Sample Configurations with SGeRAC and Oracle RAC 11gR1

Version 1.1, March 2009

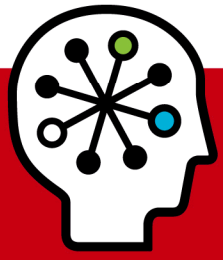

Introduction

This document discusses the various aspects of architecting, planning, and implementing an Oracle Real Application Clusterware (RAC) 11g Release 1 solution on HP-UX 11i with Serviceguard and Serviceguard Extension for RAC (SGeRAC). The document also includes a brief description about the planning considerations and presents a sample configuration for Oracle RAC 11g Release 1 with Serviceguard, Serviceguard Extension for RAC (SGeRAC) and Shared Logical Volume Manager (SLVM), and Cluster File System (CFS).

This document also provides step by step installation instructions for creating a SGeRAC cluster, creating shared storage using SLVM and CFS, installing Oracle Clusterware (OC) and RAC, creating a demo database, and creating Serviceguard packages to synchronize start and stop of the complete solution stack.

Audience

The target audiences are those who are interested in architecting, planning, and implementing SGeRAC and Oracle RAC 11g clusters.

The reader should be familiar with Serviceguard (SG), Serviceguard Extension for RAC (SGeRAC), Oracle RAC11g software, Shared Logical Volume Manager (SLVM), Symantec Veritas Cluster Volume Manager (CVM), Symantec Veritas Cluster File System (CFS), HP-UX 11i v2, and HP-UX 11i v3.

Terms and definitions

- APA – Auto Port Aggregation provides bonding of multiple networking interface cards where traffic is distributed to all interface cards.
- APA/Hot Standby – Auto Port Aggregation Hot Standby mode provides high availability through bonding of a primary and a standby interface card. Traffic is not distributed.
- CFS – Cluster File System allows multi-system shared access to common file system.
- CSS – Cluster Synchronization Service is a component of Oracle Clusterware that maintains Oracle cluster membership and heartbeat.
- CSS-HB – Cluster Synchronization Service heartbeat traffic
- CVM – Cluster Volume Manager allows multi-system shared access to volumes.
- GAB – Group Membership Service/Atomic Broadcast manages cluster membership and cluster communication for Symantec Veritas CFS 4.1/5.0 and CVM 4.1/5.0.
- GMS – Group Membership Service refers to HP's implementation of the NMAPI2 API on HP-UX with SGeRAC that provides group membership notification and process monitoring facility to monitor group status.
- HA – High Availability refers to configurations that are resilient to single failure.
- JBOD – Just a Bunch of Disks refers to single disk or a set of disk in disk enclosures that do not provide RAID capability for HA.
- LLT – Low Latency Transport provides kernel-to-kernel communications at link level and monitors network connections for Symantec Veritas CFS 4.1/5.0 and CVM 4.1/5.0. Distributes Symantec Veritas traffic amount network connections and maintains Symantec Veritas heartbeat.
- MNP – Multi-node package, a Serviceguard package that runs on multiple nodes at the same time and can be independently started and halted on individual nodes.

- NIC – Network Interface Card, host bus adapter for network communications, for example Ethernet card.
- OC – Oracle Clusterware can run in conjunction with Serviceguard Extension for RAC and provides Oracle cluster membership and resource management services.
- OCR – Oracle Cluster Registry is shared storage used to keep Oracle cluster and configuration information.
- ODM – Oracle Disk Manager is a standard API specified by Oracle for database I/O.
- RAC – Real Application Cluster enables a multi-instances concurrent shared access database.
- RAC-DB-IC – Real Application Cluster Interconnect traffic for both Global Cache Service and Global Enqueue Service.
- RAID – Redundant Array of Independent Disks refers to disk storage that provides HA through redundancy within an array of disks by internal mirroring or use of parity disks.
- RIP – Serviceguard Relocatable IP Address user for client application access and failovers with package
- SG-HB – Serviceguard Heartbeat traffic.
- SGeRAC – Serviceguard Extension for RAC extends Serviceguard to support Oracle RAC.
- SLVM – Shared Logical Volume Manger allows multi-system shared access to LVM volumes for RAC.
- VIP – Virtual IP address is used by OC to configure access to Oracle clients and for remote failover to reject client connections.
- Voting Disk – Shared storage used by Oracle Clusterware as vote tie breaker and for disk based heartbeat.

Additional information

The audience is encouraged to use additional documentation in conjunction with this document. Sources are listed in the References section.

- Serviceguard Version A.11.19 Release Notes
- Serviceguard Extension for RAC Version A.11.19 Release Notes
- Managing Serviceguard Sixteenth Edition, March 2009
- Using Serviceguard Extension for RAC, Sixth Edition, March 2009
- Oracle Clusterware Installation Guide 11g Release 1 (11.1) for HP-UX, November 2007
- Oracle Real Application Clusters Installation Guide 11g Release 1 (11.1) for Linux and UNIX
- Oracle Clusterware Administration and Deployment Guide 11g Release 1 (11.1), September 2007
- Oracle Real Application Clusters Administration and Deployment Guide 11g Release 1 (11.1), November 2007
- Symantec Veritas Storage Foundation 5.0 for Oracle RAC Configuration Guide Extracts for HP Serviceguard Storage Management Suite Extracts
- Symantec Veritas Storage Foundation 4.1 for Oracle RAC HP Serviceguard Storage Management Suite Extracts

Planning considerations

Proper planning is a requirement for high availability configurations – there are important considerations when deciding the cost-benefit tradeoffs.

Capacity planning

Proper capacity planning ensures sufficient resources are available to meet the expected service levels.

Processor capacity

In a cluster environment, cluster members maintain heartbeat traffic between nodes. Insufficient CPU processing capacity affects heartbeat processing and thus affects the heartbeat timeout. A larger timeout means a longer cluster reconfiguration time and a longer time before a node failure or network failure can be detected.

Memory

Sufficient physical memory should be available for all processes. Insufficient memory may result in swapping activities that affect the CPU processor availability to components that have timed heartbeat communications. Generally, one should consider sufficient memory to reduce paging activities, fit Oracle System Global Area (SGA) into main memory, and allow for user processes.[1]

Network: clients

Insufficient bandwidth on the client network affects availability to the client.

Network: cluster interconnect

Insufficient bandwidth from the cluster interconnect affects communication between cluster components.

Storage

Sufficient storage bandwidth and storage space are required to maintain optimal database service. Additionally, sufficient space should be allocated for database archives for recovery purposes.

Failover time requirements

In a RAC configuration with SGeRAC, each node may have concurrent access to the database. The database service is accessible from all nodes and each node provides a client connection endpoint (IP address, port, and listener). When one node fails, clients can connect to another node for services. The client connection endpoint does not need to failover for the clients to continue service. However, even though an alternate connection endpoint is available, upon certain failures (for example node or network) and until failure detection and recovery, new client connections may not connect or the database service may be unavailable while the cluster goes through reconfiguration and/or recovery.

The failover time is the time between when a failure occurs to when the service is once again available to the client. A failover time requirement is important for the following reasons:

How fast the clients reconnect.

- On local LAN failover, this depends on detection time and the local LAN failover scheme.
- On remote failover, this depends on whether clients are enabled with Oracle Fast Application Notification (FAN), how fast cluster reconfiguration happens, how soon the VIP address fails over, and how soon the client connection times out.

How fast the cluster and RAC go through reconfiguration before database service resumes.

[1] See Oracle Database Performance Tuning Guide 11g Release 1 (11.1), B282274-01, July 2008 at http://download.oracle.com/docs/cd/B28359_01/server.111/b28274.pdf.

- On node failure, the reconfiguration time depends on the Serviceguard heartbeat timeout, the number of nodes, and the type of quorum device used.
- On cluster network interconnect failures, the database service availability depends on how soon the interconnect failure is discovered, speed of recovery actions by affected components (for example SG, GMS, SLVM, CSS, and/or RAC), and database recovery.
    - On a complete SG cluster interconnect failure, SG sees the failure within the SG heartbeat timeout.
    - With SG/CFS, GAB/LLT and SG shared the same networks and SG sees the interconnect failure within SG heartbeat timeout.
    - With CSS traffic, on configurations where CSS and SG share the same interconnect network, SG sees the failure within the SG heartbeat timeout. If CSS traffic is on a SG monitored network, SG can be configured to take actions via SG packages and optionally use Cluster Interconnect Subnet Monitoring[2]. If CSS traffic is on a non-SG monitored network, CSS sees the interconnect failure within the CSS timeout.
    - With RAC traffic, if the interconnect is not configured to be monitored and acted upon by the other components, RAC discovers the interconnect failure within the Instance Membership Recovery (IMR) timeout.

The failover time requirement determines important timeouts, such as Serviceguard heartbeat timeout, network polling intervals, and cluster interconnect monitoring.

Planning for high availability

A properly configured high availability configuration should survive single failures and continue to operate.

Public network

The following describes the two levels of client public network high availability (HA): redundant components and client failover.

Redundant network interfaces and switches with local LAN failover by Serviceguard or bonding by Auto-Port Aggregation (APA) protects against single point network failures.

Client failover is needed with failures that impact existing and new client sessions. These failures include node failures (such as those caused by a power failure) and network failures (for example all redundant network interface/link failed). Protection is available at three levels: Oracle Fast Application Notification (FAN), remote VIP failover, and client connect timeout. Clients that are FAN integrated or using the FAN API may interrupt existing sessions and failover. Remote VIP failover is useful for non-FAN clients attempting to connect to the local node to avoid the TCP connect timeout. The client connect timeout is useful when client connect takes a long time for whatever reason.

VIP high availability

This section describes high availability for VIP.

[2] Cluster Interconnect Subnet Monitoring is available starting SGeRAC A.11.18

Serviceguard local LAN failover mechanism – preferred choice

For client public network HA in a SGeRAC configuration, the preferred method for network HA is to use Serviceguard primary and standby. Serviceguard monitors the redundant network and additional APA software is not required.

When the client network is configured with Serviceguard local LAN failover, Serviceguard performs the local LAN failover and Oracle Clusterware (OC) configures the VIP after Serviceguard local LAN failover. Since OC performs monitoring and manages the VIP address, client connectivity maybe unavailable until OC detects the outage and configures the VIP address on the local node. [3]

Local LAN failover using APA

When APA is used where the network interface cards are bonded, APA provides traffic distribution and load balancing capability among multiple physical network interface cards (NIC) or links. Load balancing may be a benefit which is desirable to configurations where a single interface is insufficient to handle the network traffic. When a physical NIC or link fails, APA provides HA by distributing traffic among remaining NIC or links. One virtual link is presented to OC and APA network load balancing is transparent to OC. APA requires the same type of NIC. Since APA network connections go to the same switch, a switch failure means outage of the client network.

When APA/Hot Standby is used, APA/Hot Standby provides the primary-to-hot-standby failover by rerouting traffic from failed primary link to hot standby link. APA/Hot Standby does not load balance. Serviceguard does not monitor this network. One virtual network link is presented to OC and the physical failover is transparent to OC because the same virtual network link remains available. Both NICs must be the same type as in APA.

Remote failover

OC fails over the VIP address to a surviving node on a catastrophic failure such as node failure or network failure.

Network for cluster communication

Serviceguard, OC, and each RAC instance maintain communication with peers on other nodes. When communication is broken, either through network partition or node failure, each of these components needs to reform its membership and eject non-members as needed.

In CFS and CVM (4.1/5.0) configurations, Symantec Veritas' Group Membership Service/Atomic Broadcast and Low Latency Transport (GAB/LLT) also uses the cluster interconnect for peer to peer communications.

The categories of traffic between nodes are distinguished as follows:

- SG-HB – Serviceguard heartbeat and communications traffic. Supported over single or multiple subnet networks.
- CSS-HB – Cluster Synchronization Service (CSS) heartbeat traffic and communications traffic for Oracle Clusterware. CSS-HB uses a single logical connection over a single subnet network.
- RAC-DB-IC – RAC instance peer to peer traffic and communications for Global Cache Service (GCS) and Global Enqueue Service (GES), formerly Cache Fusion (CF) and Distributed Lock Manager (DLM). Network HA is provided by the HP-UX 11i platform (Serviceguard or APA bonding).

[3] See Doc ID: Note:296874.1 Configuring the HP-UX Operating System for the Oracle 10g VIP at https://metalink.oracle.com/ (Oracle MetaLink account required)

- ASM-IC – Applicable only when using Automatic Storage Management (ASM). ASM instance peer to peer traffic. When it exists, ASM-IC should be on the same network as CSS-HB. Network HA is required either through Serviceguard failover or APA bonding.
- GAB/LLT – Applicable only when using CFS/CVM. Symantec cluster heartbeat and communications traffic. GAB/LLT communicates over link level protocol (DLPI) and supported over Serviceguard heartbeat subnet networks, including primary and standby links. GAB/LLT is not supported over APA or virtual LANs (VLAN).

Note that each category maintains its own timeout for which members may be evicted from its respective membership.

The interconnect network requires HA configurations. When a single network failure occurs, for example LAN card or switch failures, all the cluster nodes continue to operate. Without HA, a single network failure results in a network partition between the nodes and evicted nodes are halted.

Using Serviceguard primary and standby links is the preferred HA model to provide HA for the cluster communications interconnect network HA. With redundancy through Serviceguard primary and standby, Serviceguard monitors the network and performs local failover if the primary network becomes unavailable.

General principles

It is preferred to have all interconnect traffic for cluster communications on a single heartbeat network with a standby interface so that Serviceguard will monitor the network and resolve interconnect failures by cluster reconfiguration. This preferred configuration is the recommended common configuration.

The following examples are instances when it is not possible to place all interconnect traffic on the same network:

- RAC GCS (cache fusion) traffic may be very high, so a separate network for RAC-DB-IC may be needed.[4] One RAC-DB-IC may interfere with another RAC-DB-IC on the same cluster. RAC-DB-IC may also interfere with heartbeat traffic.
- Some networks are not supported by CFS/CVM, so the RAC-DB-IC traffic may be on a separate network.
- Certain configurations for fast re-configuration require at least SG heartbeat network, while CSS-HB and RAC-DB-IC do not support multiple network for HA purposes.

Note:
Starting with A.11.19, the faster failover capability is integrated as part of the base Serviceguard product.

In these cases, you may see a longer time to recover some network failures beyond those protected by primary and standby, unless special logic is developed.

[4] See Administering Multiple Cluster Interconnects on Linux and UNIX Platforms, page 3-16, Oracle Real Application Clusters Administration and Deployment Guide 11g Release 1 (11.1) B28254-04, November 2007 (http://download.oracle.com/docs/cd/B28359_01/rac.111/b28254.pdf)

Cluster interconnect configurations for SLVM

Configurations with SLVM have configurations that are combinations of SG-HB, CSS-HB and RAC-DB-IC on Ethernet. The following figures show several examples on how cluster interconnect traffic can be distributed. This is not an exhaustive list.

Figure 1. Preferred: SG-HB, CSS-HB, and RAC-DB-IC on same subnet

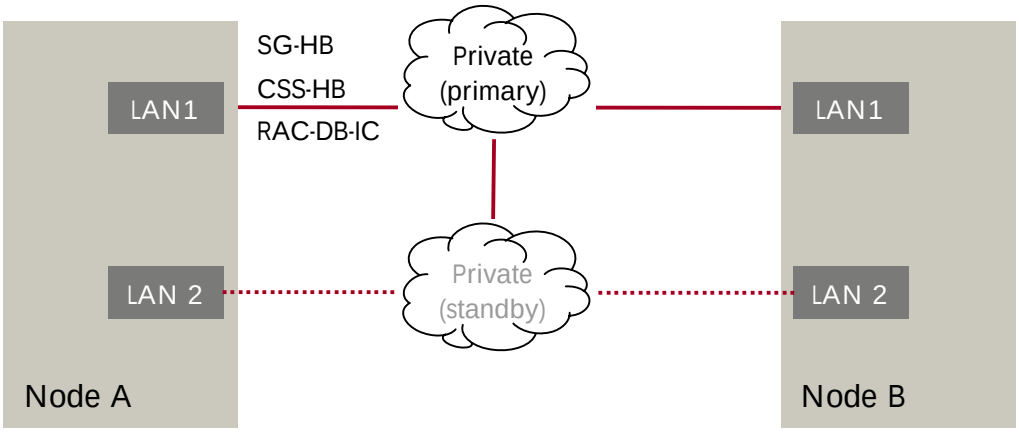

Figure 1 shows a common configuration where all inter-cluster traffic flows through the primary network. This configuration is suitable for most common installations where the RAC traffic does not saturate the network and negatively affect other components. There is one network and the network has sufficient bandwidth. If the primary network fails, Serviceguard performs a local LAN failover to use the standby network. Node failure is detected when Serviceguard misses heartbeats.

Configurations with heavy RAC-DB-IC traffic may place a limit on how aggressive the Serviceguard heartbeat timeout can be used since SG-HB may not be processed in time. Therefore, a longer Serviceguard heartbeat timeout may be needed to avoid false cluster reconfigurations.

Figure 2. Dual SG-HB with CSS-HB and RAC-DB-IC on same subnet

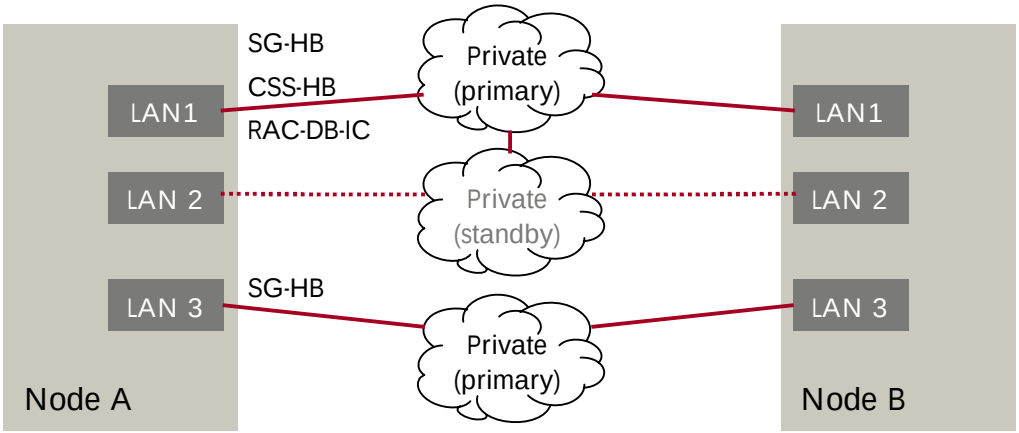

Figure 2 shows one method to overcome the heavy RAC-DB-IC traffic affecting SG-HB. This configuration uses two SG-HB networks. Two SG HB networks are required in faster failover configurations. The Serviceguard heartbeat timeout can be as aggressive as the systems allow. The subnet for CSS-HB and RAC-DB-IC has redundancy with primary and standby adapters and switches configured. If the primary (LAN1) fails, Serviceguard performs a local LAN failover to the standby (LAN2). If a node fails, Serviceguard detects the failure from the loss of the SG-HB.

If both primary (LAN1) and standby (LAN2) fail, Serviceguard logs the failure but will not take action unless Serviceguard packages with monitored subnets are configured to take action, for example node fail fast. The advantage for Serviceguard to take action on a monitored subnet is that the failure detection can be faster than the CSS-HB timeout, and thus recovery action can be quicker. Use of Serviceguard subnet monitoring has a limitation where if all interconnects fail (for example primary and standby switch failed at the same time), all the nodes are halted. If there is a concern with simultaneous failure of both switches, starting with SGeRAC A.11.18, Serviceguard supports cluster interconnect subnet monitoring.

If there are no Serviceguard packages configured to take action, CSS-HB timeout occurs and CSS goes through reconfiguration and reboot the node that is not part of the CSS cluster membership. When the node reboots, Serviceguard will reform with the new membership.

Figure 3. Single SG-HB with CSS-HB and RAC-DB-IC on separate subnet

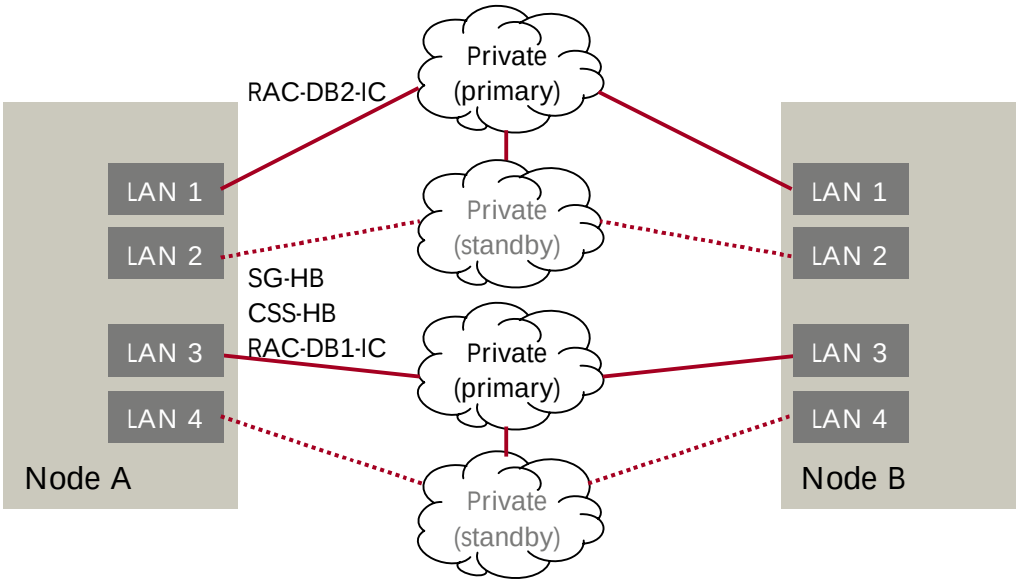

Figure 3 is a variation of figure 2 and shows the CSS-HB and RAC-DB1-IC residing on the same subnet as SG-HB. The RAC-DB2-IC is for a second database and is on a separate network and thus do not affect the HB and RAC-DB1-IC traffic. If the primary (LAN1) fails, Serviceguard performs local LAN failover. If both primary (LAN1) and standby (LAN2) fail, RAC Instance Membership Recovery (IMR) reforms and evicts suspect nodes. The IMR timeout is by default longer than the Serviceguard reconfiguration time and CSS reconfiguration time.

The advantage of this configuration is that the second RAC database instance traffic is separate from heartbeat traffic. The RAC traffic does not interfere with heartbeat traffic. A Serviceguard package

can be configured to monitor the RAC-DB-IC subnet. If the RAC-DB2-IC subnet fails (both primary and standby), the Serviceguard package can be configured to shutdown the RAC instance in order to avoid RDBMS IMR timeout, and only the specific RAC instance is affected. This configuration allows halting the RAC instance rather than evicting and halting the node. Use of Serviceguard subnet monitoring has a limitation where if all interconnect fails (for example primary and standby switch failed at the same time), all the instances are halted. If there is a concern with simultaneous failure of both switches, starting with SGeRAC A.11.18, Serviceguard supports cluster interconnect subnet monitoring.

Cluster interconnect configurations for CFS and CVM

Configurations with CFS and CVM add GAB/LLT as an additional set of traffic.

Figure 4. Preferred: single subnet with Ethernet primary and standby including GAB/LLT

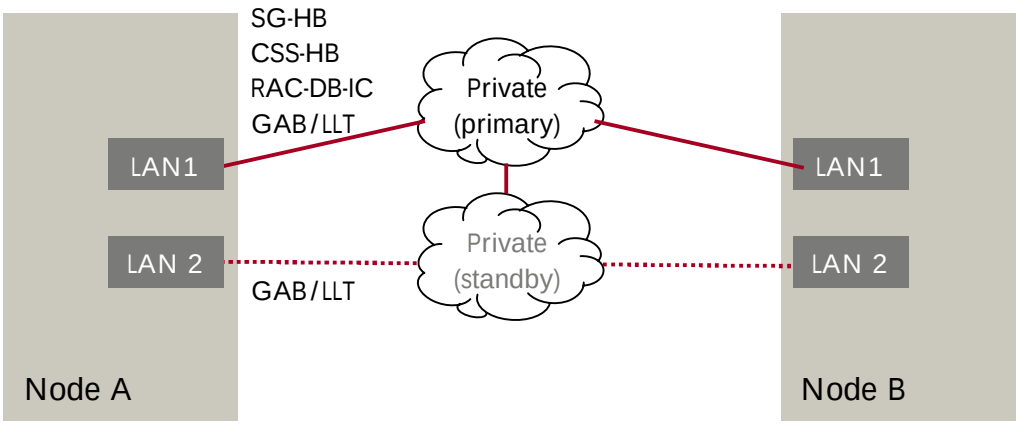

Figure 4 shows a common configuration where SG-HB, CSS-HB, RAC-DB-IC, and GAB/LLT share the same network for cluster communications. This configuration is common for most installations where the RAC traffic does not saturate the network and negatively affect other components (SG heartbeat, CSS heartbeat, and GAB/LLT traffic.) There is one network and the network has sufficient bandwidth. A failure of primary is protected by Serviceguard local LAN failover for SB-HB, CSS-HB and RAC-DB-IC. GAB/LLT uses both primary (LAN1) and standby (LAN2) for communications, so GAB/LLT can tolerate a failure of either the primary or standby.

Figure 5. Dual SG-HB with CSS-HB and RAC-DB-IC on single subnet including GAB/LLT

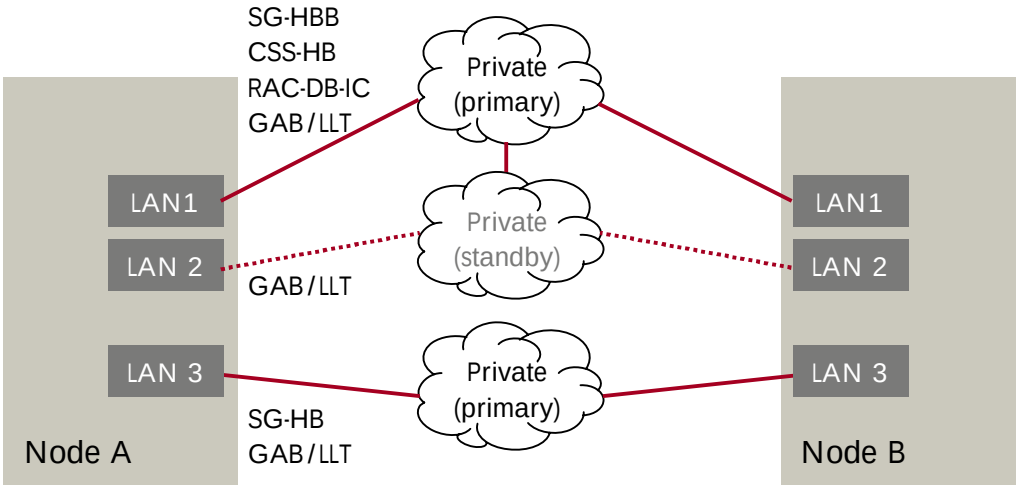

Figure 5 shows a configuration to address the possibility that heavy RAC-DB-IC traffic affects SG-HB. The pros and cons are same as figure 3.

Figure 6. Dual primary and standby Ethernet including GAB/LLT

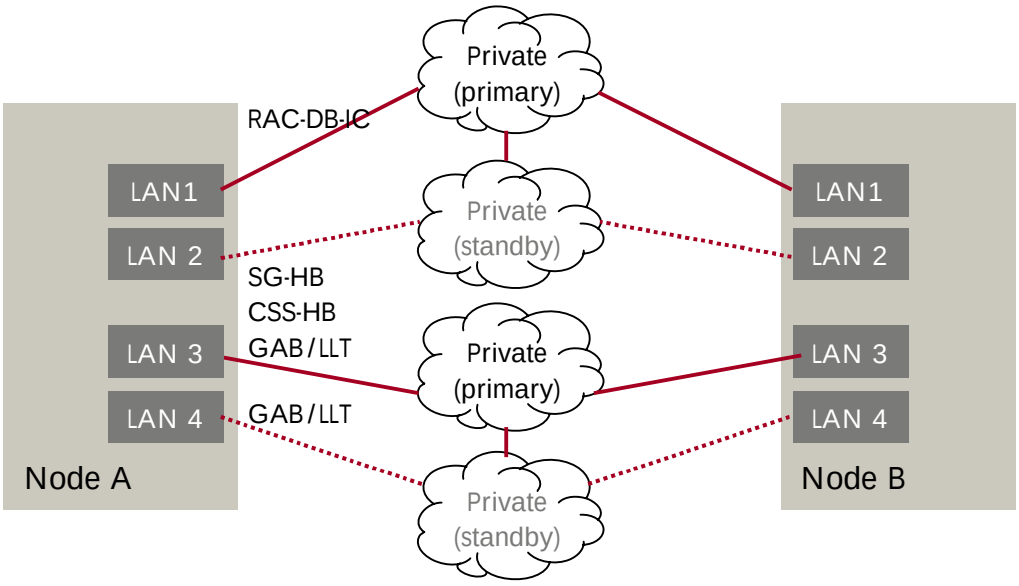

Figure 6 shows the same variation as figure 3 except this configuration is for CFS and CVM. This configuration is for heavily loaded configurations where RAC-DB-IC traffic interferes with heartbeats and other cluster communications. Placing RAC-DB-IC traffic on a separate network allows more aggressive SG-HB timeout values. The drawback is that the RAC-DB-IC should be monitored so network failures can be detected sooner than the Instance Membership Recovery (IMR) timeout and recovery can be started earlier.

Storage high availability

Storage HA is available at several levels, as follows:

- Redundant links to the same disk device through multi-pathing.
- Storage arrays provide redundancy at the disk level.
- Volume Manager mirroring to multiple devices, for example by SLVM or CVM.
- Multiple copies on multiple disks.

OC relies on the platform to provide transparent redundant links to the same device. OC provides redundancy for the Oracle Cluster Registry (OCR) by providing mirroring capability, and thus disks from JBODs ("just a bunch of disk") can be used. When JBODs are used and the disks themselves are not highly available, two physical disks are required to protect against link failure and disk failure for OCR, and three voting disks each on a separate a physical disk are required.

When redundant links and storage arrays are used, the configuration of OC is simplified by configuring a single OCR and a single voting disk with HA provided by SLVM, CVM and/or storage array.

Multiple nodes

Multiple nodes protect against failures at the system node level. For cluster HA, a minimum of two nodes is required.

Power

Redundant components should be separately powered so that a single power failure does not impact all nodes, all switches, and/or all storage.

Storage

Oracle Clusterware (OC) assumes the required storage is available when OC starts. Therefore, OC does not perform any storage activation and leaves it up to the platform or users to activate the storage prior to starting OC. For SGeRAC configurations, Serviceguard packages are used as the mechanism to activate storage prior to starting OC. For SLVM and CVM configurations, the shared storage activation is performed by the Serviceguard package that starts OC. For CFS configurations, the shared storage activation is performed by multi-node packages; therefore, the Serviceguard package that starts OC must have a dependency on the relevant multi-node packages.

Oracle Clusterware requirements: OCR and voting disk

Oracle Cluster Registry (OCR)

Each OCR requires 256 MB of disk space.[5] The OCR must be shared and accessible by all cluster nodes. OC uses the OCR to keep Oracle cluster information and configuration information regarding cluster databases. It is also used to keep track of processes that the Oracle Clusterware controls.

Voting disk

Each voting disk requires 256 MB of disk space.[6] As with the OCR, the voting disk is shared and accessible by all cluster nodes. Oracle uses the voting disk to manage Oracle cluster software membership. The voting disk is used as a health check device; in case of network failure, it is used to arbitrate cluster ownership among the instances.

Shared storage

Each OCR and voting disk can exist as follows:

[5] Configuring Storage for Oracle Clusterware Files on Raw Devices, page 1-1, 4-7, Oracle Clusterware Installation Guide 11g Release 1 (11.1) B28259-05, November 2007 (http://download.oracle.com/docs/cd/B28359_01/install.111/b28259.pdf)

[6] See footnote above.

- SLVM – The file can reside on SLVM as a raw logical volume where the whole logical volume is used as the OCR or voting disk.
- CFS – The file can reside on a cluster file system as a regular file. In this case, it is not really a disk in the traditional sense of a physical device.
- CVM – The file can reside on CVM as a raw volume similar to SLVM raw logical volume.

The advantages include multi-path support, mirroring, and controlled access.

RAC instance data files

With SGeRAC, RAC instance data files may reside on SLVM or CVM raw volumes, CFS, ASM over SLVM, ASM over raw disks, and raw disks.

Oracle binaries (Clusterware and RAC)

Oracle binaries may reside on local file system or cluster file system (CFS). Oracle Clusterware requires 650 MB and Oracle RAC software requires 4 GB of disk space[7].

Archive log files

Oracle archive log files may be needed for database recovery and should be available to the node that performs database recovery. The archive logs may reside on cluster file system (CFS), or a file system that can fail over within a Serviceguard package.

Prerequisites

In the sample configurations, the following prerequisites apply:

Software

- HP-UX 11i v2 0505 Enterprise Operating Environment with June 2007 updates
- Serviceguard A.11.17 or A.11.18 or A.11.19
- Serviceguard Extension for RAC A.11.17 or A.11.18 or A.11.19
- HP Serviceguard Management Suites Bundles A.01.00 or later.
- Oracle Clusterware and RAC 11g Release 1

Server hardware

- One two-node cluster for SLVM (node names: “eenie” and “meenie”)
- One two-node cluster for CVM/CFS (node names: “mo” and “minie”)

Network

- Public – Ethernet with two redundant NICs for primary and standby
    - OC requires one VIP address per node

[7] See Checking the Hardware Requirements, page 2-7, Oracle Clusterware Installation Guide 11g Release 1 (11.1) for HP-UX, November 2007 (http://download.oracle.com/docs/cd/B28359_01/install.111/b28259.pdf)

    - HA is provided by Serviceguard local LAN failover.
- Private –Ethernet with redundant NICs for primary and standby
    - All private cluster communications flow through the private network.
    - HA is provided by Serviceguard local LAN failover.

Storage

- Each node has internal storage for OC and RAC binaries. OC Home requires 650 MB space and RAC software requires 4 GB of disk space. Additionally, 400 MB of disk space is required from `/tmp` directory.[8]
- Shared Storage:
    - Primary Link
        - `/dev/dsk/c4t0d0  /dev/rdsk/c4t0d0`
        - `/dev/dsk/c4t1d0  /dev/rdsk/c4t1d0`
        - `/dev/dsk/c4t2d0  /dev/rdsk/c4t2d0`
        - `/dev/dsk/c4t3d0  /dev/rdsk/c4t3d0`
    - Redundant Link
        - `/dev/dsk/c5t0d0  /dev/rdsk/c5t0d0`
        - `/dev/dsk/c5t1d0  /dev/rdsk/c5t1d0`
        - `/dev/dsk/c5t2d0  /dev/rdsk/c5t2d0`
        - `/dev/dsk/c5t3d0  /dev/rdsk/c5t3d0`

[8] See Oracle Clusterware Installation Guide 11g Release 1 (11.1) for HP-UX, November 2007, page 2-7 (http://download.oracle.com/docs/cd/B28359_01/install.111/b28259.pdf)

Cluster for SLVM (eenie and meenie)

Figure 7. Cluster for SLVM

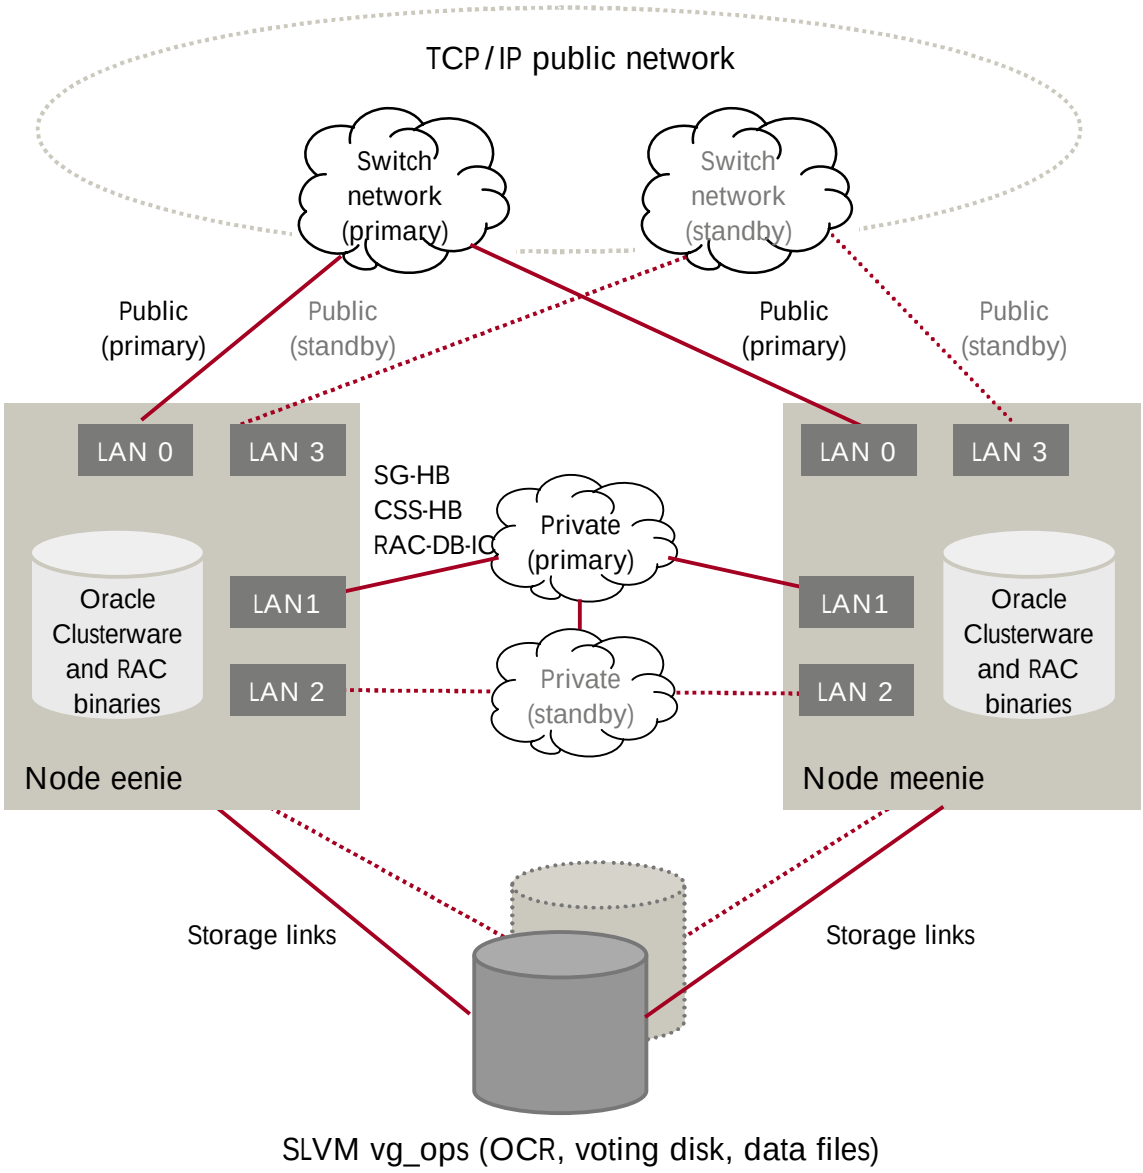

Figure 8. Cluster for CFS

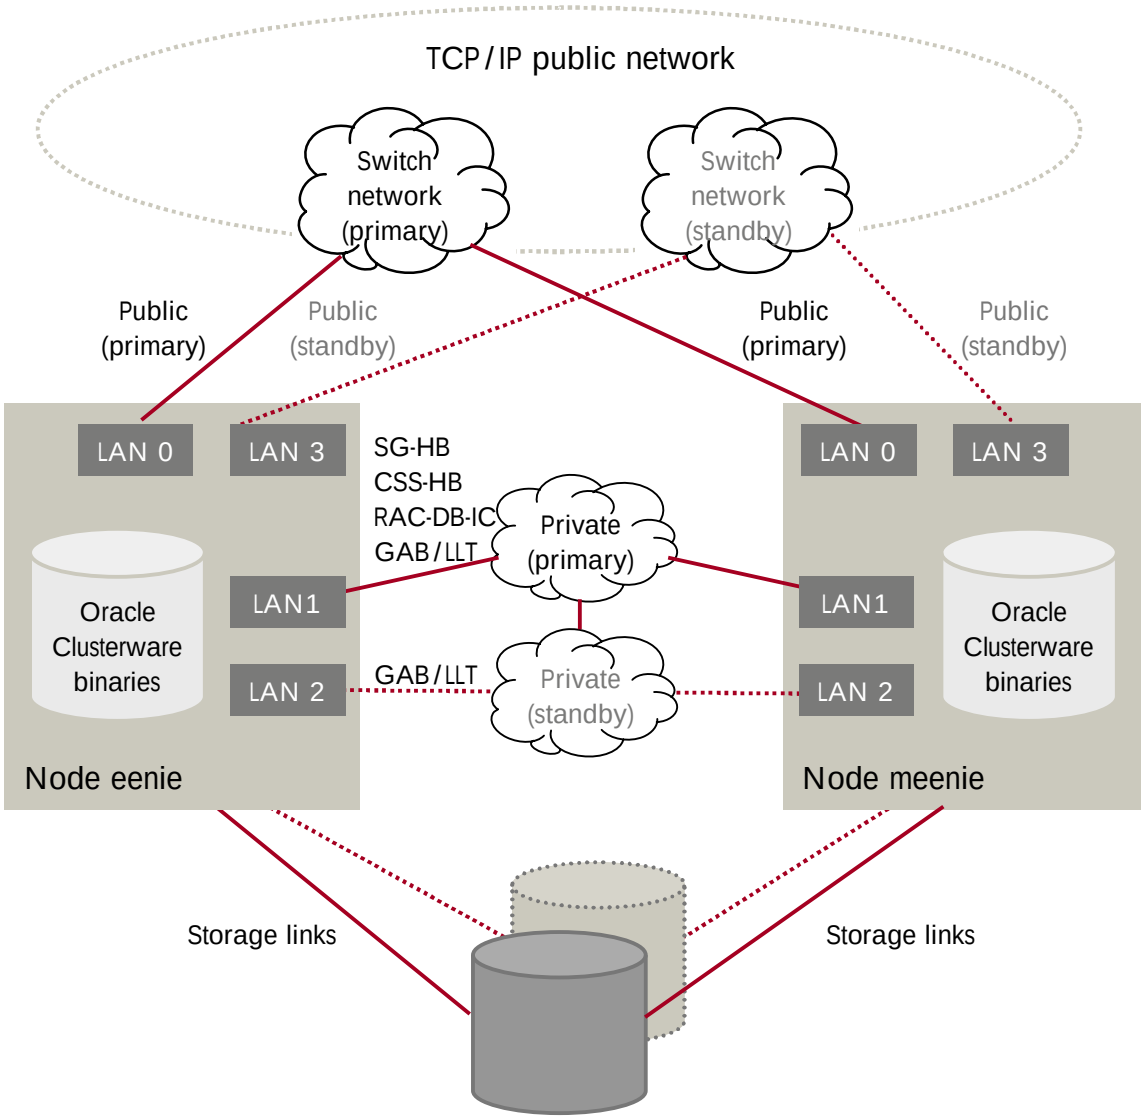

CFS (RAC binaries, OCR, voting disk, data files, flash recovery area, archive logs)

Sample Configurations

The following sections describe sample configuration steps for Oracle RAC 11g and SLVM:

Configuring Oracle RAC 11g on SLVM

The following sections describe the process for configuring Oracle RAC 11g on SLVM:

Assumptions for this sample configuration

- Cluster hardware configured.
- HP-UX 11i v2 0505 Enterprise Operating Environment with June 2007 updates.
- Serviceguard and Serviceguard Extension for RAC installed.
- Same private interconnect used for all inter-node traffic (Serviceguard, RAC, CSS)
- One shared disk for shared volume group.
- Two private disks (one disk per node for local file system for local Oracle storage)
- Kernel Parameters tuned as per Oracle documentation.

Creating a Serviceguard Extension for RAC cluster with SLVM for Oracle RAC 11g

The following sections describe the process for configuring SLVM.

SLVM logical volume planning

For the sample configuration with SLVM, the following shared logical volumes are used. Note that one SLVM volume group is used for both Oracle Clusterware (OC) required files and RAC database files. When more than one RAC instance database or when the cluster will have more than one RAC instance database, it is recommended that separate volume groups are used for OC and each RAC instance database.

Using multiple volume groups permits better maintainability, manageability, and performance. For example, with SLVM single node online volume reconfiguration (SNOR), it is possible to reconfigure SLVM volumes online while the RAC instance is running. Since SNOR requires that the target volume group be deactivated on all but one node, if only one volume group is used, RAC instances and OC on other nodes are impacted.

Figure 9. Preferred SLVM logical volume worksheet

```
                              RAW LOGICAL VOLUME NAME    SIZE (MB)
Oracle Cluster Registry: ___/dev/vg_rac/rora_ocr___256___ (one per cluster)
Oracle Cluster Voting Disk: __/dev/vg_rac/rora_vote__256___ (one per
cluster)

Oracle Control File: ____/dev/vg_rac/rorcl_control1_raw_110m__110______
Oracle Control File 2: __/dev/vg_rac/rorcl_control2_raw_110m__110______
Instance 1 Redo Log 1: __/dev/vg_rac/rorcl_redo1_1_raw_120m___120______
Instance 1 Redo Log 2: __/dev/vg_rac/rorcl_redo1_2_raw_120m___120_______
Instance 1 Redo Log: _____________________________________________________
Instance 1 Redo Log: _____________________________________________________
Instance 2 Redo Log 1: __/dev/vg_rac/rorcl_redo2_1_raw_120m____120________
Instance 2 Redo Log 2: __/dev/vg_rac/rorcl_redo2_2_raw_120m____120________
Instance 2 Redo Log: _____________________________________________________
Instance 2 Redo Log: _____________________________________________________
Data: System ___/dev/vg_rac/rorcl_system_raw_720m___720__________
Data: Sysaux ___/dev/vg_rac/rorcl_sysaux_raw_1300m___1300_________
Data: Temp ___/dev/vg_rac/rorcl_temp_raw_250m_____250_______
Data: Users ___/dev/vg_rac/rorcl_users_raw_120m____120_________
Parameter: spfile ___/dev/vg_rac//rorcl_spfile_raw_5m__5_____
Password: ______/dev/vg_rac/rorcl_pwdfile_raw_5m__5_______
Instance 1 undotbs1: /dev/vg_rac/rorcl_undotbs1_raw_500m___500___
Instance 2 undotbs2: /dev/vg_rac/rorcl_undotbs2_raw_500m___500___
Data: example1__/dev/vg_rac/rorcl_example_raw_160m________160____
```

Creating volume group and logical volumes

1. Initialize LVM disk on node (“eenie”)

   ```
   # pvcreate /dev/rdsk/c4t3d0
   ```

2. Create the volume group on node (“eenie”).

   ```
   # mkdir /dev/vg_rac
   # mknod /dev/vg_rac/group c 64 0x080000
   ```

   Note: <0x080000> is the minor number on this sample configuration.

   ```
   # vgcreate /dev/vg_rac /dev/dsk/c4t3d0
   # vgextend /dev/vg_rac /dev/dsk/c5t3d0
   ```

   Note: <c5t3d0> is a redundant link to <c4t3d0>

3. Create logical volumes on node (“eenie”). For each of the logical volumes in the worksheet, create the logical volumes.

   ```
   # lvcreate –L <size> -n <lv name> <vg name>
   ```

   Example:

   ```
   # lvcreate  -L 256 –n ora_ocr /dev/vg_rac
   ```

   Repeat this step for each logical volume specified in the worksheet.

4. Export Volume Group on node (“eenie”)

   ```
   #  vgexport –s –p –m /tmp/vg_rac.map /dev/vg_rac
   ```

5. Import Volume Group on node (“meenie”). Assuming remote shell is configured, logon on node (“meenie”) and import shared volume group.

   ```
   # rcp eenie:/tmp/vg_rac.map /tmp
   # mkdir /dev/vg_rac
   # mknod /dev/vg_rac/group c 64 0x080000
   # vgimport –s –m /tmp/vg_rac.map /dev/vg_rac
   ```

Create cluster ASCII file

```
# cd /etc/cmcluster
# cmquerycl –C clm.asc –n eenie –n meenie
```

Edit the cluster ASCII file. A two node configuration requires either a Serviceguard Quorum Server or cluster lock device. In this sample, the shared disk from the SLVM volume group is used as a cluster lock device; however it is recommended to use Quorum Server when available.

```
CLUSTER_NAME            cluster_eenie
FIRST_CLUSTER_LOCK_VG   /dev/vg_rac

NODE_NAME         eenie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP 16.89.114.170
  NETWORK_INTERFACE     lan3
```

```
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP  192.168.1.1
  NETWORK_INTERFACE     LAN2
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c4t3d0

NODE_NAME         meenie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP 16.89.114.171
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP  192.168.1.2
  NETWORK_INTERFACE     LAN2
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c4t3d0

AUTO_START_TIMEOUT      600000000
NETWORK_POLLING_INTERVAL    2000000
NETWORK_FAILURE_DETECTION   INOUT
MAX_CONFIGURED_PACKAGES     150s

OPS_VOLUME_GROUP  /dev/vg_rac
```

For A.11.18 and prior, the heartbeat timeout is

```
HEARTBEAT_INTERVAL      1000000
NODE_TIMEOUT            2000000
```

For A.11.19 and later, the heartbeat timeout is

```
MEMBER_TIMEOUT          14000000
```

Create cluster (sample)

Create and start the cluster on node.

```
# cmapplyconf –C clm.asc
# cmruncl
# cmviewcl

CLUSTER         STATUS
cluster_eenie   up

  NODE           STATUS       STATE
  eenie          up           running
  meenie         up           running
```

Activate shared storage (on each node)”

```
# vgchange –a n /dev/vg_rac
# vgchange –a s /dev/vg_rac
```

Prerequisites for Oracle 11g (sample installation)

The following are sample steps to prepare a SGeRAC cluster for Oracle 11g:

Note:
Consult Oracle documentation for Oracle installation details.

Create groups on each node

Create the Oracle Inventory group if one does not exist, create the OSDBA group, and create the Operator Group (optional).

```
# /usr/sbin/groupadd oinstall
# /usr/sbin/groupadd dba
# /usr/sbin/groupadd oper
```

Add HP-UX privileges to Oracle Inventory group members

Edit /etc/privgroup file and add the following line if it does not already exist:

```
oinstall RTPRIO MLOCK RTSCHED
```

Add these privileges to Oracle Inventory group:

```
# /usr/sbin/setprivgrp –f /etc/privgroup
```

Verify these privileges:

```
# /usr/bin/getprivgrp oinstall
oinstall: RTPRIO MLOCK RTSCHED
```

Create Oracle user on each node

In this sample configuration, the “oracle” user owns all Oracle software.

```
# /usr/bin/useradd –u 203 –g oinstall –G dba,oper oracle
```

Create oracle user home and modify oracle user to new home directory on each node.

```
# mkdir /home/oracle
# chown oracle:oinstall /home/oracle
# usermod –d /home/oracle oracle
```

Change password on each node

```
# passwd oracle
```

Enable remote access (ssh or remsh) for Oracle user on all nodes

For `remsh`, add oracle user to the `.rhosts` file or `host.equiv` file, and to the oracle user’s `.rhosts` file.

Create symbolic links

Required if Motif 2.1 Development Environment Package is not installed.

```
# ln -s /usr/lib/libX11.3 /usr/lib/libX11.sl
# ln -s /usr/lib/libXIE.2 /usr/lib/libXIE.sl
# ln -s /usr/lib/libXext.3 /usr/lib/libXext.sl
# ln -s /usr/lib/libXhp11.3 /usr/lib/Xhp11.sl
# ln -s /usr/lib/libXi.3 /usr/lib/libXi.sl
```

```
# ln -s /usr/lib/libXm.4 /usr/lib/libXm.sl
# ln -s /usr/lib/libXp.2 /usr/lib/libXp.sl
# ln -s /usr/lib/libXt.3 /usr/lib/libXt.sl
# ln -s /usr/lib/libXtst.2 /usr/lib/libXtst.sl
```

Create file system for Oracle directories

In the following samples, `/u01/app` is a mounted file system for Oracle software. Assume there is a private disk c2t0d0 of 18 GB size on all nodes. Create the local file system on each node.

```
# umask 022
# pvcreate /dev/rdsk/c2t0d0
# mkdir /dev/vg01
# mknod /dev/vg01/group c 64 0x010000
# vgcreate /dev/vg01 /dev/dsk/c2t0d0
# lvcreate –L 16000 /dev/vg01
# newfs –F vxfs /dev/vg01/rlvol1
# mkdir –p /u01/app
# mount /dev/vg01/lvol1 /u01/app
# chmod 775 /u01/app
# chown oracle:oinstall /u01/app
```

Create Oracle Cluster Software home directory

For installing Oracle Cluster Software on local file system, create the directories on each node.

```
# mkdir –p /u01/app/11.1.0/crs
# chown –R oracle:oinstall /u01/app/11.1.0/crs
# chmod -R 775 /u01/app/11.1.0
```

The Oracle Cluster Software home directory is as follows:

```
/u01/app/11.1.0/crs/
```

When installing Oracle Cluster Software, you should set the `ORACLE_HOME` environment to specify this directory. Please note at installation and before running the `root.sh` script, the parent directories of the Oracle Cluster Software home directory should be changed to permit only the `root` user to write to those directories.

Create Oracle base directory (for RAC binaries on local file system)

If installing RAC binaries on a local file system, create the oracle base directory on each node.

```
# mkdir –p /u01/app/oracle
# chown –R oracle:oinstall /u01/app/oracle
# chmod –R 775 /u01/app/oracle
```

Prepare shared storage on SLVM

This section assumes the OCR, voting disk, and database files are created on SLVM volume group `vg_rac`.

Change permission of shared logical volume group

```
# chmod 755 /dev/vg_rac
```

Change permission and ownership of Oracle Cluster Software voting disk and database files

```
# chown oracle:dba /dev/vg_rac/r*
# chmod 660 /dev/vg_rac/r*
```

Change permission of OCR device

```
# chown root:oinstall /dev/vg_rac/rora_ocr
# chmod 640 /dev/vg_rac/rora_ocr
```

Change permission of voting disk

```
# chown oracle:dba /dev/vg_rac/rora_vote
# chmod 660 /dev/vg_rac/rora_vote
```

Create raw device mapping file for Oracle database configuration assistant

In this example, the database name is "orcl".

```
# ORACLE_BASE=/u01/app/oracle ; export ORACLE_BASE
# mkdir –p $ORACLE_BASE/oradata/orcl
# chown –R oracle:oinstall $ORACLE_BASE/oradata
# chmod –R 775 $ORACLE_BASE/oradata
```

The following is a sample of the mapping file for DBCA.

```
system=/dev/vg_rac/rorcl_system_raw_720m
sysaux=/dev/vg_rac/rorcl_sysaux_raw_1300m
undotbs1=/dev/vg_rac/rorcl_undotbs1_raw_500m
undotbs2=/dev/vg_rac/rorcl_undotbs2_raw_500m
temp=/dev/vg_rac/rorcl_temp_raw_250m
example=/dev/vg_rac/rorcl_example_raw_160m
users=/dev/vg_rac/rorcl_users_raw_120m
redo1_1=/dev/vg_rac/rorcl_redo1_1_raw_120m
redo1_2=/dev/vg_rac/rorcl_redo1_2_raw_120m
redo2_1=/dev/vg_rac/rorcl_redo2_1_raw_120m
redo2_2=/dev/vg_rac/rorcl_redo2_2_raw_120m
control1=/dev/vg_rac/rorcl_control1_raw_110m
control2=/dev/vg_rac/rorcl_control2_raw_110m
spfile=/dev/vg_rac/rorcl_spfile_raw_5m
pwdfile=/dev/vg_rac/rorcl_pwdfile_raw_5m
```

In this sample, create the DBCA mapping file with user oracle and group dba ownership having read and write permission and place at:

```
/u01/app/oracle/oradata/orcl/orcl_raw.conf.
```

Installing and configuring Oracle 11g Clusterware on local file system

Login as "oracle" user.

```
$ export DISPLAY=<display>:0.0
$ cd <11g Cluster Software disk directory>
$ ./runInstaller
```

Note:

1. Specify Inventory directory as `/u01/app/oraInventory` and Operating System group name as "oinstall"
2. Specify CRS HOME as `/u01/app/11.1.0/crs`. This is a local file system.
3. Supply the VIP addresses
    a. Oracle clusterware requires one VIP address for each node.
4. Specify the public network and private network
    a. In this sample, the private network is 192.168.1.0

5. Specify Oracle Cluster Registry (OCR) Location
    a. For multi-path disk arrays, use External Redundancy and Specify OCR Location as "/dev/vg_rac/rora_ocr"
6. Specify voting disk location
    a. For multi-path disk arrays, use External Redundancy and Specify Voting Disk Location as "/dev/vg_rac/rora_vote"
7. When prompted, run `orainstRoot.sh` on each node
8. When prompted, run `root.sh` on each node

When Oracle Clusterware is installed, the Oracle cluster is also created.

Installing Oracle RAC 11g on local file system

Login as "oracle" user.

```
$ export ORACLE_BASE=/u01/app/oracle
$ export DISPLAY=<display>:0.0
$ cd <Oracle RAC 11g installation disk>
$ ./runInstaller
```

Note:

1. In this example, the path to ORACLE_HOME is on a local file system `/u01/app/oracle/product/11.1.0/db_1`
2. Select node names for cluster (eenie, meenie).
3. Select installation for database software only.
4. When prompted, run `root.sh` on each node

Creating a RAC demo database on SLVM

Export environment variables for "oracle" user.

```
export ORACLE_BASE=/u01/app/oracle
export DBCA_RAW_CONFIG=/u01/app/oracle/oradata/orcl/orcl_raw.conf

export ORACLE_HOME=$ORACLE_BASE/product/11.1.0/db_1
export ORA_CRS_HOME=/u01/app/11.1.0/crs

LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib:$ORACLE_HOME/rdbms/lib
SHLIB_PATH=$ORACLE_HOME/lib32:$ORACLE_HOME/rdbms/lib32
export LD_LIBRARY_PATH SHLIB_PATH
export PATH=$PATH:$ORACLE_HOME/bin:$ORA_CRS_HOME/bin:/usr/local/bin:
CLASSPATH=$ORACLE_HOME/jre:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib:$ORA
CLE_HOME/network/jlib
export CLASSPATH

export DISPLAY={display}:0.0
```

Setting up listeners with Oracle Network Configuration Assistant

```
$ netca
```

Notes:

5. Select Cluster Configurations
6. Select all nodes
7. Select Listener configuration
8. Select Add
9. Provide Listener name

10. Select Protocols
11. Select TCP/IP port number for listener

Creating demo database with Database Configuration Assistant

$ dbca

Unless specified, the default options are used.
Notes:

1. Unless specified, the default options are used.
2. Select Oracle Real Application Clusters database
3. Select Create a Database
4. Select all nodes
5. Select General Purpose template
6. Provide Global Database Name
    a. In this sample, the global database name and SID prefix are “orcl”.
7. Select Management Options
    a. In this sample, no management options chosen.
8. Provide passwords for user accounts
9. Select Listeners to register database.
    a. In this sample, the listeners used are “LISTENER_EENIE” and “LISTENER_MEENIE”.
10. Select Storage Options
    a. In this sample, Select the storage option for Raw Devices
11. Provide Raw Device Mapping File Location
    a. In this sample, the file is located at `/u01/app/oracle/oradata/orcl/orcl_raw.conf`.
12. Choose Recovery Configuration
    a. In this sample, use default parameters (no flash recovery and archiving.).
    b. Flash Recovery Area and archiving can be configured. When configuring archiving, choose Enable Archive Mode Parameter and specify where to place archive logs. If Flash Recovery Area is configured, archive logs default to the Flash Recovery area.
    c. Without a Cluster File System, for simplicity, the archives logs should be on a file system that can be accessed by any node that would be performing a database recovery.
13. Select Database Content
14. Configure Database Services
15. Configure Initialization Parameters
16. Select Security Settings
17. Configure Automatic Maintenance Tasks
18. Configure Database Storage as needed
19. Create Database

Configuring Serviceguard packages using Serviceguard Extension for RAC

It is recommended to start and stop Oracle Cluster Software in a Serviceguard package, as that will ensure that Oracle Cluster Software will start after SGeRAC is started and will stop before SGeRAC is halted. Serviceguard packages should also be used to synchronize storage activation and deactivation with Oracle Cluster Software and RAC instances. Additionally, the Serviceguard package also checks CSS in case CSS is halted outside the package. If CSS is halted outside the package, the package halts.

For SGeRAC A.11.17, there is the Use of Serviceguard Extension For RAC Toolkit with Oracle RAC 10g Release 2 or later White Paper (http://docs.hp.com/en/8987/sgeractoolkit-wp.pdf) to configure packages with multi-node packages and simple dependency features. The SGeRAC Toolkit

can be downloaded from the HP Software Depot (http://software.hp.com/ à High availability à Serviceguard Extension for RAC Toolkit).

SGeRAC A.11.18 and later includes the SGeRAC Toolkit along with a README document describing the use of SGeRAC Toolkit.

Note:
The SGeRAC Toolkit version A.11.18 requires a patch update before it can be used for Oracle Clusterware and RAC 11g Release 1.

Prepare Oracle Cluster Software for Serviceguard packages

Login as "root" user.

```
# export ORA_CRS_HOME=/u01/app/11.1.0/crs
# export PATH=$PATH:$ORA_CRS_HOME/bin
```

Stop Oracle Clusterware on each node

For 11g 11.1.0.6 or later:

```
# crsctl stop crs
```

Wait until Oracle Cluster Software completely stops. (Check CRS logs or check for Oracle processes, for example `ps –ef | grep ocssd.bin`)

Change Oracle Cluster Software from starting at boot time on each node

For 11g 11.1.0.6 or later:

```
# crsctl disable crs
```

Creating Serviceguard packages

In this configuration, the cluster is configured with one Serviceguard package multi-node package that will start and stop Oracle Clusterware.

Creating the Serviceguard package for Oracle Clusterware

1. Create package directory and copy toolkit files.

```
# mkdir /etc/cmcluster/crsp-slvm
# cd /etc/cmcluster/crsp-slvm
# cp /opt/cmcluster/SGeRAC/toolkit/crsp/* ./
```

2. Create package files

```
# cmmakepkg –p crsp-slvm.conf
# cmmakepkg –s crsp-slvm.ctl
```

3. Edit the package configuration file `crsp-slvm.conf`.

```
PACKAGE_NAME                    crsp-slvm
PACKAGE_TYPE                    MULTI_NODE
#FAILOVER_POLICY                CONFIGURED_NODE
#FAILBACK_POLICY                MANUAL
NODE_NAME                       eenie
NODE_NAME                       meenie
```

```
RUN_SCRIPT                      /etc/cmcluster/crsp-slvm/crsp-slvm.ctl
HALT_SCRIPT                     /etc/cmcluster/crsp-slvm/crsp-slvm.ctl
SERVICE_NAME                    crsp-slvm-srv
SERVICE_FAIL_FAST_ENABLED       NO
SERVICE_HALT_TIMEOUT            300
```

4. Edit the package control script `crsp-slvm.ctl`.

```
VGCHANGE="vgchange -a s"
VG[0]="vg_rac"
SERVICE_NAME[0]="crsp-slvm-srv"
SERVICE_CMD[0]="/etc/cmcluster/crsp-slvm/toolkit_oc.sh check"
SERVICE_RESTART[0]=""

function customer_defined_run_cmds
{
# ADD customer defined run commands.
       /etc/cmcluster/crsp-slvm/toolkit_oc.sh start
        test_return 51
}

function customer_defined_halt_cmds
{
# ADD customer defined halt commands.
       /etc/cmcluster/crsp-slvm/toolkit_oc.sh stop
        test_return 52
}
```

5. Edit the toolkit configuration file `oc.conf`.

```
ORA_CRS_HOME=/u01/app/11.1.0/crs
```

6. Add the package to the cluster. Distribute Oracle Clusterware multi-node package (MNP) directory to all nodes.

```
# cd /etc/cmcluster
# rcp –r crsp-slvm root@meenie:/etc/cmcluster
```

Add package to cluster.

```
# cd /etc/cmcluster/crsp-slvm
# cmapplyconf –P crsp-slvm.conf
Begin package verification…

Modify the package configuration ([y]/n)? y
Completed the cluster update
```

Starting and stopping Serviceguard packages and Oracle RAC

On each node, halt Oracle Clusterware if running.

```
# $ORA_CRS_HOME/bin/crsctl stop crs
```

Start the complete stack by running the Serviceguard package.

```
# cmrunpkg crsp-slvm
Running package crsp-slvm on node eenie
Successfully started package crsp-slvm on node eenie
Running package crsp-slvm on node meenie
Successfully started package crsp-slvm on node meenie
cmrunpkg: All specified packages are running
```

```
# cmviewcl
CLUSTER         STATUS
cluster_eenie   up

  NODE           STATUS       STATE
  eenie          up           running
  meenie         up           running

MULTI_NODE_PACKAGES

  PACKAGE        STATUS       STATE        AUTO_RUN     SYSTEM
  crsp-slvm      up           running      enabled      no
```

Configuring Oracle RAC 11g on CFS

The following sections describe the process for configuring Oracle RAC 11g on CFS:

Assumption for this sample configuration

1. Cluster hardware configured
2. HP-UX 11.23 0505 Enterprise Operating Environment
3. HP Serviceguard Storage Management Suite (A.02.00 or later) Installed
4. Same private interconnect used for all inter-node traffic (Serviceguard, RAC, CSS, GAB/LLT)
5. One shared disk for CFS
6. Two private disks (one disk per node for local file system for local Oracle storage)

CFS and ODM requirement

ODM is required when using Oracle RAC with CFS and SGeRAC.

Creating a Serviceguard Extension for RAC cluster with CFS for Oracle 11g

In this sample, both the Oracle RAC software and data files reside on CFS. There is a single Oracle home. Three CFS files systems are created for Oracle home, Oracle data files, and for the Oracle Cluster Registry (OCR) and vote device. The Oracle Cluster Software home is on local file system.

```
/cfs/mnt1 – for Oracle Base and Home
/cfs/mnt2 – for Oracle datafiles
/cfs/mnt3 – for OCR and Vote device.
```

Initializing the VERITAS Volume Manager

If you have not already done so, install VxVM license key on all nodes.

```
# vxinstall
```

Create Cluster ASCII file

```
# cd /etc/cmcluster
# cmquerycl –C clm.asc –n mo –n minie
```

Edit the cluster ASCII file. A two node configuration requires either a Serviceguard quorum server or cluster lock device. In this sample, since there is no SLVM volume group to be used for a cluster lock device, a quorum server is used.

```
CLUSTER_NAME            cluster_mo
QS_HOST                 white
QS_POLLING_INTERVAL     120000000
QS_TIMEOUT_EXTENSION    2000000

NODE_NAME               minie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP       15.13.170.82
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP        192.1.1.3
  NETWORK_INTERFACE     LAN2

NODE_NAME               mo
  NETWORK_INTERFACE     lan0
    STATIONARY_IP       15.13.171.137
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP        192.1.1.4
  NETWORK_INTERFACE     LAN2

AUTO_START_TIMEOUT      600000000
NETWORK_POLLING_INTERVAL 2000000
NETWORK_FAILURE_DETECTION       INOUT
MAX_CONFIGURED_PACKAGES         150
```

For A.11.18 and prior, the heartbeat timeout is

```
HEARTBEAT_INTERVAL      1000000
NODE_TIMEOUT            5000000
```

For A.11.19 and later, the heartbeat timeout is

```
MEMBER_TIMEOUT          14000000
```

Create cluster (sample)

```
# cmapplyconf –C clm.asc
```

Start the cluster:

```
# cmruncl
# cmviewcl

CLUSTER        STATUS
cluster_mo     up

  NODE           STATUS       STATE
  minie          up           running
  mo             up           running
```

Configuring Cluster Volume Manager (CVM)

Configure the system multi-node package, SG-CFS-pkg, to configure and start the CVM/CFS stack.

```
# cfscluster config -s
CVM is now configured
Starting CVM...
It might take a few minutes to complete
```

When CVM starts up, it selects a master node, and this is the node from which you must issue the disk group configuration commands. To determine the master node, issue the following command from each node in the cluster.

```
# vxdctl –c mode
mode: enabled: cluster active - SLAVE
master: minie
```

Initializing disks for CVM/CFS

You need to initialize the physical disks that will be employed in CVM disk groups. If a physical disk has been previously used with LVM, you should use the `pvremove` command to delete the LVM header data from all the disks in the volume group (this is not necessary if you have not previously used the disk with LVM)

To initialize a disk for CVM, log on to the master node, then use the `vxdiskadm` program to initialize multiple disks, or use the `vxdisksetup` command to initialize one disk at a time, as in the following example:

```
# /etc/vx/bin/vxdisksetup -i c4t1d0
```

Create disk groups for RAC

Use the `vxdg` command to create disk groups. Use the `–s` option to specify shared mode, as in the following example:

```
# vxdg –s init cfsdg1 c4t1d0
```

Create disk group multi-node package

Add the disk group to the cluster.

```
# cfsdgadm add cfsdg1 all=sw
Package name "SG-CFS-DG-1" was generated to control the resource
Shared disk group "cfsdg1" was associated to the cluster
```

Activate disk group

```
# cfsdgadm activate cfsdg1
```

Creating volumes and adding a Cluster File System

```
# vxassist -g cfsdg1 make vol1 10240m
# vxassist -g cfsdg1 make vol2 10240m
# vxassist -g cfsdg1 make vol3 300m

# newfs -F vxfs /dev/vx/rdsk/cfsdg1/vol1
    version 7 layout
    10485760 sectors, 10485760 blocks of size 1024, log size 16384 blocks
    largefiles supported

# newfs -F vxfs /dev/vx/rdsk/cfsdg1/vol2
    version 7 layout
    10485760 sectors, 10485760 blocks of size 1024, log size 16384 blocks
    largefiles supported

# newfs -F vxfs /dev/vx/rdsk/cfsdg1/vol3
    version 7 layout
    307200 sectors, 307200 blocks of size 1024, log size 1024 blocks
    largefiles supported
```

Configure mount point

```
# cfsmntadm add cfsdg1 vol1 /cfs/mnt1 all=rw
Package name "SG-CFS-MP-1" was generated to control the resource
Mount point "/cfs/mnt1" was associated to the cluster

# cfsmntadm add cfsdg1 vol2 /cfs/mnt2 all=rw
Package name "SG-CFS-MP-2" was generated to control the resource
Mount point "/cfs/mnt2" was associated to the cluster

# cfsmntadm add cfsdg1 vol3 /cfs/mnt3 all=rw
Package name "SG-CFS-MP-3" was generated to control the resource
Mount point "/cfs/mnt3" was associated to the cluster
```

Mounting Cluster File System

```
# cfsmount /cfs/mnt1
# cfsmount /cfs/mnt2
# cfsmount /cfs/mnt3
```

Check Cluster File System mount points

```
# bdf | grep cfs
/dev/vx/dsk/cfsdg1/vol1
                   10485760   36455 9796224    0% /cfs/mnt1
/dev/vx/dsk/cfsdg1/vol2
                   10485760   36455 9796224    0% /cfs/mnt2
/dev/vx/dsk/cfsdg1/vol3
                    307200    3246  284957    1% /cfs/mnt3
```

Viewing configuration

```
# cmviewcl

CLUSTER        STATUS
cluster_mo     up

  NODE           STATUS       STATE
  minie          up           running
  mo             up           running

MULTI_NODE_PACKAGES

  PACKAGE        STATUS       STATE        AUTO_RUN     SYSTEM
  SG-CFS-pkg     up           running      enabled      yes
  SG-CFS-DG-1    up           running      enabled      no
  SG-CFS-MP-1    up           running      enabled      no
  SG-CFS-MP-2    up           running      enabled      no
  SG-CFS-MP-3    up           running      enabled      no
```

Prerequisites for Oracle 11g (sample installation)

These are sample steps to prepare a SGeRAC cluster for Oracle 11g. Consult Oracle documentation for Oracle installation details.

Create groups on each node

Create the Oracle Inventory group if one does not exist, create the OSDBA group, and create the Operator Group (optional).

```
# /usr/sbin/groupadd oinstall
# /usr/sbin/groupadd dba
# /usr/sbin/groupadd oper
```

Add HP-UX privileges to Oracle Inventory group members
Edit `/etc/privgroup` file and add the following line if it does not already exist:

```
oinstall RTPRIO MLOCK RTSCHED
```

Add these privileges to Oracle Inventory group:

```
# /usr/sbin/setprivgrp –f /etc/privgroup
```

Verify these privileges :

```
# /usr/bin/getprivgrp oinstall
oinstall: RTPRIO MLOCK RTSCHED
```

Create Oracle user on each node
In this sample configuration, the “oracle” user owns all Oracle software.

```
# /usr/bin/useradd –u 203 –g oinstall –G dba,oper oracle
```

Create oracle user home and modify oracle user to new home directory on each node.

```
# mkdir /home/oracle
# chown oracle:oinstall /home/oracle
# usermod –d /home/oracle oracle
```

Change password on each node

```
# passwd oracle
```

Enable remote access (ssh or remsh) for Oracle user on all nodes
For `remsh`, add oracle user to the `.rhosts` file or `host.equiv` file.

Create symbolic links
Required if Motif 2.1 Development Environment Package is not installed.

```
# ln -s /usr/lib/libX11.3 /usr/lib/libX11.sl
# ln -s /usr/lib/libXIE.2 /usr/lib/libXIE.sl
# ln -s /usr/lib/libXext.3 /usr/lib/libXext.sl
# ln -s /usr/lib/libXhp11.3 /usr/lib/Xhp11.sl
# ln -s /usr/lib/libXi.3 /usr/lib/libXi.sl
# ln -s /usr/lib/libXm.4 /usr/lib/libXm.sl
# ln -s /usr/lib/libXp.2 /usr/lib/libXp.sl
# ln -s /usr/lib/libXt.3 /usr/lib/libXt.sl
# ln -s /usr/lib/libXtst.2 /usr/lib/libXtst.sl
```

Create File System for Oracle directories
In the following samples, `/u01/app` is a mounted file system for Oracle software. Assume there is a private disk `c2t0d0` at 18 GB size on all nodes. Create the local file system on each node.

```
# umask 022
# pvcreate /dev/rdsk/c2t0d0
```

```
# mkdir /dev/vg01
# mknod /dev/vg01/group c 64 0x010000
# vgcreate /dev/vg01 /dev/dsk/c2t0d0
# lvcreate –L 16000 /dev/vg01
# newfs –F vxfs /dev/vg01/rlvol1
# mkdir –p /u01/app
# mount /dev/vg01/lvol1 /u01/app
# chmod 775 /u01/app
# chown oracle:oinstall /u01/app
```

Create Oracle Cluster Software home directory

For installing Oracle Cluster Software on local file system, create the directories on each node.

```
# mkdir –p /u01/app/11.1.0/crs
# chown –R oracle:oinstall /u01/app/11.1.0/crs
# chmod -R 775 /u01/app/11.1.0
```

The Oracle Cluster Software home directory is as follows:

`/u01/app/11.1.0/crs/`

When installing Oracle Cluster Software, you should set the `ORACLE_HOME` environment to specify this directory. Please note at installation and before running the `root.sh` script, the parent directories of the Oracle Cluster Software home directory should be changed to permit only the `root` user to write to those directories.

Create Oracle base directory (for RAC binaries on Cluster File System)

If installing RAC binaries on Cluster File System, create the oracle base directory once since this is CFS directory visible by all nodes. The CFS file system used is `/cfs/mnt1`.

```
# mkdir –p /cfs/mnt1/oracle
# chown –R oracle:oinstall /cfs/mnt1/oracle
# chmod –R 775 /cfs/mnt1/oracle
# chmod 775 /cfs/mnt1
```

Change directory permission on each node.

```
# chmod 775 /cfs
```

Modify oracle user to new home directory on each node.

```
# usermod –d /cfs/mnt1/oracle oracle
```

Prepare shared storage on CFS

This section assumes the OCR, Vote device, and database files are created on CFS directories. The OCR and vote device reside on `/cfs/mnt3` and the demo database files reside on `/cfs/mnt2.`

Create OCR and vote device on CFS

Create directories for OCR and vote device on Cluster File System. Run commands only on one node.

```
# chmod 755 /cfs/mnt3
# cd /cfs/mnt3
# mkdir OCR
# chmod 755 OCR
# mkdir VOTE
# chmod 755 VOTE
```

```
# chown –R oracle:oinstall /cfs/mnt3
```

Change directory permission on each node (if needed).

```
# chmod 775 /cfs
```

Create directory for Oracle demo database on CFS

Create the CFS directory to store Oracle database files. Run commands only on one node.

```
# chmod 775 /cfs/mnt2
# cd /cfs/mnt2
# mkdir oradata
# chown oracle:oinstall oradata
# chmod 775 oradata
```

Change directory permission on each node (if needed).

```
# chmod 775 /cfs
```

Installing and configuring Oracle 11g Clusterware on local file system

Login as "oracle" user.

```
$ export DISPLAY=<display>:0.0
$ cd <11g Cluster Software disk directory>
$ ./runInstaller
```

Note:

1. Specify Inventory directory as "/u01/app/oraInventory" and Operating System group name as "oinstall"
2. Specify CRS HOME as /u01/app/11.1.0/crs. This is a local file system.
3. Supply the VIP addresses
    a. Oracle clusterware requires one VIP address for each node.
4. Specify the public network and private network.
    a. In this sample, the private network is 192.1.1.0.
5. Specify OCR Location as /cfs/mnt3/OCR/ocr_file if using CFS for OCR.
6. Specify Vote Disk Location as /cfs/mnt3/VOTE/vote_file if using CFS for vote device.
7. When prompted, run orainstRoot.sh on each node.
8. When prompted, run root.sh on each node.

When Oracle Clusterware is installed, the Oracle cluster is also created. For configuring Oracle VIP in a configuration with Serviceguard Local LAN failover, see Oracle MetaLink Note:296874.1.[9]

Installing Oracle RAC 11g on CFS

Login as "oracle" user. This step installs the Oracle RAC binaries on CFS. It is recommended that the RAC binary installation and demo database create to be separate steps.

```
$ export ORACLE_BASE=/cfs/mnt1/oracle
$ export DISPLAY=<display>:0.0
$ cd <Oracle RAC 11g installation disk>
$ ./runInstaller
```

[9] Doc ID: Note:296874.1 Configuring the HP-UX Operating System for the Oracle 10g VIP at https://metalink.oracle.com/ (Oracle MetaLink account required).

Note:

1. In this example, the path to ORACLE_HOME is located on a CFS directory. `/cfs/mnt1/oracle/product/11.1.0/db_1.`
2. Select installation for database software only.
3. When prompted, run `root.sh` on each node.

Configuring ODM

ODM is required when using Oracle RAC with SGeRAC and CFS. For this sample configuration, the ODM libraries are included with the HP Serviceguard Storage Management Suite bundle for RAC.

Check ODM availability

1. Check the VRTSdbed license.

```
# /opt/VRTS/bin/vxlictest -n "VERITAS Storage Foundation for Oracle" -
f "ODM"
ODM feature is licensed
```

2. Check that the VRTSodm package is installed.

```
# swlist VRTSodm

VRTSodm                     5.0.01.01 Veritas Oracle Disk Manager
VRTSodm.ODM-KRN                   5.0.01.01 Veritas ODM kernel files
VRTSodm.ODM-MAN                   5.0.01.01 Veritas ODM manual pages
VRTSodm.ODM-RUN                   5.0.01.01 Veritas ODM commands
```

3. Check that `libodm.sl` is present.

```
# ls -lL /opt/VRTSodm/lib/libodm.sl
-r-xr-xr-x 1 root sys 81464 May 15 2006 /opt/VRTSodm/lib/libodm.sl
```

Configure Oracle to use ODM library

1. Login as Oracle user
2. Shutdown database (if running)
3. Link the Oracle Disk Manager library into Oracle home for Oracle 11g

For HP 9000 systems:

```
$ rm ${ORACLE_HOME}/lib/libodm11.sl
$ ln -s /opt/VRTSodm/lib/libodm.sl ${ORACLE_HOME}/lib/libodm11.sl
```

For Integrity systems:

```
$ rm ${ORACLE_HOME}/lib/libodm11.so
$ ln -s /opt/VRTSodm/lib/libodm.sl ${ORACLE_HOME}/lib/libodm11.so
```

Configure Oracle to stop using ODM library

1. Login as Oracle user
2. Shutdown database (if running)
3. Link the original Oracle Disk Manager library into Oracle home for Oracle 11g

For HP 9000 systems:

```
$ rm ${ORACLE_HOME}/lib/libodm11.sl
$ ln -s ${ORACLE_HOME}/lib/libodmd11.sl ${ORACLE_HOME}/lib/libodm11.sl
```

For Integrity systems:

```
$ rm ${ORACLE_HOME}/lib/libodm11.so
$ ln -s ${ORACLE_HOME}/lib/libodmd11.so ${ORACLE_HOME}/lib/libodm11.so
```

Creating RAC demo database on Cluster File System

Export environment variables for "oracle" user.

```
export ORACLE_BASE=/cfs/mnt1/oracle

export ORACLE_HOME=$ORACLE_BASE/product/11.1.0/db_1
export ORA_CRS_HOME=/u01/app/11.1.0/crs

LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib:$ORACLE_HOME/rdbms/lib
SHLIB_PATH=$ORACLE_HOME/lib32:$ORACLE_HOME/rdbms/lib32
export LD_LIBRARY_PATH SHLIB_PATH
export PATH=$PATH:$ORACLE_HOME/bin:$ORA_CRS_HOME/bin:/usr/local/bin:
CLASSPATH=$ORACLE_HOME/jre:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib:$ORA
CLE_HOME/network/jlib
export CLASSPATH

export DISPLAY={display}:0.0
```

Setting up listeners with Oracle Network Configuration Assistant

```
$ netca
```

Notes:

1. Select Cluster Configurations
2. Select all nodes
3. Select Listener configuration
4. Select Add
5. Provide Listener name
6. Select Protocols
7. Select TCP/IP port number for listener

Creating demo database with Database Configuration Assistant

```
$ dbca
```

Unless specified, the default options are used.

Notes:

1. Unless specified, the default options are used.
2. Select Oracle Real Application Clusters database.
3. Select Create a Database.
4. Select all nodes.
5. Select General Purpose template.
6. Provide Global Database Name.
    a. In this sample, the global database name and SID prefix are "orcl".
7. Select Management Options.
    a. In this sample, no management options chosen.
8. Provide passwords for user accounts.
9. Select Listeners to register database.
    a. In this sample, the listeners used are "LISTENER_MO" and "LISTENER_MINIE".

10. Select Storage Options.
    a. In this sample, Select the storage option for Cluster File System.
11. Provide Database File Locations.
    a. In this sample, choose "Use Common Location for all Database Files" and enter `/cfs/mnt2/oradata` as the common directory.
12. Choose Recovery Configuration.
    a. In this sample, use default parameters (Flash Recovery Area configured) and select Enable Archiving.
    b. If Flash Recovery Area is configured, archive logs default to the Flash Recovery area.
    c. With CFS, the archives can be on a file system that is accessed by any node that would be performing a database recovery.
13. Select Database Content
14. Configure Database Services.
15. Configure Initialization Parameters.
16. Select Security Settings
17. Configure Automatic Maintenance Tasks
18. Configure Database Storage as needed
19. Create Database

Configuring Serviceguard packages with Serviceguard Extension for RAC

It is recommended to start and stop Oracle Cluster Software in a Serviceguard package, as that will ensure that Oracle Cluster Software will start after SGeRAC is started and will stop before SGeRAC is halted. Serviceguard packages should also be used to synchronize storage activation and deactivation with Oracle Cluster Software and RAC instances. Additionally, the Serviceguard package also checks CSS in case CSS is halted outside the package. If CSS is halted outside the package, the package halts.

For SGeRAC A.11.17, there is the Use of Serviceguard Extension For RAC Toolkit with Oracle RAC 10g Release 2 or later White Paper (http://docs.hp.com/en/8987/sgeractoolkit-wp.pdf) to configure packages with multi-node packages and simple dependency features. The SGeRAC Toolkit can be downloaded from the HP Software Depot (http://software.hp.com/ à High availability à Serviceguard Extension for RAC Toolkit).

SGeRAC A.11.18 and later includes the SGeRAC Toolkit along with a README document describing the use of SGeRAC Toolkit.

Prepare Oracle Cluster Software for Serviceguard packages

Login as "root" user.

```
# export ORA_CRS_HOME=/u01/app/11.1.0/crs
# export PATH=$PATH:$ORA_CRS_HOME/bin
```

Stop Oracle Clusterware on each node

For 11g 11.1.0.6 or later:

```
# crsctl stop crs
```

Wait until Oracle Cluster Software completely stops. (Check CRS logs or check for Oracle processes, for example `ps –ef | grep ocssd.bin`)

Change Oracle Cluster Software from starting at boot time on each node

For 11g 11.1.0.6 or later:

```
# crsctl disable crs
```

Creating Serviceguard packages

In this configuration, the cluster is configured with one Serviceguard multi-node package that will start and stop Oracle Clusterware.

Creating the Serviceguard package for Oracle Clusterware

1. Create package directory and copy toolkit files

```
# mkdir /etc/cmcluster/crsp
# cd /etc/cmcluster/crsp
# cp /opt/cmcluster/SGeRAC/toolkit/crsp/* ./
```

2. Create package files

```
# cmmakepkg –p crsp.conf
# cmmakepkg –s crsp.ctl
```

3. Edit the package configuration file `crsp.conf`. Since the OCR, voting disk, RAC binaries, and demo database files reside on CFS, the OC package is configured to depend on the CFS packages.

```
PACKAGE_NAME                    crsp
PACKAGE_TYPE                    MULTI_NODE
#FAILOVER_POLICY                CONFIGURED_NODE
#FAILBACK_POLICY                MANUAL
NODE_NAME                       mo
NODE_NAME                       minie
RUN_SCRIPT                      /etc/cmcluster/crsp/crsp.ctl
HALT_SCRIPT                     /etc/cmcluster/crsp/crsp.ctl

DEPENDENCY_NAME                 SG-CFS-MP-1
DEPENDENCY_CONDITION            SG-CFS-MP-1=UP
DEPENDENCY_LOCATION             SAME_NODE

DEPENDENCY_NAME                 SG-CFS-MP-2
DEPENDENCY_CONDITION            SG-CFS-MP-2=UP
DEPENDENCY_LOCATION             SAME_NODE

DEPENDENCY_NAME                 SG-CFS-MP-3
DEPENDENCY_CONDITION            SG-CFS-MP-3=UP
DEPENDENCY_LOCATION             SAME_NODE

SERVICE_NAME                    crsp-srv
SERVICE_FAIL_FAST_ENABLED       NO
SERVICE_HALT_TIMEOUT            300
```

4. Edit the package control script crsp.ctl.

```
SERVICE_NAME[0]="crsp-srv"
SERVICE_CMD[0]="/etc/cmcluster/crsp/toolkit_oc.sh check"
SERVICE_RESTART[0]=""
```

```
function customer_defined_run_cmds
{
# ADD customer defined run commands.

        /etc/cmcluster/crsp/toolkit_oc.sh start
        test_return 51
}

function customer_defined_halt_cmds
{
# ADD customer defined halt commands.

        /etc/cmcluster/crsp/toolkit_oc.sh stop
        test_return 52
}
```

5. Edit the toolkit configuration file `oc.conf`.

```
ORA_CRS_HOME=/u01/app/11.1.0/crs
```

6. Add the package to the cluster. Distribute Oracle Clusterware multi-node package (MNP) directory to all nodes.

```
# cd /etc/cmcluster
# rcp –r crsp root@minie:/etc/cmcluster
```

Add package to cluster.

```
# cd /etc/cmcluster/crsp
# cmapplyconf –P crsp.conf
Modify the cluster configuration ([y]/n)? y
Completed the cluster creation
```

Starting and stopping Serviceguard packages and Oracle RAC

On each node, halt Oracle Clusterware if running.

```
# $ORA_CRS_HOME/bin/crsctl stop crs
```

Start the complete stack by running the Serviceguard package.

```
# cmrunpkg crsp
# cmviewcl

CLUSTER        STATUS
cluster_mo     up

  NODE           STATUS       STATE
  minie          up           running
  mo             up           running

MULTI_NODE_PACKAGES

  PACKAGE        STATUS       STATE        AUTO_RUN     SYSTEM
  SG-CFS-pkg     up           running      enabled      yes
  SG-CFS-DG-1    up           running      enabled      no
  SG-CFS-MP-1    up           running      enabled      no
  SG-CFS-MP-2    up           running      enabled      no
  SG-CFS-MP-3    up           running      enabled      no
  crsp           up           running      enabled      no
```

Verify Oracle Clusterware status.

```
# $ORA_CRS_HOME/bin/crsctl check crs
CSS appears healthy
CRS appears healthy
EVM appears healthy
```

Cluster start and stop

The following sections describe the process for starting and stopping Oracle Clusterware 11g.

Start and stop Oracle Clusterware 11g

Placing the start and stop of Oracle Clusterware in Serviceguard packages ensures that the shared storage required by Oracle Clusterware is available.

When halting the Serviceguard cluster ("`cmhaltcl -f`"), the package dependencies ensure that the Oracle Clusterware packages are halted first before the cluster file systems are unmounted.

When using "`crsctl stop crs`" to stop Oracle Clusterware while Oracle RAC instances are running, stopping Oracle Clusterware causes the instances to shutdown abort. This is the default behavior since other shutdown modes may not complete in time for Oracle Clusterware to stop in a timely manner.

If a different Oracle RAC instance shutdown mode (for example normal, immediate) is desired, the instances should be halted prior to initiating Oracle Clusterware shutdown.

Start and stop Oracle RAC 11g instance

For this sample configuration, the Oracle RAC instances startup and shutdown are controlled by Oracle Clusterware and not by RAC instance Multi-node package (MNP). Therefore, manually starting and stopping Oracle RAC 11g instances are not required.

Appendix

Sample configuration for SLVM with Serviceguard Extension for RAC

Cluster configuration file for cluster (eenie and meenie)

```
CLUSTER_NAME            cluster_eenie
FIRST_CLUSTER_LOCK_VG   /dev/vg_rac

NODE_NAME          eenie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP 16.89.114.170
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP  192.168.1.1
  NETWORK_INTERFACE     LAN2
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c4t3d0

NODE_NAME          meenie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP 16.89.114.171
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP  192.168.1.2
  NETWORK_INTERFACE     LAN2
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c4t3d0

AUTO_START_TIMEOUT      600000000
NETWORK_POLLING_INTERVAL    2000000
NETWORK_FAILURE_DETECTION   INOUT
MAX_CONFIGURED_PACKAGES     150

OPS_VOLUME_GROUP  /dev/vg_rac
```

For A.11.18 and prior, the heartbeat timeout is

```
HEARTBEAT_INTERVAL      1000000
NODE_TIMEOUT            2000000
```

For A.11.19 and later, the heartbeat timeout is

```
MEMBER_TIMEOUT          14000000
```

Package configuration for SLVM (Oracle Clusterware MNP)

```
PACKAGE_NAME                    crsp-slvm
PACKAGE_TYPE                    MULTI_NODE
#FAILOVER_POLICY                CONFIGURED_NODE
#FAILBACK_POLICY                MANUAL
NODE_NAME                       eenie
NODE_NAME                       meenie
RUN_SCRIPT                      /etc/cmcluster/crsp-slvm/crsp-slvm.ctl
HALT_SCRIPT                     /etc/cmcluster/crsp-slvm/crsp-slvm.ctl
SERVICE_NAME                    crsp-slvm-srv
SERVICE_FAIL_FAST_ENABLED       NO
SERVICE_HALT_TIMEOUT            300
```

Package control script for SLVM (Oracle Clusterware MNP)

```
VGCHANGE="vgchange -a s"
VG[0]="vg_rac"
SERVICE_NAME[0]="crsp-slvm-srv"
```

```
SERVICE_CMD[0]="/etc/cmcluster/crsp-slvm/toolkit_oc.sh check"
SERVICE_RESTART[0]=""

function customer_defined_run_cmds
{
# ADD customer defined run commands.
       /etc/cmcluster/crsp-slvm/toolkit_oc.sh start
        test_return 51
}

function customer_defined_halt_cmds
{
# ADD customer defined halt commands.
       /etc/cmcluster/crsp-slvm/toolkit_oc.sh stop
        test_return 52
}
```

Serviceguard Extension for RAC configuration for Oracle Clustware

The Serviceguard Extension for RAC configuration file has the following change to specify the environment variable for Oracle Clusterware HOME:

```
ORA_CRS_HOME=/u01/app/11.1.0/crs
```

Sample configuration for CFS

The following sections describe sample configurations with CFS:

Cluster configuration file for cluster (minie and mo)

```
CLUSTER_NAME            cluster_mo
QS_HOST                 white
QS_POLLING_INTERVAL     120000000
QS_TIMEOUT_EXTENSION    2000000

NODE_NAME               minie
  NETWORK_INTERFACE     lan0
    STATIONARY_IP       15.13.170.82
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP        192.1.1.3
  NETWORK_INTERFACE     LAN2

NODE_NAME               mo
  NETWORK_INTERFACE     lan0
    STATIONARY_IP       15.13.171.137
  NETWORK_INTERFACE     lan3
  NETWORK_INTERFACE     LAN1
    HEARTBEAT_IP        192.1.1.4
  NETWORK_INTERFACE     LAN2

AUTO_START_TIMEOUT      600000000
NETWORK_POLLING_INTERVAL 2000000
NETWORK_FAILURE_DETECTION       INOUT
MAX_CONFIGURED_PACKAGES         150
```

For A.11.18 and prior, the heartbeat timeout is

```
HEARTBEAT_INTERVAL      1000000
NODE_TIMEOUT            2000000
```

For A.11.19 and later, the heartbeat timeout is

```
MEMBER_TIMEOUT          14000000
```

Package configuration for CFS (Oracle Clusterware MNP)

```
PACKAGE_NAME                    crsp
PACKAGE_TYPE                    MULTI_NODE
#FAILOVER_POLICY                CONFIGURED_NODE
#FAILBACK_POLICY                MANUAL
NODE_NAME                       mo
NODE_NAME                       minie
RUN_SCRIPT                      /etc/cmcluster/crsp/crsp.ctl
HALT_SCRIPT                     /etc/cmcluster/crsp/crsp.ctl

DEPENDENCY_NAME                 SG-CFS-MP-1
DEPENDENCY_CONDITION            SG-CFS-MP-1=UP
DEPENDENCY_LOCATION             SAME_NODE

DEPENDENCY_NAME                 SG-CFS-MP-2
DEPENDENCY_CONDITION            SG-CFS-MP-2=UP
DEPENDENCY_LOCATION             SAME_NODE

DEPENDENCY_NAME                 SG-CFS-MP-3
DEPENDENCY_CONDITION            SG-CFS-MP-3=UP
DEPENDENCY_LOCATION             SAME_NODE

SERVICE_NAME                    crsp-srv
SERVICE_FAIL_FAST_ENABLED       NO
SERVICE_HALT_TIMEOUT            300
```

Package control script for CFS (Oracle Clusterware MNP)

```
SERVICE_NAME[0]="crsp-srv"
SERVICE_CMD[0]="/etc/cmcluster/crsp/toolkit_oc.sh check"
SERVICE_RESTART[0]=""

function customer_defined_run_cmds
{
# ADD customer defined run commands.

        /etc/cmcluster/crsp/toolkit_oc.sh start
        test_return 51
}

function customer_defined_halt_cmds
{
# ADD customer defined halt commands.

        /etc/cmcluster/crsp/toolkit_oc.sh stop
        test_return 52
}
```

Serviceguard Extension for RAC configuration for Oracle Clusterware

```
ORA_CRS_HOME=/u01/app/11.1.0/crs
```

Document revision history

Revision	Date	Description	Comment
1.0	January 2008	First version	
1.1	March 2009	Minor update	Update for A.11.19

For more information

www.hp.com/go/serviceguardsolutions
www.hp.com/go/sgerac

References

HP documentation

All of the following materials can be found on the HP Technical Documentation web site at http://docs.hp.com.

- HP Serviceguard Storage Management Suite Version Release Notes
  http://docs.hp.com/en/ha.html
- Serviceguard Version A.11.19 Release Notes
  http://docs.hp.com/en/ha.html
- Serviceguard Extension for RAC Version A.11.19 Release Notes
  http://docs.hp.com/en/ha.html
- Managing Serviceguard, Sixteenth Edition
  http://docs.hp.com/en/ha.html
- Using Serviceguard Extension for RAC, Sixth Edition
  http://docs.hp.com/en/ha.html
- Symantec Veritas Storage Foundation 5.0 for Oracle RAC Configuration Guide Extracts for HP Serviceguard Storage Management Suite Extracts
  http://docs.hp.com/en/T2771-90046/T2771-90046_FINAL4-24.pdf
- Symantec Veritas Storage Foundation 4.1 for Oracle RAC HP Serviceguard Storage Management Suite Extracts
  http://docs.hp.com/en/7412/T2771-90010.pdf
- Whitepaper "Use of Serviceguard Extension For RAC Toolkit with Oracle RAC 10g Release 2 or later, March 2009"
  http://docs.hp.com/en/8987/sgeractoolkit-wp.pdf
- Whitepaper “Support of Oracle RAC ASM with SGeRAC, January 2008”
  http://docs.hp.com/en/8988/ASM-SGeRAC-tk.pdf

Oracle documentation

All of the following materials can be found on the Oracle Technical Documentation web site at http://www.oracle.com/pls/db111/homepage.

- Oracle 11g Release 1 HP-UX Installation Guides
  http://www.oracle.com/pls/db111/portal.portal_db?selected=11&frame=#hp-ux_installation_guides
- Oracle Clusterware Installation Guide 11g Release 1 (11.1) for HP-UX, November 2007
  http://download.oracle.com/docs/cd/B28359_01/install.111/b28259.pdf
- Oracle Real Application Clusters Installation Guide 11g Release 1 (11.1) for Linux and UNIX, November 2007
  http://download.oracle.com/docs/cd/B28359_01/install.111/b28264.pdf
- Oracle Clusterware Administration and Deployment Guide 11g Release 1 (11.1), September 2007
  http://download.oracle.com/docs/cd/B28359_01/rac.111/b28255.pdf
- Oracle Real Application Clusters Administration and Deployment Guide 11g Release 1 (11.1), November 2007
  http://download.oracle.com/docs/cd/B28359_01/rac.111/b28254.pdf
- Oracle Database Net Service Administrator's Guide 11g Release 1 (11.1), September 2007
  http://download.oracle.com/docs/cd/B28359_01/network.111/b28316.pdf
- Oracle Database Net Services Reference 11g Release 1 (11.1), July 2007
  http://download.oracle.com/docs/cd/B28359_01/network.111/b28317.pdf

April 2009