Implementation Guide:
Guide to Installation - Hardware Setup

*By Chris Dotson - Enterprise Engineering and
Steve Lopez - Enterprise Engineering*

Sun BluePrints™ OnLine - January 2002

http://www.sun.com/blueprints

Please
Recycle

Adobe PostScript™

MODULE 1

Hardware Setup ○━▼

Purpose The following equipment setup and cabling procedures detail how to configure the Sun Cluster 3.0 hardware to ensure that no single points of failure (SPOF) exist. Specifically, a properly configured hardware platform provides the highest level of availability that can reasonably be expected of the configuration. Important Sun Cluster 3.0 hardware setup concepts are also reviewed during the hands-on lab modules and exercises.

Objectives This module describes procedures to correctly assemble the cluster hardware, including:

❍   Identifying each Sun Cluster 3.0 hardware setup component and all equipment required for the SunPlex™ component.

❍   Identifying all cabling and interconnects.

Introduction When using this guide to construct a Sun Cluster hardware setup, participants perform configuration steps and procedures required to successfully configure each Sun Cluster 3.0 hardware component and construct a two-node Sun Cluster 3.0 hardware cluster.

Procedures include steps necessary to configure the cluster hardware (for example, each SunPlex component), along with stated examples of "key practices" (for example, configurations and methods for optimizing availability and/or performance).

A complete enterprise solution or architecture is a complex entity. Customers looking to purchase Sun Cluster hardware are often looking for more than just a pair of nodes to run a set of applications in an high availability (HA) environment. Each implementation must achieve specific identified needs, meeting goals for: serviceability, manageability, scalability, and availability.

The objective is to configure a Sun Cluster 3.0 hardware setup, including the software, to enhance data and application availability for production (for example, mission-critical) environments. The basic Sun Cluster 3.0 hardware platform being configured connects two or more servers by means of a private network. Each server can access the same application data using multi-ported or shared disk storage and share the network resources, thereby enabling either cluster node to inherit an application when its primary server becomes unable to provide services.

All configuration exercises contained within this Sun Cluster 3.0 hardware lab guide assume that you are implementing a standard cluster hardware configuration. Many steps in this lab guide refer to manual or local procedures which are performed only if you have local or physical access to the lab hardware, and you can perform visual inspection and verification. (These local, physical verification procedures are not covered in this lab guide). Final verification of the cluster hardware setup (for example, failover) can be confirmed only after the required software has been installed and configured. Formal acceptance test procedures include full verification for each system, subsystem, and component. See "References" on page 19 and Modules 2 through 6 of this lab guide for additional information.

Sun Microsystems

Figure 1-1 on page 7 describes the Sun Cluster 3.0 lab hardware implementation and Tables 1-1 through 1-5 on page 8 define each connection.

Prerequisites

This lab manual assumes that the user (participant) is a qualified Solaris Network Administrator. For installation queries, shell usage questions, patches, and packages refer to the Sun Educational Services manuals for the *Solaris System Administration 1* and *Solaris System Administration 2* courses.

Sun Cluster 3.0 Hardware Planning Considerations

Sun Cluster 3.0 Design Goal

Sun Cluster 3.0 hardware setup moves beyond Sun Cluster 2.2 hardware setup. By design, Sun Cluster 2.2 hardware setup provided a platform for availability and limited scalability. In contrast, the Sun Cluster 3.0 hardware setup platform (as configured during these lab exercises) creates a cluster environment capable of hosting HA applications that can be architected to provide near continuous availability and scalability, both horizontal and vertical.

Customer Requirements (Analysis)

Caution – For a given set of customer requirements, Sun Cluster 3.0 hardware applications are implemented across many layers, each layer having several aspects that need to be considered when proposing a solution. Only a comprehensive analysis of the customer's application requirements combined with a thorough knowledge of Sun Microsystems products and services can ensure that the cluster solution meets the customer's internal service level agreement (SLA). An incomplete analysis often leads to a poorly functioning cluster that adds little value to the customer's business and actually increases their cost of ownership. Dissatisfied customers are unlikely to purchase further clusters from Sun Microsystems. See "References" on page 19 for additional references.

Architectural Limitations

The Sun Cluster 3.0 hardware architecture is able to provide the highest levels of availability for hardware, operating system, and applications without compromising data integrity. The Sun Cluster 3.0 hardware environment (for example, hardware, operating environment, framework, API, and applications) can be customized to create HA applications.

Regardless of design, any software or hardware has physical or technical limitations. The Sun Cluster 3.0 hardware setup environment can be used in varying hardware and software configurations. However, before implementing the Sun Cluster 3.0 software, it is important to understand limitations of the product.

Configuring Clusters for HA: Primary Considerations

The primary configuration and planning considerations for HA applications and databases include identifying requirements for: software versions and features, boot environment, shared storage, and data services and their agents.

Designing a production cluster environment for a mission-critical environment is a complex task, involving the careful selection of optimum components amid a seemingly baffling array of options. We recommend that you work closely with a qualified consulting practice, such as Sun Professional Services, in making these selections. For example, determining the optimum number or mix of database instances or services per node, or ensuring that no potential agent conflicts exists and resolving any service level conflicts.

Different cluster topologies require carefully prescribed setup procedures in relation to the following cluster components:

❍ Number of logical hosts per node (including their agents, agent interoperability, and service level requirements)

❍ Type of Volume Manager

❍ Disk striping and layout

❍ File systems vs. Raw Device database storage

❍ Performance (local storage vs. GFS considerations)

❍ Network Infrastructure requirements and redundancy

❍ Client Failover Strategy

❍ Logical host Failover Method (manual vs. automatic)

❍ Naming conventions (host ID, disk label, disk groups, meta sets, mount points, etc.)

❍ Normal Operations Policies and Procedures

❍ Backup and Recovery Procedures for the SunPlex components

Note: See ''References'' on page 19, for relevant Sun BluePrints Online articles, when considering each of these important topics.

Standard Sun Cluster Configurations

The Sun Cluster 3.0 hardware configuration implemented for these hands-on labs represents an entry-level platform, providing no SPOF for the cluster pair.

No SPOF implies that a single component failure within the cluster platform will not permanently disable operations. Instead, the implemented failover mechanisms must be configured to enable the surviving node or components to continue operations. Failover and recovery times vary, depending upon the application and the configuration of the reliability, availability, and serviceability (RAS) features. Only through actual testing of failover operations and for each application/failure mode can recovery times be established. Published recovery times should reflect the estimated time of failover to occur and when services are resumed.

Study the descriptions and diagrams presented in this module, examining the Sun Cluster 3.0 hardware configuration for the cluster pair. Identify each component selected and the options implemented and consider the data paths and other RAS features configured to meet the design goals for the solution; no SPOF for the cluster-pair.

Additional design decisions and configuration options are detailed during subsequent lab modules and exercises, when important availability features are implemented.

No SPOF

Multiple faults occurring within the same cluster environment can result in unplanned downtime. A SPOF can exist within the software applications architecture or a SPOF for a single node might be an embedded controller or memory module.

This basic Sun Cluster 3.0 hardware configuration is based on the Sun Enterprise™ 220R server and should be configured as an entry-level platform for providing no SPOF for the cluster pair.

Installation and Planning Considerations: General

New installations that are well-planned and well-executed are critical to ensuring reliability and ultimately, availability. Reducing the number of unplanned outages involves using proven methods when configuring HA platforms and minimizing SPOFs. Adhering to these guidelines should result in a well-executed solid installation.

The following steps can contribute to successful configurations and assist in sustaining daily operations towards maximizing platform availability and performance:

❍ Document any potential SPOFs that could occur, including associated workarounds, troubleshooting procedures, and best practices.

❍ Ensure that Sun Cluster 3.0 hardware administrators are highly trained and able to successfully test and conduct cluster failover operations for each HA application, and associated systems and subsystems, including fault isolation/troubleshooting and recovery procedures using all available utilities and interfaces.

❍ Document site-specific and application-specific configurations and procedures as part of implementing best practices and simplifying datacenter operations and maintenance.

❍ Record all standard and non-standard configurations, implementing change management procedures for all systems and sub-systems (for auditing and tracking key systems and components throughout the life cycle of the data center).

❍ Implement well-known established configurations and techniques that minimize platform complexity and the number of active components to simplify operations and maintenance.

❍ Provide diagnostics, utilities, and interfaces that are easy to use and interpret with clearly documented error messages and procedures for resolving potential problems.

❍ Refer to the *Sun Cluster 3.0 Configuration Guide* for current restrictions on hardware, software, and applications running in the Sun Cluster 3.0 hardware environment.

Note: Please see "References" on page 19 when planning cluster installations for additional information, especially the document, *Sun Cluster Site Planning*.

Software Licenses

Sun Cluster 3.0 software requires its own user license. For bundled HA agents, obtain licenses from your local Sun Service provider. For HA agents developed by Sun Microsystems or third-party vendors requiring licenses, contact your local Sun Microsystems representative for professional services. Additionally, Sun Cluster 3.0 software does not include a VERITAS Volume Manager (VxVM) software or CVM software license (as Sun Cluster 2.2 software does) so it must be purchased.

Note: In some cases, Sun StorEdge™ arrays include a VxVM software license.

Cable Configuration Schematic   Figure 1-1 illustrates the hardware implementation and schematic for each connection.
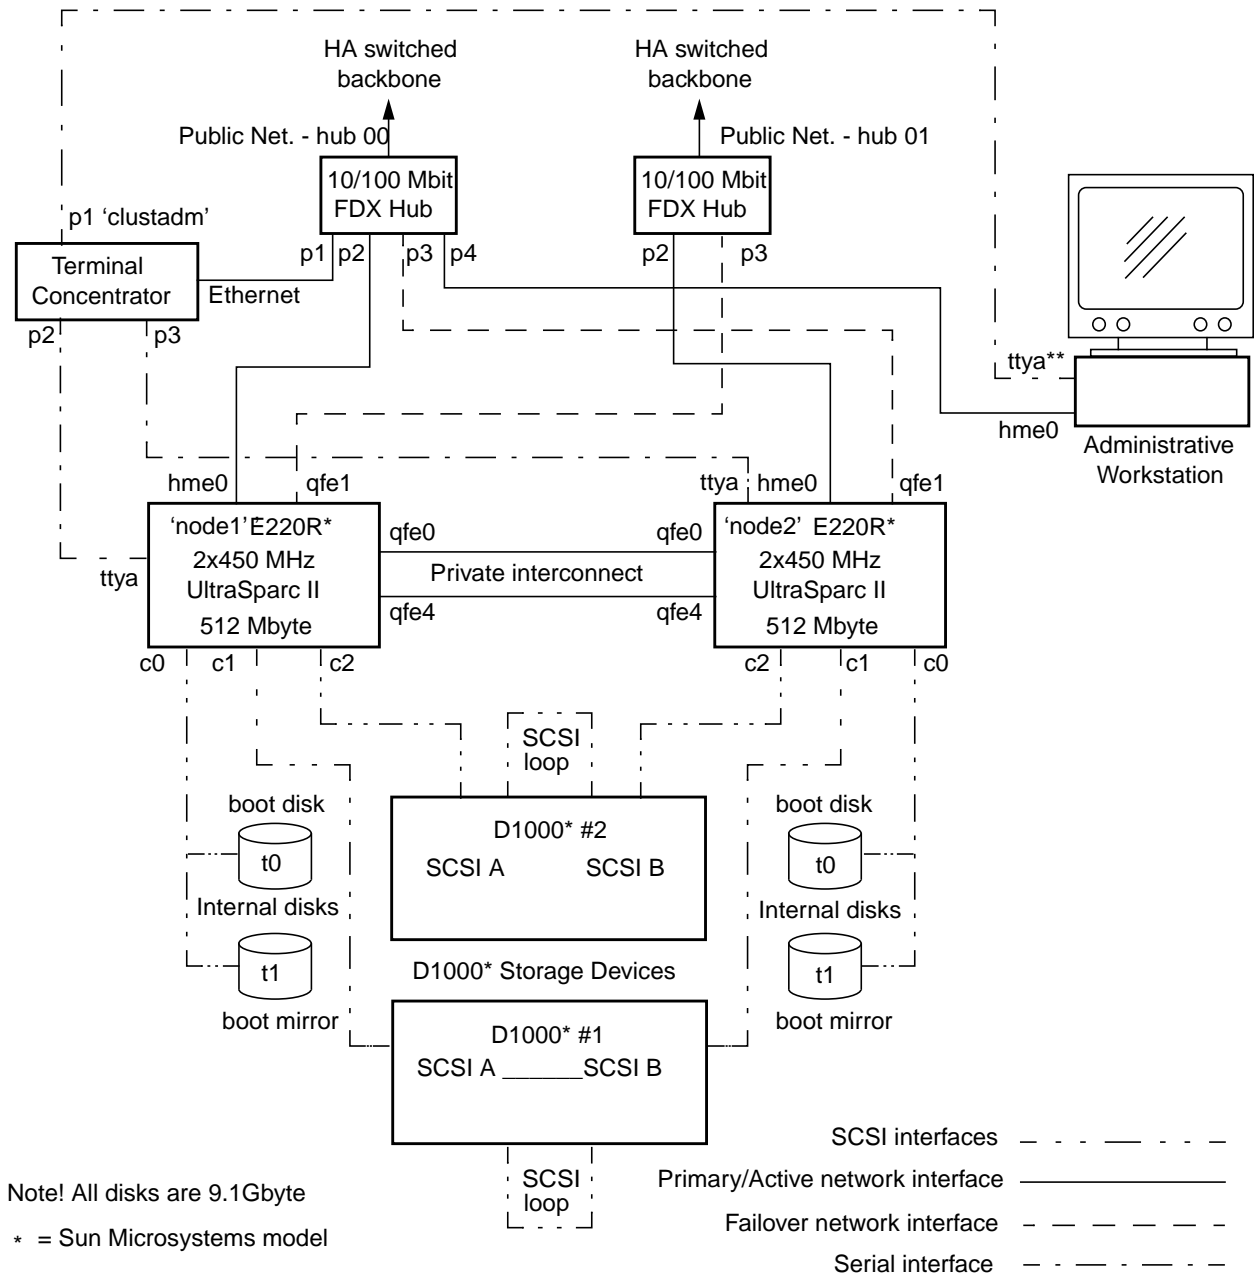
Figure 1-1 Cable Configuration Schematic

Note: See Figure 1-9 on page 16 for the proper stacking order of equipment.

Note: In Figure 1-1, c1 = PCI3, c2 = PCI4 and D1000's include t0, t1, t2, t8, t9, t10.

Cable Connection Tables

Tables 1-1 through 1-5 further define the configuration schematic illustrated in Figure 1-1 on page 7.

From Device	From Location	To Device	To Location	Cable Label
E220R # 1	SCSI A (PCI3)	D1000 #1	SCSI A	C3/1 - C3/3A
E220R # 2	SCSI A (PCI3)	D1000 #1	SCSI B	C3/1 - C3/3B
E220R # 1	SCSI A (PCI4)	D1000 #2	SCSI A	C3/2 - C3/3A
E220R # 2	SCSI A (PCI4)	D1000 #2	SCSI B	C3/2 - C3/3B

Table 1-1 Server to Storage Connections

From Device	From Location	To Device	To Location	Cable Label
E220R # 1	qfe0	E220R # 2	qfe0	C3/1 - C3/2A
E220R # 1	qfe4	E220R # 2	qfe4	C3/1 - C3/2B

Table 1-2 Private Network Connections

From Device	From Location	To Device	To Location	Cable Label
E220R # 1	hme0	Hub # 00	Port #2	C3/1 - C3/5A
E220R # 1	qfe1	Hub # 01	Port #3	C3/1 - C3/6A
E220R # 2	hme0	Hub # 01	Port #2	C3/2 - C3/5A
E220R # 2	qfe1	Hub # 00	Port #3	C3/2 - C3/6A

Table 1-3 Public Network Connections

From Device	From Location	To Device	To Location	Cable Label
E220R # 1	Serial Port A	Terminal Concentrator	Port #2	C3/1 - C3/4A
E220R # 2	Serial Port A	Terminal Concentrator	Port #3	C3/2 - C3/4A
Terminal Concentrator	Ethernet Port	Hub # 00	Port #1	C3/4 - C3/5A

Table 1-4 Terminal Concentrator Connections

From Device	From Location	To Device	To Location	Cable Label
Administration Workstation	hme0	Hub # 00	Port #4	F2/1 - C3/5A
Administration Workstation	Serial Port A	Terminal Concentrator	Port #1 **	F2/1 - C3/5B

Table 1-5 Administrative Workstation Connections

Note: Here's some additional details regarding the Cable Label column in Tables 1-1 through Table 1-5:

1. The Cable Label column assumes the equipment is located in grid location C3 (see Figure 1-8 on page 14). The number following the grid location identifies the level a piece of equipment is stacked with 1 being the lowest level. For additional details please see "Step 1.2" on page 16, and Figure 1-8 on page 14.

2. The letter at the end of the label tag indicates how many cables terminate at that level. For example, the letter A indicates one cable, B indicates two cables, etc.

3. The label tag F2 is the grid location of the administrative workstation. For additional information see "Step 1.2" on page 16.

4. ** indicates that this cable is only connected when configuring the terminal concentrator.

Equipment Summary

The hardware used for this lab example includes, two Sun Enterprise 220R servers (cluster nodes) configured with four PCI I/O cards. Figure 1-2 shows a Sun Enterprise 220R server. The lab kit may or may not have the four I/O cards installed.

Figure 1-2 Sun Enterprise 220R Server

Note: Two PCI `qfe` I/O cards (Sun Microsystems part # 1034A) should be installed in PCI slots 1 and 2. Install the Dual-Channel Differential UltraSCSI host adaptors (Sun Microsystems part # 6541A) in PCI slots 3 and 4.

Key Practices—Specifications should ensure consistent configurations, across nodes and devices, when possible. For example, the cluster boot, boot disk, and mirror should be of the same type or model, having identical device specifications. For these lab exercises, the cluster node(s) and disk subsystem configuration are configured in a consistent, symmetric manner, as represented in Figure 1-1 on page 7, and Tables 1-1 through 1-5 on page 8. As a result, configuration and sustaining operations often benefits. Ensure that the configuration of adapters and devices are consistent across all nodes. Unless the design specifications include cluster nodes that are specifically non-symmetric (due to failover and workload requirements) the components selected for use in each cluster node should be of the same type, model, or brand. During installation, ensure that all components are installed in the prescribed manner. Maintain consistent configurations for systems and subsystems. For example, if PCI slot 2 on node 1 has a PCI SCSI host adapter, then PCI slot 2 on node 2 should also have a PCI SCSI host adapter.

Note: The practice of ensuring consistent configurations across nodes can hold true for SBUS (EXX00) systems—if slot 0, 1, or 2 on system board 2 for node 1 has an SBUS `qfe` card—then the system board 2 on node 2 should have a corresponding SBUS `qfe` card in slot 0, 1, or 2. Following this practice assists in future troubleshooting and maintenance.

❍ **One 8-port Terminal Concentrator (TC).**

The terminal concentrator (see Figure 1-3) enables access to the console port on each node in the SunPlex component. Sun Cluster 3.0 hardware implementations no longer specify the TC brand and model. For some Sun Cluster 2.2 hardware implementations, the Sun Microsystems terminal concentrator was required and performed the task of *failure fencing* in configurations where the cluster had more than two nodes and a failure has occurred. This is not the case with a Sun Cluster 3.0 hardware setup.

Note: This component is also known as a *terminal server*, however, this lab guide refers to it as a *terminal concentrator (TC)* to avoid confusion with the typical meaning of the term *server*.

Figure 1-3 Typical 8 Port Terminal Concentrator (Rear view)—Ethernet Port Circled

Note: Configuration and setup of the TC requires that the connection of a serial cable (Sun Microsystems part # 530-2152) between port 1 of the TC and the serial/`tty` port on the Management Server (or administrative workstation). Furthermore, it is highly recommended that the TC be implemented using the exact type specified by Sun Microsystems according to the specific brand and model. This was a requirement for Sun Cluster 2.2 hardware setup.

❍ **Two Sun StorEdge D1000** disk storage devices.

The Sun StorEdge D1000 server (illustrated in Figure 1-4) has been designed to eliminate known SPOFs by incorporating dual power supplies and other redundancy features. In this exercise, each Sun StorEdge D1000 server is configured as a single backplane (loopback cables are installed for a non-split SCSI bus). Each Sun StorEdge D1000 server is configured with six disk drives, as target id 0, 1, 2, 8, 9, and 10. We use a volume manager (software RAID) to mirror across the two Sun StorEdge D1000 server arrays.

Figure 1-4 Sun StorEdge D1000 server

❍ **Two 10/100 BaseT Ethernet hubs**

Figure 1-5 illustrates a typical ethernet hub.

Figure 1-5 Typical Ethernet Hub

❍   **One Ultra 5™ Workstation**

Figure 1-6 shows a Ultra™ 5 workstation to be used as the administrative workstation and JumpStart™ server for the nodes. A larger workstation can also be used. Figure 1-7 shows the rear view of a Ultra 5 workstation serial port.

Figure 1-6 Ultra 5 WorkStation

Figure 1-7 Ultra 5 Workstation Serial Port (Rear View)—Serial Port Circled

Note: The serial cable is required when initially configuring a TC for use in the cluster.

Caution – According to common practice, when performing the hardware installation *do not* connect the power cords or attempt to power *on* the equipment until all hardware and components are installed and the cabling is completed. At the appropriate time, follow the recommendations for powering on subsystems and cluster nodes. One common practice suggests powering on all external (for example, storage) devices prior to powering on each node.

Step 1.1 For each cluster node, record the hostname, hostid, IP address, and Ethernet (MAC) address. Figure 1-8 is an illustration of equipment stacking detailing slot assignments and labels.
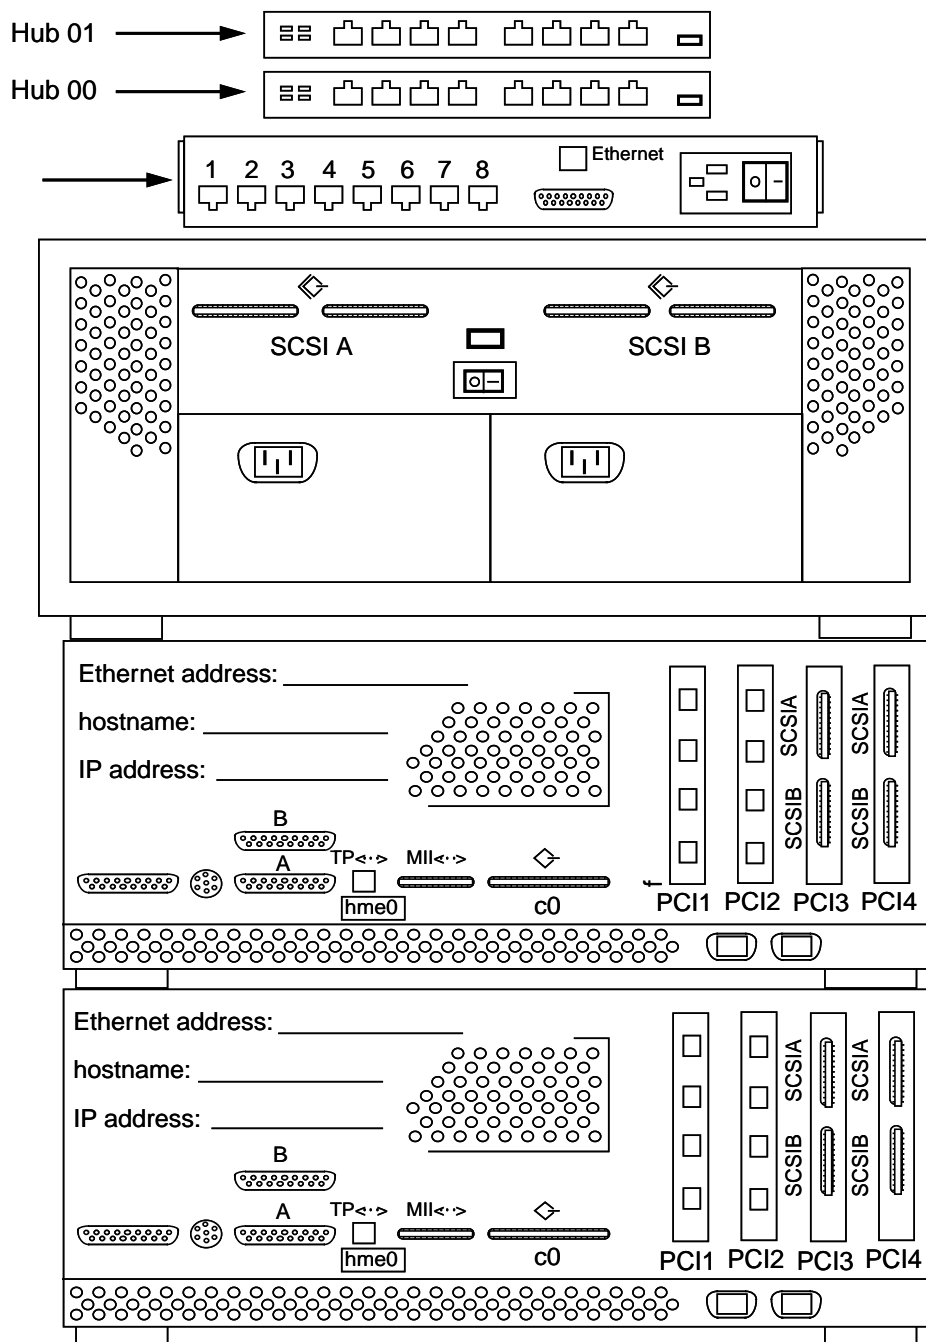
Figure 1-8 Equipment Stacking Showing Slot Assignments and Labels

Disk Subsystem Considerations	When configuring disk subsystem components, consider the following points, noting that *clustnode2's* global SCSI initiator ID is set to '6' avoiding the potential conflict with *clustnode1*:

❏ The embedded SCSI controller (`c0`) connects the root/swap disk (`c0t0`), plus the boot mirror (c0t1). PCI slots 3 and 4 are assigned to SCSI adapters `c1` and `c2`. PCI slot 3 (SCSI adapter) connects disks `c1t0` through `c1t2` and `c1t8` through `c1t10`. PCI slot 4 (SCSI adapter) connects `c2t0` through `c2t2` and `c2t8` through `c2t10`.

❏ Refer to Figure 1-1 on page 7 for the slot and controller assignments, SCSI target IDs, and datapaths.

❏ Both cluster nodes must share access to disk array #1 and #2. This means a potential conflict will arise if both *clustnode1* and *clustnode2* attempt to access the same SCSI bus. Prior to enabling the two cluster nodes to share the same SCSI bus the *clustnode2* global `scsi-initiator-id` is set to 6, while its internal assignment remains at 7. Remember that SCSI ID settings must be unique for all devices attached to the same SCSI bus, and that the SCSI ID establishes priority when multiple devices arbitrate for control of the SCSI bus. A global value of 6 establishes that *clustnode2* retains high priority. It is lower only to the *clustnode1* adapter(s), which remains set to the default value of 7. The settings for *clustnode2* are established in the `nvram`, to ensure that the configuration is saved across power outages.

Additional Sun StorEdge D1000 server and disk considerations (for example, performance, availability, partitioning, etc.) are addressed during subsequent lab modules.

Private Networks	Figure 1-1 on page 7 illustrates how each cluster node is interconnected via their respective `qfe0` and `qfe4` ports. The basic Sun Cluster 3.0 hardware platform connects two or more servers by means of a private cluster interconnect.

The cluster interconnect supports cluster application data and/or locking semantics between cluster nodes. The cluster interconnect is fundamental to cluster operations. Do not use the interconnect for routing any other traffic or data. The private network establishes exclusive use of pre-assigned or hard-coded IP addresses for each cluster node. Redundant network links (for example, fault tolerant) are implemented. Failover is transparent and immediate. The type of application determines which interconnects are supported.

Note: Using cross-over cables can minimize the use of hubs and conform to established **Key Practices** in a two-node cluster.

Terminal Concentrator (TC)	Because a cluster node does not have a monitor or keyboard, a single TC is used for console access and operations. The TC is not required for normal operation of cluster nodes. If the single TC fails, access to the cluster nodes is provided through public network connections. You may configure multiple terminal concentrators to achieve higher levels of availability.
Switches and Hubs	It is recommend that two hubs are configured in a basic cluster for redundant connections to public networks and switches. On each cluster node, Figure 1-1 on page 7 depicts how NAFO software can provide failover between `hme0` and `qfe1`, however, Sun Professional Services strongly recommends more reliable network availability solutions using Alteon Websystem or Extreme Networks switches. See "References" on page 19 for additional references to specific Sun BluePrints Online articles.
TC Settings	In the space provided, record the settings and TC cable connections. This data is used to create the /etc/serialports file during subsequent lab exercises.

Port on TC	Node Name
1	clustadm
2	clustnode1
3	clustnode2
4	

Step 1.2   One example easily identifies cables using labels. Cables used to connect equipment should be clearly marked to reflect the following information (local and manual installations only):

❍   Location within a datacenter or a grid reference

❍   From device including the level ID

❍   To device including the level ID

The datacenter floor can be divided into a grid arrangement, similar to Figure 1-9.
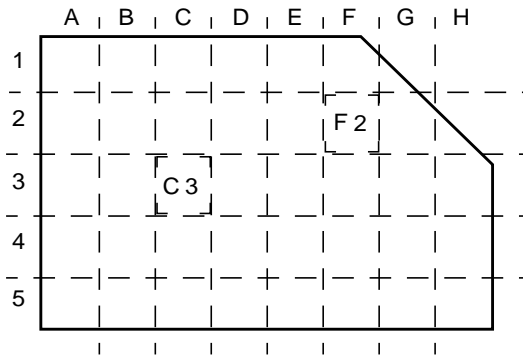

Figure 1-9 Datacenter Grid Diagram

Each location within the grid is known as a cabinet. In this exercise, we have stacked the equipment at grid reference C3, with the administrative workstation located at grid reference F2. Additionally, each piece of equipment within the cabinet has a level ID. Labelling each cable end with an identification tag makes it easier for future troubleshooting or maintenance. A cable tag may look like the follow example.

Example: C3/6 - F2/1. This is from grid reference/level ID to grid reference/level ID.

Step 1.3   For manual installations, label all major components as to their function or name.

Key Practices—Label devices and cables. Devices and cables that are easily identified make troubleshooting and maintenance easier. We have found this helpful even for systems that do not require high availability. Where high availability is a pre-requisite, cable and device identification and labelling should be high priority. Verify that the disks (boot devices, mirrors, disk quorums, clones, hot-spares, etc.), the tape drives, and the cable labels are correct. This allows service operations to correlate error messages to specific global Device ID (DID) numbers, metadevice names, or the specific controller number and sd/ssd instances and help in the interpretation of these errors to determine specific failed disks and the related components. Label tape drives with their rmt instance numbers.

Step 1.4   Connect the cables to the components as illustrated in Figure 1-1 on page 7 and as detailed in Tables 1-1 to 1-5 on page 8, paying special attention to the following information.

Key Practices—Connect cables used for the same function to the same cards and ports on both machines. This assists in troubleshooting; especially with installations where it is difficult to visually trace cabling between devices.

Example: If the private interconnect (heartbeat) #l is connected to port 0 of the Ethernet card in bay 1 on node #1—it should be connected to port 0 of the Ethernet card in bay 1 on node #2.

Key Practices—To minimize SPOFs, distribute the I/O and datapaths (cabling and connections) across as many different I/O Cards or system boards as possible and practical. Implementing this practice ensures that no single board failure will bring down the redundant hardware component enabling rapid recovery efforts.

Example: If the private interconnect (heartbeat) line #l is connected to the Ethernet card in slot 1 then the private network line #2 should be connected to the Ethernet card in slot 2. Therefore, if either Ethernet card fails, the other private network line is unaffected and the system continues to function.

Note: The private interconnect cables must be "Null" Ethernet crossover cables (Sun Microsystems part # X3837A) because they connect directly machine-to-machine. Connect `qfe4` on node 1 to `qfe4` on node 2, connect `qfe0` on node 1 to `qfe0` on node 2 as shown in Figure 1-10.
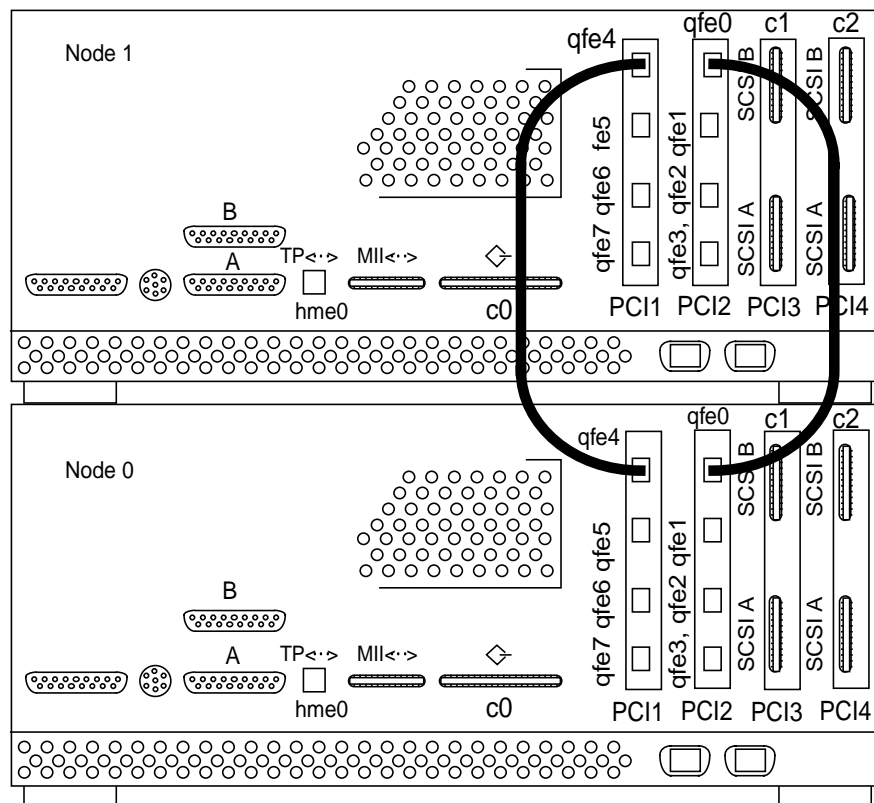

Figure 1-10 Sun Enterprise 220R Server Private Interconnect Cabling

Step 1.5   For manual installations, connect all device power cords and power on all equipment *except* the TC which is configured separately.

Summary of Key Practices

The following is a review list of the key practices that we have detailed so far in this lab guide.

> Establish Cluster Requirements (design goals) and implement configurations to meet these specifications.
>
> Plan the installation carefully, considering all components of the SunPlex solution and applications, throughout the production life-cycle.
>
> Ensure consistent configurations across all nodes and devices, when possible, benefiting sustaining operations and maintenance.
>
> Identify and label all devices and cables, clearly.
>
> Minimize SPOFs with straight-forward designs, distributing cabling and connections across multiple system boards, and IO cards.
>
> Configure fewer active components, in a less complex and consistent manner (for example, cross-over cables instead of active hubs).
>
> Do not power on to any equipment until all cables and connections have been verified.
>
> Disconnect TC port 1 when TC has been configured and verified. All access to Port 1 should be carefully managed to avoid potential security risks.

Note: The final validation of this module occurs during Module 2, after you have configured the TC software and have installed the Sun Cluster client software on the administrative workstation. At this time, it must be assumed that no equipment faults exist for any of the SunPlex components.

End of Module One

At the end of Module One, the hardware has been set up. In this module, you have identified each component required to successfully complete the hardware installation, including:

1. Management Server or Admin Workstation

2. Cluster nodes

3. Sun StorEdge D1000 server Disk Arrays (shared storage)

4. Private and Public Networks

5. Terminal Concentrator (TC)

6. Ethernet hub(s)

7. Cabling and connections

References

Contents
This section provides a list of references, as well as related URLs to online documents and Sun BluePrint Online articles to review with this Sun Cluster 3.0 Lab Guide and when implementing Sun Cluster hardware solutions.

On-Line References to Sun Cluster 3.0 Documents and Sun BluePrints Online Article*s*

Cluster Platform 220/100 Architecture—A Product from SunTone Platforms Portfolio by Enrique Vargas can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Excellent comprehensive examples of Sun Cluster installation checklists (configuration worksheets) can be found at this URL:
`http://sunweb.germany/EIS/Web/inst-support/checklists/cluster30-e.pdf`

Recommend reviewing the following chapters; Chapter 3, "Sun Cluster Site Planning," Chapter 4 "Site Preparation," and Chapter 8 "Acceptance Testing," which you can download at this URL:
`http://sunweb.germany/EIS/Web/inst-support/EDocu/server/index.html`

Sun Cluster 3.0 System Administration Guide user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Installation Guide user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Hardware Guide user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Release Notes: refer to the Patches and Firmware requirements; Appendix A "Sun Cluster Installation and Configuration Worksheets," and user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Concepts user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Error Messages manual user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 (U1) Data Services Installation and Configuration Guide user documentation can be found at this URL:
`http://suncluster.eng.sun.com/products/SC3.0/`

Sun Cluster 3.0 Cluster File System (CFS): Making the most of the global file service, by Tim Read, Senior Consultant, Sun Microsystems, UK and many valuable contributions can be found at this URL:

`http://llgweb.uk.sun.com/~timr/`

Sun Web Start Flash Archive (Flash Project) functional specification, by Matt Simmons can be downloaded at this URL:

`http://sps.central/Flash/func_spec.wp.pdf`

Sun Web Start Flash information can be found at this URL:

`http://www.sun.com/solaris/webstartflash/`

Sun BluePrints Online Articles and valuable writings are available by browsing this URL:

`http://www.sun.com/blueprints`

Availability—What It Means, Why It's Important, and How to Improve It by Richard McDougal can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Disaster Recovery Requirements by Stan Stringfellow, can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

High Availability Best Practices by Enrique Vargas can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

HA: Boot/Root/Swap by Jeannie Johnstone Kobert can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

The Art of Production Environment Engineering by Bill Walker, can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Datacenter Naming Scheme by Mark Garner, can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Robust Clustering: A Comparison of Sun Cluster 3.0 vs. Sun Cluster 2.2

`http://www.sun.com/blueprints/0901/sc30vs22.html`

Architecting a Service Provider Infrastructure for Maximum Growth by Stan Stringfellow can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Building Secure N-Tier Environments by Alex Noordergraaf can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Planning for Large Configurations of Netra t1 Servers by Stan Stringfellow can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Planning to Fail by John S. Howard can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

High Availability Fundamentals by Enrique Vargas can be found at this URL:

`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Tales from the Trenches: "The Case of the RAM Starved Cluster", by Richard McDougal can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Building a JumpStart Infrastructure, by Alex Noordergraaf can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Using NTP to control and Synchronize System Clocks (Parts 1 - 3) can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Solaris 8 Additions to "sysidcfg" by Rob Snevely can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

An Overview of Methodology, by Adrian Cockcroft can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

JumpStart™ Mechanics: Using JumpStart Application for Hands-Free Installation of Unbundled Software (Parts 1 and 2) by John S. Howard can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Setting Up A Solaris Operating Environment Install Server and the Solaris JumpStart Feature by Rob Snevely can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Sun GigaBit Ethernet vs. Alteon WebSystems, Inc.: Performance (throughout) Benchmark Analysis and Datasheet can be found at this URL:
`http://www.sun.ca/products-n-solutions/hw/networking/connectivity/sungigabitethernet/gem.html#analysis`

Scrubbing Disks using Solaris OE "format" Program, can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Fast Oracle Exports by Stan Stringfellow can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

IP Network Multipathing by Mark Garner can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Sun/Oracle Best Practices can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

MR System for Rapid Recovery by John S. Howard can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

VxVM Private Regions: Mechanics and Internals of the VxVM Configuration Database by Gene Trantham can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Veritas VXVM Storage Management Software by Gene Trantham can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

*SSP Best Practice*s by John S. Howard can be found at this URL:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Author's Bio

Chris Dotson is a member of the Enterprise Engineering staff at Sun Microsystems, Inc. Since 1992, he has provided global support and technical education for Sun Microsystems, Inc. on a wide range of topics, including the Solaris Operating Environment, Volume Management Software, Network Security, Sun Enterprise 10000 Servers, and Sizing and Capacity Planning. Chris has designed technical education and curricula for delivering worldwide training to thousands of Sun Microsystems, Inc. field engineers and global partners. He has presented his curricula throughout the United States, Canada, South America, South Africa, India, Asia, and Europe.

Author's Bio

Steve Lopez is a member of the Enterprise Engineering staff at Sun Microsystems. With over 16 years of UNIX® experience, he has served as post sales and pre sales technical support at Sun Microsystems for 14 years. He authored an Installation Guide on LDAP 4.1.2 and co-authored two installation guides on Sun Cluster 2.2 and Sun Cluster 3.0. Steve recently designed and developed the Remote Lab Environment for Sun technical staff worldwide to practice Architecting, Implementing and Maintaining Sun solutions based on Sun Blueprints Guides.