

Rules for configuring HP Serviceguard with HP 9000 and Integrity systems on HP-UX 11i v2 September 2004 update



Executive summary.....	3
Introduction.....	3
Purpose of document.....	3
Definition of terms.....	4
Audience.....	4
Related documents.....	5
Problem statement and goals for mixed clusters.....	5
Summary of the results.....	5
Summary of changes for existing clusters.....	6
Architectural rules and typical configurations.....	6
Cluster hardware architecture.....	6
Allowed servers in a mixed PA/Itanium HP Serviceguard cluster.....	7
Allowed storage systems in a mixed PA/Itanium HP Serviceguard cluster.....	7
Supported cluster software architecture.....	9
Application software requirements.....	9
Application code.....	9
Homogeneous architecture for primary and failover node.....	12
Heterogeneous architecture for both primary and failover nodes.....	12
Typical configuration examples.....	12
Adding one Itanium server to a two-node PA cluster.....	12
Original two-node PA cluster.....	13
New three-node mixed PA/Itanium cluster.....	13
Adding two Itanium servers to a two-node PA cluster.....	14
Original two-node PA cluster.....	14
New four-node mixed PA/Itanium cluster.....	15
Adding one critical application and two new Itanium nodes to existing four-node PA cluster.....	16
Original four-node PA cluster with four critical applications.....	16
New six-node mixed PA/Itanium cluster.....	17
Mixed clusters as transition aid from PA to Itanium in a multitier SAP environment.....	18
Initial four-node homogenous PA cluster.....	18
First step: adding the first Itanium node to convert to a mixed cluster.....	19
Second step: adding one more Itanium node to mixed cluster.....	20
Third step: converting the mixed cluster to a homogenous Itanium cluster.....	20

PA to Itanium transition service utilizing HP Serviceguard cluster technology	21
Specific ISV information.....	21
Oracle 9i/10g database server.....	22
SAP with HP Serviceguard Extension for SAP.....	22
How to implement a mixed PA/Itanium HP Serviceguard cluster.....	22
Starting with an existing cluster.....	23
Starting with an existing HP-UX 11i v1 PA cluster	23
Starting with an existing HP-UX 11i v1.6 or v2 Itanium cluster	23
Merging an existing PA cluster with an existing Itanium cluster.....	24
Installing a new mixed cluster	24

Executive summary

With the advent of HP-UX 11i version 2 (HP-UX 11i v2), HP Integrity (Intel® Itanium®-based) systems are becoming prevalent in customer environments. Many customers are using HP Serviceguard to protect their mission-critical applications. For their PA-RISC-based servers, customers are running HP-UX 11i version 1 (HP-UX 11i v1) and they are using HP-UX 11i version 2 (HP-UX 11i v2) for their Itanium-based systems. In the fourth calendar quarter of 2004, HP will be shipping an update of its HP-UX 11i v2 operating system that supports both PA-RISC- and Itanium-based systems. This update will be referred to in documents as HP-UX 11i v2 September 2004 update or release or as HP-UX 11i v2 update 2 (HP-UX 11i v2UD2).

Up to now, it was not possible to combine HP 9000 and Integrity servers into a single cluster because two different versions of the operating system were required for them. It has also not been possible to have two versions of HP-UX or two versions of HP Serviceguard in the same cluster, except in a rolling upgrade situation. Now, with HP-UX 11i v2UD2 and HP Serviceguard A.11.16, it is possible to have HP 9000 and Integrity servers in the same cluster.

It is expected that this mixed cluster support will:

- Provide a much easier transition from one server architecture to the other
- Provide investment protection for existing customers
- Allow a much longer transition period than what would have been possible by having each architecture running in a separate cluster

Introduction

With the September 2004 update of HP-UX 11i v2, HP provides one operating system version with similar functionality to the HP 9000 and Integrity lines of servers.

Starting with HP Serviceguard A.11.16, HP 9000 and Integrity servers that run HP-UX 11i v2UD2 can be configured to run in a single heterogeneous PA/Itanium HP Serviceguard cluster.

Purpose of document

This technical whitepaper explains how to create a new cluster with both types of servers and how to upgrade an existing PA-based cluster or an existing Itanium-based cluster to include both server architectures. The document:

- Explains the new feature of HP Serviceguard that provides the opportunity to mix Itanium and PA servers in a single HP Serviceguard cluster
 - HP-UX 11i v2UD2 provides the same feature set to both PA and Itanium platforms.
 - HP Serviceguard A.11.16 provides feature parity between PA and Itanium systems.
 - Many independent software vendors (ISVs) provide their software in two versions, one for PA and the other for Itanium.
- Informs the reader about configuration rules and best practices for:
 - Cluster hardware
 - HP-UX 11i v2UD2
 - nPARs and vPARs
 - HP Serviceguard A.11.16
 - Protected applications
 - Rules when mixing architectures with a single application
 - Rules when mixing architectures with multiple applications
 - Rules when mixing architectures for multitier applications

Definition of terms

Table 1.

Term	Definition
Aries	Aries is the HP-UX PA-RISC to HP-UX Itanium binary emulator. Aries transparently emulates both 32- and 64-bit HP-UX PA-RISC applications when running on HP Integrity servers.
DTS	Disaster Tolerant Solutions; multi data center HP Serviceguard configurations that protect against the loss of one complete data center.
HP-UX 11i v1	HP-UX 11i version 1; HP 9000 server operating system (PA-RISC) only; internally also know as 11.11 (output of <code>uname -r</code> command).
HP-UX 11i v2	HP-UX 11i version 2; HP Integrity server operating system (Itanium) only; internally also know as 11.23 (output of <code>uname -r</code> command).
HP-UX 11i v2UD2	HP-UX 11i version 2 September 2004 update; first common operating system for HP 9000 and Integrity servers; also know as 11i v2 update 2 (output of <code>swlist</code> for the Operating Environment will show B.11.23.0409).
ISV	independent software vendor
MCOE	mission-critical operating environment; the HP-UX 11i MCOE is a high-availability operating environment for HP servers.
nPAR	nPartition; the HP hard partition offering
PA	PA-RISC, Precision Architecture Reduced Instruction Set Computer; an HP proprietary processor architecture, predecessor of the current Itanium architecture.
SG A.11.16	HP Serviceguard release A.11.16; the first release of HP Serviceguard to support mixed clusters built with HP 9000 and HP Integrity systems
SGeRAC	HP Serviceguard Extension for Oracle Real Application Cluster (RAC); a software package that enables you to integrate Oracle 9i and 10g RAC into HP Serviceguard
SGeSAP	HP Serviceguard Extension for SAP is a software package that enables you to integrate SAP applications into Serviceguard to make them highly available.
vPAR	vPartition; the HP virtual/soft partition offering
WLM	Workload Manager; a resource management application for controlling the usage of system resources per application.

Audience

This white paper is intended for the reader who is considering the implementation of a mixed PA/Intel® Itanium® HP Serviceguard cluster on HP-UX 11i v2UD2—the first PA/Itanium release of HP-UX. The reader should be familiar with HP-UX and HP Serviceguard in general. For further information, see the “Related documentation.” This whitepaper explains the various supported configurations and the rules that must be followed when mixing the two server architectures in a single HP Serviceguard cluster.

Related documents

The following list of offers links to documents that provide detailed information on the technology discussed in this white paper:

- Release notes of HP-UX 11i v2 September 2004 update
 - <http://docs.hp.com/hpux/os/11iv2/index.html#Release%20Notes>
- HP 9000 and HP Integrity Server-Storage Support Matrices
 - http://www.hp.com/products1/serverconnectivity/support_matrices.html
- Current HP Serviceguard release notes and manual
 - <http://docs.hp.com/hpux/pdf/B3935-90078.pdf> (release notes)
 - <http://docs.hp.com/hpux/pdf/B3936-90079.pdf> (manual, 11th edition)
- Aries
 - http://devresource.hp.com/drc/STK/docs/refs/ieee_aries.pdf
 - <http://devresource.hp.com/drc/STK/docs/refs/Aries.jsp>
 - <http://docs.hp.com/cgi-bin/otsearch/hpsearch>; search for “Aries”
- SGeSAP release notes
 - <http://docs.hp.com/hpux/pdf/T2357-90004.pdf>
- Oracle 9i PA to Itanium transition
 - http://download.oracle.com/docs/pdf/B10566_02.pdf

Problem statement and goals for mixed clusters

Because both HP 9000 and Integrity servers will be supported on HP-UX 11i v2UD2, HP Serviceguard customers expect to be able to configure clusters that contain both types of systems. Providing this feature will:

- Assist customers who want to
 - Transition from PA to Itanium
 - Have HP 9000 and Integrity systems coexist in their environments
- Support early Itanium adopters
- Provide a proven and stable path for the early majority
- Allow the late majority to remain on subsequent PA releases while introducing Itanium systems to their operating environments at the same time

Supporting mixed PA/Itanium clusters in HP-UX 11i v2 provides an ideal platform for customers to transition to Itanium at their own pace with optimal support from the HP-UX environment. With mixed clusters:

- HP Integrity servers can be easily added to an existing PA cluster.
- Customers can “roll” their applications to Integrity using a failover operation.
- Applications can be “test driven” on Itanium nodes inside an existing cluster and easily failed back to an original PA node as a fast recovery/rollback option in cases in which a problem is discovered with the application on the Itanium servers that does not appear on PA servers.

Summary of the results

With the HP-UX 11i v2 September 2004 update, HP offers one common operating system for both server lines—HP 9000 PA-RISC based and HP Integrity IA-64 based servers—and enables mixing them in a single HP Serviceguard A.11.16 cluster. HP-UX 11i v2UD2 provides customers with the opportunity to gain first experience with HP Integrity servers within their known environments. It also

promotes transition to Itanium by enabling customers to gradually add Integrity servers into an all-PA clustered environment.

Mixed PA/Itanium clusters are easy to set up and no different in terms of configuration and operation. Mixed PA/Itanium cluster implementations have only three basic requirements:

- All nodes must have HP-UX 11i v2UD2 installed as a minimum.
- All nodes must use HP Serviceguard version A.11.16 or later.
- The applications are configured to follow one of the ISV's configuration rules:
 - Failover between homogeneous nodes in a mixed cluster only.
 - Failover between heterogeneous nodes in a mixed cluster. Following this rule, either of the two options are available for the application using:
 - One set of PA binaries and Aries emulator on Itanium
 - Two native sets of binaries (one for PA and one for Itanium), either all 32-bit or all 64-bit

Summary of changes for existing clusters

Existing HP-UX 11i v1 Serviceguard customers with HP 9000 systems will see the following changes when transitioning to HP-UX 11i v2 September 2004 update:

- No support is available for the HP-UX System Administration Manager (SAM) GUI to configure the cluster; the command line interface and HP Serviceguard Manager are available to configure the cluster.
- No support is available for advanced tape sharing (ATS).
- No support for M by N Metrocluster/Symmetrix Remote Data Facility (SRDF) configurations will be available at initial release; HP plans to add support at a later date.

Customers already using HP-UX 11i v2 to cluster HP Integrity systems will see the following change when transitioning to HP-UX 11i v2 September 2004:

- The maximum cluster size increases from eight to 16 nodes.

Customers familiar with homogeneous HP Serviceguard clusters on either HP-UX 11i v1 or HP-UX 11i v2 will see the following change when transitioning to HP-UX 11i v2 September 2004:

- No support for vPARs will exist at initial release of HP-UX 11i v2 September 2004; vPAR support will be added later. Customers who currently deploy vPARs on HP-UX 11i v1 with HP 9000 servers should synchronize their upgrades to HP-UX 11i v2 until vPAR support is added to it.
- Mixed clusters are not currently supported with SGeRAC. HP is investigating the support of this environment in the future.

Architectural rules and typical configurations

This section explains the architectural rules for mixed PA/Itanium HP Serviceguard clusters. It lists the requirements on the hardware and software sides.

Cluster hardware architecture

The HP-UX 11i v2UD2 Release Notes (available at <http://docs.hp.com>) specify the HP 9000 and Integrity servers supported with the first common operating system for servers of both architectures. Starting with HP Serviceguard A.11.16, they can be configured in a mixed cluster.

Not all storage arrays—especially older systems—are supported on new Integrity servers. Most often, customers start with an existing homogeneous HP Serviceguard cluster consisting of PA systems, including shared storage. Before adding Itanium nodes to the cluster, support for the shared storage

array on Itanium systems must be verified. The hardware configuration guide of the particular server and the “Allowed storage systems in a mixed PA/Itanium HP Serviceguard cluster” section provides this information.

Allowed servers in a mixed PA/Itanium HP Serviceguard cluster

The following servers are supported in a mixed PA/Itanium HP Serviceguard cluster:

- PA-RISC-based servers (PCI or PCI-X based backplane only)
 - Entry level—rp24xx, rp34xx, rp44xx
 - Mid-range—rp54xx, rp74xx, rp84xx
 - High end—Superdome
 - Current A-, L- and N-class servers
- Itanium-based servers
 - Entry level—rx16xx, rx26xx, rx46xx, rx56xx
 - Mid-range—rx76xx, rx86xx
 - High end—Superdome

Allowed storage systems in a mixed PA/Itanium HP Serviceguard cluster

The storage system must be supported on both HP 9000 and Integrity Servers with HP-UX 11i v2UD2 to build a mixed cluster. Refer to “HP 9000 and HP Integrity Server-Storage Support Matrices” for detailed information regarding supported server/storage combinations. Currently, individual matrices exist for Integrity (HP-UX 11i v2) and HP 9000 (HP-UX 11i v1) servers. Table 2 is an extract from http://www.hp.com/products1/serverconnectivity/support_matrices.html both server types.

Note

It is assumed that HP will post a combined matrix for HP 9000 and Integrity servers at the time the HP-UX 11i v2 September 2004 update becomes available.

Table 2. HP-UX 11i v2 storage connectivity

	Integrity servers	HBA	Storage
Fibre Channel	Superdome	A6795A	HP StorageWorks XP128/1024 Disk Arrays
	rx8620-32	A9782A	
	rx7620-16	A9784A	HP StorageWorks Enterprise Virtual Array 5000 (EVA5000) ³
	rx5670	A6826A	
	rx4640-8		HP StorageWorks Enterprise Virtual Array 3000 (EVA3000) ³
	rx2600-2		
	rx1600-2		HP StorageWorks Virtual Array 7410 (VA7410)
			HP StorageWorks Virtual Array 7110 (VA7110)
			HP StorageWorks Modular Smart Array 1000 (MSA1000) ¹
			EMC Symmetrix ⁴

	Integrity servers	HBA	Storage
SCSI	Superdome	A6828A	DS2100
	rx8620-32	A6829A	DS2300
	rx7620-16	A5838A ²	HP StorageWorks Modular Smart Array 30 (MSA30)
	rx5670	Core I/O	MI
	rx4640-8		
	rx2600-2		
	rx1600-2		

Notes

- All configurations provide boot support unless otherwise noted.
- All configurations provide HP Serviceguard support unless otherwise noted.

¹ Single controller support is only available with HP-UX.

² No boot support/HP Serviceguard support of A5838A is available with Integrity servers.

³ No boot support is available with EVA5000 and EVA3000.

⁴ Symmetrix 5.0, 5.5, 6.0, and 6.5 are supported in mixed PA/Itanium clusters.

Table 3. HP-UX 11i v1 storage connectivity

	HP 9000 Servers	HBA	Storage
Fibre Channel	Superdome	A6795A	XP128/1024
	rp8420-32	A9782A	EVA5000
	rp8400-16	A9784A	EVA3000
	rp7420-16	A6826A	VA7410
	rp7410-8		VA7110
	rp4440-8		MSA1000 ¹
	rp3440-4		EMC Symmetrix ²
	rp3410-2		
	rp5400 series		
	rp2400 series		
SCSI	Superdome	A6828A	MSA30 MI
	rp8420-32	A6829A	DS2100
	rp8400-16	A5838A	DS2300
	rp7420-16	Core I/O	
	rp7410-8		
	rp4440-8		
	rp3440-4		
	rp3410-2		
	rp5400 series		
	rp2400 series		

Notes:

- All configurations provide boot support unless otherwise noted.
- All configurations provide HP Serviceguard support unless otherwise noted.

¹ Single controller support is available only with HP-UX.

² Symmetrix 5.0, 5.5, 6.0, and 6.5 are supported in mixed PA/Itanium clusters.

Supported cluster software architecture

HP-UX 11i v2UD2 is the first HP-UX operating system release that supports both HP 9000 and HP Integrity servers. Having this version installed on all nodes is a prerequisite for configuring a mixed PA/Itanium cluster. The following requirements exist on the cluster software:

- HP-UX 11i v2UD2
 - Mission-critical operating environment (MCOE) installed on all nodes

Note

The “Typical configuration examples” section of this whitepaper explains how to prepare an existing homogeneous cluster for mixed cluster support.

- A non-MCOE based system with HP Serviceguard version A.11.16 or newer
- Same or similar patch set for operating system, HP Serviceguard, and ISV applications
- Same version of volume manager software (if used for shared storage)
 - Logical Volume Manager (LVM)
 - Veritas Volume Manager (VxVM)
 - Veritas Cluster Volume Manger (CVM)
- Same version of file system software (if used for shared storage)

HP journal file system (JFS)/Veritas file system (VxFS) version 3.5 with file system layout 5 will be shipped with HP-UX 11i v2UD2. This software version is compatible to file systems created with a previous layout. If a VxFS file system was created on HP-UX 11i v1, it is layout 4 based. It can be used as-is on HP-UX 11i v2UD2 or converted to layout 5. However, if it is converted to layout 5, it cannot be used with an HP-UX 11i v1 system anymore.

Application software requirements

HP can only provide guidelines on how to configure applications in a mixed PA/Itanium cluster. It is up to the individual application vendor or the customer, in the case of self-developed applications, to actually approve and support the configuration.

Application code

HP-UX 11i v2UD2 supports applications that have been compiled in 32- and 64-bit mode on both PA and Itanium systems, which leads to four possible formats of executables (also called binaries). Each format is identified by the operating system through its magic number (two bytes at the beginning of a binary file identifying the type of the file). All four forms of executables can run on Itanium systems—the Itanium executables run natively, and the PA executables run through the Aries binary emulator on Integrity servers. Only native PA executables can run on HP 9000 systems. The following table shows the possible combinations.

Table 4.

	HP 9000 systems with HP-UX 11i v2UD2	HP Integrity systems with HP-UX 11i v2UD2
PA 32-bit application	Supported	Supported with Aries
PA 64-bit application	Supported	Supported with Aries
Itanium 32-bit application	N/A	Supported
Itanium 64-bit application	N/A	Supported

Even though Itanium systems support all four types of applications at the same time, specific applications will either be all 32-bit or all 64-bit. The mix is between PA and Itanium and not between 32-bit and 64-bit.

When preparing an application for a mixed PA/Itanium cluster with the ability to fail over between nodes of a different architecture, the application executables must be available in a form that runs on both architectures. Providing only Itanium executables would not be sufficient, as Table 4 shows. Two options are available:

- Native executables for both architectures are available for a specific version of the application—either from an ISV or self-developed.
- The application is available as PA version only and runs natively on PA nodes and through the Aries binary emulator on Itanium nodes.

Note

Although HP fully supports the use of the Aries emulator for running PA executables on an Itanium system, some ISVs might require the use of executables that have been compiled directly for a given architecture. These ISVs will not support the use of a PA executable running with the Aries emulator on an Itanium system.

Native PA and Itanium executables

HP has been working with many ISVs and individual customers over the last couple of years to help them transition their applications to Itanium. As a result of this effort, many individual applications and those of major ISVs are available as native Itanium versions on HP Integrity servers.

If a particular application is available natively for both architectures, HP recommends taking advantage of the two sets of native executables and configuring the cluster in a way that the appropriate set is used on each node. To do this, you can either:

- Store the applications executables locally on each node of the cluster and only have the application data on shared volumes. In this way, you can ensure that PA executables reside on HP 9000 nodes and Itanium executables reside on Integrity server nodes. Also, having one set of executables locally available on each node increases the overall availability in the event one set is inadvertently corrupted by human error. On the other hand, if there is more than just one server of each architecture in the cluster, the versions of the executables must be kept in sync manually.
- Store both sets of executables, along with the applications data on shared volumes, and add logic to the applications startup scripts for choosing the appropriate executable, depending on the system on which it is started. Although this procedure might impose additional administration effort while setting up the cluster, it reduces the burden of keeping the executables in sync between the nodes because they are only present once on a shared disk.

Note

`model(1)` or `getconf(1)` commands can be added to the application startup script to distinguish between PA and Itanium nodes for choosing the appropriate set of executables.

The following figure illustrates sample script to pick correct set of executables in a mixed cluster.

Figure 1.

```
typeset ARCH=`model | cut -c 1-4`
case "${ARCH}" in
    "9000")
        echo "use PA binaries"
        ;;
    "ia64")
        echo "use IPF binaries"
        ;;
    *)
        echo "Error: system is neither PA nor IPF"
        ;;
Esac
```

Native PA executables and Aries

Aries is a binary emulator included with HP-UX 11i v2 that transparently emulates 32- and 64-bit HP-UX PA-RISC executables on HP-UX Itanium-based machines. Aries is transparent in that the:

- HP-UX PA-RISC application does not have to be recompiled
- User does not explicitly invoke Aries

The HP-UX kernel running on an Itanium system recognizes an HP-UX PA-RISC executable and invokes Aries to emulate the application automatically without any user intervention.

Most PA-RISC applications run under Aries emulation without any problem. Refer to the Aries documentation for details about the features and limitations of Aries. This information is especially important for customers planning to deploy self-developed applications or applications for which the source code is not available. Performance under Aries emulation will vary according to the specific application. Further information can be found at:

- http://devresource.hp.com/drc/STK/docs/refs/ieee_aries.pdf
- <http://devresource.hp.com/drc/STK/docs/refs/Aries.jsp>

In the case of an ISV application deployment, details about the support under the Aries emulator can be acquired from the particular ISV.

Application data

The application data resides on shared storage, and the format must be transparent to PA and Itanium systems, which is the case for 32-bit (ILP 32) PA/Itanium applications and 64-bit (LP 64) PA/Itanium applications. ILP 32 means Integer, Long, and Pointer are 32 bits wide, and LP 64 means Long and Pointer are 64 bits wide.

Database applications, for instance, often store their data in an architecture-independent manner for which it does not matter on which architecture the data was generated. It will be usable by systems of both architectures, PA and Itanium.

Note

If not explicitly allowed and supported by the application vendor, the application must be of the same type (32-bit or 64-bit) on both architectures.

Homogeneous architecture for primary and failover node

A mixed HP Serviceguard cluster consists of both PA and Itanium nodes. The application, however, could be running only on nodes of the same architecture. Some applications would run on the PA-based nodes, and different applications would run on the Itanium-based nodes.

This condition could apply when, for business or technical reasons, the IT department is not ready to deploy the application on both architectures. Additionally, the application might only be available as an Itanium executable and unable to run on PA nodes, or the application might exist as a PA executable only and is either not supported or unable to run using the Aries emulator.

When running an application on homogeneous nodes only, no specific impact on the application or on the support of it is expected.

The system administrator must configure the application package so that the applications can only run on servers of the same platform.

Heterogeneous architecture for both primary and failover nodes

In this scenario, an application can run on nodes of both architectures in a mixed PA/Itanium cluster. The following constraints must be met to deploy a specific application in a mixed PA/Itanium cluster and allow the application to fail over between nodes of different architectures:

- The application must be supported on both platforms, either by:
 - Allowing the PA version of the application to run under the Aries emulator
 - Providing native executables of the application for both architectures
- The application data must be compatible between PA and Itanium systems.

HP provides only guidelines on how to configure applications in a mixed PA/Itanium cluster. The ISV, or the customer in the case of a self-developed application, must actually approve and support the configuration and allow failover between nodes of both architectures.

Typical configuration examples

The most typical use case is an existing PA cluster to which a customer wants to add new Itanium nodes. These examples are based on an interim assumption that database vendors do not allow failover between nodes of heterogeneous architecture. On-going discussions with Oracle® regarding this topic had not been finalized by the time this paper was written.

The examples in this section just give an overview of possible scenarios. The “How to implement a mixed PA/Itanium HP Serviceguard cluster” section gives a more detailed overview of the steps involved in implementing a mixed cluster.

Adding one Itanium server to a two-node PA cluster


This example shows how one Itanium server is added to an existing two-node homogeneous PA cluster configuration that hosts a database and an application server.

Original two-node PA cluster

Before the Itanium node is added to the homogeneous PA cluster, each node hosts a critical application. The database and application are configured to fail over to the other node in case their primary node goes down. The failover scenarios are:

- Node A fails.
 - Node B runs both database and application servers.
- Node B fails.
 - Node A runs both database and application servers.

Figure 2. Original two-node PA cluster

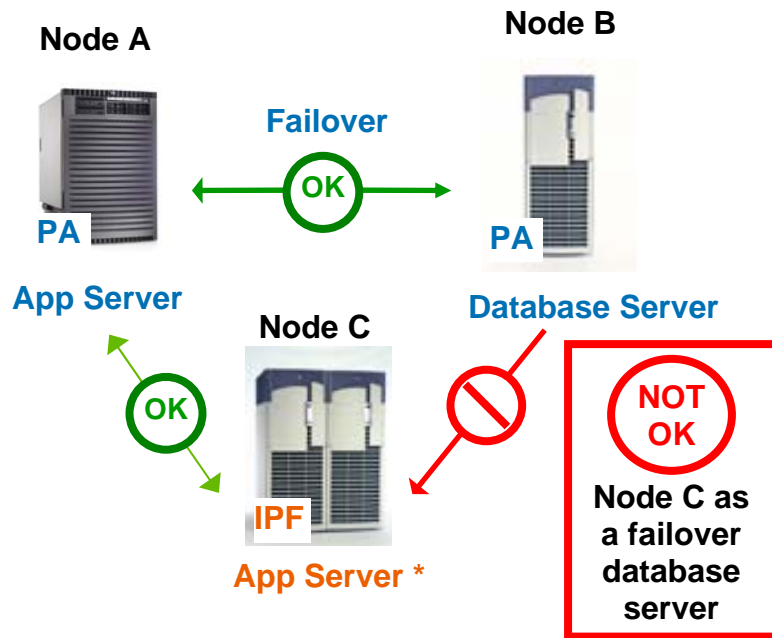


New three-node mixed PA/Itanium cluster

In this example, a new Itanium node has been added to the cluster for horizontal scaling on the application server level. The failover scenarios are:

- Node A fails.
 - No failover occurs; node C takes over application server load for all users.
 - Optionally, the application server instance from node A fails over to node C if supported and more preferred by the application.
- Node B fails.
 - Node A runs the database and application servers (node C cannot act as failover node for database service of node B).
 - Optionally, the application server instance from node A can be halted or reduced in resource entitlements through Workload Manager (WLM) if the remaining instance on node C can handle the application server load alone.
- Node C fails.
 - No failover occurs; node A takes over the application server load for all users.
 - Optionally, the application server instance from node C fails over to node A if supported and more preferred by the application.

Figure 3. New three-node mixed PA/Itanium cluster



Adding two Itanium servers to a two-node PA cluster

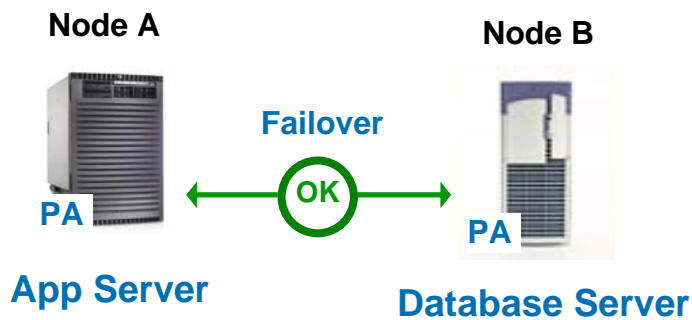
This example shows how mixed clusters can help customers to transition from HP 9000 to Integrity servers. Two new Itanium servers are added to a two-node homogeneous PA cluster to transition the database tier to HP Itanium servers.

Original two-node PA cluster

Before the Itanium nodes are added to the homogeneous PA cluster, each original cluster node hosts one critical application. The database and application servers are configured to fail over to the other node in case the primary node goes down. The failover scenarios are:

- Node A fails.
 - Node B runs both database and application servers.
- Node B fails.
 - Node A runs both database and application servers.

Figure 4. Original two-node PA cluster



New four-node mixed PA/Itanium cluster

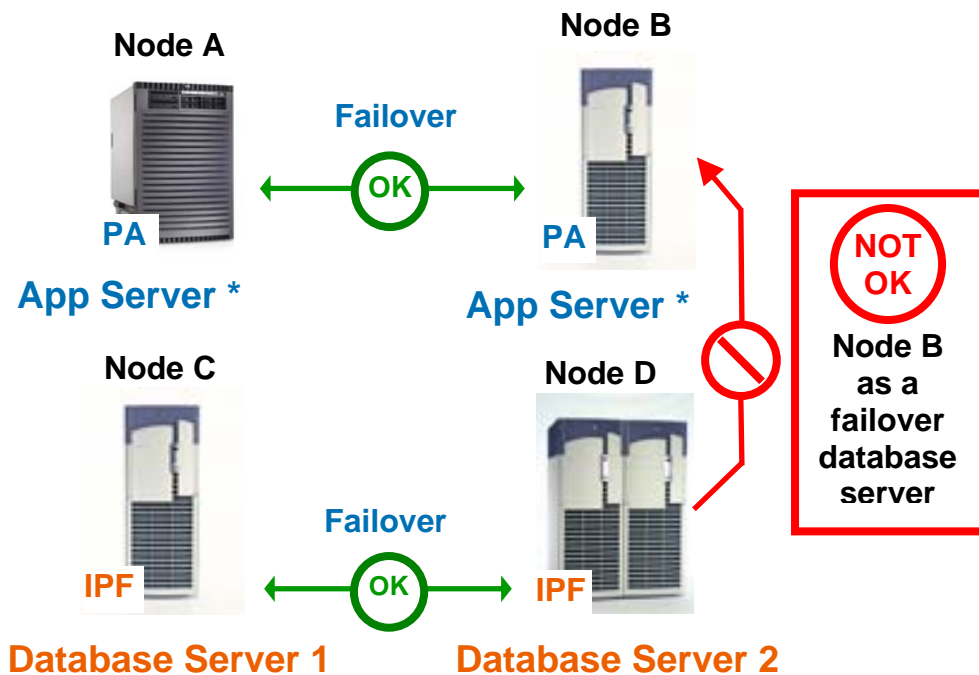
After the two new Itanium servers are added to the cluster, the database and application server packages are reconfigured. The database server is completely transitioned to the new Integrity servers. The original primary database server (node B) gets redeployed as an application server. The failover scenarios are:

- Node A fails.
 - No failover occurs; node B takes over application server load for all users;
 - Optionally, the application server instance from node A fails over to node B if more preferred by the application.
- Node B fails.
 - No failover occurs; node A takes over the application server load for all users.
 - Optionally, the application server instance from node B fails over to node A if more preferred by the application.
- Node C fails.
 - Database server 1 fails over to node D; node D runs both database server 1 and 2.
- Node D fails:
 - Database server 2 fails over to node C; node C runs both database server 1 and 2.

Note

Because the new Itanium nodes C and D run databases, no failover can happen to node A and B, which are HP 9000 systems.

Figure 5. New four-node mixed PA/Itanium cluster



Adding one critical application and two new Itanium nodes to existing four-node PA cluster

This example shows a homogeneous four-node PA cluster with one mission-critical application on each node. Because of resource constraints and new application deployment, the cluster must be upgraded.

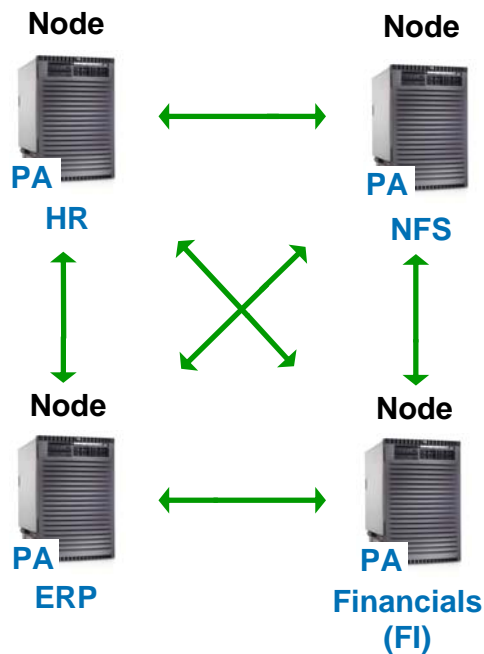
Original four-node PA cluster with four critical applications

Each node is configured with one mission-critical application:

- Human Resource (HR)
- Enterprise Resource Planning (ERP)
- Financial (FI)
- Network File Sharing (NFS) system

All applications, except NFS, have their own databases. The cluster is configured symmetrically, which, in theory, allows any host to run any application. In reality, each node can hardly meet its performance expectation by just running one application. Reduced performance must be accepted in a failover scenario when just two applications are running on one node.

Figure 6. Existing four-node PA cluster



After being in production for some time, the customer faces the following challenges:

- The current PA cluster is missing its service level agreements (SLA) during failover times because all four nodes became fully utilized over time by their primary applications.
- One additional mission-critical application—a Customer Relationship Management (CRM) system—must be introduced and made highly available.
- Further investment in HP 9000 servers is prohibited in favor of HP Integrity servers.

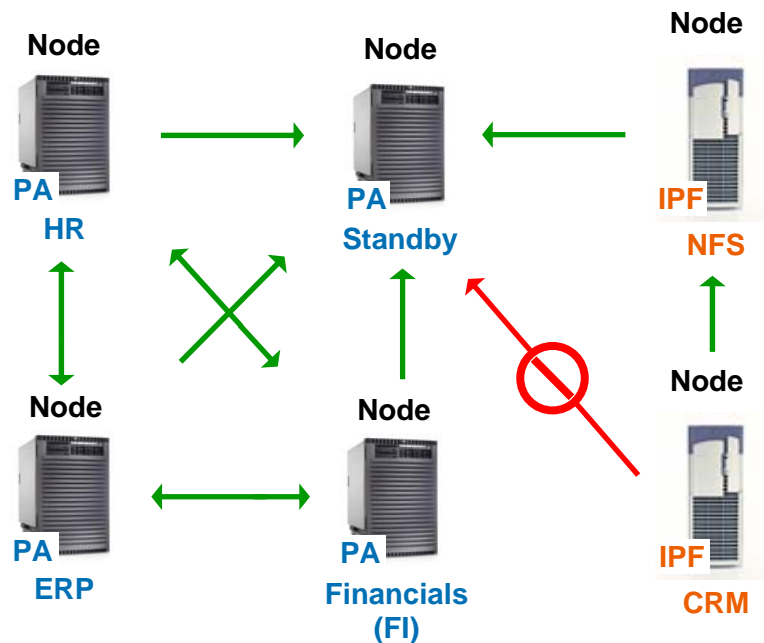
The challenges are met by adding two new HP Integrity servers to the two-node PA cluster.

New six-node mixed PA/Itanium cluster

Two new HP Integrity servers are added to the four-node homogeneous PA cluster. The following high-level changes are performed:

- The highly available NFS application is moved from node C to node E—the first HP Integrity server added to the cluster. This addition frees node C to provide all resources to a HR, ERP, FI, or NFS in case either of the nodes A, B, D, or E fail. Alternatively, a non-critical application could be placed on node C, which will be either shut down or limited in resource entitlements in case of a primary node failure (A, B, D, or E).
- The newly introduced CRM application is placed on node F—the second HP Integrity server added to the cluster. Because this application consists of a database, it cannot fail over to node C, an HP 9000 system.
- The first HP Integrity server (node E) is configured to be the failover node for the CRM application. Initially, it is expected that node E is powerful enough to meet the resource requirements, even in a failover scenario when CRM and NFS run on it together. If the combined resource requirements of NFS and CRM later increase above the level at which node E can successfully handle them, NFS can be pushed to node C.

Figure 7. New six-node mixed PA/Itanium cluster



Mixed clusters as transition aid from PA to Itanium in a multitier SAP environment

This example shows how a customer who initially uses a homogeneous cluster of four HP 9000 nodes and takes advantage of the mixed cluster support to transition to an Itanium cluster in several steps.

Initial four-node homogenous PA cluster

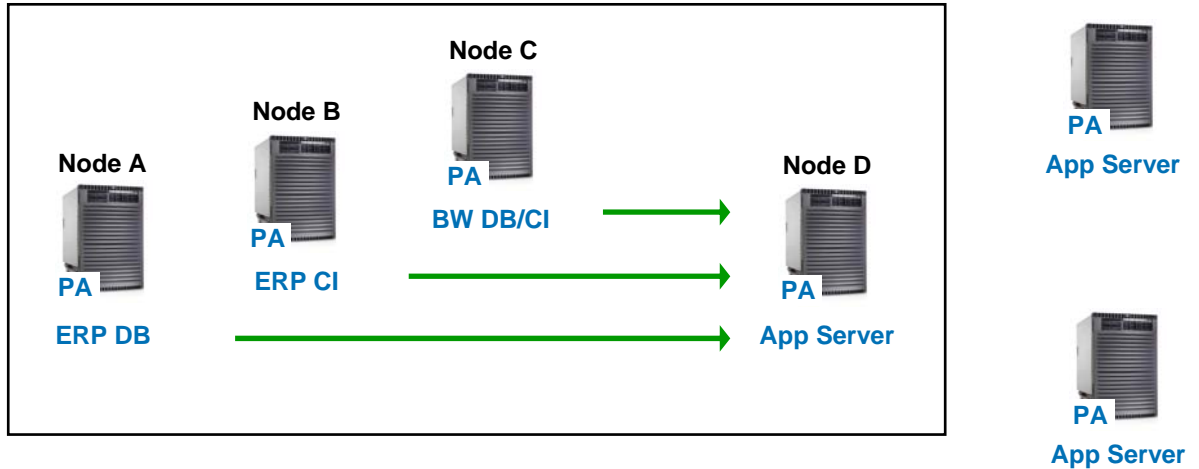
The initial cluster consists of four HP 9000 nodes, of which three are running mission-critical SAP systems:

- Node A runs the database of the ERP system.
- Node B runs the central instance of the ERP system.
- Node C runs the database and central instance of the business warehouse (BW)—the SAP name for the data warehouse application, which is used for online reporting—in a single HP Serviceguard package.

The fourth node (node D) runs a non-critical application server of the ERP system at normal times and is configured to be the failover node for the applications running on node A, B, and C.

Before one of the three critical applications fails over to node D, the non-critical application server on node D will be shut down or optionally, its resource entitlements will be reduced using WLM.

Figure 8. Initial four-node homogeneous PA cluster



First step: adding the first Itanium node to convert to a mixed cluster

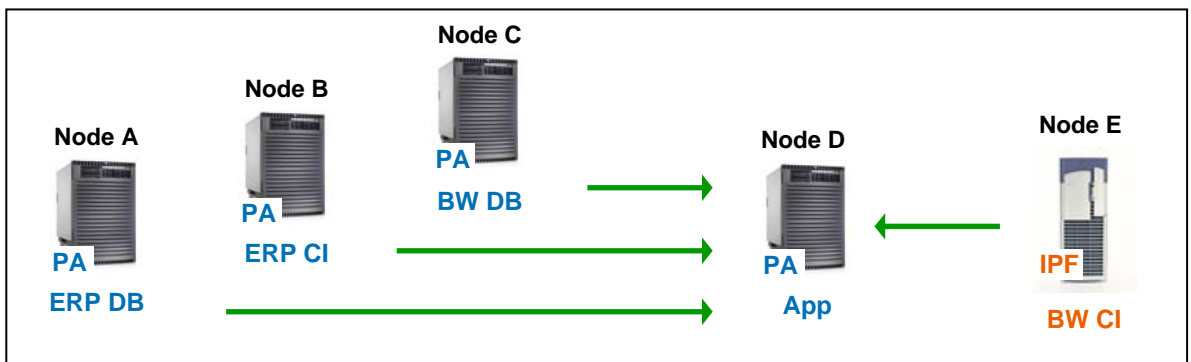
After being in production for some time, the customer faces the following challenges:

- The resource requirements of the BW systems did increase to a level at which node C cannot satisfy the requests in a timely manner.
- Node C cannot be upgraded because further investment in HP 9000 servers is not allowed.

The requirements are met by adding one new HP Integrity servers to the four-node PA cluster and splitting the BW central instance from the BW database.

The following figure shows the new mixed cluster. The additional application servers that reside outside the cluster still remain in the configuration but are not shown.

Figure 9. New five-node PA/Itanium SAP cluster with ERP and BW system



The configuration of node A and B, with the ERP database and central instance, remains unchanged.

The database and central instance of the BW system are split in two packages. The BW database stays on node C, while the BW central instance is configured to run on the newly added Itanium node (node E). Because the BW central instance does not run any database, it can fail over between PA and Itanium nodes.

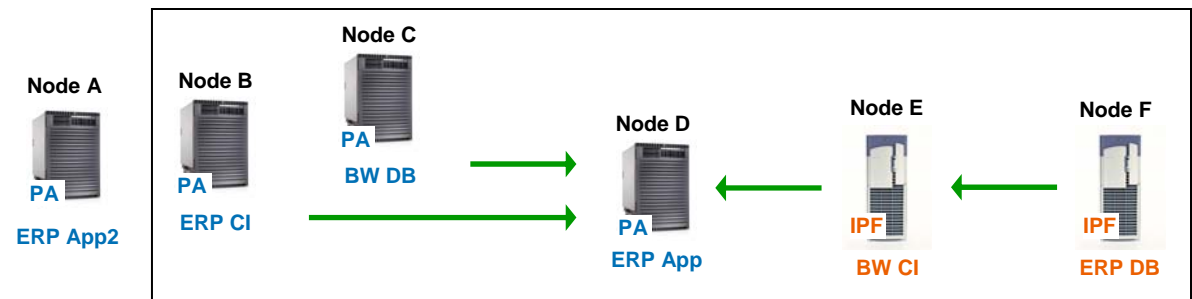
Node D remains as the common failover node for all applications that normally run on node A, B, C, or E. Its own non-critical application will be shut down or limited in resource entitlements before a mission-critical application fails over.

Second step: adding one more Itanium node to mixed cluster

After a couple of months running the mixed cluster in production, the resource requirements of the ERP database system increase beyond the capabilities of node A. The positive experience gained with the first HP Integrity server make it a natural choice to continue adding Itanium systems to the mixed cluster.

The following figure shows the new mixed cluster. The ERP database is moved from node A to the newly added node F. Node A—now without a mission-critical application—is removed from the cluster and redeployed as additional application server outside the cluster.

Figure 10. New five-node mixed PA/Itanium cluster



The configuration of nodes B, C, and D, including their applications, are unchanged.

The newly added HP Integrity server (node F) hosts the ERP database. Because this database now runs on an Itanium system, node D, which is an HP 9000 system, can no longer function as a failover node for the ERP database.

The previously added HP Integrity server (node E), which runs the BW central instance at normal times, will be configured as a failover node for the ERP database.

The WLM configuration can be extended to node E to grant resources to the ERP database when it fails over to node E. Alternatively, the BW central instance could be moved to node D before the ERP database fails over from node F to E to completely eliminate performance impact on the BW system.

Nodes E and F could perform bidirectional failover; however, to avoid performance problems on the BW central instance server that could affect all BW users, HP recommends failing over the BW central instance to the less utilized ERP application server (node D), rather than the highly active ERP database server (node F).

Third step: converting the mixed cluster to a homogenous Itanium cluster

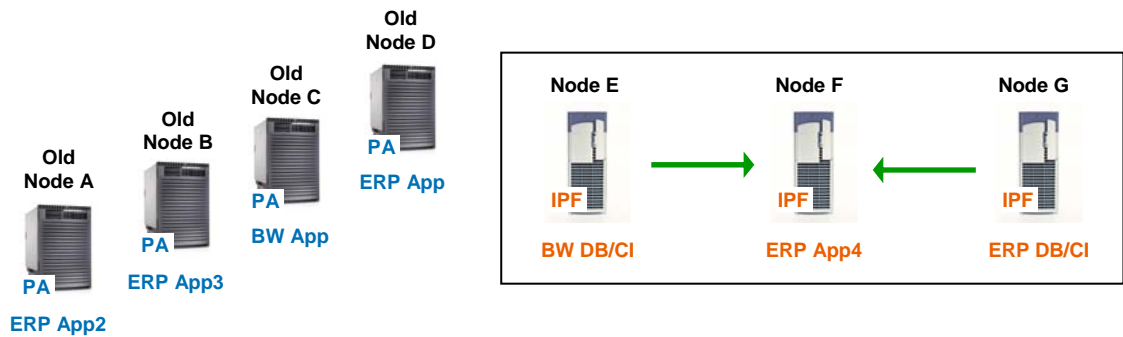
The new HP Integrity servers perform well and have plenty of resources left idle. On the other side, the application layer of both the ERP and the BW system start showing resource shortages.

The customer would also like to consolidate the database and the central instance for each of the systems into a single package—one for the ERP system and one for the BW system—to simplify operation and reduce the number of systems running mission-critical applications that pose a single point of failure (SPOF) in the overall configuration.

The customer achieves this goal by:

- Consolidating the BW database and central instance in one package and placing it on node E
- Consolidating ERP database and central instance in one package and placing it on node F
- Removing the old HP 9000 servers nodes B, C, and D from the cluster and redeploying them as application servers of the ERP and BW system outside the cluster to add capacity on their application layers
- Adding one more HP Integrity server (node G) to the cluster, which will host a non-critical application server of the ERP instance and function as a failover node for the consolidated BW database and central instance from node E and the consolidated ERP database and central instance from node F

Figure 11. Final three-node homogeneous IPF cluster



This example shows how the support of mixed clusters enables the customer to transition from PA- to Itanium-based systems over time at the pace fitting business needs.

PA to Itanium transition service utilizing HP Serviceguard cluster technology

Another scenario exists in which a customer explicitly uses HP Serviceguard cluster technology for transitioning an application from HP 9000 to Integrity servers.

A transition process would include the following steps:

- Create a one-node cluster with the existing HP 9000 server and package the application.
- Add a new HP Integrity server to the cluster.
- Move the application package from the PA node to the Itanium node.
- Test the various failover scenarios and configurations.
 - If the tests fail, move the package back to the PA node.
 - Resolve the issue and start another iteration (move and test package).
- Remove the PA node from the cluster.

The benefit of this process is that the application can be quickly moved back to the HP 9000 system in case of issues showing up with the application on the Integrity server.

Specific ISV information

ISV-specific information with regard to mixed PA and Itanium clusters is currently available for Oracle and SAP.

Oracle 9i/10g database server

Our understanding of the current Oracle position on mixed clusters is:

- Oracle supports **migration** of a database from PA to Itanium. The process includes backing up the data on the HP 9000 system, restoring it on the Integrity system, and replacing the PA binaries with Itanium binaries.
- Oracle does not support **failing** over a database from PA to Itanium or from Itanium to PA.
- Oracle has not tested applying archive logs from PA systems to Itanium systems (a necessary step in database failover).
- Oracle does not support running the PA database binaries under the Aries emulator on HP Integrity servers.
- Only Oracle 9iR2 or later will be supported, with the PA version running on PA systems and the Itanium version running on Itanium systems.

Unless Oracle supports failover between PA and Itanium nodes in the future, HP suggests the following mixed cluster “rules” for Oracle databases:

- Oracle database support may be restricted to allow single-instance Oracle failover between PA nodes only and failover between Itanium nodes only. However, a database on either architecture should support application servers on both architectures, which is currently the case if the application is SAP, no matter if the database or the application servers run inside a cluster or not.
- Mixed clusters are not currently supported with Oracle Real Application Cluster (RAC) and HP Serviceguard Extension for RAC (SGeRAC). HP is investigating support of this environment in the future.

SAP with HP Serviceguard Extension for SAP

The HP/SAP Competence Center has already validated and tested mixed clusters and will add support for it to HP Serviceguard Extension for SAP (SGeSAP) B.03.11.

SAP Application Server and SAP Central Instance packages (APP and CI package types) are supported in mixed environments with SGeSAP B.03.11. They can be configured to fail over between heterogeneous nodes.

A database package (DB and DBCI package types) must fail over between homogenous nodes—either HP 9000 only or Integrity server only.

The same applies for Livecache packages (LC package type) because it also contains a database part.

The details about the supported configurations depend largely on the Oracle support statement, particularly the configurations in which a single SGeSAP package includes both the Oracle database instance and SAP applications—like the package types DB, DBCI, and LC.

How to implement a mixed PA/Itanium HP Serviceguard cluster

There are several paths to get to a supported mixed PA/Itanium HP Serviceguard A.11.16 cluster on HP-UX 11i v2UD2. The starting point is always a currently supported configuration. The transition process details depend on the current customer configuration. Two major steps are involved:

- All systems must run HP-UX 11i v2UD2.
 - Upgrade the systems to HP-UX 11i v2UD2 before adding them to the cluster.

- If the systems are already part of an existing cluster, use the rolling upgrade process to upgrade the systems to HP-UX 11i v2UD2 (from any previous version of HP-UX 11i).
- All nodes must have HP Serviceguard A.11.16 installed.
 - Install HP Serviceguard A.11.16 before adding a new node to a mixed PA/Itanium cluster.
 - If the systems are already part of an existing cluster, use the rolling upgrade process to bring all nodes to HP Serviceguard A.11.16.

Starting with an existing cluster

Customers interested in mixed PA/Itanium HP Serviceguard clusters most probably have one or more clusters in place already. Depending on whether the customer uses only systems of a single architecture or PA and Itanium servers in different clusters, the procedure for how to achieve a mixed cluster varies, as described in the following sections.

Starting with an existing HP-UX 11i v1 PA cluster

The most typical scenario in which customers would like to deploy mixed clusters is when existing PA clusters with HP Serviceguard 11.15 on HP-UX 11i v1 are already in use and the customer wants to transition to Itanium with individual applications. The following high-level process describes the necessary steps:

- Upgrade the server to HP-UX 11i v2UD2 and HP Serviceguard A.11.16 through the rolling upgrade process as described in Appendix E of the HP Serviceguard user manual:
 - Shut down the application packages running on a specific node and restart them on another node.
 - Perform HP-UX and HP Serviceguard upgrades.
 - Move applications back to their original node and test failure scenarios.
 - Continue the process through the other nodes in the cluster.
- Add new Itanium systems (with HP-UX 11i v2UD2 and HP Serviceguard A.11.16 already installed) to the existing cluster.
- Reconfigure the cluster/package to consider new requirements:
 - New and/or changed primary and failover nodes
 - Special requirements caused by two sets of application binaries
- Test the various failover scenarios and configurations.

Starting with an existing HP-UX 11i v1.6 or v2 Itanium cluster

This use case is much less likely than the one described previously. Because the HP-UX version is probably already up to date, it is only the HP Serviceguard version that might be upgraded to A.11.16 through the rolling upgrade process. In the rare circumstance that the nodes are on HP-UX 11i v1.6 or HP-UX 11i v2, they must be upgraded to HP-UX 11i v2UD2. This upgrade process is similar to the rolling upgrade process of PA nodes from HP-UX 11i v1 to HP-UX 11i v2UD2 as described previously. The process of converting an Itanium cluster to a mixed PA/Itanium cluster is similar to converting a PA cluster to a mixed PA/Itanium cluster, which involves the following steps:

- Upgrade the server to HP-UX 11i v2UD2 and HP Serviceguard A.11.16 through the rolling upgrade process as described in Appendix E of the HP Serviceguard user manual.
 - Shut down the application packages running on a specific node and restart them on another node.
 - Perform HP-UX and HP Serviceguard upgrades.
 - Move applications back to their original node and test failure scenarios.
 - Continue the process through the other nodes in the cluster.
- Add new PA systems (with HP-UX 11i v2UD2 and HP Serviceguard A.11.16 already installed) to the existing cluster.

- Reconfigure the cluster/package to consider new requirements:
 - New and/or changed primary and failover nodes
 - Special requirements caused by two sets of application binaries
- Test the various failover scenarios and configurations.

Merging an existing PA cluster with an existing Itanium cluster

In the scenario that a customer already deploys different clusters of both architectures and would like to merge a PA with an Itanium cluster, a few additional steps are required beyond the ones listed previously. In this scenario, the PA cluster could be on HP-UX 11i v1 (most likely) or HP-UX 11i v2UD2 already., and the Itanium cluster could be on HP-UX 11i v2 (most likely), HP-UX 11i v1.6, or HP-UX 11i v2UD2 already. The following steps are involved:

- If nodes of both clusters are not already on HP-UX 11i v2UD2 and HP Serviceguard A.11.16, upgrade them to HP-UX 11i v2UD2 and HP Serviceguard A.11.16 through the rolling upgrade process within their homogeneous cluster, as described in Appendix E of the HP Serviceguard user manual:
 - Shut down the application packages running on a specific node and restart them on another node.
 - Perform HP-UX and Serviceguard upgrades.
 - Move applications back to their original node.
- Choose one of the clusters as the remaining cluster. Which of the original clusters will be chosen is expected to be case-specific. The decision can be based on the cluster with the:
 - Most critical application
 - Largest number of nodes
 - Least outstanding changes
- Move the nodes from the cluster that will not survive to the intended mixed PA/Itanium cluster, one by one, following these steps:
 - Shut down the application packages running on a specific node and restart them on another node.
 - Remove the node from the to-be-dismissed cluster.
 - Add the node to the intended mixed cluster.
 - Decide, on a per-package level after each node is moved to the mixed cluster, whether to leave the package on the to-be-dismissed cluster or transition it to the mixed cluster.
- Reconfigure the package and move it to the mixed cluster if not already done in the step before.
- Consider mixed requirements. The application might have two sets of binaries, if available.
- Test the various failover scenarios and configurations.

Installing a new mixed cluster

When no cluster exists, there is little to consider because nodes of both architectures form the cluster. One point to address is the requirements caused by two sets of application binaries.

You must have HP-UX 11i v2UD2 and HP Serviceguard A.11.16 or newer installed on all systems, which can be achieved by installing the MCOE on all systems before creating the cluster or by installing the base operating system, HP Serviceguard, and desired ISV individually.

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

5982-8980EN, 10/2004

