /var/adm/cmcluster/frdump.cmcl.d.x (where x is a digit)

NOTE: If you did a swremove of an older version of Serviceguard before the swinstall, a zero-length binary configuration file (/etc/cmcluster/cmclconfig) may be left on your system. Remove this file before you issue the swinstall command. If you do not remove the zero-length binary configuration file, the installation will proceed correctly, but you may see error or warning messages such as:

```
Bad binary config file directory format.
Could not convert old binary configuration file.
```

These messages may safely be ignored.

If You Need To Disable identd



CAUTION: HP does not recommend disabling this security feature, as it maintains the integrity and high availability of your data.

If you must disable `identd`, do it *after* installing Serviceguard but *before* each node rejoins the cluster (for example, just before issuing `thecmrunnodeor` or `cmruncl` command).

Instructions are in Chapter 5 of the *Managing Serviceguard* manual, under the heading “Disabling `identd`” in the section “Managing the Running Cluster”.

Upgrading from an Earlier Serviceguard Release

- For information about supported Serviceguard upgrade paths, see the support matrix at docs.hp.com -> High Availability -> Support Matrixes.
- This release of Serviceguard requires HP-UX 11i v2 or 11i v3. For information about HP-UX upgrade paths, see the *HP-UX Installation and Upgrade Guide* for the target version of HP-UX at <http://www.docs.hp.com> -> <HP-UX target version>-> Installing and Updating.
- If you are upgrading from a release earlier than A.11.16, see “Access changes as of A.11.16” (page 42).

Veritas Storage Management Products

Veritas Volume Manager (VxVM) 3.5 is not supported on HP-UX 11i v3. If you are running VxVM 3.5 as part of HP-UX (VxVM-Base), version 3.5 will be upgraded to version 4.1 when you upgrade to HP-UX 11i v3. For more information about installing and updating VxVM, see the *Veritas 4.1 Installation Guide* at <http://www.docs.hp.com> -> 11i v3 -> VxVM.



NOTE: There is a new default Disk Layout Version for VxVM 4.1, and not all earlier Disk Layout versions are supported on HP-UX 11i v3. See the *Veritas 4.1 Installation Guide* for details.



IMPORTANT: Veritas Cluster Volume Manager (CVM) and Cluster File System (CFS) are not yet supported on HP-UX 11i v3. If you are currently running CVM or CFS, do not upgrade to 11i v3 until support for these Veritas products is available.

Rolling Upgrade

In some cases you can upgrade Serviceguard and HP-UX without bringing down the cluster; you do this by upgrading one node at a time, while that node's applications are running on an alternate node. The process is referred to as a **rolling upgrade**, and is described in Appendix E of the *Managing Serviceguard* manual.

Requirements for Rolling Upgrade

Rolling upgrade to Serviceguard A.11.18 is supported only if you are upgrading from:

- Serviceguard A.11.15 or greater on HP Integrity systems running HP-UX 11i v2 or 11i v3;
or
- Serviceguard A.11.16 or greater on HP 9000 systems running HP-UX 11i v2 or 11i v3.

Rolling Upgrade Exceptions

HP-UX Cold Install

A rolling upgrade cannot include a cold install of HP-UX on any node. A cold install will remove configuration information; for example, device file names (DSFs) are not guaranteed to remain the same after a cold install.

HP Serviceguard Storage Management Suite and standalone CVM product

In many cases you cannot do a rolling upgrade if you are upgrading the HP Serviceguard Storage Management Suite.

Specifically you cannot do a rolling upgrade if you are using the Veritas clustering capabilities - CVM and CFS. In the case of CVM, this applies whether you purchase it as part of a Storage Management bundle, or as a standalone product.

For more information, see the *Veritas 5.0 Installation Guide* and the *HP Serviceguard Storage Management Suite A.02.00 Release Notes*, or (for version 4.1) *HP Serviceguard Storage Management Suite A.01.01 Release Notes* at docs.hp.com -> High Availability -> HP Serviceguard Storage Management Suite.

Migrating to Agile Addressing if Using Cluster Lock Disk

You cannot migrate to the HP-UX 11i v3 agile addressing scheme for device files during a rolling upgrade if cluster lock disks are used as a tie-breaker, because that involves changing the cluster configuration. See "Updating the Cluster Lock Configuration" in Chapter 7 of *Managing Serviceguard* for instructions in this case. See "About Device Special Files (DSFs)" (page 37) of these Release Notes for more information about agile addressing.

Upgrading from an Earlier Release if You Are Not Using Rolling Upgrade (Non-Rolling Upgrade)

If your cluster does not meet the requirements for a rolling upgrade, or you decide not to use rolling upgrade for some other reason, you must bring down the cluster (`cmhaltcl`) and then upgrade Serviceguard and HP-UX on the nodes. See Appendix E of the *Managing Serviceguard* manual.

You can perform a **non-rolling** upgrade (that is, an upgrade performed while the cluster is down) from any supported HP-UX/Serviceguard release to any supported HP-UX/Serviceguard release. See the support matrix at docs.hp.com -> High Availability -> Support Matrixes.

Uninstalling Serviceguard

To uninstall the Serviceguard software, run the SD-UX `swremove` command. Before removing software, note the following:

1. Serviceguard must be halted (not running) on the node from which the `swremove` command is issued.
2. The system from which the `swremove` command is issued must be removed from the cluster configuration. (If the node is *not* removed from the cluster configuration first, `swremove` will cause the current cluster to be deleted.)
3. The `swremove` command should be issued from one system at a time. That is, if Serviceguard is being de-installed from more than one system, it should be removed from one system at a time.

If a zero-length binary configuration file (`/etc/cmcluster/cmclconfig`) is left your system you should remove it.

Patches for this Version

The table below lists patches required or recommended for Serviceguard A.11.18 on HP-UX 11i v2 or 11i v3.

In addition, a patch is needed to activate some of the new features described earlier in this document; see “New Features for April 2008 Patch” (page 17) (for 11i v2) or (page 18) (for 11i v3).

Before installing Serviceguard, you should also check the Hewlett-Packard IT Resource Center web page for any new patch requirements:

<http://itrc.hp.com> (Americas and Asia Pacific)

<http://europe.itrc.hp.com> (Europe)



NOTE: One quick way to see which patches that have been applied to your system is a command such as the following:

```
swlist -l patch | grep applied | more
```

For complete information, see the section “Which Patches Are on a System?” in the *Patch Management User Guide for 11.x Systems*, at <http://docs.hp.com>: choose 11i v3 under Operating Environments and then choose Patch Management.

Table 1-7

Patch Number	Description
PHNE_35894	Patch to enable online replacement of LAN cards on HP-UX 11i v3. Fixes JAGag27186; see “JAGag27186 (SR8606472298): OLR of a LAN Card in SG cluster fails on HP-UX 11i v3” (page 68).
PHNE_34671 (or later) cumulative transport patch	11i v2 patch. Fixes JAGaf75382. See “JAGaf75382 (SR8606415528): lan2 has an IPv6 address (fec0:0:0:f08::36) but not in config. file” (page 59).

Fixed in This Version

This section shows which defects listed as known problems in the Release Notes for the previous release have now been fixed. Specifically, the fixes below are to defects reported in the A.11.17 and A.11.17.01 Release Notes.



NOTE: Serviceguard A.11.18 also includes all fixes already included in patches to earlier Serviceguard versions; these fixes are not necessarily documented here.

JAGaf46654 (SR8606386500): cmcld: software causes connection abort if standby LAN card fails

What was the problem? `cmcld` aborted during startup if a LAN card has failed, or cannot be seen by the system because of other hardware failures.

What is the solution? `cmcld` now handles missing LAN cards during startup. The code has been changed so that a local switch will occur if a LAN card fails and there is a standby available in that bridged net.

JAGaf67601 (SR8606407696): cmcld does not consider newer drivers e.g. c8xx and mpt for cluster lock

What was the problem? `cmcld` had not been updated to handle the drivers for the newer supported cluster lock interface cards such as the Ultra160 and Ultra320 cards which use the `c8xx` and `mpt` drivers respectively. This meant `cmcld` defaulted the cluster lock timings to the default (worst case) leading to longer failover times than would be expected, approximately 60 seconds, rather than 30 seconds which would be expected for the `c720` driver in a simple two-node cluster with a two second `NODE_TIMEOUT`.

What is the solution? The code has been modified to support the Ultra160 and Ultra320, which use the `c8xx` and `mpt` drivers.

JAGaf75382 (SR8606415528): lan2 has an IPv6 address (fec0:0:0:f08::36) but not in config. file

What was the problem? This happened only with an IPv6 configuration. When a primary network interface configured with IPv6 address failed, and the switching of an IPv6 address encountered `ENOMEM`, a message such as the following appeared in `syslog`:

```
There is(are) 1 incidence(s) of :
Sep 11 16:34:38 ptest90 cmclfd[11930]: Failed to switch
fec0:0:0:f08::36 from lan1:34(0) to lan2:27(3):
Not enough space
```

And then any subsequent Serviceguard configuration operations (such as `cmapplyconf` etc.) failed.

What is the solution? Fixed in patch PHNE_34671, “Cumulative ARPA Transport Patch” for HP-UX 11i v2 (see “Patches for this Version” (page 58)). Now if `ENOMEM` is encountered, the IP is left on the source interface and cleaned up on the destination interface.

JAGag09818 (SR8606453031): Cannot modify `SERVICE_HALT_TIMEOUT` to default value (0 seconds)

What was the problem? You could not change the `SERVICE_HALT_TIMEOUT` parameter in the package configuration file back to the default value (0 seconds) if the entire line had been deleted from the file.

What is the solution? The code has been modified to handle the deletion of the entire line correctly.

JAGag11398 (SR8606454772): `cmclfd` aborts with Failed to set keep alive: Invalid argument

What was the problem? `cmclfd` aborted with
Failed to set keep alive: Invalid argument

What is the solution? The code has been modified to close the TCP connection upon failure to set keep-alive with `EINVAL`, instead of aborting `cmclfd`. This should result in re-establishing the connection; in the worst case a cluster re-formation could result.

JAGag12644 (SR8606456223): `cmsrvassistd` doesn't check script permissions

What was the problem? `cmsrvassistd` looped when the script or program specified in a package `service_cmd` parameter did not exist or did not have execute permission, attempting to restart the service until the configured maximum `service_restart` count had been reached.

What is the solution? The code has been modified to allow service restarts only if the service script exists and has execute permission.

JAGag20225 (SR8606464542): Serviceguard probe to CD, DVD devices

What was the problem? `cmquerycl`, `cmcheckconf` and `cmapplyconf` might log errors if CD and DVD drives from TEAC and other manufacturers were connected to a node, although the command succeeded. Messages such as the following might appear in `syslog.log`:

```
cmclconfd: Error looking up device /dev/dsk/c17t1d0: /dev/config is not open.cmclconfd: Unable to open disk
/dev/rdisk/c0t0d0: Error 0
```

TEAC CD/DVD devices were not being excluded from probing because of their unique peripheral descriptions, and as a result were not being detected as CDs or DVDs. Some CD and DVD devices from other manufacturers were being probed because their descriptions were not configured into `cmclconfd`.

What is the solution? `cmclconfd` has been modified to exclude more CD and DVD devices, including specific TEAC devices, from probing.

JAGag21443 (SR8606465899): `cmclfd` aborts when `select()` is interrupted

What was the problem? `cmclfd` aborted when the `select()` system call was interrupted by a signal.

What is the solution? The `select()` system call is now retried a maximum of ten times when it fails because of an interrupted system call.

JAGag25508 (SR8606470417): Log file permissions set to 666

What was the problem? Permissions for log files defined in `/etc/cmcluster.conf` were set to 666 (allowing anyone to read or write to the file).

What is the solution? These permissions are now set to 644, restricting writes to the owner of the file.

JAGag35297 (SR8606481305): cfsmount failed since the package was already started by cmcld

- *What was the problem?* This affected the CFS mount point package. If the package was already down on all nodes when you ran the `cfsmount` command, `cfsmount` could fail, complaining the package has already been started.
- *Resolution:* If `cfsmount` finds any package whose per-node state is `blocked`, it will run `cmhaltpkg -n` to clear the state to `halted`, then run `cmmodpkg` to allow `cmrunpkg` to bring up the package.

JAGag41937 (SR8606489376): Hanging node problem causing 2 out of 4 nodes to die

- *What was the problem?* If a node experienced multiple hangs during cluster reformation, this could cause two nodes to fail: the node experiencing the hang and the node that was the candidate for cluster coordinator.
- *Resolution:* The code has been modified so that the cluster coordinator does not fail.

JAGag48756 (SR8606496838): Serviceguard minimum lan requirements not updated for VM detection information

- *What was the problem?* When a cluster uses virtual machine guests as nodes, the following warning may appear during the network validation phase of the Serviceguard commands `cmapplyconf`, `cmcheckconf` and `cmquerycl`:

```
Minimum network configuration requirements for the cluster have not been met. Minimum network configuration requirements are:
```

- 2 or more heartbeat networks OR
- 1 heartbeat network with local switch (HP-UX Only) OR
- 1 heartbeat network using APA with 2 trunk members (HP-UX Only) OR
- 1 heartbeat network using bonding (mode 1) with 2 slaves (Linux Only).

The warning really means that the minimum requirements are not met *at the immediate network level*. If you see this warning, you need to make sure that the minimum requirements are met in your overall network configuration.

- *Resolution:* You can ignore the warning in a VM configuration if there is one heartbeat network on the guest, backed by a network on the host using APA with two trunk members as in the third case listed in the warning text (HPVM), or using NIC teaming (bonding) as in the fourth case listed in the warning text (VMware ESX Server).

JAGag46475 (SR8606494542): cmquerycl -f line output may display incorrect VG disk device files

- *What was the problem?* If a shared volume group used different device file names in `/etc/lvmtab` on different nodes, the output displayed by `cmquerycl -f` line showed the device files for all nodes as configured for the first node referenced in the `-n <node>` option of `cmquerycl`. This could result in Serviceguard Manager creating a cluster configuration file which had the wrong device file names for the cluster lock volume group for some of the cluster nodes.
- *Resolution:* The code has been modified to display the correct device file names for each node.

JAGag47195 (SR8606495273): Improper log message from cmmodnet on a system on which IPv6 is not enabled

- *What is the problem?* When IPv6 was not enabled on a system, using `cmmodnet` to add an IPv6 address to an interface produced the following incorrect error message:

```
<address> might already be configured as a heartbeat IP or stationary IP address. Or it does not belong to subnet <subnet>.
```

The real problem is that IPv6 is not enabled: `cmmodnet` did not check whether IPv6 is enabled on the system before trying to add the address.

- *Resolution:* `cmmodnet` now checks whether IPv6 is enabled on the system before trying to add the address.

JAGag48606 (SR8606496688): cmmigratepkg missing PEV and external script permission

- *What was the problem?* The `-e` option of `cmmigratepkg` should generate Package Environment Variables (PEVs) but it did not. (See “Migrating Legacy to Modular Packages” (page 31) for more information about `cmmigratepkg`.)

In addition, the permissions on external scripts output by `cmmigratepkg` were wrongly set to 555, giving everyone read and execute permission; 744 is correct, giving only read permission to users other than the owner, who must be the root user.

- *Resolution:* The code has been modified to correct these problems.

QXCR1000594093: cmcheckconf failed to detect cluster lock error

- *What was the problem?* `cmcheckconf` fails to report an error if a lock disk configured in the cluster configuration file is missing or its volume group is not activated. `cmapplyconf` *does* detect and report the error.
- *Resolution:* `cmcheckconf` and `cmapplyconf` both now behave correctly.

QXCR1000739586: Serviceguard does not log .cm_start_time message correctly in syslog after abort

- *What was the problem?* If `cmclld` aborted and you restarted the cluster without first rebooting the node on which the abort occurred, the following message appeared in `syslog.log`:

```
It appears that package applications or resources may be active on this node.Re-starting the cluster could cause data corruption.To recover from this situation reboot this system:After ensuring that no package applications or resources are active, you can override this data. For CFS customers, it is highly recommended that this node be rebooted.
```

This message is missing important information. It should say:

```
It appears that package applications or resources may be active on this node.Re-starting the cluster could cause data corruption. To recover from this situation reboot this system:
```

```
shutdown -r (stops package components)
```

```
After ensuring that no package applications or resources are active, you can override this data integrity protection by issuing the following commands (which allow the daemon to start without rebooting):
```

```
rm /var/adm/cmcluster/.cm_start_time
```

```
touch /var/adm/cmcluster/.cm_start_time
```

```
For CFS customers, it is highly recommended that you reboot the node instead of using the data override mechanism
```

- *Resolution:* The error message now includes the complete version of the directions provided above.

QXCR1000740150: cmcld abort when deleting node online

- *What is the problem?* When you run `cmapplyconf` to delete a node from the cluster configuration while the cluster is running, `cmcld` aborts if the name of the node `cmapplyconf` is run from is a subset of the name of the node being deleted.

For example, if you run `cmapplyconf` from `node1` to delete `node10`, `cmcld` aborts with the following message in `syslog.log`:

```
cmcld: Aborting! Deleting local node that is online(file: rcomm/comm_config.c, line: 583)The cluster daemon
aborted our connection (231).cmapplyconf: Unable to apply the configuration
```

- *Resolution:* The code has been modified to prevent this problem.

QXCR1000740395: Should display warnings about subnets only for related packages

- *What was the problem?* `cmapplyconf` and `cmdeleteconf` issue a warning if a `monitored_subnet` configured for a package is not available on all nodes in a cluster. (The configuration may be correct if this is a cross-subnet configuration; see “About Cross-Subnet Configurations” (page 33).) The problem was that this warning was given for all packages configured for the cluster, not just for the package that was being configured or deleted via `cmapplyconf` or `cmdeleteconf`.
- *Resolution:* The warnings are now given correctly — only for a package that is being added or modified.

QXCR1000741571: Hanging node problem causing cluster to go down

- *What is the problem?* A hang on one node can cause multiple cluster re-formations, until eventually the entire cluster fails.
- *Resolution:* The code has been modified so that hung node no longer causes this problem.

QXCR1000746305: Site controller pkg failed to start after `cmhaltnode -f`

- *What was the problem?* When all the nodes on one site were halted with `cmhaltnode -f`, the site controller package might not run successfully on the remote site. (This problem could occur in **site-aware extended-distance clusters**; see the subsection “For More Information” (page 34), in the section “About Cross-Subnet Configurations”.)
- *Resolution:* The site controller package now runs successfully.

QXCR1000747823: `cmviewconf` does not display managed packages for SC package

- *What was the problem?* `cmviewconf` did not report `Managed Packages` information for the site controller package. (This problem could occur in **site-aware extended-distance clusters**; see the subsection “For More Information” (page 34), in the section “About Cross-Subnet Configurations”.)
- *Resolution:* `cmviewconf` now reports the information correctly. You can also get this information by running `cmviewcl -v -f line`.

QXCR1000749527: `cmmakepkg -l` does not list non-SG modules

- *What was the problem?* The `-l` option of `cmmakepkg` should list all modules that can be used to build a package, but it listed only Serviceguard modules; partner modules such as the MetroRAC `dtS/mcsc` module were not listed.
- *Resolution:* `cmmakepkg -l` now lists all installed modules.

QXCR1000749924: Dependent packages started on non-primary nodes

- *What was the problem?* Failover packages which are dependent on a multi-node package may start on a node other than their primary node even though the dependency is met on the primary node.
- *Resolution:* The code behaves as designed. To ensure a dependent package runs on its primary node whenever possible, set its *failback_policy* to `automatic`, so that if the package starts on a non-primary node, it will be automatically failed back as soon as the primary node is available. For more information, see “About Multi-Node Packages” (page 32).

QXCR1000766057: cmviewcl package state “blocked” is not documented

- *What was the problem?* `cmviewcl` shows the per-node state of multi-node packages as `blocked`, but the manpage does not explain what “blocked” means.
- *Resolution:* “Blocked” means that the multi-node package has never run on the node. This may be because a dependency has not been met, or *auto_run* is set to `no`. This information has been added to the `cmviewcl` manpage.

Dependencies are configured in the package configuration file; see Chapter 6 of *Managing Serviceguard* for more information.

QXCR1000766101: cmhaltpkg reports that auto_run is disabled for multi-node packages but it is not

- *What was the problem?* When you used `cmhaltpkg` to halt packages, `cmhaltpkg` printed the following message:

```
One or more packages or package instances have been halted. These packages have AUTO_RUN disabled and no new instance can start automatically. To allow automatic start, enable AUTO_RUN via cmmodpkg -e <package_name>
```

This is true of the failover packages in the command list but not true of the multi-node packages.
- *Resolution:* `cmhaltpkg` now behaves correctly.

Problems Fixed in this Version of the Serviceguard Manager Plug-in

This section describes defects that have been fixed in Version B.01.01 of Serviceguard Manager.

JAGag05829: Operations log window should point to the operation just performed.

- *What was the problem?* The Operations Log did not distinguish between these two types of refresh actions:
 1. *Refresh actions triggered by the user* In this case, the top summary table should focus on the row selected by the user, and automatically scroll to the row position.
 2. *Refresh actions triggered by an administrative or configuration operation* In this case, the top summary table should focus on the operation row for the operation which just got executed, and automatically scroll to the very top.
- *Resolution:* The operation(s) which occurred in the past 20 seconds (or an arbitrary period of time of ~1 minute) will be visually distinguished by a bold font. You will be able to tell at a glance the most recent operation, even when the focus is on a different row.

JAGag05831: A few operations do not immediately show in the Operations Log

- *What was the problem?* An Operations Log browser window refresh might not be triggered by a non-Administrative or non-Configuration operation.
- *Resolution:* Redesigned DSAU configuration to use standard Operations Log redirect mechanisms.

JAGag07820: Handling invalid cluster name with * / \ space characters is broken

- *What was the problem?* When you specified an invalid character such as * \ / or a space in the Cluster Name field and tried to jump to other tabs, Serviceguard Manager did not warn the user of the problem.
- *Resolution:* Serviceguard Manager now checks for valid cluster name before attempting to create a new cluster.

JAGag07875: Edit cluster summary page: warning about pkg config should be restricted

- *What was the problem?* If a user added a new node to a cluster (which already had packages configured to run on all nodes), the summary page displayed a warning message "The following packages will be automatically configured on the new cluster member nodes:" followed by list of packages that will be configured on the new node.
This warning message always appeared even when there were no packages in the cluster or even when cluster contained only those packages which did not need to be configured on the newly added node.
- *Resolution:* Serviceguard Manager now checks to see if any packages are present before displaying the warning message.

JAGaf97851: Serviceguard Manager hangs when launching System Log Viewer

- *What was the problem?* Serviceguard Manager would hang when launching System Log Viewer using "Display Consolidated syslog" menu item if System Log Viewer is already open.
- *Resolution:* The problem has been fixed.

JAGag11832 - Launching consolidated syslog first time causes SMH logout

- *What was the problem?* SMH automatically logged you out of the system when you attempted to launch the consolidated log window, immediately after configuring it for the first time.
- *Resolution:* The problem has been fixed.

QXCR1000581069: Deleting multiple packages failed if a selected package was not down

- *What was the problem?* Deleting multiple packages in Serviceguard Manager failed if a selected package was not down.
- *Resolution:* Serviceguard Manager prints a message saying which of the selected packages can be deleted and which packages cannot be deleted. You can request Serviceguard Manager to delete the subset of packages which can be deleted.

QXCR1000591254: Tooltip value incorrect

- *What was the problem?* The Tooltip over the Service Status column in a Serviceguard Manager property page listed uninitialized as one of the possible values, which is not correct.
- *Resolution:* The incorrect value has been removed.

QXCR1000741159: Incorrect error message

- *What was the problem?* Serviceguard Manager displayed a misleading error message if you attempted to edit a Serviceguard modular package:
Cannot modify A.11.18 package <package name>
- *Resolution:* The error message now correctly says:
Cannot modify modular (UPCC) package <package name>
For more information, see “About Modular Packages” (page 22), and “Limitations” (page 44) (under “About the Serviceguard Manager SMH Plug-In”).

QXCR1000752875: Using Serviceguard Manager to edit a CFS package with environment variables corrupts the package configuration file

- *What was the problem?* Serviceguard Manager failed to prevent users from modifying or deleting Veritas Cluster File System (CFS) packages.
- *Resolution:* Serviceguard Manager no longer allows you to modify or delete CFS packages, but does correctly display CFS environment variables in the package property sheets.

QXCR1000765143 SGMgr failed to discover Oracle DB and instance on 10G and 11G.

- *What was the problem?* In Serviceguard Extension for Oracle RAC (SGeRAC) installations, Serviceguard Manager failed to discover Oracle 10g RAC and 11g RAC databases and instances.
- *Resolution:* The Oracle database discovery script has been updated to fix this problem.

Known Problems and Workarounds

This section lists known problems in Serviceguard Version A.11.18 at the time of publication. This list is subject to change without notice.

More-recent information may be available on the Hewlett-Packard IT Resource Center:

<http://www.itrc.hp.com> (Americas and Asia Pacific) or

<http://www.europe.itrc.hp.com> (Europe).

Contact your HP support representative for up-to-the-moment information.

JAGad39695 (SR8606170431): User error can result in “ghost” services:

- *What is the problem?* A package was shut down when references to one of its services had already been removed, and the service failed while the package was halting.
In this case, since the service no longer exists, it cannot be halted manually. Any attempt to do so results in an error such as this:

```
cmhaltserv : Service name oasmon is not running.
```


Because the service has failed, `cmsrvassistd` tries to re-start the service (since it was never halted). But when `cmsrvassistd` goes to register the re-started service, its status entry has been removed by `cmclld` because of the package shutdown. When the user issues the `cmapplyconf`, this service is permanently removed from `cmclld` although `cmsrvassistd` is still trying to re-start it. After the `cmapplyconf` the “ghost” service still cannot be stopped because `cmclld` denies its existence.
- *What is the workaround?* One workaround is to make sure you follow the correct procedure when removing a service: halt the package *before* you start editing any file.
However, if you do get into this situation:
 - Follow these steps:
 1. Halt the package
 2. Add the service back into the ASCII file

3. Re-apply the configuration
 4. Manually halt the service with `cmhaltserv`
 5. Re-remove the service from the ASCII file
 6. Re-apply the configuration
 7. Re-start the package
- Another solution is to add a new package with a service name matching that originally deleted and then halt the service with `cmhaltserv`. This would allow the problem to be resolved without halting either the package, node or cluster.
 - Alternatively, if the cluster/node were halted and re-started the problem would go away. This resolution requires cluster/node downtime, however, and that is not always possible in 24 x 7 environments.

JAGad54887 (SR8606185685): `cmquerycl` command may hang when probing disabled disks

- *What is the problem?* If a Serviceguard command which probes disks, such as `cmquerycl`, encounters a disabled disk, the command may hang.
- *What is the workaround?* If the overall disk configuration (number of LUNs) is not large, wait for the command to complete; otherwise terminate the command with **CTRL-C**.

JAGaf32484 (SR8606372066): Problems doing a rolling upgrade when moving from CVM 3.2 to 3.5

- *What is the problem?* In a Serviceguard cluster with CVM 3.2, it is not possible to perform a rolling upgrade from CVM 3.2 to CVM 3.5. After a node is upgraded to CVM 3.5, when it attempts to rejoin the CVM cluster, it will fail, possibly causing a system reset.

This problem exists regardless of whether the HP-UX operating system revision was changed during the rolling upgrade. "Rolling upgrade" in this case means halting one node in the cluster, upgrading CVM to 3.5 and then starting the node back in the cluster, then performing the same steps for each node in the cluster. The cluster is never completely halted. However, this issue does require the cluster to be completely halted for at least some period of time, since it is not possible to run the cluster with a mix of CVM 3.2 and 3.5.

- *What is the workaround?*

A workaround exists to do a semi-rolling upgrade to a cluster being upgraded from CVM 3.2 to 3.5. It requires a small amount of downtime though.

This workaround involves managing the nodes in the cluster so that all nodes in the clusters are running only CVM 3.2 or only CVM 3.5, but never a combination of the two. To do this:

1. Halt Serviceguard on one or more of the cluster nodes while the other nodes remain in the cluster, still running CVM 3.2.
2. Perform the upgrade from CVM 3.2 to 3.5 on the nodes which are not running in the Serviceguard cluster, but do not restart them in the Serviceguard cluster when the CVM upgrade is complete.
3. Now, when you are able to take the entire cluster down, run `cmhaltcl -f` on one of the nodes that is still running in the Serviceguard cluster.
4. After the cluster has successfully halted, start the Serviceguard cluster on the nodes that were upgraded to CVM 3.5. To do this, you must not run the normal `cmrunccl` command to start up the Serviceguard cluster, since that will attempt to start all cluster nodes. Instead, run `cmviewcl -n node1 -n node2 -n noden`, where the nodes specified are the nodes that were upgraded to CVM 3.5.
5. Now perform the CVM upgrade from 3.2 to 3.5 on the remaining nodes in the cluster.
6. After the upgrade to CVM 3.5 has been completed on each node, the node can rejoin the running Serviceguard cluster by running the command, `cmrunnode` on that node.

For example, in a four-node cluster, you could take down two nodes in the cluster and upgrade CVM to 3.5, while the other two nodes continue running the cluster with CVM 3.2. Then halt the entire cluster, and start up the cluster on the two nodes that were upgraded to CVM 3.5. Now upgrade the remaining CVM 3.2 nodes to CVM 3.5; once the upgrade is complete, they can be added back into the cluster.

JAGaf79389 (SR8606419559): Package not started on adoptive node after `cmhaltnode -f`

What is the problem? A package is not started on an adoptive node after `cmhaltnode -f`.

This happens when a node is shut down and a halt script for a package with `node_fail_fast_enabled` set to `yes` exceeds its `halt_script_timeout`. Serviceguard kills the halt script's process, and the node performs a system reset. The package is not restarted on an adoptive node.

What is the workaround? There is no workaround.

JAGag27186 (SR8606472298): OLR of a LAN Card in SG cluster fails on HP-UX 11i v3

What is the problem? Online replacement (OLR) of a LAN card in a Serviceguard cluster fails on a system running HP-UX 11i v3 because the Critical Resource Analysis (CRA) performed as part of the OLR operation returns `CRA_SYS_CRITICAL`. You will encounter this problem on HP-UX 11i v3 whether you use the Peripheral Device Tool (`pdweb`) or the HP-UX `olrad` command.

What is the workaround? Apply patch PHNE_35894. (See "Patches for this Version" (page 58) for more information about patches.)



NOTE: You can apply the patch without a reboot.

On a system to which the patch has been applied, you will be able to perform online replacement of hot-swappable cards (without bringing down the cluster). See "Replacing LAN or Fibre Channel Cards" in Chapter 8 of *Managing Serviceguard* for more information. (You can find the manual at <http://docs.hp.com> -> High Availability -> Serviceguard.)



NOTE: If for some reason you need to proceed without the patch, you must follow the "Off-Line Replacement" procedure under "Replacing LAN or Fibre Channel Cards" in Chapter 8 of *Managing Serviceguard*.

JAGag31490 (SR8606477006): Modular package fails with the error: 'Unable to retrieve package attribute'.

What is the problem? In clusters with over 100 packages configured, when the cluster and packages start simultaneously, some packages may fail to start. Errors may appear in the log files of modular package log files, saying Unable to retrieve package attributes, and that the package failed to start. Messages such as the following appear in the package log file:

```
Feb  7 09:00:50 root@aialik.cup.hp.com master_control_script.sh[11121]: ERROR: Unable to retrieve package attributes. Feb  7 09:00:50 root@aialik.cup.hp.com master_control_script.sh[11121]: ERROR: Function sg_source_pkg_env
```

The messages indicate that the `cmgetpkgenv` command had difficulty obtaining information about the package that is required for the package scripts to run. This does not mean that the package cannot be started after the other packages are stable.

What is the workaround? If you see this problem, you can issue a `cmrunpkg` on the packages that have not started.

JAGag43145 (SR8606490767): cmcheckconf fails for non-existent deferred filesystem EMS resources

- *What is the problem?* The `cmcheckconf` command fails if you are trying to apply a new or existing package which contains a non-existent deferred EMS resource. An example of such a resource is a file system that will not be mounted until the package has started and so cannot be verified as a resource during the `cmcheckconf` and `cmapplyconf`.
- *What is the workaround?* Manually create the resource before running `cmcheckconf` and `cmapplyconf`. For example, for a file system, do the following:
 1. On the first node, activate the volume group and mount the filesystem.
 2. On the other nodes, activate the volume group in read-only mode (`vgchange -r`) and mount the filesystem read-only.
 3. Run `cmcheckconf` and/or `cmapplyconf`.
 4. Unmount the filesystem(s) on each node.
 5. Deactivate the volume group on all nodes.

JAGag44319 (SR8606492148): cmcluster boot script does not print cmrunnode errors

- *What is the problem?* If `AUTOSTART_CMCLD` is set to 1 and `cmrunnode` starts the cluster successfully after a node reboot, connectivity errors on startup, if any, are not logged to `/etc/rc.log`.
- *What is the workaround?* The `cmrunnode` errors (but not the command log) can be found in the `syslog` file.

QXCR1000768063: Serviceguard A.11.18 network validation errors missing

- *What is the problem?* When network validation fails when a cluster node is started, the errors are not recorded. For example, the following is the output from `cmruncl` when it detects a network configuration error:

```
cmruncl: Validating network configuration...cmruncl: Failed to
validate the network configuration as reported above but will try
to start the cluster anyway.cmruncl: Network validation
completecmruncl: Validating cluster lock disk .... DoneWaiting for
cluster to form ..... doneCluster successfully formed.Check the
syslog files on all nodes in the cluster to verify that no warnings
occurred during startup.
```

Similar output appears in `syslog`. The errors “reported above” are not reported in either case. `cmrunnode` *does* report network validation errors, but it, too, fails to log them to `syslog`.

- *What is the workaround?* There is no workaround at present.

QXCR1000770172: cmgetconf fails if some PV Links are missing

- *What is the problem?* If one of the PV (physical volume) links to an LVM volume group is missing, `cmgetconf` (and Serviceguard Manager) will fail to identify the cluster’s volume groups.
- *What is the workaround?* There is no workaround at present.

Known Problems and Workarounds for the Serviceguard Manager Plug-in

This section documents known problems and suggested workarounds for Serviceguard Manager B.01.01.

JAGag35535: Adding a node to a cluster with Lock LUN fails

- *What is the problem?* Adding node to an existing cluster that has Lock LUN configured fails.
- *What is the workaround or available patch?* Use the command `cmquerycl` directly, specifying each new node with the `-L` option.

JAGag34285: Deleting multiple packages will fail if a selected package is not down

- *What is the problem?* Although Serviceguard Manager displays the appropriate popup warning message that the non-down packages cannot be deleted, when you click OK to proceed with the down packages, the operation fails.
The Operations Log indicates that Serviceguard Manager attempted to delete all selected packages, including those packages that were not down.
- *What is the workaround or available patch?* Make sure that all selected packages are down (not running) before attempting to delete them.

JAGag04773: Clicking on the browser refresh button takes you back to main page

- *What is the problem?* Clicking the browser refresh button from any Serviceguard Manager page takes you back to the Serviceguard Manager main page.
- *What is the workaround or available patch?* There is no workaround for this problem.

JAGag32353: With Mozilla or Firefox browser, configuring a SGeRAC package is confusing

- *What is the problem?* When creating or editing a SGeRAC package using a Mozilla or Firefox browser, if a subnet is marked as monitored then changing the `Cluster Interconnect Subnet` flag unmarks the subnet.
- *What is the workaround or available patch?* If the subnet becomes unmarked after changing `Cluster Interconnect Subnet`, select it again by clicking on the check box.

JAGag22514: Operations Log focus returns to previous selection after operation is over

- *What is the problem?* If the Operations Log window lists more than one operation, and you select a radio button corresponding to an operation, other than the first one in the list, and you perform another administration or configuration operation, the operations log shows the operation that was performed and the corresponding description in the bottom table. However, once the operation is over, the focus comes back to the previously selected radio button.
- *What is the workaround or available patch?* Scroll to the top of the operations log and locate the operation you want to review.

JAGag05524: When editing roles and monitored resources, some buttons may be disabled

- *What is the problem?* When configuring a cluster or a package:
 1. If you forget to press the Add button after entering values in the form, the data entered in the form is lost.
 2. If an incorrect value is entered in the form, no error message is displayed when Add is pressed. The error message is displayed only when a different tab is selected. In some cases, the line that contained an invalid value is removed from the list.
 3. In some cases, if there is an error message caused by incorrect input, some buttons become disabled.
- *What is the workaround or available patch?* Do one of the following:
 - Make sure to click the Add button before moving to another tab.
 - Make sure you enter valid information before clicking Add or moving to another tab.
 - If the Remove button becomes disabled, switch to another tab and back.

- If the Edit button is disabled, remove the line and add it again.
- If some entered data gets lost after an error message, re-enter the data.

JAGag29458: In matrix views, empty columns are not displayed

- *What is the problem?* Matrix views do not display columns that are empty.
- *What is the workaround or available patch?* There is no workaround for this problem.

JAGag32176: With Mozilla 1.5, Auto Refresh button appears after first auto-refresh

- *What is the problem?* In Mozilla 1.5, the Auto Refresh button on the main page does not appear until the first automatic refresh. This problem occurs only in the Mozilla 1.5 browser. It does not occur in later versions of Mozilla or any other supported browsers.
- *What is the workaround or available patch?* There is no workaround for this problem.

JAGag31375: SGeSAP missing from cluster extensions

- *What is the problem?* You may have the SGeSAP extension installed or configured, but the Extension tab does not display SGeSAP in the list.
- *What is the workaround or available patch?* There is no workaround for this problem.

JAGag28015: All dropdown menu options become enabled when clog is re-launched

- *What is the problem?* All dropdown menu options become available (enabled) when the clog package is re-launched. This problem occurs when the clog package is running on a node other than the one running the SMH (System Management Homepage).
- *What is the workaround or available patch?* There is no workaround for this problem.

About Serviceguard Releases

Types of Releases and Patches

Versions of Serviceguard are provided as **platform releases**, **feature releases**, or **patches**.

Platform Release

A platform release is a stable version of Serviceguard, which is the preferred environment for the majority of Serviceguard customers. Platform releases may also contain new Serviceguard features. These releases are supported for an extended period of time, determined by HP. Patches will be made available within the extended support time frame even though a newer version of Serviceguard is available.

Serviceguard A.11.18 is a platform release.



NOTE: You are encouraged to move to the A.11.18 release. For compatibility information for this and earlier releases, see the Serviceguard Compatibility and Feature Matrix at <http://www.docs.hp.com> -> High Availability -> Serviceguard -> Support Matrixes. See also “Compatibility Information and Installation Requirements” (page 53) of this document.

Feature Release

A feature release contains new Serviceguard features. Feature releases are for customers who want to use the latest features of Serviceguard. In general, feature releases will be supported until a newer version of Serviceguard becomes available. In order to receive fixes for any defects found in a feature release after a newer version is released, you will need to upgrade to the newer, supported version.

Patch

Normally a patch to a release is issued in response to a critical business problem found by a Serviceguard customer, but occasionally a patch may enable new features. (Such features do not affect the running of an existing Serviceguard cluster until you activate them, for example by reconfiguring the cluster or reworking a package.)

In the case of a patch, the following is guaranteed:

- Patch-specific release testing is performed prior to posting the patch.
- Existing functionality, scripts, etc. will continue to operate without modification.
- All fixes from the previous patch are incorporated.

Certification testing for a patch is recommended only for those fixes that are important to your specific installation, or for new features that you intend to use.

Supported Releases

For information about the support life of a given release and its compatibility with versions of related products, see the *Serviceguard Compatibility and Feature Matrix*, on `docs.hp.com` -> High Availability, under the heading `Serviceguard` -> Support Matrixes.

Version Numbering

Beginning with A.11.14.01, Serviceguard releases employ a version numbering string that helps you identify the characteristics of the release. The version string has 4 parts:

- Alphabetic Prefix
- First Numeric Field
- Second Numeric Field
- Third Numeric Field

When a new release is issued, different portions of the version string are incremented to show a change from a previous version of the product.

Software Availability in Native Languages

Serviceguard Manager provides Native Language Support, but the command-line interface does not.



IMPORTANT: Even though the command-line interface does not provide Native Language Support, Serviceguard functions correctly whether or not the `LANG` variable is set to `C`.

The Serviceguard Manager plug-in supports the following languages in addition to American English:

- Japanese
- Simplified Chinese
- Korean
- Traditional Chinese
- Standard French
- Standard German
- Standard Italian
- Standard Spanish

Serviceguard documentation is translated into the following languages:

- Japanese
- Simplified Chinese
- Korean
- Traditional Chinese