

Session 2274

Implementing Blade Clusters with the MSA1000 Using Microsoft Cluster Server

Lou Lawrence

AJ Moore

Sean Beard

Systems Engineers
High Availability Systems Engineering
Hewlett Packard



Phase 1 – Additional Information

- Getting Started
 - Use CTRL+ALT+INS to log in to the VM
 - Use the soft copy of stationinfo.txt
- Section 1
 - 4. Pay close attention to directory names (ex. sps.40b)
 - 6. Copy the contents of spinstall to the DS directory
- Section 2
 - 1c. Rename the cluster file to *stationxclus.ini* not *blclus.ini*
 - 3b. Physical NIC enumeration is unique in BL20p
 - 6. Do not append the domain suffix in the cluster file
- Section 5
 - Complete Section 5 and STOP!

Phase 2 and Phase 3 – Additional Information



■ Section 6

- 5. If this task takes longer than 5 minutes to complete, refer to the **Note**.

■ Section 7

- 10. Create a quorum drive that is at least 510MB.
- 11. If you would like more information on SSP, please see an instructor

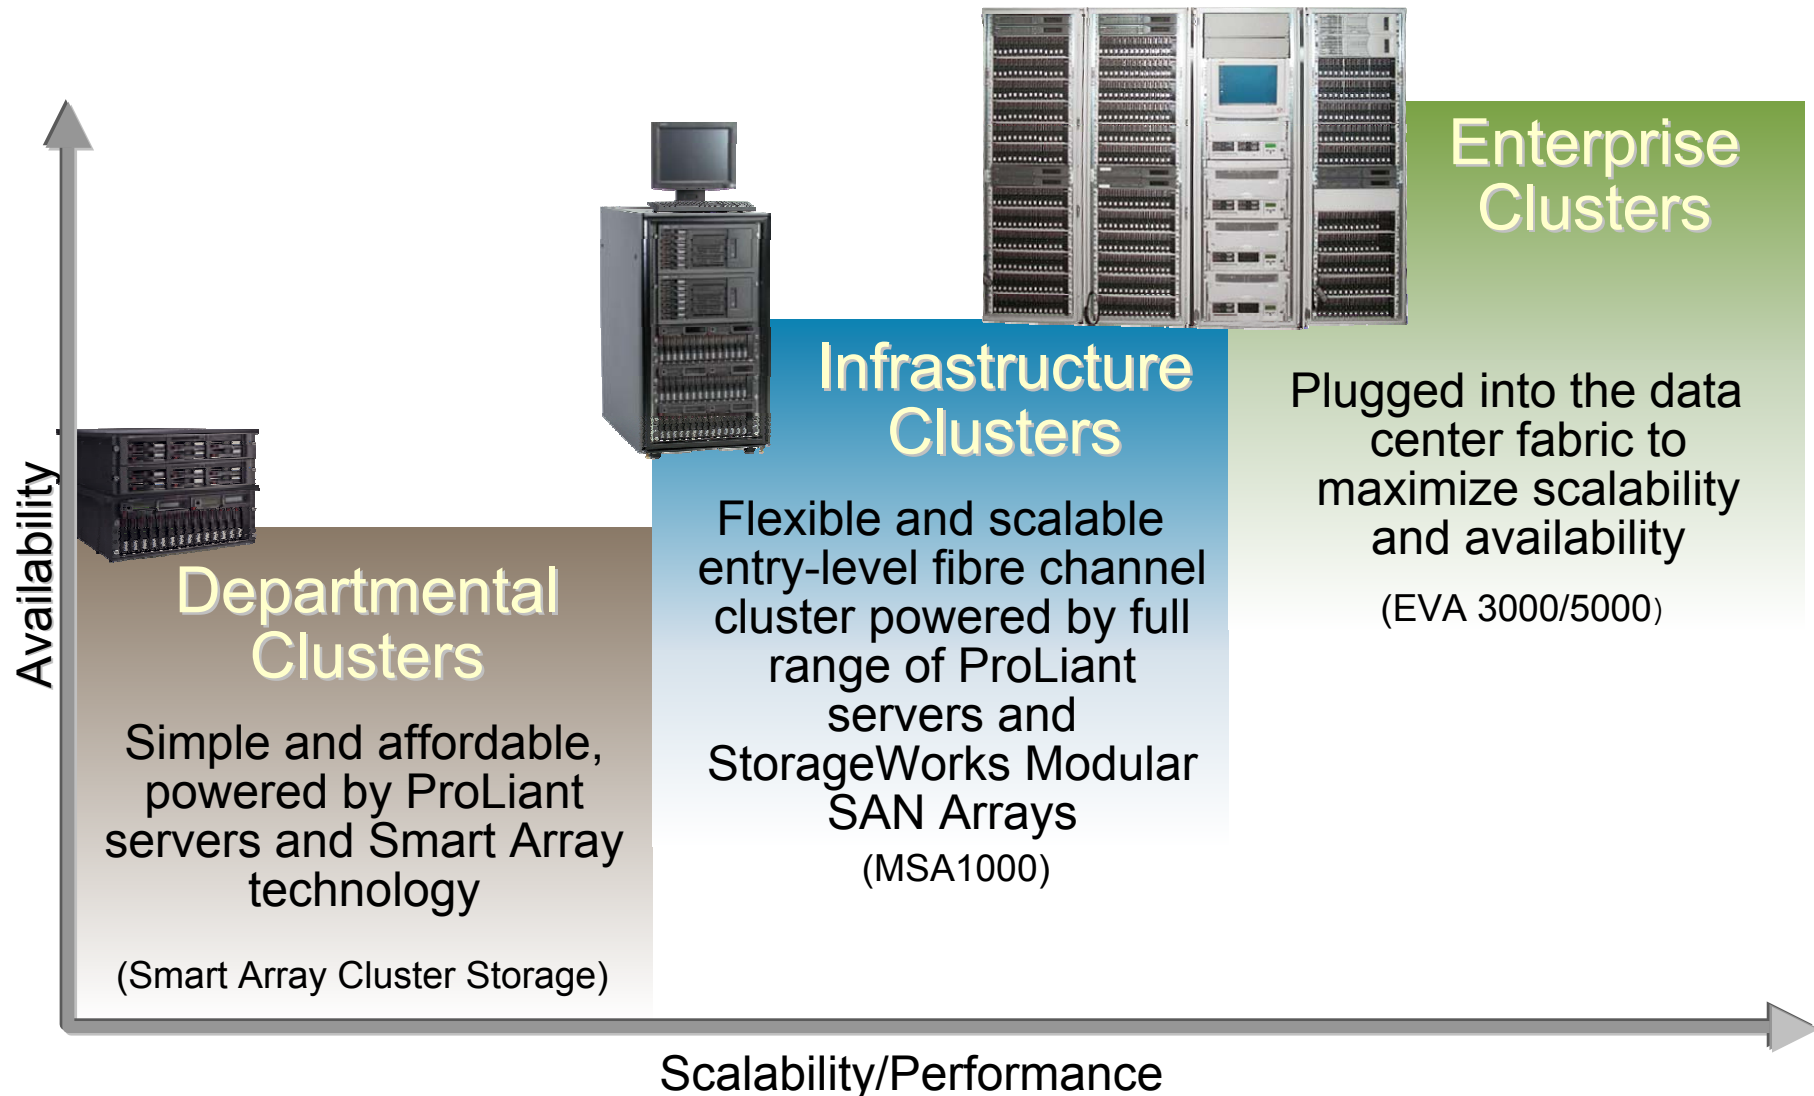
■ Finishing up

- Please complete the survey for this lab (2274) online at hpworld.com
- Thank you for attending!

Agenda

- ProLiant Cluster product offerings
- ProLiant Blade Cluster SAN configurations
- ProLiant Blade Clusters overview
- ProLiant Blade Cluster deployment overview
- ProLiant Essential Rapid Deployment Pack
- ProLiant Blade Cluster deployments using RDP 1.40
- ProLiant Blade Cluster deployment walkthrough
- Questions

ProLiant Cluster Product Offerings



ProLiant Cluster HA/F200 (MSA1000)

- Multi-cluster support
- 2Gb support
- 6 TB raw storage capacity
- Max 20 cluster nodes
 - 4 total clusters
- Interconnect devices
 - 1Gb / 2Gb SAN switch
 - 8 port internal switch
- Secure Path support



ProLiant Cluster HA/F500 (Enterprise Virtual Array)

- Multi-cluster support
- Virtualization support
- 2Gb support
- Max 32 cluster nodes
- Interconnect devices
 - 1Gb / 2Gb SAN Switch
- SAN Script support
- Secure Path support



ProLiant Cluster HA/F500 (MA8000/EMA12000)

- Multi-cluster support
- 2Gb support
- Interconnect devices
 - 1Gb / 2Gb SAN Switch
- Command Scripter support
- SAN Management Appliance support
- Secure Path support



ProLiant Blade Cluster SAN Configurations

- ProLiant BL Storage Connectivity
- Fibre Channel Adapters
- BL20p G2 Interconnect Options
 - G2 Patch Panel 2
 - GbE2 Switch
- BL40p Interconnect Options




ProLiant BL Storage Connectivity

- ProLiant BL10e and BL20p
 - Network attached storage
 - No cluster support
- ProLiant BL20p G2 (Cluster Support)
 - Fibre channel (FC) via mezzanine card (with a pass-through connectivity option)
 - GbE2 interconnect switch or G2 Patch Panel 2
 - Network attached storage
- ProLiant BL40p (Cluster Support)
 - (2) PCI-X slots that can be used for storage HBAs
 - Network attached storage

Fibre Channel Adapters – BL20p G2 Fibre Adapters

- Two ports on each FC mezzanine card provide a redundant fibre connection to each server blade
- Supports up to 2Gb speed
- Based on QLogic chipset
 - Supports HP StorageWorks arrays
 - Option for ProLiant BL20p G2 (not ProLiant BL20p G1)



ProLiant BL20p G2

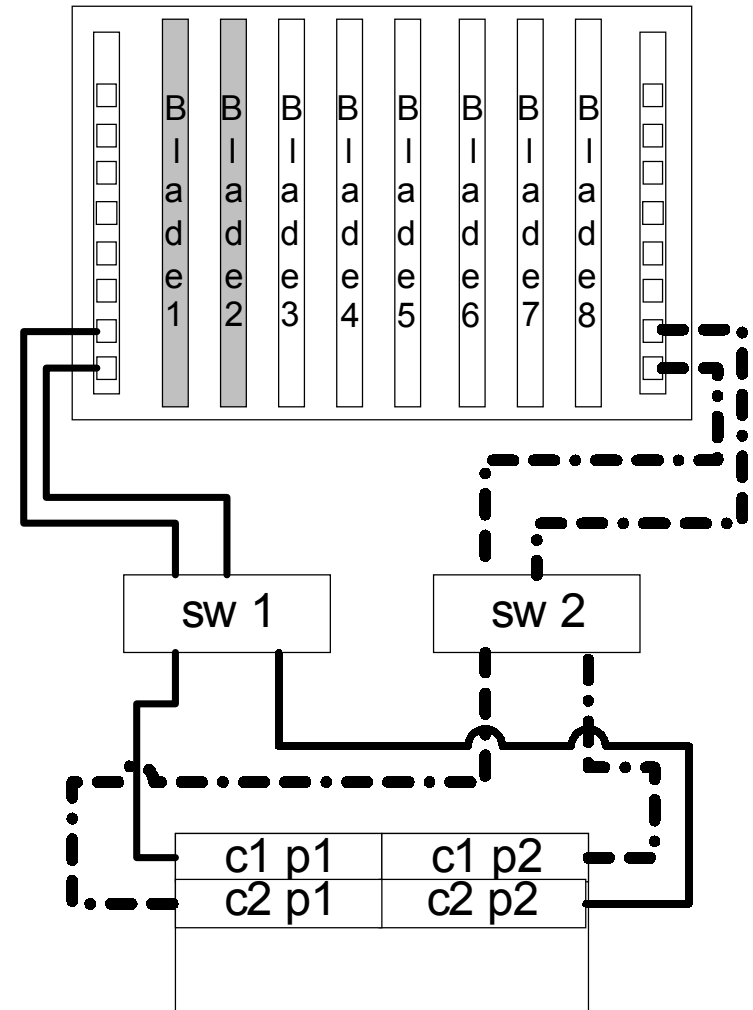


Fibre Channel Adapters – BL40p Fibre Adapters

- Adapters support up to 2Gb speed
- Standard single port PCI Fibre Channel Adapters (FCA)
- FCA2101 support
 - Based on the Emulex chipset
 - Supports HP StorageWorks arrays
- FCA2214 for Linux support
 - Based on the QLogic chipset
 - Supports HP StorageWorks arrays
- 64-Bit/33-MHz PCI-to-Fibre Channel HBA for Windows support
 - Based on the Emulex chipset
 - Supports HP StorageWorks arrays


BL20p G2 Interconnect Options

- Connect to SAN switches
 - G2 Patch Panel 2 FC pass through
 - Optional GbE2 Integrated Switch with FC pass through
- Redundant FC path using the embedded FC mezzanine




BL p-Class G2 Patch Panel 2

- Allows pass through of both LAN and SAN signals
- LAN Ethernet signals (16) at rear
- SAN FC signals (8) at front
 - FC cables can be routed through the a channel in the server blade enclosure for rear cable management




Server Blade
Enclosure (front)



G2 Patch Panel 2

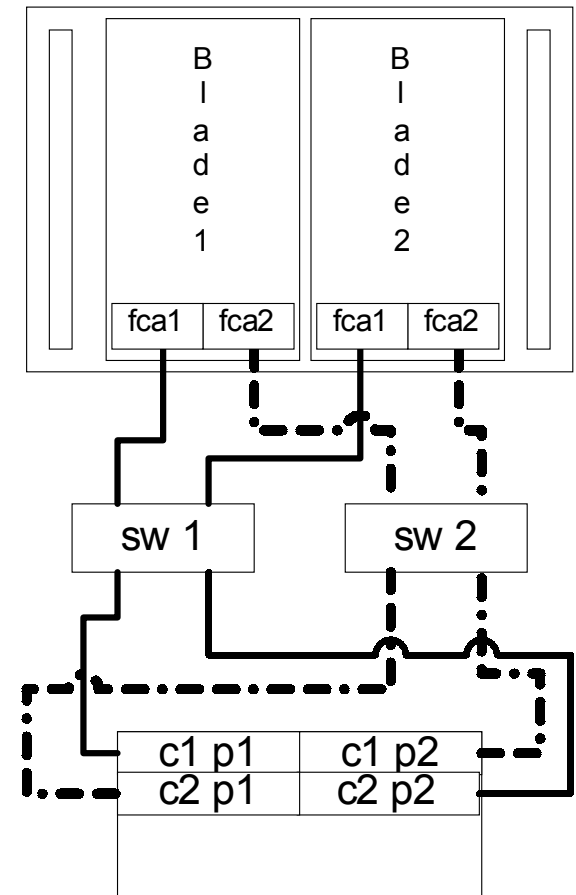
BL p-Class GbE2 Interconnect Switch

- Next generation ProLiant BL p-Class integrated switch
- Pass-through of the ProLiant BL20p G2 FC signals
- SAN interconnect module at rear of each switch
 - 8 FC GBICs (LC connectors) per switch



BL40p Interconnect Options

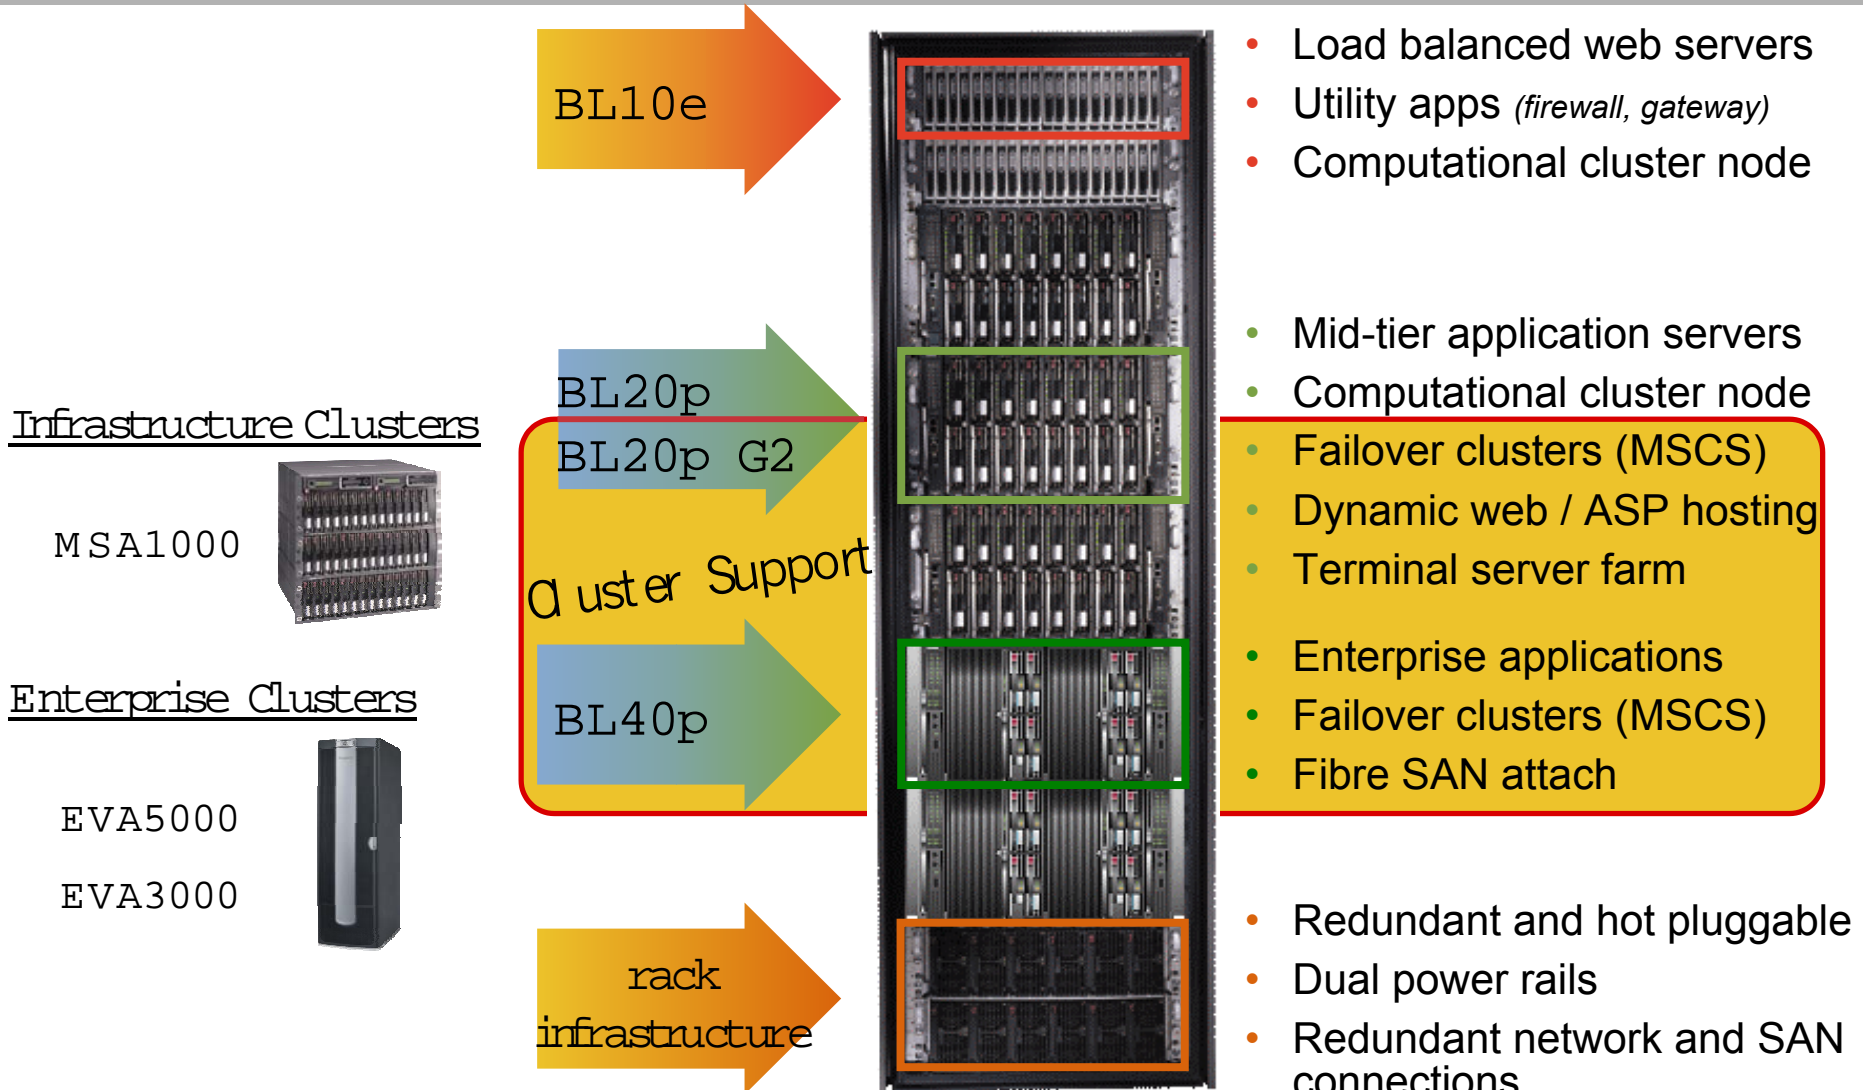
- Direct connection to SAN switches from FCA
 - Uses standard PCI Fibre Channel Adapters
- (2) PCI-X slots allow redundant Fibre Channel paths
- Uses standard SAN switches



ProLiant Blade Clusters Overview

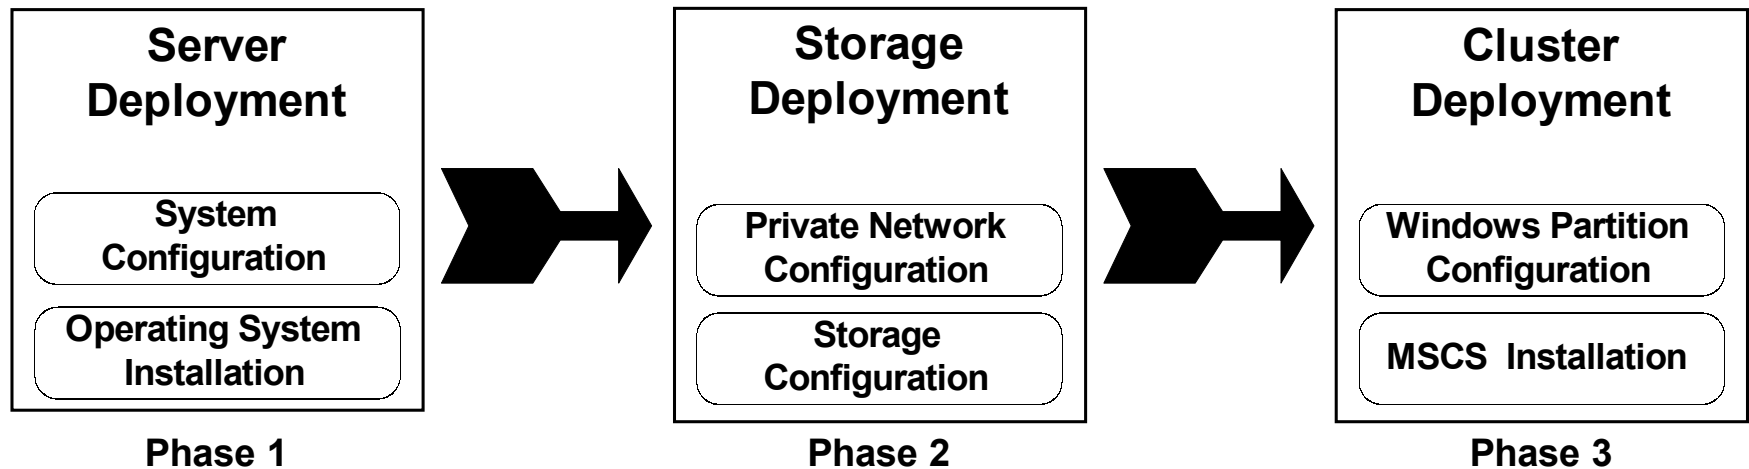
- ProLiant BL20p G2
 - Uses dual port Fibre Channel mezzanine card
- ProLiant BL40p
 - Uses standard Fibre Channel adapters
- Supported external storage enclosures
 - StorageWorks Enterprise Virtual Array (EVA 3000/5000)
 - StorageWorks Modular Array (MA8000/EMA12000)
 - StorageWorks Modular San Array (MSA1000)
- MSA1000, EVA 3000/5000, MA8000
- Will use standard HA/F200 and HA/F500 Cluster Kits
- Investigating Lifekeeper for Linux on Blades

ProLiant Blade Failover Cluster Support



ProLiant Blade Cluster Deployment

Conceptual Overview of Cluster Deployment



- Phase 1 - Automated server configuration
- Phase 2 - Manual private **network** and storage configuration
- Phase 3 - Automated Windows partition and MSCS configuration

ProLiant Blade Cluster Deployment - Phase 1

Automated server configuration with a deployment job

- Configure the blade server hardware
 - CONREP
 - ACR
 - Create partition
- Install and configure Windows
 - Install Windows
 - Install PSPs
- Deploy SAN Deliverables
 - Fibre Channel Adapter (FCA) Driver
 - FCUtil (Fibre Channel Setup Utility)
 - StorageWorks Secure Path

ProLiant Blade Cluster Deployment – Phase 2



- Manual private **network** and storage configuration
 - Use iLO, TS, or RDC for these tasks
- Configure the cluster interconnect
 - Choose from 2 free connections in BL20p G2
 - Choose from 4 free connections in BL40p
- Configure the external storage

ProLiant Blade Cluster Deployment – Phase 3

Automated Windows partition and MSCS configuration with a deployment job

- Configure logical drives for use by Windows
 - Create Windows partitions
 - Format partitions
- Create or join the cluster
 - Primary node creates the cluster while the secondary node(s) waits
 - Secondary node(s) joins when the primary node creates a flag

ProLiant Essentials Rapid Deployment Pack

Features of RDP

- GUI, console-based deployment server
- Built-in PXE services and PXE image tools
- Network booting for headless deployment
- Deploy via scripting or imaging
- Built in script generation and editing
- Remote power control (via WOL, RILOE II, iLO)
- Server configuration on-the-fly
- Drag and drop tasks to create configs
- Scalable deployment without network degradation



ProLiant Essentials Rapid Deployment Pack



What's New in RDP 1.40?

- ProLiant DL380 Packaged Cluster Deployment
 - Imaging and scripting deployment jobs for Packaged Cluster
 - Support for Windows 2000 and Windows Server 2003
- Altiris Deployment Solution for Servers 5.6
- Role and scope-based security
 - Control who can perform tasks/jobs
 - Control who can manage groups of servers
 - Supports NT/AD authentication
- Switch Management Support
 - Windows 32 interface for discovering switches and attached clients
 - Ability to manually set the VLAN for the ports on the switch
 - Sample event showing how to use the CLI
 - Support for Cisco, HP ProLiant Blade, and 3 COM switches

ProLiant Blade Cluster Deployment Using RDP 1.40



- What's new for cluster deployment in RDP v1.40?
 - One unattended answer file per OS
 - One configuration file per cluster
 - One job for all blade clusters
 - More robust error checking and status reporting

- Advantages of using RDP to deploy clusters
 - Unattended installation of operating system
 - Unattended configuration of shared partitions and MSCS
 - Consistent cluster configurations
 - Deploy multiple clusters simultaneously

Blade Deployment Features for RDP 1.40

- Jobs are provided for Server Deployment...
 - Windows 2000 scripted installs
 - Scripted Windows 2000 BL20p G2 for SAN
 - Scripted Windows 2000 BL40p for SAN
 - Windows Server 2003 scripted installs
 - Scripted Windows 2003 BL20p G2 for SAN
 - Scripted Windows 2003 BL40p for SAN
- ...and for Cluster Deployment
 - BL40p/20p G2 Windows Partitions and MSCS Deployment
- These jobs can be used for all blade cluster deployments (1 OS install + the MSCS job)

Blade Deployment Features for RDP 1.40

Also provided for blade deployment RDP 1.40

- Cluster configuration file template
- Unattended text file template
- Computer import file template
- Documentation for the entire process
 - Modifying and using the configuration files
 - Modifying and using the deployment jobs
 - Configuring the private network
 - Configuring the storage system

Using the Deployment Jobs with RDP 1.40

- Complete pre-deployment configuration tasks
 - Prepare deployment environment
 - Edit cluster configuration file
 - Edit OS answer files
 - Import nodes into deployment console
 - Modify the MSCS deployment job for the domain
- Launch the Phase 1 deployment job
- Complete the manual configuration tasks for Phase 2
- Launch the Phase 3 deployment job
- Validate cluster deployment

Pre-Deployment Steps - Preparing the Environment

- Download and run setup for blade deployment files
 - Download SP24893 (see link at end of presentation)
 - In `.\extras\bladecluster` on RDP 1.40 CD
- Cable the network for all blades to be deployed
 - Connect PXE-capable NIC to deployment server
- Create and populate *san* directory on deployment server with most current SAN deliverables
 - FCA2101 (`.\san\kgpsa`)
 - Embedded FC mezzanine in BL20p G2 (`.\san\qla`)
 - Secure Path Server (`.\san\sps.40b`)
 - FCUtil (`.\san\kgpsa` and `.\san\qla`)

Cluster Configuration File

- A cluster configuration file is required for each cluster deployed
- Name the configuration file after the cluster
 - Ex: *mycluster.ini*
- Edit the `cluster` section with cluster specific variables
 - Removes admin credentials from the deployment job

```
;===Cluster Section=====
[Cluster]
account=administrator
password=password
domain=rdp
ipaddr=30.31.0.50
subnet=255.255.0.0
```

Cluster Configuration File

- Edit the `Network` section with cluster network information
 - Used to specify network for cluster
 - Allows a choice for NIC to host private network
- The private network addresses specified are not used

```
;===Network Section=====
;PRIVATE0=1.1.1.1
;PRIVATE1=1.1.1.2
;ClusPublicNet=Local Area Connection
;ClusPrivateNet=Local Area Connection 2
```

Pre-Deployment Steps – Unattended OS Answer File

- Only one Windows unattended answer file is need for all cluster nodes (*w2kclus.txt* and *wnetclus.txt*)
- Verify all required values

```
[GuiUnattended]
  AdminPassword=*

[UserData]
  ProductKey=* (needed if not using a Microsoft Select distribution)
  FullName=*
  OrgName=*

[Identification]
  DomainAdmin=*
  DomainAdminPassword=*
  JoinDomain=*

[Components]
  Cluster=On (only needed for Windows 2000 AS deployments)
```

Pre-Deployment Steps – Importing Computers

- Servers must be in the deployment console before execution of jobs
- 3 methods for computer import available
 - Import file
 - Manual import via GUI
 - Configure node details after automatic discovery via PXE

Pre-Deployment Steps – Method 1: Import File


- Deployment server import mechanism using a formatted computer import file
- Required values
 - Server name/console name
 - Server serial number (or MAC address of PXE NIC)
 - Cluster name (computer group name)
- Don't forget the commas!

Computer Import File Sample Text

```
PNODE1,,D129FRW1K361,,PNODE1,1,CLUSDEMO,,1,,,,,,,,,,,,,HASE,hp,,,,  
,,,,,,,,,PCLUS,,,,,15.15.15.1,255.0.0.0
```


Pre-Deployment Steps – Method 2: Import GUI

- Use *New Computers GUI* to add each blade to console
- For each node, provide
 - Server name/console name
 - Server serial number (or MAC address of PXE NIC)
- Once all nodes are in the Altiris console:
 - Create a group with the same name as the cluster
 - Move the renamed cluster nodes into the new group



Pre-Deployment Steps – Method 3: After PXE Boot

- Once a server boots using PXE, it shows up in the console
- Once all nodes are in the Altiris console:
 - Create a group with the same name as the cluster
 - Rename each machine in the console
- Move the renamed cluster nodes into the new group



Launching the Deployment Jobs – Phase 1

- Complete all pre-deployment tasks
- Drag-and-drop the appropriate OS deployment job onto the cluster group
- After the manual configuration tasks are complete, drag-and-drop the storage and cluster deployment job onto the cluster group

Manual Task: Configuring the Networks – Phase 2



- Manual configuration of the network cards
 - Configure the cluster private interconnect
- ProLiant BL20p G2 has a total of 4 network connections
- ProLiant BL40p has a total of 6 network connections

Manual Tasks: Storage Configuration – Phase 2

- Connect the storage box to the SAN switches
- Configure the SAN appliance
- Create the zones on the SAN for the cluster
- Create the logical units
- Add the logical units to the zone
- Present the drives to the Blade servers

Launching the Deployment Jobs – Phase 3




- Complete the manual configuration tasks in Phase 2
- Drag-and-drop the **BL40p/20p G2 Windows Partitions and MSCS Deployment** job onto the cluster group

Validation of Cluster Deployment

- Initiate a failover of all of the cluster groups
 - In a Microsoft Windows Server 2003, Enterprise Edition n-node configuration, failover to all nodes
- Verify all resources come online
- Initiate a failback of all cluster groups
- Verify all resources come online

ProLiant HA Web Site

<http://www.hp.com/servers/proliant/highavailability>



ProLiant high availability

» **ProLiant home**

- » adaptive infrastructure
- » server management
- » new products
- » retired products

purchase assistance

- » **how to buy**
- » financial services
- » locate a reseller

» **ProLiant high availability home**


related resources

- » high availability services
- » white papers
- » full-line web-based training
- » ProLiant clusters feedback

collateral/multimedia

- » fibre channel packaged clusters


» special promotions



**cost-effective availability
for the volume enterprise**


» [learn more](#)

» products	» solutions	» information center
» Microsoft Clusters	» ProLiant fibre channel Cluster solutions and bundles	» Cluster basics
» Novell Clusters	» Mail, Messaging and Collaboration	» Planning
» Linux Clusters	» Databases	» Deployment and installation
» Parallel Database Clusters for Oracle RAC	» ERP	» Management and maintenance
» DL380 Packaged Cluster		» Migration
		» Cluster Aware applications



Clustering Spotlight

Get the Latest Information on News, Success Stories, and more



Microsoft **Windows Server 2003**

"Find out more about HP ProLiant Cluster support for Microsoft Windows Server 2003, Enterprise Edition. Everything you need to know for migrating your existing clusters or installing new clusters".

» [details](#)

Other Information

- Blade Cluster Deployment Files
 - SP24893 available at HP High Availability Website
<http://h18004.www1.hp.com/solutions/enterprise/highavailability/whitepapers/proliant-bl.html>
 - Or on the RDP-WE 1.40 CD
[.\extras\bladecluster](#)
- HP High Availability Solutions
 - www.hp.com/servers/proliant/highavailability
- ProLiant Essentials Rapid Deployment Pack
 - www.hp.com/servers/rdp
- Session 2186 - ProLiant Clusters: Deploying a Well Managed ProLiant Cluster for Windows
- Session 2187 – Deploying Blade Clusters

Questions?





Interex, Encompass and HP bring you a powerful new HP World.



BL20p G2 Interconnect Options

		Provides network cable reduction	Supports FC pass-through for BL20p G2	Supported NIC speeds from server blade
Patch Panels	G2 Patch Panel	No	Yes	10/100/1000
Switches	GbE2 switch (3Q03 Availability)	Yes	Yes with GbE2 switch FC Option	10/100/1000