

HP ProLiant Cluster F100 Installation Guide



May 2004 (First Edition)
Part Number 364768-001

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

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.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

HP ProLiant Cluster F100 Installation Guide

May 2004 (First Edition)

Part Number 364768-001

Audience Assumptions

This document is for the person who installs, administers, and troubleshoots servers and storage systems. HP assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Contents

Overview of the HP ProLiant Cluster F100	5
F100 Cluster Components.....	5
Guidelines for Multiple Clusters.....	6
HP ProLiant Servers.....	7
HP StorageWorks Storage Systems	7
HP StorageWorks Controllers.....	8
HP StorageWorks SAN Switches	8
Host Bus Adapters	8
Optical Transceivers	9
Cluster Interconnect.....	9
Client Network	9
Private or Public Interconnect.....	10
Interconnect Adapters.....	10
Redundant Interconnects	10
Cables	11
Microsoft Software	12
HP SmartStart CD.....	13
Array Configuration Utility	13
ProLiant Support Packs	13
HP Modular Smart Array 1000 Support Software CD.....	14
Systems Insight Manager.....	14
Resources for Application Installation.....	14
Setting Up the HP ProLiant Cluster F100	17
Preinstallation Overview.....	17
Preinstallation Guidelines	18
Installing the Hardware.....	20
Setting Up the Nodes.....	20
Installing Host Bus Adapters	21
Installing the Cluster Interconnect.....	21
Setting Up the Storage System	22
Setting Up a Dedicated Interconnect.....	22
Setting Up a Public Interconnect	23
Redundant Interconnect.....	24
Multiple Cluster Setup.....	24
Installing the Software	24
SmartStart Installation.....	25

4 HP ProLiant Cluster F100 Installation Guide

Cluster-Specific SmartStart Installation	26
SmartStart Installation Steps.....	26
Validating the Cluster	30
Technical Support	31
Before You Contact HP	31
HP Contact Information.....	31
Acronyms and Abbreviations	33
Glossary	37
Index	45

Overview of the HP ProLiant Cluster F100

In This Section

F100 Cluster Components	5
Guidelines for Multiple Clusters	6
HP ProLiant Servers	7
HP StorageWorks Storage Systems.....	7
HP StorageWorks Controllers	8
HP StorageWorks SAN Switches.....	8
Host Bus Adapters	8
Optical Transceivers	9
Cluster Interconnect.....	9
Microsoft Software	12
HP SmartStart CD	13
HP Modular Smart Array 1000 Support Software CD	14
Systems Insight Manager.....	14
Resources for Application Installation	14

F100 Cluster Components

The F100 cluster includes the following hardware solution components:

- Two ProLiant servers ("HP ProLiant Servers" on page [7](#))
- One or more storage systems ("HP StorageWorks Storage Systems" on page [7](#))
 - MSA1000
 - RA4100
- One storage system controller ("HP StorageWorks Controllers" on page [8](#)) per storage system
 - MSA1000 Controller
 - RA4000 Controller
- One supported switch ("HP StorageWorks SAN Switches" on page [8](#)) or hub

- One supported HBA ("Host Bus Adapters" on page [8](#), "host bus adapter" on page [40](#)) per server
- NICs
- Optical transceivers (on page [9](#))
- Cables (on page [11](#)):
 - Multi-mode Fibre Channel cable
 - Ethernet crossover cable
 - Network (LAN) cable
- Software solution components:
 - Microsoft® operating system ("Microsoft Software" on page [12](#))
 - HP SmartStart CD (on page [13](#)) (included in the Server Setup and Management Pack)
 - HP Modular Smart Array 1000 Support Software CD (on page [14](#)) (if using the MSA1000 storage system)
 - Systems Insight Manager (on page [14](#)) (optional)

IMPORTANT: Refer to the MSCS support matrix for the supported hardware, software, cluster-supported servers, and firmware version levels the cluster requires. To obtain the matrix, access the High Availability website (<http://www.hp.com/servers/proliant/highavailability>), click **Cluster configuration support matrices** in the Related Information pane. Select the operating system, then the storage system for your configuration to display the matrix.

Guidelines for Multiple Clusters

Multiple-cluster configuration is available for the MSA1000 storage systems. The F100 for MSA1000 can have the following:

- Two-node clusters for Microsoft® Windows® 2000 Advanced Server
- Eight-node clusters for Microsoft® Windows® Server 2003, Enterprise Edition
- Five MSA1000 storage systems per cluster

- Up to five two-node clusters sharing a single MSA1000
- A maximum of 20 total nodes in unique clusters, not to exceed five unique clusters if using two-node clusters

The F100 for MSA1000 has the following guidelines for multiple-cluster setup:

- Each two-node cluster must use the same model of ProLiant server.
- Each server in a cluster must use the same operating system.
- Each two-node cluster can use a different operating system from the other clusters.
- The MSA1000 must have the appropriate firmware. Use the MSAFlash utility to upgrade the firmware.
- The ACU and the HBA drivers must be installed or updated to the supported versions. Use the HP Modular Smart Array 1000 Support Software CD to install or upgrade these items.
- When cabling multiple clusters, it is important to cable the first HBA of each cluster node to the first switch.

HP ProLiant Servers

HP industry-standard servers are a primary component of all models of ProLiant clusters. For more detailed information, refer to the server documentation provided with the ProLiant server.

To obtain a comprehensive list of cluster-supported servers, refer to the High Availability website (<http://www.hp.com/servers/proliant/highavailability>).

HP StorageWorks Storage Systems

The F100 must have at least one storage device set up as external shared storage. Consult the F100 ProLiant cluster website to determine the maximum supported cluster configuration.

The supported storage systems for the F100 are:

- MSA1000
- RA4100

For more information on the shared storage system, refer to the shared storage documentation.

HP StorageWorks Controllers

Each storage system is shipped with one HP StorageWorks Controller installed. In an F100 cluster, each controller is connected to the servers through a single switch or Fibre Channel storage hub.

The supported storage system controllers for the F100 are:

- MSA1000 Controller
- RA4000 Controller

For more information, refer to the controller documentation.

For more information about shared storage clustering, refer to the Microsoft® clustering documentation.

HP StorageWorks SAN Switches

A StorageWorks SAN Switch is an interconnect component built on Fibre Channel technology and has 2 Gb/s connectivity for an entry-level SAN. It is fully non-blocking and can provide up to 32 Gb/s switching capacity for uncongested sustained 2 Gb/s full duplex throughput.

For more information, refer to the appropriate SAN switch documentation.

Host Bus Adapters

The HBA is the interface between the servers and the F100 storage system. At least two HBAs, one for each cluster node, are required in the ProLiant Cluster F100 configuration.

For more information, refer to the HBA documentation.

Optical Transceivers

The F100 cluster uses optical transceivers, which convert data that is transmitted and received over a fibre-optic cable. The F100 cluster uses either SFP for the MSA1000 or GBIC-SW for the RA4100.

The transceivers hot plug into switches, Fibre Channel storage hubs, array controllers, and HBAs. The SFP transceiver provides 200 Mb/s performance and the GBIC-SW provides 100-Mb/s performance. The transceivers support distances up to 500 meters using multimode fiber-optic cable.

For more information, refer to the transceiver documentation.

Cluster Interconnect

The cluster interconnect is a data path over which nodes of a cluster communicate. This type of communication is termed intracluster communication. At a minimum, the interconnect consists of two network adapters (one in each server) and a crossover cable connecting the adapters.

The cluster nodes use the interconnect data path to:

- Communicate individual resource and overall cluster status
- Send and receive heartbeat signals
- Update modified registry information

IMPORTANT: TCP/IP must be used as the cluster communication protocol. When configuring the interconnects, be sure to enable TCP/IP.

Client Network

Every client/server application requires a LAN over which client machines and servers communicate. The components of the LAN are no different than with a stand-alone server configuration.

Because clients desiring the full advantage of the cluster will now connect to the cluster rather than to a specific server, configuring client connections will differ from those for a stand-alone server. Clients will connect to virtual servers, which are cluster groups that contain their own IP addresses.

Private or Public Interconnect

There are two types of interconnect paths:

- Private interconnect
- Public interconnect

For more information on recommended interconnect strategies, refer to the white paper, *Best Practices Checklist—Increasing Network Fault Tolerance in a Microsoft® Windows® Server 2003, Enterprise Edition High Availability Server Cluster*, available from the ProLiant High Availability website (<http://www.hp.com/servers/proliant/highavailability>).

Interconnect Adapters

Ethernet adapters and switches are supported as interconnects in ProLiant clusters. Either a 10-Mb/s, 100-Mb/s, or 1000-Mb/s Ethernet adapter can be used.

Ethernet adapters can be connected together using an Ethernet crossover cable or a private Ethernet hub. Both of these options provide a dedicated interconnect.

Implementing a direct Ethernet connection minimizes potential single points of failure.

NOTE: An Ethernet crossover cable is provided in the *HP ProLiant Cluster Starter* kit.

Redundant Interconnects

To reduce potential disruptions of intracluster communication, use a redundant path over which communication can continue if the primary path is disrupted.

HP recommends configuring the client LAN as a backup path for intracluster communication. This provides a secondary path for the cluster heartbeat in case the dedicated primary path for intracluster communications fails. This is configured when installing the cluster software, or it can be added later using the MSCS Cluster Administrator.

HP offers a feature that configures two HP Ethernet adapters (or two ports on a single adapter) so that one is a hot backup for the other. There are two ways to achieve this configuration, called NIC Teaming, and the method you choose depends on the hardware. One way is through the use of the Redundant NIC Utility available on all HP 10/100/1000 Fast Ethernet products. The other option is through the use of the Network Fault Tolerance feature designed to operate with the HP 10/100/1000 Intel® silicon-based NICs.

For more information on recommended interconnect strategies, refer to the white paper, *Best Practices Checklist—Increasing Network Fault Tolerance in a Microsoft® Windows® Server 2003, Enterprise Edition High Availability Server Cluster*, available from the ProLiant High Availability website (<http://www.hp.com/servers/proliant/highavailability>).

Cables

Three general categories of cables are used for the F100 cluster:

- Server to storage
- Cluster interconnect
- Network interconnect

Server to Storage

Shortwave (multi-mode) fiber optic cables are used to connect the servers, switches or Fibre Channel storage hubs, and storage systems in a Fibre Channel configuration.

Cluster Interconnect

When Ethernet cluster interconnect cables are used with Ethernet NICs to implement the interconnect, there are three options:

- Dedicated interconnect using an Ethernet crossover cable—An Ethernet crossover cable (supplied in the *HP ProLiant Cluster Starter* kit) can be used to connect the NICs directly to create a dedicated interconnect. This option is applies to only two-node clusters.
- Dedicated interconnect using standard Ethernet cables and a private Ethernet hub—Standard Ethernet cables can be used to connect the NICs through a private Ethernet hub to create another type of dedicated interconnect. Do not use an Ethernet crossover cable when using an Ethernet hub because the hub performs the crossover function.
- Shared interconnect using standard Ethernet cables and a public hub—Standard Ethernet cables can also be used to connect the NICs to a public network to create a non-dedicated interconnect.

Network Interconnect

Standard Ethernet cables are used to provide this type of connection.

Microsoft Software

Microsoft® Windows® Server 2003, Enterprise Edition and Microsoft® Windows® 2000 Advanced Server are the operating systems for the HP ProLiant Cluster F100.

The Microsoft® clustering software provides the underlying technology to:

- Send and receive heartbeat signals between the cluster nodes.
- Monitor the state of each cluster node.
- Initiate failover and failback events.

Microsoft® Cluster Administrator enables you to:

- Define and modify cluster groups.
- Manually control the cluster.
- View the current state of the cluster.

For more information, refer to the Microsoft® documentation.

HP SmartStart CD

SmartStart is a software package that provides a streamlined process for the installation of operating systems, and provides key system software, such as drivers, utilities, diagnostic tools, and ROM updates. SmartStart also provides automated methods for configuring server settings.

HP SmartStart is located on the HP SmartStart CD included in the *ProLiant Essentials Foundation Pack* shipped with ProLiant servers. HP recommends using SmartStart to configure the F100 cluster nodes. SmartStart uses a step-by-step process to configure the operating system and load the system software. For information concerning SmartStart, refer to the *ProLiant Essentials Foundation Pack*.

For more information about using SmartStart to install the F100 cluster nodes, refer to SmartStart installation (on page [25](#)) in this guide.

Array Configuration Utility

The ACU is a Web-based configuration utility available for some servers that makes it easy to configure and expand the disk drive arrays. The ACU is available on the SmartStart CD or the HP Modular SAN Array 1000 Support Software CD.

ProLiant Support Packs

PSPs represent operating system specific bundles of ProLiant optimized drivers, utilities, and management agents. Refer to the PSP website (<http://h18000.www1.hp.com/products/servers/management/psp.html>).

HP Modular Smart Array 1000 Support Software CD

The HP Modular Smart Array 1000 Support Software CD contains drivers and utilities required for the MSA1000 storage system. The CD contains items, such as the Fibre Channel HBA drivers, which are required for the HBAs to communicate with the MSA1000 storage system. The CD also contains the ACU, which enables you to view, set up, and configure HP array controllers and storage systems.

For more information on the HP Modular Smart Array 1000 Support Software CD, refer to the *HP Modular Smart Array 1000 Setup and Management* kit.

Systems Insight Manager

Systems Insight Manager is a Web-based application that enables system administrators to accomplish normal administrative tasks from any remote location, using a Web browser. Systems Insight Manager provides device-management capabilities that consolidate and integrate management data from HP and third-party devices.

For additional information, refer to the Management CD in the *ProLiant Essentials Foundation Pack*.

Resources for Application Installation

The client/server software applications are among the key components of any cluster. HP is working with its key software partners to ensure that cluster-aware applications are available and that the applications work seamlessly on HP ProLiant clusters.

HP provides a number of checklists and white papers to assist you with installing these applications in an HP ProLiant cluster environment.

To download current versions of these technical documents, refer to the High Availability website (<http://www.hp.com/servers/proliant/highavailability>).

IMPORTANT: The software applications might need to be updated to take full advantage of clustering. Contact the software vendors to verify whether their software supports MSCS and to ask whether any patches or updates are available for MSCS operation.

Setting Up the HP ProLiant Cluster F100

In This Section

Preinstallation Overview	17
Preinstallation Guidelines.....	18
Installing the Hardware	20
Installing the Software.....	24
SmartStart Installation	25
Validating the Cluster	30

Preinstallation Overview

The ProLiant Clusters F100 for MSA1000 and F100 for the RA4100 are combinations of several individually available products. Have the following documents available as you set up the cluster:

- Documentation for the clustered ProLiant servers
- Shared external storage documentation
- HBA documentation
- Installation guide for the NIC
- Installation guide for the switch or Fibre Channel storage hub
- Documentation received with the operating system
- SmartStart poster
- Systems Insight Manager Installation Poster
- Microsoft® Windows® Server 2003, Enterprise Edition or Microsoft® Windows® 2000 Advanced Server clustering documentation
- Zoning documentation (if installing more than one two-node cluster to a switch)

For more information, refer to the *HP Zoning User's Guide* Version 3.0 at the HP High Availability website (<http://www.hp.com/servers/proliant/highavailability>).

The installation and setup of the ProLiant cluster is described in the following sections:

- Preinstallation guidelines (on page [18](#))
- Installing the hardware (on page [20](#)), including:
 - Cluster nodes
 - Storage system
 - Cluster interconnect
- Installing the software (on page [24](#)), including:
 - HP SmartStart
 - Microsoft® Windows® Server 2003, Enterprise Edition or Microsoft® Windows® 2000 Advanced Server
 - Systems Insight Manager (optional)
- Validating the cluster (on page [30](#))

IMPORTANT: Refer to the MSCS support matrix for the supported hardware, software, cluster-supported servers, and firmware version levels the cluster requires. To obtain the matrix, access the High Availability website (<http://www.hp.com/servers/proliant/highavailability>), click **Cluster configuration support matrices** in the Related Information pane. Select the operating system, then the storage system for your configuration to display the matrix.

Preinstallation Guidelines

Before installing MSCS on a cluster node, write down the answers to the following questions.

- Are you forming or joining a cluster?
- What is the cluster name?

- What is the username, password, and domain for the domain account under which MSCS will run?
- What disks will you use for shared storage?
- Which shared disk will you use to store permanent cluster files?
- What are the adapter names and IP addresses of the network adapter cards you will use for client access to the cluster?
- What are the adapter names and IP addresses of the network adapter cards you will use for the dedicated interconnect between the cluster nodes?
- What is the IP address and subnet mask of the address you will use to administer the cluster?
- What are the slot numbers of the adapters to be managed by the cluster?

Installing clustering software requires several specific steps and guidelines that might not be necessary when installing software on a single server. Read and understand the following items before proceeding with any software installation:

- Be sure that you have sufficient software licensing rights to install the Microsoft® Windows® operating system and software applications on each server.
- Be sure the Fibre Channel storage hub or Fabric switch has AC power.
- Power up the storage system before the cluster nodes are powered up.
- Log on to the domain using an account that has administrative permissions on both cluster nodes. When installing MSCS, both cluster nodes must be in the same Microsoft® Windows® Server 2003, Enterprise Edition or Microsoft® Windows® 2000 Advanced Server domain. The cluster nodes can be members of an existing Microsoft® Windows® Server 2003, Enterprise Edition or Microsoft® Windows® 2000 Advanced Server domain.
- Configure the logical drives in the storage system using the ACU.
- When the ACU runs on the first cluster node, configure the shared drives in the storage system. If the utility is used to configure internal storage on the second cluster node, it will display information on the shared drives that was entered when the ACU was run on the first node.

- Only NTFS and basic disks are supported on shared drives.
- MSCS software requires drive letters to remain constant throughout the life of the cluster. Therefore, you must assign permanent drive letters to the shared drives.
- Microsoft® Windows® Server 2003, Enterprise Edition and Microsoft® Windows® 2000 Advanced Server makes dynamic drive letter assignments (when drives are added or removed, or when the boot order of drive controllers is changed), but Disk Management enables you to make permanent drive letter assignments.
- Cluster nodes can be members of only one cluster.
- When you set up the cluster, select TCP/IP as the network protocol. MSCS requires the TCP/IP protocol. The cluster interconnect must be on its own subnet. The IP addresses of the interconnects must be static, not dynamically assigned by DHCP.

Installing the Hardware

The following installation steps detail a new installation and setup of an F100 for MSA1000 or F100 for the RA4100.

Setting Up the Nodes

Each cluster node requires at least two network adapters: one to connect to a public network and one to connect to a private network. Physically preparing the nodes (servers) for use in a cluster is not very different from preparing them for individual use. The primary difference will be in setting up the shared storage.

1. Install all necessary adapter cards and insert all internal hard drives.
2. Attach network cables and plug in Fibre Channel cables.
3. Set up one node completely, then set up the second node.

IMPORTANT: Do not load any software on either cluster node until all the hardware has been installed in both cluster nodes.

NOTE: HP recommends that ASR be left at the default values for clustered servers.

Follow the installation instructions in the ProLiant server documentation to set up the hardware. To install the HBAs and any NICs, follow the instructions in the next section.

IMPORTANT: Refer to the MSCS support matrix for the supported hardware, software, cluster-supported servers, and firmware version levels the cluster requires. To obtain the matrix, access the High Availability website (<http://www.hp.com/servers/proliant/highavailability>), click **Cluster configuration support matrices** in the Related Information pane. Select the operating system, then the storage system for your configuration to display the matrix.

Installing Host Bus Adapters

The HBAs that connect the two servers to the storage through Fabric switch or Fibre Channel storage hub is installed in each server like any other PCI card.

Follow the installation instructions in the HBA documentation and the ProLiant server documentation to install the HBAs in the servers.

Installing the Cluster Interconnect

There are a number of methods to physically set up an interconnect. Refer to the cluster interconnect section for a description of the types of interconnect strategies.

If you are using a dedicated interconnect, install an Ethernet interconnect adapter card in each cluster node.

NOTE: HP recommends a dedicated NIC for the LAN and a dedicated NIC for the interconnect of each cluster node.

For specific instructions on how to install an adapter card, refer to the documentation for the interconnect card you are installing or the ProLiant server you are using. The cabling of interconnects is outlined later in this chapter.

Setting Up the Storage System

Follow the instructions in the shared external storage documentation to set up the storage system, the supported StorageWorks switches or hubs, the controller, and the Fibre Channel cables.

HP shared external storage documentation explains how to install these devices for a single server. Because clustering requires shared storage, install these devices for two servers. This will require running an extra Fibre Channel cable from the switch or Fibre Channel storage hub to the second server.

Powering Up

Before applying power to the storage system, be sure that all components are installed and connected to the Fabric switch or Fibre Channel storage hub.

Power up the cluster in the following order:

1. Fabric switches or Fibre Channel storage hubs
2. Storage systems
3. Servers

Configuring Shared Storage

The ACU sets up the hardware aspects of any drives attached to an array controller, including the drives in the shared storage systems. The ACU can initially configure the array controller, reconfigure the array controller, add additional disk drives to an existing configuration, and expand capacity. The ACU stores the drive configuration information on the drives themselves. Therefore, after you have configured the drives from one of the cluster nodes, it is not necessary to configure the drives from the other cluster node.

For detailed information about configuring the drives, refer to the section on the ACU in the shared external storage documentation.

Setting Up a Dedicated Interconnect

There are two ways to set up a dedicated interconnect:

- Ethernet direct connect using a crossover cable
- Ethernet direct connect using a switch or hub

Ethernet Direct Connect Using a Crossover Cable

An Ethernet crossover cable is included with the *HP ProLiant Cluster Starter* kit. This cable directly connects two NICs that have been designated as the dedicated interconnect. Connect one end of the cable to the NIC in node 1 and the other end of the cable to the NIC in node 2.

IMPORTANT: Connect the cable to the dedicated interconnect NICs and not to the Ethernet connections used for the network clients (the public LAN).

NOTE: The crossover cable will not work in conjunction with a network hub or switch.

Ethernet Direct Connect Using a Switch or Hub

An Ethernet hub or switch requires standard Ethernet cables. Ethernet crossover cables will not work with a switch or hub. To cable the server interconnect using an Ethernet switch or hub:

1. Connect the end of one of the Ethernet cables to the NIC in node 1.
2. Connect the other end of the cable to a port in the switch or hub.
3. Repeat steps 1 and 2 for the NIC in node 2.
4. Assign static IP addresses.

IMPORTANT: Connect the cable to the dedicated interconnect NICs and not to the Ethernet connections used for the network clients (the public LAN).

Setting Up a Public Interconnect

It is possible, but not recommended, to use a public network as the dedicated interconnect path. To set up a public Ethernet interconnect, connect the NICs, switch or hub, and cables as you would in a non-clustered environment. Then configure the NICs for both network clients and for the dedicated interconnect.

Microsoft® recommends using static IP addresses on the public LAN. Do not use DHCP-assigned addresses.

IMPORTANT: Using a public network as the dedicated interconnect path is not recommended because it represents a potential single point of failure for cluster communication.

Redundant Interconnect

MSCS enables you to configure any supported network card as a possible path for intracluster communication. If you are employing a dedicated interconnect, use MSCS to configure the LAN network cards to serve as a backup for the interconnect.

Multiple Cluster Setup

For a multiple cluster setup, be sure that:

- Each two-node cluster is using the same model of ProLiant server.
- The storage system has the latest supported firmware version.
- The same HBAs are being used.

NOTE: The MSA1000 storage system can be upgraded using the MSAFlash utility. To learn more about MSAFlash and how to use it, refer to the documentation provided within the download of the MSA1000 firmware upgrade package.

Installing the Software

The following sections describe the software installation steps for the:

- F100 for MSA1000 with Microsoft® Windows® Server 2003, Enterprise Edition
- F100 for MSA1000 with Microsoft® Windows® 2000 Advanced Server
- F100 for the RA4100 with Microsoft® Windows® Server 2003, Enterprise Edition
- F100 for the RA4100 with Microsoft® Windows® 2000 Advanced Server

Proceed with these steps after you have all equipment installed and the switches or hubs, storage system, and one server powered up.

IMPORTANT: Refer to the MSCS support matrix for the supported hardware, software, cluster-supported servers, and firmware version levels the cluster requires. To obtain the matrix, access the High Availability website (<http://www.hp.com/servers/proliant/highavailability>), click **Cluster configuration support matrices** in the Related Information pane. Select the operating system, then the storage system for your configuration to display the matrix.

You need the following during installation:

- HP SmartStart CD
- HP SmartStart poster
- HP Modular Smart Array 1000 Support Software CD (if using the MSA1000 storage system)
- One of the following operating systems:
 - Microsoft® Windows® Server 2003, Enterprise Edition software and documentation
 - Microsoft® Windows® 2000 Advanced Server software and documentation
- Microsoft® Service Packs

NOTE: Windows® 2000 Advanced Server requires the use of Service Pack 4 or later.

- Systems Insight Manager software and documentation (optional)

SmartStart Installation

IMPORTANT: Before the installation of Microsoft® Windows® Server 2003, Enterprise Edition or Microsoft® Windows® 2000 Advanced Server, upgrade the system ROM on each node with the latest system ROM update from the HP support website (<http://www.hp.com/support>).

Use the SmartStart procedure to configure the servers (nodes) in the ProLiant cluster. You will set up two nodes during this process. Proceed through all of the steps on each of the nodes, with noted exceptions.



CAUTION: Installation using SmartStart assumes that SmartStart is being installed on new servers. Any existing data on the boot drive of the server will be erased.

Cluster-Specific SmartStart Installation

The SmartStart Poster describes the typical procedure for configuring and installing software on a single server. The difference between running SmartStart on a stand-alone server and running SmartStart for a cluster are noted as follows:

- For the RA4100 using the ACU, you can configure the shared drives on both servers. For cluster configuration, configure the drives on the first server, then accept the same settings for the shared drives when given the option on the second server.
- When configuring drives through the ACU, create a logical drive with 510 MB of space to be used as the quorum disk.

IMPORTANT: Microsoft® recommends at least 500 MB for the cluster quorum drive. The extra space for the logical drive size specified in the ACU is to account for external disk size calculations used by the ACU. Specifying 510 MB will ensure that the size of the disk will be at least 500 MB of formatted drive space for use as the quorum drive. Refer to Microsoft® KB Article EN-280345 (<http://support.microsoft.com/default.aspx?scid=kb:EN-US:280345>) or the help documentation on the cluster node for more information on cluster disk sizes.

SmartStart Installation Steps

Use the following sections to install the operating system on nodes 1 and 2 using SmartStart.

Installing the Node 1 Operating System

IMPORTANT: Power down node 2 when setting up node 1.

1. Power up the hardware in the following order:
 - a. Power on the switch or Fibre Channel storage interconnect.

- b. Power on the shared storage. It might take up to two minutes for the storage system to completely power up.
 - c. Power on and boot node 1 with the SmartStart CD in the CD-ROM drive. The CD will automatically run.
2. Click **Launch Setup** and follow the on-screen instructions. SmartStart will discover all hardware.
3. On the **Hardware Configuration** screen, configure the boot drive array and click **Next**.
4. After the boot drive array completes configuration, select the appropriate operating system, and follow the SmartStart on-screen instructions and prompts.

After the operating system installation is complete, SmartStart will automatically install the HP support software.

5. Configure the TCP/IP settings for the public network connection.

If the network adapter can transmit at multiple speeds, then manually specify a speed and duplex mode. The speed for the network adapter should be hard set (manually set) to be the same on all nodes according to the card manufacturer's specification.
6. Configure the TCP/IP settings for the private network connection.

To eliminate possible private network cluster communication issues, refer to Microsoft® KB article EN-US258750 (<http://support.microsoft.com/default.aspx?scid=kb;EN-US;258750>) to properly set up the private network.
7. Join the Windows® domain and reboot when prompted.
8. After rebooting, log into the domain.
9. For an F100 installation, insert the HP Modular Smart Array Support Software CD. Click **Install Online Array Configuration Utility**. For all other installations, skip to step 11.

NOTE: The F100 for the RA4100 HBA drivers are installed as part of the HP support software installation in the previous steps.

10. After the ACU is installed, select the appropriate option to install the HBA drivers for the operating system.
11. Run the ACU to configure shared storage:

- a. From the desktop of node 1, run the **HP Array Configuration Utility**.
 - b. Configure the shared storage drives.

IMPORTANT: Microsoft® recommends at least 500 MB for the cluster quorum drive. The extra space for the logical drive size specified in the ACU is to account for external disk size calculations used by the ACU. Specifying 510 MB will ensure that the size of the disk will be at least 500 MB of formatted drive space for use as the quorum drive. Refer to Microsoft® KB Article EN-280345 (<http://support.microsoft.com/default.aspx?scid=kb;EN-US;280345>) or the help documentation on the cluster node for more information on cluster disk sizes.
 - c. Reboot node 1 to complete the proper discovery of all the disk drives.
 - d. Log into the node and wait for PnP to complete the discovery of new drives.
 - e. After the discovery is complete, select **Start, Programs, Administrative Tools, Computer Management**. Then select **Disk Management** to create volumes out of the logical drives.

NOTE: Do not upgrade the logical drives from basic to dynamic. MSCS does not support dynamic disks.
 - f. Assign drive letters and format the volumes as NTFS.
 - g. Close Disk Management.
12. Create the cluster on node 1 for Microsoft® Windows® Server 2003, Enterprise Edition. If creating a cluster in Microsoft® Windows® 2000 Advanced Server, use step 13.
 - a. Select **Start, Programs, Administrative Tools, Cluster Administrator**.
 - b. Select **Create New Cluster** from the Action dropdown box. Click **OK**.
 - c. Follow the on-screen instructions to create the cluster.
 - d. Select **Start, Settings, Control Panel, HP Management Agents**. In the list of Inactive Agents, select **Clustering Information** and click **Add** to move this agent to the list of active agents. Click **OK**.
 - e. Restart the agents when prompted.
 - f. Upgrade the node with the appropriate Windows® service pack.

13. Create the cluster on node 1 for Microsoft® Windows® 2000 Advanced Server.
 - a. Install the Cluster Service (MSCS) component in Add/Remove Programs. For more information on installing and configuring MSCS, refer to the Microsoft® Windows® 2000 Advanced Server documentation.
 - b. Follow the on-screen instructions to create the cluster.
 - c. Select **Start, Settings, Control Panel, HP Management Agents**. In the list of Inactive Agents, select **Clustering Information** and click **Add** to move this agent to the list of active agents. Click **OK**.
 - d. Restart the agents when prompted.
 - e. Upgrade the node with the appropriate Windows® service pack.

Installing the Node 2 Operating System

1. Power on and boot node 2 with the SmartStart CD in the CD-ROM drive.
2. Repeat steps 2 through 10 of Installing the Node 1 Operating System (on page [26](#)) to configure node 2.
3. When the installation is complete, restart node 2.
4. Join node 2 to the cluster for Microsoft® Windows® Server 2003, Enterprise Edition. If creating a cluster in Microsoft® Windows® 2000 Advanced Server, use step 5.
 - a. Log in to the domain.
 - b. From node 2, select **Start, Programs, Administrative Tools, Cluster Administrator**.
 - c. Select **Add nodes to cluster** from the Action dropdown box. Enter the name of the cluster and click **OK**.
 - d. Follow the on-screen instructions to create the cluster.
 - e. Select **Start, Settings, Control Panel, HP Management Agents**. In the list of Inactive Agents, select **Clustering Information** and click **Add** to move this agent to the list of active agents. Click **OK**.
 - f. Restart the agents when prompted.
 - g. Upgrade the node with the appropriate Windows® service pack.

For the F100 for MSA1000 with Microsoft® Windows® Server 2003, you can cluster additional nodes using the previous steps.

5. Create the cluster on node 2 for Microsoft® Windows® 2000 Advanced Server.
 - a. Install the Cluster Service (MSCS) component in Add/Remove Programs. For more information on installing and configuring MSCS, refer to the Microsoft® Windows® 2000 Advanced Server documentation.
 - b. Follow the on-screen instructions to create the cluster.
 - c. Select **Start, Settings, Control Panel, HP Management Agents**. In the list of Inactive Agents, select **Clustering Information** and click **Add** to move this agent to the list of active agents. Click **OK**.
 - d. Restart the agents when prompted.
 - e. Upgrade the node with the appropriate Windows® service pack.

The installation is complete.

Validating the Cluster

From the desktop of either node:

1. Select **Start, Programs, Administrative Tools, Cluster Administrator**, and connect to the cluster.
2. Right-click one of the cluster groups and select **Move Group**.
3. Verify that the group fails over and all resources come online.
4. Right-click the same cluster group and select **Move Group**.
5. Verify that the group fails over and all resources come online.
6. Repeat the Validating the Cluster steps, if desired, for each group.

Technical Support

In This Section

Before You Contact HP	31
HP Contact Information.....	31

Before You Contact HP

Be sure to have the following information available before you call HP:

- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

HP Contact Information

For the name of the nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- In other locations, refer to the HP website (<http://www.hp.com>).

For HP technical support:

- In North America, call the HP Technical Support Phone Center at 1-800-633-3600. This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored.

- Outside North America, call the nearest HP Technical Support Phone Center. For telephone numbers for worldwide Technical Support Centers, refer to the HP website (<http://www.hp.com>).

Acronyms and Abbreviations

ACU

Array Configuration Utility

ASR

Automatic Server Recovery

CPU

central processing unit

DHCP

Dynamic Host Configuration Protocol

DNS

Domain Name System

FQDN

Fully Qualified Domain Name

GBIC-SW

Gigabit Interface Converter-Shortwave (on page [39](#))

HBA

host bus adapter ("Host Bus Adapters" on page [8](#), on page [40](#))

I/O

input/output

LAN

local-area network

MSA1000

HP StorageWorks Modular Smart Array 1000 (on page [40](#))

MSA1000 Controller

HP StorageWorks Modular Smart Array 1000 Controller (refer to "HP StorageWorks Modular Smart Array 1000 Controller" on page [40](#))

MSCS

Microsoft® Cluster Server/Service ("Microsoft Cluster Server/Service" on page [41](#))

NIC

network interface controller

NTFS

NT File System (on page [41](#))

PnP

plug and play

POST

Power-On Self-Test

PSP

ProLiant Support Pack

RA4000 Controller

HP StorageWorks RAID Array 4000 Controller (on page [40](#))

RA4100

HP StorageWorks RAID Array 4100 (on page [40](#))

RAID

redundant array of inexpensive (or independent) disks

ROM

read-only memory

SAN

Storage Area Network

SCSI

small computer system interface

SFP

small form-factor pluggable (on page [43](#))

TCP/IP

Transmission Control Protocol/Internet Protocol (on page [44](#))

Glossary

adapter

A device that converts the protocol and hardware interface of one bus type into another without changing the function of the bus.

array

All the physical disk drives in a storage system that are known to and under the control of a controller pair.

availability

A measure of how well a computer system or cluster can continuously deliver services to its clients. Availability is typically expressed as a percentage, with 100 % being the best possible rating.

cluster

A group of systems that work collectively as a single system to provide fast, uninterrupted computing service. Clustering is a way to increase availability, processing capacity, and I/O bandwidth.

cluster group

A collection of interdependent resources that logically represents a clustered client/server function. This is a user-definable entity used by Microsoft® Cluster Server software.

controller

A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array.

dedicated interconnect

A type of interconnect that is used solely for intracluster (node-to-node) communication. Communication to and from network clients does not occur over this type of interconnect. *Also called* private interconnect.

disk group

A physical disk drive set or pool in which a virtual disk is created. A disk group can contain all the physical disk drives in a controller pair array or a subset of the array.

driver

A hardware device or a program that controls or regulates another device. For example, a device driver is a driver developed for a specific device that enables a computer to operate with that device, such as an HBA or a disk drive.

Ethernet

A standard network protocol that operates mostly on a physical level, using network interface cards and cabling to transmit data between computers. Transfer rates are normally 1,000 or 10,000 Mb/s.

fabric

The multiple Fibre Channel switches interconnected and using Fibre Channel methodology for linking nodes and routing frames in a Fibre Channel network.

failback (cluster)

1. The process that takes place when a previously failed controller is repaired or replaced and reassumes the workload from a companion controller.
2. The process that takes place when the operation of a previously failed cluster group moves from one cluster node back to its primary node.

failover (cluster)

1. The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.
2. The process that takes place when the operation of a cluster group moves from one cluster node to another node in the same cluster.

fault tolerance

The ability of a system or component to continue normal operation when a fault (or failure) is encountered. Tolerance is achieved primarily by designing redundant elements into the system.

Fibre Channel

An IEEE standard for providing high-speed data transfer among workstations, servers, mainframes, supercomputers, desktop computers, storage devices, and display devices.

Gigabit Interface Converter-Shortwave

A device that converts electrical signals to optical signals at the point where the fiber cables connect to the Fibre Channel elements, such as controllers or adapters.

heartbeat

A signal transmitted between cluster nodes to indicate whether the nodes are operating.

high availability

A term used to identify a computer system that can continuously deliver services to its clients 99.9 % of the time (no more than 8.5 hours of downtime per year).

host

The primary or controlling computer in a system of computers connected by communication links.

host bus adapter

A card used to connect a peripheral device to a host server.

HP StorageWorks Modular Smart Array 1000

A storage device including disk drives and one or more resident array controllers.

HP StorageWorks Modular Smart Array 1000 Controller

A hardware device that facilitates communication between a host and one or more devices organized on an MSA1000 storage system.

HP StorageWorks RAID Array 4000 Controller

A hardware device that facilitates communications between a host and one or more devices organized on an RA4100 storage system.

HP StorageWorks RAID Array 4100

A storage device including disk drives and one or more resident array controllers.

input/output

A term that pertains to input and output functions.

interconnect

A physical connection between cluster nodes that transmits intracluster communication.

intracluster communication

The type of communication in which the cluster interconnect is a data path over which nodes of a cluster communicate. At a minimum, the interconnect consists of two network adapters (one in each server) and a cable connecting the adapters.

IP address

Internet Protocol address. An address assigned to a network interface card, which computer entities use to locate and communicate with each other. IP addresses can be statically or dynamically assigned.

logical unit

Commonly called a LUN (which is the acronym for logical unit number). A physical or virtual device addressable through a target ID number. Logical units use the target bus connection to communicate on the SCSI bus. The host sees a virtual disk as a logical unit.

logical unit number

1. A value that identifies a specific logical unit belonging to a SCSI target ID number. LUN is commonly used in reference to a logical unit.
2. A number associated with a physical device unit during the I/O operations of a task. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.

Microsoft Cluster Server/Service

The software needed for clustering servers.

MSAFlash utility

An HP utility that can be used to upgrade the firmware of the MSA1000, environmental monitoring unit, and the MSA Fabric Switch 6.

network interface controller

A board that enables a computer to be connected to a network and that works with the network operating system to control the flow of information over the network.

node

An individual server in a cluster.

NT File System

A file organization system by which data is stored and accessed in a Windows® operating system.

partition

A logical division of a container, represented to the host as a logical unit.

port

1. In general terms, a logical channel in a communication system.
2. The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus.

Power-On Self-Test

A set of operations executed every time a system is turned on that verifies components are present and operating.

private interconnect

A type of interconnect that is used solely for intracluster (node-to-node) communication. Communication to and from network clients does not occur over this type of interconnect. *Also known as* dedicated interconnect.

public interconnect

A type of interconnect that enables communication between the cluster nodes and shares the data path with communication between the cluster and its network clients.

quorum disk

A device managed by the Microsoft® cluster software that provides a means for persistent storage of the cluster configuration information required for failover and failback events as well as for arbitrating ownership of cluster resources.

redundancy

The provision of multiple, interchangeable components to perform a single function to cope with failures and errors. A RAID set is considered to be redundant when user data is recorded directly to one member and all of the other members include associated parity information.

Redundant Array of Inexpensive Disks

A method of using hard disk drives in an array to provide data redundancy to increase system reliability and performance.

reliability

The continuous integrity of a system (server, storage, network, or cluster).

resource

A software or hardware entity on which a client/server application or service is dependent. As it pertains to Microsoft® Cluster Server, a cluster resource must have the ability to be managed by the cluster and must reside on one of the cluster nodes. A resource can be a member of only one group.

shared storage clustering

The cluster architecture in which clustered servers share access to a common set of hard drives. Microsoft® Cluster Software is based on shared storage clustering and requires all clustered (shared) data to be stored in an external storage system.

Small Computer System Interface

A standard parallel interface for rapid data transmission.

Small Form-Factor Pluggable

A device that converts electrical signals to optical signals at the point where the fiber cables connect to the Fibre Channel elements, such as controllers or adapters.

system

A complete computer system capable of operating independently.

Transmission Control Protocol/Internet Protocol

A suite of communication protocols developed to enable communication between different types of computers and networks.

Index

A

ACU (Array Configuration Utility) 6, 13, 18, 22, 26, 33
ACU, installation 26
additional information 31
ASR (Automatic Server Recovery) 20, 33
authorized reseller 31

C

cables, overview 11
client network, features 9
cluster components 5
cluster interconnect, features 9
cluster interconnect, installing 21
cluster interconnect, overview 9
cluster, installing 26
cluster, installing node 1 26
cluster, installing node 2 29
cluster, validating 30
clustering software guidelines 18
communication protocol, TCP/IP 9
contacting HP 31
controllers, overview 8

D

dedicated interconnect 10, 11, 22
DHCP (Dynamic Host Configuration Protocol) 18, 23, 33

E

Ethernet adapters 10
Ethernet connections 22
Ethernet crossover cable 11, 23
Ethernet direct connect 22

F

F100, overview 5
F200, overview 5
failback 12, 38, 42
failover 6, 12, 38, 42
firmware, updating 6

G

GBIC-SW (Gigabit Interface Converter-Shortwave) 9, 39

H

hardware requirements, F100 5
hardware supported 5
HBA, installation 21
HBA, overview 8
high availability website 7
HP ProLiant Servers 7
HP StorageWorks controllers 8
HP StorageWorks SAN switches 8
HP StorageWorks storage systems 7
HP Technical Support 31

I

installing hardware 20
installing software 24
interconnect, cluster 9, 21
intracluster communication 9

L

LAN 9, 10, 21, 23, 34
LAN, features 9

M

Microsoft software 12
MSA1000 (Modular Smart Array 1000) 6, 7
MSA1000 Support Software CD 14

MSAFlash utility 24, 41
MSCS (Microsoft Cluster Server/Service) 18,
24, 34
MSCS support matrix 24
multiple cluster guidelines 6
multiple cluster setup 24

N

network protocols, TCP/IP 18
network, local area 9
networking, cluster interconnect 9
networking, TCP/IP protocol 9
NIC (network interface controller) 11, 21, 23,
34
node setup 20
nodes, maximum 6

O

optical transceivers, overview 9
overview, F100 5
overview, F200 5

P

phone numbers 31
powering up 22
powerup sequence 22
preinstallation, guidelines 18
preinstallation, overview 17
private interconnect 10, 37, 42
ProLiant server, overview 7
ProLiant Support Packs 13
public interconnect 10, 23, 42

Q

quorum drive, recommended size 26

R

RA4000 Controller (RAID Array 4000
Controller) 8, 35
RA4100 (RAID Array 4100) 7, 35

required information 31
resources 14

S

SAN switch, overview 8
SFP (small form-factor pluggable) 9, 35, 43
shared storage configuration 26
single point of failure 10
SmartStart software 25
SmartStart, installation steps 26
SmartStart, overview 13
software installation 24
storage system, configuring 22
storage system, overview 7
storage system, set up 22
support 31
Systems Insight Manager, overview 14

T

TCP/IP (Transmission Control Protocol/Internet
Protocol) 9, 18, 26, 35
TCP/IP protocol 9, 18
technical support 31
telephone numbers 31

Z

zoning documentation 17