BELLSOUTH®

BellSouth HP Serviceguard DR Recovery Job Aid

AMENDMENT HISTORY

Version	Date	Revision Description
1.0	04/13/2003	Initial Issuance
2.0	04/20/2004	Updated Header and Footer Information

This document contains information on how to recover a Service Guard Cluster in the event of a Disaster. The steps presented here are steps that are generally used in a Disaster Recovery Drill simulation. In the event of an actual disaster, conditions and system availability may not be optimal and slight deviation from the steps presented here may be necessary.

BACKGROUND INFORMATION

MC/ServiceGuard is a HP software product for HP9000/800 computers that allows you to create a high availability system by providing a level of network and systems redundancy where if a system component fails, the application can be migrated over to the backup system almost instantaneously.

In order to provide this redundancy, a minimum of 2 and a maximum of 8 computers are networked together (usually by a private network) into a cluster. Each of the single computer systems is called a node in the cluster and the application software is encapsulated into a package which is controlled by the MC/ServiceGuard cluster daemon (cmcld). In the event of a single service, node, network, or other resource failure, MC/ServiceGuard will automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption.

Most ServiceGuard configurations are comprised of the following:

- 2 lan connections to the enterprise network (primary and secondary)
- 2 lan connections to a “private network”
- 1 set of shared disk storage
- A minimum of 2 HP/9000 servers

Note: the above is not the minimum configuration for a Service Guard Cluster.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

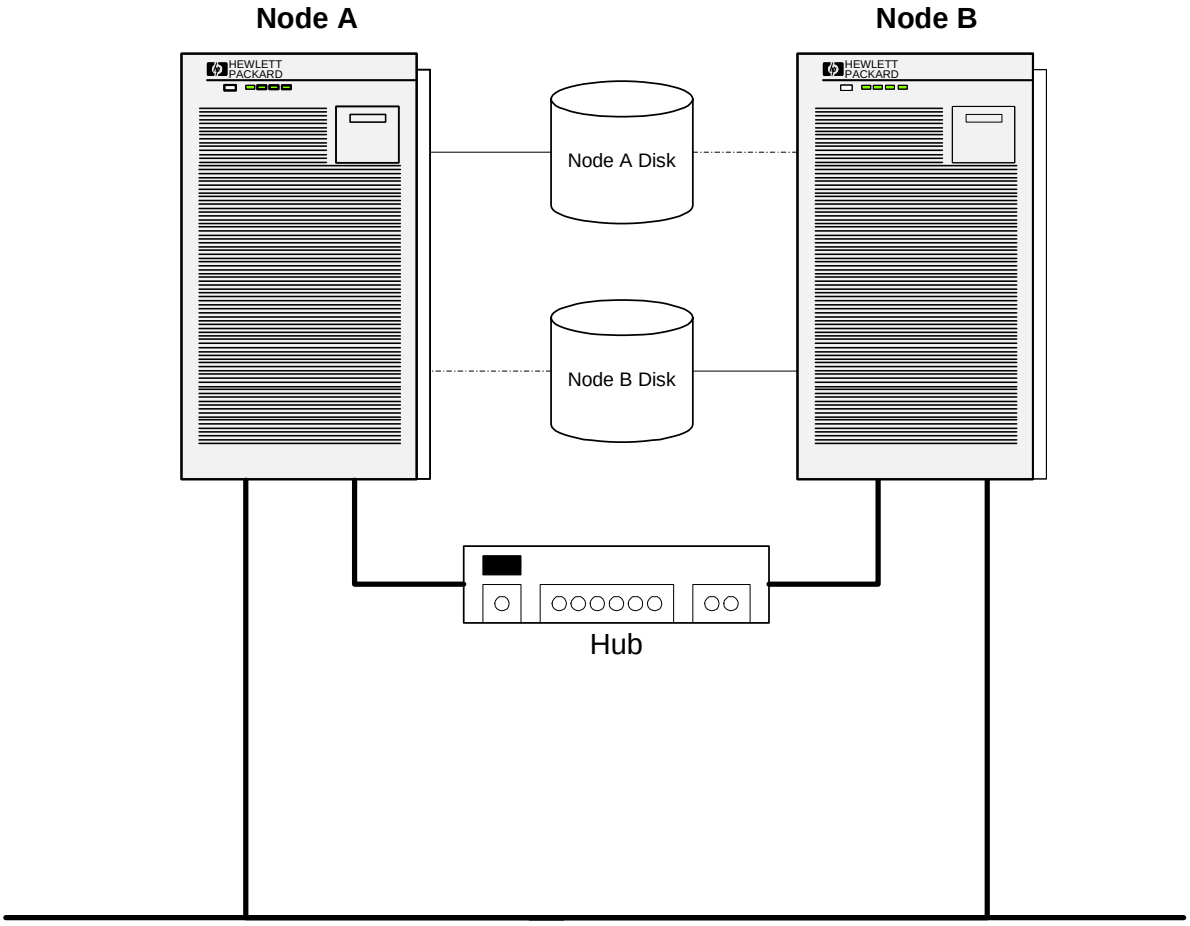

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

Disaster Recovery Instructions

Ignite

In a recovery situation, if you have been able to successfully ignite the system, and have not created your file systems for the EBR restore, you should take the following steps to recover your service guard cluster:

1 Disable the ServiceGuard application from starting on reboot:

   a. Edit ***/etc/rc.config.d/cmcluster*** and set **AUTOSTART_CMCLD=0**.

2 **Build** your Volume Groups, Logical Volumes and File Systems normally according to your recovery specs.

3 **Deactivate** and **unmount** the volume groups that you are planning on distributing across your cluster.

> **Note :** At the time you create the volume group, it is active on the primary node and before setting up the volume group for use on other nodes, you must first unmount any file systems that reside on the volume group, then deactivate it. At run time, volume group activation and file system mounting are done through the package control script.

```
# /umount /fs_name
# /vgchange –a n /vg_name
```

4 **Distribute** the volume groups to the other node(s) and make them shareable.

   a. Export the volume group map file information.

   ```
   # vgexport -p -s -m /tmp/vg_name.map  /dev/vg_name
   ```

   b. Copy the map file to the other nodes.

   ```
   # rcp /tmp/vg_name.map nodeX:/tmp/vg_name.map
   ```


PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

c. Create the volume group directory
   # mkdir /dev/vg_name

d. Create the volume group "group" special file on the new node.
   # mknod /dev/vg_name/group c 64 0xhh0000

e. Import the volume group map file from the parent node.
   # vgimport -s -m /tmp/vg_name.map /dev/vg_name

   Note the –s option which is to make the volume group ServiceGuard shareable.

f. Create the volume group mount point.
   # mkdir /vg_name_mount_point

g. Enable the volume group.
   # vgchange -a y /dev/vg_name

h. Mount the volume group in order to verify it.
   # mount /dev/vg_name/lvolx /vg_name_mount_point
   Note: You may have multiple logical volumes to mount depending on your disk layout

i. Unmount the volume group.
   # umount /vg_name_mount_point

j. Deactivate the volume group
   # vgchange -a n /dev/vg_name

Once you have completed these steps for all of the nodes in your cluster, your shared disk will be able to be seen and float from one node to the other.

If you have not already restored your system data via EBR you can begin the restore. **Caution** is noted here, that you have the volume group mounted on the system that it will be restored on.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

Cluster Reconfiguration

Depending on your Disaster Recovery environment and the hardware available, you may need to reconfigure your ServiceGuard cluster by modifying the cluster configuration files and/or the package configuration files. The details will vary from system to system based on the naming conventions uses and the level of package integration, but we will outline the primary trouble spots to look at and how to reconfigure the cluster.

Once you have restored your system and rebooted, the cmcluster daemon should be up but your cluster will not be formed because you should have turned off that option (see the beginning of this document) on all of the nodes prior laying out the volume groups and file systems.

1. Make a **backup** copy of your cluster configuration file.

   #cp /etc/cmcluster/clusterconfigfile.ascii /etc/cmcluster/clusterconfigfile.ascii.backup

2. **Query** your cluster and create a new cluster configuration file for comparison

   # cmquerycl -v -C /etc/cmcluster/cmclconf_tempfile.ascii -n node1 -n node2 … -n node8

3. **Compare** your original cluster configuration file against the queried version and make the appropriate modification.

   Note : In this section we have provided for you a sample configuration file with embedded notes and highlighted pertinent sections of interest.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

```
# **********************************************************************
# ********* HIGH AVAILABILITY CLUSTER CONFIGURATION FILE ***************
# ***** For complete details about cluster parameters and how to    ****
# ***** set them, consult the cmquerycl(1m) manpage or your manual. ****
# **********************************************************************

# Enter a name for this cluster.  This name will be used to identify the
# cluster when viewing or manipulating it.

CLUSTER_NAME            cluster1

# Cluster Lock Device Parameters.  This is the volume group that
# holds the cluster lock which is used to break a cluster formation
# tie.  This volume group should not be used by any other cluster
# as cluster lock device.

FIRST_CLUSTER_LOCK_VG   /dev/vglock    check existence of lock disk if needed.


# Definition of nodes in the cluster.
# Repeat node definitions as necessary for additional nodes.

NODE_NAME               node1          validate your network connections
  NETWORK_INTERFACE     lan0           all of your nodes!
    HEARTBEAT_IP        90.128.121.12
  NETWORK_INTERFACE     lan3
    HEARTBEAT_IP        128.6.7.3
  NETWORK_INTERFACE     lan4
  NETWORK_INTERFACE     lan1
    HEARTBEAT_IP        90.45.143.10
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c1t2d0


# List of serial device file names
# For example:
# SERIAL_DEVICE_FILE    /dev/tty0p0


# Primary Network Interfaces on Bridged Net 1: lan0.
#   Warning: There are no standby network interfaces on bridged net 1.
# Primary Network Interfaces on Bridged Net 2: lan3.
#   Possible standby Network Interfaces on Bridged Net 2: lan4.
# Primary Network Interfaces on Bridged Net 3: lan1.
#   Warning: There are no standby network interfaces on bridged net 3.

NODE_NAME               node2
  NETWORK_INTERFACE     lan0
    HEARTBEAT_IP        90.128.121.10
```


PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

```
  NETWORK_INTERFACE     lan3
    HEARTBEAT_IP        128.6.7.4
  NETWORK_INTERFACE     lan4
  NETWORK_INTERFACE lan1
    HEARTBEAT_IP        90.45.143.20
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c1t2d0


# List of serial device file names
# For example:
# SERIAL_DEVICE_FILE    /dev/tty0p0


# Primary Network Interfaces on Bridged Net 1: lan0.
#   Warning: There are no standby network interfaces on bridged net 1.
# Primary Network Interfaces on Bridged Net 2: lan3.
#   Possible standby Network Interfaces on Bridged Net 2: lan4.
# Primary Network Interfaces on Bridged Net 3: lan1.
#   Warning: There are no standby network interfaces on bridged net 3.


# Cluster Timing Parameters (microseconds).

HEARTBEAT_INTERVAL       1000000
NODE_TIMEOUT             2000000


# Configuration/Reconfiguration Timing Parameters (microseconds).

AUTO_START_TIMEOUT       600000000
NETWORK_POLLING_INTERVAL 2000000

# Package Configuration Parameters.
# Enter the maximum number of packages which will be configured in the
cluster.
# You can not add packages beyond this limit.
# This parameter is required.
MAX_CONFIGURED_PACKAGES           8


#
# List of cluster aware Volume Groups. These volume groups
# will be used by clustered applications via the vgchange -a e command.
# For example:
# VOLUME_GROUP          /dev/vgdatabase
# VOLUME_GROUP          /dev/vg02
VOLUME_GROUP            /dev/vg01     make sure these exist and correct
VOLUME_GROUP            /dev/vg02
```


PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

4. If you have a 2 node cluster you will need to **verify your cluster lock disk**.

Note: The volume listed in your configuration file under FIRST_CLUSTER_LOCK_VG should exist and be properly configured.

5. Check your cluster configuration.

cmcheckconf -v -C /etc/cmcluster/cmclconf.ascii

The following items will be checked:

- Network addresses and connections.
- Cluster lock connectivity.
- Validity of configuration parameters for the cluster and packages.
- Uniqueness of names.
- Existence and permission of scripts specified in the command line.
- If all nodes specified are in the same heartbeat subnet.
- If you specify the wrong configuration filename.
- If all nodes can be accessed.
- No more than one CLUSTER_NAME, HEARTBEAT_INTERVAL, and AUTO_START_TIMEOUT are specified.
- The value for HEARTBEAT_INTERVAL is at least one second.
- The value for NODE_TIMEOUT is at least twice the value of HEARTBEAT_INTERVAL.
- The value for AUTO_START_TIMEOUT variables is >=0.
- Heartbeat network minimum requirement. The cluster must have one heartbeat LAN configured with a standby, or two heartbeat LANs, or one heartbeat LAN and an RS232 connection.
- At least one NODE_NAME is specified.
- Each node is connected to each heartbeat network.
- All heartbeat networks are of the same type of LAN.
- The network interface device files specified are valid LAN device files.
- VOLUME_GROUP entries are not currently marked as cluster-aware.

6. Redistribute your cluster binary.

- Activate the cluster lock volume group so that the lock disk can be initialized:
- # vgchange -a y /dev/vglock

- Generate the binary configuration file and distribute it across the nodes.
- # cmapplyconf -v -C /etc/cmcluster/cmclconf.ascii -P /etc/cmcluster/pkg1/pkg1conf.ascii

- Deactivate the cluster lock volume group.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

- # vgchange -a n /dev/vglock

7. Start your cluster

cmruncl cluster_name

Your cluster should come up cleanly. If there are problems in your cluster configuration, you will not be able to compile and distribute you cluster binary. In the event of errors during your cluster validation that you can not fix, consult with a recovery group or ServiceGuard specialist.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

Reference : MC/ServiceGuard Commands

The following is a list of commands used for MC/ServiceGuard configuration and maintenance. Man pages for these commands are available on your system *after installation*.

Command	**Description**
cmapplyconf	Verify and apply MC/ServiceGuard and MC/LockManager cluster configuration and package configuration files. cmapplyconf verifies the cluster configuration and package configuration specified in the cluster_ascii_file and the associated pkg_ascii_file(s) , creates or updates the binary configuration file, called cmclconfig, and distributes it to all nodes. This binary configuration file contains the cluster configuration information as well as package configuration information for all packages specified. This file, which is used by the cluster daemons to manage the entire cluster and package environment, is kept in the /etc/cmcluster directory. If changes to either the cluster configuration or to any of the package configuration files are needed, first update the appropriate ASCII ile(s) (cluster or package), then validate the changes using the cmcheckconf command and then use cmapplyconf again to verify and redistribute the binary file to all nodes. The cluster and package configuration can be modified whether the cluster is up or down, although some configuration requires either the cluster or the package be halted. Please refer to the manual for more detail. The cluster ASCII file only needs to be specified if configuring the cluster for he first time, or if adding or deleting nodes to the cluster. The package ASCII file only needs to be specified if the package is being added, or if the package configuration is being modified.
cmapplyconf	It is recommended that the user run the cmgetconf command to get either the

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

(continued)	cluster ASCII configuration file or package ASCII configuration file whenever changes to the existing configuration are required.Note that cmapplyconf will verify and distribute cluster configuration or package files. It will not cause the cluster daemon to start or removed from the cluster configuration. The same kind of processing will apply to the package configuration to determine whether to add or delete package nodes, package subnet, etc. Not all package configuration changes require the package to be halted.
cmcheckconf	Check high availability cluster configuration and/or package configuration files. cmcheckconf verifies the cluster configuration as specified by the cluster_ascii_file and/or the package configuration files specified by each pkg_ascii_file in the command. If the cluster has already been configured previously, the cmcheckconf command will compare the configuration in the cluster_ascii_file against the previously configuration information stored in the binary configuration file and validates the changes. The same rules apply to the pkg_ascii_file. It is not necessary to halt either the cluster or any of the packages to run the cmcheckconf command.
cmdeleteconf	Delete either the cluster or the package configuration. cmdeleteconf deletes either the entire cluster configuration, including all its packages, or only the specified package configuration. If neither cluster_name nor package_name is specified, cmdeleteconf will delete the local cluster's configuration and all its packages. If only the package_name is specified, the configuration of package_name in the local cluster is deleted. If both cluster_name and package_name are specified, the package must be configured in the cluster_name, and only the package package_name will be deleted. The local cluster is the cluster that the node running the cmdeleteconf command belongs to.
cmgetconf	Get cluster or package configuration information. Cmgetconf obtains either the cluster configuration, not including the package configuration, or the specified package's configuration information, and writes to either the output_filename file, or to stdout. This command can be run whether the cluster is up or down. If neither cluster_name nor package_name is

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

	specified, cmgetconf will obtain the local cluster's configuration. If both cluster_name and package_name are specified, the package must be configured in the cluster_name, and only the package configuration for package_name will be written to output_filename or to stdout.
cmhaltcl	Halt a high availability cluster. cmhaltcl causes all nodes in a configured cluster to stop their cluster daemons, optionally halting all packages or applications in the process.This command will halt all the daemons on all currently running systems. If the user only wants to shutdown a subset of daemons, the cmhaltnode command should be used instead.
cmhaltnode	Halt a node in a high availability cluster. cmhaltnode causes a node to halt its cluster daemon and remove itself from the existing cluster.When cmhaltnode is run on a node, the cluster daemon is halted and, optionally, all packages that were running on that node are moved to other nodes if possible.
cmhaltpkg	Halt a high availability package. cmhaltpkg performs a manual halt of high availability package(s) running on MC/ServiceGuard or MC/LockManager clusters. This command may be run on any node within the cluster and may operate on any package within the cluster.
cmmodpkg	Enable or disable switching attributes for a high availability package. cmmodpkg enables or disables the ability of a package to switch to another node upon failure of the package, and it enables or disables a particular node from running specific packages.Switching for a package can be enabled or disabled globally. For example, if a globally disabled package fails, it will not switch to any other node, and if a globally enabled package fails, it will attempt to switch to the first available node on which it is configured to run.
cmquerycl	Query cluster or node configuration information. cmquerycl searches all specified nodes for cluster configuration and Logical

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

	Volume Manager (LVM) information. Cluster configuration information includes network information such as LAN interface, IP addresses, bridged networks and possible heartbeat networks. LVM information includes volume group (VG) interconnection and file systemmount point information. This command should be run as the first step in preparing for cluster configuration. It may also be used as a troubleshooting tool to identify the current configuration of a cluster.
cmruncl	Run a high availability cluster. cmruncl causes all nodes in a configured cluster or all nodes specified to start their cluster daemons and form a new cluster.This command should only be run when the cluster is not active on any of the configured nodes. If a cluster is already running on a subset of the nodes, the cmrunnode command should be used to start the remaining nodes and force them to join the existing cluster.
cmrunnode	Run a node in a high availability cluster. cmrunnode causes a node to start its cluster daemon to join the existing cluster.Starting a node will not cause any active packages to be moved to the new node. However, if a package is DOWN, has its switching enabled, and is able to run on the new node, that package will automatically run there.
cmrunpkg	Run a high availability package. cmrunpkg runs a high availability package(s) that was previously halted.This command may be run on any node within the cluster and may operate on any package within the cluster. If a node is not specified, the node on which the command is run will be used. This will result in an error if the current node is not able to run the package or is not in the list of possible owners of the package. When a package is started on a new node, the package's run script is executed.
cmscancl	Gather system configuration information from nodes with MC/ServiceGuard or MC/LockManager installed. cmscancl is a configuration report and diagnostic tool which gathers system software and hardware configuration information from a list of nodes, or from

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.

	all the nodes in a cluster. The information that this command displays includes LAN device configuration, network status and interfaces, file systems, LVM configuration, link-level connectivity, and the data from the binary cluster configuration file. This command can be used as a troubleshooting tool or as a data collection tool.
cmviewcl	View information about the current high availability cluster. cmviewcl displays the current status information of a cluster. Output can be displayed for the whole cluster or it may be limited to particular nodes or packages.
cmviewconf	View MC/ServiceGuard or MC/LockManager cluster configuration information. cmviewconf collects and displays the cluster configuration information, in ASCII format, from the binary configuration file for an existing cluster. Optionally, the output can be written to a file. This command can be used as a troubleshooting tool to identify the configuration of a cluster.

PRIVATE/PROPRIETARY/SECURE
Contains Private and/or Proprietary Information. May Not Be Used Or Disclosed Outside
The BellSouth Companies Except Pursuant To A Written Agreement. Must Be Securely Stored When Not In Use.