High Availability Choices in HP-UX 11i Versions 2 and 3

Tutorial # 3517

Bob Sauers / Chief Architect
Availability Clusters Solutions Lab
HP High Availability Advanced Technology Center (HA ATC)

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice

Agenda

- HP High Availability Solutions
- HP Disaster Tolerant Solutions
- Serviceguard and Oracle 10g
- High Availability Solutions in HP-UX 11i v2UP2
- Comparing Serviceguard and TruCluster Features
- High Availability Choices in HP-UX 11i v3
- Summary

HP High Availability Solutions

High Availability Products

- Serviceguard
  - Serviceguard Clusters
  - Serviceguard Manager (GUI)
  - HA & Partitions
  - HA & Resource Management with Workload Manager
  - ECM Toolkit
- Serviceguard Extension for Faster Failover (SGeFF)
- Serviceguard Extension for SAP (SGeSAP)
- Serviceguard Extension for RAC (SGeRAC)

Serviceguard

- Single SG Cluster up to 16 nodes
- For use when all nodes are in a single Data Center
- Automatic failover, up to 150 application packages (up to 900 services total); supports up to 200 relocatable package IP addresses per cluster
- Cluster within a single data center
- SCSI or Fibre Channel for disks
- Single IP subnet for heartbeat networks (Ethernet, FDDI, Token Ring) (NOTE: DWDM will provide extended distances for Ethernet)
- IPv6 support (data links only)
  Heartbeat must still use an IPv4 network.
  In SG version 11.15, package IPs may be either IPv4 or IPv6)
- LVM or VxVM, MirrorDisk/UX (for LVM) optional
- Cluster lock disk required for 2 nodes, optional for 3-4 nodes only
- Alternatively can use Quorum Server in place of cluster lock disk with up to 16 nodes

Failover Models

- Active / Active
  - All nodes are running (different) applications
  - Upon failover, choice of
    - Reduced capacity when multiple applications run on the same node
    - Shutdown less critical applications
    - Optional use of PRM and WLM to guarantee resource entitlements
- Active / Standby
  - One or more nodes are reserved for failover use
  - Upon failover, the applications maintain performance due to spare capacity
- Rotating Standby
  - Upon failover, the standby system becomes the new production system and the repaired system becomes the new standby system

Serviceguard for Linux (SG/LX)

- Package Failover (for node, network, or service)
- Support up to 150 packages and 900 services per cluster (previously up to 30 packages with 30 services per package)
- Heartbeat over Ethernet, supporting up to 7 Heartbeat subnets
- Local Network Switch (Bonding)
- Online Reconfiguration
- Rolling Upgrade Infrastructure
- Quorum Service — supporting maximum of 100 nodes/50 clusters (previously maximum of 2 clusters, 8 nodes)
- Toolkits: NFS, Apache, Samba, SendMail, and Oracle, MySQL
- SG/Mgr (localizable, and capable of monitoring and polling Linux and HP-UX clusters) + administration, multi-subnets and alert summary
- Disk Monitor
- Support reiser file system (a Journaled file system)
- Parallel fsck
- Improved installation process (dynamically loadable modules)

SG Linux FC Configuration: Supports up to 16 nodes

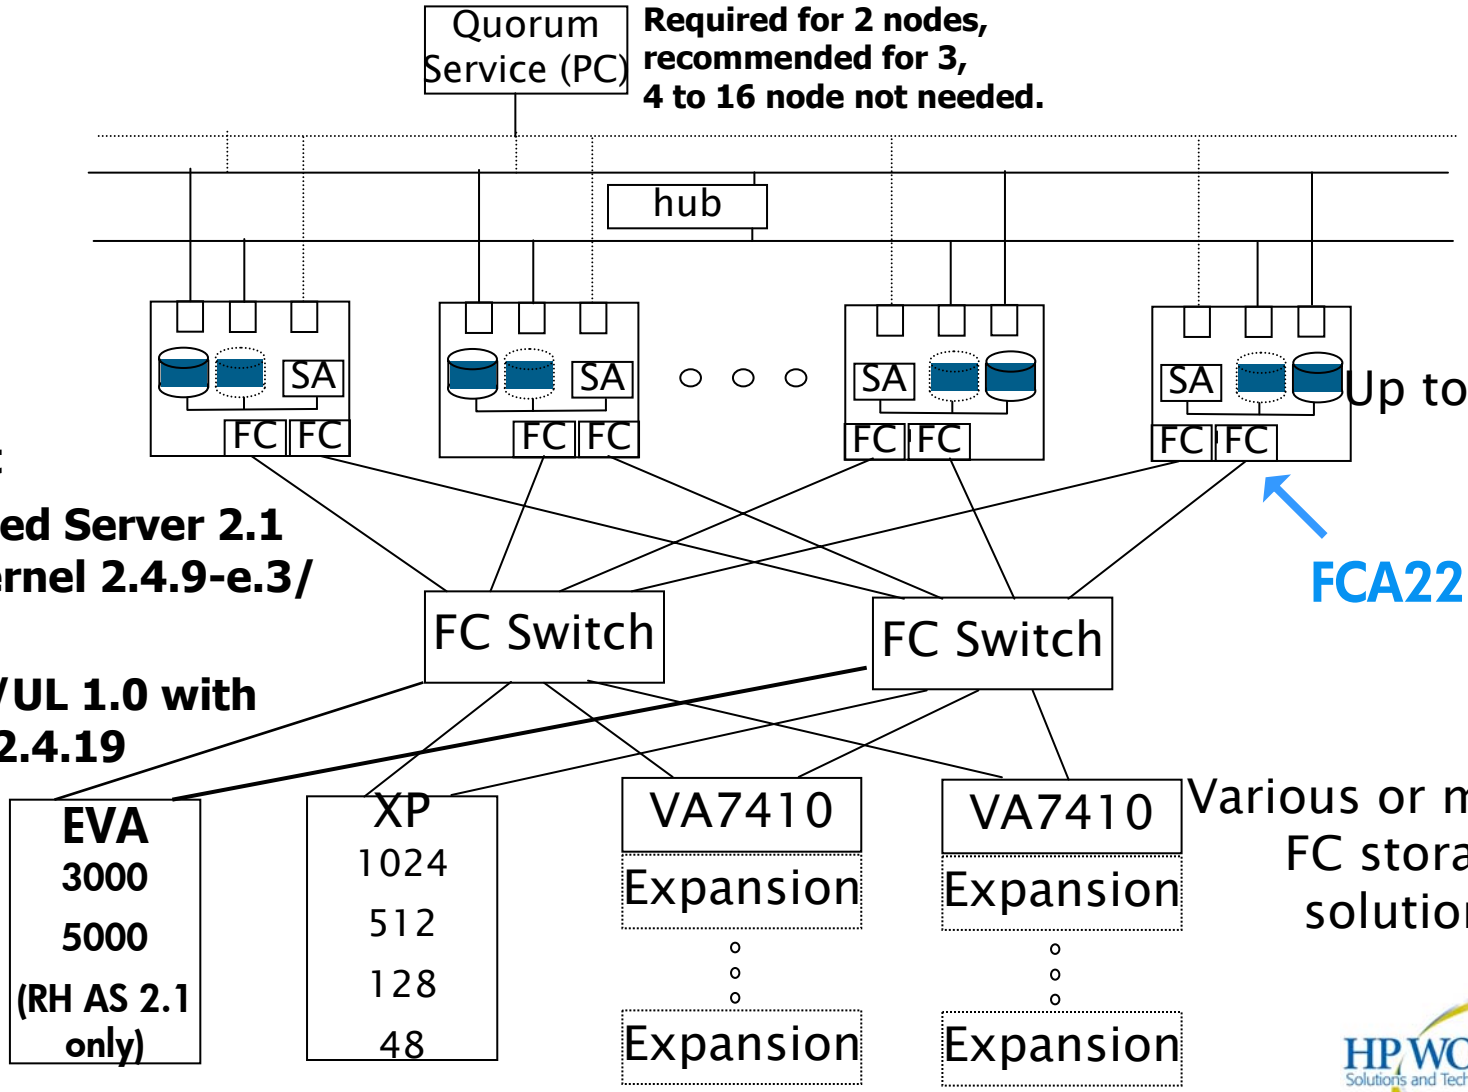

RedHat

Advanced Server 2.1 with kernel 2.4.9-e.3/ -e.25,

SLES 8/UL 1.0 with kernel 2.4.19

Protecting against split brain and data corruption

- Possible tie-breakers:
  - Cluster lock disk
    - single cluster lock disk (when all servers are in a single data center)
    - dual cluster lock disks (when the servers are distributed across two data centers)
  - Arbitrator systems
    - one or two Arbitrator systems which are fully participating members of the cluster (when the servers are distributed across three data centers)
  - Quorum Server
    - a (small) server that is outside of the cluster

Serviceguard Manager

- An intuitive and easy-to-use Java©-based GUI to display HP Serviceguard clusters on both HP-UX and Linux
  - Integrated with Service Control Manager and OpenView/Operations & NNM
  - Uses color-coded icons to show status information about a cluster, node or package
  - Provides GUI for:
    - Monitoring
    - Administration / management (start/stop/move)
    - Configuration (cluster and package)
  - Role-based access rules for segregating users – Example:
    - Operators can only monitor the clusters
    - Level 2 Support can start/stop/move application packages
    - Administrators have full capabilities
  - It's free !
  - New version planned for integration with System Insight Manager on HP-UX 11i v3

Serviceguard Manager example

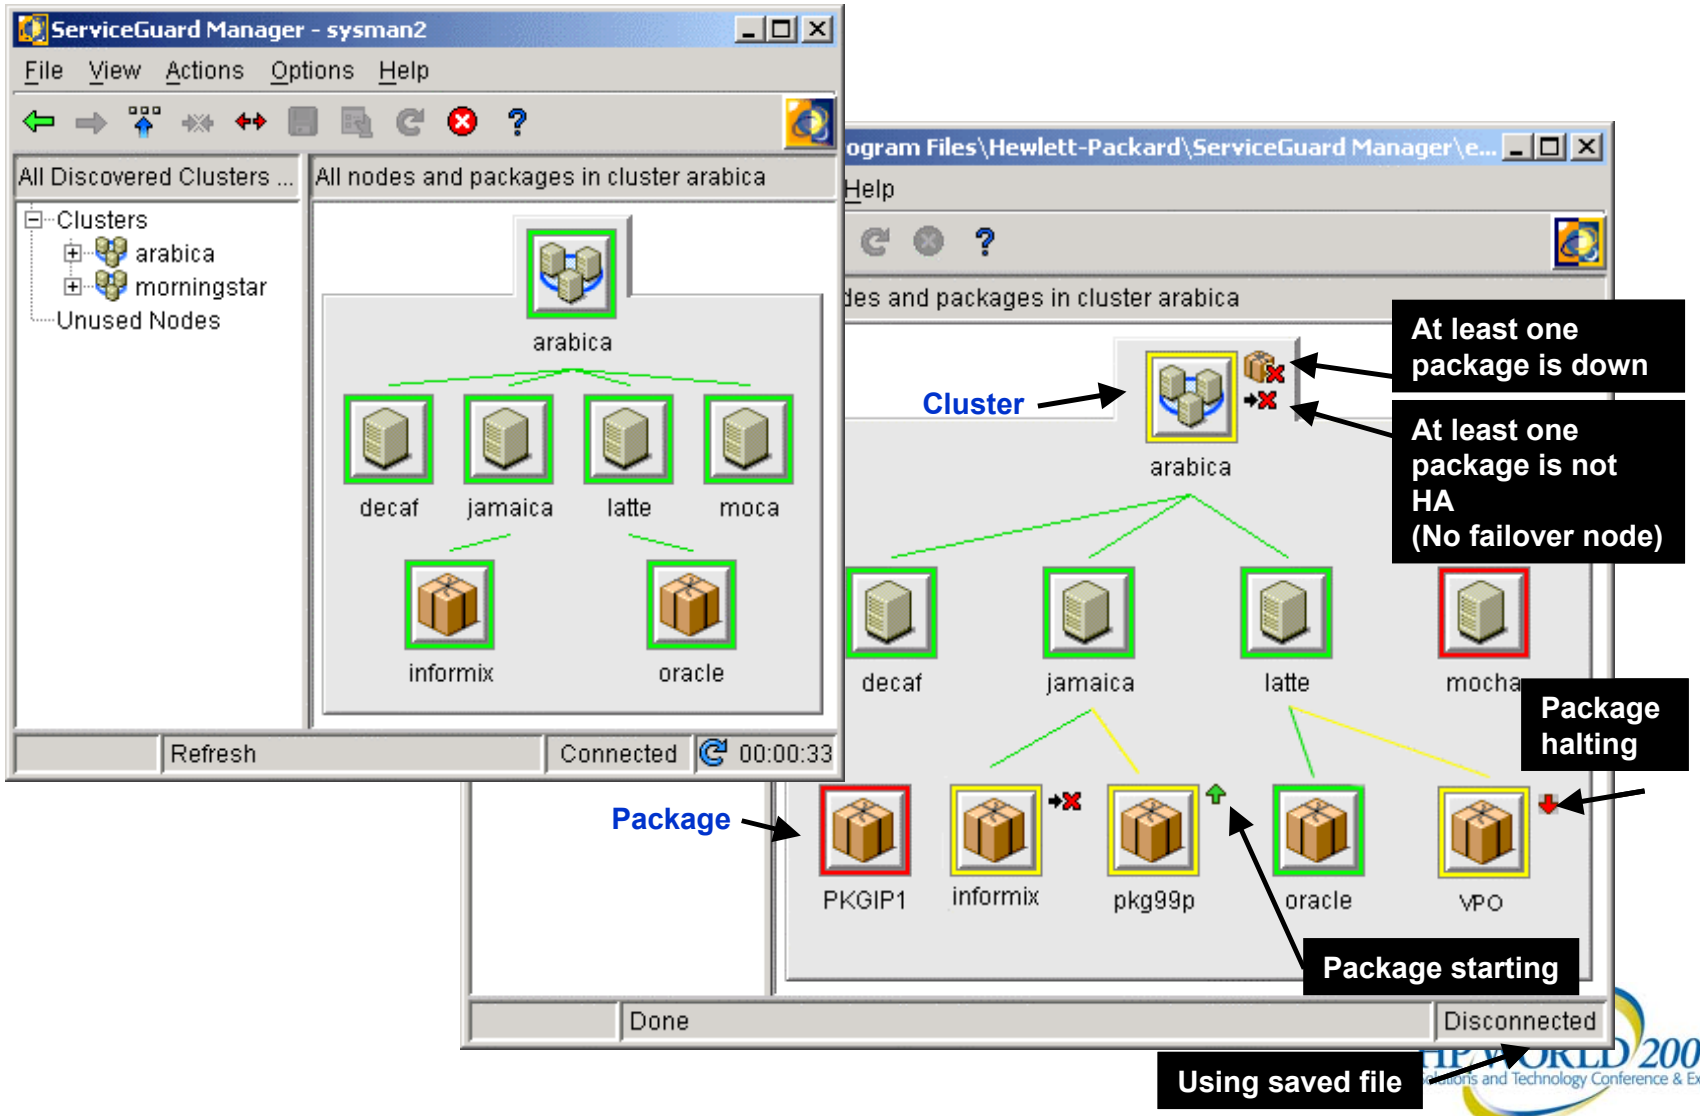

HA and Virtual Server Environment

- Serviceguard integrates with the Virtual Server Environment (VSE):
  - nPARs (hard partitions with electrical isolation)
  - vPARs (soft or virtual partitions with software isolation)
  - Workload Manager (WLM) - automatically adjust resource (CPU, Memory, Disk) allocation in normal and failover situations

Workload Manager (WLM) and Process Resource Manager (PRM) with SG

- Service Level Management solutions that can be easily integrated into an Serviceguard environment to provide management of computing resources (CPU, memory and I/O utilization) according to desired SLOs both before and after application failover among nodes

- Applications can be assigned maximum resources when they run on their "preferred" server.

- After a failover to a different node, the resource utilization for the existing and new applications can be dynamically reset to:
  - favor important applications
  - restrict resource utilization of less important applications

Components of Failover Time

- Node Failover Time
  1. failure detection
  2. node selection
  3. node failover

Serviceguard-Specific

- Application Failover Time
  1. volume group activation
  2. file system recovery (if applicable)
  3. file system mount (if applicable)
  4. application startup
  5. application recovery

Application-Specific

(Not affected by SG)

HP Serviceguard Extension for Faster Failover (SGeFF)

NEW ! – July 2004

- Optimized Serviceguard and Serviceguard Extension for RAC environments may achieve improved node failover times of 5 seconds (current failover 30-45 seconds)
- Expected to be used with
  - An application or a database that has fast (or no) startup time
  - An application that has fast recovery (tens of seconds at most), e.g., Oracle RAC
  - Systems with predictable loads
- Restricted Configuration
  - SG A.11.16 and later
  - 2-node cluster only
  - Multiple heartbeat networks
  - Requires quorum server as tie-breaker
  - LVM, SLVM or VxVM (not CVM)
  - Need to qualify network latency over time

Typical applications for SGeFF

- Oracle RAC:
  - The RAC instances are already running on all of the nodes, so there is no application startup time
  - RAC can be tuned for rapid recovery, often in 10-60 seconds
- Other application factors that are well-suited for SGeFF include:
  - An application where there is no disk-based database
  - An application where you can have the processes already started, but idle and waiting until a failover occurs (i.e., Hot Standby) such as in some Telco applications
  - Systems with fewer volume groups
  - Applications using raw volumes rather than file systems -OR- systems with fewer file systems
  - VxFS file systems with the intent log enabled at least for meta-data (note: this is the default for VxFS, and it should not be changed)

Toolkits

- Enterprise Cluster Master Toolkit (ECMT)
  - Fully-tested and supported collection of integration templates for certain popular third-party applications
- HA NFS Toolkit
  - Pre-tested and supported templates to make NFS servers highly available
- SGeSAP – Serviceguard Extension for SAP

Serviceguard extensions for SAP (SGeSAP)

Integrate SAP R3 with:

- Serviceguard
  - Toolkit template for easily configuring SAP with Serviceguard (HP-UX and Linux)
  - Options to on how to configure the Central Instance (CI) and the Database (DB) servers
- Metrocluster (HP-UX only)
  - Optional template to create a disaster tolerant architecture for SAP

(Configuration example shown in the Metrocluster section of this presentation)

ServiceGuard Extension for RAC (SGeRAC)

- Same protection & functionality for applications as Serviceguard
- Parallel database environment for increased availability and scalability with Oracle Real Application Cluster (RAC)
- Up to 16 nodes with Shared Logical Volume Manager (SLVM) and 4 nodes with Cluster Volume Manager (CVM)

Disaster Tolerant Products

Extended/Campus Cluster, Metrocluster and Continentalclusters

Extended/Campus Clusters

- Single SG Cluster up to 4 nodes with 2 data centers, 16 nodes with 3 data centers (NOTE: 4 nodes with CVM near future)
- Automatic failover, up to 150 application packages
- Campus and city distances
  - Between 2 data centers (up to 10 km)
  - Among 3 data centers
    - Up to 10 km between A & B with FibreChannel Hubs
    - Up to 50 km with FDDI networking & Finisar Long Haul FibreChannel GBICs
    - Up to 100 km using DWDM/CWDM
- Fibre Channel for disks
- Single IP subnet for heartbeat networks (Ethernet or FDDI)
- Dedicated links for network and physical data replication (dark fiber) –
  - no telco switched networks
  - no conversion to other protocols like ATM
- MirrorDisk/UX or Veritas Mirroring REQUIRED
- LVM or VxVM for Serviceguard; SLVM (2-nodes) or CVM (4-nodes) for SGeRAC
- Dual cluster lock disks (LVM) required for up to 4 nodes with 2 data centers
- QS or Arbitrators with 3-data center architecture

Extended/Campus Cluster Architectures

- Two data center architecture
  - Limited to 4 nodes
  - Requires dual cluster lock disks
  - Minute risk of split brain
  - Physical data replication using MirrorDisk/UX

- Three data center architecture
  - Limited to 16 nodes
  - Cluster lock disks not used
  - Arbitrators or Quorum Server act as tie breakers
  - No chance of split brain
  - Physical data replication using MirrorDisk/UX or Veritas Mirroring

Extended Campus Cluster (3 data centers)

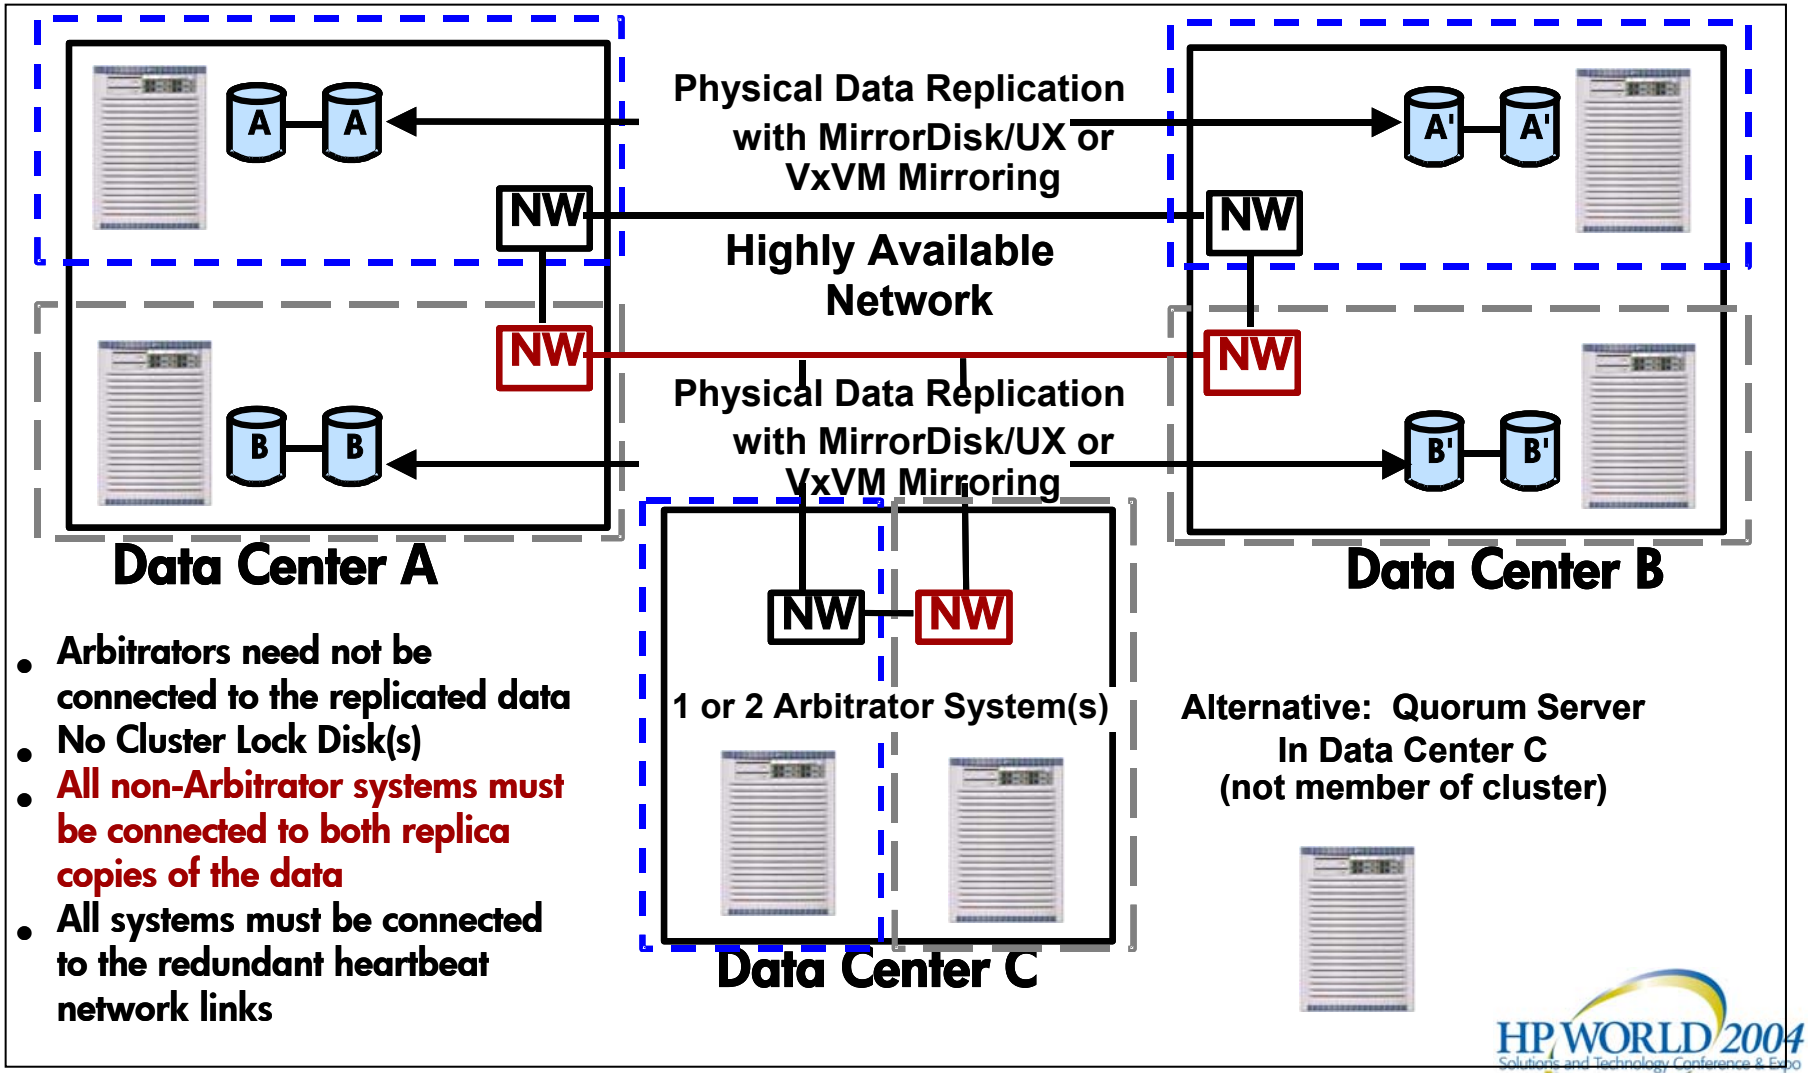

- Arbitrators need not be connected to the replicated data
- No Cluster Lock Disk(s)
- All non-Arbitrator systems must be connected to both replica copies of the data
- All systems must be connected to the redundant heartbeat network links

HP Extended Cluster for RAC
Disaster tolerance with continuous application availability

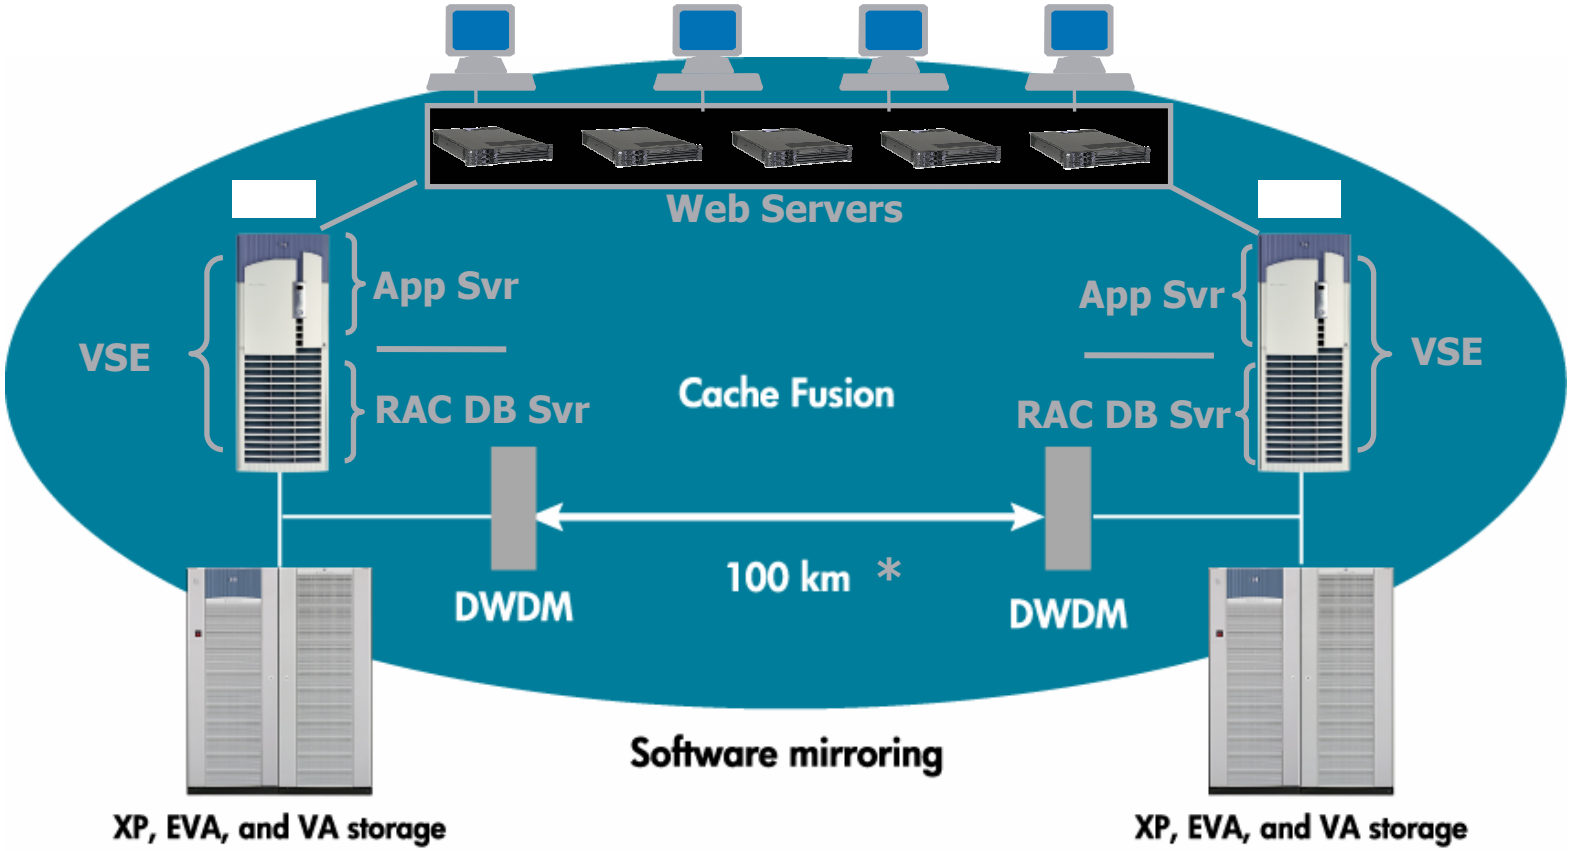

Extended Cluster for RAC

Active/active environment with transparent access to applications and data!

* Plans to extend distance > 100 km

Two Data Center Extended SGeRAC Cluster – up to 10km

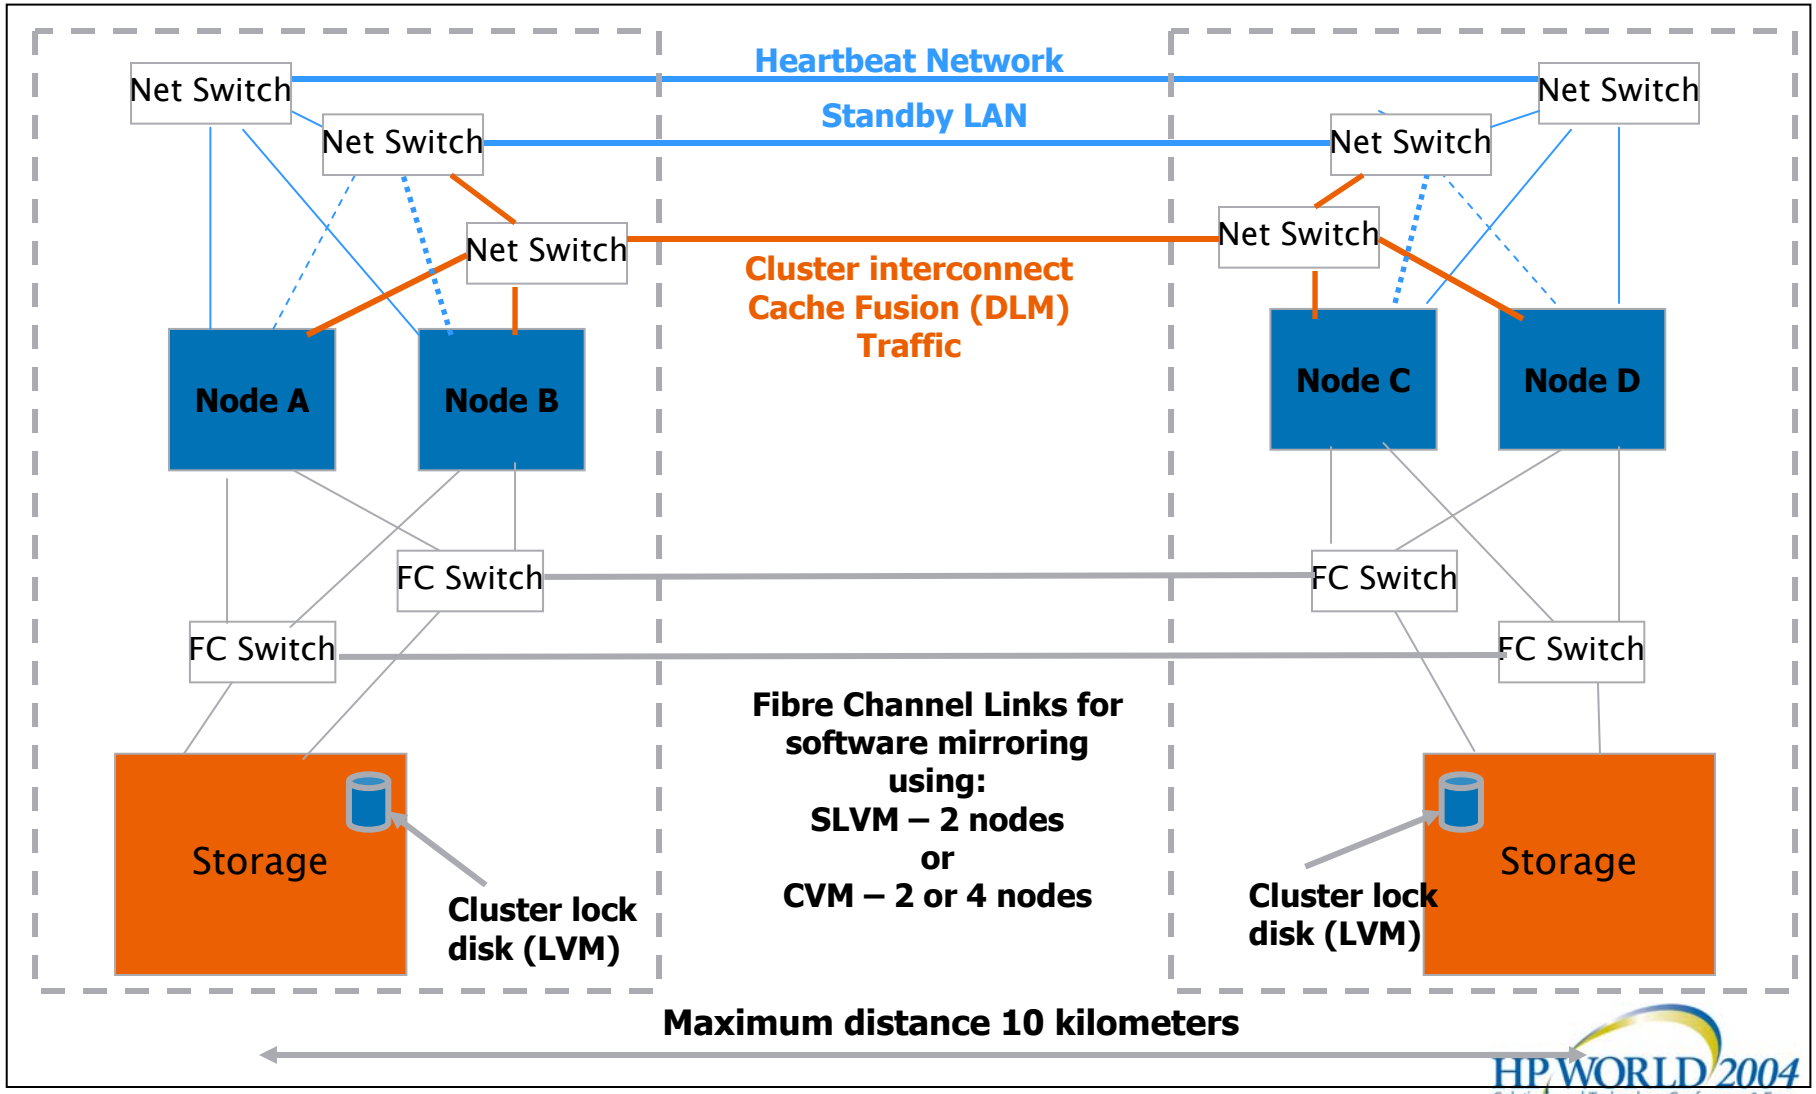

Two Data Center Extended SGeRAC Cluster – up to 100km

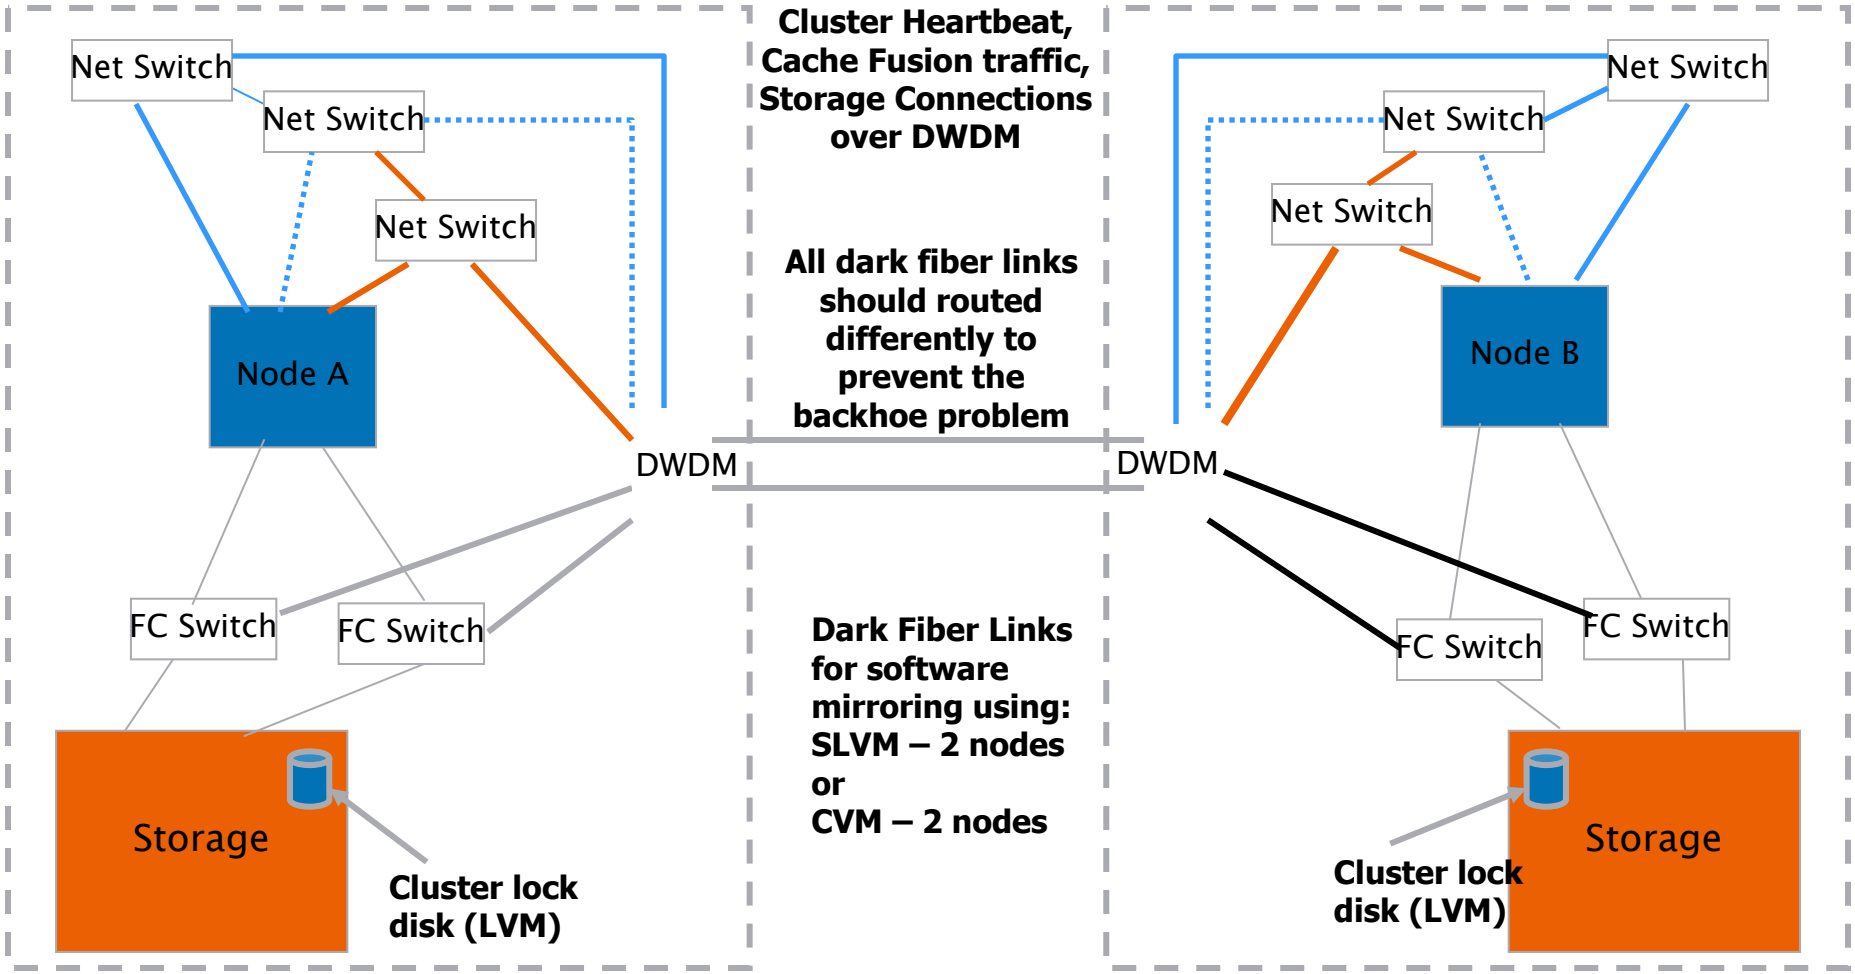

Metrocluster

- Single SG Cluster up to 16 nodes (NO SG OPS)
- Automatic failover, up to 150 application packages among all nodes
- 3 data centers required – Metropolitan (same city) distances
  - up to 50 km between A & B
  - 100 km with DWDM or CWDM
- SCSI or Fibre Channel for local disk connectivity
- Data replication between Data Centers A & B in hardware
  - HP Continuous Access XP (Asynchronous and Synchronous)
  - EMC SRDF (Synchronous ONLY)
  - Toward the end of 2004: EVA CA
- Dedicated links for network and physical data replication (dark fiber) –
  - no telco switched networks
  - no conversion to other protocols like ATM
  - Single IP subnet for heartbeat networks (Ethernet, FDDI)
- MirrorDisk/UX optional for root disks only
- LVM or VxVM ONLY
- No cluster lock disks

Metrocluster

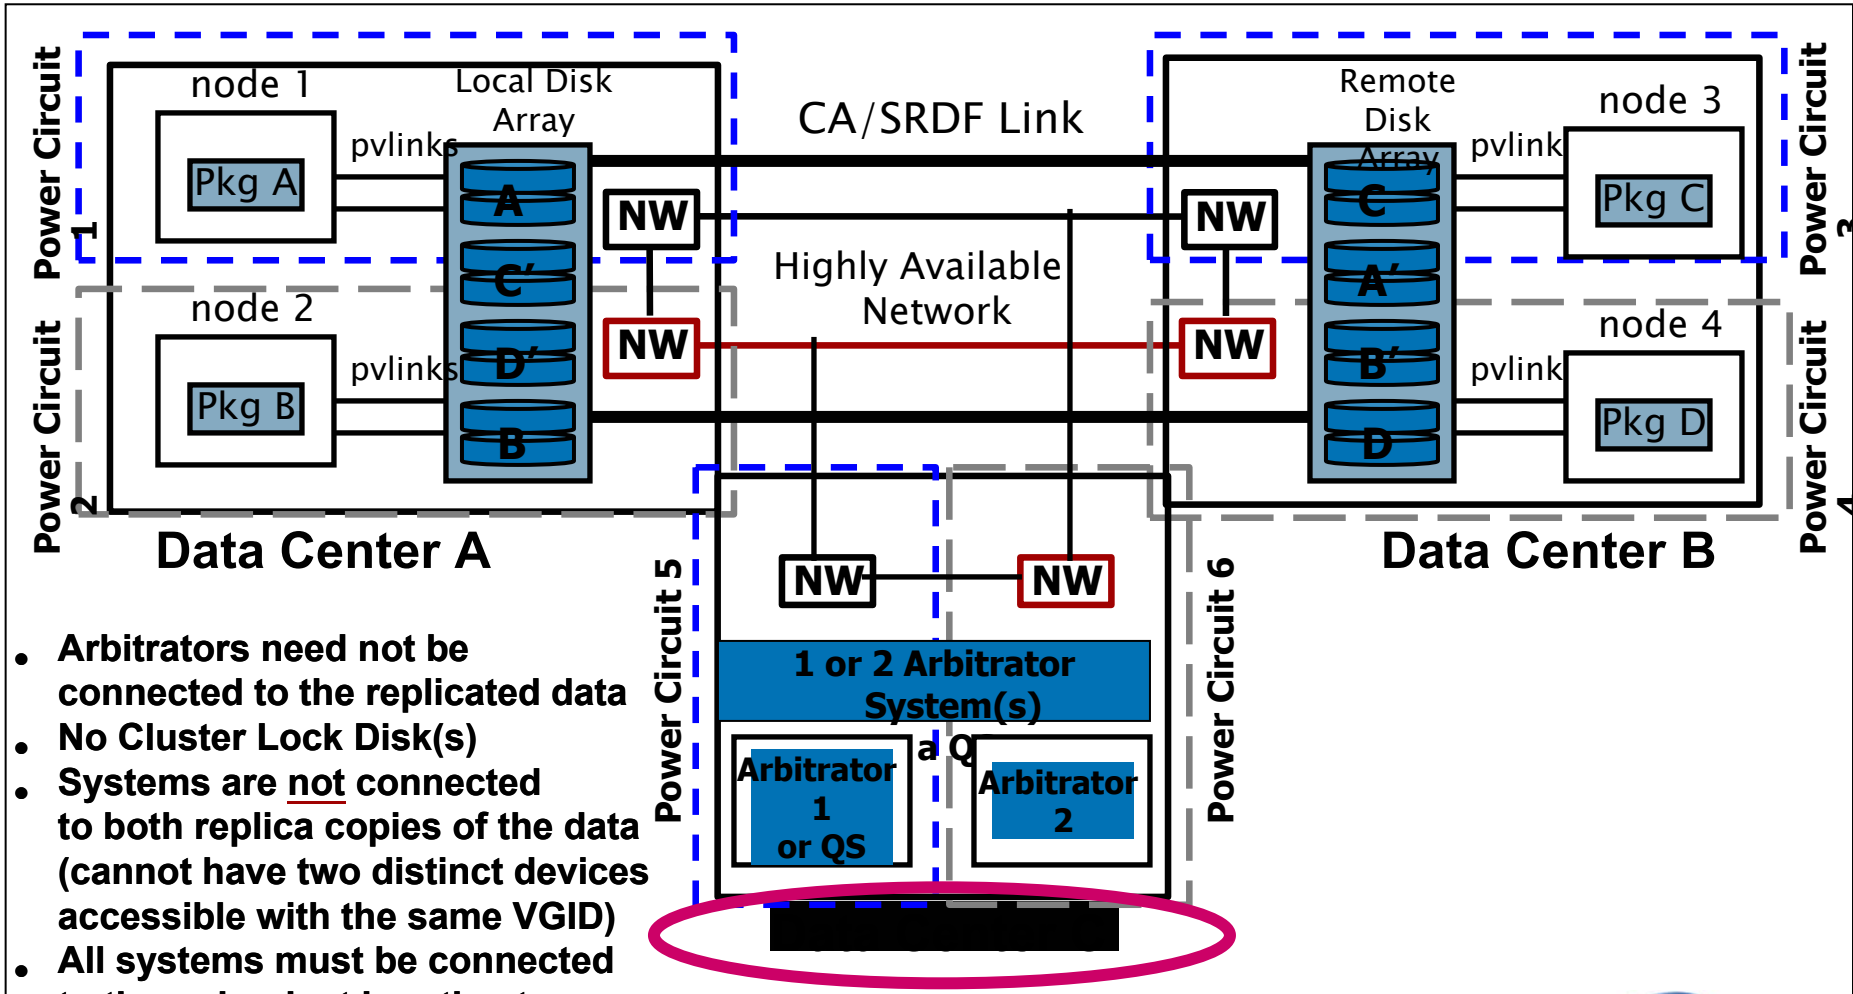

- **Arbitrators need not be connected to the replicated data**
- **No Cluster Lock Disk(s)**
- **Systems are not connected to both replica copies of the data (cannot have two distinct devices accessible with the same VGID)**
- **All systems must be connected to the redundant heartbeat network links**

Metrocluster example: 2 Main data centers & 3rd location instead of 3 data centers

Customer requirements for disaster tolerance dictates the distance to the third location

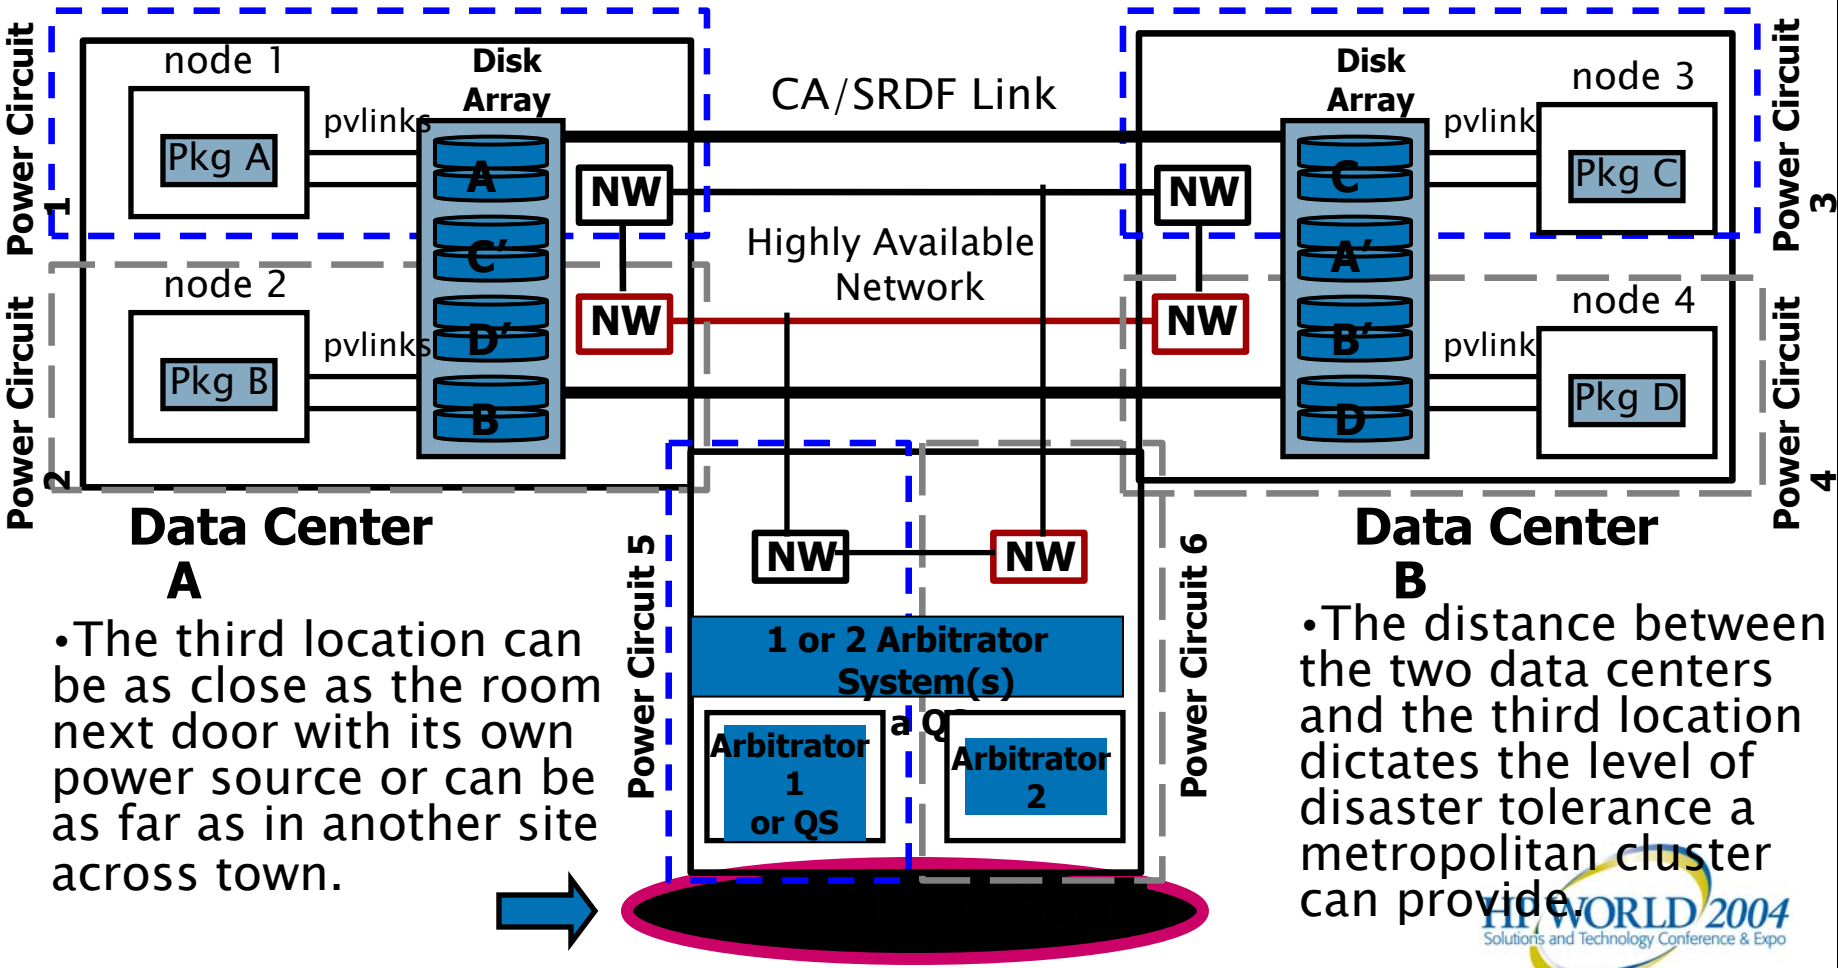

- The third location can be as close as the room next door with its own power source or can be as far as in another site across town.
- The distance between the two data centers and the third location dictates the level of disaster tolerance a metropolitan cluster can provide.

3 active site configuration (with 1-1 data replication) for Metrocluster CA XP

- valid for all supported Metrocluster CA XP versions

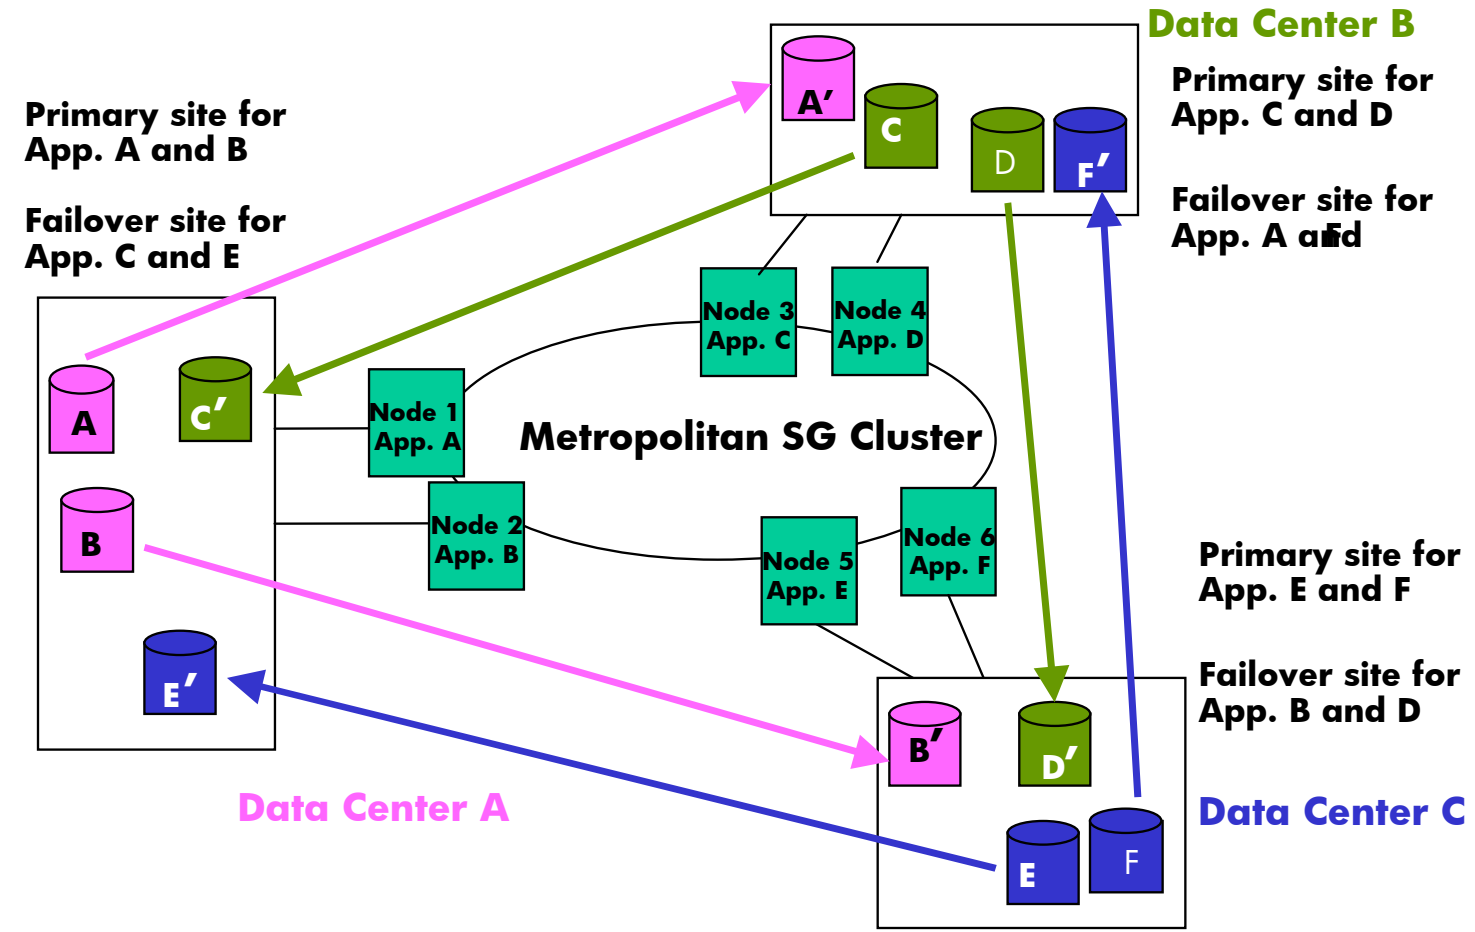

Optional Metrocluster/ XP CA architectures

- Support for switched network links (ATM and IP) for XP/CA data replication
  - Besides dark fiber link, MC/CA now also supports switched network links (ATM and IP) for XP/CA data replication.
  - Enables implementation of a Metrocluster/CA solution in an environment where dark fiber is not available.
  - Basic Requirements
    - Maximum distance between the XP disk arrays is 100 kilometers
    - Network latency for cluster network no more than 200 milliseconds
    - Business Copy at both sites for point-in-time copies to recover from a rolling disaster (due to decreased reliability of these links)
    - Cluster network must be pure Ethernet (i.e., no ATM encapsulated Ethernet packets) with all nodes in the same subnet
    - There must be at least two (redundant) supported converter boxes installed at each site
    - For IP configuration, there must be at least two IP routers or LAN switches installed at each site
    - The Device Group Monitor, which detects link failure and splits the BC, must be configured as an application package

Continentalclusters

- Dual SG Clusters up to 16 nodes each
- Automatic failover, up to 150 application packages within each cluster, push-button failover of pre-configured set of application packages between the clusters (bi-directional)
  - Group based recovery can be used to recover a selected recovery group and gives the option of recovering applications in a desired sequence and time
  - Option to switch primary and recovery package roles in the recovery groups for which the specified cluster is defined as the primary cluster
- 2 data centers required with up to Intercontinental distances
- SCSI or Fibre Channel for local disk connectivity
  - HP Continuous Access XP physical data replication
  - EMC SRDF physical data replication
  - Oracle Standby Database logical replication
- Telco switched networks and protocol conversion to other protocols like ATM OK between clusters ONLY
- Single IP subnet for heartbeat networks (Ethernet, FDDI) within each cluster
- MirrorDisk/UX optional for root disks only when using physical replication
- LVM, VxVM, SLVM or CVM
- Normal cluster lock rules within clusters

Continentalclusters

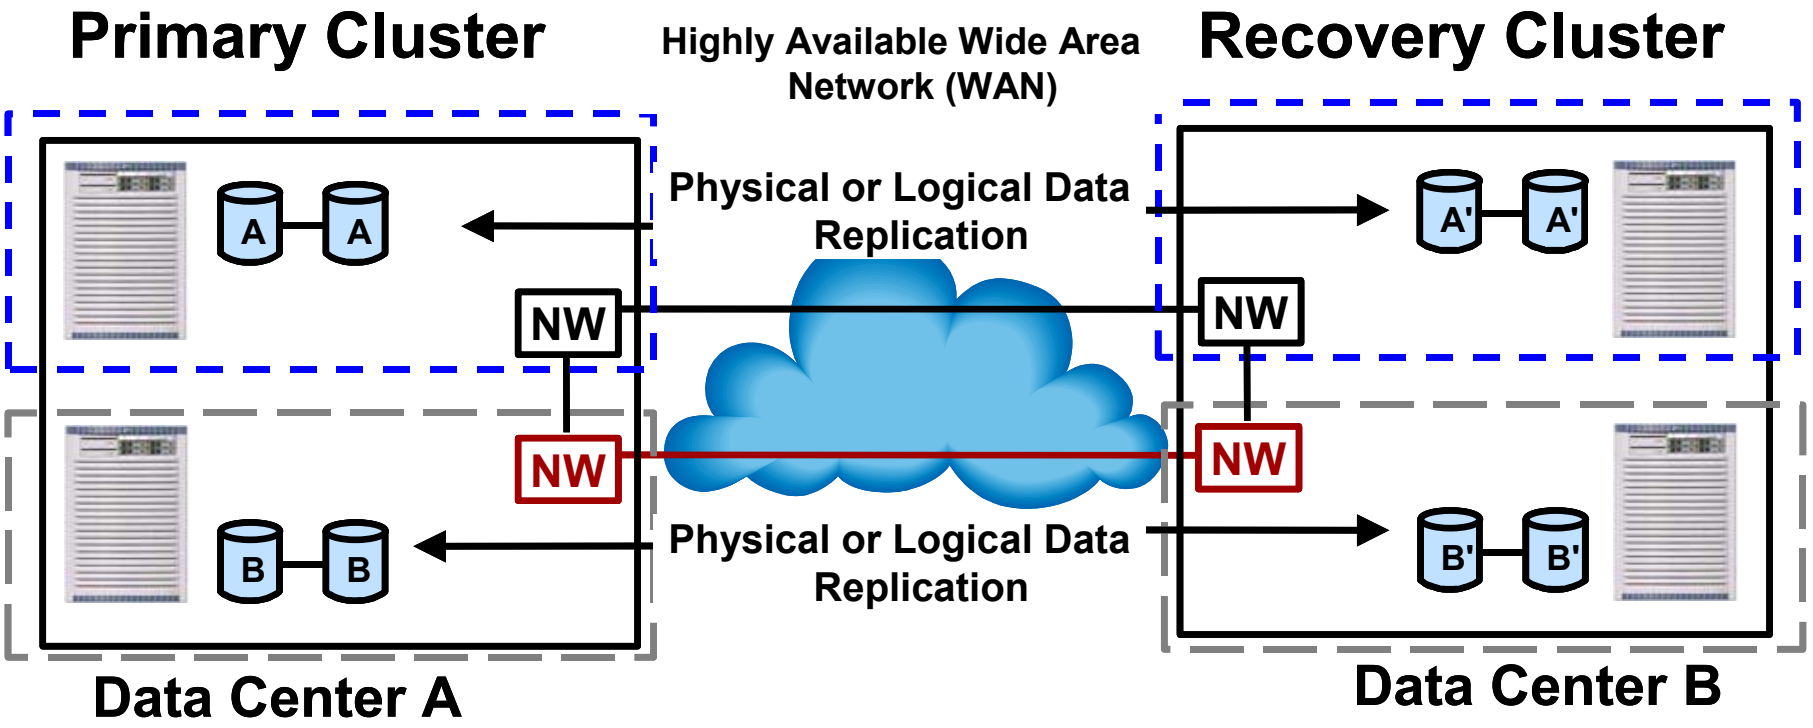

- **Systems are not connected to both replica copies of the data (hosts in each cluster are connected to only one copy of the data)**
- **Each cluster must separately conform to heartbeat network requirements**
- **Each cluster must separately conform to quorum rules (cluster lock disks, Arbitrators or Quorum Server)**
- **Use of cluster lock disks requires three power circuits in each cluster**

* **HA network is used for both data replication and inter-cluster monitoring**

Continentalclusters (continued)

- Each cluster can backup the other cluster (bi-directional)
- Human decision is necessary to initiate the failover
- Cluster problem notification:
  - Text files
  - System console
  - e-mail
  - SNMP trap
  - opcmsg (OpenView/Operations)
- Local failover still occurs within the primary cluster
- Remote failover is used only when the entire primary cluster fails or is unreachable (not if a single node or package fails)
- Choice of data replication methods (physical or logical):
  - HP XP Continuous Access (Async and Sync)
  - EMC SRDF (Sync only)
  - Oracle Standby Database
  - **"Allow model" for other data replication methods to be integrated**

Continentalclusters and Oracle 9i RAC

- Oracle 9i RAC instances running on the primary cluster will be restarted/recovered by CC on the recovery cluster upon primary cluster failure
- This is a bi-directional (mutual) active/standby model where DIFFERENT RAC databases may run in each cluster
- Requirements
  - SLVM for volume management
  - XP CA or EMC SRDF for data replication.
  - VxVM/CVM volume management is not currently supported in this configuration

Oracle 9i RAC instances running in CC environment before failover:

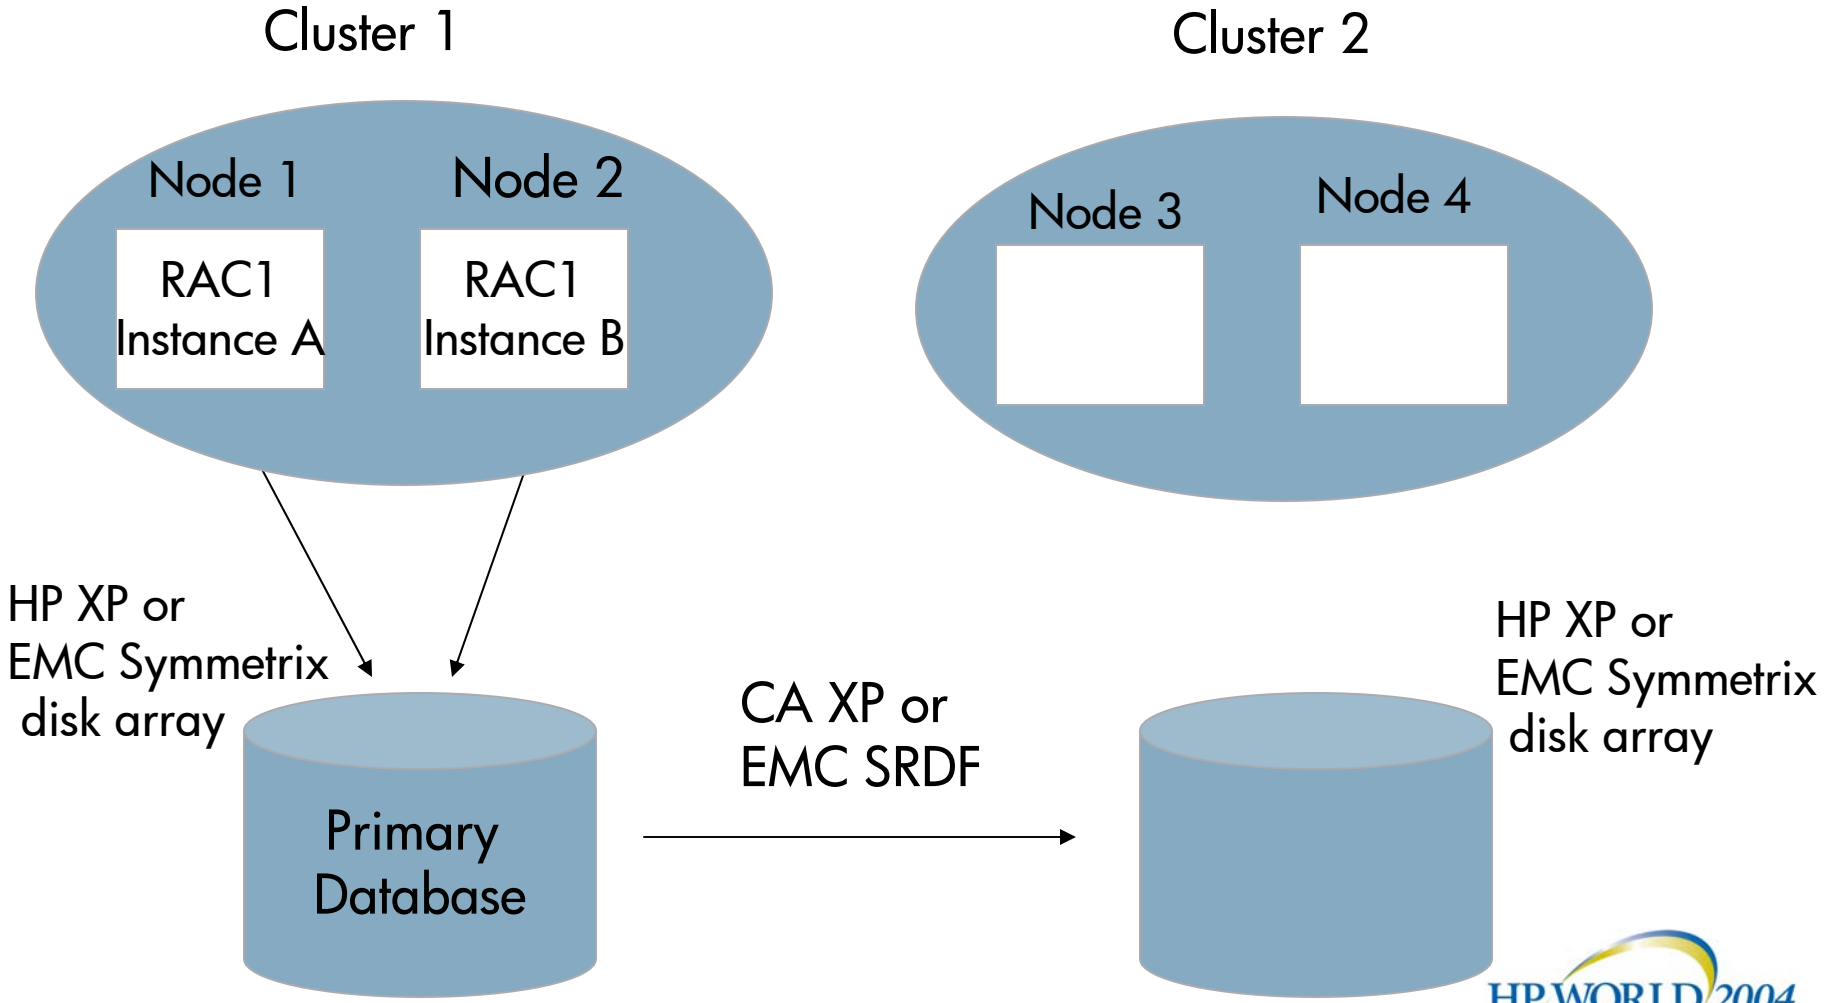

Oracle 9i RAC instances running in CC environment after failover:

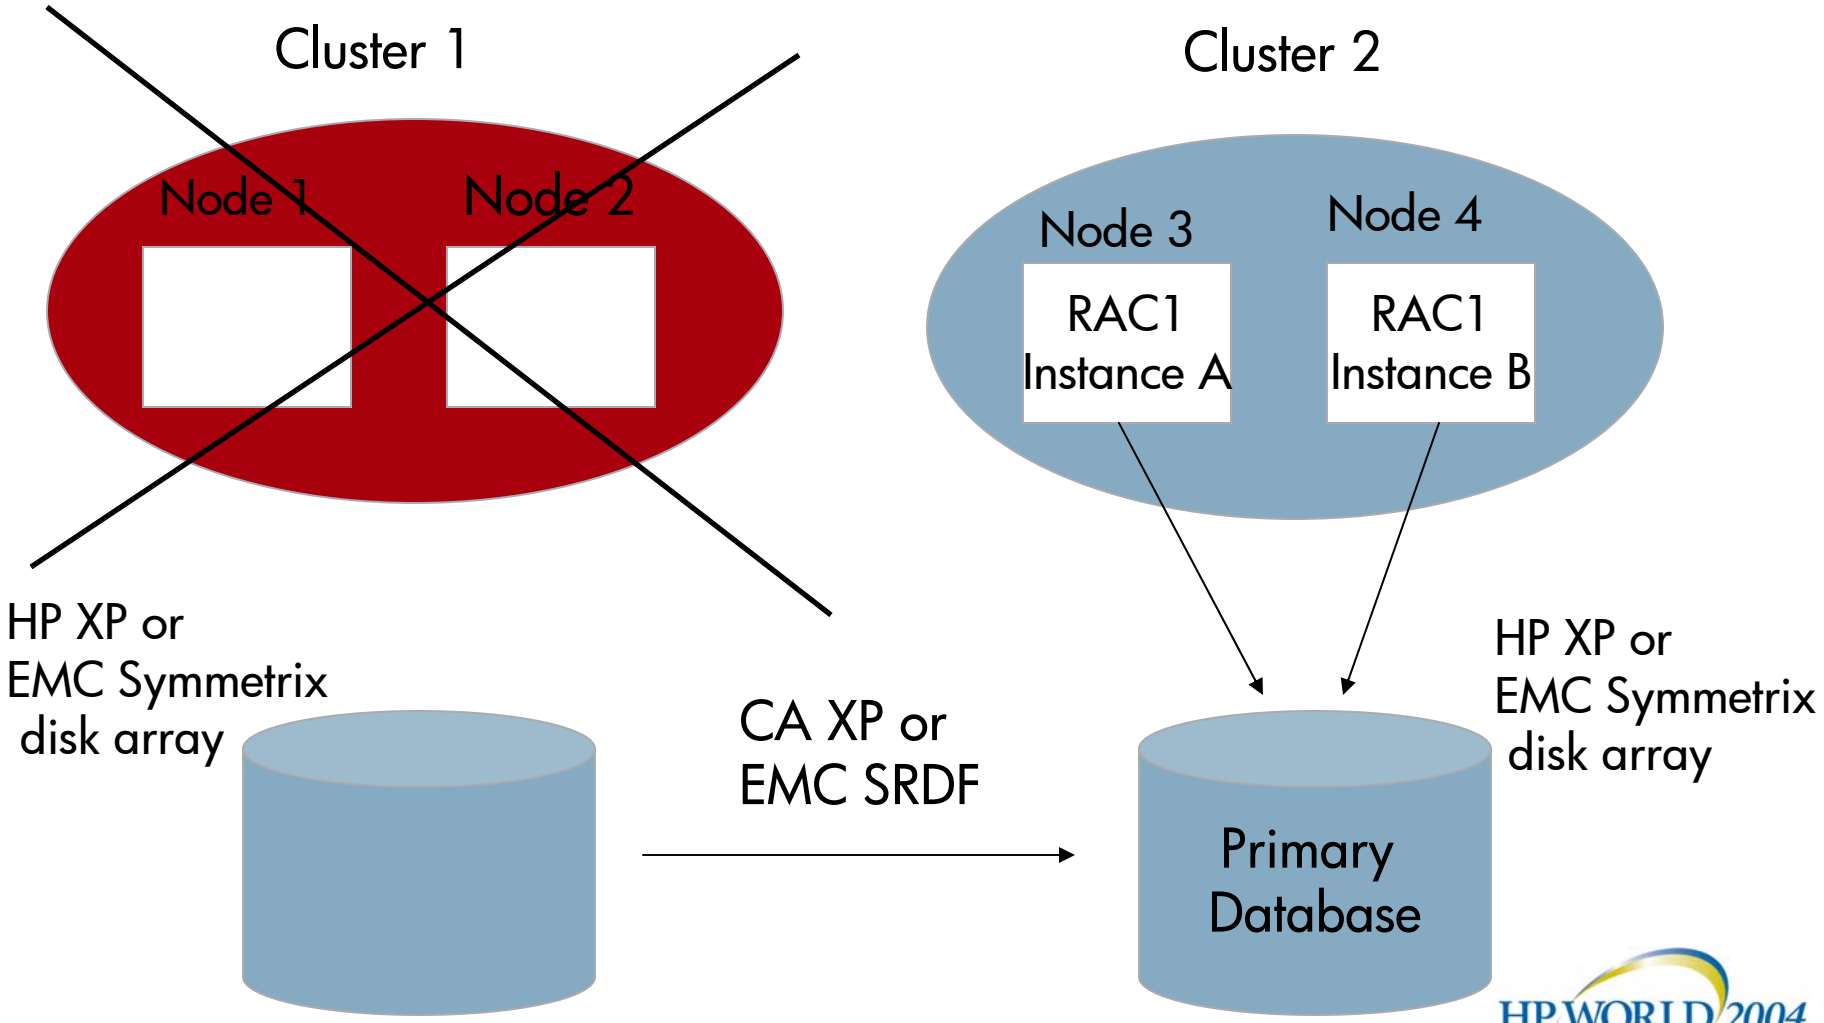

HP Continentalclusters for RAC Disaster tolerance with failover at unlimited distances

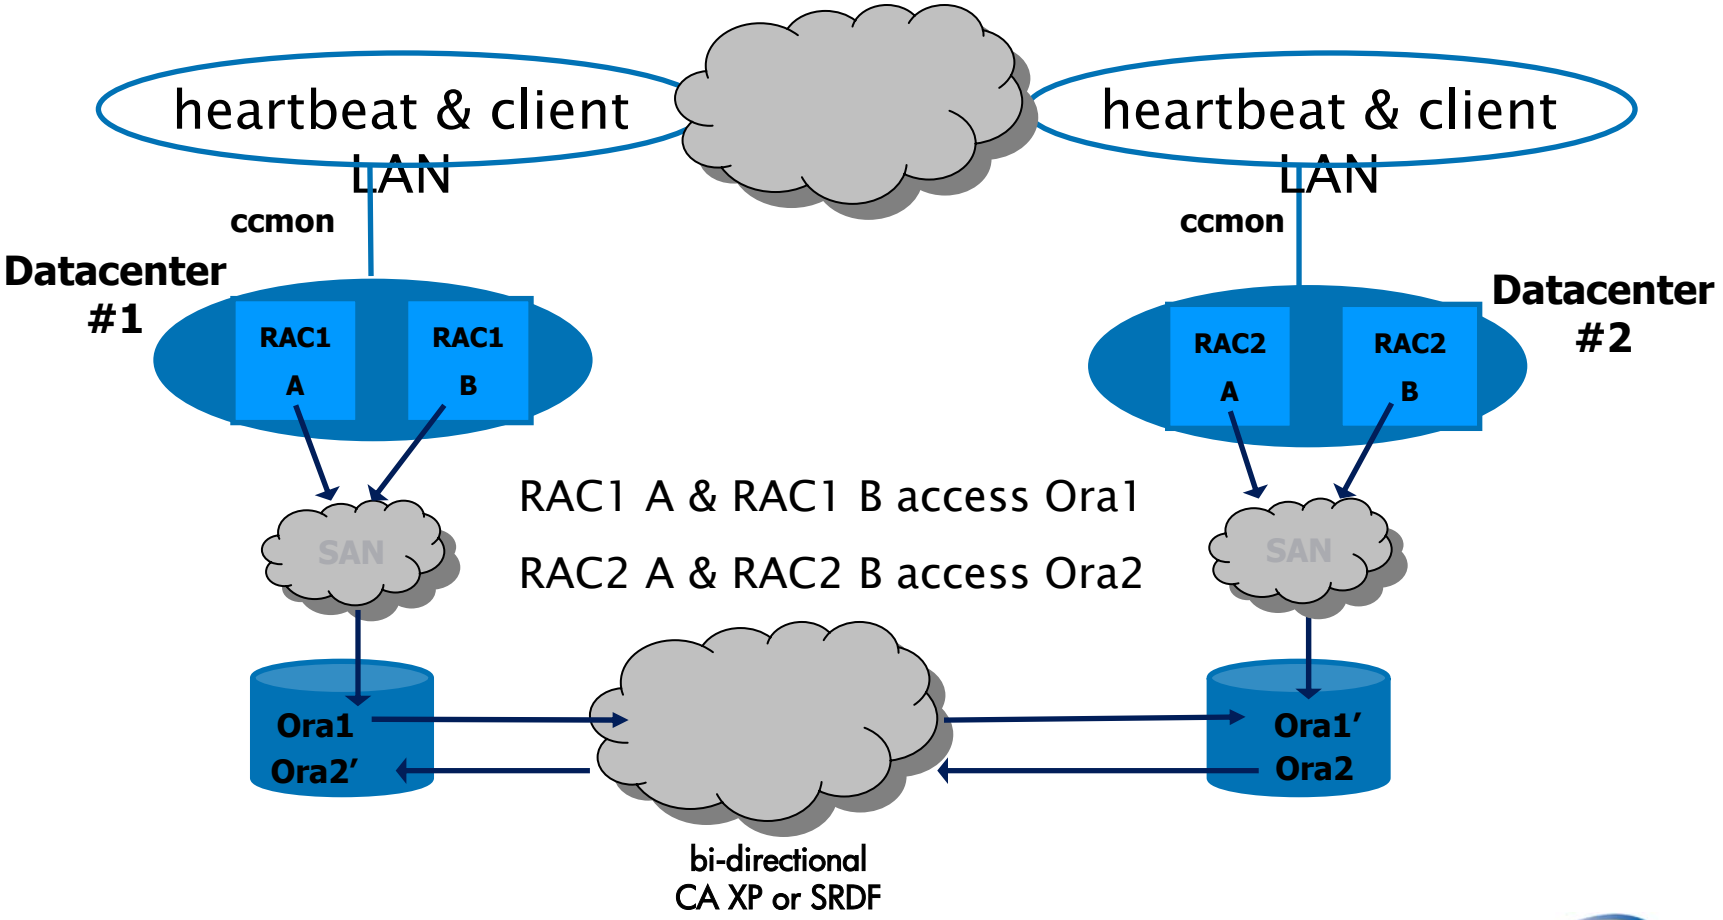

Combining Metrocluster & Continentalclusters

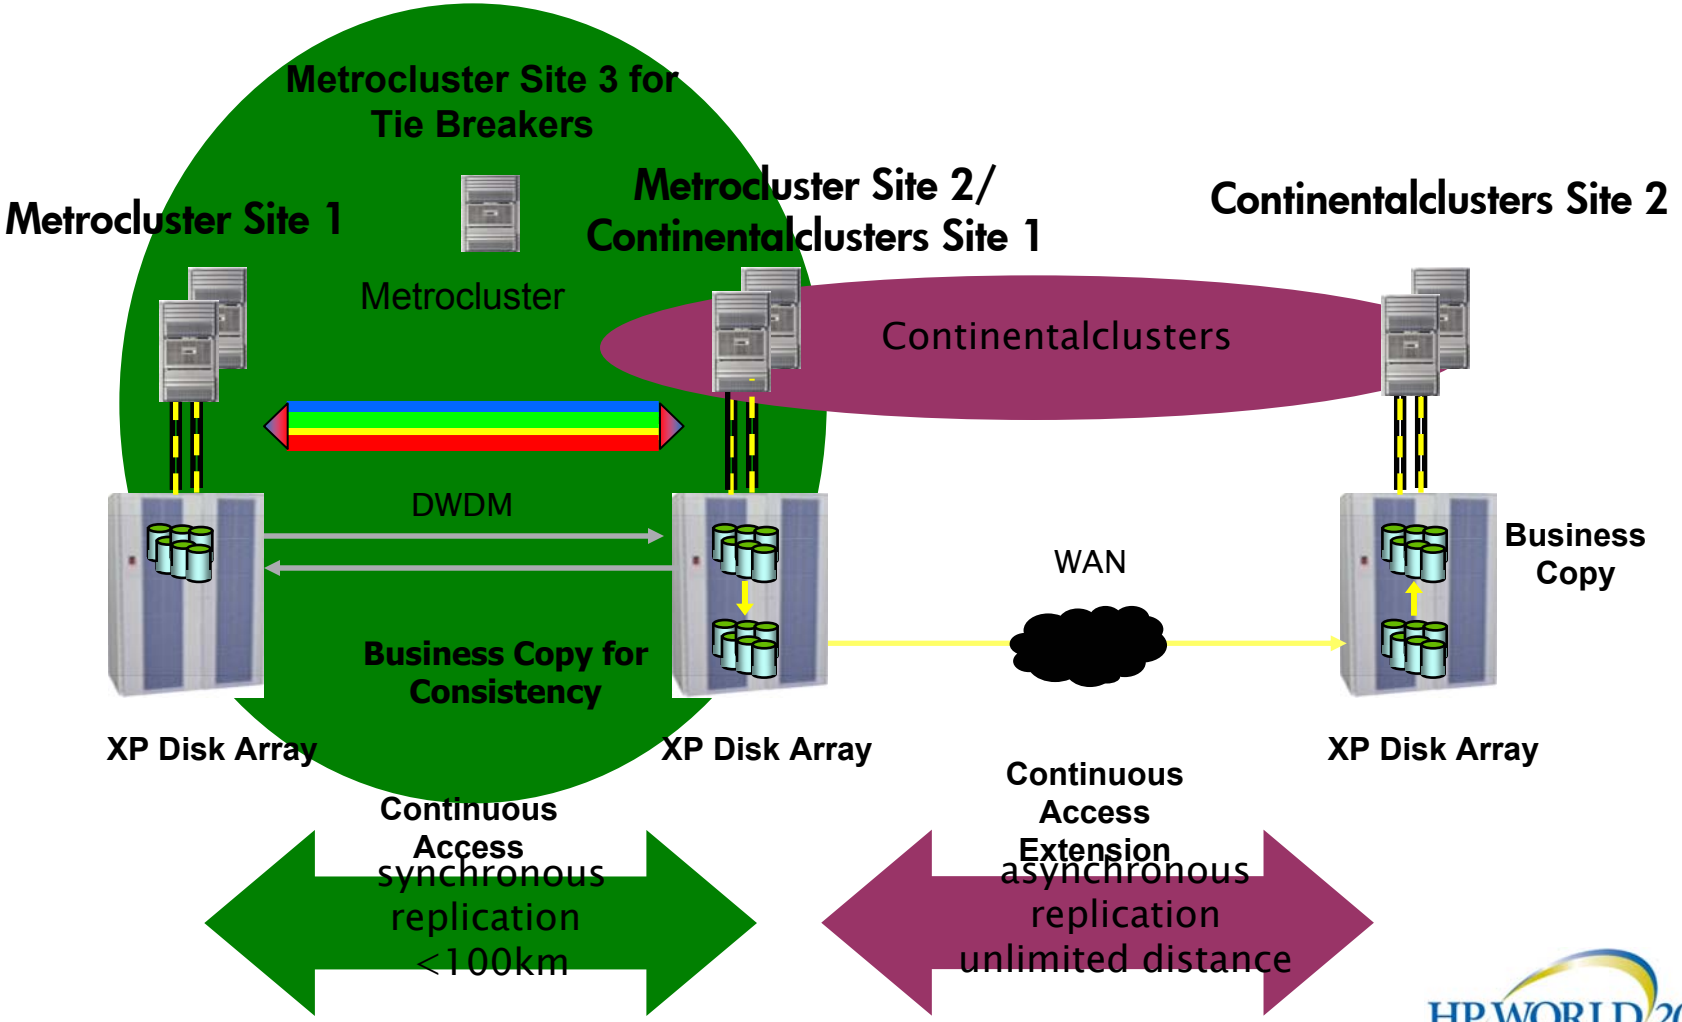

Criteria for Choosing a Campus Cluster Architecture

- Short distances (campus environment)
  - High throughput links for maximum performance
  - Relatively low cost of implementation
    - Software mirroring is much less expensive than hardware array-based replication
    - Network infrastructure is already in place
    - Separate links for FibreChannel and TCP/IP
  - Require an active/active RAC disaster tolerant solution
  - Meets Disaster Tolerance business requirements
    - Not often subject to natural disasters that would impact the entire campus (perhaps have separate electric service or backup generators)
    - Impact of data and service loss does not warrant the cost of implementing a distant data center

Criteria for Choosing an Extended Cluster Architecture

- Metropolitan distances (up to 100 km)
  - Dark fiber already exists or is readily available
  - High throughput links for maximum performance
  - Medium cost of implementation
    - Software mirroring is much less expensive than hardware array-based replication
    - Network infrastructure costs increase due to
      - Dark fiber
      - Redundant DWDM/CWDM converters
    - Share links for TCP/IP and FibreChannel
  - Require an active/active RAC disaster tolerant solution
  - Meets Disaster Tolerance business requirements
    - Fully automated failover without human intervention
    - Greatest availability
    - Not often subject to natural disasters that would impact the entire metropolitan area (e.g., not in earthquake country)
    - Impact of data and service loss does not warrant the cost of implementing a more distant data center

Criteria for Choosing a Metrocluster Architecture

- Metropolitan distances (up to 100 km)
  - Dark fiber already exists or is readily available
  - High throughput links for maximum performance
  - Medium cost of implementation
    - Must invest in enterprise disk arrays that support array-based replication
    - Network infrastructure costs increase due to
      - Dark fiber
      - Redundant DWDM/CWDM converters
    - Share links for TCP/IP and FibreChannel
  - Meets Disaster Tolerance business requirements
    - Fully automated failover without human intervention
    - Greatest availability
    - Not often subject to natural disasters that would impact the entire metropolitan area (e.g., not in earthquake country)
    - Impact of data and service loss does not warrant the cost of implementing a more distant data center

Criteria for Choosing a Continentalclusters Architecture

- Continental and Intercontinental distances
  - Dark fiber is not readily available or the distance is too great
  - Link cost is the largest component
    - Network infrastructure costs increase due to
      - Bandwidth needed to support both TCP/IP and data replication
      - Long distance
  - Want control over the initiation of the failover
  - Failover only when entire data center fails meets DR model and availability requirements
  - Active/standby RAC configuration is acceptable
  - Meets Disaster Tolerance business requirements
    - Not often subject to natural disasters that would impact the entire metropolitan area (e.g., not in earthquake country)
    - Impact of data and service loss does not warrant the cost of implementing a more distant data center

Serviceguard and Oracle 10g

Serviceguard and Oracle 10g RAC

- Oracle Real Application Clusters (RAC) 10g is an option to Oracle 10g Enterprise Edition
- Differences from the previous Oracle9i RAC product include:
  - Integrated clusterware (Cluster Ready Services - CRS)
  - Application Storage Management (ASM) software can be used to manage storage for the RAC database
  - Performance improvements
  - Zero downtime patches for Oracle RAC environments
- NOTE: the use of CRS is OPTIONAL; you can continue to use SGeRAC

Oracle 10g Cluster Ready Services (CRS)

- Introduced with 10g, CRS provides:
  - A standard cluster interface on all platforms for high availability and workload management
  - Management of cluster DB functions, including:
    - Node membership (connectivity)
    - Group services (messaging and locking)
    - Global resource management for Oracle processes
    - High availability (cluster control and recovery)
- Services are created for each application or for major components within complex applications as a method for controlling availability and workloads

Example RAC 10g cluster with two services

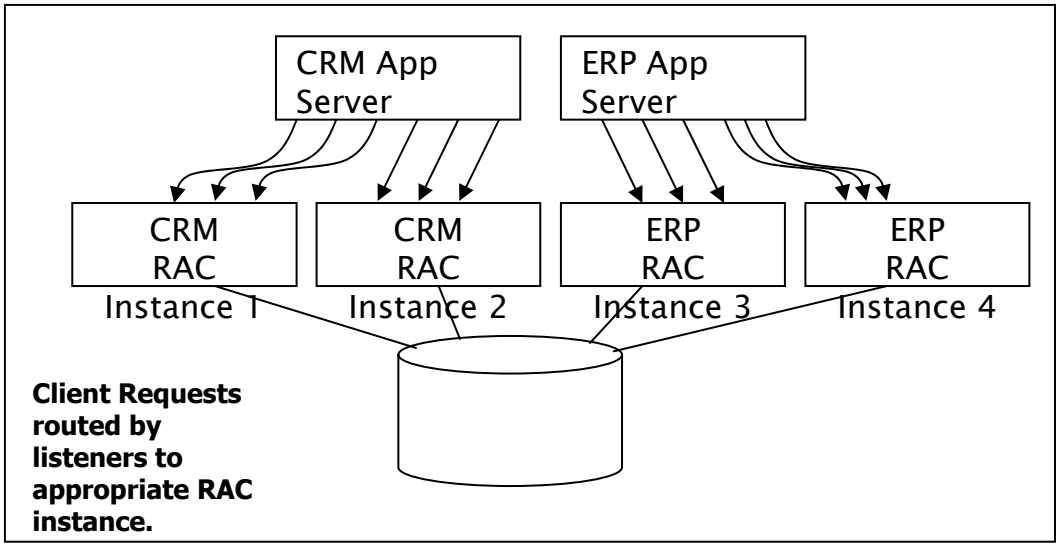

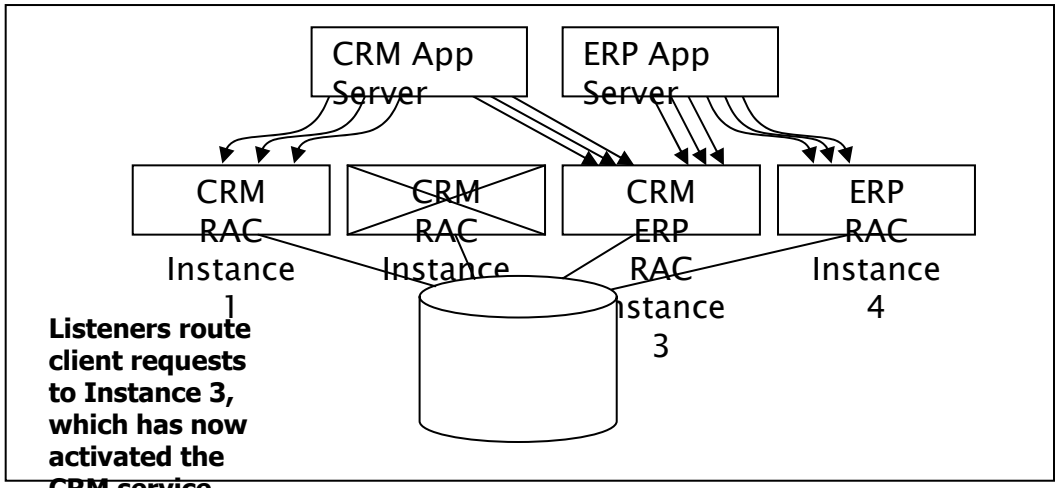

Oracle 10g Application Storage Management (ASM)

- Software that can be optionally used to
  - Manage storage for the RAC database
  - Collects LUNs together
  - No file system or volume manager
  - No cluster file system (CFS)

- Use of an HP volume manager and file system
  - Will still be required for non-RAC-database files
  - Can still be used for RAC database files
  - Will be needed if the HP-UX 11i v3 Cluster File System will be used
  - Allows a single tape backup solution

SGeRAC with RAC 10g

- Increases availability for non-RAC processes and applications running on nodes within the RAC cluster
  - since CRS focuses only on Oracle-specific processes and resources
- Is required when using a volume manager
  - Greater flexibility when working with LUNs
  - Features such as mirroring, striping, etc.
  - SLVM or CVM provides concurrent access to the same storage by multiple nodes
- Provides reliable node membership information through
  - Tight kernel integration
  - Real Time priority execution
- Improves network reliability through monitoring and failover management of
  - User LAN
  - Cluster interconnect
  - IPv6 networks

WLM with Oracle RAC 10g

- HP-UX WLM can be used to supplement the workload management capabilities of CRS:
  - CRS manages workloads of only Oracle RAC-related processes
  - WLM can manage resource entitlements of non-Oracle processes and applications
    - Based on defined SLAs and relative priorities within a server's partitioning schemes (e.g., vPARs, nPARs, and psets)
    - Automatically manages Pay Per Use and Temporary iCOD processors for temporary capacity increases
    - For example, WLM can be used to manage resource entitlements between development, test and production environments

Serviceguard and Oracle 10g

- There are several Technical Whitepapers being written that are expected to be available later in 2004
- These whitepapers are intended to help HP field people and customers to choose the best HA solutions for their Oracle 10g environment

High Availability Solutions in HP-UX 11i v2UD2

High Availability Solutions in HP-UX 11i Version 2 Update 2

- Planned to ship November 2004
- Support for PA-RISC and Integrity (IPF-based) servers
  - Single OS with the same features
  - Serviceguard 11.16
  - Metrocluster and Continentalclusters
  - Both families of servers in the same cluster

New!

  - Technical Whitepaper will be available to fully describe the supported architectures and restrictions

Mixed PA and IPF clusters are not supported today

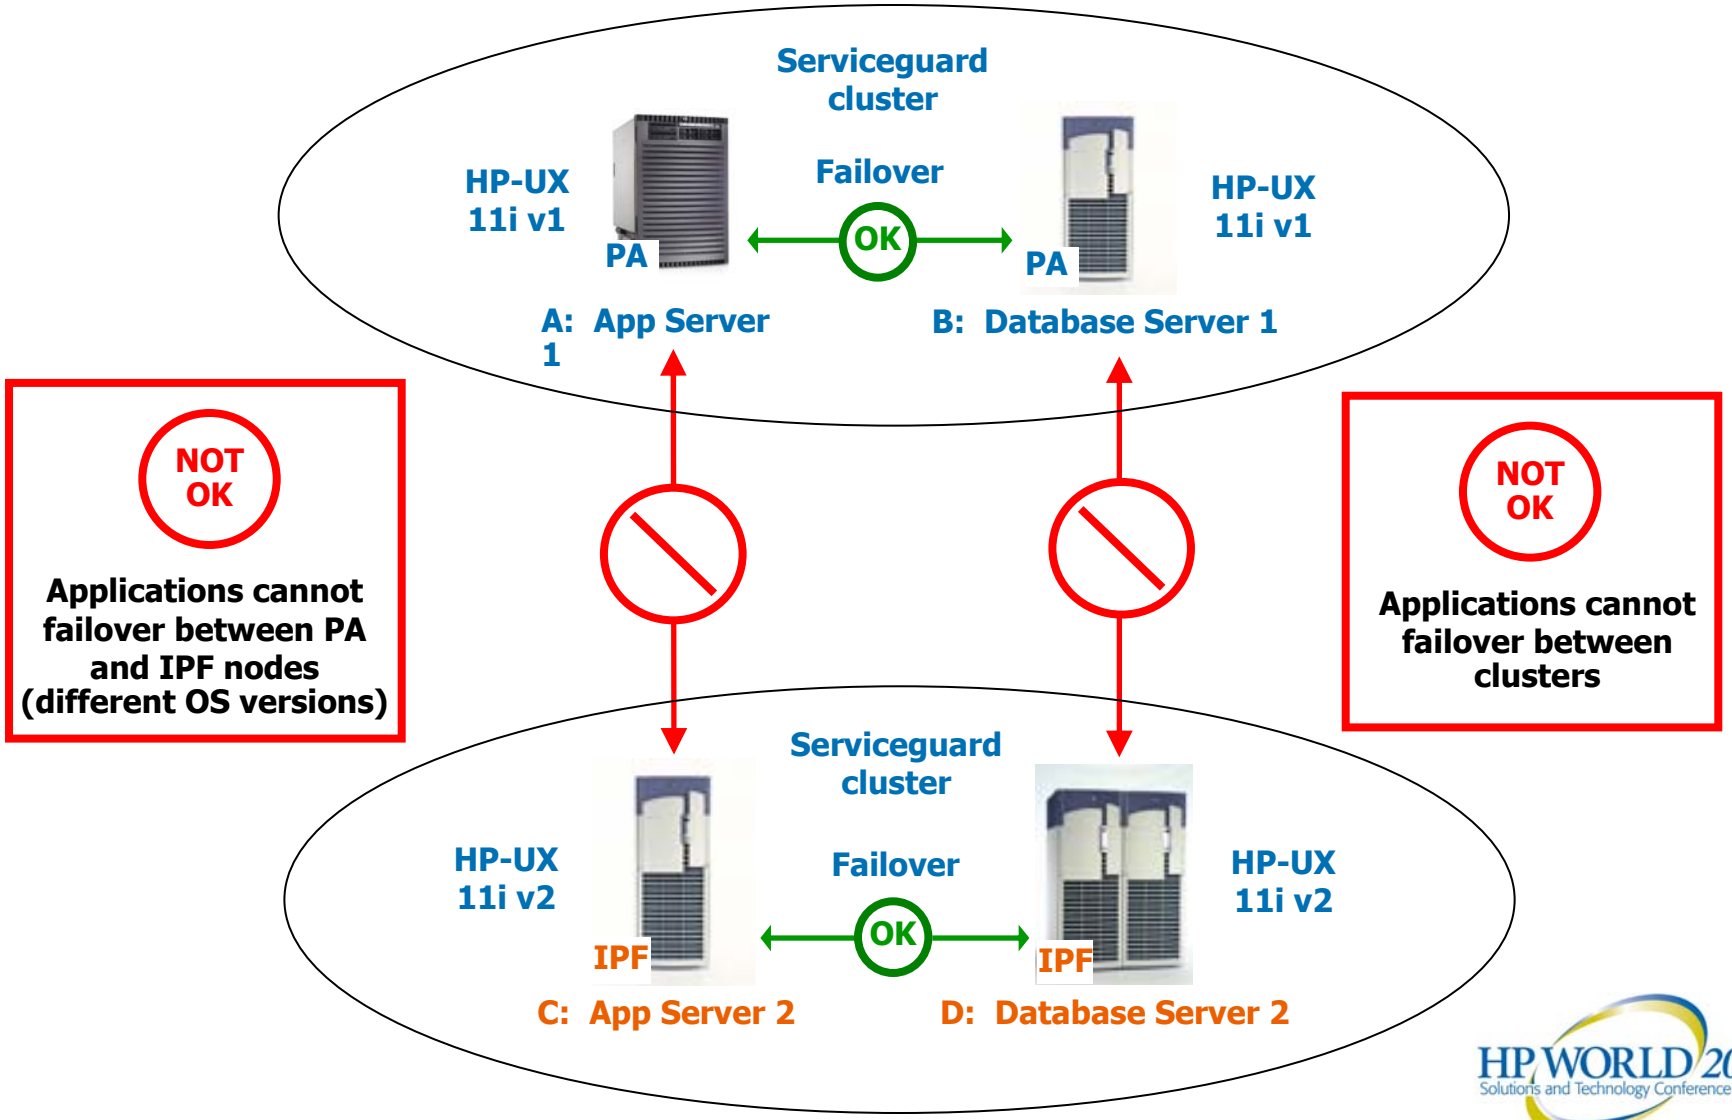

Mixed PA and IPF system clusters tomorrow (Q4'2004)

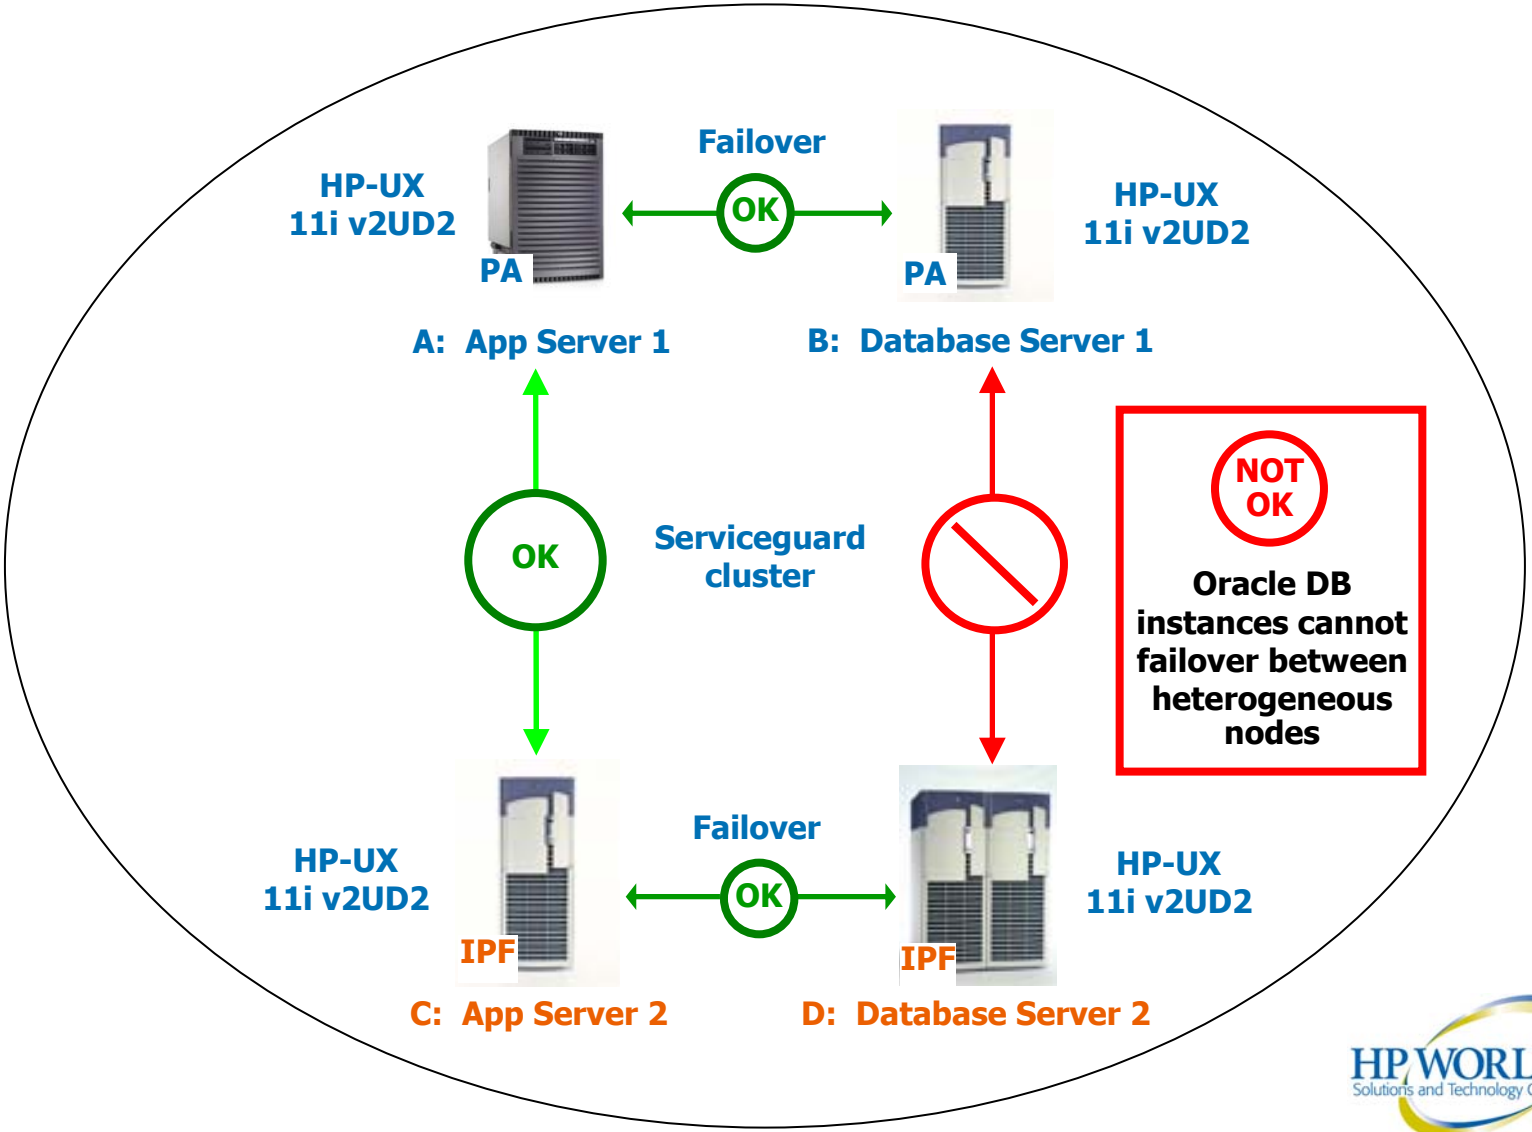

Mixed PA and IPF system clusters tomorrow (4CQ 2004)

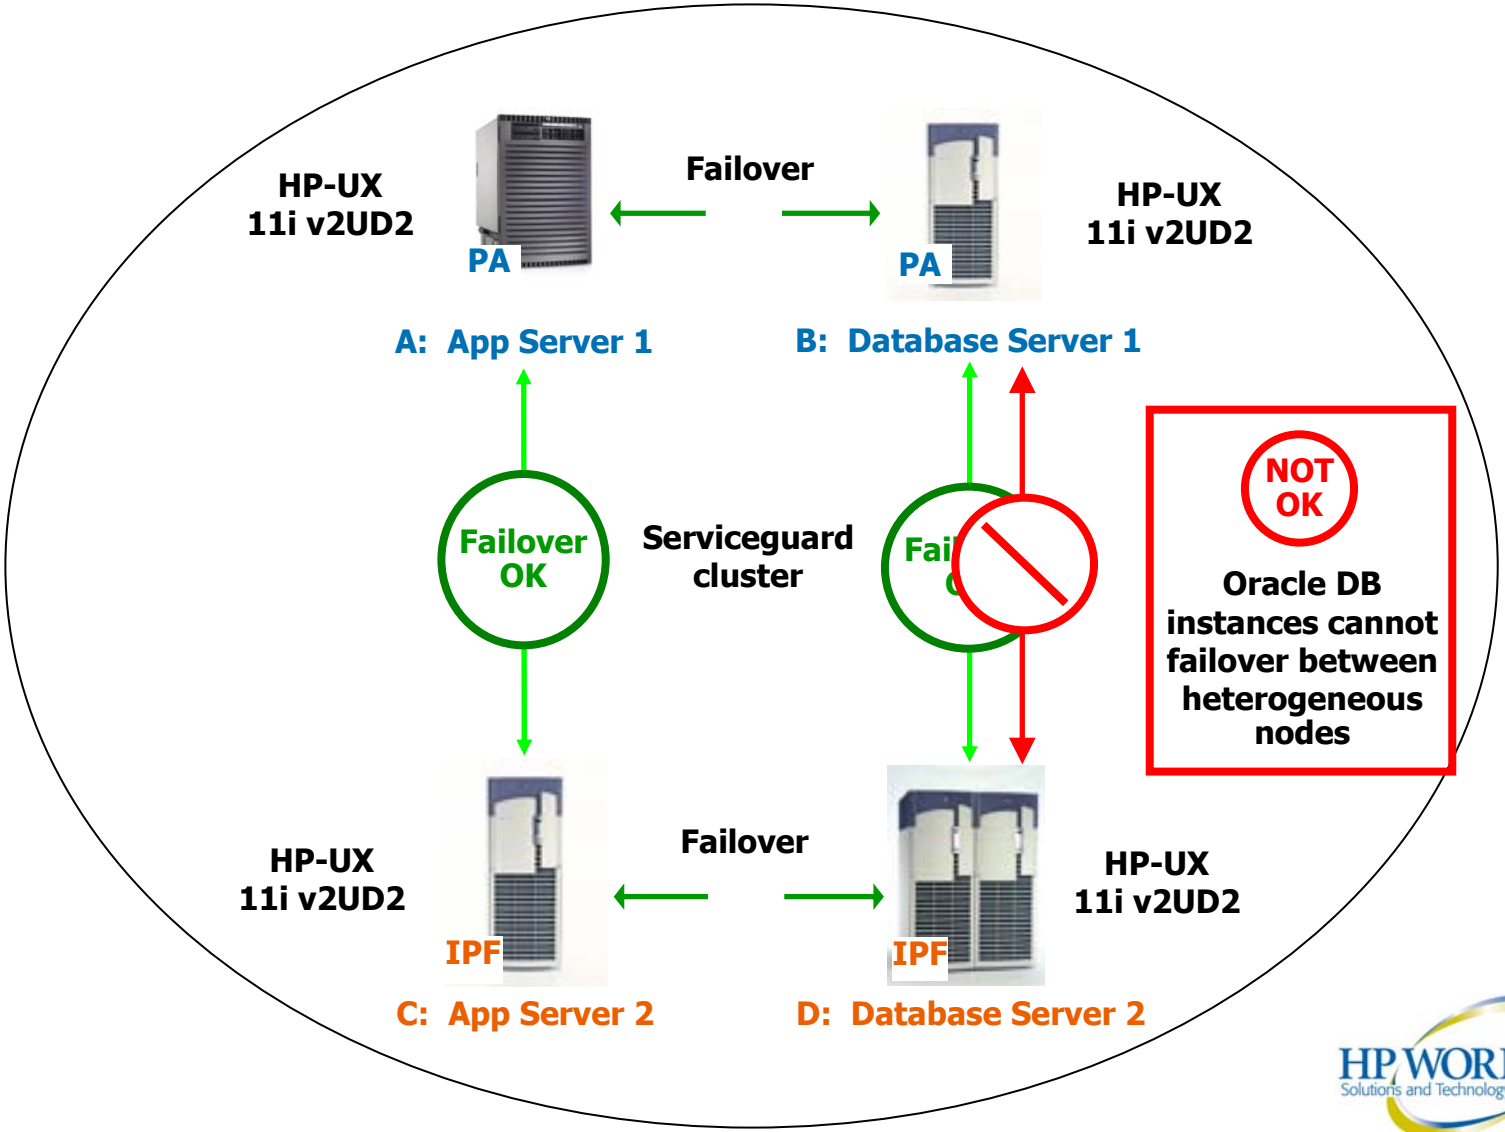

Mixed cluster advantages

- Simplified administration for both PA and IPF-based systems
  - One common OS (HP-UX 11i v2UD2)
  - one common cluster software (Serviceguard 11.16)
- Easy addition of IPF-based systems to existing PA-based SG clusters
- Ability to gain experience with IPF-based systems in known environments
- Transition from PA- to IPF-based systems can be more easily accomplished

Native Executables and Aries

- Normally, executables are compiled for a particular architecture
  - PA executables for use on a PA-based system
  - IPF executables for use on an IPF-based system
  - Magic numbers in the a.out file are used to code the executable for a particular architecture

- Aries is a binary emulator included with HP-UX 11i v 2
  - transparently emulates 32-bit and 64-bit HP-UX PA-RISC executables on HP-UX Itanium®-based machines
  - HP-UX PA-RISC applications do not have to be recompiled
  - The user does not explicitly invoke Aries
    - The HP-UX kernel running on an IPF system recognizes an HP-UX PA-RISC executable and automatically invokes Aries to emulate the application without any user intervention
    - Note: Application performance under Aries emulation may vary

Aries application emulator

- Consists of 4 shared libraries
  - /usr/lib/hpux32/aries[32/64].so
    - 32/64-bit Aries application emulator
  - /usr/lib/hpux32/pa_boot[32/64].so
    - Loads and executes the appropriate emulator when the HP-UX Itanium-based kernel recognizes a HP-UX PA-RISC executable
- Supports all HP-UX PA-RISC applications
  - Debugging of PA applications on IPF systems using PA gdb
  - Inter-process communication (IPC)
    - Semaphores, shared memory, sockets, etc.
    - Signal/exception behavior
    - IPC between emulated PA applications and native IPF applications
- Includes performance-enhanced Dynamic Translator (beta)
  - Beta translator not invoked by default (enabled via resource file)
  - Measured translation time improvement by 20%, on average
  - Some risk associated with using beta translator on mission-critical apps

Notes on Aries emulation

- Aries PA-RISC emulation performance will vary depending on the application
- Aries included with HP-UX 11i v2 supports PA applications running on HP-UX 11i v1.6 and below
- No support for PA programs that load IPF shared libraries
- Does not support privileged PA-RISC instructions (device drivers, loadable kernel modules)
- Does not guarantee correct emulation of PA applications that make assumptions about the time required to execute certain sections of application code or system calls
- Emulated processor-related system calls return information pertinent to a PA-RISC 2.0 processor
- See Aries(5) man pages and release notes for detailed information on Aries emulation

Requirements for mixed PA/IPF cluster implementations

- All nodes must have HP-UX 11i v2UD2 installed
- All nodes must have Serviceguard 11.16 installed
- The applications are configured to follow the ISV's configuration requirements:
  - Failover between homogeneous nodes in a mixed cluster only, OR
  - Failover between heterogeneous nodes in a mixed cluster using either of the following (not all ISVs will support this):
    - One set of PA binaries running native on PA and the same set using the Aries emulator on IPF, OR
    - Two sets of native binaries (one for PA, one for IPF)
      - either all 32-bit or all 64-bit binaries

Supported servers in a mixed PA/IPF 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
- IPF-based servers:
  - Entry Level: rx16xx, rx26xx, rx46xx, [rx56xx]
  - Mid Range: rx76xx, rx86xx
  - High End: Superdome

Supported storage in a mixed PA/IPF Serviceguard cluster

Disk Arrays:	Interface
XP48	1Gb FC
XP128	2Gb FC
XP128	1Gb iSCSI
XP256 – Fibre Channel ports	1Gb FC
XP512	1Gb FC
XP1024	2Gb FC
XP1024	1Gb iSCSI
XP enhancement	2Gb FC
VA7100	1Gb FC
VA7400	2Gb FC
VA7110	2Gb FC
VA7410	2Gb FC
FC60	1Gb FC
MSA1000	2Gb FC
MSA1500	2Gb FC
SabreJet	1Gb iSCSI
MA8000 (HSG80)	1Gb FC
EMA 12000 & 16000 (HSG80)	1Gb FC
EVA1500 XL	2Gb FC
EVA5000 (HSV110)	2Gb FC
EVA3000 (HSV100)	2Gb FC
JBODs:	
SC10	LVD SCSI U2
FC10	1Gb FC
DS2100	LVD U160
DS2120	LVD U320
DS2300	LVD SCSI U160
DS2405	2Gb **FC-AL only**
Storage works 4314T & 4314R	LVD U160
Storage works 4354R	LVD U160
Storageworks Northstar	LVD U320
Storageworks Southern Cross	LVD U320

SAN Switches and Hubs	
B Series (Brocade)	
Storageworks SAN switch 8-EL	1Gb FC
Storageworks SAN switch 16-EL	1Gb FC
Storageworks SAN switch /8	1Gb FC
Storageworks SAN switch /16	1Gb FC
Storageworks SAN switch Intregrated 32	1Gb FC
Storageworks SAN switch Intregrated 64	1Gb FC
Storageworks SAN switch 2/8-EL	2Gb FC
Storageworks SAN switch 2/8	2Gb FC
Storageworks SAN switch 2/16-EL	2Gb FC
Storageworks SAN switch 2/16	2Gb FC
Storageworks SAN switch 2/32	2Gb FC
Storageworks Core Switch 2/64	2Gb FC
M Series (McDATA)	
McDATA ED5000	1Gb FC
EMC Connectrix DS-16B	1Gb FC
EMC Connectrix ED-1032	1Gb FC
Storageworks SAN Director 64	1Gb FC
Storageworks Edge switch 2/32	2Gb FC
Storageworks Director 2/64	2Gb FC
C Series (CISCO)	
MDS 9216	2Gb FC
MDS 9509	2Gb FC
SR2122-2 iSCSI storage router	iSCSI
Other Switch/Hub	
Surestore L10/S10 hub	1Gb FC

Supported software in a mixed PA/IPF Serviceguard cluster

- HP-UX UX 11i v2UD2
  - Mission Critical Operating Environment (MCOE) installed on all nodes
  - A non-MCOE 11i v2UD2-based system with Serviceguard 11.16
- Same or similar patch set for OS, Serviceguard and ISV's applications
- Same version of volume manager software (if used for shared storage)
  - LVM
  - VxVM
  - CVM
- Same version of file system software (if used for shared storage) Note: The default file system layout versions of VxFS will differ based on the OS release:
  - VxFS created on 11i v1 use file system version 4 and is compatible with older layouts
  - VxFS 3.5 created on all 11i v2 use file system version 5 and is compatible with version 4
  - A migration tool is available to convert version 4 to version 5
- nPARs
  - No special mixed cluster requirements on Serviceguard configurations with nPARs
- vPARs
  - Currently, Serviceguard is supported with vPARs on PA systems only for 11i v2UD2
  - A new version of vPARs should be available in 1H05 on 11i v2 that will add vPAR support for IPF systems
  - Configurations with Serviceguard and virtual partitions within the same server are not recommended
- Campus/Extended Cluster, Metrocluster and Continentalclusters (subject to constraints above)

Note: SGeRAC in mixed clusters are not currently supported (under investigation for a future release)

Application code support in mixed clusters

- HP-UX 11i v2UD2 supports applications that have been compiled in 32-bit mode as well as 64-bit mode on both PA and IPF systems

	PA systems with HP-UX 11.i v2UD2	IPF systems with HP-UX 11.i v2UD2
PA 32-bit application	Supported	Supported with Aries
PA 64-bit application	Supported	Supported with Aries
IPF 32-bit application	N/A	Supported
IPF 64-bit application	N/A	Supported

Mixed cluster configuration example:
Adding 1 IPF server to a 2-node PA cluster

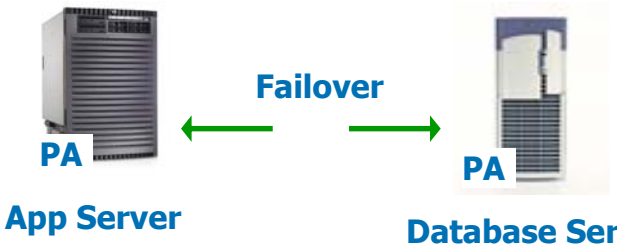

Original 2-node PA cluster

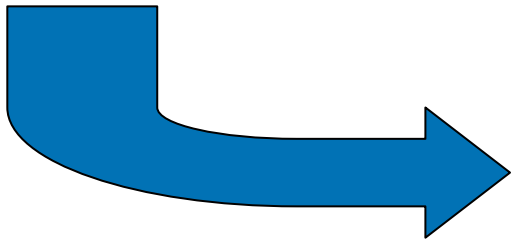

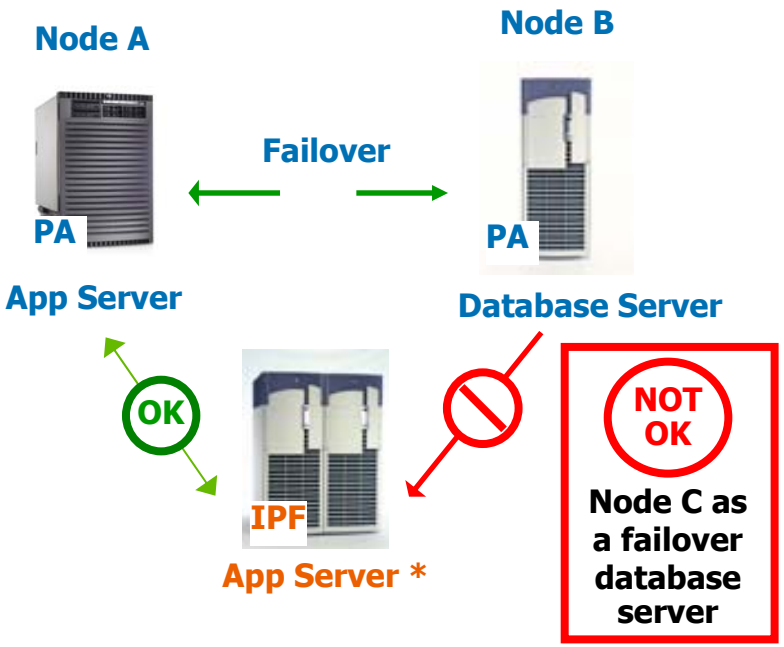

New 3-node mixed PA/IPF cluster

* not all app servers need failover functionality

Mixed cluster configuration example:
Adding 2 IPF servers to a 2-node PA cluster

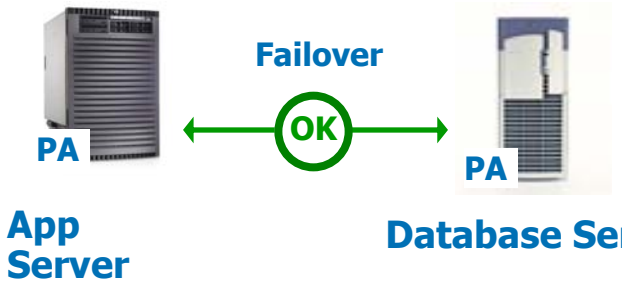

Original 2-node PA cluster

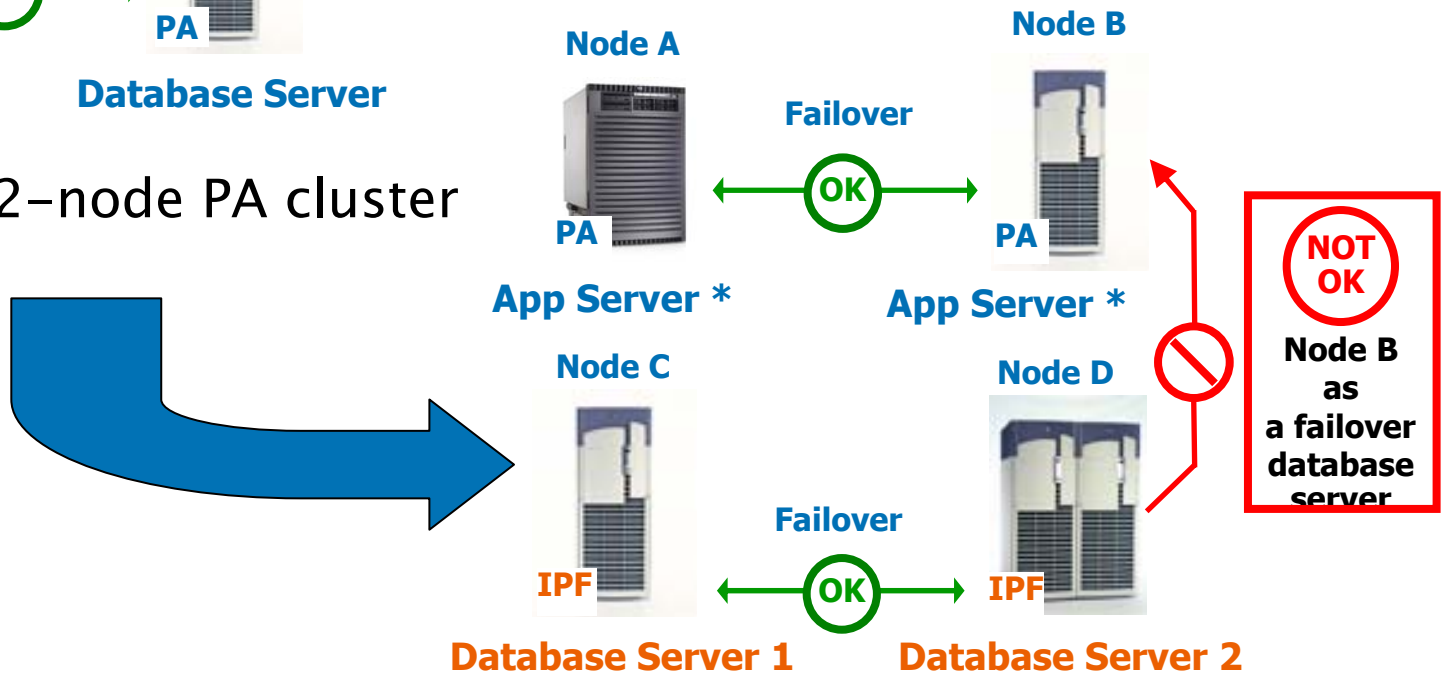

New 4-node mixed PA/IPF cluster

* not all app servers need failover functionality

Mixed cluster configuration example:

Adding one critical application and two IPF nodes to existing 4 node PA cluster

Existing 4-node PA cluster with 4 critical applications

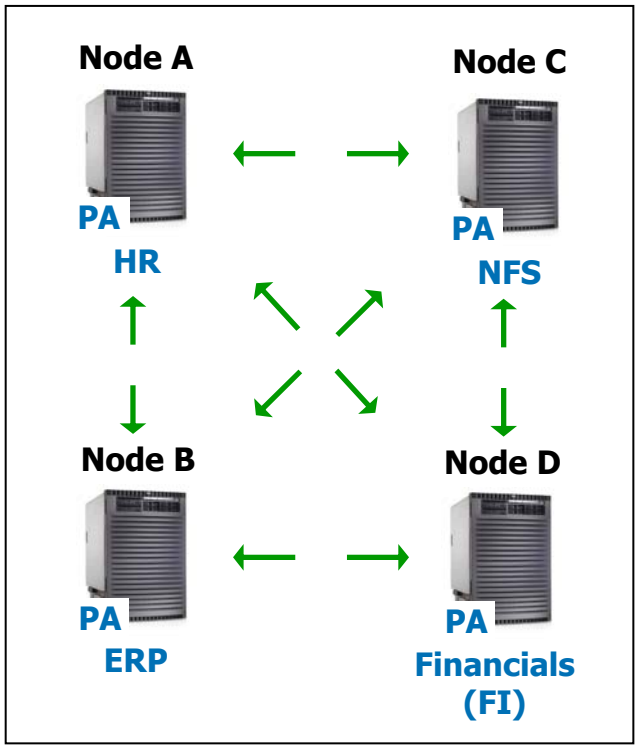

New mixed cluster with:

- **Additional new critical app (CRM)**
- **2 additional IPF nodes**
- **Maintained performance in failover situation**

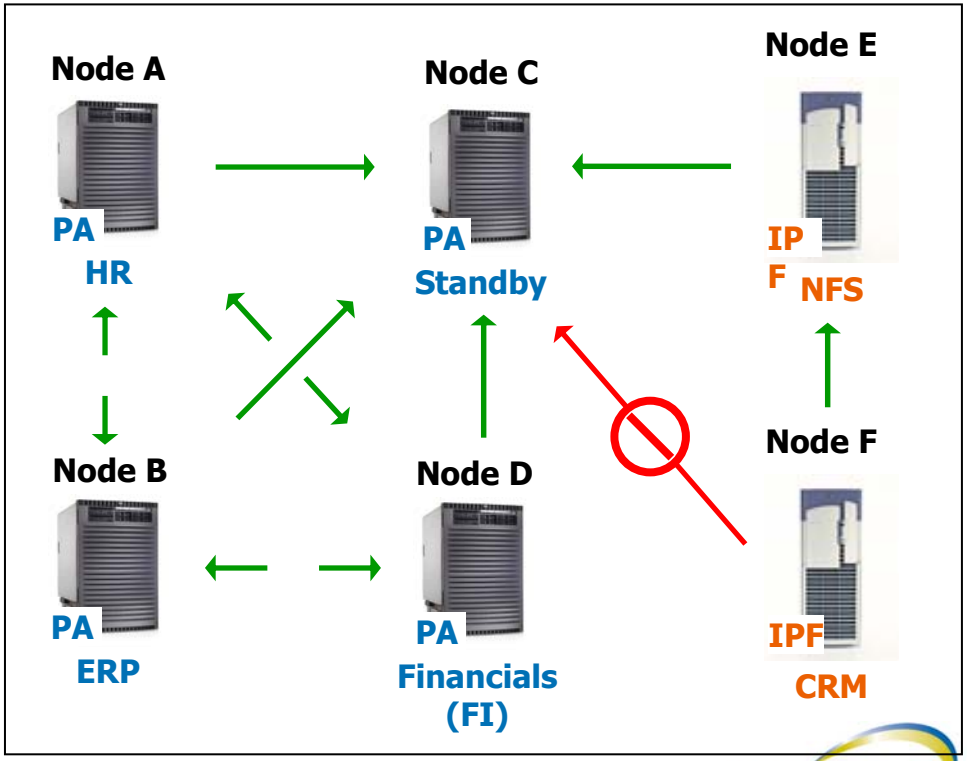

HP World 2004 - Seminar # 3517 - Copyright 2004 Hewlett-Packard Company

Mixed cluster configuration example:
Add 2 IPF nodes to PA telco cluster for more capacity

Existing 3-node PA telco billing cluster

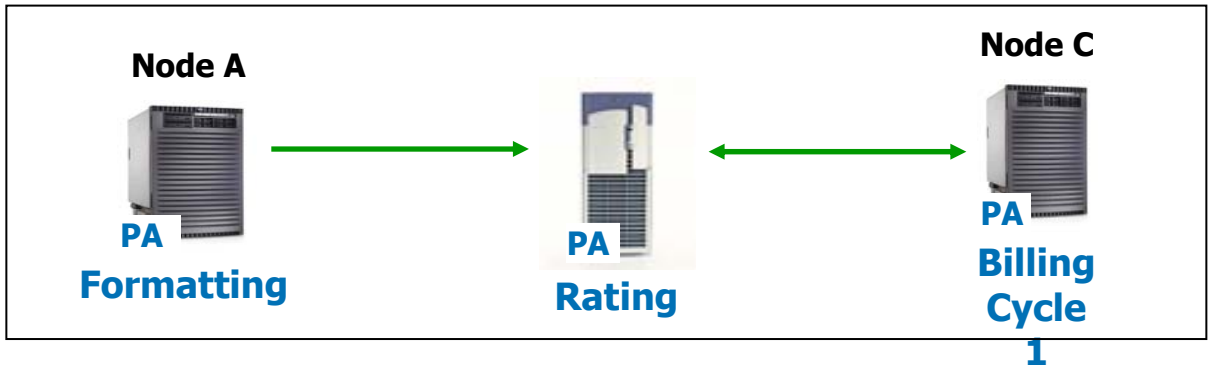

New 5-node PA/IPF cluster with 2 additional billing cycle systems

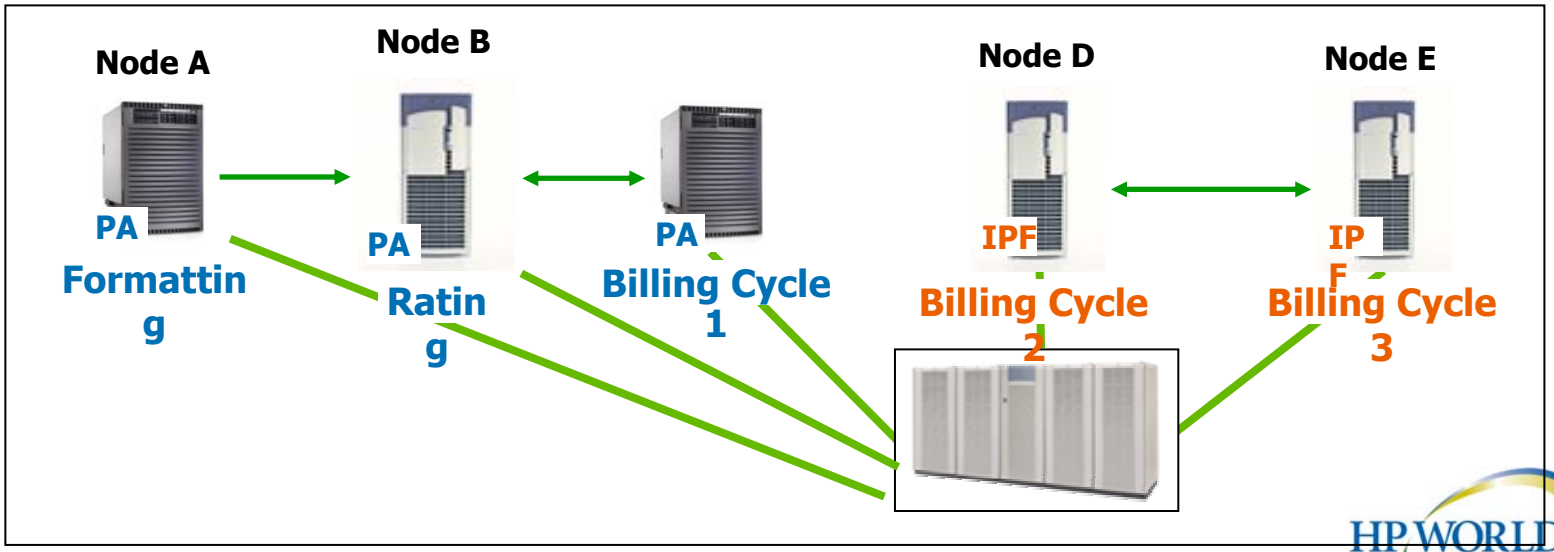

HP World 2004 - Seminar # 3517 - Copyright 2004 Hewlett-Packard Company

Transition from PA to IPF via mixed clusters

Add 1 IPF node to 4 node SAP PA cluster

Existing 4-node all-PA SAP cluster with ERP and Business Warehouse (BW) systems

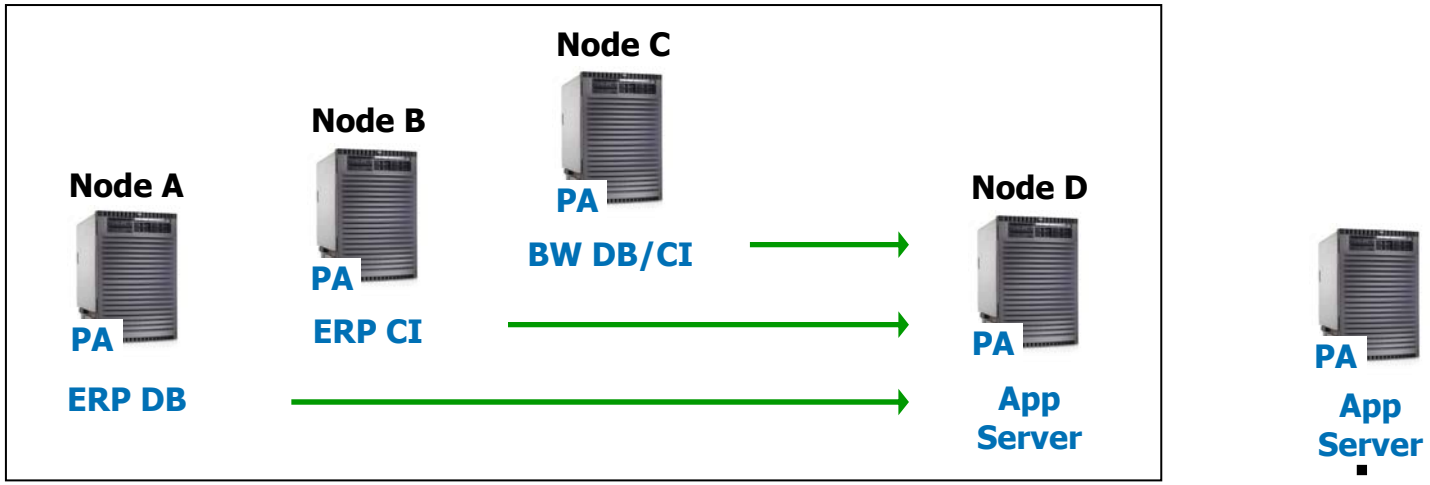

New 5-node PA/IPF SAP cluster with ERP and BW systems

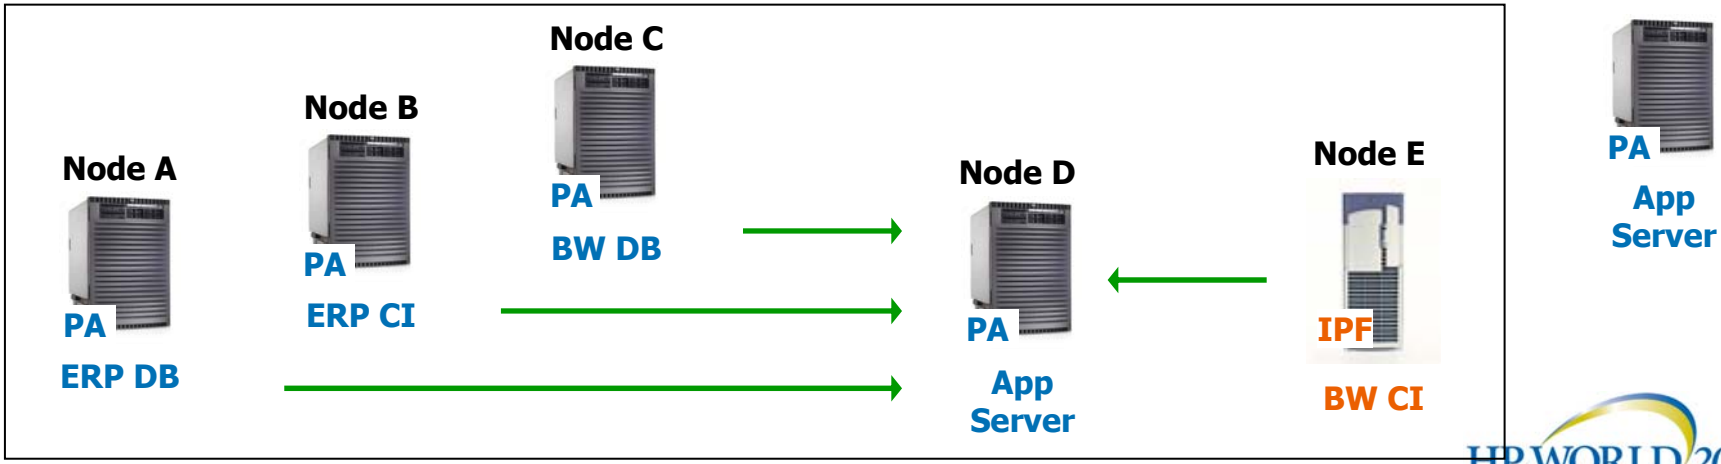

Transition from PA to IPF via mixed clusters

Add 2nd IPF node to 5 node SAP cluster

Existing 5-node PA/IPF SAP cluster with ERP and BW systems

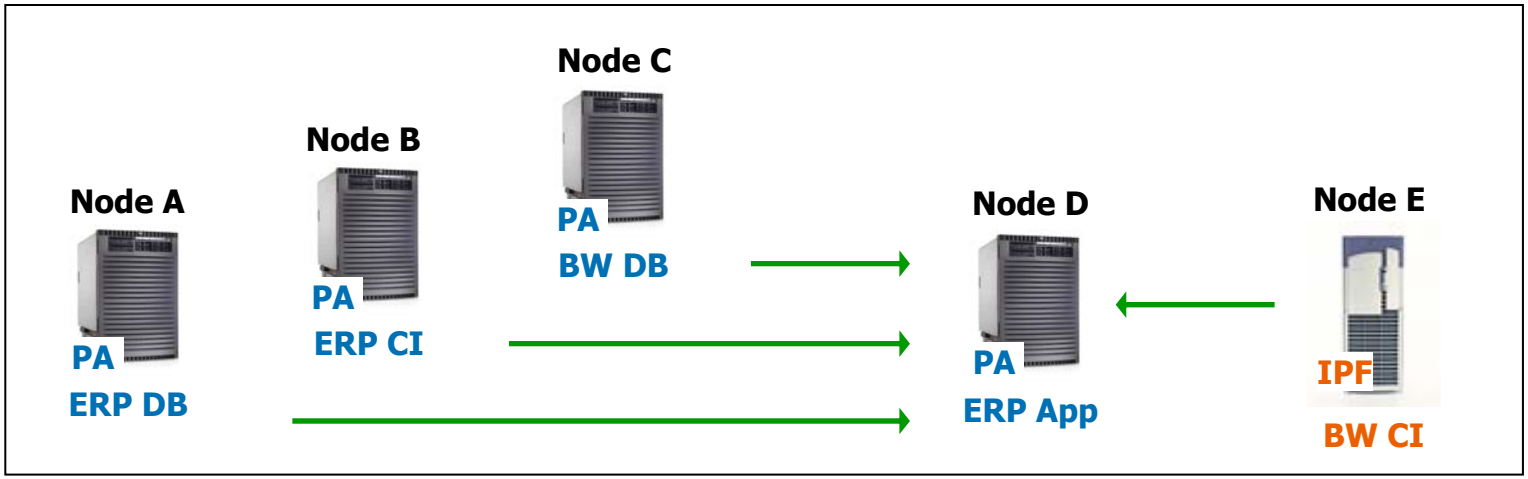

New 5-node PA/IPF SAP cluster with ERP and BW systems

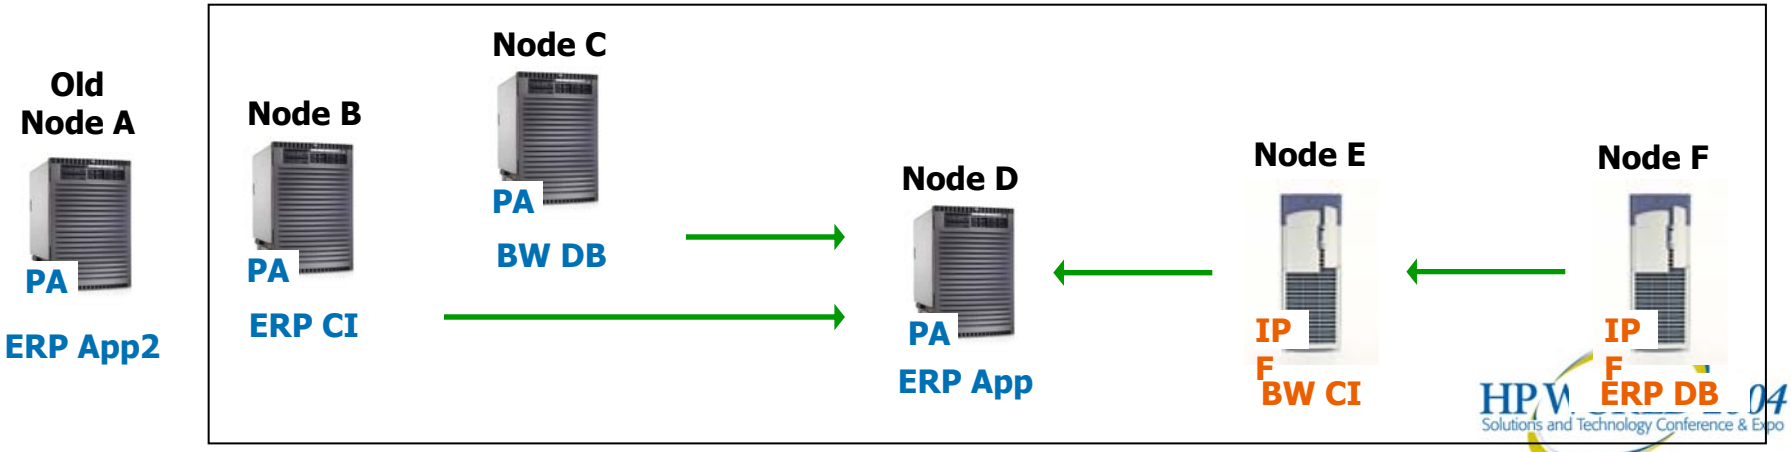

Transition from PA to IPF via mixed clusters

Add 3rd IPF node and recycle all PA nodes as Apps

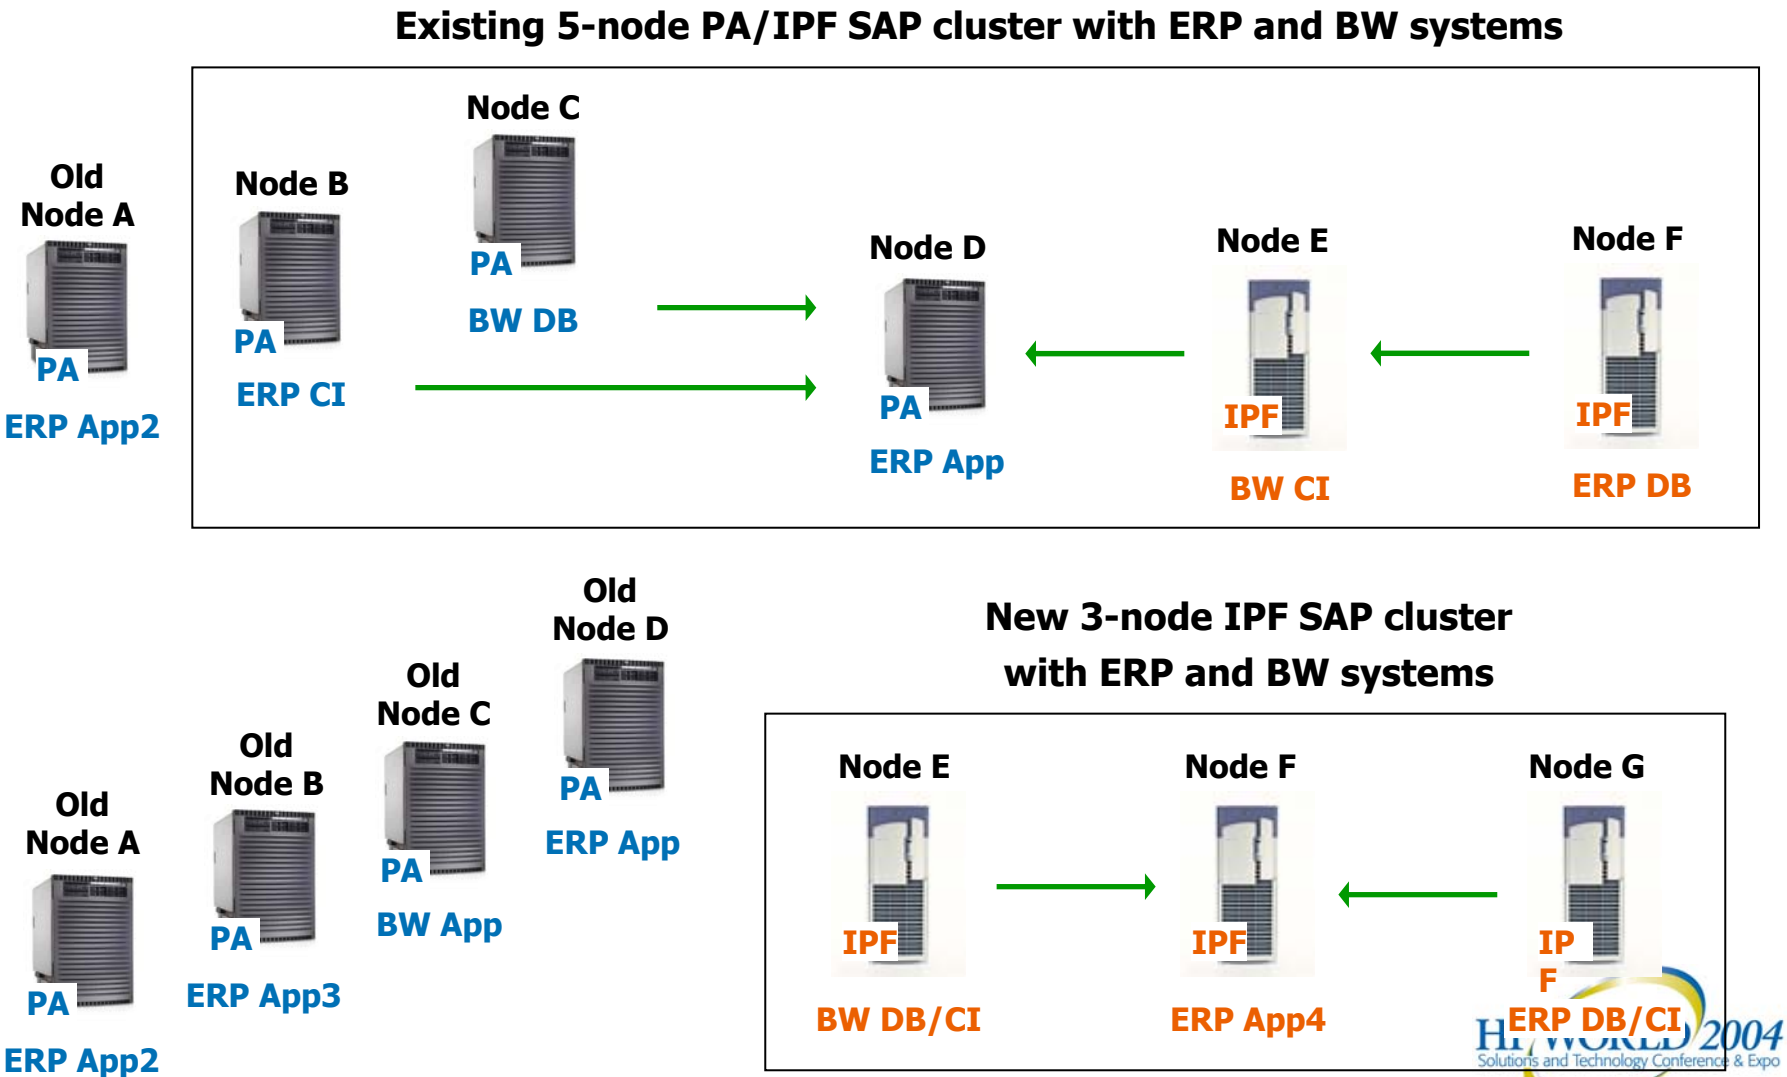

Comparing Serviceguard and TruCluster Features

Target Audience

- The target audience for this presentation includes Customers:
  - with existing Serviceguard or SGeRAC clusters
  - interested in understanding the differences between Serviceguard and TruCluster on HP-UX
  - who want to learn about the High Availability choices in HP-UX 11i Version 3
  - trying to make educated decisions about
    - whether to migrate from Serviceguard to TruCluster on HP-UX
    - when to migrate from Serviceguard to TruCluster on HP-UX
  - who want to understand the process of migrating from Serviceguard to TruCluster on HP-UX

Comparing Serviceguard and TruCluster

Serviceguard and TruCluster

- Clustering solutions developed independently and enhanced over the years to provide premiere high availability capabilities
- There are differences today
  - Each product has its advantages and disadvantages
  - Both products have been successful in the marketplace and meet definite market needs
- The market is evolving – the landscape is changing – customers expect HP to grow with them and to continue to provide premier HA solutions that meet their changing needs
- HP is incorporating TruCluster technology into HP-UX 11i Version 3 and this will be the premier HA solution for the future that will combine the advantages of high availability, scalability and manageability for applications of the future

Serviceguard – a Loosely-Coupled Cluster

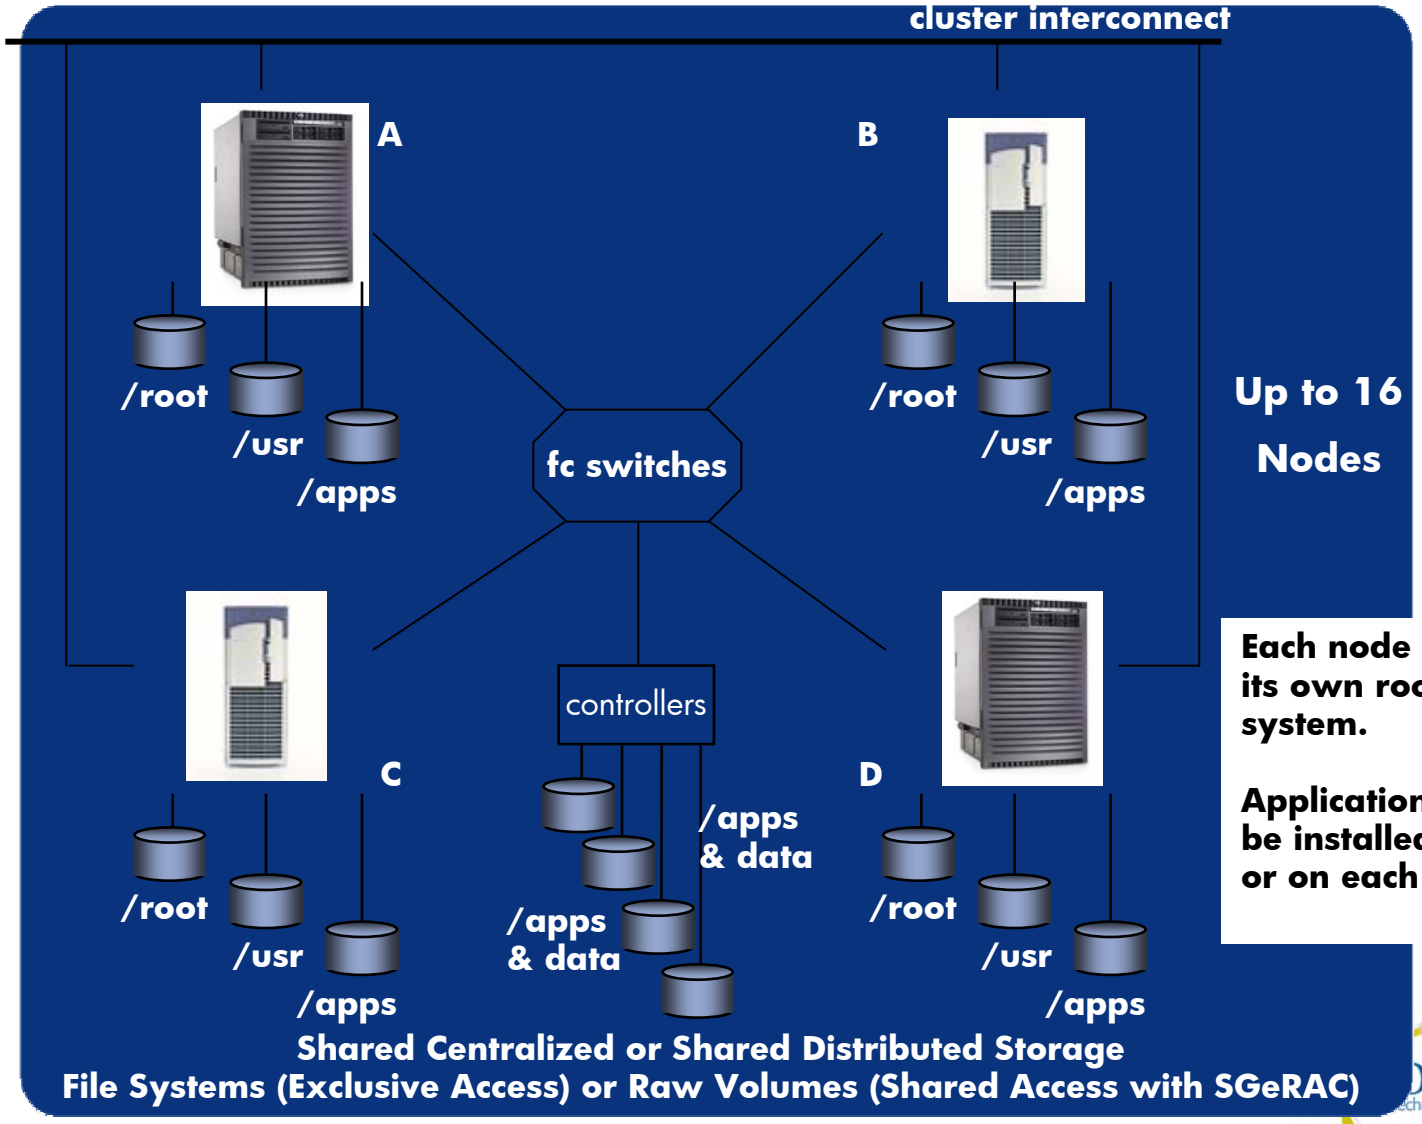

TruCluster – a Tightly-Coupled Cluster

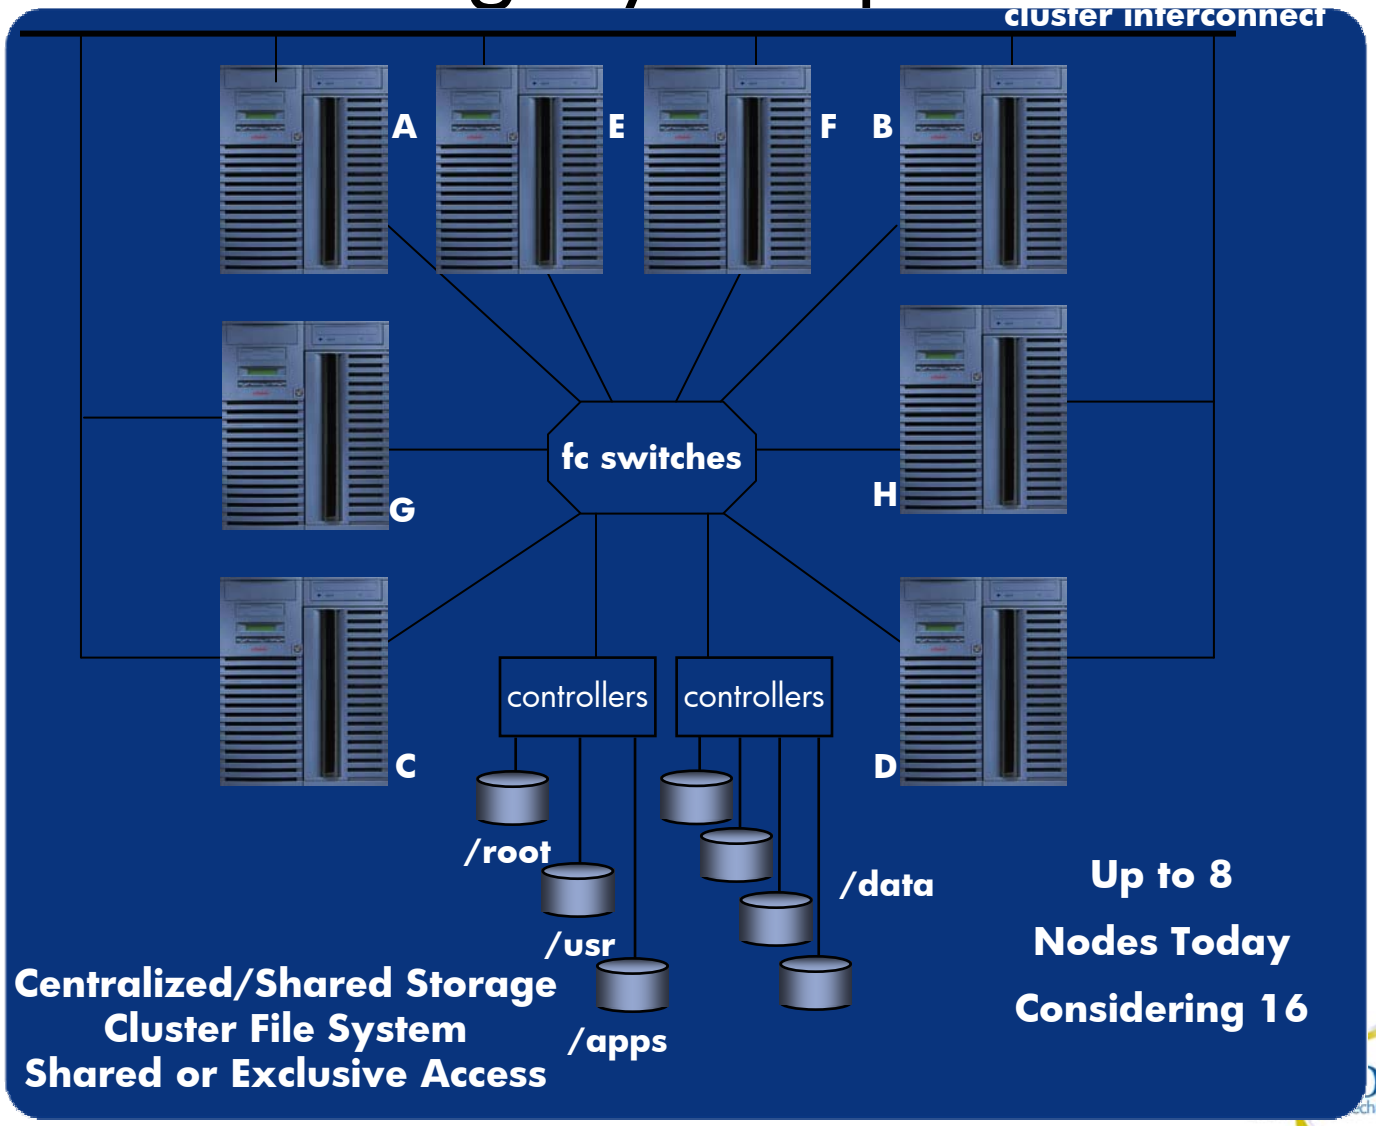

Serviceguard and TruCluster
Maximum Cluster Sizes and Operating Systems

- Serviceguard:
  - 16-node clusters
  - HP-UX and Linux
  - HP-UX 11i version 3 in the future

- TruCluster:
  - 8-node clusters (for HA)
  - Up to 256-node clusters for technical computing
  - Tru64 today
  - HP-UX 11i version 3 in the future
    (8 nodes committed, considering 16)

Both products protect against failures

Serviceguard	Trucluster
The OS	The OS
Hosts	Hosts
Networks	Networks
Applications (services)	Applications
User-defined resources	Defined set of resources

Comparing the Application Interface

The SG Application Package	The TruCluster Application integration with CAA (Cluster Application Availability)
Servers	Servers
Network Identity – Relocatable IP	Network Identity – Cluster Alias
Data	Data is dealt with at the OS level; Volume groups and File Systems are activated/mounted globally
Applications	Applications
User-defined Resources	Other resources
Dependencies require use of optional, contributed toolkit	Dependencies are a built-in feature

Serviceguard and TruCluster

Integration steps:

1. How to startup the application
2. How to shutdown the application
3. How to monitor the application
4. What resources are used by the application

Serviceguard and TruCluster Operating System and File System

- Serviceguard:
  - OS is installed on each node in the cluster; patched separately; during rolling upgrade, can have two versions (can use Ignite server or OV/Software Distributor to ease the effort)
  - Volume groups are activated and file systems are mounted by the application package control script
  - User and password information is administered on each node of the cluster (central management with LDAP server or Service Control Manager)
  - Application software may be installed once on a shared disk or on a private disk of each node
  - File systems are accessed exclusively by one node for read/write access
  - Device naming is usually different on each node
  - Multi-pathing is done by the volume manager or by disk-specific add-on software
  - Each node manages its own PID assignment
  - Volume Groups must be created and kept up to date on every node

Serviceguard and TruCluster Operating System and File System

- TruCluster:
  - OS is installed once for the cluster; patched once; during rolling upgrade, can have two versions installed concurrently
    - Single root file system shared by all cluster members
  - Volume groups are activated and file systems (AdvFS) are mounted globally at boot time; no need for the package manager (CAA) to be involved with this process
  - Single security and management domain
    - e.g., user and password information is administered once, globally
  - Application software is usually installed once and shared by all nodes
  - Cluster File System is available to provide shared read and write access to an AdvFS file system
  - Device naming is identical on all nodes
  - Multi-pathing is done at the OS/driver level, independent of volume managers or disks
  - Additional features when all links on a node fail to access file systems remotely (DRD)
  - Provides cluster-wide unique PIDs and IPC objects

Serviceguard and TruCluster File Systems and Volume Managers

- Serviceguard (including HP-UX 11i version 3):
  - HFS and VxFS file systems
  - LVM, SLVM, VxVM and CVM volume managers

- TruCluster:
  - Tru64:
    - ufs, AdvFS file systems
    - LSM volume manager and raw partitions (no volume manager)
  - TruCluster on HP-UX (HP-UX 11i version 3):
    - HFS/ufs, VxFS and AdvFS file systems
    - LVM and cluster-aware LVM in a post-11.31 release

Serviceguard and TruCluster Failover models

- Serviceguard:
  - Single instance applications in failover mode (restart the application upon failover)
  - Multi-instance parallel application support for specific applications only
    - Oracle RAC (OPS)
    - Informix XPS
- TruCluster:
  - Single instance applications in failover mode (restart the application upon failover)
  - Multi-instance parallel application support for any cluster aware applications with specific support for
    - Oracle RAC

Serviceguard and TruCluster Oracle RAC Support

Serviceguard	TruCluster
Up to 16-node RAC cluster; requires SGeRAC add-on product	Up to 8-node RAC clusters
SGeRAC requires raw logical volumes	Supports use of raw partitions
No cluster file system (CFS)	Supports use of AdvFS file system with the Cluster File System (CFS)

Serviceguard and TruCluster Cluster Interconnect

Serviceguard	TruCluster
Requires redundant heartbeat networks that may be dedicated or shared with other network traffic (highly recommend at least one dedicated heartbeat network)	Requires dedicated cluster interconnect; recommend redundant links Cluster interconnect is used for heartbeat, CFS, remote device I/O
May be shared with Oracle RAC for database lock transfer	May be shared with Oracle RAC for database lock transfer

Serviceguard and TruCluster Networks for Oracle RAC Database Lock Transfer

Serviceguard	TruCluster
– Ethernet	– Ethernet
– HyperFabric 2 High throughput, low latency proprietary interconnect	– Memory Channel High throughput, low latency proprietary interconnect
– Infiniband (with HP-UX 11i v3) High throughput, low latency standards-based interconnect	– Infiniband (with HP-UX 11i v3) High throughput, low latency standards-based interconnect

Serviceguard and TruCluster Network Identity

Serviceguard	TruCluster
–Uses relocatable IP addresses that failover with the application package –Serviceguard can monitor the subnet and failover due to loss of connectivity to the subnet –APA can aggregate NICs for load balancing and failover	–Uses Cluster Aliases to provide transparency to the user –Cluster Alias can be a monitored resource (dependency) with CAA
	–Can provide additional features for parallel or multi-instance applications

Serviceguard and TruCluster Resources and Event Management

Serviceguard	TruCluster
Uses EMS monitors for event management and failover decisions	Uses EVM for event management and notifications
SG has "a priori" knowledge of the state of resources on all nodes through polling	TruCluster has knowledge of the initial state of a resource through polling at boot and gains later knowledge through state changes (events)
Free software development kit available for creation of user-defined resources	–Only pre-defined resources are available –Resource library must be linked into CAA

Protecting against split brain and data corruption

- Serviceguard uses a "tie-breaker" to prevent "Split-Brain" of the cluster
  - Cluster lock disk (LockLUN for SG/Linux)
  - Quorum service
  - Arbitrator nodes for disaster tolerant, multi-data center
  - Uses Safety Timer in kernel to prevent data corruption from a system that was hung and removed from the cluster
- TruCluster uses a "tie-breaker" to prevent "Split-Brain" of the cluster
  - Quorum disk
  - Uses I/O fencing to prevent data corruption from a system that was hung and removed from the cluster

Features to force failover even when a system hangs

- Serviceguard offers the ability to remove a node from the cluster in case of service or package failures (SERVICE & NODE FAILFAST) to prevent data corruption
  - NODE FAILFAST is often used when a package fails to halt due to the inability to cleanly unmount a file system
  - SERVICE FAILFAST would be used in application-specific cases
- TruCluster does not have a similar capability
  - However, file systems are usually mounted globally, and therefore, this feature should not be needed
  - I/O fencing prevents a previously hung node from issuing I/Os after the cluster has reformed

Failover Time Definition

- Cluster failover time:
  - Node failure detection
  - Cluster reformation
  - Node selection
  - Trigger application startup
- Application failover time
  - Volume group activation (if applicable)
  - File system recovery (if applicable)
  - File system mount (if applicable)
  - Configure application IP addresses
  - Application startup
  - Application recovery (if applicable)
- NOTE: application failure without node failure is somewhat different

Cluster Failover Times

- Serviceguard:
  - Typical 30- to 45-second failover (dependent upon disk link technology and use of cluster lock disk)
  - Fast Failover feature for 5-second failover with restricted configurations
    - 2-node clusters
    - multiple heartbeat links
    - LVM or VxVM
    - quorum server
    - SCSI or FibreChannel storage)
- TruCluster
  - Typical less than 60-second failover on Tru64
  - Targeting 30- to 45-second failover for TruCluster on HP-UX

Managing a cluster

	Serviceguard	TruCluster	TruCluster on HP-UX (11i v3)
Command Line Interface (CLI)	Yes	Yes	Yes
Terminal-based User Interface (TUI)	Yes (SAM, also Motif-based)	Yes (sysman)	No
Java-enabled GUI for:			
Monitoring	Yes	Yes (sysman)	Yes
Administration	Yes	Yes	Yes
Configuration	Yes	Yes	Yes
Role-Based Access (RBA)	Yes	Yes	Yes

HP World 2004 - Seminar # 3517 - Copyright 2004 Hewlett-Packard Company

HA and Partitions

- Serviceguard can be used with hard and virtual partitions
- TruCluster can be used with hard partitions (Tru64/Alpha)
- TruCluster on HP-UX:
  - Initial HP-UX 11i version 3 support will include hard partitions
  - Support for virtual partitions will soon follow

Workload Management and resource allocation

- Serviceguard
  - integrated with Workload Manager (WLM) and Process Resource Manager to automatically change resource (CPU, memory, disk) allocations upon failover
  - Will be integrated with the new Global WLM (gWLM)
- Tru64 (PRM-like product from ISV)
- TruCluster on HP-UX:
  - TruCluster on HP-UX will be integrated with the new gWLM

Workload Manager (WLM) and Process Resource Manager (PRM) with SG

- Service Level Management solutions that can be easily integrated into an Serviceguard environment to provide
  - management of computing resources (CPU, memory and I/O utilization) according to desired SLOs both before and after application failover among nodes
- Applications can be assigned maximum resources when they run on their "preferred" server
- After a failover to a different node, the resource utilization for the existing and new applications can be dynamically reset to:
  - favor important applications
  - restrict resource utilization of less important applications

Disaster Tolerant Solutions

- Serviceguard provides a family of fully supported disaster tolerant solutions
  - Extended Campus Cluster (up to 100 km)
  - Extended RAC Cluster (up to 100 km)
  - MetroCluster (up to 100 km)
  - ContinentalClusters (unlimited distance)
  - Multi-site disaster tolerance (combination of MetroCluster and ContinentalClusters)
- TruCluster provides
  - a consulting solution for the campus (up to 7 km) environment
  - Geoplex (a consulting solution similar to MetroCluster for up to 100km
- TruCluster on HP-UX will have support for disaster tolerant solutions similar to those with Serviceguard

Application integration and toolkits

- Serviceguard provides
  - A collection of fully-supported integration templates for certain popular third-party applications (ECMT)
  - A supported product for integration with SAP (SGeSAP)
  - A supported template for HA NFS
  - Hundreds of validated ISV applications
- TruCluster provides
  - a collection of examples for certain popular third-party applications
  - A consulting solution for integration SAP
  - HA NFS and DNS is inherently integrated into TruCluster
- HP-UX will continue to offer supported toolkits for major ISV applications

High Availability Choices in HP-UX 11i Version 3

HP-UX 11i Version 3 High Availability

- Two Cluster Personalities, One Application Package Manager
  - Serviceguard
    - Compatible, enhanced Serviceguard with new features (same SG version runs on HP-UX 11i Versions 1, 2 and 3, and on Linux)
      - Same familiar package manager with enhancements such as
        - package dependencies
      - Fast failover (available soon with Serviceguard 11.16)
      - AdvFS support on HP-UX 11i v3
  - Serviceguard with TruCluster Technology
    - Compatible with TruCluster for customers migrating from Alpha and Tru64 Unix
      - Single system image with cluster-wide root file system
      - Cluster file system for AdvFS and other file systems
      - Consistent cluster-wide device naming
    - Same package manager as with Serviceguard (may require some conversion)

HP-UX 11i Version 3 High Availability

- Two Cluster Personalities, One Application Package Manager

  - One template for integrating applications into either cluster personality
  - One supported toolkit for both cluster personalities
  - Consistent application support
    - Easier validation and maintenance of ISV applications
    - Easy movement of hundreds of ISV applications to HP-UX 11i v3
  - Enhancements such as
    - Package dependencies

Reasons for Choosing a Particular Solution

Why Stay With Serviceguard?

- Serviceguard and Serviceguard Extension for RAC (SGeRAC) will be available on HP-UX 11i v3

  - I have an existing cluster environment and architecture that works well with my application
    - I use Serviceguard Extension for Fast Failover
    - 16-node clusters
    - Quorum Service
    - HyperFabric
    - LVM Disk Volume Groups
    - High Availability Reference Architectures
    - EMS monitors integrated with Serviceguard
    - SGeRAC and SGeSAP

Why Stay With Serviceguard?

- Serviceguard and Serviceguard Extension for RAC (SGeRAC) will be available on HP-UX 11i v3
  - My application is not yet supported or certified with TruCluster on HP-UX
  - I need independent OS images (some applications, e.g., Telco, desire a looser coupling)
  - My application is not compatible with a Single System Image environment
  - I need to deploy a disaster tolerant solution on 11.31 immediately
  - I want a consistent environment for both HP-UX and Linux
  - I am satisfied with Serviceguard and I don't want to retrain my staff
  - I need to use Oracle RAC version 9.2
  - The application package manager will be the same for both personalities, as well as Linux, and will continue to be enhanced
  - You can choose to migrate from Serviceguard to TruCluster on HP-UX in the future when your needs require it, when you are ready, or when your ISV application is validated

Why Transition from Serviceguard to Serviceguard withTruCluster Technology?

- TruCluster has some features that are not available with Serviceguard
  - I want the cluster manageability improvements that come with Single System Image
    - Single root file system
    - Single cluster-wide security domain
    - Administration and configuration from any cluster member, such as
      - User and password administration (single security and management domain)
      - Installing applications
      - Software and patch installation
    - Consistent cluster-wide device naming
    - Unique, cluster-wide name space for PIDs and IPC objects
    - Cluster-aware LVM (available in a later *Fusion* release)

Why Transition from Serviceguard to Serviceguard withTruCluster Technology?

- TruCluster has some features that are not available with Serviceguard
  - I want to use the Cluster Alias capability that provides the appearance of a single system to network clients and peers
  - I want to use a Cluster File System
    - instead of raw volumes for RAC
    - Instead of NFS for other applications
  - I have a cluster aware application that is compatible with the single system image and will be able to scale across nodes using the lock manager capabilities of TruCluster
  - I want an out-of-the-box clustered solution for printing, NIS, NFS, DHCP, e-mail, etc.
  - I want to deploy Oracle RAC with a CFS and SSI architecture

When to Upgrade to HP-UX 11i Version 3

HP-UX 11i v3 HA Adoption Profile

- Both Serviceguard and TruCluster will be offered on HP-UX 11i v3
- Customers can choose either personality
- Customers will NOT be forced to transition from Serviceguard to TruCluster
- HP will continue to enhance Serviceguard in the future

When to Upgrade

- When implementing a new cluster with HP-UX 11i version 3
- When implementing a new application on a new cluster
- During a technology upgrade for an existing application and want the benefits of the Single System Image (SSI) and Cluster File System (CFS)
- During a major application (including ISV) software upgrade for an existing application and want the benefits of SSI and CFS
- When implementing an ISV application that requires SSI and/or CFS
- When planning the development or porting of a parallel, cluster-aware application to HP-UX
- When upgrading an existing Serviceguard cluster to 11i v3 (to either Serviceguard or to TruCluster on HP-UX)

If You Choose to Stay With Serviceguard

High-Level Process

- Using the normal rolling upgrade process:
  - Upgrade HP-UX from an earlier 11i version to 11i Version 3
  - Make any needed changes to the Serviceguard control script
  - Convert from existing file system to AdvFS, if desired
    NOTE: you will not be able to do this with the application package running
  - Upgrade Serviceguard to the 11i Version 3-compatible version
  - Testing, staff training, and documentation

If You Choose to Transition From Serviceguard to Serviceguard with TruCluster Technology

Transitioning from Serviceguard to Serviceguard with TruCluster Technology

- There are many great reasons to transition
- HP wants to make it as easy as possible
- Transition Assistance
  - High-level transition process presentation and documentation
  - Detailed transition process documentation
  - Analysis and conversion tools
  - Consulting services

What Can I Do to Ease My Transition?

- Upgrade your cluster to HP-UX 11i v2
  - Compatible OS for both PA-RISC and IPF systems
- AdvFS will be available for use with Serviceguard on HP-UX 11i v3
  - AdvFS has many features that Serviceguard customers may want to take advantage of
  - Migrate my data from HFS or VxFS to AdvFS in my Serviceguard environment as soon as possible
  - The Cluster File System will only be available once you fully migrate to TruCluster on HP-UX
- VxVM will not be supported with TruCluster on HP-UX
  - There will be a cluster-aware LVM in HP-UX 11i v3 which can be used with AdvFS
  - Migrate from VxVM to LVM in my Serviceguard environment as soon as possible
  - Doing so will significantly ease the full migration to TruCluster on HP-UX
- Older PA-RISC systems will not be supported on HP-UX 11i v3
  - Start upgrading my existing clusters to the newest PA-RISC systems
  - Start creating new clusters with IPF systems

Transition Timeline – Easing the Transition

Today **4Q04** **2H05** **2H06**

- Analyze environment
- Make plans for transition

- Upgrade existing clusters to HP-UX 11i Version 2
- Possible mixed cluster support

- Upgrade existing clusters to HP-UX 11i Version 3
- Upgrade older PA-RISC systems to supported PA or IPF systems
- Convert volume managers to LVM
- Convert file systems to AdvFS

- Upgrade to HP-UX 11i Version 3 Fusion release
- Transition from Serviceguard to TruCluster

Transitioning from Serviceguard to Serviceguard with TruCluster Technology on HP-UX 11i v3

High-Level Transition Process

- Pre-analysis of hardware, architecture and software
- Documentation of impact of changes (if any) to chosen HA paradigm
- Conversion
- Installation
- Testing
- Staff Training
- Documentation

High-Level Transition Process

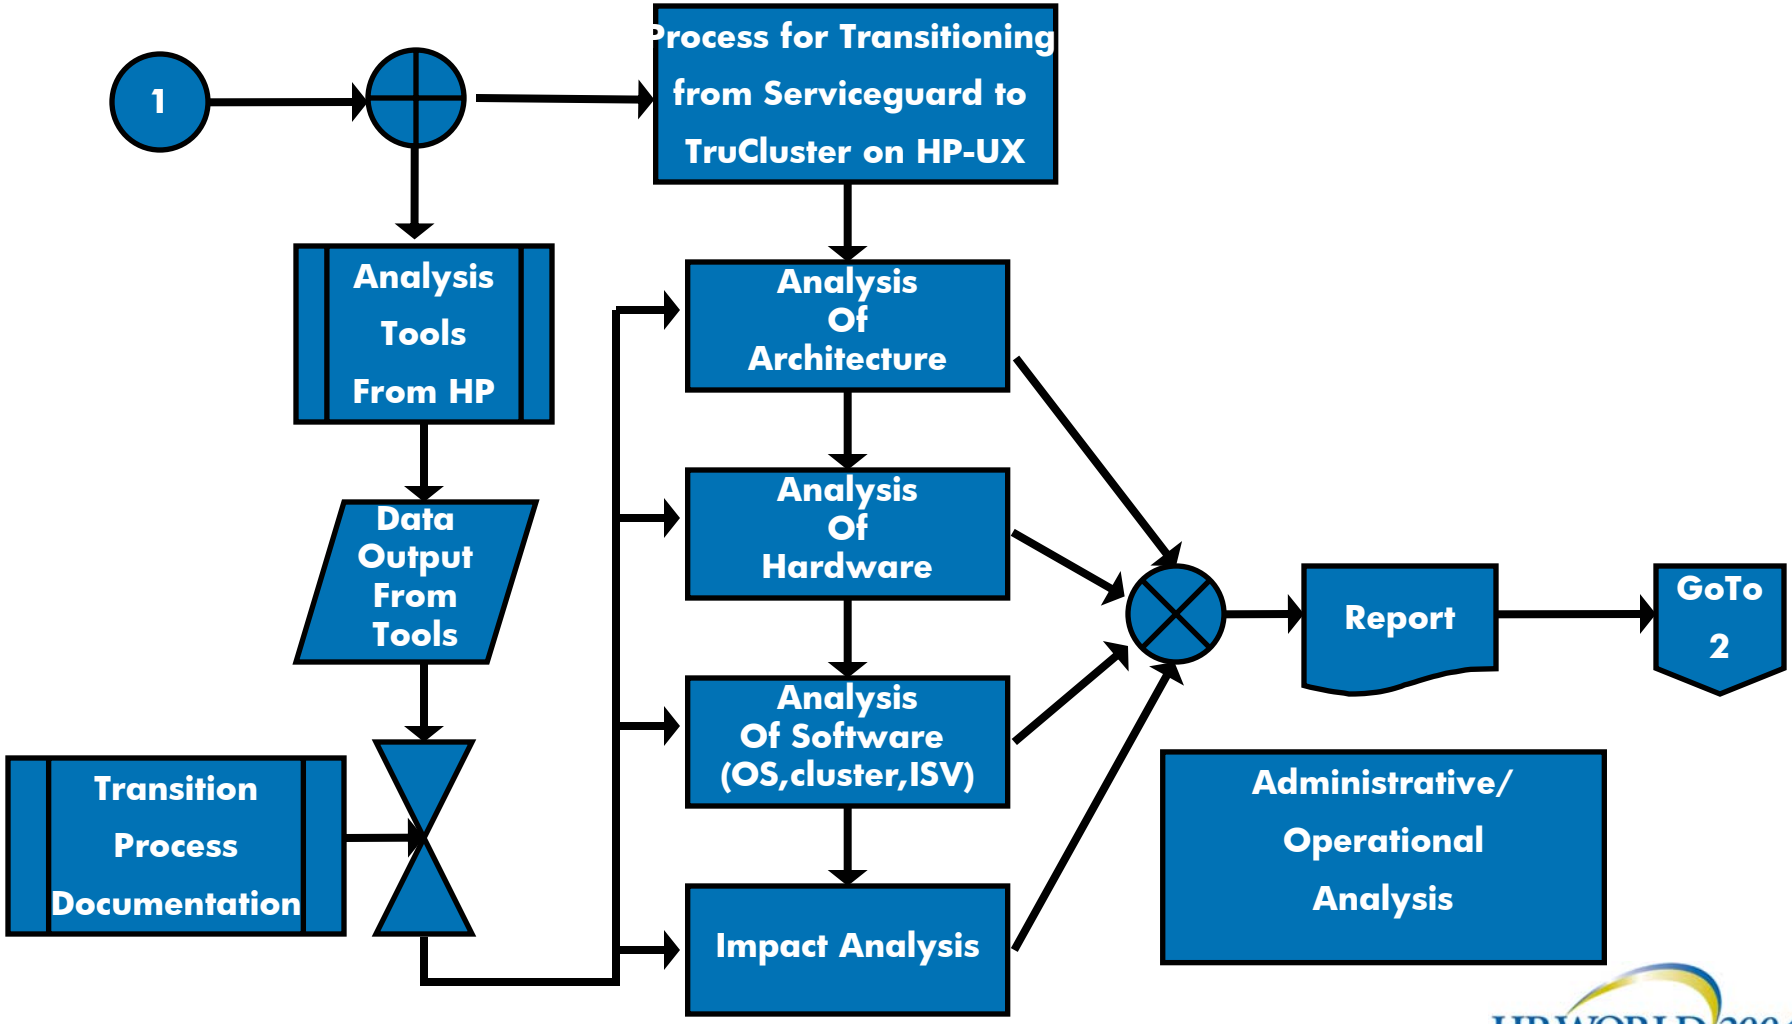

High-Level Transition Process

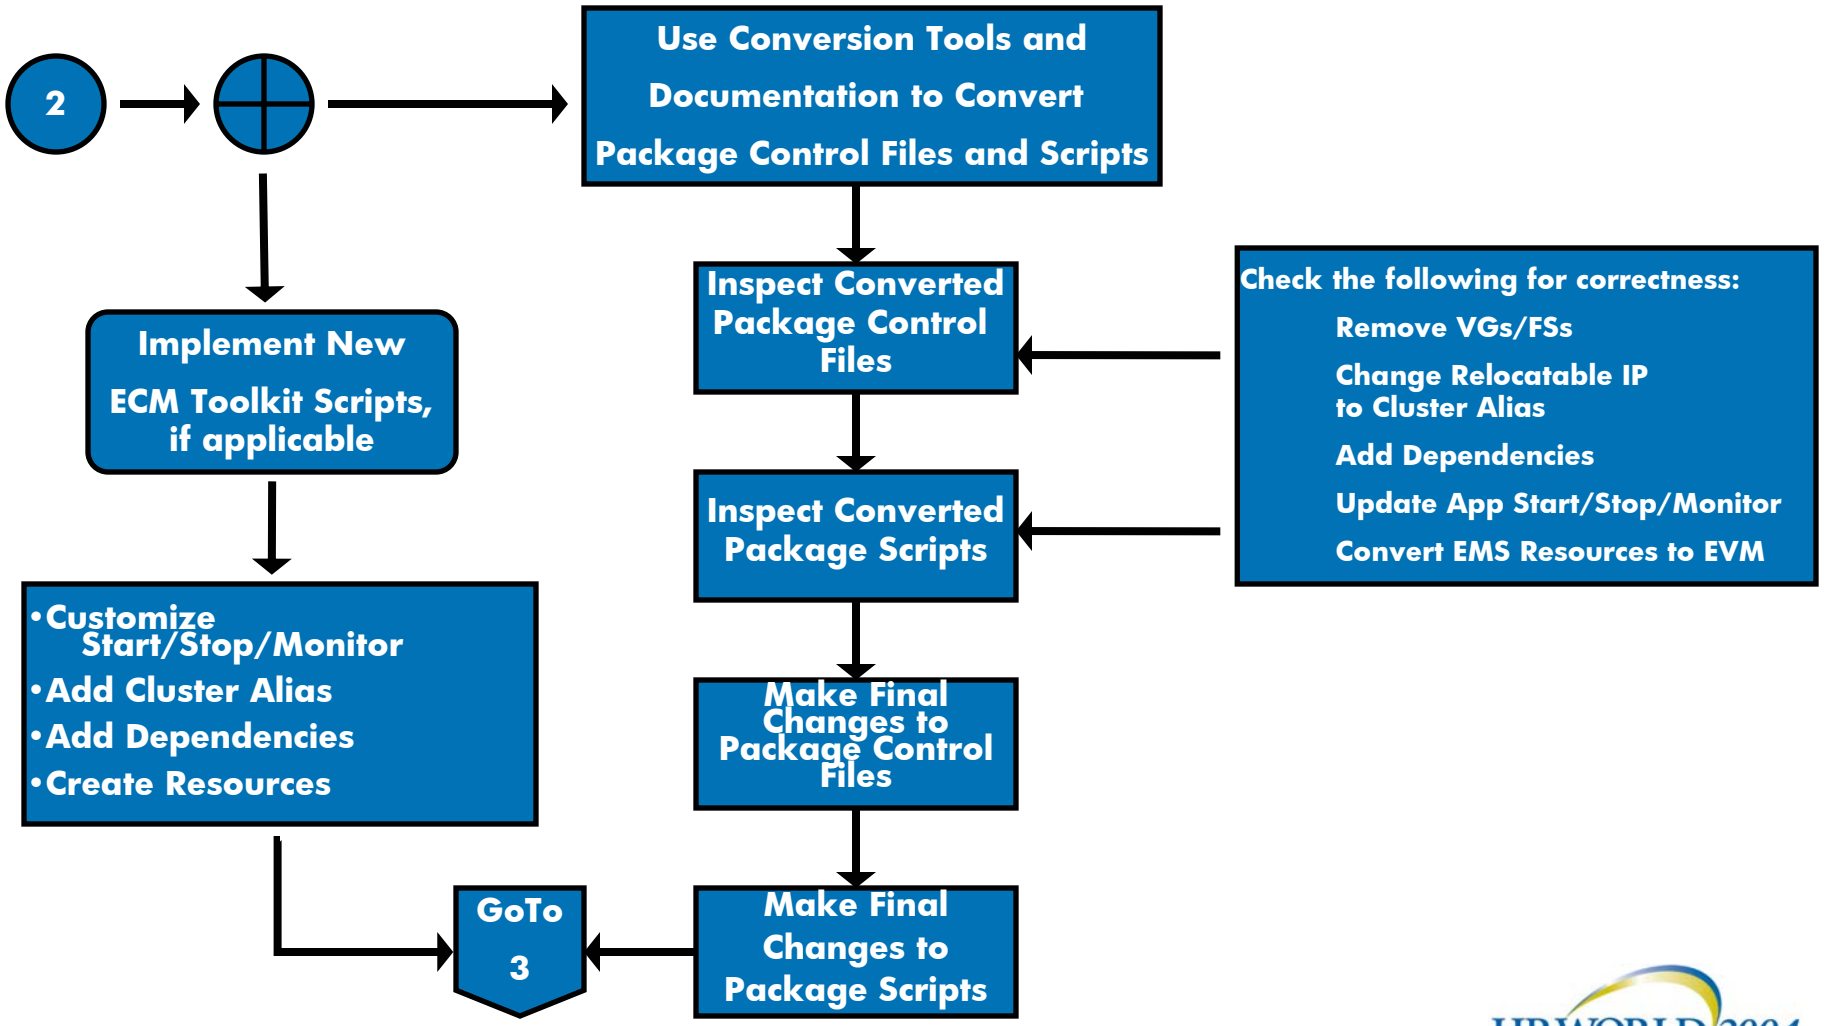

High-Level Transition Process

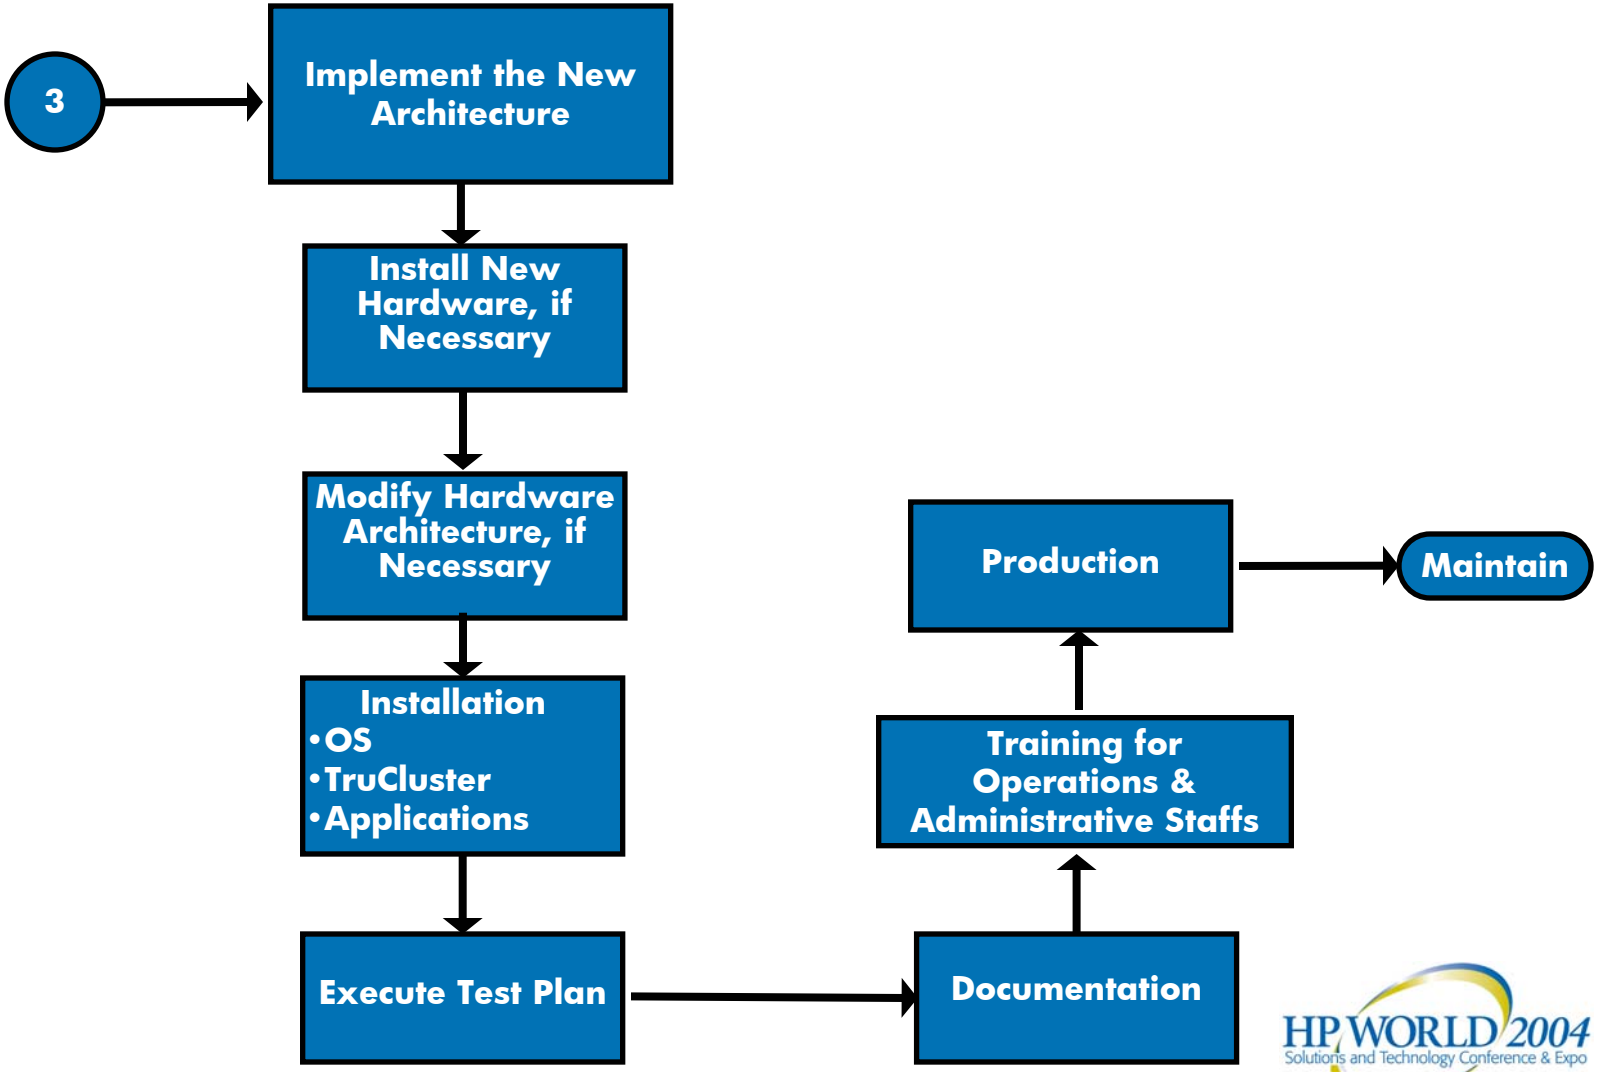

Questions?

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional