Overview

Prior to the name change to Serviceguard Extension for RAC (SGeRAC for short) revision A.11.14.01, the Serviceguard OPS Edition product bundle was a superset that included the Serviceguard product bundle. Beginning with Serviceguard Extension for RAC A.11.14.01 on the HP Integrity servers (Intel Itanium 2 processor servers), the Serviceguard product bundle was decoupled from the SGeRAC product bundle, and the SGeRAC product is now an add on product to Serviceguard. The decoupled version of SGeRAC A.11.15 for PA RISC will also be introduced in the June 2003 Application Release. This subchapter will use the names, SGeRAC and Serviceguard OPS Edition interchangeably, but the information here applies to both equally, unless otherwise specified.

As an extension to the Serviceguard clustering product, Serviceguard Extension for RAC allows up to 8 HP 9000 Enterprise Servers to be configured as an enterprise cluster that supports up to 8 nodes with the Oracle 9i Real Application Clusters (RAC) or Oracle Parallel Server (OPS) relational databases. Support for up to 16 nodes will be available with an SLVM patch coming out mid 2003. These two products work together to provide the best aspects of HP's enterprise clusters and Oracle's relational databases: high availability, data integrity, flexibility, scalability and reduced database administration costs.

Serviceguard Extension for RAC is an important component in creating a complete and robust environment for highly available mission critical applications in a parallel database environment. In order to create a complete solution, the application environment should be designed to eliminate all single points of failure and to minimize the impact of various component failures. HP recommends that Serviceguard Extension for RAC should be used also in conjunction with other high availability products and services available from HP.

Please refer to the end of this QuickSpec for information on SGeRAC support for the HP Integrity servers. Additional Itanium based information can also be found in the Serviceguard QuickSpec.

Serviceguard Extension for RAC Software and Documentation

For HP UX 11.0 using Serviceguard OPS Edition A.11.14 or earlier:

Serviceguard OPS Edition product bundle includes the Serviceguard product filesets.

- B5161FA: License and media for one system or CPU. Each node (system) or CPU of an enterprise cluster must license one copy of Serviceguard Extension for RAC. For 16 node support, a 16 node SLVM patch must be installed.
  NOTE: SLVM patch will be released in mid 2003.
- B5158FA: Manual for Serviceguard Extension for RAC are ordered separately from the license.

For HP UX 11.11 using SGeRAC A.11.15 and HP UX 11i v1.6 using SGeRAC A.11.14 or earlier:

Serviceguard prerequisite:

- B3935DA: License and media for one system or CPU. Each node (system) or CPU of an enterprise cluster must license one copy of Serviceguard. Each node must run HP UX version 11.11.
- B3936EA: Manual for Serviceguard

Serviceguard Extension for RAC:

- T1859BA License and media for one system or CPU-Each node (system) or CPU of an enterprise cluster must license one copy of Serviceguard Extension for RAC. For 16 node support, a 16 node SLVM patch must be installed.
  NOTE: SLVM patch will be released in mid 2003.
- T1862BA Manual for Serviceguard Extension for RAC

For HP UX 11i version 2 using Serviceguard Extension for RAC A.11.15:

Serviceguard A.11.15 prerequisite:

- T1905BA: License and media for one CPU (no tier licenses available). Each CPU of an enterprise cluster must license one copy of Serviceguard. Each processor must run HP UX version 11i v2
- T1906BA: Manual for Serviceguard

Serviceguard Extension for RAC:

- T1907BA: License and media for one CPU. Each CPU of an enterprise system must license one copy of Serviceguard Extension for RAC.
- T1908BA: Manuals for Serviceguard Extension for RAC

Quick Reference Ordering Matrix for Serviceguard Extension for RAC

	HP UX 11.0	HP UX 11.11 and 11i v1.6	HP UX 11i v2
	Serviceguard Extension for RAC 11.14 and earlier	Serviceguard Extension for RAC 11.14.01 & 11.15	Serviceguard Extension for RAC 11.15
Serviceguard software and LTU	B3935DA	B3935DA	T1905BA
Serviceguard manual	B3936EA	B3936EA	T1906BA
Serviceguard Extension for RAC software and LTU	B5161FA	T1859BA (Must also order B3935DA)	T1907BA (Must also order T1905BA)
Serviceguard Extension for RAC manual	B5158FA	T1862BA	T1908BA

Oracle Real Application Cluster and Oracle Parallel Server are purchased separately from Oracle Corporation.

The following outlines the status of each Serviceguard OPS Edition A.11.XX version including the end of support:

Version	Introduction Date	Patch Phase End Date	Support Phase/ End of Support	Oracle Version Supported
A.11.09	December 1999	December 31, 2002	December 31, 2003	OPS 8.1.7 RAC 9.0.1
A.11.12	December 2000	September 30, 2001	October 31, 2003	OPS 8.1.7
A.11.13	September 2001	June 30, 2003	September 30, 2004	OPS 8.1.7 RAC 9.0.1 RAC 9.2.0
A.11.14	March 2002	December 31, 2004	March 31, 2006	OPS 8.1.7 RAC 9.0.1 RAC 9.2.0
A.11.14.01 for HP UX 11i v1.6	July 2002	December 31, 2003	December 31, 2005	RAC 9.2.0
A.11.15 for HP UX 11.11	June 2003	December 31, 2004	June 30, 2006	RAC 9.2.0
A.11.15 for HP UX 11i v2.0	June 2003	December 31, 2004	June 30, 2006	RAC 9.2.0

Definitions for the matrix:

- Introduction-The date on which software was available for sale.
- Patch Phase-HP will create patches for the software during this period.
- Support Phase-Support contracts will continue to be sold during this period. HP Response Center will continue to take calls and patches will be created for critical and serious issues related to security and data integrity.
- End of Support-End date of the Support Phase in which time HP has no further obligation to provide support or patches to the software. Customers are encouraged to upgrade to the next recommended product version.

In certain cases, a specific version of Oracle OPS or RAC may be discontinued before Serviceguard OPS Edition reaches end of support. In these instances, Serviceguard OPS Edition will end support for that version of Oracle the same time as Oracle. Please visit http://www.software.hp.com/RELEASES MEDIA for the latest information on product support.

For more information about designing and configuring enterprise clusters using Serviceguard OPS Edition, see the HP manual titled, "Configuring OPS Clusters with Serviceguard OPS Edition", (based on Serviceguard OPS Edition Version A.11.14) and the appropriate Oracle manual for the OPS revision you plan to use. You may be able to find Oracle manuals at the Oracle web site: http://docs.oracle.com/database_mp.html

Serviceguard OPS Edition Revision Compatibility Matrix

Serviceguard OPS Edition Version	Main Features of the Release	Compatible HP UX Releases	Compatible OPS/RAC Release
A.11.15	IPv6 support, multi OS quorum service support (HP UX & Linux)	11.11 and 11i v2	Please contact HP or Oracle rep for supported RAC versions
A.11.14.01	First release on Itanium Processor Family based servers	11i v1.6	9.2.0
A.11.14	Serviceguard Manager administration support. Quorum Service support. VLAN, IPv6 dual kernel stack support. Performance enhancements in control scripts. Defect fixes.	11.0 and 11.11	8.1.7, 9.0.1 and 9.2
A.11.13	Increased package count, VERITAS VxVM and CVM support, ATS support, LVM mirroring. Online node addition for OPS 8.1.6 and later. Defect fixes	11.0 and 11.11	8.1.7 and 9.01
A.11.12	Serviceguard Manager support	11.0 only	8.1.7
A.11.09	De coupling of EMS CORE fileset. Bug fixes	11.0 and 11.11	8.1.7 and 9.0.1

- Effective November 1, 2002, HP discontinued support of OPS 8.1.6 on Serviceguard OPS Edition (11.09, 11.12, 11.13, 11.14).
- Oracle will discontinue support of OPS 8.1.7 on 1/1/2004.

Serviceguard Extension for RAC supports failover of RAC and non RAC applications to an alternate node. Upon failure of a RAC cluster node, Serviceguard Extension for RAC will automatically failover the applications to another node within the RAC cluster. Any IP addresses associated with the applications for client connectivity can also be failed over.

Recent Enhancements

Version A.11.15

- Fast detection-When the clients of *cmgmsd* Oracle's primary and slave members die, cmgmsd is able to detect and take appropriate action of informing other nodes faster.
- Database and Group Membership providers enable the latest version of Serviceguard Manager to display the database instances and group membership instances.
- Rolling upgrade enables maintenance of nodes within a cluster while cluster remains up and running, minimizing planned downtime. This is a pre enabled feature for future releases of the product and not previous releases.
- It is not supported to have a cluster containing servers running in a mix of both HP UX PA RISC and Itanium. The nodes in a single cluster must be all PA RISC or all Itanium.

Version A.11.14.01

- Refer to the Serviceguard subchapter for some caveats on this version supporting the HP Integrity servers.

Version A.11.14

Serviceguard OPS Edition RAC A.11.14 supports the same configurations as previous versions, with the following qualifications:

- VERITAS Volume Manager 3.2 (and 3.5) (VxVM), including VERITAS Cluster Volume Manager (CVM) are now supported. Refer to the comparison matrix in this subchapter for restrictions and limitations of VxVM and CVM storage support. VxVM and CVM support are also fully documented in the manual *Configuring OPS Clusters with Serviceguard OPS Edition*.
- Oracle Parallel Server Version 8.0.6 is not supported with Serviceguard OPS Edition Version A.11.14. Users of OPS 8.0.6 are encouraged to migrate to a later Oracle release.
- Support for the HyperFabric2 fiber adapter product (A6386A) is provided for the following configurations:
  - Two node point to point configuration
  - Two node, three node and four node configurations with HyperFabric2 LAN interfaces linked through HyperFabric2 switches.
  - Heartbeat and local LAN failover over HyperFabric are not supported.
  - For OPS 8i, only UDP protocol is supported for Oracle DLM traffic. For Oracle 9i, UDP and HMP protocol both are supported for Oracle DLM traffic.
  - For more information, refer to the HyperFabric2 Release Notes (A6386A with option AVN)
- A total of 127 OPS instances per cluster are supported.

Version A.11.13

- VERITAS Volume Manager 3.1 (VxVM) and the VERITAS Cluster Volume Manager (CVM) are now supported in Serviceguard clusters.
- Support was added later for VxVM 3.2 and 3.5. (See table below.)
- Support for Shared storage devices for up to 16 (with SLVM patch expected to be released mid 2003) cluster nodes is provided using storage with HP SLVM, and up to 4 nodes with VERITAS Cluster Volume Manager (CVM). Previous Serviceguard OPS Edition supports up to 8 nodes maximum with SLVM.

- Online node addition for OPS 8.1.6 and later when VERITAS CVM is used.
- Support of Oracle 9i Real Applications Clusters (RAC). Currently, Oracle RAC 9.0.1 is supported.

Comparison between HP LVM, Base VERITAS Volume Manager, Full VERITAS Volume Manager, and Full VERITAS Cluster Volume Manager

	HP LVM	Base VxVM (B7961AA)	Full VxVM (B9116AA)	Full CVM (B9117AA)
Pros	• Legacy system; robust and familiar • Supports up to 8 nodes per cluster • Supports use of PV links (for multiple data paths) • Support exclusive activation as well as read only activation	• Free with HPUX 11i and later • Java based admin GUI • Striping (RAID 0) • Concatenation • Online resizing of volumes • Task monitor • Path failover support, DMP (active/passive)	Includes all of the Base VxVM features plus… • Load balancing - DMP (active/active) • Supports up to 16 nodes • Hot relocation / unrelocation • Mirroring (RAID 1) • Supports up to 32 mirrors • Mirrored Stripes (RAID 1+0) • Striped Mirrors (RAID 0+1) • Online data migration • Online relayout • RAID 5	Includes Full VxVM features plus… • Supports shared activation of disk groups Faster failover and package startup time Volume configuration propagation • Online reconfiguration for shared disk groups activated on ANY node
Cons	• Lacks flexibility or extended functionality of other volume managers • Supports only 2 nodes per cluster if software mirroring using MirrorDisk/UX is used. • PV links are active/standby only	• Cannot be used for cluster lock or HPUX root/boot disks • Limited feature set • Supports exclusive read or write activation only • Package startup delay (due to lengthy vxdg import)	• Requires purchase of additional license • Does not support shared or read only activation modes • Package startup delay • Cannot be used for cluster lock or root devices	• Requires purchase of additional license • Does not allow more than a single heartbeat subnet to be defined for the cluster. • Support for SG and SGeRAC clusters (up to 4 nodes). • Dependency on SMN package • Does not support RAID 0+1 or RAID 5 in the first release With CVM, all nodes must be connected to all shared disks • Can not be used for cluster lock or root devices

NOTE: Base (Lite) VERITAS Cluster Volume Manager (CVM) is included in Serviceguard and Serviceguard OPS Edition A.11.13 or later. The features available with this CVM version include: 4 node support for non shared activation, 2 node shared write activation for Serviceguard OPS Edition, no software mirroring, DMP (active/passive), online reconfiguration of private disk groups (DG activated on at most one node)

Configuration

This section describes basic configurations for enterprise clusters constructed using Serviceguard OPS Edition.

System Configuration Requirements

- Use only systems listed in the "Supported Servers" section of the Serviceguard subchapter. Mixing of servers in a Serviceguard OPS Edition cluster is not recommended.
- An individual node can be a uniprocessor system or a SMP system. Uniprocessor and SMP systems can be mixed within a Serviceguard OPS Edition cluster. However, HP strongly recommends that the processing power and memory capacity of all nodes be the same or similar.
- Care should be taken to ensure that nodes have sufficient CPU capacity to accept additional workload in the case of a node failure. Capacity planning and performance measurement should be done as part of the installation and verification process. Process Resource Manager (PRM) or Workload Manager (WLM) can be used in Serviceguard clusters to ensure that each package application receives the system resources it needs.
- Serviceguard OPS Edition can be configured with a maximum of 8 nodes. There can be a maximum of 4 nodes connected to a single shared SCSI bus. With the release of Serviceguard OPS Edition on A.11.13, a maximum of 16 nodes can be configured using storage with SLVM, and a maximum of 4 nodes with VERITAS Cluster Volume Manager. To support more than 8 nodes with SLVM, an SLVM patch (patch ID is TBD) is required for HP UX 11.0 and 11.11. An individual node can belong to no more than one cluster.

- When HP UX 11.0 was introduced, software mirroring (using MirrorDisk/UX) was not supported for use with Shared LVM. This meant that hardware mirroring was required for Serviceguard OPS Edition. Beginning with the introduction of the patch, PHKL_22267, software mirroring is supported for Shared LVM for two node clusters only. If you require more than two nodes in your Serviceguard OPS Cluster, then you must use a hardware mirroring solution. MirrorDisk/UX on HP-UX 11.11 and 11.22 supports Shared LVM on up to two nodes (no patches are required).
- The maximum distance between nodes will be determined by the disk bus being used. See the "Disk Configuration Requirements and Information" section below.
- Consult the Oracle OPS manual for memory sizing information.

LAN Configuration Requirements

In a Serviceguard OPS Edition cluster individual nodes send status information (heartbeat) and Distributed LockManager (DLM) traffic to each other as well as sending data to attached clients.

Serial Line (RS 232) Connection

An RS 232 connection can be used as a backup heartbeat link between two nodes. (At least one LAN network is required in addition to the RS 232 Heartbeat.) The RS 232 link can be used for heartbeat only, not for data traffic or DLM traffic. Note that only two nodes in a cluster can be connected with RS 232. For additional information, see the Serviceguard subchapter. If there are two or more LANs in the cluster, then the RS 232 heartbeat must not be used.

Data and heartbeat traffic can be carried on a single LAN.

All nodes must be attached to the heartbeat/DLM LAN and must reside on the same IP subnet (i.e. repeaters or bridges can separate nodes, but routers cannot be used to separate nodes).

In cases where the application is expected to generate heavy LAN traffic, separating the data and heartbeat/DLM onto different LANs will maximize response time for the LAN attached clients.

Systems with personality cards can make use of the LAN interface contained on the personality card for use as a primary or standby LAN.

LAN types can be mixed within a cluster. For example, a legal configuration can contain a primary and standby Ethernet LAN for carrying the heartbeat along with a primary and standby FDDI LAN for carrying data/DLM.

Disk Configuration Requirements and Information

General Disk Configuration Requirements

1. Use only disk devices listed in the "Supported Disks" section of the Serviceguard subchapter.
2. Internal or external disk drives can be used for Root, Swap, and Dump disks.
3. It is strongly recommended that Root and Swap disks be mirrored (using MirrorDisk/UX) to ensure protection from a disk mechanism failure. If Root and Swap are internal disks they should be mirrored to external disks. External Root or Swap disks can reside on a dedicated bus or on a shared bus. If Root or Swap reside on a shared bus they should be configured using the Logical Volume Manager (LVM).
4. Database disks must be external disks and must be attached to a shared bus that is connected to all nodes. The database disks on the shared bus contain data, control, and log files. The Oracle instance for each node runs off of local disks that may or may not reside on a shared bus.
5. ATS shared tape devices must not be connected to the same shared bus as any shared disks.
6. All database disks must be configured using Shared LVM or CVM.
7. Hardware paths to data disk devices can be different on different nodes of the cluster. However, for simplicity, HP recommends that the hardware paths be identical.
8. At the introduction of HP UX 11.0, MirrorDisk/UX was not supported for Shared LVM software mirroring, meaning that the only supported mirroring solution for Serviceguard OPS Edition was to use hardware mirroring. With the introduction of the patch, PHKL_22267, MirrorDisk/UX can now be supported with Shared LVM for two node clusters only. MirrorDisk/UX on HP UX 11.11 and 11.22 supports Shared LVM on up to two nodes (no patches are required). RAID arrays are the only supported storage solution for Serviceguard OPS Edition clusters with more than 2 nodes. Hardware mirroring is still the preferred mirroring solution.
9. Serviceguard OPS Edition makes use of a special "cluster lock" disk on a shared bus to ensure integrity of the cluster after certain failures. A cluster lock disk must be designated for 2 node clusters, and is also supported but not required for 3 and 4 node clusters and is not supported for clusters with more than 4 nodes. Beginning with Serviceguard OPS Edition A.11.14, the Serviceguard Quorum Server is available as an alternative to a cluster lock disk.
10. Different mass storage bus types, such as SCSI and Fibre Channel can be mixed within a single enterprise cluster. Different bus types can be mixed within a single volume group, however, for simplicity and performance reasons HP recommends that this not be done.
11. For optimum performance, the I/O buses should not be heavily loaded. I/O workloads may vary depending on the system configuration and the customer application. Using performance tools, the system manager should validate that the I/O buses are not saturated. If they are, the system configuration should be modified.

SCSI specific Configuration Requirements

Single Ended SCSI

- Single ended SCSI devices are not supported for shared disk buses but can be used as non shared disks

Fast/Wide/Differential SCSI

- This is High Voltage Differential (HVD) SCSI and supports a maximum cable length of 25 meters.
- F/W SCSI disks must be equipped with an uninterruptible power supply (UPS) to avoid loss of data on powerfail.
- Configure no more than 8 stand alone disks per F/W SCSI bus.
- Configure no more than 4 F/W SCSI host bus adapters per HP PB bus.
- For online repairability, In line terminated SCSI cables are required. See the Serviceguard subchapter for more detail about the use of In line terminated SCSI cables.
- HVD and LVD SCSI devices cannot be mixed on the same shared SCSI bus.

Wide Ultra2 SCSI

- This is Low Voltage Differential (LVD) SCSI and supports a maximum cable length of 12 meters.
- LVD SCSI disks must be equipped with an uninterruptible power supply (UPS) to avoid loss of data on powerfail.
- Configure no more than 8 standalone disks per LVD SCSI bus.
- Configure no more than 4 LVD SCSI Host bus adapters per LVD SCSI bus.
- In line terminated SCSI cables are not supported, with the sole exception of configurations containing only the SC10.
- If online repairability is required, the DS2300 can be used if it is configured with two BCCs in Full Bus mode, and each host is connected to a different BCC, and the remaining connector on each BCC must have a SCSI terminator on it. To allow applications to continue to run after the failure of a Host Bus Adapter, it is required to have two DS2300s mirrored together, each connected to a separate SCSI HBA.
- HVD and LVD SCSI devices cannot be mixed on the same shared SCSI bus

Ultra160 and Ultra320 SCSI

- There are currently no supported configurations for Ultra160 and Ultra320 SCSI.

NOTE: Ultra160 SCSI will be available on DS2xxx JBODs on HP UX 11i v2 2H03.

Fibre Channel Configuration Requirements

- No more than two Fibre Channel hubs can exist in a single Fibre Channel Arbitrated Loop.
- Fibre Channel supports a maximum cable length of 500 meters (10,000 meters if a long wave port is used).
- Fibre Channel SAN configurations using Direct Fabric Attach (DFA) are supported, however, there are some restrictions for HA environments. (NOTE: These restrictions may change in the future.)
- Fibre Channel Direct Fabric Attach is preferred over Arbitrated loops for the additional performance and security it allows.
- Disk device LUNS for HA Cluster nodes must not be visible to any nodes outside the HA cluster. Similarly, disk device LUNS visible to the nodes in a cluster must not also be visible to nodes in other clusters. This LUN isolation can be accomplished through the use of fabric zoning on the switches or LUN security products, such as Secure Manager XP.
- The Primary and Alternate PVLinks must not use the same Fibre Channel Switch or Hub.

Configuration Examples

For configuration examples for a two node Serviceguard Extension for RAC cluster, see the Serviceguard QuickSpec for examples of the layout of power, disks, and LANs for three different enterprise clusters.

Supported hardware and software

This section describes the hardware and software components that can be used in conjunction with Serviceguard Extension for RAC.

Supported Servers

Refer to the Serviceguard subchapter for more information on supported servers with Serviceguard Extension for RAC.

Supported LANs

Serviceguard OPS Edition supports the following physical LAN media and network protocols:

LANs supported for Heartbeat and Data Networks:

Ethernet

- IEEE 802.3 (Ethernet and IEEE 802.3 subnets cannot be mixed. Subnets need to be either all Ethernet or all IEEE 802.3)
  FDDI *

- Token Ring *
- 10Base T
- 100VG *
- 100Base T
- 100Base FX
- 1000Base T
- 1000Base SX
- Fibre Channel (supported as a data network only on Serviceguard A.11.12 or later), not available on HP UX 11i).

Networks:

- TCP/IP
- SNA
- OSI

LANs supported for DLM traffic:

- Ethernet
- IEEE 802.3 (Ethernet and IEEE 802.3 subnets cannot be mixed. Subnets need to be either all Ethernet or all IEEE 802.3)
- FDDI *
- 10Base T
- 100VG *
- 100Base T
- 100Base FX
- 1000Base T
- 1000Base SX
- HyperFabric

* These LAN types are not supported with Serviceguard on HP UX 11i version 1.6

Other network protocols and LAN/WAN media, such as ATM or HyperFabric can be attached to nodes within a cluster, but they will not be protected by the high availability features of Serviceguard Extension for RAC. HyperFabric links can be used for DLM traffic on some OPS and RAC versions (please refer to Oracle documentation for more information).

Large Decision Support applications, especially with complex internode queries, may warrant dedicated high speed networks due to heavy internode traffic of OPS or RAC.

For Automatic Port Aggregation and Virtual LAN support refer to the Serviceguard subchapter.

VLAN Configuration Restrictions

HP UX allows up to 1,024 VLANs to be created from a physical NIC port. This obviously requires a large pool of system resources to accommodate such a configuration. With the availability of VLAN technology, Serviceguard OPS Edition may face potential performance degradation, high CPU utilization and memory shortage issues if there are a lot number of network interfaces configured in each cluster node. For best Serviceguard OPS Edition solutions with enough flexibility in networking, the following VLAN and general network configuration requirements must be adhered:

- Maximum of 30 network interfaces per node is supported in the cluster ASCII file. The interface could be physical NIC ports, VLAN interfaces, APA aggregates, or any of combinations.
- The physical LAN interfaces, upon which the VLAN interfaces are created, are recommended to configure in the ASCII file to meet Serviceguard heartbeat requirements. While VLAN interfaces can be configured for Serviceguard heartbeat, it is not recommended.
- Maximum of 14 VLAN interfaces per physical NIC port or maximum of 28 VLAN interfaces per node in the cluster ASCII file is supported.
- Local failover of VLANs must be on same link types. The primary and standby VLANs must have same VLAN id (or tag id).
- VLAN configurations are only supported on HP UX 11i.
- Only port based and IP subnet based VLANs are supported. Protocol based VLAN will not be supported because Serviceguard does not support any transport protocols other than TCP/IP.
- Each VLAN interface must be assigned an IP address in a unique subnet in order to work properly unless it is used as a standby of a primary VLAN interface, which shares the same VLAN ID.
- VLAN interfaces over APA aggregates are not supported.
- Failover from physical LAN interface to VLAN interfaces or vice versa is not supported due to restrictions of VLAN software.

Enhanced Heartbeat Requirements for VLAN

VLAN technology allows greater flexibility in network configurations in the enterprise. In order to allow Serviceguard OPS Edition to be running on such

dynamic environments successfully while maintaining its reliability and availability features simultaneously, the existing heartbeat rules must be adopted when VLAN interfaces are presented in the cluster:

- The existing minimum heartbeat requirements are unchanged, but they will be checked against the physical networks instead of any LAN interfaces, which may include VLAN interfaces, to avoid the single point of failure.
- The physical LAN interfaces, upon which the VLAN interfaces are created, are recommended to configure in the ASCII file to meet Serviceguard heartbeat requirements. While VLAN interfaces can be configured for Serviceguard heartbeat, it is not recommended.
- Heartbeats are still recommended on all physical networks interfaces.

Supported Disks

Refer to Serviceguard subchapter for more information on supported disk devices with Serviceguard Extension for RAC.

NOTE: Beginning with Serviceguard OPS Edition A.11.02 and HP UX 11.0 with the patch, PHKL_22267, MirrorDisk/UX software mirroring is supported on 2 node clusters. MirrorDisk/UX software mirroring on HP UX 11i supports Shared LVM on up to two nodes (no patches are required). Clusters with greater than 2 nodes should deploy hardware mirroring as offered in disk arrays for data protection.

Power Requirements

Power planning is critical for eliminating single points of failure in an enterprise cluster. Power circuits must be designed so that no single power circuit outage causes a loss of enough resources to prevent the cluster from operating. Consult the "Configuring OPS Clusters with Serviceguard OPS Edition" manual for a detailed discussion of this topic. General power planning rules for an enterprise cluster follow:

- Power must be provided so that no more than half of the nodes of a cluster can be lost at one time. For maximum availability each node should have its own power circuit.
- The cluster lock disk for up to a four node configuration should reside on a power circuit separate from the circuits providing power to the nodes of the cluster.
- The Quorum Server node used by any cluster should be connected to a separate power circuit, which is not shared by any equipment in the cluster.
- Each half of a mirrored disk pair should reside on different power circuits. Fibre Channel switches and hubs being used for Primary and Alternate PVLinks must be connected to separate power circuits.
- Newer disk arrays, HP 9000 servers, Fibre Channel switches and network switches have two or three separate power supplies. For clusters that contain devices with redundant power supplies, you should connect each device's power supply to a separate power circuit, such that the failure of a single power circuit will not cause the complete failure of any critical device in the cluster. For example, if all devices in a cluster have three power supplies, it will require a minimum of three separate power circuits to eliminate the power as a single point of failure for the cluster.
- Use of uninterruptible power supplies (UPS) is strongly recommended. Certain disk types require use of a UPS for powerfail support. Consult this guide for additional details

Terminal Server Requirements

Terminal Server devices, such as the DTC 16TN can be directly attached to a node, but these devices will not be supported by the high availability services of Serviceguard OPS Edition.

Other General Requirements

- Requires STREAMS/UX to be installed on the hosts.

Serviceguard Extension for RAC Consulting Requirements

Since the Serviceguard Extension for RAC product is very complex, and the technology is new to many customers, consulting is recommended. For the first installation of Serviceguard Extension for RAC, consulting is needed in the planning, design, installation, and configuration phases of the Oracle Real Applications Clusters (RAC) or Oracle Parallel Server (OPS) or cluster. The overall solution requires the correct design and setup of the cluster hardware, the Serviceguard Extension for RAC software, and the RAC or OPS database and application. Consulting contracts are optional for the second and subsequent installations at the same customer site. HP and Oracle account teams are responsible for properly setting customer expectations for high availability in an OPS or RAC environment. Both Hewlett Packard and Oracle technical consultants need to be involved to ensure a successful OPS or RAC installation.

The HP technical consultant can provide consulting in these areas:

- Analyzing high availability requirements
- Analyzing performance requirements and application characteristics to determine the appropriate number of nodes and class/type of nodes for the cluster
- Planning and designing the cluster hardware configuration
- Planning failure/recovery strategy
  Installing and configuring the Serviceguard Extension for RAC product, including:
  - Configuring the cluster manager
  - Creating shared volume groups and logical volumes
    Configuring the Distributed Lock Manager

- Creating run/halt scripts

The Oracle technical consultant can provide consulting in these areas:

- Designing the shared Oracle database
- Installing and configuring Oracle Real Application Clusters or Oracle Parallel Server
- Planning backup/recovery strategy
- Planning a strategy for client failover after failure.
- Tuning the RAC or OPS application for optimal performance.

Serviceguard Extension for RAC (SGeRAC) on HP Integrity Servers

For the first release of SGeRAC, A.11.14.01 on servers with Intel Itanium 2 processors, HP introduced a de coupled version of the product, which requires the purchase of Serviceguard (B3935DA) and Serviceguard Extension for RAC (T1859BA), along with the appropriate options. One license of each product must be ordered per processor (2AH option). Tier based pricing is not available. Note new SGeRAC product number for HP UX 11i v2.

Please refer to the Serviceguard QuickSpec for supported features and functionality on HP Integrity servers.