Partition Security and Isolation for HP Integrity and HP 9000 Servers

Last Updated August 3, 2007

This document provides information on HP-UX 11i Common Criteria Certification,

as well as the isolation and flexibility tradeoffs for the various types of partitioning and virtualization on HP Integrity and HP 9000 servers, including:

- HP nPartitions (hardware partitions)
- HP Integrity Virtual Machines (Integrity VM)
- HP Secure Resource Partitions (a combination of Resources Partitions from HP Process Resource Manager and HP-UX 11i Security Containment)
- HP-UX Virtual Partitions

and information of how to increase isolation by using the HP high availability tool, HP Serviceguard.

Executive summary

The various types of partitions on HP Integrity and HP 9000 servers vary in the amount of granularity, flexibility and isolation they provide to customers. In addition, when HP-UX 11i is used, its own certification comes into play in this overall discussion.

HP-UX 11i CERTIFICATION:

HP-UX 11i Operating Systems Common Criteria Certified

Common Criteria for Information Technology Security Evaluation was developed as an international standard recognized by multiple countries in response to more effectively certifying IT security assurance once for recognition by many governments and enterprises. ISO/IEC uses CC 2.3 in the continued development/maintenance of the ISO/IEC 15408 international standard. The governments and organizations contributing to the standard are: Australia, New Zealand, Canada, France, Germany, Japan, Netherlands, Spain, United Kingdom, and United States.

The enterprise can use the results of evaluations to help decide whether an evaluated product or system fulfils their security needs. These security needs are typically identified as a result of both risk analysis and policy direction.

Hewlett-Packard HP-UX 11i has been successfully evaluated and certified to the Common Criteria evaluation assurance level EAL4, against the functional requirements in the Controlled Access Protection Profile (CAPP).

HP-UX 11i v2 running on HP 9000 and Integrity platforms has been successfully evaluated against the requirements for the EAL4 Common Criteria (ISO 15408) Assurance Level, augmented by ALC_FLR.3 (flaw remediation), using the Controlled Access (CAPP) and Role-Based Access Control (RBAC) Protection Profiles. EAL4+ is sometimes used as the abbreviated form for additional assurances.

HP-UX 11i v3 Operating System in the process of Certification

HP-UX 11i v3 The latest release of HP's strategic UNIX operating environment makes virtualization easier to deploy and delivers mission-critical virtualization by combining virtualization and mainframe-class availability. Also, with HP-UX 11i v3, important applications, including business intelligence and data warehousing, are now simpler to secure and manage. This release announced in mid-February 2007 incorporates an integrated quality control and certification process so that certification is available in the shortest time possible. Certification evaluations are only conducted by the evaluation agency after release of the product. Their evaluation is expected to take approximately six-months. HP expects the certification for EAL4+ (EAL4 augmented with ALC_FLR.3 Systematic Flaw Remediation)/CAPP/RBACPP and NPars to be completed in the 4th quarter 2007.

Certification of HP-UX 11iv3 with multiple hardware platforms and nPartitions

The Target of Evaluation (TOE) for common criteria certification includes both the HP-UX 11iv3 operating system, the server platforms on which it operates and hard partitions, nPartitions supported by those platforms. The HP 9000 and HP Integrity platforms listed are those on which the product is supported and are also those for which the multi-platform TOE claims are made. The products include the entire range of HP 9000 and HP Integrity platforms as of the date of the HP-UX 11i v3 certification: rp3410, rp4410, rp4440, rp7420, rp8420, Superdome PA-8900, BL60p, rx1620, rx2620, rx3600, rx4640, rx6600, rx7620, rx7640, rx8620, rx8640, and Superdome Integrity.

Hard Partitions (nPartitions, or nPars)

Hard partitions (nPartitions) are also included in the TOE as a certified product operating with the HP-UX 11iv3 operating system.

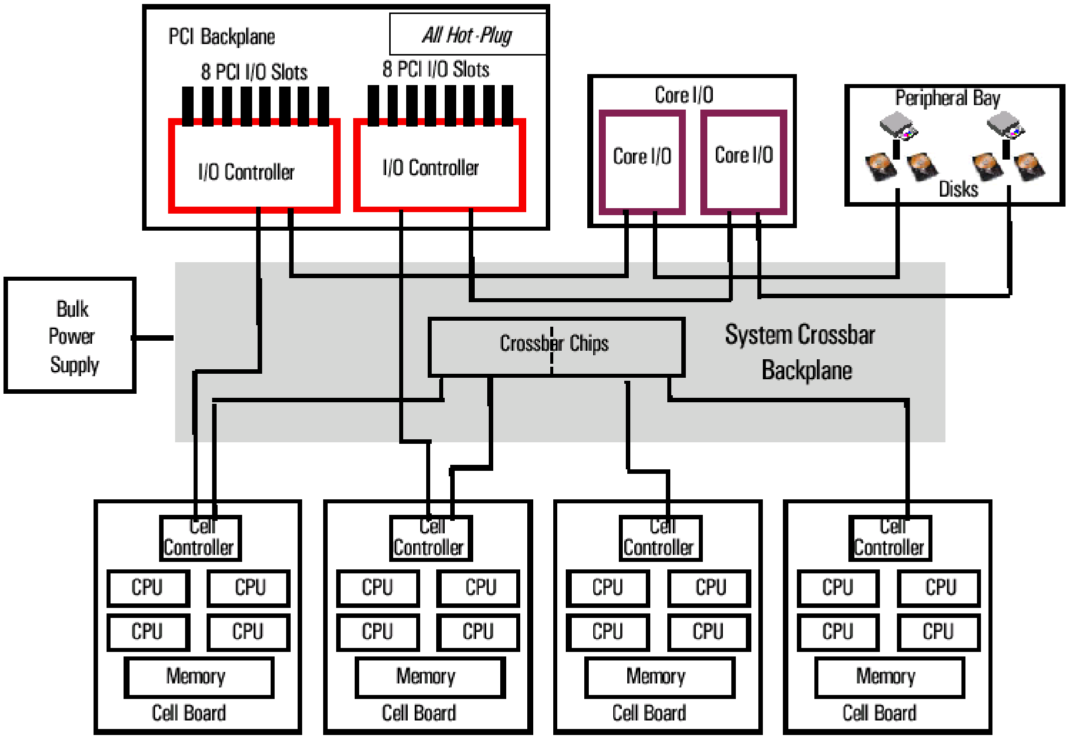

Hard partitions (nPartitions) are available on cell-based servers such as rp7420, rp8420, rx7620, rx7640, rx8620, rx8640, and Superdome. The server is split into a number of cells that can be allocated to the nPartitions. Each cell contains processor(s) and system RAM and may be associated with its own peripheral devices.

A cell-based server may be configured as one single large system or as multiple smaller systems by configuring nPartitions. Each nPartition defines a subset of server hardware resources to be used as an independent system environment. An nPartition includes one or more cells assigned to it and all I/O chassis connected to those cells. All processors, memory, and I/O in an nPartition are used exclusively by the software running in that nPartition. Thus, each nPartition has its own system boot interface and each nPartition boots and reboots independently. Each nPartition provides both hardware and software isolation, so that hardware or software faults in one nPartition do not affect other nPartitions within the same server complex.

GUIDANCE ON WHEN TO USE WHICH PARTITIONING/VIRTUALIZATION TECHNOLOGY – A TRADEOFF BETWEEN ISOLATION AND FLEXIBILITY:

nPartitions (with PARPERM set to "Restricted"): maximum availability and security isolation in a single box:

If customers are interested in the maximum availability and security isolation in a single box (e.g. multi-tiered firewall-separated architecture, or multi-department consolidation), HP recommends that they implement hard partitions (HP nPartitions) with its electrical isolation.

nPartition Configuration Privilege

Because it is not always desirable to allow a user on one nPartition to make changes that affect other nPartitions, HP provides the nPartition Configuration Privilege on sx1000-based or sx2000-based servers.

You can control the nPartition Configuration Privilege by using the PARPERM command at the service processor Command menu.

The nPartition Configuration Privilege has two settings:

- Unrestricted — The default setting, which allows the behavior described above. (A user in one nPartition can affect other nPartitions.)
- Restricted — Restricts use of the IPMI BT interface to the following capabilities:
    - Retrieving information about the server. Everything that is normally displayed by Partition Manager and the parstatus command is still available.
    - Making changes to the local nPartition's Partition Configuration Data. (Details on local versus remote nPartitions is provided later.)
    - Manipulating any of the attention indicators (LEDs).

- Powering on/off cells and I/O chassis that belong to the local nPartition.

Restricting the nPartition Configuration Privilege does not restrict deallocation of processors across nPartition boundaries.

By restricting the nPartition Configuration Privilege, you limit someone with superuser privileges on an nPartition to doing things that affect only that nPartition. However, when the nPartition Configuration Privilege is restricted, certain changes can only be made by using the nPartition management tools in the mode that utilizes IPMI over LAN.

Note: the other partitioning and virtualization solutions (HP-UX Virtual Partitions, HP Integrity Virtual Machines, and HP (secure) resource partitions) can be added to nPartitions and stand-alone servers for added granularity and flexibility.

Integrity Virtual Machines: Security isolation of separate (virtual) OS instances as well as: shared processors and I/O, and built-in dynamic resource allocation

HP Integrity Virtual Machines (Integrity VM), is a soft partitioning and virtualization technology within HP's Virtual Server Environment, which enables you to create multiple virtual servers or machines within a single HP Integrity server or nPartition. (The VM Host acts as the underlying hypervisor on that server or nPartition.)

Integrity Virtual Machines security-related features:

1. The VM Host includes HP-UX 11i v2, which has a Common Criteria EAL4 certification.
2. Integrity VM utilizes the HP Integrity security rings for increased security, by placing the Integrity VM Host in ring 0, while the individual guest OS kernels are placed in ring 1. (See Figure 2.)

   The VM Host, which runs at a higher security level than any individual guest, coordinates resources to the guests. In conjunction with other security mechanisms, this prevents a guest from interfering with any other guests.
3. Security review processes within development:
   a. Each new version of Integrity VM is required to go through an internal lab security review process (called a "threat analysis") in which HP examines all interfaces and components for security issues.
   b. We follow strict CATA (Commercial Application Threat Analysis) guidelines in Integrity VM development, and each release undergoes a CATA review in the design phase.
4. Security features for usage within Integrity VM:

a. Dedicated identity of each I/P-Frame by the implemented MAC-ID.
b. Certain Integrity VM customers are running Role-Based Access Control (RBAC) + Compartments on Integrity VM, and are even securing the guest using Compartments and RBAC.
c. Integrity VM provides secure virtual console access in that users may access the virtual machine's console without effective access to the VM Host system.
d. Integrity VM supports VLAN (virtual LAN) technology, which provides customers network isolation at port granularity for a virtual switch.

More information about #2 and 4d follow.

Integrity VM utilizes the HP Integrity security rings for increased security:

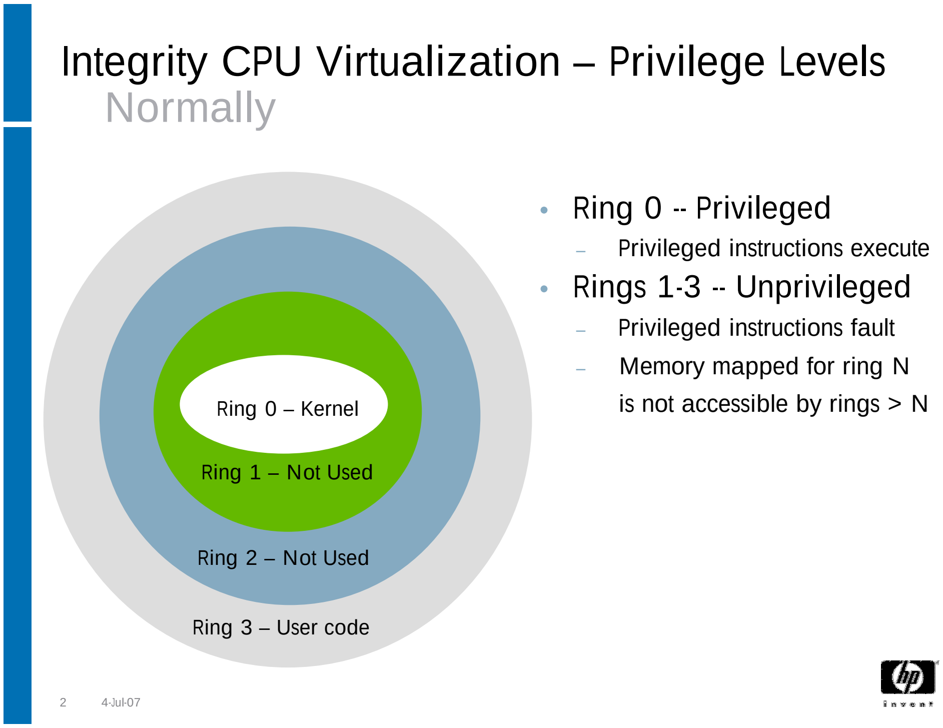

Integrity CPU Virtualization – Privilege Levels

Normally

- Ring 0 -- Privileged
  - Privileged instructions execute
- Rings 1-3 -- Unprivileged
  - Privileged instructions fault
  - Memory mapped for ring N is not accessible by rings > N

2 4-Jul-07

Figure 2

- In most operating systems, there are two modes of operation: OS kernel code typically runs in “privileged mode” (or security ring level 0) and user code runs in “unprivileged mode” (security rings 1,2, or most probably, 3). Itanium and many other processor architectures actually provide more than two modes. In the case of Itanium there are 4 levels of privilege: level 0 is the most privileged and level 3 is the least. These levels are often referred to as rings. Only code running in ring 0 can perform privileged operations such as enabling or disabling system interrupts, or managing virtual memory translations. Normally operating system kernels run at ring 0 and user code runs in ring 3.

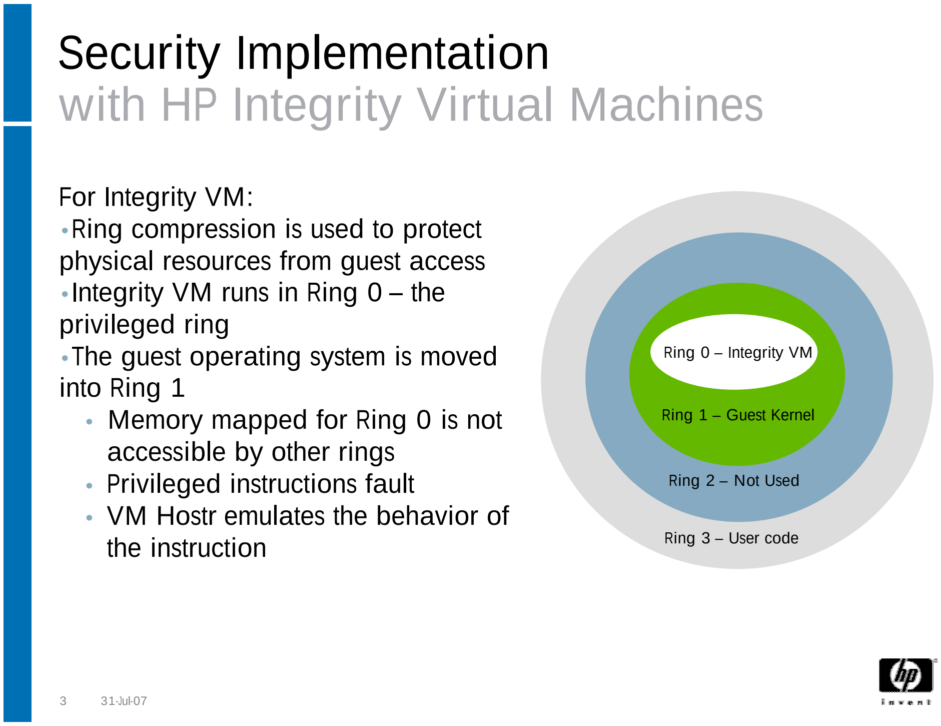

Security Implementation
with HP Integrity Virtual Machines

For Integrity VM:
- Ring compression is used to protect physical resources from guest access
- Integrity VM runs in Ring 0 – the privileged ring
- The guest operating system is moved into Ring 1
  - Memory mapped for Ring 0 is not accessible by other rings
  - Privileged instructions fault
  - VM Hostr emulates the behavior of the instruction

Ring 0 – Integrity VM

Ring 1 – Guest Kernel

Ring 2 – Not Used

Ring 3 – User code

3 31-Jul-07

- In Integrity Virtual Machines, only the virtual machine monitor runs in ring 0. Operating system kernels run in ring 1 – but are tricked into thinking they are running in ring 0. Since the guest operating system cannot affect the system state of the processor any more than a user application, the VM Host can ensure that the isolation between guests is just as strong as the isolation between two applications. One guest cannot damage the state of another guest, just as an application cannot damage the state of other applications. This model of running guest kernels at a lower privilege is called “ring compression”.
- When a guest attempts to perform a privileged operation, a fault is generated and the virtual machine monitor takes control to emulate it. Clearly emulating a privileged operation takes more time that executing it directly in hardware, so workloads that have a higher percentage of privileged operations tend to see more overhead when running in a virtual machine.

This can be summarized per the following picture.

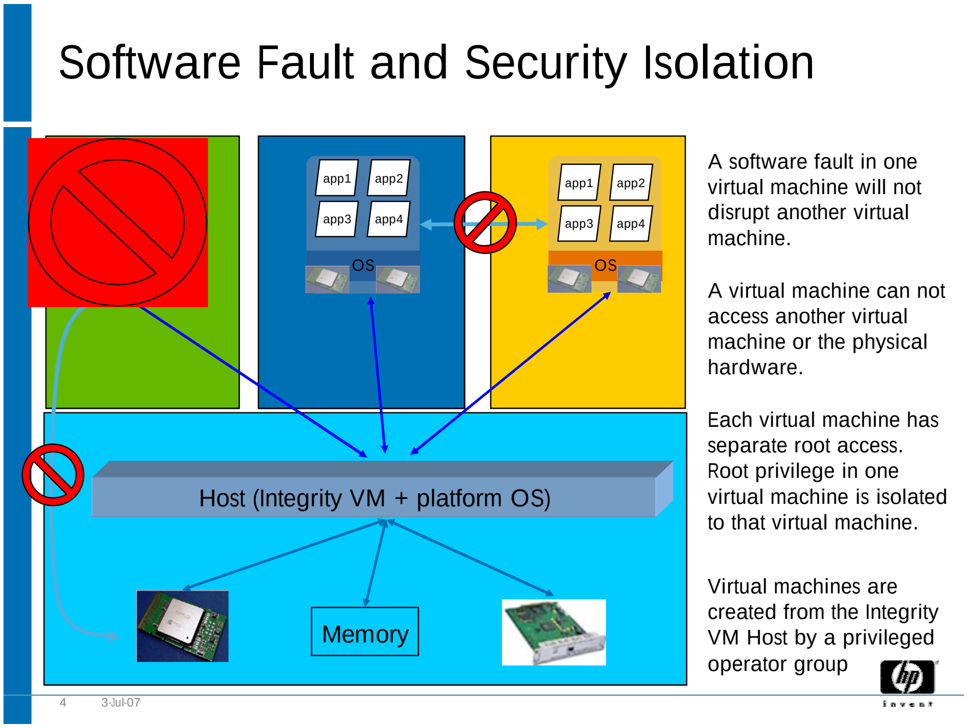

Software Fault and Security Isolation

A software fault in one virtual machine will not disrupt another virtual machine.

A virtual machine can not access another virtual machine or the physical hardware.

Each virtual machine has separate root access. Root privilege in one virtual machine is isolated to that virtual machine.

Virtual machines are created from the Integrity VM Host by a privileged operator group

Additionally, VLANs isolate the network traffic on separate virtual LANs

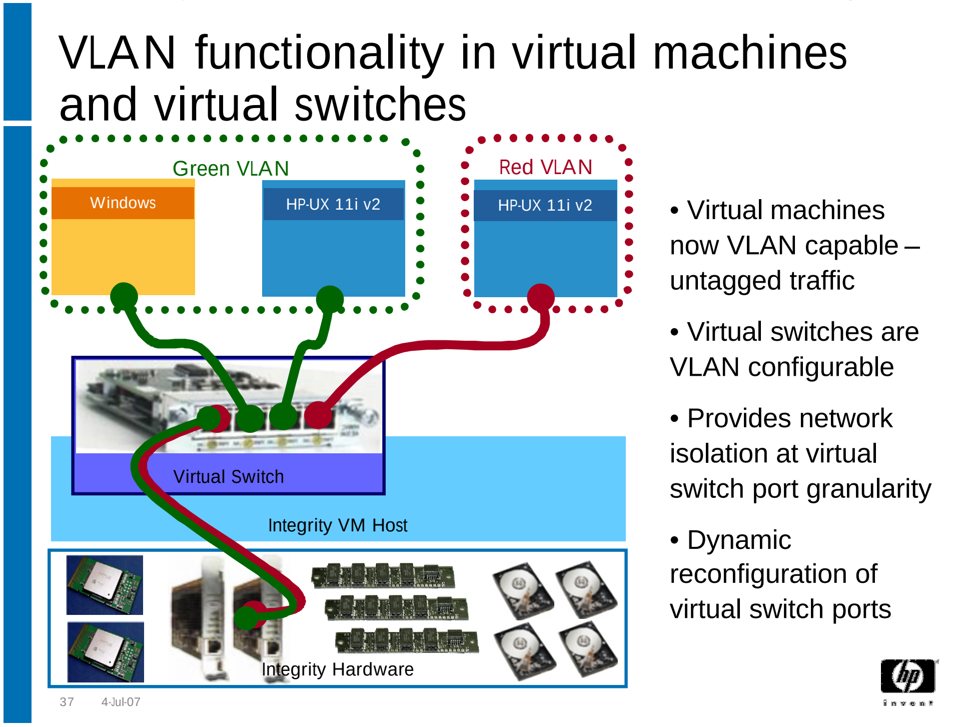

VLAN functionality in virtual machines and virtual switches

- Virtual machines now VLAN capable – untagged traffic
- Virtual switches are VLAN configurable
- Provides network isolation at virtual switch port granularity
- Dynamic reconfiguration of virtual switch ports

- VLAN support for Integrity VM's virtual switch works much as it would for a physical switch
- You can now assign ports from a virtual switch to separate, isolated virtual LANs (as you can see in the red and green virtual LANs above)
- This provides network isolation at port granularity for a virtual switch
- Note that the virtual switch ports can be dynamically reconfigurable

Some suggestions to further secure Integrity VM environments:

- Since the Integrity VM Host is a SPOF (like any hypervisor), use Serviceguard to protect a guest from a system and/or Integrity VM failure and fail the guest over to another system.
- Have separate virtual switches to effectively reduce exposure to denial-of-service issues.
- Isolate VMs to individual LUNs and/or card slots to dramatically reduce any denial-of-service issues that might arise from having two VMs share a single LUN (e.g., by mapping virtual disks to logical volumes or files residing on the same LUN).

Secure Resource Partitions (& Protected Systems): Maximum flexibility of security isolation in a single OS instance (greater initial setup cost, but reduced maintenance cost due to single OS instance), sub CPU granularity

Secure Resource Partitions combine kernel level security (of the security containment features of HP-UX 11i v2 and beyond) and proven resource management to stack multiple applications within the same operating system.

HP Process Resource Manager and HP-UX Workload Manager can dedicate specific resources to specific applications within an operating system image, avoiding resource contention issues (and thus Denial of Service cross-partition attacks). Using HP-UX 11i v2 and beyond Security Containment, organizations can ensure that application instances cannot access processes or files from other applications or the system. This ensures that multiple application instances run securely in a consolidated environment, providing the benefits of consolidation while preserving the security of a scale out environment. For software licensing purposes, the maximum number of CPUs in a Secure Resource Partition can be capped.

HP-UX 11i Virtual Partitions (vPars): High performance consolidation where administrators and processes with admin privileges are trusted across vPars.

vPars is a high performance, highly scalable, soft partitioning solution that provides granularity and flexibility to cell-based servers or nPartitions.

Admin privileges are typically trusted across vPars.

However, vPar Flexible Administration restricts vPar administrative capabilities to one or more privileged vPars:

- Before the release of vPars A.03.03, root on any vPar could administer any other vPar. Therefore all administrators had to be trusted for all vPars.
- With Flexible Administration (available in vPars release A.03.03):

    - You can restrict vPar administrative capabilities to one or more privileged vPars
    - Non-privileged vPar admin can only operate on local vPar
    - Any operation affecting another vPar is only allowed from the privileged vPar
    - Can only be enabled or disabled from the monitor
        - The list of designated vPars can be configured from the monitor
        - Note: does not restrict administrative functionality when working from the MON> prompt; it is still very important that you protect console access
        - Note: does not increase security isolation of root users in different vPar instances from each other

Serviceguard: Maximum availability and failover capabilities, single trust domain

Although HP Serviceguard is not a partitioning technology, it provides maximum availability and failover capabilities to servers, partitions and applications, with a single trust domain.

Since the monitor or hypervisor of any type of partitions is a Single Point of Failure, HP Serviceguard can be used to protect a partition and its applications from a system and/or partitioning failure, and fail the guest OS over to another system.

More information:

External partitioning continuum information library – brochures, white papers, documentation at: www.hp.com/go/PartitionLibrary:

For advice on choosing between different types of HP Integrity and HP 9000 Partitions – see section “Partitioning Continuum”:

- o “HP Partitioning Continuum for HP Integrity and HP 9000 servers: Flexible for optimal IT consolidation”, brochure (http://h71028.www7.hp.com/ERC/downloads/4AA0-1469ENW.pdf)
- o Under “White Papers”: “HP Partitioning Continuum for HP-UX 11i on HP 9000 and HP Integrity servers white paper” (http://h71028.www7.hp.com/ERC/downloads/5982-9141EN.pdf)

HP-UX 11i Security Solutions:

- https://www.hp.com/go/hpux11isecurity: This site contains information on all HP-UX 11i security features. Links to more detailed information and free software downloads are obtained through this site.
- https://www.hp.com/go/SoftwareDepot: Internally to HP, you can access software downloads and obtain links to presentations, documents, and white papers.

Summary

The various types of partitions on HP Integrity and HP 9000 servers vary in the amount of granularity, flexibility and isolation they provide to customers. In addition, when HP-UX 11i is used, its own certification comes into play in this overall discussion.

Since the monitor or hypervisor of any type of partitions is a Single Point of Failure, HP Serviceguard can be used to protect a partition and its applications from a system and/or partitioning failure, and fail the guest OS over to another system.

4AA0-5913ENW, August 2007