Serviceguard Version A.11.16
Release Notes
Second Edition

Manufacturing Part Number: B3935-90078
September 2004

Legal Notices

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Serviceguard, Serviceguard OPS Edition, Serviceguard Extension for RAC, Serviceguard Extension for Faster Failover, and Serviceguard Manager are products of Hewlett-Packard Development Company, L.P., and all are protected by copyright.

Corporate Offices:

Hewlett-Packard Co.
3000 Hanover St.
Palo Alto, CA 94304

Use, duplication or disclosure by the U.S. Government Department of Defense is subject to restrictions as set forth in paragraph (b)(3)(ii) of the Rights in Technical Data and Software clause in FAR 52.227-7013.

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Use of this manual and flexible disc(s), compact disc(s), or tape cartridge(s) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs in their present form or with alterations, is expressly prohibited.

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

1 Serviceguard Version A.11.16 Release Notes, Second Edition

Printing History

Table 1-1 Printing History for Serviceguard Version A.11.16 Release Notes (HP part number B3935-90075)

Printing Date	Revision
Sept 2004	Updated for HP-UX 11i v2, September 2004
June 2004	Initial release of version A.11.16 of Serviceguard

The second edition is revised with content for HP-UX 11i v2 September 2004 update (externally also known as HP-UX 11i v2 update 2)

Announcements

Serviceguard is a specialized software product that protects mission-critical applications from a wide variety of hardware and software failures, and ensures data integrity.

The following Serviceguard versions are now available:

- For HP-UX 11i v2 (B.11.23):
  - Product T1905BA — A.11.16 — software and license
  - Product T1906BA — A.11.16 — documentation
- For HPUX 11i v1 (B.11.11):
  - Product B3935DA — A.11.16 — software and license
  - Product B3936EA — A.11.16 — documentation

Serviceguard includes the following component product:

- B8324BA (Cluster Object Manager, or COM B.03.00.01) — provides the ability to query Serviceguard cluster status and convey commands and information to and from Serviceguard Manager.

The following products are delivered free with Serviceguard, on the Distributed Components CD:

- Serviceguard Manager, the graphical user interface for Serviceguard, and the *Serviceguard Manager Version A.04.00 Release Notes*. The Serviceguard Manager software can also be downloaded directly from `http://software.hp.com`.
- Quorum Server, and the *Quorum Server Version A.02.00 Release Notes*

Serviceguard A.11.16 is being released for use with the HP-UX version 11i v1 (11.11) and 11i v2 (11.23) operating system, including the Mission Critical Operating Environment (OE).

Since Serviceguard configurations may be complex to configure and maintain, it is strongly recommended that you use Hewlett-Packard's high availability consulting services to ensure a smooth installation and rollout. Please contact your HP representative to inquire about high

availability consulting. In addition, you should work with your HP representative to ensure that you have the latest firmware revisions for disk drives, disk controllers, LAN controllers, and other hardware.

What's in this Version

The A.11.16 version of Serviceguard is a platform release. This release contains new functionality, defect repairs, and support for new hardware configurations.

This second edition is revised with content for HP-UX 11i v2 September 2004 update (externally also known as HP-UX 11i v2 update 2).

Highlights of the release are as follows:

- Support for servers and clusters:
  - With the HP-UX 11i v1 release, Serviceguard A.11.16 can be installed on HP 9000 servers, and can have clusters up to 16 nodes.
  - With the HP-UX 11i v2 release, Serviceguard A.11.16 can be installed on HPIntegrity servers and can have clusters up to 8 nodes.
  - With the HP-UX 11i v2 September 2004 update, Serviceguard A.11.16 can be installed on either HP Integrity servers or HP 9000 servers, and can have clusters up to 16 nodes. The nodes in each cluster can be HP 9000 servers, HPIntegrity servers, or a combination of both.
- A new method for non-root access for Serviceguard commands, both on command line and with Serviceguard Manager, the graphical user interface. Non-root access to view or to issue administration commands is now defined in the new Access Control Policy parameter in the configuration files of the cluster or one of its packages.

  If a Serviceguard A.11.16 configuration has been applied to a node, Serviceguard will no longer look at the `.rhosts` or `cmclnodelist` files. Instead, it will check for Access Control Policies in the cluster and package configuration files. Root user on a node is always allowed access, but all other users must be listed in at least one Access Control Policy.

- Clusters and packages can now be configured through Serviceguard's graphical user interface, Serviceguard Manager. This interface now replaces all the functionality of SAM Cluster Tool. The Cluster Tool in SAM has been discontinued with the A.11.16 version of Serviceguard.
- A new parameter, Network Failure Detection, gives users a choice about how a network monitor decides to declare a LAN card down.
- A new product, Serviceguard Extension for Faster Failover, can be used with Serviceguard version A.11.16. The Release Notes are posted at `http://docs.hp.com/hpux/ha`. At the same location, there is a new technical white paper, *Optimizing Failover Time in a Serviceguard Environment,* that tells you how you might reduce failover time with or without the purchased product.

Serviceguard A.11.16 supports the same configurations as previous versions, with these exceptions:

- Single-ended SCSI disk devices are no longer supported as cluster lock devices. Configurations which require a single-ended SCSI cluster lock device cannot upgrade beyond Serviceguard version A.11.15.
- `cmclnodelist` is no longer used for access, except as a bootstrap to first configure a cluster. It has been replaced by Access Control Policies in the configuration files. Once a cluster configuration file exists, all non-root users must be listed in an access policy in that file or in a package configuration file.
- If Serviceguard Extension for Faster Failover (SGeFF) is installed, there is a new option in the cluster configuration file to enable or disable SGeFF.
- A new parameter, Network Failure Detection, will let you change the way network failures are determined. The default is the same method used in earlier versions.
- On the September 2004 update to HP-UX 11i v2, a cluster can have nodes with HPIntegrity servers or HP 9000 servers, or both.

For more information, see the Serviceguard manual, *Managing Serviceguard, 11th Edition*, posted at `http://docs.hp.com/hpux/ha`

Access Control Policies

Non-root access to Serviceguard is now defined in the cluster and package configuration files, in a parameter called Access Control Policy.

You can have up to 200 policies in a cluster. Policies can be added, modified, or deleted from the configuration without halting the cluster or the package. Conflicting or redundant policies will cause an error at `cmapplyconf`, and the configuration change will fail.

Each policy has three parts:

- USER_NAME This can be any user that is defined in the USER_HOST's /etc/passwd file.
- USER_HOST This is the node where the user will log in to issue commands (not necessarily the node where the commands take effect).
- USER_ROLE This is the role, or capabilities granted to the user:
  - Monitor: The user can view the cluster objects (read-only). It is defined in the cluster configuration file.

    On the command line, users can issue `cmviewcl, cmgetconf, cmviewconf`, and `cmquerycl`.

    In the graphical user interface, this user can see information about the Serviceguard cluster on the map and tree, and in the Properties.
  - Package Admin: Includes Monitor privileges. The user can issue commands to administer the package.

    On the command line, users can issue: `cmrunpkg`, `cmhaltpkg`, and `cmmodpkg`.

    In the graphical user interface, these menu choices are offered: run or halt a package, move a package from one node to another, and change the node- and package-switching flags.
    - If defined in a package configuration file, the user can administer that specific package.
    - If defined in a cluster configuration file, the user can administer all packages in the cluster.

— Full Admin: Includes Monitor and Package Admin privileges. This user can issue commands to administer the cluster. It is defined in the cluster configuration file.

On the command line, users can issue the `cmhaltcl`, cmruncl, `cmhaltnode`, and `cmrunnode` commands.

In the graphical user interface, these menu choices are offered: run or halt a cluster, run or halt a node, and run or halt a System Multi-node Package.

Configuring

Configuration still requires root permission (UID=0) on a cluster node. Root users do not need an Access Control Policy. On the command line, you must log in as root to a node. In Serviceguard Manager, you will be prompted for a root password when you begin configuration.

Now you can do configuration through Serviceguard Manager, Serviceguard's graphical user interface. With the new configuration capabilities, Serviceguard Manager has replaced the functionality of the SAM Cluster tool; the tool no longer appears in SAM.

Serviceguard Manager is available free from software.hp.com, and is delivered on the Distributed Components CD that comes free with Serviceguard. For more information, see the Serviceguard Manager Version A.04.00 Release Notes at `http://docs.hp.com/hpux/ha`.

Network Failure Detection

A new cluster configuration parameter, Network Failure Detection, now offers two choices for detecting when a LAN card is down:

- INOUT, the present method is the default. With this method, the Network Manager polls between cards, and watches the message count. It will declare a card down only if both the inbound and outbound count stops. Then it will attempt failover.
- INOUT_OR_INONLY, the new method, includes INOUT, so it will declare a card down if both the inbound and outbound count stops. It also adds INONLY, which will also declare a card down, if only the inbound traffic count stops. It will attempt failover in either case.

If you are not sure which method is best for your environment, consult the technical white paper, *Serviceguard Network Manager*, posted at http://docs.hp.com/hpux/ha -> Serviceguard -> White Papers.

What Documents are Available for This Version

The following documents relate to Serviceguard A.11.16 and related products. They can be found on the web at `http://docs.hp.com/hpux/ha`.

- *Managing Serviceguard*, A.11.16 (B3936-90079). This manual has been revised for the A.11.16 release.
- *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 Extension for RAC Version A.11.16 Release Notes, Second Edition* (T1907-900011)
- *Serviceguard Extension for SAP Version B.03.11 Release Notes* (T2357-90004)
- *Managing Serviceguard Extension for SAP* (T2357-90005)
- *Metrocluster with EMC SRDF Version A.04.13 Release Notes*, Second Edition (B6264-90020)
- *Metrocluster with Continuous Access XP Version A.04.22 Release Notes, Second Edition* (B8109-90017)
- *Continental Clusters Version A.04.02 Release Notes, Second Edition* (T2346-90005)
- *Serviceguard Extension for Faster Failover Version A.01.00 Release Notes* (T2388-90001)
- *ECM Toolkit* (B5139-90070)
- *Rules for Configuring Serviceguard with HP 9000 and HP Integrity Systems on HP-UX 11i v2 Septemper 2004 update* posted at `http://docs.hp.com/hp-ux/ha` -> Serviceguard -> White Papers

For information about the Event Monitoring Service, refer to:

- *Using the Event Monitoring Service* (B7612-90015)
- *Using High Availability Monitors* (B5736-90025)

- *Writing Monitors for the Event Monitoring Service* (B7611-90016) available from http://software.hp.com in the "High Availability" area.

Other relevant information for HP-UX systems is found in:

- *Managing Systems and Workgroups* (5187-2216)
- *HP Auto Port Aggregation Release Notes* (J4240-90024)
- *Using HP-UX VLAN* (T1453-90001)
- *Managing Serviceguard NFS* (B5140-90021)

For information about on-line replacement of PCI LAN cards, refer to

- *Configuring HP-UX for Peripherals* (B2355-90698)

Before attempting to use VxVM storage with Serviceguard, please refer to the following:

- *VERITAS Volume Manager Administrator's Guide*. This contains a glossary of VERITAS terminology.
- *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- *VERITAS Volume Manager for HP-UX Release Notes*

Further Information

Additional information about Serviceguard and related high availability topics may be found on Hewlett-Packard's HA web page:

`http://www.hp.com/go/ha`

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

`http://docs.hp.com/hpux/ha` and `http://docs.hp.com/linux`

Support information is available from the Hewlett-Packard Electronic Support Centers:

`http://www2.itrc.hp.com`

To receive the latest news about recommended patches, product support matrices, and recently supported hardware, go to the IT Resource Center site above, and subscribe to the *High availability programs tips and issues digest*.

Compatibility Information and Installation Requirements

Read this entire document and any other Release Notes or READMEs you may have before you begin an installation.

Compatibility

Serviceguard version A.11.16 is compatible with HP-UX 11i v1 (11.11) and 11i v2 (11.23) The first release of 11i v2 includes the Cluster Object Manager (COM) Version B.03.00. The 11i v2 September 2004 update, however, includes COM Version B.03.00.01. The COM is the API (application programmable interface) that enables communication between Serviceguard objects and Serviceguard Manager, the graphical user interface.

All the nodes in a cluster should have the same version of the COM. Version B.03.00.01 is recommended. Use the `swlist` command to see the version of COM (B8324).

For more complete compatibility information, including older versions, see the Serviceguard Compatibility and Support Matrix posted on `http://www.docs.hp.com/hpux/ha` -> Serviceguard -> white papers.

Types of Releases and Patches

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

Platform Release

A platform release may contain new Serviceguard features. Platform 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.

Feature Release

A feature release adds new Serviceguard features to the platform release. Feature releases are for customers who desire 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, the customer will need to upgrade to the newer, supported version.

Patch

A patch to a release is issued in response to a critical business problem found by a Serviceguard customer. With 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.

For a patch, certification testing is recommended only for those fixes that are important to the customer.

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 four 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 following these guidelines, to show particular types of change from a previous version of the product:

Before Installing Serviceguard A.11.16

Before you install Serviceguard A.11.16, you need to make sure that your cluster has the correct hardware upgrades. If you are upgrading Servcieguard on an older system, have your HP representative review the firmware levels of F/W SCSI controller cards and install the latest versions.

Required Firmware Upgrades for FibreChannel SCSI Multiplexer Model A3308A

The Model A3308A FibreChannel SCSI Multiplexer should be upgraded to Firmware revision 3810 (980611) or newer. This firmware revision supports SCSI II Reserve/Release functionality. This is required in order to enforce exclusive access to tape devices.

Setting of SCSI Auto-Termination

Highly available shared SCSI buses require the use of SCSI Inline Termination cables and SCSI V-cables. The use of these cables requires that SCSI Auto-Termination be set OFF on the SCSI host bus adapter. On some SCSI host bus adapters, auto-termination is configured by switches on the card; while on other SCSI host bus adapters, auto-termination is configured programmatically through the HP 9000 PDC (**processor dependent code**, or firmware).

On recent HP 9000 systems, the PDC has been modified so that SCSI auto-termination is enabled by default for some SCSI host bus adapters. In addition, many newer SCSI host bus adapters now have a default setting of auto-termination enabled. In order to use the SCSI Inline Terminator and V-cables, you must ensure that the auto-termination is disabled on the SCSI host bus adapters. Be sure to disable the auto-termination for all SCSI host bus adapters connected to a shared SCSI bus. This is done through the PDC configuration utility (refer to the manuals for your SCSI host bus adapter for instructions for how to do this). For SCSI cards where this configuration is set manually using switches on the card, ensure that auto-termination is disabled. If you are upgrading to a newer PDC version, you must check the auto-termination parameter and re-set it to disabled for each SCSI host bus adapter connected to a shared SCSI bus.

NOTE Cards that cannot be programmatically configured may show an auto-termination setting of either OFF or UNKNOWN. Either of these is acceptable.

Restriction on LUNs on 11i v2 September 2004 update

Configurations using VxVM or CVM with 11i v2 September 2004 update are restricted to a maximu of 256 LUNs, pending fixes to JAGaf36081 and JAGaf36760. For updates on these issues, please check the web page `http://www2.itrc.hp.com`. Search the Technical Knowledge base for bug reports using the defect number as the keyword.

Access changes in version A.11.16

Serviceguard version A.11.16 introduces a new access method. Instead of looking at the `cmclnodelist` or `.rhosts` files, Serviceguard will check the configuration files for an Access Control Policy.

If you want any non-root access to a cluster, you must configure an Access Control Policy in the cluster configuration or in a package configuration file.

There is a brief description of policies in "What's in This Version" above. For more information about role-based access, see the Managing Serviceguard manual (`http://docs.hp.com/hpux/ha`), in the Serviceguard Manager help (under Configuring Clusters -> Roles), and in the cluster and package configuration files themselves.

Upgrading Serviceguard

Access policies are not supported in a mixed revision cluster.

If you relied on `.rhosts` for access, be sure to manually configure policies for those users.

When you upgrade from an earlier version, Serviceguard converts `cmclnodelist` entries into Access Control Policies when every node in the cluster finishes updating. All are given the role of Monitor in the cluster configuration file.

After you complete a rolling upgrade, be sure to create and save a copy of the new configuration, using the `cmgetconf` command. If an `cmapplyconf` is issued, you want to be sure it applies the newly migrated Access Control Policies.

Newly installed Serviceguard

When you install Serviceguard for the first time on a node, you do not have a cluster. Without cluster configuration, however, you cannot have an Access Control Policy.

You can, however, create a `cmclnodelist` file to act as a "bootstrap" monitor access. Bootstrap files are useful if you are doing a rolling upgrade, so the nodes with older versions can still access the newer cluster nodes. Monitor access in a bootstrap file also means the node can appear in Serviceguard Manager; you can see information about it in Properties, and you can configure it into a cluster after you give the root password.

To create a bootstrap file:

1. Create the file `/etc/cmcluster/cmclnodelist` on the node.
2. Using any ascii editor, add a comment like this:

   ```
   ####################################################
   # Do not try to configure access in this file.
   # This is for bootstrapping only, before a cluster is configured.
   # Once a cluster exists, Serviceguard will ignore this file.
   ####################################################
   ```

3. Below the comment, create monitor access. You can add a number of <host_node> <user> pairs, but it may be easiest to add a wildcard + (plus) below the comment. This is equivalent to granting the view-only Monitor role to Any User from Any Serviceguard Node.

Using Serviceguard A.11.16 to reach earlier versions

For Serviceguard clusters with versions earlier than A.11.16, access is granted in the `cmclnodelist` or `.rhosts` file. Only a root user can modify these files to grant access.

In pre-A.11.16 clusters, the only role for a non-root user is Monitor.

To monitor a cluster, modify a (pre-A.11.16) cluster node's `cmclnodelist` file. Read-only access is granted by entering the pair *<user_hostname> <NonRootUser_name>*. Or, you can enter a + (plus) wild card to allow any user.

A command line user can issue the `cmviewcl` command with this entry.

A Serviceguard Manager user can view a cluster in the map and tree, and read Properties of all the cluster's objects. A Serviceguard Manager user can issue administrative commands if they log in to a Session Server as root. They can issue configuration commands if they give the root password for one of the cluster's nodes.

Serviceguard Manager checks two places for access: once when the user logs in to the Session Server, and again when the Session Server contacts the target node. For more information about Serviceguard Manager policies, see the Serviceguard Manager Release Notes, or the online help.

For versions earlier than A.11.16, the `/.rhosts` file must not allow write access by `group` or `other`. If `/.rhosts` file write permission is enabled for `other` or `group`, Serviceguard commands will fail, logging a "Permission denied for user" message. This situation can arise when the Serviceguard remote security file, `/etc/cmcluster/cmclnodelist`, is not used and remote node security is resolved with the `.rhosts` facility instead. (These rules apply only to target nodes with versions earlier than Serviceguard A.11.16.)

Upgraded Serviceguard

If you have a cluster with entries in cmclnodelist, those entries are also updated when you update to version A.11.16.

Every *<host_node>* *<user_name>* pair is now in the cluster configuration file as an Access Control Policy triplet, with:

- USER_NAME *<user_name>*
- USER_HOST *<host_node>*
- USER_ROLE Monitor

If you had a wild-card + (plus), you see an access policy with wildcards:

- USER_NAME ANY_USER
- USER_HOST ANY_SERVICEGUARD_NODE
- USER_ROLE Monitor

Memory Requirements

The memory requirement for Serviceguard depends partly on the number of packages configured in the cluster. The following equation provides a rough estimate of how much lockable memory is needed:

```
Total Memory = 6MB + 100 KB/package
                          in cluster
```

The total amount is needed on all nodes in the cluster, regardless of whether a given package is on that node or not.

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.

- hacl-hb 5300/tcp
- hacl-gs 5301/tcp
- hacl-cfg 5302/tcp
- hacl-cfg 5302/udp
- hacl-probe 5303/tcp
- hacl-probe 5303/udp
- hacl-local 5304/tcp
- hacl-test 5305/tcp
- hacl-dlm 5408/tcp
- Serviceguard also uses port 9/udp discard during network probing setup when running configuration commands such as `cmcheckconf` or `cmapplyconf` and `cmquerycl`. If it is disabled, the command will fail and you will get an error message in `syslog`.

In addition, 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, specific communications must be allowed to ensure proper cluster operation. Specific IPFilter rules required by Serviceguard are documented in the HP-UX IPFilter Release Notes, available from `http://www.docs.hp.com -> Internet and Security Solutions`.

General guidelines for using a system firewall with Serviceguard are listed below.

- To enable intra-cluster communications, each HEARTBEAT_IP network on every node within the cluster must allow the following communications in both directions with all other nodes in the cluster:
  - tcp on port numbers 5300-5304, and 5408 - and allow only packets with the SYN flag
  - udp on port numbers 5300 and 5302
  - tcp and udp on dynamic ports (typically 49152-65535)
- If your Serviceguard configuration uses a quorum server, all nodes within the cluster must allow the following communication to the quorum server IP address:
  - tcp on port 1238 - and allow only packets with the SYN flag

  Any node providing quorum service for another cluster must allow the following communication from that cluster's nodes:
  - tcp on port 1238 - and allow only packets with the SYN flag
- Running the `cmscancl` command requires the "shell" port be open.

There are additional firewall considerations to enable execution of Serviceguard commands from nodes outside the cluster, such as those listed in `cmclnodelist`. To allow execution of Serviceguard commands, follow the guidelines below.

All nodes in the cluster must allow the following communications:

- from the remote nodes:
  - tcp on ports 5302 - and allow only packets with the SYN flag
  - udp on port 5302
- to the remote nodes:
  - tcp and udp on port numbers 49152-65535

The remote nodes must allow the following communications:

- from the cluster nodes
  - tcp and udp on port numbers 49152-65535

- to the cluster nodes
  - — tcp on ports 5302 - and allow only packets with the SYN flag
  - — udp on port 5302

Cluster Object Manager (COM) nodes

If you are using a Cluster Object Manager (COM) on a node outside of the cluster to provide connections to Serviceguard Manager or Continental Clusters clients, follow these rules.

Each node in the cluster must allow the following communications:

- from the COM node
  - — tcp on ports 5302 - and allow only packets with the SYN flag
  - — udp on port 5302
- to the COM node
  - — tcp and udp on port numbers 49152-65535 from the cluster nodes

The node running the COM must allow the following communications:

- from COM clients, such as Serviceguard Manager and Continental Clusters:
  - — tcp on port 5303 - and allow only packets with the SYN flag
- from the cluster nodes:
  - — tcp and udp on port numbers 49152-65535
- to the cluster nodes:
  - — tcp on ports 5302 - and allow only packets with the SYN flag
  - — udp on port 5302

Installing Software on HP-UX 11i v1

Following are the directions for installing Servcieguard A.11.16 on HP-UX 11i v2 (`uname -a = 11.11`). For instructions on installing on HP-UX 11i v2, see page 29.

To install your software, run the SD-UX `swinstall` command. It will invoke a user interface that will lead you through the installation. Separate installation of Serviceguard is not necessary if you are installing the Mission Critical Operating Environment.

After updating to HP-UX 11i v1, you may find that kernel memory is low due to post 11i v1 software configuration. This can result in insufficient memory to run Serviceguard commands, such as `cmrunnode`, after Serviceguard A.11.16 is installed. To avoid problems, reboot the node after the update has completed, then issue the `cmrunnode` command.

NOTE The EMS-CORE fileset is no longer a component of Serviceguard, but it is a dependency, which means that it must be present on the installation media or depot when Serviceguard is installed. If you create your own depots as a part of your installation process, be sure to include the EMS-CORE fileset, which is part of the Event Monitoring Services product (B7609BA).

Installation steps should be performed in the following order:

1. You need to have installed HP-UX 11.11 before loading Serviceguard Version A.11.16.
2. Install your Serviceguard product — version A.11.16. The filesets that make up the Serviceguard product are:
   - Serviceguard.CM-SG
   - Cluster-Monitor.CM-CORE
   - Cluster-Monitor.CM-CORE-MAN
   - Package-Manager.CM-PKG
   - Package-Manager.CM-PKG-MAN
   - ATS-CORE.ATS-MAN

- ATS-CORE.ATS-RUN

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/lib/Q4/cmfr.pl`
- `/var/adm/cmcluster/frdump.cmcld.`*x* (where *x* is a digit)

The Cluster Object Manager product (B8324BA) is installed along with Serviceguard. The filesets for this product are:

- Cluster-OM.CM-DEN-MOF
- Cluster-OM.CM-DEN-PROV
- Cluster-OM.CM-OM
- Cluster-OM.CM-OM-AUTH
- Cluster-OM.CM-OM-MAN
- Cluster-OM.CM-OM-TOOLS
- CM-Provider-MOF.CM-MOF
- CM-Provider-MOF.CM-PROVIDER
- OPS-Provider-MOF-OPS-MOF
- OPS-Provider-MOF-OPS-PROVIDER

NOTE

If you did a `swremove` of an older version of Serviceguard before the `swinstall`, your system may be left with a zero-length binary configuration file (/etc/cmcluster/cmclconfig). This file should be removed 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 with the following text:

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

These messages may safely be ignored.

Problems Installing EMS Software

- *What is the problem?* If you have already installed EMS software with a later version number than the version you are trying to install, you may see errors similar to the following in the swinstall logfile, `/var/adm/sw/swagent.log`:

```
ERROR       A later revision (one with a higher
            revision number) of fileset
            EMS-Core.EMS-CORE,r=A.04.00.01
            has already been installed.
            Either remove this fileset
            or change the allow_downdate option to
            true.
```

- *What is the workaround?* If this occurs, choose to ignore the error messages, change the "Enforce Dependency Analysis Errors" option to False, and proceed with the installation. Serviceguard can safely use the later EMS version.

NOTE Do *not* use `-x allow_downdate=true` as suggested by the error message. This would over-write filesets installed on the system with older filesets from the install media.

— To change the `Enforce Dependency Analysis Errors` through the command line, you can invoke Software Distributor as follows:

```
# swinstall -x enforce_dependencies=false
```

— To change the `Enforce Dependency Analysis Errors` through the `swinstall` GUI, use the following procedure:

1. Cancel out of the Analysis.

2. From the `Software` menu, select `Options`, then `Change`. Click `OK`.
3. Scroll to view and un-click the box for `Enforce Dependency Analysis Errors in Agent`. Click `OK`.
4. From the `Action` menu, select `Install`. The Status returns:

   ```
   Ready with Errors. Products scheduled: less
   than the full set. Excluded: older version of
   EMS
   ```

5. Click `OK`. Install begins.

You will see the following error messages during the execution phase:

```
Summary of Analysis Phase:

ERROR             Skipped (in analysis)
                  EMS-Config.EMS-GUI, r=A.04.00

ERROR             Skipped (in analysis)
                  EMS-Core.EMS-CORE, r=A.04.00

ERROR             2 of 5 filesets had Errors.

3 of 5 filesets had no Errors or Warnings:

ERROR          The Execution Phase had errors. See
               the above output for details.
```

NOTE Do not use the `-x enforce_dependencies=false` option under normal circumstances. Since the option applies to all filesets with this analysis error, it would potentially prevent other required filesets from being installed.

For more information about installation procedures and related issues, refer to the following manuals:

- *Managing HP-UX Software with SD-UX* (B2355-90154)
- *swinstall* (1M) in the *HP-UX Reference* (B2355-90166)

De-Installing Serviceguard

To deinstall your 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.
3. The `swremove` command should be issued from one system at a time. That is, if Serviceguard is being deinstalled from more than one system, it should be removed from one system at a time.

   If your system is left with a zero-length binary configuration file (/etc/cmcluster/cmclconfig), you should remove it.
4. Before doing an `swremove` of both the Serviceguard and the Event Monitoring Service (EMS) products, make sure that all user-written resource monitors are terminated.

NOTE

When you remove Serviceguard from a system, the Event Monitoring Service (EMS) is not automatically removed. If you intend to re-install Serviceguard at a later time, or if you are installing a different version of Serviceguard, you should check if this version of Serviceguard requires a newer version of EMS, and if so, remove the existing EMS version separately before re-installing Serviceguard.

Installing Serviceguard on HP-UX 11i v2

This section has directions for installing Serviceguard A.11.16 on HP-UX 11i v2 (`uname -a` = `11.23`). To install on HP-UX 11i v1, see page 24

After you install HP-UX 11i version 2, use the `swinstall` command to install Serviceguard, product number T1905BA. After installation, you can use the following command to display a list of all installed components:

```
# swlist -R T1905BA
```

NOTE The EMS-CORE fileset is no longer a component of Serviceguard, but it is a dependency, which means that it must be present on the installation media or depot when Serviceguard is installed. If you create your own depots as a part of your installation process, be sure to include the EMS-CORE fileset, which is part of the Event Monitoring Services product (B7609BA).

Installation steps should be performed in the following order:

1. You need to have installed HP-UX 11i v2 (11.23) before loading Serviceguard Version A.11.16.
2. Install your Serviceguard product — version A.11.16. The filesets that make up the Serviceguard product are:
   - Serviceguard.CM-SG
   - Cluster-Monitor.CM-CORE
   - Cluster-Monitor.CM-CORE-MAN
   - Cluster-Monitor.CM-CORE-COM
   - Package-Manager.CM-PKG
   - Package-Manager.CM-PKG-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/lib/Q4/cmfr.pl`
- `/var/adm/cmcluster/frdump.cmcld.`*x* (where *x* is a digit)

The Cluster Object Manager (COM) product (B8324BA) is installed along with Serviceguard. The September 2004 update to HP-UX 11i v2 installs the COM version A.03.00.01. Earlier versions install COM version A.03.00. If one node in a cluster has version A.03.00.01, all nodes should have it. You can install COM only using the disk from the 11i v2 September 2004 release; after issuing the `swinstall` command, select COM B8324BA.

The filesets for this product are:

- 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
- Cluster-OM.CM-OM-TOOLS
- CM-Provider-MOF.CM-MOF
- CM-Provider-MOF.CM-PROVIDER
- OPS-Provider-MOF.OPS-MOF
- OPS-Provider-MOF.OPS-PROVIDER

Serviceguard will also install one file into the EMS fileset:

- EMS-CORE-COM

NOTE If you did a `swremove` of an older version of Serviceguard before the `swinstall`, your system may be left with a zero-length binary configuration file (/etc/cmcluster/cmclconfig). This file should be removed 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 with the following text:

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

These messages may safely be ignored.

De-Installing Serviceguard

To deinstall your 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.
3. The `swremove` command should be issued from one system at a time. That is, if Serviceguard is being deinstalled from more than one system, it should be removed from one system at a time.

   If your system is left with a zero-length binary configuration file (/etc/cmcluster/cmclconfig), you should remove it.
4. Before doing an `swremove` of both the Serviceguard and the Event Monitoring Service (EMS) products, make sure that all user-written resource monitors are terminated.

NOTE When you remove Serviceguard from a system, the Event Monitoring Service (EMS) is not automatically removed. If you intend to re-install Serviceguard at a later time, or if you are installing a

different version of Serviceguard, you should check if this version of Serviceguard requires a newer version of EMS, and if so, remove the existing EMS version separately before re-installing Serviceguard.

If You Are Upgrading from Earlier Releases

NOTE A cold install of the operating system during a rolling upgrade is *not supported*. This is because the Serviceguard conversion tools cannot be guaranteed to match the same network PPA numbers that were in use before the upgrade. Also, disk device file names can be changed by a cold install, causing mismatches between the Serviceguard configuration and the system configuration.

Rolling upgrade to A.11.16 is supported from all the following:

Table 1-2 Rolling Upgrade Paths

Serviceguard Release	HP-UX Release
Serviceguard A.10.12	HP-UX 10.20
Serviceguard A.11.01	HP-UX 11.00
Serviceguard A.11.03	HP-UX 11.00
Serviceguard A.11.04	HP-UX 11.00
Serviceguard A.11.05	HP-UX 11.00
Serviceguard A.11.07	HP-UX 11.00
Serviceguard A.11.08	HP-UX 11.00
Serviceguard A.11.09	HP-UX 11.00, HP-UX 11.11
Serviceguard A.11.12	HP-UX 11.00, HP-UX 11.11
Serviceguard A.11.13	HP-UX 11.00, HP-UX 11.11
Serviceguard A.11.14	HP-UX 11.0, HP-UX 11.11

Table 1-2 Rolling Upgrade Paths (Continued)

Serviceguard Release	HP-UX Release
Serviceguard A.11.15	Serviceguard A.11.15 is supported on HP-UX 11i v1 (11.11) It is supported on HP-UX 11i v2 (11.23) original release. It is supported on 11iv2 September 2004 update on Integrity servers, but not on HP 9000 servers.

Patches and Fixes in this Version

The contents of Serviceguard releases A.11.01 through A.11.15 have been incorporated into A.11.16. This section describes patches that are required and defects that have been fixed in version A.11.16 of Serviceguard.

Required and Recommended Patches for HP-UX 11i v2

The following table lists patches required or recommended for Serviceguard A.11.16 on HP-UX 11i v2 (11.23). This list is subject to change without notice. Contact your HP support representative for up-to-the-moment information. Patches can be superseded or withdrawn at any time, so always be sure to check the status of any patch before downloading it.

Table 1-3 Patches for HP-UX 11i v2 (11.23)

Patch Number	Description
PHCO_29605	Required for clusters using VxVM. Not required for 11i v2 September 2004 release. s700_800 11.23 VxVM 3.5-IA.004 Command Patch 01
PHKL_30686	Required for clusters using MirrorDisk/UX mirroring. s700_800 11.23 LVM chooses metadata in incorrect PV
PHSS_30687	Required if using the SG SNMP cluster subagent. s700_800 11.X OV EMANATE15.3 Agent Consolidated Patch
PHSS_30688	Required if using the Serviceguard SNMP cluster subagent. s700_800 11.23 OV EMANATE 15.3 Agent Consolidated Patch
PHSS_30966	s700_800 11.11 ld(1) and linker tools cumulative patch

Required and Recommended Patches for HP-UX 11i v1

The following table lists patches required or recommended for Serviceguard A.11.16 on HP-UX 11i v1 (11.11). This list is subject to change without notice. Contact your HP support representative for up-to-the-moment information. Patches can be superseded or withdrawn at any time, so always be sure to check the status of any patch before downloading it.

Table 1-4 Patches for HP-UX 11i v1 (11.11)

Patch	Description
PHCO_27957	Required if using the parallelized package control script functionality. (Note that this patch has dependencies on a number of other patches.) s700_800 11.11 umount(1M) cumulative patch, Dev IDs Enab.
PHCO_28693	Required if VxVM 3.5 is used s700_800 11.11 VERITAS VM Mgmt Service Provider Patch
PHCO_29379	Required if HP Virtual Array is used. s700_800 11.11 LVM commands cumulative patch
PHCO__30834	Required if VxVM 3.5 is used. s700_800 11.11 VxVM 3.5m Command Cumulative Patch 05
PHKL_ 26719	s700_800 11.11 KEPD_PRINTF ioctl fails from 32-bit process
PHKL_ 27532	s700_800 11.11 Serviceguard/vsar incompatibility removed
PHKL_ 27727	This is a required patch for all Serviceguard clusterss700_800 11.11 Fix clock sync for SD, fix negative useconds
PHKL_28114	This is a required patch for all Serviceguard clusters s700_800 11.11 timeout; Serviceguard TOC
PHKL_28695	This is a required patch for all Serviceguard clusters. s700_800 11.11 Cumulative VM, Psets, Preemption, PRM, MRG
PHKL_ 28984	Required if using Fibre Channel Mass Storage s700-800 11.11 Fibre Channel Mass Storage Patch

Table 1-4 Patches for HP-UX 11i v1 (11.11) (Continued)

Patch	Description
PHKL_29527	This is a required patch for all Serviceguard clusters s700_800 11.11 filesystem buffer cache performance fix
PHKL_ 29704	Required if using PRM or WLM s700_800 11.11 Psets Enablement, SCHED_NOAGE, FSS
PHKL_29981	Required if VxVM 3.5 is used. s700_800 11.11 VxVM 3.5m Kernel Cumulative Patch 05
PHKL_30032	Required if using PRM or WLM s700_800 11.11 detach, NOSTOP, Abort; Psets; slpql perf; FSS
PHKL_30033	Required if using PRM or WLM s700_800 11.11 Core PM, vPar, Psets Cumulative; slpql; FSS
PHKL_30035	Required if using PRM or WLM s700_800 11.11 Psets Enablement; FSS iCOD; callback; FSS
PHKL_30036	Required if using PRM or WLM s700_800 11.11 FSS cumulative; FSS capping
PHKL_30510	Required if using VxVM s700_800 SCSI IO cumulative patch
PHKL_ 30037	Required if using PRM or WLM, and PSETs are installed s700_800 11.11 RTE Psets Enablement; FSS, pset perf; FSS cap
PHKL_30511	Required if using Ultra160 SCSI Host bus adaptors for application data on shared SCSI buses s700_800 11.11 crash, vpars, timeout; SG TOC, nPar Config
PHKL_30512	This is a required patch for all Serviceguard clusters. s700_800 11.11 crash, vpars, timeout; SG TOC, nPar Config
PHKL_30622	Required if using MirrorDisk/UX s700_800 11.11 LVM cumulative patch; 16 node SLVM cluster

Table 1-4 Patches for HP-UX 11i v1 (11.11) (Continued)

Patch	Description
PHKL__30833	Required if using JFS 3.3 (VxFS). s700_800 11.11 dmapi; fsadm; ACL; locking order
PHNE_24384	s700_800 11.11 gated (1M) patch
PHNE_28328	This is a required patch for all Serviceguard clusters. s700_800 11.11 inetd (1M) cumulative patch
PHNE_28778	Required for Auto-Port Aggregation s700_800 11.11.[00-07] Auto-Port Aggregation cumulative patch
PHNE_28799	Required for all 100BaseT LAN cards, including multiport cards and built-in LAN ports. s700_800 11.11100BT unified driver cumulative patch
PHSS_28880	Required for all Serviceguard A.11.16 nodes.on 11i v1 s700_800 11.11 HP aC++ -AA runtime libraries (aCC A.03.50)
PHNE_29887	Required if using VLAN functionality s700_800 11.11 cumulative ARPA Transport patch
PHNE_29945	Required if using IGELAN 1000Base-SX/T driver versions prior to B11.11.07 s700_800 11.11 IGELAN 1000Base-SX/T B.11.11[02-11] patch
PHSS_28509	Required if using A51158A or A6795A Tachhyon TL Fibre Channel cards s700_800 11.11 Tachyon TL Fibre Channel Driver Patch
PHSS_31015	Use only for Serviceguard Version A.11.14 or Serviceguard OPS Edition A.11.14. This is a required patch.
PHSS_26056	Introduces the ability to configure up to 150 packages in a cluster. s700-800 11.x MC/Serviceguard and SG-OPS Edition A.11.14

Fixed in This Version

The following defects have been fixed in Serviceguard A.11.16:

JAGaf13778 cmcld SIGSEV on systems with serial heartbeat

- *What was the problem?* In 2 node Serviceguard cluster using a serial heartbeat link, in one code path an uninitialised pointer is referenced which can result in a cmcld abort. The code path is only executed if a serial heartbeat is configured.
- *What was the resolution?* Removed uninitialised pointer reference and use another variable which serves the same purpose

JAGaf12369 cmcld stuck on partially opened cluster lock device

- *What was the problem?* The `cmcld` opens a physical link before cluster lock acquisition so that bus reset can be done to clear any pending I/0. For a Fiber Channel cluster lock disk, this open returns successfully but the disk can only get partially open (its LUN size is 0) and all subsequent tries of `cmcld` to access the disk would fail. To fix this problem at HP-UX level the device needs to be closed by any process.
- *What was the resolution?* For Fiber Channel Storage, bus reset are not supported. Therefore, `cmcld` does not open Fiber Channel cluster lock devices anymore.

JAGaf05904 Serviceguard probing do not recognize virtual bus interfaces

- *What was the problem?* When SG does disk probing during configuration process, it tries to query the I/O interface of the disk and only expects to see type "INTERFACE". When it is another type, "VIRTBUS" in this case, SG tries to query the parent node in the I/O tree, but can’t.
- *What was the resolution?* Made change so that SG recognizes type "VIRTBUS"

JAGaf04449 memory leak in cmlvmd when `vgdisplay` is run

- *What was the problem?* The Serviceguard `cmlvmd` daemon has a memory leak such that every time a vgdisplay of a cluster volume group is executed, cmlvmd malloc’s memory doesn’t free it. The memory leak is a small one, but it nonetheless needs to be fixed. When vgdisplay -v <VG> is issued, `cmlvmd`’s clv_local_get_vgdisplay_info() is called. Here `cmlvmd` creates some

memory, fills in the required information and sends it back to the command. Later it is supposed to free up the memory that was created. In this case, freeing of the memory is not done.

- *What was the resolution?* The fix is to free up the memory, if it was created successfully.

JAGaf03966 2 node cluster with no QS or quorum server:

- *What is the problem?* `cmapplyconf` does not properly check for adding cluster lock while the cluster is running.
- *What is the workaround?* Modified the `cmapplyconf` code to properly check and not allow addition of a cluster lock while the cluster is running.

JAGaf03957 Control script executing cmhaltpkg command fails

- *What was the problem?* In this case the `cmhaltpkg` command was in progress while waiting for the package control script to finish. If reconfiguration happens during that time, the command does not know if the halt of the control script was successful or not. The retry to learn the status might have found that the package is not running even though it might have halted successfully. So the command exits with an error.
- *What was the resolution?* As this is a race condition and multiple things are happening at a particular time, the error cannot be avoided in all cases. So a clearer message is displayed if this condition is encountered.

JAGaf01415: SG 11.15 `cmclconfd` reduces RLIMIT_NOFILE, preventing application startup

- *What was the problem?* The `cmclconfd` and `cmcld` daemon recalculates the number of file descriptors for its own environment and sets its value in the system (currently 1024 set by `cmclconfd`). The same value gets passed to `cmcld` and from `cmcld` to the applications starting from Serviceguard, thus creating a problem.

  (See JAGae94512: SG 11.15 reduces the RLIMIT_NOFILE causing applications to fail to start)

- *What was the resolution?* Even if cmclconfd & cmcld daemon recalculates and sets new file descriptors limit, the original value will be restored for child processes so that they will have the original system limit.

JAGae98584 local lan failover problems when a single network switch fails

- *What was the problem?* Because of a logic error and a time-sensitive issue, only stationary IP address(es) were switched to the second standby interface while the relocatable IP address(es) remained in the primary interface after the shared network switch was powered down.
- *What was the resolution?* The logic error has been fixed so that Serviceguard will allow all IP address(es) to switch to the second standby properly.

JAGae94575 AUTO_START_TIMEOUT ignored by cmcluster RC script

- *What was the problem?* AUTO_START_TIMEOUT ignored by cmcluster RC script. The `/sbin/init.d/cmcluster` script has a hard-coded timeout of 600 seconds.
- *What was the resolution?* The script was modified to get timeout value from CDB via cmviewconf. The AUTO_START_TIMEOUT will be honored even when a node is powered down.

JAGae94512: SG 11.15 reduces the RLIMIT_NOFILE causing applications to fail to start

- *What was the problem?* The `cmclconfd` and `cmcld` daemon recalculates the number of file descriptors for its own environment and sets its value in the system (currently 1024 set by `cmclconfd`). The same value gets passed to `cmcld` and from `cmcld` to the applications starting from Serviceguard, thus creating a problem.

  (See JAGaf01415: SG 11.15 cmclconfd reduces RLIMIT_NOFILE preventing application startup)
- *What was the resolution?* Even if `cmclconfd` & `cmcld` daemon recalculates and sets new file descriptors limit, the original value will be restored for child processes so that they will have the original system limit.

JAGae91702 - CONCURRENT_VGCHANGE_ OPERATIONS is not useful

- *What was the problem?* There is a perception that the CONCURRENT_VGCHANGE_OPERATIONS functionality is not useful in some situations. However, when package failover test was run with 40 volume groups on an 8 processor system, a significant improvement was found in the time it takes package to failover. So more explanation is needed on how/when to use this option to gain benefit.
- *What was the resolution?* More text added to explain the benefits of concurrent operations during package startup/shutdown. Package control script created through `cmmakepkg` will have additional comments.

JAGae91411 if subnet mask changes, online cmapplyconf does not detect this

- *What was the problem?* If customers changed netmask of a subnet on all cluster nodes where the subnet is present, the `cmcheckconf` and `cmapplyconf` commands do not detect it. The commands succeed without updating the CDB.

  Thus `cmviewconf` shows a different netmask vs. the new one in the system configuration. In the `cmcheckconf/cmapplyconf` code, we only read in LAN interface name and IP address from the ASCII file and let the config daemon verify these two types of data. So, even if the netmask have been changed on all nodes of the cluster, we would not detect it because we have no existing data to compare against.
- *What was the resolution?* Since we do not support online networking change at this time, the commands should not proceed if a netmask on all nodes has been changed.

  If the netmasks are mismatched while the cluster is running, the command will report an error. Otherwise, the commands will succeed and change the netmask in the CDB. The fix is to copy the old netmasks from the existing configuration (from the *`old_cl`* in `cf_read_cluster_ascii()` function) to the *`new_cl`* after data is read from the ASCII file, and use to them to compare with the new netmasks obtained from the system so that we can detect the differences during `gather_network_config` phase of the commands.

The `cmcheckconf` and `cmapplyconf` have been enhanced to not allow verification of a cluster if at least one of the netmasks in the cluster has been changed in the system. This is to prevent users from changing their network masks, which is an unsupported configuration change, after the cluster has been up and running.

JAGae91402 Cmscancl 'network connection checking' is wrong

- *What was the problem?* The PPA values of the HyperFabric interfaces are the same as ethernet network interfaces thus causing the linkloop command used in `cmscancl` to give incorrect results. Clic0(HyperFabric interface) has PPA 0, same as lan0 (Ethernet interface). Linkloop uses the card PPA number, and so can't really distinguish between cards.
- *What was the resolution?* Skip the network connectivity check for non-LAN hardware (i.e. HyperFabric, ATM etc.), if any, since linkloop command is supported only for LAN hardware.

JAGae91372 WARNING statements from config cmds cause undue concern

- *What was the problem?* Serviceguard configuration commands (cmquerycl, cmcheckconf, cmapplyconf and cmgetconf) all generate the following sort of messages when they detect a disk which does not belong to an LVM volume group or a VxVM disk group:

  ```
  Warning: Disks which do not have IDs cannot be included
  in the topology description. Use pvcreate(1M) to
  initialize a disk for LVM or, use vxdiskadm(1M) to
  initialize a disk for VxVM.
  ```

  and

  ```
  Warning: The disk at /dev/dsk/c1t2d0 on node eon does not
  have an ID, or a disk label.
  ```

- *What was the resolution?* (1) Removed the frequent message: `Unable to determine a unique identifier for physical volume %s on node %s. Use pvcreate to give the disk an identifier.`

  (2) Reworded this one-time message:

From: `Disks which do not have IDs cannot be included in the topology description. Use pvcreate(1M) to initialize a disk for LVM or, use vxdiskadm(1M) to initialize a disk for VxVM.`

To: `Disks were discovered which are not in use by either LVM or VxVM. Use pvcreate(1M) to initialize a disk for LVM or, use vxdiskadm(1M) to initialize a disk for VxVM.`

JAGae80291 Wrong order in vxdg deport for cluster node shutdown

- *What was the problem?* Wrong order in vxdg deport for cluster node shutdown in `/sbin/init.d/cmcluster`: It first deports the disk groups and then halts the cluster nodes.

  Since packages might still be running and file systems might still be mounted when a shutdown command is issued, the deport commands will fail. During shutdown, the `cmcluster.init` deports vxvm disk group before halting the node. Deport can fail sometimes because the package using it might be still running and file systems still mounted. Sometimes, this can impact shutdown.

- *What was the resolution?* Fix is to move vxdg deport to after `cmhaltnode` in `halt_cluster_node`.

JAGae69654servicefail fast occurs even on single node clusters. - better error msg

- *What was the problem?* This is not a defect but a designed behavior.

  Since VxVM-CVM-pkg uses FAILFAST option for both service and node, once it fails, the node will TOC.

- *What was the resolution?* Logging messages are added to notify customer when FAILFAST is enabled and that node can TOC when package fails, even if it's the only node in the cluster.

JAGae64683: System Multi-Node (SMN) packages aren't handled correctly by `cmsnmpd`

- *What was the problem?* The `cmsnmpd` does not send traps when a SMN package is started, because the `owner=0` in the Serviceguard event is considered unowned. When a SMN goes down, the

coordinator node does not send an SMN down trap. The SNM package current node MIB variable is also incorrect in the SNMP MIB table.

- *What was the resolution?* Enhancements allow the subagent and its API to identify SMN packages and treat them differently than regular packages.

JAGae61889: cmGetsatus for packages returns -10 after online reconfiguration.

- *What is the problem?* When packages are added or deleted concurrently, cmsnmpd is trying to retrieve a status of a package that has been deleted. This logfile output is seen:

```
cmGetstatus (CM_PDKG_STATUS) returns: -10 (errno=2)
***Error: retrieveing package statuses: -10
```

- *What is the resolution?* Added logging returns the package name not found, and logs the errorno from cln_get_health().

JAGae61376 config commands (`cmcheckconf` and `cmapplyconf`) are slow when not using the -k

- *What was the problem?* This is a follow-on defect to previous Enhancement Requests submitted concerning the slow speed of the SG configuration commands: `cmquerycl, cmcheckconf and cmapplyconf`. These enhancements had been added to SG A.11.15 by adding a new option -K to the configuration commands to bypass all disks that are not part of the cluster lock. Using this new option makes the configuration commands much faster. However, a second part of the enhancements had not been provided in SG A.11.15.

  This is to make the configuration commands faster when doing full disk probing. As described in the previous Enhancement Requests, there are two major areas where the ability to probe the disks in parallel and streamline the opening and closing of the Volume Groups to read the cluster attributes.

- *What was the resolution?* Config library now queries a few disks at a time and can connect to multiple cmclconfds on a node to improve performance. Changes made in config library and cmclconfd to improve feedback to the user during disk data gathering, to remove redundant opens and closes of devices, and to have multiple cmclconfd processes gathering data at the same time.

JAGae60038: cmapplyconf prints bogus message when cl_disk_init fails

- *What is the problem?* A misleading error message is printed when trying to apply a configuration that needs to initialize the cluster lock disk when the cluster lock vg is not activated. The message is:

  ```
  Protocol failure talking with cmclconfd on grcdg319 (1) :
  Error 0
  ```

- *What is the resolution?* Corrected code uses the correct variable when logging the error message.

JAGae57819: `cmviewcl` displays incorrect status for package during failfast event.

- *What is the problem?* `cmviewcl` displays incorrect status for package during failfast event.
- *What is the resolution?* The command changed to display a package as unknown if the owning node is down or unknown. Also, config lib, used by the cmprovider, modified to return correct package status in this case.

JAGae49363: `vgchange` op. fails after certain operations related to online node reconfig

- *What is the problem?* After doing online node reconfiguration involving more than one node, `vgchange -a s` fails, saying lvm is busy.
- *What is the resolution?* The counter increments just once, no matter how many online adds or deletes are being done.

Known Problems and Workarounds

JAGaf36760: Uncorrectable write errors on vxvm volumes.

- *What is the problem?* Configurations using greater than 256 LUNs with VxVM on 11i v2 September 2004 update experience uncorrectable write errors.
- *What is the workaround?* Restrict configurations to a maximum of 256 LUNs. Please continue to check the status of this problem at http://www2.itrc.hp.com for updated information. Search the Technical Knowledge base for bug reports with the keyword JAGaf36760.

JAGaf36352: Possible incomplete package configuration.

- *What is the problem?* In some rare situations, when using the Serviceguard Manager interface to do package configuration with Serviceguard A.11.16, you may see a message saying the configuration failed but the operation log indicates that cmapplyconf succeeded. This is due to a race condition in the shell script that SGmgr uses.
- *What is the workaround?* User should repeat the configuration operation. If the user is still in the package configuration screens in Serviceguard Manager, clicking on the apply button again should cause the package to be created or modified correctly.

JAGaf36081: Long boot times.

- *What is the problem?* Configurations using greater than 256 LUNs with CVM on 11i v2 September 2004 update experience boot times of greater than 1 hour.

  September 2004 update experience boot times of greater than 1 hour.

- *What is the workaround?* Restrict configurations to 256 LUNs or less. Please continue to check the status of this problem at http://www2.itrc.hp.com for updated information. Search the Technical Knowledge base for bug reports with the keyword JAGaf36081.

JAGaf32447: Windows: Avoid security risk posed by JRE bundle installed with Serviceguard Manager A.04.00.

- *What is the problem?* A Denial of Service risk is described in a Security Bulletin issued for JRE 1.4.2.00. This is the JRE bundled with, and automatically installed with, Serviceguard Manager A.04.00.

  For details, please reference the HP ITRC database security bulletin: SST4749 rev.0 HP=UX Java Runtime Environment (JRE) remote DoS, and `http://sunsolve.sun.com` -> Search: DoS JRE or Search: 57555.

- *What is the workaround?* For instructions on replacing the JRE on Windows, search the technical knowledge base for keyword = JAGaf32447 at your support site:

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

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

JAGaf32443: HPUX: Avoid security risk posed by JRE bundle installed with Serviceguard Manager A.04.00

- *What is the problem?* A Denial of Service risk is described in a Security Bulletin issued for JRE 1.4.2.00. This is the JRE bundled with, and automatically installed with, Serviceguard Manager A.04.00.

  For details, please reference the ITRC database security bulletin: SST4749 rev.0 HP=UX Java Runtime Environment (JRE) remote DoS, and `http://sunsolve.sun.com` -> Search: DoS JRE or Search: 57555.

- *What is the workaround?* For instructions on replacing the JRE on HP-UX, search the technical knowledge base for keyword = JAGaf32443 at your support site:

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

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

JAGaf32449: Linux: Avoid security risk posed by JRE bundle installed with Serviceguard Manager A.04.00

- *What is the problem?* A Denial of Service risk is described in a Security Bulletin issued for JRE 1.4.2.00. This is the JRE bundled with, and automatically installed with, Serviceguard Manager A.04.00.

  For details, please reference the ITRC database security bulletin: SST4749 rev.0 HP=UX Java Runtime Environment (JRE) remote DoS, and `http://sunsolve.sun.com` -> Search: DoS JRE or Search: 57555.

- *What is the workaround?* For instructions on replacing the JRE on Linux search the technical knowledge base for keyword = JAGaf32449 at your support site:

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

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

JAGaf31895: Cluster configuration: Standby interfaces will be unconfigured by Serviceguard Manager interface for more than one standby interface.

- *What is the problem?* When using Serviceguard Manager to create or modify cluster configuration, if multiple standby interfaces are specified, Serviceguard Manager will only take the first one. The remaining interfaces will be ignored.
- *What is the workaround?* If multiple standby interfaces are required, please use the Serviceguard commands to perform the cluster configuration opertion.

JAGaf24444: Unable to receive device query message

- *What is the problem?* Serviceguard configuration commands (cmquerycl, etc) may fail if you have no LVM volume groups configured on one or more of your systems.
- *What is the workaround?* The work around is to create a volume group and import it on all nodes that have no LVM volume groups. There is a patch for Serviceguard A.11.14: PHSS_31015. Patch numbers for A.11.15 and A.11.16 will be posted when testing is complete.

JAGaf32484: 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 the node to be TOC'ed.

  This problem exists regardless of whether the HP-UX operating system revision was changed during the rolling upgrade. By "rolling upgrade", we mean 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 `cmruncl` 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 4 node cluster, you could take down 2 nodes in the cluster and upgrade CVM to 3.5, while the other 2 nodes continue running the cluster with CVM 3.2. Then halt the entire cluster, and start up the cluster on the 2 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.

JAGaf91413: cmmodnet does not check SUBNET mask when adding an IP

- *What is the problem?* If the subnet mask of a network interface has been changed from the value originally used when the cluster configuration was last applied, any package relocatable IP addresses added to that network interface (when a package starts) will still use the original subnet mask, resulting in an inconsistent network configuration
- *What is the workaround?* Check error messages for `cmapplyconf`. Error message will prompt the user to fix the subnet mask problem

JAGaf08686: It is not possible to configure some combinations of roles

- *What was the problem?* Duplicate roles and conflicting roles are not allowed in Access Control Policies. This is especially problematic when wild cards are used. For example, if ANY_USER from ANY_SERVICEGUARD_NODE has a role, no other Access Control Policy can be created that would not conflict or be redundant. Every user possible user already has a role.

  But, what if you want everyone in the lab to have monitor access, one smaller group to have package admin, and the manager, Jon, to have full cluster admin. Until you remove the double wild card, you cannot define another role.

- *What was the workaround?* Avoid broadly defined policies, especially those with wildcards for both users and nodes. Instead define roles for groups and individuals, or specify only certain nodes. For example, the following policies have no conflicts or redundancies:

  — You can create an /etc/passwd entry for a user ITlab, and give everyone in the lab the passwords to log in as ITlab

    — USER_NAME ITlab

    — USER_HOST ANY_SERVICEGUARD_NODE

    — USER_ROLE monitor

  — For a smaller group or an individual, you can list individual names:

    — USER_NAME *admin1 admin2 admin3 admin4 admin5 admin6 admin7 admin8*

    — USER_HOST CLUSTER_MEMBER_NODE

    — USER_ROLE package_admin

    — USER_NAME jon

    — USER_HOST ANY_SERVICEGUARD_NODE

    — USER_ROLE full_admin

JAGae87101: LOG_PERIODIC messages should not be logged at log level 1

- *What is the problem?* Because periodic messages show up very often, enabling this sort of debug logging will fill up `syslog.log` very quickly.
- *What is the workaround?* There is no workaround.

JAGae62205: Serviceguard package cannot be restarted after hardware monitoring is re-enabled.

- *What is the problem?* When hardware monitors are disabled using the `monconfig` tool, all associated hardware monitor requests are removed from the persistence files; when hardware monitoring is re-enabled, these requests are re-created. However, hardware monitor requests that were created using SAM, or established when Serviceguard is started, are related to the `psmmon` hardware monitor; these requests are not re-created.

  Here is one example:

  1. A package is configured in SAM with a hardware monitor request. The cluster starts, Serviceguard creates a request, and the package starts.
  2. Later, an administrator stops hardware monitoring, using `monconfig -K`. All monitors stop, and all related requests are removed. The package halts, to move to another node.
  3. The administrator re-enables the hardware monitors using `monconfig -E`. However, the requests that were created outside of `monconfig` are not re-created. The package can no longer start, because the monitor request is lost.

- *What is the workaround?* If you have created a hardware monitor request without using the `monconfig` tool and then you stop the monitors, you must re-start the requests yourself. Here are two different ways to do that.

  — While the cluster is running and after hardware monitoring has been re-enabled, manually re-add the monitor requests using the `cmstartres` command. Then start the package.

— Halt Serviceguard on the node, then restart it. This will create the monitor request. Then start the package.

JAGad39695 User error can result in "ghost" services:

- *What is the problem?* When the package was shutdown, references for this service had already been removed, and the service failed while the package was halting. Since the service no longer exists, it cannot be halted manually. Any attempt to do so results in this error:

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

  Due to the failure of the service, `cmsrvassistd` tries to re-start the service since it was never halted. However, when `cmsrvassistd` goes to register the re-started service, it's status entry has been removed by cmcld because of the package shutdown. When the customer issues the `cmapplyconf`, this service is permanently removed from `cmcld` although `cmsrvassistd` is still trying to re-start it. Therefore after the `cmapplyconf` the "ghost" service still cannot be stopped because cmcld 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.

  Here are two ways to fix the problem, however, if you do get into this solution:

  — 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 24x7 environments.

SR 8606185685 (JAGad54887): `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.

Software Availability in Native Languages

The command line interface for Serviceguard Version A.11.16 does *not* provide Native Language Support. However, separate native language versions of documentation are available as a part of products B3935DA and B3936EA with the following options:

- ABA: English B3935-90078
- ABJ: Japanese B3935-90079
- AB1: Korean B3935-90080
- AB2: Simplified Chinese B3935-90081
- AB0: Traditional Chinese B3935-90082

Serviceguard Manager, the graphical user interface, is translated into all of these languages.