

Upgrading to HP Serviceguard 11.17 with Cluster Volume Manager/Cluster File Systems



Abstract.....	2
Overview.....	2
Packaging and installation considerations.....	2
Networking considerations	4
HP Serviceguard/VERITAS interoperability considerations	5
New HP Serviceguard commands.....	6
VERITAS storage considerations.....	7
HP Serviceguard configuration considerations.....	8
Data migration to CVM/CFS.....	9
Other considerations	10
Upgrade roadmap.....	10
Appendix A	12
For more information.....	16

Abstract

HP Serviceguard A.11.17 tightly integrates with technology from Symantec to enable Cluster File Systems (CFS). Seven Serviceguard Storage Management bundles, which include Serviceguard and VERITAS software corresponding to different degrees of capability, are available. Network requirements for some of these configurations are somewhat different than with previous versions of Serviceguard. Several new cluster and package configuration options are available, and new commands are introduced to enable management of the new capabilities. When the selected new configuration radically differs from existing one, such as when moving data from raw volumes to file systems, the data must be migrated. This paper explores the considerations in upgrading to these environments and provides guidelines for recommended solutions.

Overview

When upgrading to Serviceguard A.11.17 and VERITAS Storage Foundation 4.1, there are multiple areas in which changes might be appropriate. This document addresses each of the following areas of interest:

- Packaging and installation considerations
- Networking considerations
- Serviceguard/VERITAS interoperability considerations
- New Serviceguard commands
- VERITAS storage considerations
- Serviceguard configuration considerations
- Data migration to Cluster Volume Manager (CVM)/CFS
- Other considerations
- Upgrade roadmap

This document is intended to give an overview of the upgrade process and provide guidance for selecting from among the recommended solutions, but it does not provide all the necessary details of the involved processes. Many other documents offer detailed explanations of these processes. References to these documents appear throughout the paper, and a comprehensive list is included at the end.

At the time of the initial release, Serviceguard A.11.17 is only available as part of Serviceguard Storage Management bundles (described in the next section). Application Release (AR) and Operating Environment (OE) releases that include Serviceguard A.11.17 are expected to be available in late 2005. The following sections include material related to non-bundle installations of Serviceguard to assist with future upgrade planning.

Packaging and installation considerations

In addition to the traditional packaging of Serviceguard as a stand-alone product, part of AR, or part of the Mission Critical Operating Environment (MCOE), Serviceguard Storage Management Suite software bundles are available. These include Serviceguard, VERITAS Storage Foundation 4.1, and additional products that provide enhanced storage management and database integration capabilities. Each of these bundles is designed to provide a comprehensive set of software tools suited to a particular application environment.

All bundles that include CVM also provide CFS support. While CVM support is also available by installing a stand-alone version of CVM, CFS support is only available by installing one of the Serviceguard CFS bundles.

Additional software in some of the bundles includes enterprise tools, such as Quality of Storage Service (QoS), FlashSnap, and Checkpoints, and additional database tools, including Quick I/O and Online Data Manager (ODM) for Oracle® and HP Serviceguard Extension for RAC (SGeRAC). The bundle product numbers are as follows:

- **T2771BA** HP Serviceguard Storage Management
- **T2772BA** HP Serviceguard Storage Management Premium
- **T2773BA** HP Serviceguard Storage Management for Oracle
- **T2774BA** HP Serviceguard Storage Management for Oracle Premium
- **T2775BA** HP Serviceguard Cluster File System
- **T2776BA** HP Serviceguard Cluster File System for Oracle
- **T2777BA** HP Serviceguard Cluster File System for RAC

Bundle contents are described in detail in *Serviceguard Storage Management Suite Release Notes*. Following is a summary of the contents.

Bundle	All features from	VxVM	CVM	VxFS	VxCFS	Management utilities: - Quicklog - Flashsnap	DB accelerators: - Quick I/O - ODM
HP Serviceguard Storage Management	Storage Foundation	Yes	No	Yes	No	Yes: Quicklog No: Flashsnap	No
HP Serviceguard Storage Management Premium	Storage Foundation	Yes	No	Yes	No	Yes	No
HP Serviceguard Storage Management for Oracle	Storage Foundation for Oracle	Yes	No	Yes	No	Yes: Quicklog No: Flashsnap	Yes
HP Serviceguard Storage Management for Oracle Premium	Storage Foundation for Oracle	Yes	No	Yes	No	Yes	Yes
HP Serviceguard Cluster File System	Storage Foundation Cluster File System	Yes	Yes	Yes	Yes	Yes	No
HP Serviceguard Cluster File System for Oracle	Storage Foundation Cluster File System (PLUS QuickIO and ODM)	Yes	Yes	Yes	Yes	Yes	Yes

Bundle	All features from	VxVM	CVM	VxFS	VxCFS	Management utilities: - Quicklog - Flashsnap	DB accelerators: - Quick I/O - ODM
HP Serviceguard Cluster File System for RAC	Storage Foundation Cluster File System/ RAC	Yes	Yes	Yes	Yes	Yes	Yes

None of the bundles can be installed at the same time with or on top of stand-alone CVM (either version 3.5 or 4.1). If CVM 3.5 (product number **B9117AA**) is already installed, it must be removed before installing any of the bundles. Non-bundle installations of Serviceguard A.11.17 continue to interoperate with stand-alone CVM 3.5 in the same way as with previous versions of Serviceguard. Stand-alone CVM 4.1 can also be installed with non-bundle distributions of Serviceguard A.11.17 (that is, stand-alone, AR, or MCOE) and is supported without any changes to Serviceguard/CVM capabilities (a small configuration change is required in this case, as described later in this document).

CVM and CFS do not support rolling upgrade. But because CVM and CFS cannot be configured until the upgrade is complete, it is possible to perform a rolling upgrade to any of the bundles as long as all packages that use CVM are removed first and CVM is uninstalled, as described later in this document.

Networking considerations

With CVM 3.5, only one heartbeat network (with standby LANs or Auto Port Aggregation [APA]) is supported by Serviceguard because of restrictions in CVM 3.5. With Serviceguard A.11.17 and CVM 4.1 (when used either only for CVM or for CVM and CFS), at least two Ethernet networks are required. Because APA is not supported for heartbeat with CVM 4.1, one or more of these networks must be configured as heartbeat LANs, and any remaining networks may be configured as standby LANs. Neither Infiniband nor serial connections are supported for heartbeat with CVM. The need for at least two networks could require modifications to the current Serviceguard network configuration, which might require additional network hardware to be installed. The only configuration that is supported with both CVM 3.5 and CVM 4.1 without changes is a single heartbeat network with standby LANs.

The following table summarizes the supported heartbeat configurations for CVM 3.5 and 4.1.

	1 HB LAN	1 HB LAN + standbys	1 HB LAN with APA	>1 HB LAN + optional standbys
CVM 3.5	Not recommended	OK	OK	No
CVM 4.1	No	OK	No	OK

The Infiniband and APA restrictions apply only to heartbeat and only with CVM. Using Infiniband, APA, or both for data and application connectivity (and for heartbeat, when CVM is not used) is fully supported.

HP recommends making the needed networking changes before configuring the new Serviceguard installation. The Serviceguard cluster ASCII configuration file might need to be modified to reflect these changes. Network configuration recommendations and configuration steps are discussed in detail in the *Managing Serviceguard Twelfth Edition* manual.

HP Serviceguard/VERITAS interoperability considerations

With Serviceguard A.11.17, as with previous versions of Serviceguard, Logical Volume Manager (LVM) and VERITAS Volume Manager (VxVM) volumes can coexist on the cluster and can continue to be used in the same packages, if desired. Because cluster lock disks can only be configured on LVM Volume Groups (VGs), this functionality provides an alternative to using a Quorum Server for cluster lock functionality.

Before version A.11.17, Serviceguard supported CVM by running VERITAS daemons in a `SYSTEM_MULTI_NODE` package (SMNP) **VxVM-CVM-pkg**. In A.11.17, Serviceguard continues to use the same method for working with CVM 3.5. When CVM 4.1 or above is installed, Serviceguard closely interacts with CVM in establishing and managing cluster membership. The immediately visible change with CVM 4.1 or higher is that the SMNP has a new name, **SG-CFS-pkg**, and when one of the CFS bundles is installed, a new command, `cfsccluster`, is available to simplify configuring and managing this package.

Before Serviceguard A.11.17, CVM storage was specified with the `STORAGE_GROUP` attribute in the package ASCII configuration file and by configuring script variables `CVM_DG[*]` and `CVM_DISK_ACTIVATION` in the package control script. This usage continues to be supported with Serviceguard A.11.17 for backward compatibility. In addition, with A.11.17, Serviceguard adds a new model of management for CVM/CFS by adding several new features and commands to make storage management both easier and more flexible. This model allows the storage infrastructure to be configured independently of the application package configuration for improved manageability. These new capabilities and commands are available only when one of the CFS bundles is installed.

Cluster-wide disk groups, file systems mount points, checkpoints, and snapshots can be defined in their own Serviceguard packages, which have instances running on multiple nodes; these are called `MULTI_NODE` packages (MNP). MNPs differ from System Multi-Node packages in that the package instance on individual nodes can be stopped and started without affecting cluster membership. The relationships between an application and its storage are defined by declaring in the package ASCII file that an application package **depends** on an MNP. This is done instead of specifying any `STORAGE_GROUP` attributes and `CVM_` script variables. The following is an example from an application package's ASCII configuration file that describes its dependency on an MNP named `SG-CFS-DG-01`:

```
DEPENDENCY_NAME           SG-CFS-DG-01
DEPENDENCY_CONDITION      SG-CFS-DG-01 = UP
DEPENDENCY_LOCATION       SAME_NODE
```

This declares that an instance of MNP named `SG-CFS-DG-01` must be running on a node for the application to start and keep running on that node. If the MNP represents a CVM disk group (as it does in this example, according to the naming convention used), this means that availability of the disk group is enforced by Serviceguard on the same node where the application package is running. An application could declare its dependency on a file system MNP in a similar manner.

Dependencies are also used to establish relationships between file system MNPs (including mount points, checkpoints, and snapshots) and disk group MNPs. A set of new commands, described in the following section, is provided to configure and manage these specialized disk group and file system MNPs. Configuration of MNPs other than with these commands is not supported in Serviceguard A.11.17.

There are no changes in supported storage management capabilities for non-CVM storage with Serviceguard A.11.17. Several data migration scenarios from non-CFS storage to CVM/CFS are explored in later sections of this paper.

New HP Serviceguard commands

Several new Serviceguard commands are available when one of the CFS bundles is installed to help manage the storage infrastructure. Following is an overview of their syntax and functionality. All of these commands are described in detail in their man pages and in *Managing Serviceguard Twelfth Edition*.

The functionality offered by these commands is also accessible through the VERITAS Enterprise Administrator (VEA) tool. VEA offers an intuitive graphical user interface and can simplify many VERITAS and Serviceguard storage-related configuration tasks.

Normally, all the MNPs administered with these commands will be created with the `AUTO_RUN` attribute set to `YES`, so they will be started and stopped automatically when nodes join and leave the cluster.

The system MNP, `SG-CFS-pkg`, which provides the CVM infrastructure, is managed with the following commands:

```
cfscluster config
cfscluster unconfig
cfscluster start
cfscluster stop
cfscluster status
```

No additional parameters are required by these commands. They allow the SMNP to be created, deleted, started, stopped, and monitored in a simple manner. These commands always affect all nodes of the cluster.

CVM disk groups are administered with the `cfsdgadm` command. The following is a syntax summary for this command.

```
cfsdgadm add <shared_disk_group> [<package_name>]
           <node_name>=<activation_mode> ...
cfsdgadm add <shared_disk_group> [<package_name>]
           all=<activation_mode> [[<node_name>] ...]
cfsdgadm delete <shared_disk_group>
cfsdgadm modify <shared_disk_group> <node_name>=<activation_mode> ...
cfsdgadm modify <shared_disk_group> all=<activation_mode>
cfsdgadm modify <shared_disk_group> add <node_name>=<activation_mode> ...
cfsdgadm modify <shared_disk_group> delete <node_name> ...
cfsdgadm modify <shared_disk_group> dg <new_shared_disk_group_name>
cfsdgadm move <source_shared_disk_group> <target_shared_disk_group>
cfsdgadm show_package <shared_disk_group>
cfsdgadm show_autorun <shared_disk_group>
cfsdgadm display [-v] [<shared_disk_group>]
cfsdgadm set_autorun <shared_disk_group> {YES|NO}
cfsdgadm activate <shared_disk_group> [[<node_name>] ...]
cfsdgadm deactivate <shared_disk_group> [[<node_name>] ...]
```

To configure Serviceguard to manage a CVM disk group with the preceding commands, the disk group must be created first with `vxdg -s init` on the CVM master node (as reported by `cfscluster status` or by `vxdctl -c mode`). CVM disk groups can also be configured with VEA, which provides a convenient graphical interface for this task.

CFS mount points, checkpoints, and snapshots are administered using the `cfsmntadm` command. The following is a syntax summary, as applicable to mountpoints:

```

cfsmntadm add <shared_disk_group> <shared_volume> <mount_point>
    [<package_name>] <node_name>=[<mount_options>] ...
cfsmntadm add <shared_disk_group> <shared_volume> <mount_point>
    [<package_name>] all=[<mount_options>] [[<node_name>] ...]
cfsmntadm delete <device_path>
cfsmntadm delete <mount_point>
cfsmntadm modify <mount_point> <node_name>=[<mount_options>] ...
cfsmntadm modify <mount_point> <node_name>+=[<mount_options>] ...
cfsmntadm modify <mount_point> <node_name>-=[<mount_options>] ...
cfsmntadm modify <mount_point> all=[<mount_options>]
cfsmntadm modify <mount_point> all+=[<mount_options>]
cfsmntadm modify <mount_point> all-=[<mount_options>]
cfsmntadm modify <mount_point> add <node_name>=[<mount_options>] ...
cfsmntadm modify <mount_point> delete <node_name> ...
cfsmntadm modify <device_path> vol <new_volume_name>
cfsmntadm modify <mount_point> vol <new_volume_name>
cfsmntadm show_package <mount_point>
cfsmntadm show_autorun <mount_point>
cfsmntadm display [-v] {<mount_point>| [<node_name>]}
cfsmntadm setpolicy <mount_point> [[<node_name>] ...]
cfsmntadm set_autorun <mount_point> {YES|NO}
cfsmount <mount_point> [[<node_name>] ...]
cfsu mount <mount_point> [[<node_name>] ...]

```

Checkpoints and snapshots are also managed by the `cfsmntadm` command with similar syntax. The file systems to be managed must already be created and initialized before these commands are used, as described in the *VERITAS File System Administrator's Guide*.

VERITAS storage considerations

According to the *VERITAS Storage Foundation Installation Guide* (which includes a detailed discussion of version, capability, and compatibility considerations), Storage Foundation 4.1 supports existing disk groups with minimal version of 90. Storage Foundation 4.1 also supports VERITAS File System (VxFS) disk layout version 4 and above. These restrictions correspond to VERITAS 3.2 or later. Any disk groups and disks that were created with earlier versions of VERITAS software must be upgraded following the procedures detailed in the *VERITAS Storage Foundation Installation Guide*. The installation guide also contains a detailed description of different versions of disk groups and file system layouts. The recommendation is to upgrade the disk groups to version 120 and VxFS disk layout to version 6, the most current versions to be able to take advantage of all the capabilities available with Storage Foundation 4.1.

Disk group versions can be identified with the `vx dg list diskgroup` command, for example:

```

# vx dg list rootdg | grep version
version: 120

```

The disk groups can be upgraded using the `vx dg upgrade diskgroup` command.

To determine the VxFS disk layout version of a file system, run the `fstyp` command on the file system volume device, for example:

```

# fstyp -v /dev/vx/dsk/rootdg/lvol3 | grep version
version: 6

```

To upgrade disk layouts from version 4 and above, use the `vxfsconvert <path-to-raw-device>` command. To upgrade from earlier versions, use the `vxupgrade` command.

These operations are described in detail in the *VERITAS Storage Foundation Installation Guide*.

VERITAS Storage Foundation 4.1 no longer requires *rootdg* to be configured. The existing *rootdg* disk group will continue to work and does not need updating other than as mentioned earlier, but it is no longer required and may be deleted, if desired.

VERITAS 4.1 also introduces a new Cross-platform Data Sharing (CDS) disk layout as the default disk layout for any new disks. Because existing disk layouts continue to be supported, this will only be of consideration when creating new disks, unless you want to make existing data portable across different UNIX® platforms. Detailed information about disk layouts and tools that can be used to convert between them appears in the *VERITAS Volume Manager Administration Guide*.

With Storage Foundation 4.1, new disk groups are created with version 120 and new file systems are created with disk layout version 6, unless explicitly specified otherwise, which makes new disk groups and file systems incompatible with versions of VERITAS Storage Foundation earlier than 4.1. This condition is a limitation only if you intend to use them with earlier versions of VERITAS software.

Any user-developed applications that make file system calls in the *statvfsdev()* family of functions must be re-linked to enable them to work with file systems that have disk layout version 6.

HP Serviceguard configuration considerations

If a stand-alone version of CVM is installed on the cluster, it must be uninstalled before installing any of the HP Serviceguard Storage Management Suite bundles. This operation requires halting and deleting all CVM related packages: the application packages that specify *STORAGE_GROUP* attributes in their configuration and the SMNP package (VxVM-CVM-pkg for CVM 3.5 or SG-CFS-pkg for CVM 4.1). After a CFS bundle is installed, these packages can be reapplied or upgraded, as described in this section.

If the installation included installing CVM 4.1 (either stand-alone or as part of a bundle), the new SG-CFS-pkg SMNP must be used to manage CVM. No changes are required to existing application packages before reapplying them. The SG-CFS-pkg package can be configured and started using the following commands, if one of the CFS bundles is installed.

```
# cfscluster config
# cfscluster start
```

If stand-alone CVM is installed, these commands are unavailable. The SMNP for CVM can be applied with the *cmapplyconf* command in this case. The package name is VxVM-CVM-pkg for CVM 3.5 and SG-CFS-pkg for CVM 4.1. The package ASCII files and control scripts for these packages are located in */etc/cmcluster/cvm* and */etc/cmcluster/cfs* directories for CVM 3.5 and CVM 4.1, respectively.

If one of the CFS bundles is installed and you want to upgrade the existing packages to take advantage of the new dependency features, HP recommends that you use the following procedure:

1. Configure and start the CVM SMNP.
2. Configure and start disk group MNPs with the following commands:

```
# cfsdgadm add <dg-name> all=sw      # for shared/write everywhere
# cfsdgadm activate <dg-name>
```
3. Edit application packages' ASCII configuration files to remove *STORAGE_GROUP* attributes and add dependencies on the disk group packages created in step 2.
4. Generate new package scripts for these packages with *cmmakepkg -s*. Copy any custom code present in the old scripts to the new scripts. Because any disk groups used by the package are already defined in the MNPs, no variables that describe CVM disk groups used (for example, *CVM_DG[*]* and *CVM_DISK_ACTIVATION*) must be set in the script.
5. Apply and start the upgraded application packages.

Appendix A presents a set of illustrations of various upgrade scenarios that describe different possibilities of configuring Serviceguard to manage storage.

Data migration to CVM/CFS

Many sites will want to upgrade the way the application data is stored to take advantage of the CFS capabilities available with some of the bundles. Existing LVM volumes can continue to be used in the current form as long as desired, before upgrading to take advantage of the new features. Because there is no rolling upgrade process available from earlier versions of CVM under Serviceguard control to CFS, when CVM is used, the data cannot be migrated without restarting the applications.

As with earlier versions of Serviceguard and VERITAS, VxVM- and LVM-based storage can be used simultaneously both within a cluster and within a single package.

Because Serviceguard supports only LVM disks for cluster lock disk use, it does not make sense to upgrade these disks, unless the cluster is being reconfigured to use a Quorum Server instead.

Current data might reside on file systems or raw volumes. The raw data might be on CVM (with pre-4.1 VERITAS), LVM, or SLVM volumes. The file system is typically OnlineJFS (VxFS) on LVM or VxVM disk volumes.

Migrated data may reside on either raw CVM volumes or CFS file systems.

To convert Shared Logical Volume Manager (SLVM) data to CVM, the volume must be deactivated and unshared. The instructions for unsharing and deactivating an SLVM volume appear in the *Using Serviceguard Extension for RAC Third Edition* manual.

LVM volume groups must also be deactivated and taken out of the cluster to be converted. This can be accomplished with the following command (after halting any packages that might be using the volume groups):

```
# vgchange -a n <vg-name>; vgchange -c n <vg-name>
```

The volumes can then be converted “in place” to VxVM using the `vxvmconvert` command. Detailed documentation on this process can be found in the *VERITAS Volume Manager Migration Guide*. After the migration is complete, the resulting VxVM disk groups can be configured for Serviceguard management with `cfsdgadm` and used with applications as described in the “HP Serviceguard configuration considerations” section. If the data currently resides in a VxFS file system, this procedure can also be used to convert the underlying volume to VxVM, so it can be used as a CFS file system.

If the data is currently on a raw volume, the volume might need to be dumped to an intermediate location and restored to a new file system that is created for the data. Data could also be moved directly with `dd`, as detailed in the *VERITAS Storage Foundation Administration Guide*.

Oracle provides a variety of tools to move data—some of which are version-independent, and others are available with certain versions of Oracle Database. The canonical tools are the `exp` and `imp` commands, which work for all versions of the database and can be done without taking down the database, but might not perform as well as some of the other options. Oracle 10g features the new Data Pump mechanism that is designed for high performance and is recommended. A white paper describing this feature can be found on Oracle Technology Network website under the name *DATA PUMP IN ORACLE® DATABASE 10g: FOUNDATION FOR ULTRA HIGH-SPEED DATA MOVEMENT UTILITIES*. For earlier versions of Oracle, the database instance must be stopped for conversion. For Oracle 9i, HP recommends converting to Oracle Disk Manager files. Conversion processes for Oracle 9i are described in detail in the *VERITAS Storage Foundation for Oracle Administrators Guide*.

Other applications that support both CVM raw volumes and CFS file systems define their own procedures for data conversion. Consult application-specific documentation for the recommended procedures.

Other considerations

Configuring SGeRAC with Oracle RAC 9i and 10g is described in detail in the *Using Serviceguard Extension for RAC Third Edition* manual. Most of the upgrade considerations applicable to SGeRAC have already been described previously.

One desirable application property that becomes feasible with CFS is the capability of cluster-wide installation, configuration, or both. Because this typically requires some support from the application and explicit awareness of the cluster environment, it will be up to vendors of individual applications to provide these capabilities. Oracle is a good example of this type of application because it supports placing ORACLE_HOME on a CFS file system.

Migration from Tru64 UNIX and TruCluster platforms is being addressed and covered separately. Detailed information is available from the Tru64 Transition site at <http://www.hp.com/go/tru64transition/>.

Upgrade roadmap

To perform the upgrade:

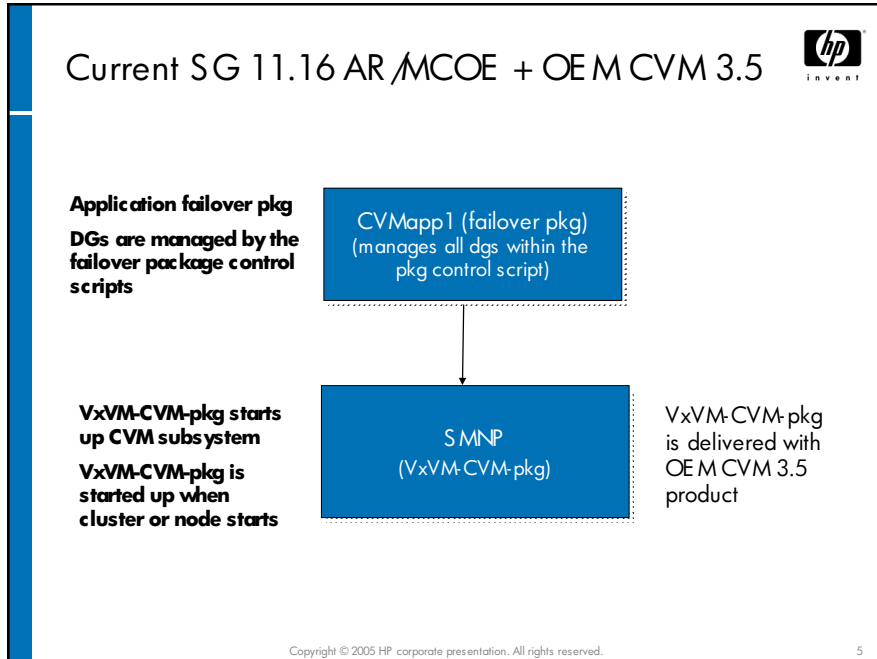
1. Perform a complete backup of existing configuration.
2. If one of the Serviceguard Storage Management bundles or stand-alone CVM will be installed:
 - a. Halt and delete from the cluster any application packages that use CVM. Retain the package configuration files to be applied after the upgrade.
 - b. Halt and delete from the cluster the CVM package **VxVM-CVM-pkg**, if present.
3. Halt the Serviceguard cluster.
4. Prevent Serviceguard from starting automatically during boot by changing the line `AUTOSTART_CMCLD=1` in `/etc/rc.config.d/cmcluster` to `AUTOSTART_CMCLD=0`.
5. Upgrade network hardware components as needed.
6. If a Serviceguard Storage Management bundle will be installed and CVM is currently installed on the cluster, remove it from each node with the following command:
`# swremove B9117AA`
7. Install the new software.
8. Make and apply any changes needed to the Serviceguard cluster ASCII configuration file to reflect network modifications.
9. If CVM was installed above as part of a CFS bundle, configure and apply the new CVM package **SG-CFS-pkg** with the following command:
`# cfscluster config`
10. If stand-alone CVM 4.1 was installed, apply the CVM 4.1 SMNP with:
`# cmapplyconf /etc/cmcluster/cfs/SG-CFS-pkg`
11. If CVM 3.5 is installed, apply the CVM 3.5 SMNP with:
`# cmapplyconf /etc/cmcluster/cvm/VxVM-CVM-pkg`
12. Migrate existing package configurations to use the new dependency infrastructure:
 - a. Create CVM disk group MNPs for the disk groups used by the packages.
 - b. Remove any `STORAGE_GROUP` attributes from the application packages' ASCII configuration files.
 - c. Add dependencies on the disk group packages to the application packages' ASCII configuration files.
 - d. Create new control scripts for the packages with `cmmakepkg -s`. Edit the resulting scripts to add any custom script functions from the previous version of the control scripts.
 - e. Apply the resulting packages to the cluster.

13. Perform any desired data migration steps.
14. Configure any CFS file systems needed.
15. Add new packages to represent any new file systems and disk groups created in step 12.
16. Re-enable automatic startup of Serviceguard by changing the line `AUTOSTART_CMCLD=0` in `/etc/rc.config.d/cmcluster` to `AUTOSTART_CMCLD=1`.
17. Restart the Serviceguard cluster or reboot.

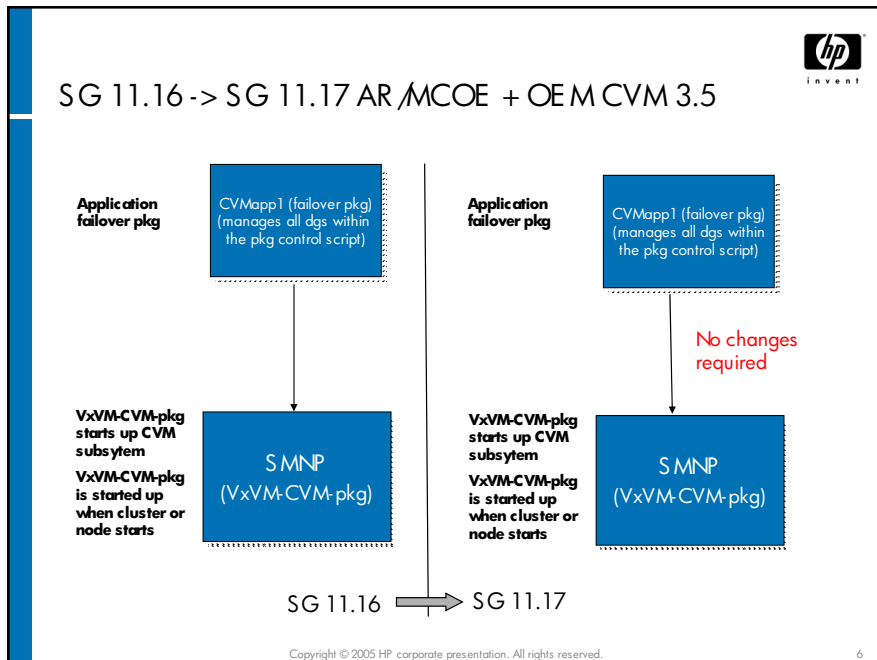
Appendix A

The following slides illustrate various upgrade scenarios and the relevant storage management options.

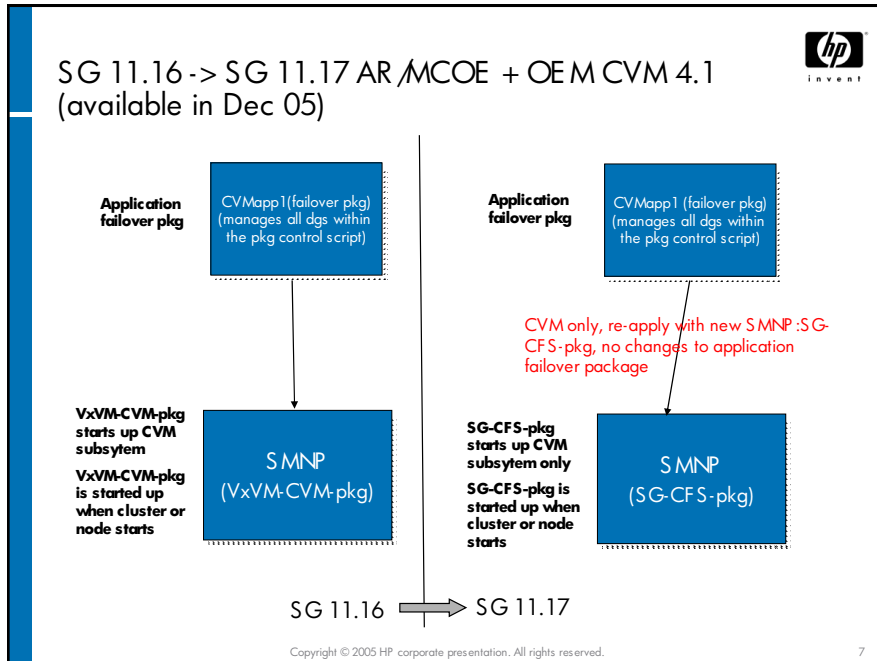
This is a typical pre-upgrade scenario. CVM support is provided by the VxVM-CVM-pkg SMNP, the application package is configured with STORAGE_GROUP attributes, and the package control script declares CVM_DG[*] and CVM_DISK_ACTIVATION variables to specify storage.



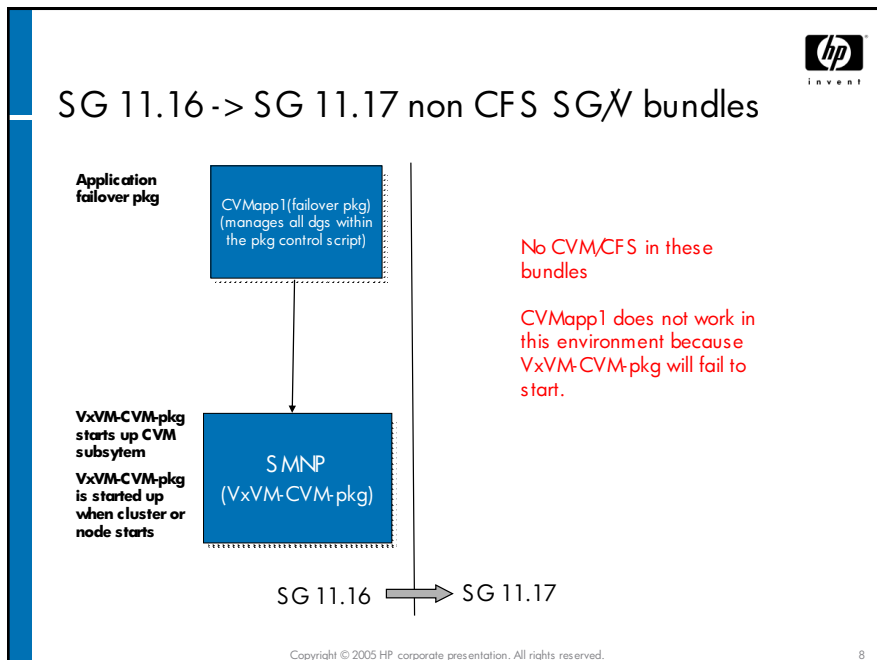
This is the scenario after upgrading Serviceguard only to A.11.17. There are no changes in the way storage is specified from Serviceguard A.1.16.



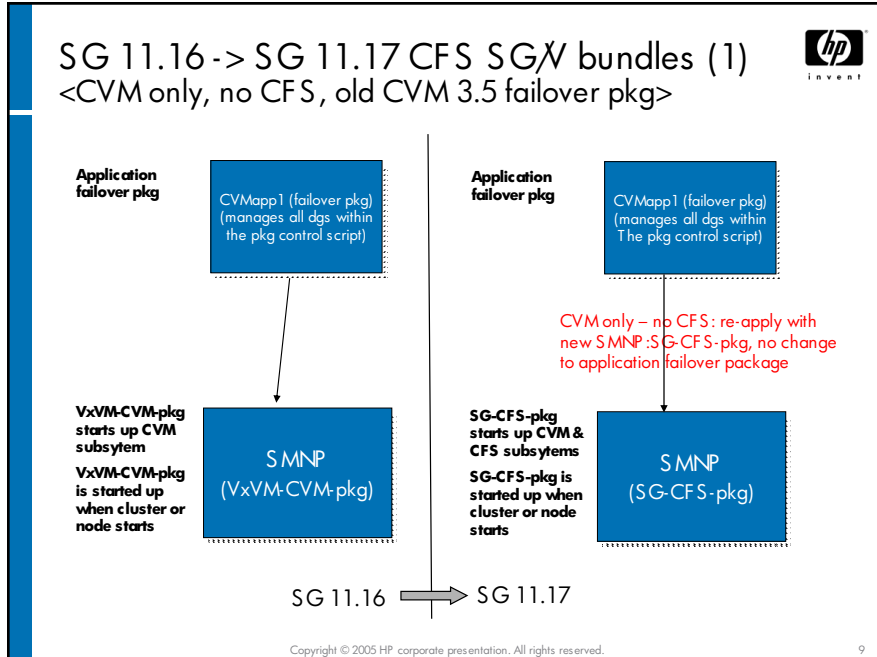
This is similar to the previous scenario, but CVM 4.1 is installed. The new SG-CFS-pkg SMNP is used; otherwise, configuration and management are unchanged.



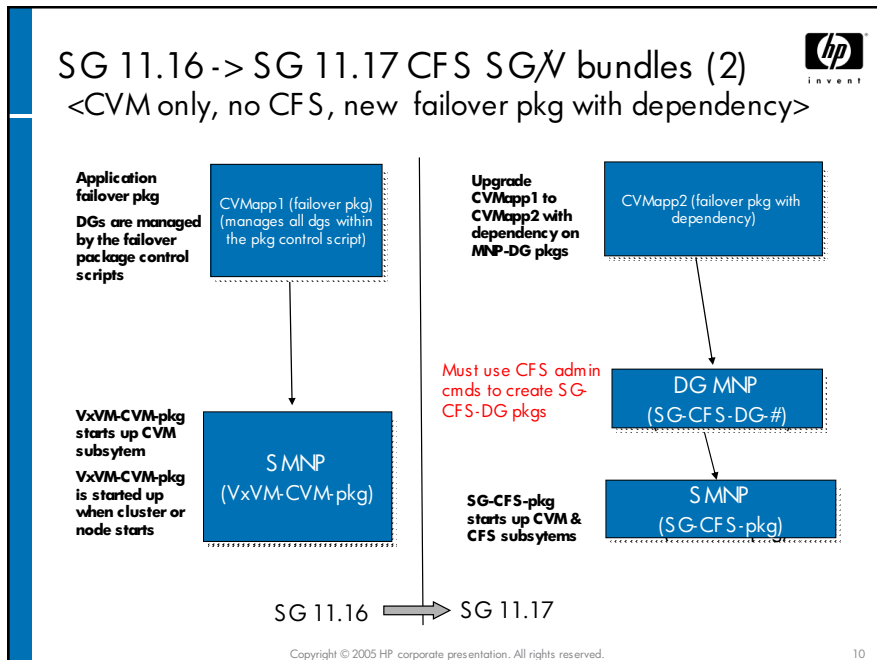
When any HP Serviceguard Storage Management bundle is installed, CVM must be removed before the installation and cannot be later installed without removing the bundle. This makes CVM unavailable to any application packages.



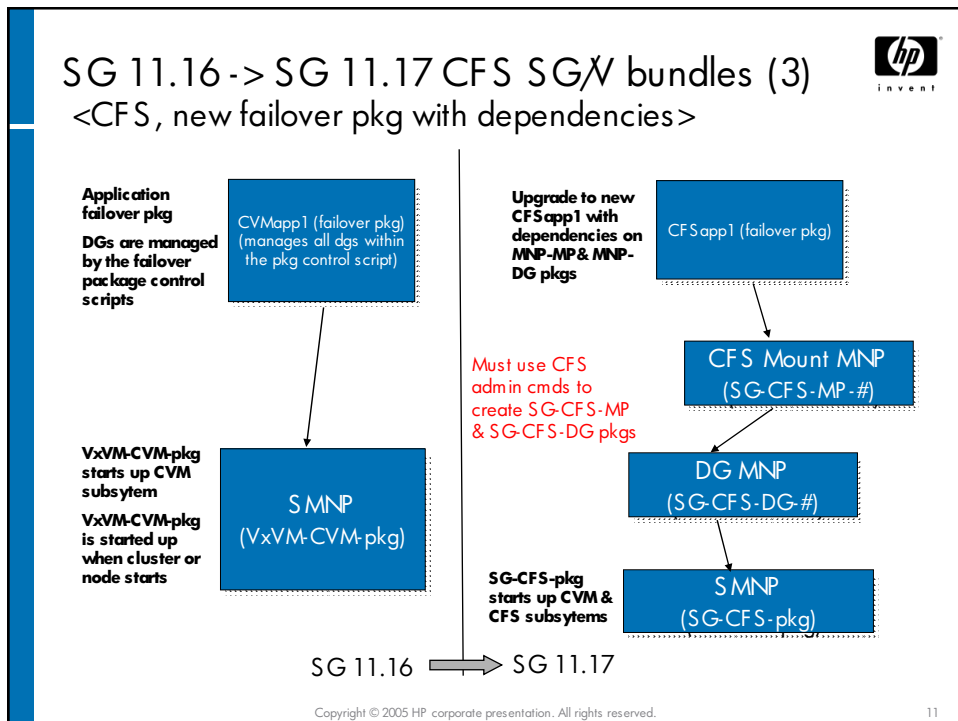
If a CFS bundle is installed and CFS is not used, the application package can be managed in a similar way to Serviceguard A.11.16. The new SMNP must be configured. This can be done with the `cfsccluster` command, which is part of the bundle.



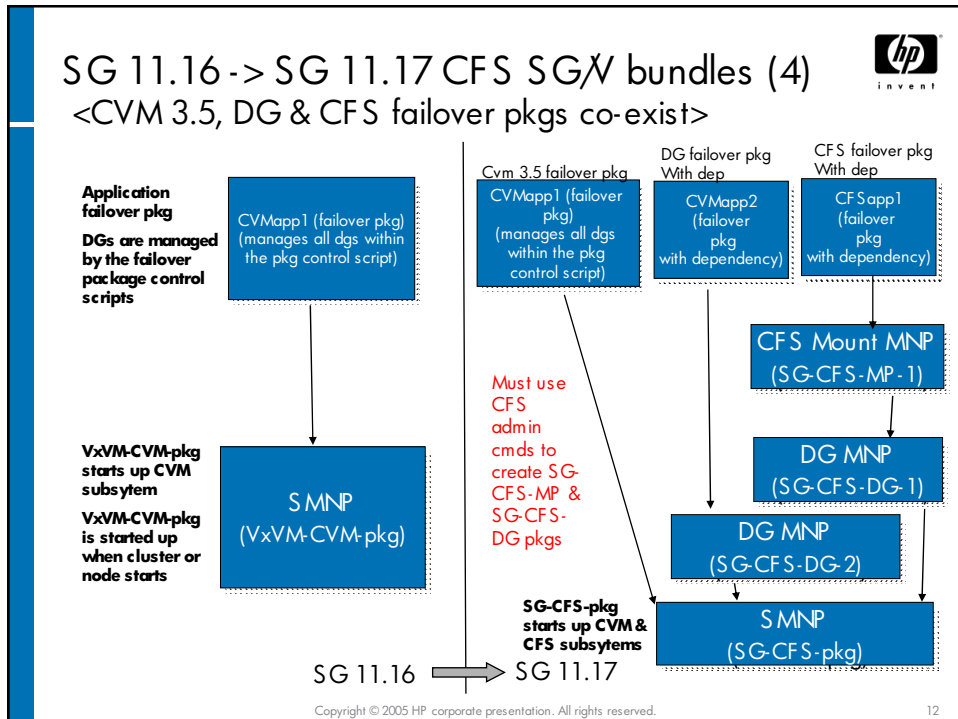
An application package can be upgraded to use the dependency infrastructure even if CFS is not used. The new commands `cfsccluster` and `cfsgadm` are used to configure and manage storage in this scenario.



This slide illustrates an application package that uses a file system.



This case illustrates using a pre-A.11.17 package, a new style package that uses a disk group directly, and another that uses CFS—all on the same cluster.



For more information

For more information, refer to the following references used in this document:

- *Managing Serviceguard Twelfth Edition*
- *Using Serviceguard Extension for RAC Third Edition*
- *VERITAS Storage Foundation Installation Guide*
- *VERITAS Volume Manager Administration Guide*
- *VERITAS File System Administrator's Guide*
- *VERITAS Storage Foundation for Oracle Administrators Guide*
- *DATA PUMP IN ORACLE® DATABASE 10g: FOUNDATION FOR ULTRA HIGH-SPEED DATA MOVEMENT UTILITIES*
- Tru64 Transition site at <http://www.hp.com/go/tru64transition/>

© 2005 Hewlett-Packard Development Company, L.P. 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.

Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California. UNIX is a registered trademark of The Open Group.

4AA0-1401ENW, August 2005

