

HP AlphaServer SC

Installation Guide

December 2005

This document describes how to install an HP AlphaServer SC system from the Hewlett-Packard Company.

Revision/Update Information:	This version supersedes the <i>HP AlphaServer SC Installation Guide</i> issued in September 2004 for HP AlphaServer SC Version 2.6 (Update Kit 1).
Operating System and Version:	HP Tru64 UNIX Version 5.1B-3 (also known as HP Tru64 UNIX Version 5.1B Patch Kit 5)
Software Version:	Version 2.6 (Update Kit 2)
Maximum Node Count:	1024 nodes
Node Type:	HP AlphaServer ES45 HP AlphaServer ES40 HP AlphaServer DS20L

Legal Notices

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.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY

3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Use of this manual and media is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs, in their present form or with alterations, is expressly prohibited.

Copyright Notices

© Copyright 2003, 2005 Hewlett-Packard Development Company, L.P

Some information in this document is based on Platform documentation, which includes the following copyright notice:
Copyright 1994 - 2003 Platform Computing Corporation.

The Load Sharing Facility (LSF) software product is developed by Platform Computing Corporation ("Platform"). LSF has been licensed by Platform to HP for inclusion in HP AlphaServer SC systems.

The HP MPI software that is included in this HP AlphaServer SC software release is based on the MPICH V1.2.4 implementation of MPI, which includes the following copyright notice:

© 1993 University of Chicago
© 1993 Mississippi State University

Permission is hereby granted to use, reproduce, prepare derivative works, and to redistribute to others. This software was authored by:

*Argonne National Laboratory Group
W. Gropp: (630) 252-4318; FAX: (630) 252-7852; e-mail: gropp@mcs.anl.gov
E. Lusk: (630) 252-5986; FAX: (630) 252-7852; e-mail: lusk@mcs.anl.gov
Mathematics and Computer Science Division, Argonne National Laboratory, Argonne IL 60439*

Mississippi State Group

N. Doss and A. Skjellum: (601) 325-8435; FAX: (601) 325-8997; e-mail: tony@erc.msstate.edu

Mississippi State University, Computer Science Department & NSF Engineering Research Center for Computational Field Simulation, P.O. Box 6176, Mississippi State MS 39762

GOVERNMENT LICENSE

Portions of this material resulted from work developed under a U.S. Government Contract and are subject to the following license: the Government is granted for itself and others acting on its behalf a paid-up, nonexclusive, irrevocable worldwide license in this computer software to reproduce, prepare derivative works, and perform publicly and display publicly.

DISCLAIMER

This computer code material was prepared, in part, as an account of work sponsored by an agency of the United States Government. Neither the United States, nor the University of Chicago, nor Mississippi State University, nor any of their employees, makes any warranty express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Trademark Notices

Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

The following product names refer to specific versions of products developed by Quadrics Ltd ("Quadrics"). These products combined with technologies from HP form an integral part of the Supercomputing Systems produced by HP and Quadrics. These products have been licensed by Quadrics to HP for inclusion in HP AlphaServer SC systems.

- *Interconnect hardware developed by Quadrics, including switches and adapter cards*
- *Elan, which describes the PCI host adapter for use with the interconnect technology developed by Quadrics*
- *PFS or Parallel File System*
- *RMS or Resource Management System*

Contents

Preface	xxi
----------------------	-----

1 Installation Overview

1.1	Installation Overview	1-2
1.2	License Requirements.....	1-3
1.3	HP AlphaServer SC Subsets.....	1-4
1.4	Base System Subsets.....	1-6
1.5	General Considerations.....	1-7

2 Pre-Installation Planning

2.1	Review the Release Notes	2-2
2.2	Assign the External Network IP Addresses	2-2
2.3	Assign the System Name and Default Cluster Aliases.....	2-3
2.3.1	System Name.....	2-3
2.3.2	Domain Types	2-4
2.3.3	Default Cluster Aliases.....	2-4
2.4	Plan the Local and Global Storage	2-7
2.4.1	Confirm the Disk Layout.	2-9
2.4.2	Plan the Partition Sizing for Local Storage.....	2-11
2.4.3	Review the System Storage Rules	2-12
2.4.4	Plan the RAID Configuration.....	2-13
2.4.4.1	Planning the HSG80 RAID Configuration.....	2-14
2.4.4.2	Planning the HSV110 RAID Configuration.....	2-14
2.4.4.3	Planning the MSA1000 RAID Configuration	2-15
2.4.4.4	Managing a Fibre Channel Switch.....	2-16
2.5	Choose the Root Password	2-16
2.6	Record the External Gateway IP Address	2-16

3 Physical Installation

3.1	Management Server	3-2
3.1.1	Standalone Management Server	3-2
3.1.2	Clustered Management Server	3-3
3.2	Compute Nodes and File Serving Nodes	3-4
3.3	Physical Installation Overview.....	3-4

3.4	Connect the Management Network.	3-11
3.4.1	Cable the Management Network	3-11
3.4.2	Configure the Summit Switch.	3-12
3.4.3	Configure the ProCurve Switch	3-14
3.5	Populate the HP AlphaServer SC PCI Slots	3-17
3.5.1	HP AlphaServer ES40 PCI Slots	3-17
3.5.2	HP AlphaServer ES45 PCI Slots	3-18
3.5.3	HP AlphaServer DS20L PCI Slots	3-20
3.6	Configure Hardware for a Dual-Rail Configuration	3-20
3.7	Connect the HP AlphaServer SC Interconnect	3-20
3.8	Connect the Node Console Port	3-21
3.9	Configure the HP AlphaServer SC Interconnect Control Card with an IP Address.	3-22
3.10	Configure the Terminal Servers with an IP Address	3-22
3.11	Connect the Fibre Channel Switches	3-23
3.12	Verify Storage Component Revisions.	3-24
3.13	Configure the Fibre Channel Switch.	3-24
3.13.1	Configure Fibre Channel Switch Network Addressing	3-24
3.13.2	Switch Zoning	3-26
3.14	Configure the System Storage on the HSG80.	3-27
3.14.1	HSG80 Storage Management Tools	3-29
3.14.1.1	Managing Storage Arrays and Associated HSG80 RAID Controllers.	3-29
3.14.1.2	Using SANworks Command Scripiter.	3-30
3.14.2	Configure the HSG80 RAID Controllers.	3-31
3.14.3	Configure the HSG80 RAID Storage	3-36
3.15	Configure the System Storage on the HSV110.	3-44
3.15.1	HSV110 Storage Management	3-46
3.15.1.1	Licenses	3-46
3.15.1.2	The HSV Element Manager	3-46
3.15.1.3	SANWorks Scripting Utility	3-47
3.15.2	Configure the HSV110 RAID Controllers.	3-47
3.15.2.1	Locate the Example EVA (HSV110) Configuration Scripts	3-48
3.15.2.2	Calculate the Minimum Storage Requirement	3-48
3.15.2.3	Create Virtual Disk Units for the Clustered Management Server	3-52
3.15.2.4	Edit the Configuration Script	3-55
3.15.2.5	Run the Configuration Script	3-61
3.16	Configure the System Storage on the MSA1000	3-62
3.16.1	Check the MSA1000 Controller Firmware.	3-64
3.16.2	Configure the MSA1000 RAID Storage	3-64
3.16.3	Rename the Connections	3-66

4 Upgrade Installation Procedure

4.1	Understanding the Upgrade Process	4-2
4.1.1	Upgrade Versus Installation.	4-2

4.1.2	Upgrade Procedure Overview	4-3
4.1.3	Upgrade Restart	4-4
4.1.4	Upgrade States	4-4
4.1.4.1	Check the Upgrade State	4-5
4.1.5	Preserved and Unpreserved Files	4-5
4.1.6	Special Upgrade Mechanisms on Domains	4-6
4.1.7	Approximate Timing Guideline	4-7
4.2	Pre-Upgrade Audit	4-8
4.3	Preparing RIS for C2 Security	4-10
4.4	Upgrade with a Management Server	4-10
4.4.1	Back Up the Management Server	4-11
4.4.2	Upgrade the Management Server	4-11
4.4.2.1	Prepare for Upgrade	4-11
4.4.2.2	Perform the Pre-Upgrade Check	4-13
4.4.2.3	Back Up the SC Database	4-13
4.4.2.4	Remove the HP AlphaServer SC Software	4-14
4.4.2.5	Install the Operating System Patch Software	4-14
4.4.2.6	Build the New Kernel and Reboot	4-17
4.4.2.7	Configure the RIS Server	4-17
4.4.2.8	Install the HP AlphaServer SC Software	4-19
4.4.2.9	Restore the SC Database	4-19
4.4.2.10	Migrate the SC Database	4-20
4.4.2.11	Register RIS Clients for the New RIS Environment	4-21
4.4.2.12	Build the New Kernel and Reboot	4-21
4.5	Upgrade with a Clustered Management Server	4-21
4.5.1	Back Up the Clustered Management Server	4-22
4.5.2	Upgrade the Clustered Management Server	4-22
4.5.2.1	Prepare for Upgrade	4-22
4.5.2.2	Perform the Pre-Upgrade Check	4-24
4.5.2.3	Back Up the SC Database	4-24
4.5.2.4	Remove the HP AlphaServer SC Software	4-25
4.5.2.5	Install the Operating System Patch Software	4-25
4.5.2.6	Configure the RIS Server	4-28
4.5.2.7	Install the HP AlphaServer SC Software	4-29
4.5.2.8	Restore the SC Database	4-29
4.5.2.9	Migrate the SC Database	4-30
4.5.2.10	Register RIS Clients for the New RIS Environment	4-31
4.5.2.11	Build the New Kernel and Reboot	4-31
4.6	Upgrade without a Management Server	4-32
4.6.1	Back Up the First Domain	4-32
4.6.2	Upgrade the First Domain	4-32
4.6.2.1	Prepare for Upgrade	4-33
4.6.2.2	Perform the Pre-Upgrade Check	4-34
4.6.2.3	Back Up the SC Database	4-34

4.6.2.4	Install the Upgrade Subset on the First Domain	4-35
4.6.2.5	Disable Cookies	4-35
4.6.2.6	Configure the RIS Server	4-36
4.6.2.7	Register RIS Clients for the New RIS Environment	4-37
4.6.2.8	Upgrade the First Domain	4-38
4.7	Upgrading the Domains	4-41
4.7.1	Back Up the Domains	4-41
4.7.2	Upgrade all of the Domains	4-42
4.8	Recover from Failures during an Upgrade	4-45
4.8.1	Minor Failures: Re-add Members that Failed to Upgrade	4-45
4.8.2	Medium Failures: Re-run the sra upgrade Command	4-46
4.8.3	Serious Failures: Recover from Backup	4-46
4.8.4	Severe Failures: Full System Restoration	4-48
4.9	Post-Upgrade Tasks	4-49
4.9.1	Create Alternate Boot Partitions	4-49
4.9.2	Upgrade the Database Revision	4-49
4.9.3	Re-Add Members Deleted prior to Upgrade.	4-50
4.9.4	Re-enable MSQL Cookie Mechanism	4-50
4.9.5	Add the cmf Archive entry into crontab	4-50
4.9.6	Add the rms Archive entry into crontab	4-50
4.9.7	Verify Network Configuration	4-50
4.9.8	Verifying Local Files and Customizations	4-51
4.9.9	Check that Members 2 and 3 Have a Vote	4-51
4.9.10	Configure Nodes into RMS Partitions	4-51
4.9.11	Setting Up the SC Monitor System	4-51
4.9.12	Assigning Cabinets in the SC Database.	4-51
4.9.13	Reinstall Uninstalled Kits	4-51
4.9.14	Build and Deploy Generic Kernels	4-52

5 Installing: When the System Has a Management Server

5.1	Set Up the Management Server	5-2
5.1.1	Set the Console Variables	5-3
5.1.2	Check the System Firmware	5-4
5.1.3	Creating Bootable Devices on EMA/EVA Storage	5-4
5.1.4	Install the Tru64 UNIX Operating System	5-6
5.1.5	Customize the System Configuration	5-10
5.1.5.1	Register Licenses (PAKs)	5-10
5.1.5.2	Set Up Networks	5-14
5.1.5.3	Configure DNS (BIND)	5-17
5.1.5.4	Configure NTP	5-18
5.1.5.5	Configure NFS	5-18
5.1.5.6	Configure NIS	5-19
5.1.5.7	Configure Mail	5-21

5.1.5.8	Configure Printers	5-21
5.1.6	Install the Operating System Patch Software	5-22
5.1.7	Install and Configure the Clustered Management Server	5-22
5.1.7.1	Register the TruCluster Server Software License	5-23
5.1.7.2	Load the TruCluster Server Subsets	5-24
5.1.7.3	Install the TruCluster Server Patch Software	5-24
5.1.7.4	Configure the EVA Storage System Disks for the Management Cluster	5-25
5.1.7.5	Run the clu_create Command	5-25
5.1.7.6	Back Up Important Configuration Files on Member 1	5-27
5.1.7.7	Prevent Other Members Booting During Installation	5-27
5.1.7.8	Run the clu_add_member Command	5-28
5.1.7.9	Boot the New Member	5-29
5.1.7.10	Back Up Important Configuration Files on New Members	5-30
5.1.7.11	Shut Down the Non-Lead Member of the Clustered Management Server	5-31
5.1.8	Configure the RIS Server	5-31
5.1.9	Install the HP AlphaServer SC System Software	5-32
5.1.10	Install the HP Fortran Run-Time Libraries	5-33
5.1.11	Install Layered Products (Optional)	5-33
5.1.12	Install the SANworks Storage System Scripting Utility	5-33
5.1.13	Define the RMS Master Node (rmshost)	5-34
5.1.14	Build the New Kernel and Reboot	5-35
5.2	Set Up the SC Database	5-37
5.3	Configure Out All Nodes During Installation	5-45
5.4	Check All Nodes in the HP AlphaServer SC System	5-45
5.4.1	Check the State of the Nodes	5-45
5.4.2	Check the System Firmware	5-46
5.5	Configure and Diagnose the HP AlphaServer SC Interconnect	5-47
5.5.1	Upgrading the HP AlphaServer SC Interconnect Control Processor Software	5-48
5.5.2	Creating an Interconnect Configuration Using SC Viewer	5-48
5.5.3	Confirming the Operation of the HP AlphaServer SC Interconnect	5-49
5.6	Set Up the SC Monitor System	5-50
5.6.1	Add a SAN Appliance as Monitored Devices	5-50
5.6.2	Change Terminal Servers Monitoring Distribution	5-50
5.6.3	Change Extreme Switches Monitoring Distribution	5-51
5.6.4	Activate the Changes	5-52
5.7	Assign Cabinets in the SC Database	5-52
5.7.1	Load the Physical Relationships	5-53
5.7.2	Populate Cabinets with Nodes	5-54
5.7.3	Populate Cabinets with Other Hardware	5-54

6 Installing: When the System Does Not Have a Management Server

6.1	Set Up Node 0	6-2
6.1.1	Set the Console Variables	6-3

6.1.2	Check the System Firmware	6-4
6.1.3	Install the Tru64 UNIX Operating System	6-4
6.1.4	Customize the System Configuration	6-8
6.1.4.1	Register Licenses (PAKs)	6-8
6.1.4.2	Set Up Networks	6-12
6.1.4.3	Configure DNS (BIND)	6-14
6.1.4.4	Configure NTP	6-15
6.1.4.5	Configure NFS	6-16
6.1.4.6	Configure NIS	6-16
6.1.4.7	Configure Mail	6-17
6.1.4.8	Configure Printers	6-18
6.1.5	Install the Latest Operating System Patch Software	6-18
6.1.6	Configure the RIS Server	6-19
6.1.7	Install the HP AlphaServer SC System Software	6-20
6.1.8	Install the HP Fortran Run-Time Libraries	6-21
6.1.9	Install Layered Products (Optional)	6-21
6.1.10	Install the SANworks Storage System Scripting Utility	6-21
6.1.11	Add sysconfigtab Parameters	6-22
6.1.12	Define the RMS Master Node (rmshost)	6-23
6.2	Set Up the SC Database	6-24
6.3	Check All Nodes in the HP AlphaServer SC System, Except Node 0	6-31
6.3.1	Check the State of the Nodes	6-31
6.3.2	Check the System Firmware	6-31
6.4	Configure and Diagnose the HP AlphaServer SC Interconnect	6-33
6.4.1	Upgrading the HP AlphaServer SC Interconnect Control Processor Software	6-33
6.4.2	Creating an Interconnect Configuration Using SC Viewer	6-34
6.4.3	Confirming the Operation of the HP AlphaServer SC Interconnect	6-34
6.5	Set Up the SC Monitor System	6-35
6.5.1	Add a SAN Appliance as Monitored Devices	6-36
6.5.2	Change Terminal Servers Monitoring Distribution	6-36
6.5.3	Change Extreme Switches Monitoring Distribution	6-37
6.5.4	Activate the Changes	6-38
6.6	Assign Cabinets in the SC Database	6-38
6.6.1	Load the Physical Relationships	6-39
6.6.2	Populate Cabinets with Nodes	6-39
6.6.3	Populate Cabinets with Other Hardware	6-40
6.7	Review the SC Database Disk Settings	6-41
6.8	Transform Node 0 into a Single Node Domain	6-43
6.9	Configure Out All Nodes During Installation	6-44
6.10	Run the HP AlphaServer SC Interconnect Tests on Node 0	6-44

7 Building the Domains

7.1	Understanding the Automated Installation Process	7-1
-----	--	-----

7.1.1	SC Database Installation Tables	7-2
7.1.2	The sra install Command.	7-2
7.1.3	The Installation Daemon	7-2
7.1.4	Installation States.	7-3
7.1.5	Monitoring the Installation State	7-4
7.1.5.1	Displaying Information on sra Commands	7-5
7.1.5.2	Querying Node Installation Status	7-5
7.1.5.3	Starting the Monitor Utility	7-5
7.1.5.4	Filtering Views of Installation Progress	7-11
7.1.5.5	Text Based sramon Utility	7-13
7.2	Review the SC Database System Settings	7-14
7.3	Add sysconfigtab Parameters	7-23
7.4	Create the Domains.	7-23
7.5	Boot the System	7-25
7.6	Complete the Setup of the Domains	7-25

8 Completing the Installation

8.1	Boot the Non-Lead Member of the Clustered Management Server	8-2
8.2	Configure the External Network Interfaces	8-2
8.2.1	Set Up NFS to Allow Mounts from External Machines	8-2
8.3	Improve Cluster Availability	8-3
8.3.1	Cluster Quorum	8-4
8.3.2	Add Votes	8-4
8.3.3	Side Effects of Using Quorum	8-6
8.3.3.1	Shutting Down a Domain.	8-6
8.3.3.2	Bootting a System	8-7
8.4	Load File System Configuration Data in the SC Database	8-7
8.5	Initial Setup of Monitoring for HSG80 RAID Systems	8-8
8.6	Configure the RMS Database	8-10
8.6.1	Set Up RMS Partitions	8-10
8.6.2	Customize RMS Partitions	8-11
8.7	Provide RMS with CAA Failover Capability	8-12
8.8	Enable CMF as a CAA Application	8-15
8.9	Run the Example MPI Program	8-17
8.10	Verify the HP AlphaServer SC Interconnect	8-17
8.11	Configure LSM	8-18
8.12	Verify Swap Mode	8-18
8.13	Add a Second Rail to an HP AlphaServer SC System after Domain Creation	8-18
8.13.1	HP AlphaServer SC System Composed of HP AlphaServer ES40 Nodes	8-19
8.13.2	HP AlphaServer SC System Composed of HP AlphaServer ES45 Nodes	8-21

9 Installing LSF for HP AlphaServer SC

9.1	LSF Overview	9-2
9.1.1	Preparing for LSF Installation	9-2
9.2	Installing a New LSF for HP AlphaServer SC	9-3
9.2.1	Before You Install.	9-3
9.2.2	Installing LSF for HP AlphaServer SC	9-3
9.2.2.1	Example install.config	9-5
9.2.3	Setting up LSF hosts (hostsetup)	9-8
9.2.3.1	What hostsetup does	9-8
9.2.3.2	Running hostsetup	9-9
9.2.4	Next Steps	9-10
9.3	Configuring LSF for HP AlphaServer SC and Starting the LSF	9-10
9.3.1	Configuring the LSF hosts File for Multiple Network Interfaces	9-10
9.3.1.1	Configuring LSF_SERVER_HOSTS	9-11
9.3.1.2	When to Update the Hosts File	9-12
9.3.2	Starting LSF on the SC Cluster.	9-12
9.3.3	Verifying that the Configuration is Correct	9-13
9.3.3.1	Sample lsload -l output	9-14
9.3.3.2	Sample bhosts output	9-14

10 Post-Installation Tasks

10.1	State of System Immediately After Installation	10-2
10.2	General Post-Installation Tasks	10-3
10.2.1	Restrict Access to Login Nodes	10-3
10.2.2	Back Up and Restore the Management Server Root (/), /usr, and /var File Systems	10-4
10.2.2.1	Backup the Management Server	10-4
10.2.2.2	Restore the Management Server	10-6
10.2.3	Back Up the Cluster Root (/), /usr, and /var File Systems	10-7
10.2.4	Implement Security Recommendations	10-10
10.3	Console Network	10-10
10.4	Storage and File Systems	10-11
10.5	User Administration	10-12
10.6	Cluster Aliases	10-13
10.7	RMS	10-13
10.7.1	Mandatory RMS Administration Tasks	10-13
10.7.2	Optional RMS Administration Tasks	10-14

11 Troubleshooting

11.1	Tips for Installing an HP AlphaServer SC System	11-2
11.2	InstallSC Errors	11-5
11.3	Interpreting Problems During Software Installation	11-5
11.3.1	Uninstalled State	11-7

11.3.2	UNIX_Installed State	11-8
11.3.3	UNIX_Config State	11-9
11.3.4	UNIX_Patched State	11-10
11.3.5	SC_Installed State	11-12
11.3.6	SC_Patched State.	11-12
11.3.7	NHD_Installed State.	11-12
11.3.8	CLU_Create State	11-14
11.3.9	CLU_Added State	11-15
11.3.10	Bootp_Loaded State	11-16
11.3.11	Member_Added State.	11-17
11.4	Boot Errors	11-17
11.4.1	RIS Boot Failures	11-18
11.4.2	RIS Boot Reports a Bootstrap Failure.	11-19
11.4.3	RIS Boot Failure: Cannot Determine RIS Home Directory.	11-19
11.4.4	Failed to RIS Boot a Node When Attempting to Add It to the Cluster.	11-20
11.4.5	Cannot Communicate with the CAA Daemon	11-21
11.4.6	New Member Fails to RIS Boot	11-21
11.4.7	New Member Fails After RIS Boot	11-21
11.4.8	RIS: Boot Error.	11-22
11.4.9	sra boot Has Timeout Error	11-22
11.4.10	Elan Error During Node Boot.. . . .	11-23
11.4.11	Error During Node Boot	11-23
11.4.12	Node Hang During RIS Boot	11-24
11.4.13	RIS: File Open Failure for BOOTP	11-24
11.4.14	RIS: Boot Failure Access Violation	11-24
11.5	Adding RIS Client With No Default Route	11-25
11.6	Corrupt .member.list File Causes Core Dumps	11-25
11.7	Adding a New Member Fails	11-26
11.8	sra setup Errors	11-26
11.8.1	The Node is Not at the SRM Prompt.	11-26
11.8.2	The Console Logger is Misconfigured	11-27
11.9	Terminal Server Errors	11-27
11.9.1	Terminal Server Configuration Has Changed	11-27
11.9.2	Terminal Server Refuses a Connection	11-27
11.10	rinfo Command Displays UID or Wrong User Name	11-28
11.11	How to Drop and Rebuild the RMS Database	11-29
11.12	rcontrol Reports Errors During Node Boot	11-29
11.13	rcontrol Reports Error When Starting Partitions	11-30
11.14	clu_get_info Prints CONFIGURATION_ERROR.	11-30
11.15	How to Powercycle an HP AlphaServer SC System	11-30
11.15.1	Shutting Down	11-31
11.15.2	Powering Off	11-31
11.15.3	Powering On	11-31
11.15.4	Starting Up.	11-32

11.16	Error When Installing the Elan Subset on a Management Server.	11-32
11.17	Database Access Denied Errors	11-32
11.18	Database Access Denied Error on Some Domains	11-33
11.19	Diagnosing Federated Network Routing Problems	11-34
11.20	Wakeup on LAN (WOL) Problem	11-34
11.21	scfsmgr/pfsmgr report Could Not Open Socket	11-35
11.22	Problem with HSG Devices in Installation Process	11-35
11.23	Upgrade Errors	11-37
11.23.1	Upgrade Setup: RIS Host Appears to Be Invalid	11-37
11.23.2	Problem with dupatch Failure when Upgrading an HP AlphaServer SC System	11-37
11.23.3	Restarting a Failed Upgrade	11-38
11.23.4	Problem with Dependency on First Cluster	11-38
11.23.5	Problems when clu_quorum does not Complete	11-38
11.23.6	Upgrade Backup: node failed with cfs_find_drv_handle panic	11-39
11.23.7	Restoring from Backup Can Sometimes Fail	11-39
11.23.7.1	Incorrect Special Device Number for the Backup Disk	11-39
11.23.7.2	Persistent Reservations in the HSG for the backup disk	11-40
11.24	Interpreting Problems During Software Upgrade	11-40
11.24.1	Pre_Upgrade State	11-41
11.24.2	Upg_Installed State.	11-42
11.24.3	Checked State	11-43
11.24.4	Setup State	11-44
11.24.5	Installed State.	11-44
11.24.6	Upgraded State.	11-45
11.25	Increasing the Number of ptys	11-45
11.26	Increasing Socket Listen Queue Limits	11-46

Appendix A Installation Overview and Checklist: When the System Has a Management Server A-1

A.1	Installation Overview	A-2
-----	---------------------------------	-----

Appendix B Installation Overview and Checklist: When the System Does Not Have a Management Server B-1

B.1	Installation Overview	B-2
-----	---------------------------------	-----

Appendix C Checklist: Adding a Management Server to a Cluster C-1

C.1	Update SRA	C-3
C.2	Disable RIS on Node 0	C-6

C.2.1	Mandatory Tasks	C-6
C.2.2	Optional Tasks	C-6
C.3	Stopping the RMS System and mSQL	C-7
C.3.1	Manually Starting RMS	C-8
Appendix D Information Checklists		D-1
Appendix E Example Installation Output		E-1
E.1	sra setup	E-2
E.2	clu_create	E-20
E.3	clu_add_member	E-23
E.4	clu_quorum	E-25
E.5	upgrade_check	E-27
Appendix F Cluster-Related Messages in System Log Files		F-1
F.1	Startup Messages After Creating a Domain	F-2
F.2	Startup Messages After Adding a Domain Member	F-9
Appendix G Configuring Networker for HP AlphaServer SC		G-1
G.1	Domains as a Client of a Corporate Networker Server	G-2
G.1.1	Installing the Networker Client Software on the Cluster	G-2
G.1.2	Adding Domains as a Client of Corporate Networker Server	G-3
G.2	Domains as the Networker Server	G-5
G.3	Sample Output	G-8
Appendix H Cluster-Aliases and External Networks in the 10.x.x.x Range		H-1
Appendix I Configuring DNS Servers		I-1
I.1	Management Server as DNS Server and Cluster as Client	I-2
I.1.1	On the Management Server	I-2
I.1.2	On the Cluster	I-3
I.2	Management Server as Master DNS Server and Cluster as Slave DNS Server	I-5
I.2.1	On the Management Server	I-5
I.2.2	On the Cluster	I-5

I.2.3	Test the Slave DNS Server	I-7
I.3	Other Configurations.	I-8

Appendix J Configuring MSA1000 J-1

J.1	Preparing to Upgrade the MSA1000 Firmware	J-2
J.2	Upgrading MSA1000 Firmware	J-3
J.3	MSA1000 Sample Configuration	J-7

Index

List of Figures

Figure 2–1: Disk Layout in an HP AlphaServer SC Domain	2–9
Figure 2–2: Disk Layout in an HP AlphaServer DS20L Domain	2–10
Figure 2–3: Recommended Internal Storage Partitions	2–11
Figure 3–1: Single Management Server	3–2
Figure 3–2: Dual Management Server	3–3
Figure 3–3: HP AlphaServer SC Configuration for a 16-Node System	3–5
Figure 3–4: Node Network Connection When Using an HP AlphaServer SC 16-Port Switch	3–6
Figure 3–5: Node Network Connection When Using an HP AlphaServer SC 128-Port Switch	3–7
Figure 3–6: Node Network Connections: HP AlphaServer SC 16-Port Switch, HP AlphaServer DS20L Nodes	3–8
Figure 3–7: Node Network Connections: HP AlphaServer SC 128-Port Switch, HP AlphaServer DS20L Nodes	3–9
Figure 3–8: Federated HP AlphaServer SC Interconnect Configuration	3–10
Figure 3–9: Example System Storage Configuration — Cabling	3–28
Figure 3–10: Example System Storage Configuration — RAID Storage Units	3–29
Figure 3–11: Block Diagram of HSV110 Component Connections	3–45
Figure 3–12: Example System Storage Configuration — MSA1000	3–63
Figure 5–1: Tru64 UNIX Custom Setup Menu	5–10
Figure 5–2: Tru64 UNIX Network Setup Wizard Menu	5–14
Figure 6–1: Tru64 UNIX Custom Setup Menu	6–8
Figure 6–2: Tru64 UNIX Network Setup Wizard Menu	6–12
Figure 7–1: Status of Commands and Nodes	7–6
Figure 7–2: Log Menu Items	7–8
Figure 7–3: Sample Log File	7–9
Figure 7–4: Console Menu Items	7–10
Figure 7–5: Console File Details	7–10
Figure 7–6: Filtering Options Based on Status	7–11
Figure 7–7: Display Options Based on Nodes	7–12

List of Tables

Table 0–1: Abbreviations	xxv
Table 0–2: Documentation Conventions	xxvii
Table 0–3: HP-Specific Names and Part Numbers for Quadrics Components	xxviii
Table 0–4: Network Adapters and Device Names	xxix
Table 1–1: HP AlphaServer SC Subset Contents	1–4
Table 1–2: HP AlphaServer SC Disk Space Requirements	1–5
Table 2–1: HP AlphaServer SC IP Addresses	2–5
Table 2–2: Domain Naming Scheme	2–6
Table 2–3: Example (Not Default) Local Storage Layout	2–12
Table 3–1: How to Connect the Components of an HP AlphaServer SC System	3–4
Table 3–2: Populating the HP AlphaServer ES40 PCI Slots	3–17
Table 3–3: Populating the HP AlphaServer ES45 PCI Slots	3–18
Table 3–4: Populating the HP AlphaServer DS20L PCI Slots	3–20
Table 3–5: Minimum System Driver and Firmware Versions	3–24
Table 3–6: Identifying the WWN of the HBAs — HSG80	3–33
Table 3–7: Identifying the Connections from the HSG80 Controllers to atlas0	3–34
Table 3–8: Renaming the Connections — HSG80	3–35
Table 3–9: Creating the Mirrorsets	3–38
Table 3–10: Initializing the Storagesets — HSG80	3–38
Table 3–11: Introducing the Storage Units (Virtual Disks) — HSG80	3–39
Table 3–12: Setting the Identifiers — HSG80	3–40
Table 3–13: Disabling Access to the Units	3–41
Table 3–14: Enabling Access to the Units for Connected Nodes — HSG80	3–42
Table 3–15: Allocating RAID Storagesets — HSG80	3–43
Table 3–16: Storage Requirements for a Single HP AlphaServer SC domain — HSV110	3–49
Table 3–17: Storage Requirement for the Clustered Management Server — HSV110	3–49
Table 3–18: SC45 Minimum Storage Requirements — HSV110	3–50
Table 3–19: SC20 Minimum Storage Requirements — HSV110	3–50
Table 3–20: Disk Group Virtual RAID1 Available Space per Number of Disks — HSV110	3–50
Table 3–21: Summary Allocation of Disks on the HSV110	3–51
Table 3–22: World Wide Node and Port Names — HSV110	3–57
Table 3–23: Introducing the Storage Units (Virtual Disks) — MSA1000	3–65
Table 3–24: Setting the Identifiers — MSA1000	3–65
Table 3–25: Identifying the WWN of the HBAs — MSA1000	3–66
Table 3–26: Identifying the WWN of the HBAs — MSA1000 — Management Server	3–67
Table 3–27: Renaming the Connections — MSA1000	3–69
Table 3–28: Enabling Access to the Units for Connected Nodes — MSA1000	3–69
Table 4–1: Upgrade States	4–4

Table 4–2: Upgrade Time Estimates	4–8
Table 5–1: Setting the Console Variables	5–3
Table 5–2: Minimum System Firmware Versions	5–4
Table 5–3: Recommended Partition Layout for Management Server System Disk	5–8
Table 5–4: Minimum Kernel Options	5–9
Table 5–5: Network Interface Cards on a Management Server	5–15
Table 5–6: Routing Services	5–15
Table 5–7: Hosts File When Configuring a Management Server	5–16
Table 5–8: Configuration of New Member Network Interfaces	5–29
Table 5–9: Recommended Boot Disk Partition Configuration	5–40
Table 5–10: SRM Console Variables	5–42
Table 5–11: Minimum System Firmware Versions	5–46
Table 6–1: Setting the Console Variables	6–3
Table 6–2: Minimum System Firmware Versions	6–4
Table 6–3: Recommended Partition Layout for Node 0 System Disk	6–6
Table 6–4: Minimum Kernel Options	6–7
Table 6–5: Network Interface Cards on Node 0	6–12
Table 6–6: Routing Services	6–13
Table 6–7: Hosts File When Configuring Node 0	6–13
Table 6–8: Recommended Boot Disk Partition Configuration	6–27
Table 6–9: SRM Console Variables	6–29
Table 6–10: Minimum System Firmware Versions	6–31
Table 7–1: Installation States	7–3
Table 7–2: Status of Commands and Nodes Menu Items	7–6
Table 7–3: Status of Commands and Nodes Window	7–7
Table 9–1: Required install.config Variables	9–4
Table 9–2: Variables that Require an Absolute Path	9–5
Table 11–1: Manually Adjusting ptys	11–45
Table A–1: Installation Process: When the System Has a Management Server	A–2
Table B–1: Installation Process: When the System Does Not Have a Management Server	B–2
Table C–1: Installation Checklist	C–1
Table D–1: Summit Switch and Terminal Server Port Numbers	D–1
Table D–2: Tru64 UNIX System Attributes	D–1
Table D–3: Domain Attributes	D–2
Table D–4: Member Attributes of Domain D0	D–3
Table D–5: Member Attributes of Domain D1	D–4
Table D–6: Member Attributes of Domain D2	D–5
Table D–7: Member Attributes of Domain D32	D–6
Table E–1: Upgrade Check Output	E–27
Table H–1: Alternative HP AlphaServer SC IP Addresses	H–1

Preface

Purpose of this Guide

This document describes how to install an AlphaServer SC system from the Hewlett-Packard Company ("HP").

Intended Audience

This document is for those who install and set up HP AlphaServer SC systems. Some sections will be helpful to end-users. Before starting any installation, you must:

- Read the current version of the *HP AlphaServer SC Release Notes*, particularly all sections relating to installation.
- Understand how to load and unload the installation media and know which disks are needed during the installation.
- Have a basic understanding of the file system and commands.

New and Changed Features

This section describes the changes in this guide since Version 2.6 (UK1).

Changed Information

The following chapters have been revised:

- Chapter 1: Installation Overview
- Chapter 2: Pre-Installation Planning
- Chapter 3: Physical Installation
- Chapter 4: Upgrade Installation Procedure
- Chapter 5: Installing: When the System Has a Management Server
- Chapter 6: Installing: When the System Does Not Have a Management Server

- Chapter 7: Building the Domains
- Chapter 8: Completing the Installation
- Chapter 9: Installing LSF for HP AlphaServer SC
- Chapter 11: Troubleshooting
- Appendix E: Example Installation Output

Structure of This Guide

This document is organized as follows:

- Chapter 1: Installation Overview
- Chapter 2: Pre-Installation Planning
- Chapter 3: Physical Installation
- Chapter 4: Upgrade Installation Procedure
- Chapter 5: Installing: When the System Has a Management Server
- Chapter 6: Installing: When the System Does Not Have a Management Server
- Chapter 7: Building the Domains
- Chapter 8: Completing the Installation
- Chapter 9: Installing LSF for HP AlphaServer SC
- Chapter 10: Post-Installation Tasks
- Chapter 11: Troubleshooting
- Appendix A: Installation Overview and Checklist: When the System Has a Management Server
- Appendix B: Installation Overview and Checklist: When the System Does Not Have a Management Server
- Appendix C: Checklist: Adding a Management Server to a Cluster
- Appendix D: Information Checklists
- Appendix E: Example Installation Output
- Appendix F: Cluster-Related Messages in System Log Files
- Appendix G: Configuring Networker for HP AlphaServer SC
- Appendix H: Cluster-Aliases and External Networks in the 10.x.x.x Range

- Appendix I: Configuring DNS Servers
- Appendix J: Configuring MSA1000

Terms Used in This Guide

Clusters of nodes within an HP AlphaServer SC system are called domains (except in the case of clustered management servers). However, as domain management in HP AlphaServer SC systems is based on cluster management in TruCluster Server, many of the terms used in domain management include the word “cluster” — for example, cluster alias, cluster quorum, cluster application availability (CAA), cluster root, cluster disk.

The term “LSF cluster” is used to describe a group of machines controlled by LSF. An LSF cluster can consist of one or more HP AlphaServer SC systems, or part of an HP AlphaServer SC; it can also include additional machines of arbitrary architectures.

TruCluster Server Documentation

The HP TruCluster Server documentation set provides a wealth of information about clusters and cluster management; however, there are a number of differences between the management of domains in an HP AlphaServer SC system, and the management of TruCluster Server clusters. These differences are described in the *HP AlphaServer SC System Administration Guide*.

You should use the HP TruCluster Server documentation set to supplement the HP AlphaServer SC documentation set — if there is a conflict of information, use the instructions provided in the HP AlphaServer SC documentation set.

Related Documentation

You should have a hard copy or soft copy of the following documents:

- *HP AlphaServer SC Release Notes*
- *HP AlphaServer SC System Administration Guide*
- *HP AlphaServer ES40 Owner’s Guide*
- *HP AlphaServer ES45 Owner’s Guide*
- *HP AlphaServer DS20L*
- *HP StorageWorks HSG80 Array Controller CLI Reference Guide*
- *HP StorageWorks HSG80 Array Controller Configuration Guide*
- *HP StorageWorks Fibre Channel Storage Switch User’s Guide*

- *TruCluster Server Technical Overview*
- *TruCluster Server Hardware Configuration*
- *TruCluster Server Release Notes*
- *TruCluster Server Cluster Installation*
- *HP Tru64 UNIX Installation Guide*
- *HP Tru64 UNIX Network Administration*
- *HP Tru64 UNIX System Administration*
- *HP Tru64 UNIX Software License Management*
- *HP Tru64 UNIX System Configuration and Tuning*
- *HP AlphaServer SC User Guide*
- *HP AlphaServer SC RMS Reference Manual*
- *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*
- *HP AlphaServer SC Platform LSF® Reference*
- *Administering HP AlphaServer SC Platform LSF®*
- *HP AlphaServer SC Platform LSF® Reference*
- *Running Jobs with HP AlphaServer SC Platform LSF®*
- *HP AlphaServer SC Installing and Configuring SCIP*
- *HP Management and Configuration Guide for the HP ProCurve Series 4100GL Switches, Series 2600 Switches, and Switch 6108*
- *Summit Hardware Installation Guide, Extreme Networks, Inc.*
- *ExtremeWare Software User Guide, Extreme Networks, Inc.*

Abbreviations

Table 0–1 lists the abbreviations that are used in this document.

Table 0–1 Abbreviations

Abbreviation	Description
AdvFS	Advanced File System
ATM	Asynchronous Transfer Mode
AUI	Attachment Unit Interface
BIND	Berkeley Internet Name Domain
CD-ROM	Compact Disc — Read-Only Memory
CDSL	Context-Dependent Symbolic Link
CFS	Cluster File System
CPU	Central Processing Unit
DNS	Domain Name Service
FastFD	Fast, Full Duplex
FC	Fibre Channel
FDDI	Fiber-optic Digital Data Interface
GBIC	Gigabit Interface Connector
HiPPI	High-Performance Parallel Interface
IP	Internet Protocol
JBOD	Just a Bunch Of Disks
LIM	Load Information Manager
LSF	Load Sharing Facility
LSM	Logical Storage Manager
MAU	Multiple Access Unit
MPI	Message Passing Interface
NFS	Network File System
NIS	Network Information Service

Table 0–1 Abbreviations

Abbreviation	Description
NTP	Network Time Protocol
PAK	Product Authorization Key
PCMCIA	Personal Computer Memory Card International Association
PFS	Parallel File System
RAID	Redundant Array of Independent Disks
RIS	Remote Installation Services
RMS	Resource Management System
SC	SuperComputer
SCSI	Small Computer System Interface
SMP	Symmetric Multiprocessing
UTP	Unshielded Twisted Pair
WOL	Wakeup On LAN
WWID	WorldWide Identification

Documentation Conventions

Table 0–2 lists the documentation conventions that are used in this document.

Table 0–2 Documentation Conventions

Convention	Description
atlas	atlas is an example system name.
%	A percent sign represents the C shell system prompt.
\$	A dollar sign represents the system prompt for the Bourne and Korn shells.
#	A number sign represents the superuser prompt.
P00>>>	A P00>>> sign represents the SRM console prompt.
Monospace type	Monospace type indicates file names, commands, system output, and user input.
Boldface type	Boldface type indicates the first occurrence of a special term, or a complete sentence of important information. Boldface type in interactive examples indicates typed user input.
<i>Italic type</i>	Italic (slanted) type indicates emphasis, variable values, placeholders, menu options, function argument names, HP trademarks, and complete titles of documents and CD-ROMs.
<u>Underlined type</u>	Underlined type emphasizes important information.
[]	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
“ ”	Quotation marks are used to highlight words or phrases that have a specific meaning or significance in the context in which they are used.
...	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
:	A vertical ellipsis indicates that a portion of an example that would normally be present is not shown.
cat(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat (1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
Ctrl/x	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash.
Note	A note contains information that is of special importance to the reader.

HP-Specific Names and Part Numbers for Quadrics Components

Several HP AlphaServer SC Interconnect components are created by Quadrics. HP documents refer to Quadrics components using HP-specific names. Several Quadrics components also have a (different) Quadrics name. Table 0–3 shows how the HP-specific names and part numbers map to the equivalent Quadrics names.

Table 0–3 HP-Specific Names and Part Numbers for Quadrics Components

HP Part#	HP Name	Quadrics Name
3X-CCNBA-AA	HP AlphaServer SC 16-Port Switch	QM-S16
3X-CCNXA-BA	HP AlphaServer SC 128-Way Switch (new-type)	QM-S128F ¹
3X-CCNXE-CA	HP AlphaServer SC Top-Level Switch	QM-S128F ¹
3X-CCNXA-CA	HP AlphaServer SC Node-Level Switch	QM-S128F ¹
3X-CCNXA-AA	HP AlphaServer SC 128-Way Switch (old-type)	QM-S128
3X-CCNNA-AA	HP AlphaServer SC Elan Adapter Card	QM-400
3X-CCNXF-BA	HP AlphaServer SC 16-Port Switch Card (new-type)	QM-401X ²
3X-CCNXF-AA	HP AlphaServer SC 16-Port Switch Card (old-type)	QM-401X ²
3X-CCNXR-AA	HP AlphaServer SC High-Level Switch Card	QM-402
3X-CCNCR-BA	HP AlphaServer SC Clock Card (new-type)	QM-408
3X-CCNCR-AA	HP AlphaServer SC Clock Card (old-type)	QM-403
3X-CCNXN-AA	HP AlphaServer SC 16-Link Null Card	QM-407
3X-CCNXP-AA	HP AlphaServer SC Interconnect Control Processor	QM-410

¹The Quadrics part number QM-S128F corresponds to several components. The Quadrics part number refers to the basic empty chassis. Use the HP part numbers to distinguish between the different ways in which the chassis may be populated.

²The Quadrics part number QM-401X was not updated when this component was updated. Use the HP part numbers to distinguish between the new-type and old-type versions of this component.

Supported Network Adapters

Table 0–4 lists the associated device names for each supported network adapter. The examples in this guide refer to the DE602 network adapter.

Table 0–4 Network Adapters and Device Names

Network Adapter	SRM Device Name	UNIX Device Name
DE60x	eia0	ee0
DE50x	ewa0	tu0
Gigabit Ethernet	SRM cannot use this device	alt0
Gigabit Ethernet	SRM cannot use this device	bcm0
HiPPI ^{1, 2}	SRM cannot use this device	hip0
ATM ²	SRM cannot use this device	lis0
FDDI	SRM cannot use this device	fta0

¹HiPPI is only available if you install an additional HiPPI subset — for Tru64 UNIX Version 5.1B, the minimum supported version is HiPPI kit 222.

²The `sra install` command does not configure HiPPI and ATM interfaces — you must configure such interfaces manually.

Supported Node Types

HP AlphaServer SC Version 2.6 Update Kit 2 (UK2) supports the following node types:

- HP AlphaServer ES45
- HP AlphaServer ES40
- HP AlphaServer DS20L

Multiple Domains

The example system described in this document is a 1024-node system, with 32 nodes in each of 32 domains. Therefore, the first node in each domain is Node 0, Node 32, Node 64, Node 96, and so on. If you wish to set up a different configuration, substitute the appropriate node name(s) for Node 32, Node 64, and Node 96 in this guide.

Note:

HP AlphaServer SC Version 2.6 (UK2) supports large systems with up to 128 domains. For example, a 4096-node system may have 32 nodes in each of 128 domains. In this case, the domains are D0 to D127, and the naming convention used is as follows: atlasD0, atlasD1 and so on.

For information about the domain types supported in HP AlphaServer SC Version 2.6 (UK2), see Section 2.3.2 on page 2–4.

Location of Code Examples

Code examples are located in the `/Examples` directory of the *HP AlphaServer SC System Software* CD-ROM.

Location of Online Documentation

Online documentation is located in the `/docs` directory of the *HP AlphaServer SC System Software* CD-ROM.

Comments on this Document

HP welcomes any comments and suggestions that you have on this document. Please send all comments and suggestions to your HP Customer Support representative.

Installation Overview

This manual describes how to install and configure an HP AlphaServer SC system. Before you begin the installation, please read this manual. Becoming familiar with the general sequence of installation steps can save time and prevent problems later.

Note:

The procedures in this guide assume that each system's hardware and firmware are installed and configured as described in the following manuals:

- *HP AlphaServer ES40 Owner's Guide, HP AlphaServer ES45 Owner's Guide, HP AlphaServer DS20L Owner's Guide*
- *Summit Hardware Installation Guide*
- *HP Management and Configuration Guide for the HP ProCurve Series 4100GL Switches, Series 2600 Switches, and Switch 6108*
- *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*

Do not begin the software installation until the hardware is installed and configured.

Read the *HP AlphaServer SC Release Notes* before installing the HP AlphaServer SC software. You should also read the *HP TruCluster Server Technical Overview* to become familiar with the HP AlphaServer SC architecture and terminology. Chapter 1 of the *HP AlphaServer SC System Administration Guide* describes the main components of an HP AlphaServer SC system.

This chapter provides the following information:

- Installation Overview (see Section 1.1 on page 1–2)
- License Requirements (see Section 1.2 on page 1–3)
- HP AlphaServer SC Subsets (see Section 1.3 on page 1–4)

Installation Overview

- Base System Subsets (see Section 1.4 on page 1–6)
- General Considerations (see Section 1.5 on page 1–7)

1.1 Installation Overview

The software installation process involves the following steps:

1. If this is a full installation, plan the installation and ask the site network administrator for the required IP information. This process is described in Chapter 2.

If you are upgrading from an earlier release, plan the upgrade. This process is described in Chapter 4.

2. Install and configure the software, and set up the domains, as follows:
 - a. Manually install and configure the Tru64 UNIX operating system software — on either the management server or Node 0.¹
 - b. Configure either the management server or Node 0¹ as a Remote Installation Services (RIS) server for the system, using the `ris` command.
 - c. Manually install and configure the HP AlphaServer SC system software — on either the management server or Node 0.¹
 - d. Set up the SC database — on either the management server or Node 0¹ — using the `sra setup` command.
 - e. Check the state of all of the other nodes in the HP AlphaServer SC system, by running commands (`sra diag`, `sra command`, `sra ethercheck`, and `sra elancheck`) on either the management server or Node 0¹.
 - f. Install and configure the software on the system, by running the `sra install` command to create a fully installed domain. The `sra install` command automates the following steps:
 - RIS-boots the specified nodes from the RIS server (see step b above)
 - Installs and configures the Tru64 UNIX operating system software and the HP AlphaServer SC system software on these nodes.
 - Copies the license database (including TruCluster and ASC licences) to the first node of each domain.
 - Transforms each specified single-node system into a single-node domain.
 - Adds new members to the domains.
 - Performs additional configuration tasks on the specified nodes.

This process is described in Chapter 7.

-
1. Perform this task on the management server, if used. If not using a management server, perform this task on Node 0.

3. Perform the final installation tasks: configure the external network interfaces, improve domain availability by adding votes, configure the RMS database on Node 0 or on the management server if used, create the RMS partitions. Then run the example MPI program.

This process is described in Chapter 8.

Once hardware setup is complete, it typically takes up to two working days to install HP AlphaServer SC System Software Version 2.6 (UK2) on a 1024-node system. A detailed overview of the complete installation process is provided in the following appendixes:

- Appendix A (*Installation Overview and Checklist: When the System Has a Management Server*)
- Appendix B (*Installation Overview and Checklist: When the System Does Not Have a Management Server*)

Each of these appendixes also provides an installation checklist to help you to complete all of the installation tasks in the correct order.

1.2 License Requirements

There are no system-wide licenses. Each HP AlphaServer SC member must have its own licenses installed. For the first member, after installing and licensing Tru64 UNIX, load and register an HP AlphaServer SC license and TruCluster license (TCS-UA) to create a single-member cluster.

Each time you add an additional member to the HP AlphaServer SC system, the entire license database (including the TCS-UA and ASC licenses) is propagated from the first member of the domain (Node 0, 32, 64 and so on) to the new member.

Note:

The `sra install` command propagates the license database.

After domain creation, any additional required application licenses must be registered on each member that will run the application.

Note:

You can boot a system that does not have an HP AlphaServer SC license. The system joins the domain and boots to multiuser mode, but only the `root` user can log in (with a maximum of two users).

HP AlphaServer SC Subsets

The system displays a license error message reminding you to load the license. This policy enforces license checks while making it possible to boot, license, and repair a system during an emergency.

For more information about registering licenses, see Section 5.1.5.1 on page 5–10 or Section 6.1.4.1 on page 6–8, and the *Tru64 UNIX Software License Management* manual.

1.3 HP AlphaServer SC Subsets

In addition to the operating system software (that is, Tru64 UNIX), several HP AlphaServer SC subsets are installed. Table 1–1 lists the HP AlphaServer SC subsets installed by the HP AlphaServer SC installation tools.

Table 1–1 HP AlphaServer SC Subset Contents

Subset Name	Description
SRAOSFPATCH320	HP AlphaServer SC patch software for Tru64 UNIX software
SRATCRBASE320	HP AlphaServer SC patch software for TruCluster Server software
ELNMOD340	HP AlphaServer SC Interconnect software
JTGMOD320	JTAG Switch Management software
SRABASE320	HP AlphaServer SC Installation Utility (SRA) software
SRATCL320	
SRAGXEXEC320	HP AlphaServer SC Global Execution
SRACFENGINE320	HP AlphaServer SC cfengine
SRAUPG320	HP AlphaServer SC upgrade software
TCRBASE540	TruCluster Server software
TCRMAN540	
TCRMIGRATE540	
SRASYSMAN320	HP AlphaServer SC Sysman software
SCFSBASE320	HP AlphaServer SC File System (SCFS) software
SCFSBIN320	
QSWNETDIAGS340	HP AlphaServer SC Interconnect diagnostic software
SRAMPI320	HP AlphaServer SC Message Passing Interface (MPI) software
PPMBASE320	Performance Visualizer performance monitoring software

Table 1–1 HP AlphaServer SC Subset Contents

Subset Name	Description
RMSBASE320 RMSD320 RMSELAN320 RMSLSFSUP320 RMSMAN320 RMSPANDORA320	HP AlphaServer SC Resource Management System (RMS) software
SWMSUP120	HP AlphaServer SC Interconnect Controller Switch Management software
PFSMOD320	HP AlphaServer SC Parallel File System (PFS) software
MSQLSUP340	HP AlphaServer SC mSQL Support

Table 1–2 shows the approximate disk space requirements for each HP AlphaServer SC subset in the root (/), /usr, and /var file systems on the Tru64 UNIX system.

Table 1–2 HP AlphaServer SC Disk Space Requirements

Subset Name	Root File System (bytes)	/usr File System (bytes)	/var File System (bytes)	Total (bytes and MB)
SRAOSFPATCH320	67 918	0	0	67 918
SRATCRBASE320	194 537	831 173	0	1 025 710
ELNMOD340	0	3 095 769	0	3 095 769
JTGMOD320	0	372 648	0	372 648
SRABASE320	40 120	3 631 731	41 132	3 712 983
SRATCL320	0	3 745 906	0	3 745 906
SRAGXEXEC320	369	10 410 362	8 192	10 418 923
SRACFENGINE320	16 929	2 098 775	88 708	2 204 412
SRAUPG320	8 192	452 166	0	460 358
TCRBASE540	613 261	52 224 811	10 310 776	63 148 848
TCRMAN540	0	882 353	0	882 353
TCRMIGRATE540	0	1 069 348	0	1 069 348
SRASYSMAN320	31 030	4 926 406	61 795	5 019 231
SCFSBASE320	0	53 558	0	53 558
SCFSBIN320	26 934	1 595 245	0	1 622 179
QSWNETDIAGS340	0	5 422 376	0	5 422 376

Base System Subsets

Table 1–2 HP AlphaServer SC Disk Space Requirements

Subset Name	Root File System (bytes)	/usr File System (bytes)	/var File System (bytes)	Total (bytes and MB)
SRAMPI320	0	58 479 084	0	58 479 084
PPMBASE320	0	2 593 288	0	2 593 288
RMSBASE320	0	10 077 674	41 311	10 118 985
RMSD320	23 744	21 478 373	0	21 502 117
RMSELAN320	396 377	15 468 002	0	15 864 379
RMSLSFSUP320	0	914 372	0	914 372
RMSMAN320	0	540 150	0	540 150
RMSPANDORA320	0	4 334 055	0	4 334 055
SWMSUP120	1 990 800	5 091 388	0	7 082 188
PFSMOD320	537 394	1 214 495	0	1 751 889
MSQLSUP340	27 561	1 774 258	42 062	1 843 881

1.4 Base System Subsets

When you install Tru64 UNIX, you must select AdvFS as the file system type for the root (/), /usr, and /var file systems.

If your HP AlphaServer SC system will contain different types of file systems, make sure to load the optional Tru64 UNIX subsets needed to support different hardware configurations. For example, because keyboards and graphics cards require specific subsets in order to work properly, load all keyboard and font subsets.

Regardless of the types of file systems in the HP AlphaServer SC system, HP strongly recommends that, unless prohibited by site policy, you load all subsets from the base operating system media when installing the Tru64 UNIX system.

You may add different file systems at a later date, or you may install an application that has a dependency on a subset you did not install. You can add additional software subsets later, but, because you are dealing with only one system, it is easier to install them before you create a cluster.

Note:

If you are installing subsets from the Associated Product CD-ROMs, please do not install the TruCluster software contained, as HP AlphaServer SC uses its own version of the TruCluster software.

1.5 General Considerations

Note the following general installation considerations:

- If you build your own kernels, be aware that in an HP AlphaServer SC system, `/vmunix` is a context-dependent symbolic link (CDSL):

```
/vmunix -> cluster/members/{memb}/boot_partition/vmunix
```

Treat a CDSL as you would any other symbolic link: remember that copying a file follows the link, but moving a file replaces the link. If you were to move (instead of copy) a kernel to `/vmunix`, you would replace the symbolic link with the actual file.

The *HP TruCluster Server Technical Overview* describes CDSLs; Chapter 24 of the *HP AlphaServer SC System Administration Guide* provides information on using and repairing CDSLs.

- You must have at least 512MB of memory available on each member system.
- HP AlphaServer SC supports the UNIX® File System (UFS) as a read-only file system only.

Pre-Installation Planning

This chapter describes the activities that you should perform when planning the installation.

Note:

Use the checklist provided in Appendix A (if using a management server) or Appendix B (if not using a management server), to ensure that you complete all installation tasks in the correct order.

If adding a management server to an HP AlphaServer SC system after domain creation, use the checklist provided in Appendix C to ensure that you complete all installation tasks in the correct order.

The information in this chapter is structured as follows:

- Review the Release Notes (see Section 2.1 on page 2–2)
- Assign the External Network IP Addresses (see Section 2.2 on page 2–2)
- Assign the System Name and Default Cluster Aliases (see Section 2.3 on page 2–3)
- Plan the Local and Global Storage (see Section 2.4 on page 2–7)
- Choose the Root Password (see Section 2.5 on page 2–16)
- Record the External Gateway IP Address (see Section 2.6 on page 2–16)

Review the Release Notes

2.1 Review the Release Notes

Please review the Release Notes for Tru64 UNIX Version 5.1B and for Tru64 UNIX Version 5.1B-3, as well as those for HP AlphaServer SC Version 2.6 (UK2). These documents are located as follows:

- *Tru64 UNIX Version 5.1B Release Notes* are located on the *HP Tru64 UNIX Version 5.1B Documentation* CD-ROM. This document describes significant new and changed features in Version 5.1B of the Tru64 UNIX operating system, and lists features and interfaces scheduled for retirement in future releases.

For a description of the new and changed features of Tru64 UNIX Version 5.1B-3, refer to the *Tru64 UNIX Release Notes for Version 5.1B-3*, available at:

http://h30097.www3.hp.com/docs/base_doc/DOCUMENTATION/V51B-3/HTML/TITLE.HTM

- *HP AlphaServer SC Version 2.6 (UK2) Release Notes* are located in the `/docs` directory on the *HP AlphaServer SC System Software* CD-ROM. This document provides information on restrictions to the software and documentation for HP AlphaServer SC system software.

2.2 Assign the External Network IP Addresses

The process to assign the external network IP addresses is performed when you run the `sra setup` command as described in Section 5.2 and Section 6.2. The `sra setup` command automatically allocates an IP address and subnet mask to each external network interface — you must have at least two external network interfaces on each domain (one on the first two nodes of each domain: Nodes 0 and 1; Nodes 32 and 33; Nodes 64 and 65 and so on), and one on the management server (if used).

A typical scheme uses a base IP address for the external interface for node 0, and increments this IP address by 1 for the external interface to node 1, node 32, node 33, and so on. It also uses a different base IP address for the cluster alias for domain 0, and increments this IP address by 1 for each subsequent cluster alias. If you follow this scheme, the `sra setup` process automatically enters the IP addresses in the database, based on the base IP addresses that you enter. Otherwise, the `sra setup` process prompts you for each external node IP address and each cluster alias IP address in the system.

Ask your site network administrator for this information, and record it in Appendix D.

Note:

Use the checklist provided in Appendix D to record all site-specific information.

2.3 Assign the System Name and Default Cluster Aliases

When planning the system name, consider the future size of your system, as described in Section 2.3.1 on page 2–3.

The names of the default cluster aliases are derived from the system name. However, you must plan other attributes of the default cluster aliases, as described in Section 2.3.2 on page 2–4 and Section 2.3.3 on page 2–4.

The system name is the only name that you need to choose — all other names are then derived from the system name, as shown in Table 2–2 on page 2–6.

Note:

The system name, and the names of the default cluster aliases, must not end with a digit. The system name must not contain dashes or other forms of punctuation.

2.3.1 System Name

The system name must consist of letters and numbers only and it must not end in a number.

The HP AlphaServer SC installation process uses the system name to derive both the cluster alias names and the host names of each member in each domain.

For example, in a 64-node system with system name `atlas`, the cluster aliases are `atlasD0` and `atlasD1` while the node host names are `atlas0`, `atlas1`, ..., `atlas63`.

If you prefer the cluster alias naming convention used in previous releases, you can select this option during the installation process.

However, if the number of nodes is less than or equal to 32, and you therefore choose to have only one domain, the cluster alias name is the same as the system name — in the above example, the cluster alias name would be `atlas`. If you later increase the number of nodes to 33 or more, you must create a second domain (`atlasD1`), and rename the original cluster alias (from `atlas` to `atlasD0`).

If your system currently has 32 nodes or less but you plan to have 33 nodes or more, enter the proposed total number of nodes when prompted by the `sra setup` command (see either Section 5.2 on page 5–37 or Section 6.2 on page 6–24). The number of domains required for the proposed total system are then created, with the correct naming convention. Record the system name in Appendix D.

If the system has a management server, the name of the management server is free format. However, it is advisable to derive a management server name from the system name, for example, `atlasms` in a system called `atlas`.

Assign the System Name and Default Cluster Aliases

For systems with clustered management servers, HP recommends that the names of the individual management server nodes are of the format `atlasms0` and `atlasms1`, and the management server cluster alias is `atlasms`. The alias is then the name of the management server as defined in the `sra setup` dialog.

2.3.2 Domain Types

In HP AlphaServer SC Version 2.6 (UK2), there are two domain types:

- File-Serving (FS) domain
- Compute-Serving (CS) domain

HP AlphaServer SC Version 2.6 (UK2) supports a maximum of four FS domains.

An FS domain can be any domain. It is not mandatory to create an FS domain, but you will not be able to use SCFS if you have not done so. For more information about SCFS, see Chapter 7 of the *HP AlphaServer SC System Administration Guide*.

If you choose to have a small FS domain to maximize the number of nodes in the remaining CS domains, you should avoid gaps in the numbering sequence. The first node in the first CS domain should start directly after the last node in the FS domain. From a cabling perspective, the ports used on the HP AlphaServer SC Interconnect switch should also be contiguous.

If you intend to add more nodes to the FS domain later, it is acceptable to leave gaps in the numbering sequence. In such cases, you must leave corresponding gaps in the cabling sequence. You must also ensure that the total node count set in the SC database reflects the final node count; that is, that it includes those nodes not yet added.

2.3.3 Default Cluster Aliases

The **default cluster alias** is a name, with an associated IP address, that makes all of the systems in a domain look like a single system to the outside world. Depending on the number of nodes in your HP AlphaServer SC system, you need up to 32 default cluster aliases — that is, one for each domain. The default cluster alias is assigned during cluster installation, and all members of the domain automatically join the default cluster alias at boot time.

You must plan the following attributes of each default cluster alias:

- Name (the domain name) — the default value is derived from the system name (see Table 2–2 on page 2–6)
HP recommends that you accept the default values for the domain names. The default names are derived by appending the letter D and an incrementing number to the system name. For example, if the system name is `atlas`, assign the following names to the domains: `atlasD0`, `atlasD1`, and so on (see Section 5.2 on page 5–37 or Section 6.2 on page 6–24).

Assign the System Name and Default Cluster Aliases

- IP address — note that the cluster alias IP address should not be a 10 address. However, if your system setup requires you to use a 10 address, refer to Appendix G.
- Subnet mask

Note:

The domain naming scheme (atlasA, atlasB) used in earlier HP AlphaServer SC releases is supported in HP AlphaServer SC Version 2.6 (UK2).

Ask your site network administrator for the necessary information, and record it in Appendix D.

Table 2–1 lists the convention used to assign IP addresses in an HP AlphaServer SC system. Table 2–2 on page 2–6 shows the domain naming scheme.

Table 2–1 HP AlphaServer SC IP Addresses

Component	IP Address Range
Net mask	255.255.0.0
Cluster Interconnect (IP suffix: -ics0)	10.0.x.y
System Interconnect (IP suffix: -eip0)	10.64.x.y
Management network interface card	10.128.x.y
Terminal server <i>t</i> , where <i>t</i> is 1–254	10.128.100.t
Management server <i>m</i> on management LAN, where <i>m</i> is 1, or 2 for the second member of a clustered management server	10.128.101.m
Management server <i>m</i> Cluster Interconnect (IP suffix: -ics0)	10.32.0.m
Management server Gigabit Interconnect (IP suffix: -icstcp0)	10.33.0.m
Ethernet switch <i>g</i> , where <i>g</i> is 1–11	10.128.103.g
HP SANworks Management Appliance or Fibre Channel switch <i>f</i> , where <i>f</i> is 1–254	10.128.104.f
RAID array controller <i>a</i> , where <i>a</i> is 1, 2, and so on	10.128.105.a
HP AlphaServer SC Interconnect Control Card for Node-Level switch <i>N</i> , where <i>r</i> is the rail number and <i>N</i> is 0–31	10.128.(128+r).(N+1)
HP AlphaServer SC Interconnect Control Card for Top-Level switch <i>T</i> , where <i>r</i> is the rail number and <i>T</i> is 0–15	10.128.(128+r).(T+128)
Preferred server cluster alias addresses (where <i>FS</i> is 0–3, thus accommodating a maximum of four FS domains, and <i>y</i> is the member ID within the FS domain)	10.128.(106+FS).y

Assign the System Name and Default Cluster Aliases

Table 2–2 Domain Naming Scheme

Element	First Domain		Second Domain	
	Name	IP Address	Name	IP Address
Default Cluster Alias (domain name)	atlasD0	site-dependent	atlasD1	site-dependent
Nodes (node names on management network)	atlas0, atlas1, ..., atlas31	10.128.0.1, 10.128.0.2, ..., 10.128.0.32	atlas32, atlas33, ..., atlas63	10.128.0.33, 10.128.0.34, ..., 10.128.0.64
Cluster Interconnect	atlas0-ics0, ..., atlas31-ics0	10.0.0.1, ..., 10.0.0.32	atlas32-ics0, ..., atlas63-ics0	10.0.0.33, ..., 10.0.0.64
System Interconnect	atlas0-eip0, ..., atlas31-eip0	10.64.0.1, ..., 10.64.0.32	atlas32-eip0, ..., atlas63-eip0	10.64.0.33, ..., 10.64.0.64
External Network ¹	atlas0-ext1, ..., atlas31-ext1	site-dependent	atlas32-ext1, ..., atlas63-ext1	site-dependent
Terminal Server	atlas-tc1	10.128.100.1	atlas-tc2	10.128.100.2

Element	Third Domain		Last Domain	
	Name	IP Address	Name	IP Address
Default Cluster Alias (domain name)	atlasD2	site-dependent	atlasD31	site-dependent
Nodes (node names on management network)	atlas64, atlas65, ..., atlas95	10.128.0.65, 10.128.0.66, ..., 10.128.0.96	atlas992, atlas993, ..., atlas1023	10.128.7.97, 10.128.7.98, ..., 10.128.7.128
Cluster Interconnect	atlas64-ics0, ..., atlas95-ics0	10.0.0.65,..., 10.0.0.96	atlas992-ics0,..., atlas1023-ics0	10.0.7.97,..., 10.0.7.128
System Interconnect	atlas64-eip0, ..., atlas95-eip0	10.64.0.65,..., 10.64.0.96	atlas992-eip0, ..., atlas1023-eip0	10.64.7.97 ..., 10.64.7.128
External Network ¹	atlas64-ext1, ..., atlas95-ext1	site-dependent	atlas992-ext1,..., atlas1023-ext1	site-dependent
Terminal Server	atlas-tc3	10.128.100.3	atlas-tc32	10.128.100.32

¹The suffix is `ext1` for the first external network interface, `ext2` for the second, and so on.

2.4 Plan the Local and Global Storage

System storage is provided by a RAID subsystem, which is connected to the first two nodes in each domain, as shown in Figure 2–1 on page 2–9.

The amount of data storage, and the number of file-serving nodes deployed, is site-specific. The first two nodes in each domain must have external storage capability, to serve the system storage files. As system storage requires only a small amount of space, the external storage on these nodes can also hold user data.

HP AlphaServer SC permits three kinds of shared storage technology for the domain system storage: HSG80 from the Enterprise Storage Array (ESA) and Enterprise Modular Array (EMA) families; HSV110 from the Enterprise Virtual Array (EVA) family; and the Modular SAN Array 1000 (MSA1000). For more information about system storage and supported RAID products, see Chapter 6 of the *HP AlphaServer SC System Administration Guide*.

The HP StorageWorks Modular SAN Array 1000 (MSA1000) is the entry-level storage subsystem for HP AlphaServer SC systems and is only supported for systems that have only one domain. If you have more than one domain, it is more cost-effective to use an alternative type of storage subsystem. The MSA1000 is configured with two RAID controllers, with one controller in standby mode.

The HSG80 and HSV110 storage subsystems can be used in systems that have multiple domains. In the examples of HSG80 and HSV110 storage subsystems in this chapter, the 128-node HP AlphaServer SC system is divided into four domains. The first two nodes in each domain have two fibre channel Host Bus Adapter (HBA) cards each, and the storage is configured with two RAID controllers.

For 1024-node systems, the HSG80 or HSV110 system storage configuration is typically built from eight individual and independent 128-node storage blocks. Each 128-node storage block will have a separate RAID controller pair; therefore, every configuration step described in this chapter must be repeated on each 128-node storage block. As you progress through the eight individual blocks, the node numbers will increase but all other settings (disks, units, identifiers) remain unchanged.

If your system has a different configuration (different number of domains, nodes, or HBA cards per node), please modify the relevant commands accordingly. Note that the number of fibre channel switches configured does not affect the commands in this chapter.

Note:

HP AlphaServer SC Version 2.6 (UK2) supports a maximum of four FS domains. The SCFS file system exports file systems from an FS domain to the other domains. Although the FS domains can be located anywhere in the HP AlphaServer SC system, HP recommends that you configure either the first domain(s) or the last

Plan the Local and Global Storage

domain(s) as FS domains — this provides a contiguous range of CS nodes for MPI jobs. Contiguous nodes can take advantage of hardware broadcast on the data network.

To plan the storage, perform the following tasks:

- Confirm the Disk Layout (see Section 2.4.1 on page 2–9)
- Plan the Partition Sizing for Local Storage (see Section 2.4.2 on page 2–11)
- Review the System Storage Rules (see Section 2.4.3 on page 2–12)
- Plan the RAID Configuration (see Section 2.4.4 on page 2–13)

Note:

The diagrams in this section refer to File-Serving (FS) domains and Compute-Serving (CS) domains. Refer to Clustered Management Server (see Section 3.1.2 on page 3–3) for information on clustered management servers.

2.4.1 Confirm the Disk Layout

You must assign disks for both local and global storage in an HP AlphaServer SC system. Figure 2–1 shows the disk layout for the first domain, using the standard recommended configuration.

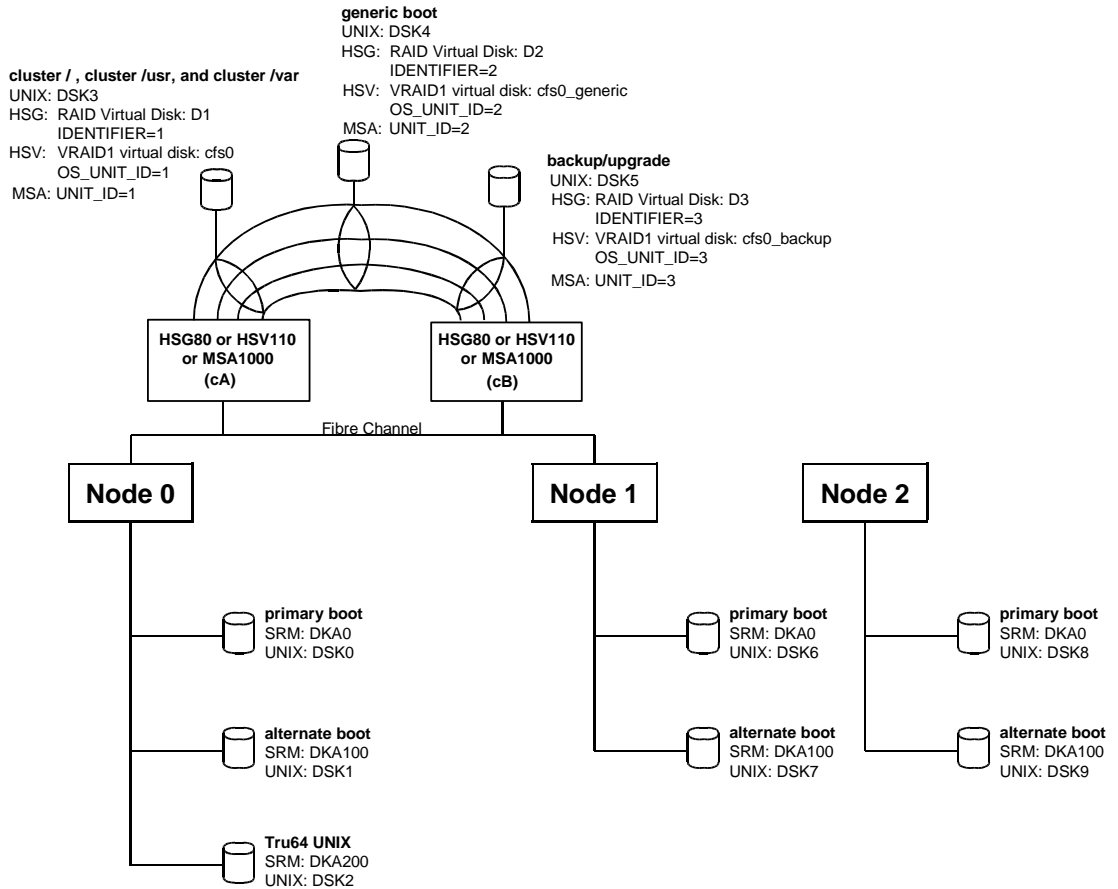


Figure 2–1 Disk Layout in an HP AlphaServer SC Domain

Plan the Local and Global Storage

Figure 2–2 shows the disk layout for the first domain for the HP AlphaServer DS20L, using the standard recommended configuration.

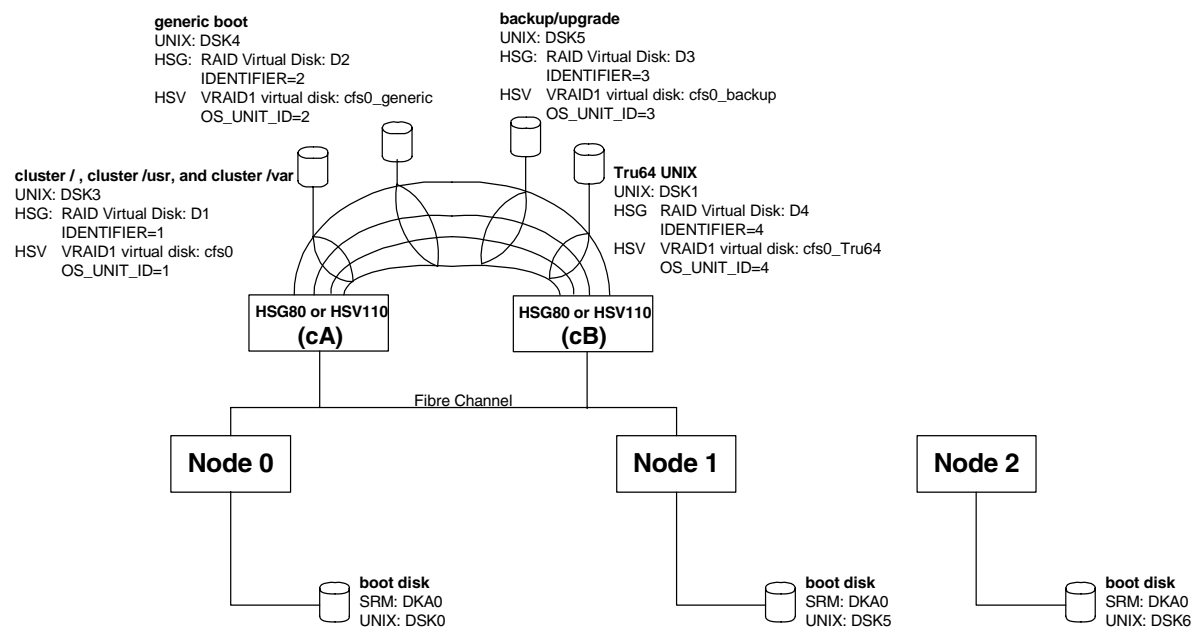


Figure 2–2 Disk Layout in an HP AlphaServer DS20L Domain

2.4.2 Plan the Partition Sizing for Local Storage

You should size the internal storage partitions to suit the configuration of the node and your temporary requirements. The system installation process suggests default sizes, which you can modify.

The operating system boot partition requires only 262MB.

The swap space partition size should be set to at least 1.25 times the size of the physical memory. This allows an entire job and system daemons to be swapped out in case of high priority or emergency work. For example, if the nodes that make up your HP AlphaServer SC system each have 4GB of physical memory, you should set the swap space to 5GB.

The local partition is mounted as `/cluster/members/memberM/local`. The installation process creates a CDSL that equates this to `/local` on every node. Similarly, the CDSL `/tmp` equates to `/cluster/members/memberM/tmp`.

The `/local` file system is used by RMS to hold job-specific files; for example, core dumps. The `/local` file system can also be used to hold component file systems for the parallel file system (PFS). Note, however, that `/local` is local to a node and, therefore, not highly available.

Figure 2–3 shows the recommended internal storage partitions for an HP AlphaServer SC system whose member nodes each have 4GB of physical memory.

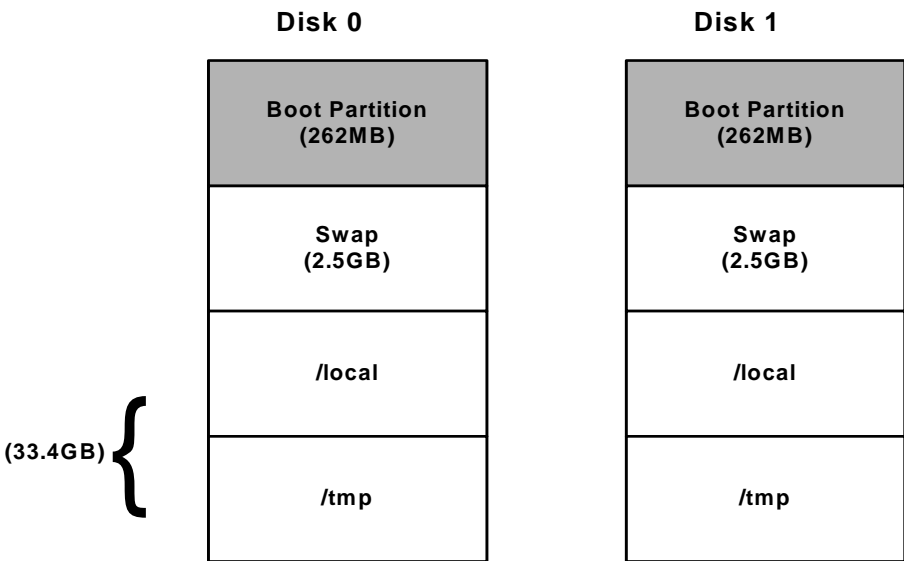


Figure 2–3 Recommended Internal Storage Partitions

Plan the Local and Global Storage

Local storage is configured at system installation time. The installation process suggests a default subdivision, which you can accept or modify. Any modifications become the default, and are applied to all nodes.

It is also possible to configure selected nodes with layouts that are different from the default layout, as shown in the example in Table 2–3.

Table 2–3 Example (Not Default) Local Storage Layout

Node	Swap	/local	/tmp
0 and 1	44% (16G)	28% (10G)	28% (10G)
2 to 31	28% (10G)	36% (13G)	36% (13G)

To configure an alternate layout, rerun the `sra setup` command and select the size to be used by the majority of the nodes (for example Node 2 to 31 in Table 2–3). Run the `sra edit` command to alter the sizes used by the remaining nodes (Node 0 and 1).

2.4.3 Review the System Storage Rules

System and data storage uses external RAID storage systems connected by one or more fibre channel switches. While there is considerable flexibility and many possible configurations, the system must obey the following storage rules so that the system software will operate correctly:

- All data storage must be connected to the FS domain(s).
- In all storage subsystems (HSG80, HSV110, and MSA1000), multiple-bus failover must be used, to improve availability and performance. (In MSA100 storage subsystems, the second bus is in standby mode.) All hosts must have operating-system software that supports multiple-bus failover mode.
- Each domain must have unique access to its own system storage devices (that is, disks) at a location in the storage infrastructure.
- The storage infrastructure that supports system storage must meet the following criteria:
 - The first two members of each domain must have a path to the storage.
 - The storage must be highly available; that is, dual HSG80 controllers, dual HSV110 controllers, or dual MSA1000 controllers, with each port of each controller connected to the Fibre Channel switch.

Note:

Dual redundant fibre fabric (that is, two fibre channel switches) can be configured, but is not required. For HP AlphaServer SC Version 2.6 (UK2), dual redundant fibre fabric is available as a *QuickSpec* option. Dual redundant fibre fabric is not available on systems composed of HP AlphaServer DS20L nodes (the HP AlphaServer SC20 product), as such nodes have only 2 PCI slots.

- Each system storage must have a configured spareset, to automatically source a replacement disk in the case of a failed Mirrorset member.
- System storage uses a combination of Mirrorsets and JBOD (just a bunch of disks), as described later in this chapter.

Note:

The criteria listed above apply to system storage only. Data storage can use any storage mechanism — RAIDset, Mirrorset, and so on — and can be configured in any way that satisfies customer requirements.

The SCFS file system is used to serve storage from the FS domain(s) to one or more CS domains.

System storage does not require a dedicated controller pair (and associated fibre channel switch). Instead, system storage disks may be placed anywhere in the available storage infrastructure. In effect, the CS domains can “piggyback” on the storage infrastructure that has been put in place to support the FS domains.

2.4.4 Plan the RAID Configuration

HP AlphaServer SC permits three kinds of shared storage technology for the domain system storage: HSG80 from the Enterprise Storage Array (ESA) and Enterprise Modular Array (EMA) families; HSV110 from the Enterprise Virtual Array (EVA) family; and the Modular SAN Array 1000 (MSA1000). Please consult the appropriate section to plan the RAID configuration for the appropriate technology.

- Planning the HSG80 RAID Configuration (see Section 2.4.4.1 on page 2–14)
- Planning the HSV110 RAID Configuration (see Section 2.4.4.2 on page 2–14)
- Planning the MSA1000 RAID Configuration (see Section 2.4.4.3 on page 2–15)

Plan the Local and Global Storage

2.4.4.1 Planning the HSG80 RAID Configuration

When creating an HSG80 RAID configuration, consider the following points:

- Number of disks
The smaller the number of disks (N) the higher the parity overhead (1/N). Using more than six disks, while economical in terms of parity overhead, will require two SCSI transfers to complete, because the RAID array only has six back-end SCSI buses. Therefore, the use of six disks will provide optimal performance.
- Number of RAID sets
The HSG80 gives optimal performance for sequential transfers with two concurrently active RAID sets, one per controller.
- Preferred controller for a unit
Units (the OS-visible equivalent of a storage set) should be balanced across the different controllers using the SET UNIT command. When the preferred controller has been set for the units, you must restart both controllers before the assignment will take effect.
- Maximum cache transfer size
When setting up a unit, set this parameter to a large number — the maximum value is 2048 blocks = 1MB. This allows the controller to aggregate data for a sequential stream in cache and to perform parity calculations in memory, avoiding read-modify-write cycles on the disk.

For more information, see the *HP StorageWorks HSG80 Array Controller CLI Reference Guide*.

2.4.4.2 Planning the HSV110 RAID Configuration

When planning an HSV110 RAID configuration, consider the following points:

- Virtual RAID (VRAID)
The HSV110 uses virtualized RAID sets instead of conventional RAID.
- Disk Groups
The HSV110 allows you to separate virtual disks into separate disk groups, distinct from and independent of the other disk groups in the subsystem. Disk group boundaries do allow some control over data placement and can be used to influence availability and performance. However, the benefits of virtualization are diminished by disk group boundaries.

Advantages:

- Creating virtual disks in different disk groups divides them into separate failure domains. In situations where the application retains two copies of the data, placing each copy in a separate disk group ensures that both copies of the data are not lost in case of failure.

Disadvantages:

- Capacity Utilization: separate disk groups create boundaries across which spare capacity cannot be used.
- Spare Capacity: the spare capacity used to recover redundancy in the event of disk failure cannot be shared between disk groups. The more disk groups that are created, the more spare capacity must be reserved.
- Performance: By creating different disk groups, you reduce the number of spindles to which a VRAID disk is distributed. This may impact performance.
- General Guidelines:
 - HP recommends that you do not attempt to calculate capacity requirements too finely. While virtualization improves capacity efficiency, the virtualization algorithms require free space. HP recommends maintaining a free capacity above 10%.
 - Create the least number of disk groups consistent with failure and performance isolation requirements.
 - Arrange disk groups vertically (visually).
 - Keep an even multiple of 2 for disk group size (for example 8, 12, 16)
 - Do not attempt to mix disk size and speed within a disk group. You can, however, mix disk sizes and speeds within different disk group sets.
 - Install disks from left to right. For system storage, add disks in pairs (that is, two at a time), and install half the disks on the top shelf and half on the bottom shelf.

2.4.4.3 Planning the MSA1000 RAID Configuration

The HP StorageWorks Modular SAN Array 1000 (MSA1000) storage controller is the entry-level storage subsystem for HP AlphaServer SC systems and is only supported for systems that have one domain. When planning an MSA1000 RAID configuration, consider the following points:

- Disk Capacity

The MSA1000 storage building block provides two 14-drive enclosures (dual Ultra-3 SCSI disk enclosures for a total of 28 disk drives on four Ultra-3 SCSI buses).
- Redundancy

The MSA1000 controllers work in active/standby mode. Only one controller is active at any time.
- Management

The MSA1000 is configured using the Command Line Interface (CLI) and through a serial cable. For more information, refer to Section 3.16 on page 3–62.

Choose the Root Password

2.4.4.4 Managing a Fibre Channel Switch

The most basic way of managing a StorageWorks Fibre Channel switch is by telnetting to the switch and using a command line interface. Please see the HP *StorageWorks Fibre Channel Storage Switch User's Guide* for details.

The only suggested StorageWorks Fibre Channel Switch management task is switch zoning, to isolate different domains sharing a switch — see Section 3.13.2, page 3–26.

2.5 Choose the Root Password

You must select a root password before you perform any installation steps.

When the installation is complete, you can change the root password.

2.6 Record the External Gateway IP Address

Ask your site network administrator for the IP address of the external gateway, and record this value in Appendix D.

Physical Installation

This chapter describes how to physically install the HP AlphaServer SC system.

Use the checklist provided in Appendix A (if using a management server) or Appendix B (if not using a management server), to ensure that you complete all installation tasks in the correct order. If you are adding a management server to an HP AlphaServer SC system after system creation, use the checklist provided in Appendix C to ensure that you complete all installation tasks in the correct order. The information in this chapter is structured as follows:

- Management Server (see Section 3.1 on page 3–2)
- Compute Nodes and File Serving Nodes (see Section 3.2 on page 3–4)
- Physical Installation Overview (see Section 3.3 on page 3–4)
- Connect the Management Network (see Section 3.4 on page 3–11)
- Populate the HP AlphaServer SC PCI Slots (see Section 3.5 on page 3–17)
- Configure Hardware for a Dual-Rail Configuration (see Section 3.6 on page 3–20)
- Connect the HP AlphaServer SC Interconnect (see Section 3.7 on page 3–20)
- Connect the Node Console Port (see Section 3.8 on page 3–21)
- Configure the HP AlphaServer SC Interconnect Control Card with an IP Address (see Section 3.9 on page 3–22)
- Configure the Terminal Servers with an IP Address (see Section 3.10 on page 3–22)
- Connect the Fibre Channel Switches (see Section 3.11 on page 3–23)
- Verify Storage Component Revisions (see Section 3.12 on page 3–24)
- Configure the Fibre Channel Switch (see Section 3.13 on page 3–24)
- Configure the System Storage on the HSG80 (see Section 3.14 on page 3–27)
- Configure the System Storage on the HSV110 (see Section 3.15 on page 3–44)

Management Server

- Configure the System Storage on the MSA1000 (see Section 3.16 on page 3–62)

3.1 Management Server

A system can optionally be configured with a front-end management server. If the front-end management server is configured, certain housekeeping functions run on this server. This management server is not connected to the high-speed interconnect. If the management server is not configured, the housekeeping functions run on Node 0 (zero).

You can configure a single standalone management server, as described in Section 3.1.1, or you can have two management servers configured as a cluster, as described in Section 3.1.2.

3.1.1 Standalone Management Server

This is a non-clustered Tru64 UNIX server. The standalone management server, similar to the first two members of each domain, is configured with two Ethernet interfaces: one for the management LAN and one for an external LAN. The standalone management server is not connected to the HP AlphaServer SC Interconnect. The server types supported in this role are HP AlphaServer DS10, HP AlphaServer DS20E, and HP AlphaServer ES45. The management server node should be configured as shown in Figure 3–1, which shows an example HP AlphaServer SC standalone management server.




Figure 3–1 Single Management Server

3.1.2 Clustered Management Server

This is a standard TruCluster Server software implementation operating over a Gigabit Ethernet Interconnect, and should not be confused with the HP AlphaServer SC system, which operates over the HP AlphaServer SC Interconnect. In HP AlphaServer SC Version 2.6 (UK2), the clustered management server is qualified at two nodes. The server types supported in this role are HP AlphaServer DS10, HP AlphaServer DS20E, and HP AlphaServer ES45. Both servers should have identical hardware specifications. The implementation uses fibre channel shared storage provided by HSG80, HSV110, or MSA1000 technology.

Figure 3–2 shows an example HP AlphaServer SC dual management server and how the management server nodes are connected.

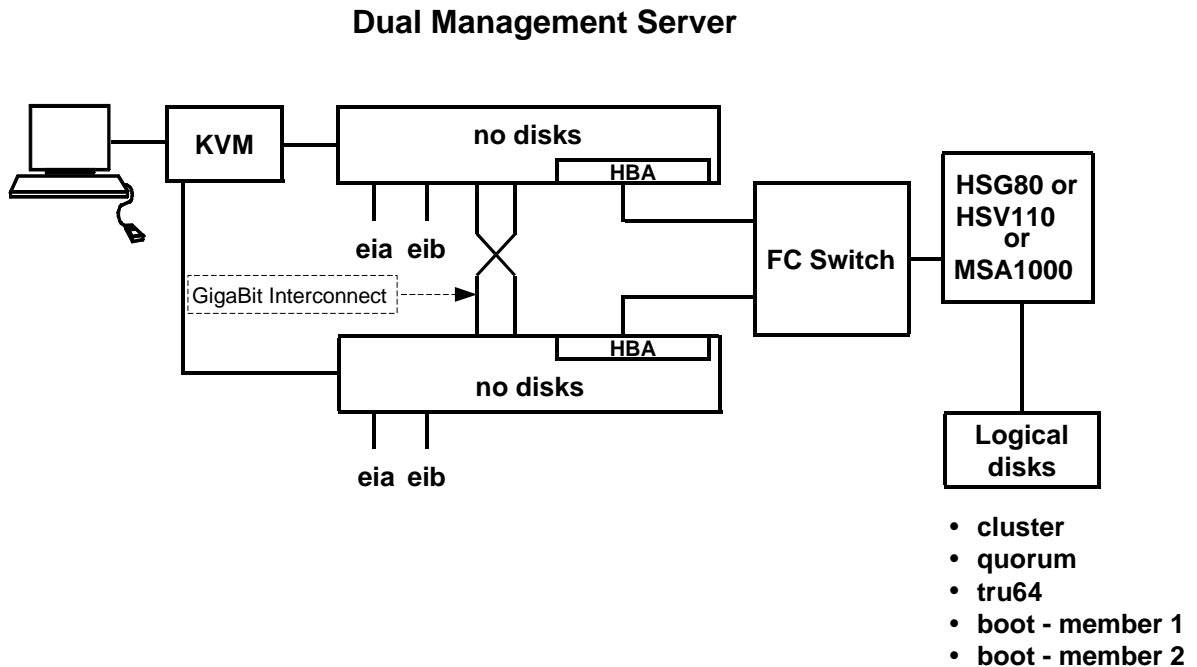


Figure 3–2 Dual Management Server

3.2 Compute Nodes and File Serving Nodes

The HP AlphaServer SC system is comprised of multiple domains. There are two types of domains: File-Serving (FS) domains and Compute-Serving (CS) domains. Nodes within these domains are called Compute Nodes or File-Serving Nodes.

3.3 Physical Installation Overview

Figure 3–3 on page 3–5 shows an example HP AlphaServer SC configuration for a 16-node system.

Figure 3–4 to Figure 3–7 show how the first three nodes are connected to the networks of the HP AlphaServer SC system, depending on the type of HP AlphaServer SC Interconnect switch used. See Table 3–1 to identify which figure applies to your system.

Table 3–1 How to Connect the Components of an HP AlphaServer SC System

If you are using...	See...
HP AlphaServer SC 16-port switch	Figure 3–4 on page 3–6
HP AlphaServer SC 128-port switch	Figure 3–5 on page 3–7
HP AlphaServer SC 16-port switch and DS20L Nodes	Figure 3–6 on page 3–8
HP AlphaServer SC 128-port switch and DS20L Nodes	Figure 3–7 on page 3–9
HP AlphaServer SC 128 Switches in a federated configuration.	Figure 3–8 on page 3–10

Note:

These diagrams are not to scale.

The nodes in these diagram have been re-arranged to show the cables more clearly — in a network cabinet, node numbers increase from bottom to top.

Figure 3–3 shows an example HP AlphaServer SC configuration, for a 16-node system. In this diagram, the ES4x value is used to represent either HP AlphaServer ES40, HP AlphaServer ES45, or HP AlphaServer DS20L nodes. KVM switch represents a Keyboard-Video-Mouse switch. Video-Mouse switch.

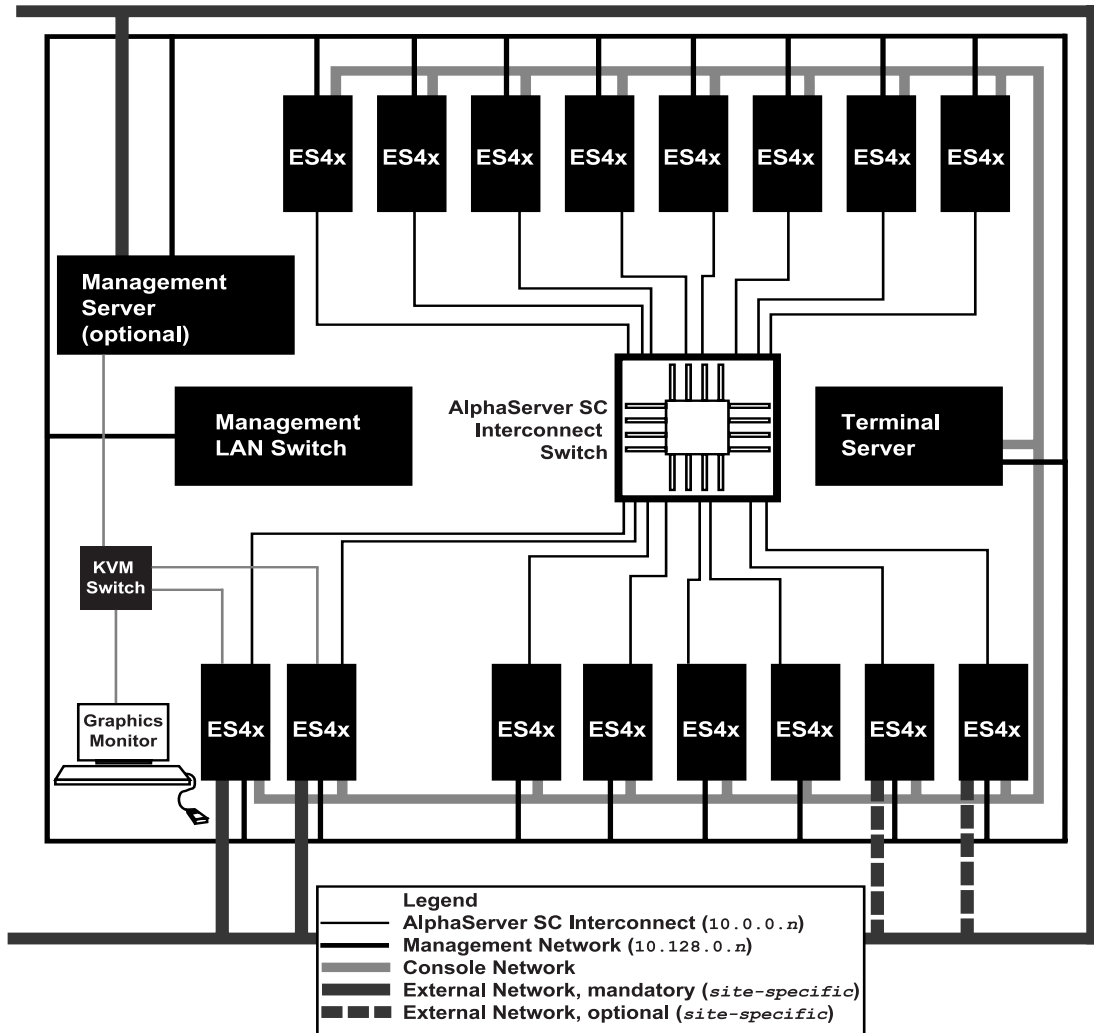


Figure 3–3 HP AlphaServer SC Configuration for a 16-Node System

Physical Installation Overview

Figure 3–4 shows how the first three HP AlphaServer ES40 nodes are connected to the networks of an HP AlphaServer SC system containing an HP AlphaServer SC 16-port switch, an optional management server, and an optional second rail.




Figure 3–4 Node Network Connection When Using an HP AlphaServer SC 16-Port Switch

Physical Installation Overview

Figure 3–5 shows how the first three nodes are connected to the networks of an HP AlphaServer SC system containing an HP AlphaServer SC 128-port switch, an optional management server, and an optional second rail.




Figure 3–5 Node Network Connection When Using an HP AlphaServer SC 128-Port Switch

Physical Installation Overview

Figure 3–6 shows how the first two HP AlphaServer DS20L nodes are connected to the networks of the HP AlphaServer SC system containing an HP AlphaServer SC 16-port switch and a management server.




Figure 3–6 Node Network Connections: HP AlphaServer SC 16-Port Switch, HP AlphaServer DS20L Nodes

Figure 3–7 shows how the first two HP AlphaServer DS20L nodes are connected to the networks of the HP AlphaServer SC system containing an HP AlphaServer SC 128-way switch and a management server.




Figure 3–7 Node Network Connections: HP AlphaServer SC 128-Port Switch, HP AlphaServer DS20L Nodes

Physical Installation Overview

Figure 3–8 shows the hardware connections when using a federated HP AlphaServer SC Interconnect configuration.

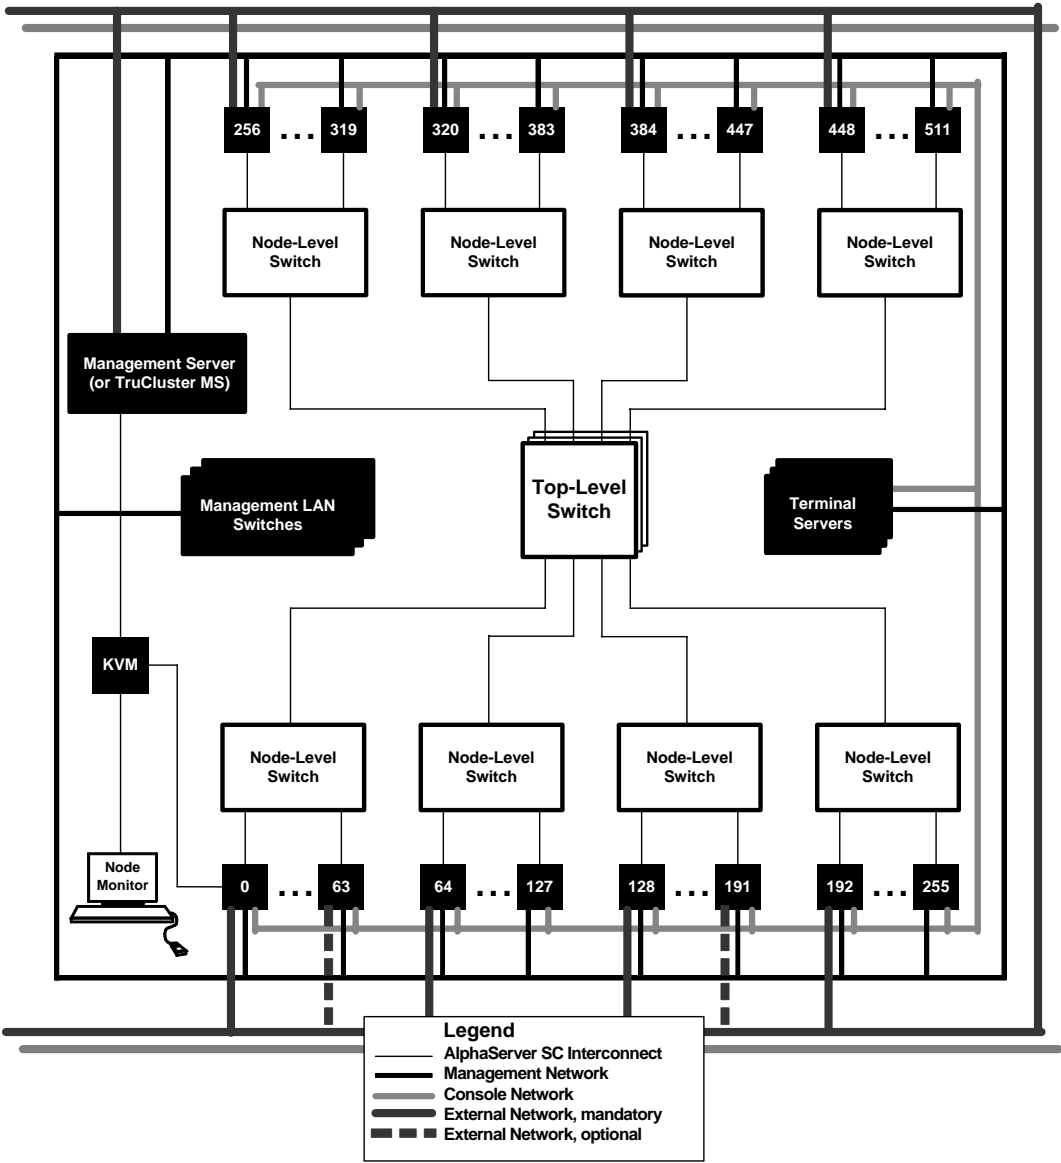


Figure 3–8 Federated HP AlphaServer SC Interconnect Configuration

3.4 Connect the Management Network

This section describes how to connect the management network and subsequently how to configure each kind of Ethernet switch supported by HP AlphaServer SC Version 2.6 (UK2). The information in this section is structured as follows:

- Cable the Management Network (see Section 3.4.1 on page 3–11)
- Configure the Summit Switch (see Section 3.4.2 on page 3–12)
- Configure the ProCurve Switch (see Section 3.4.3 on page 3–14)

When configuring the switch, please follow the instructions in the appropriate section for your switch type.

3.4.1 Cable the Management Network

Cable the management network as follows:

1. The management network interface adapter is the first network adapter on a node (in UNIX it is called `ee0`, in the SRM Console it is called `eia0`).

Note:

For the HP AlphaServer DS20L product, the second network adapter is used as the management network (in UNIX it is called `ee1`, in the SRM Console it is called `eib0`). This adaptation is required in order to implement the WOL functionality.

Note:

In Tru64 UNIX Version 5.1B, the device probe order is system-dependent. On systems with more than one network interface, the SRM equivalent of `ee0` may not be `eia0`.

Connect the appropriate adapter — on each node, and on the management server (if used) — with an appropriate cable to a port on the switch. The adapter may be connected to any port.

2. Connect the HP AlphaServer SC Interconnect Control Processor to a port on the switch. Record the port number in Appendix D; you will need this number in step 8 of Section 3.4.2 and step 7 of Section 3.4.3.
3. Connect the terminal server to a port on the switch. Record the port number in Appendix D; you will need this number in step 9 of Section 3.4.2 and step 8 of Section 3.4.3.
4. Connect the Fibre Channel switch to a port on the switch. Record the port number in Appendix D; you will need this number in step 10 of Section 3.4.2 and step 9 of Section 3.4.3.

Connect the Management Network

3.4.2 Configure the Summit Switch

Note:

This section describes how to connect the management network using the 48-port Summit switch from Extreme Networks. If you are using a 24-port Summit switch, substitute “24” for “48” in the commands in this section.

To configure a Summit switch, perform the following tasks:

1. Connect a terminal to the Summit switch.
2. Power on the Summit switch. When the switch has powered up, the login prompt appears. Enter the username `admin`, as follows:
login: **admin**
3. When prompted for the password, press the Return key on the terminal. When the system responds with the Summit prompt, enter the following commands (where `10.128.103.1` is the IP address of the Summit switch, and `10.128.0.1` is the IP address of the first node — see Table 2–1 on page 2–5 for more information about IP addresses):
* Summit 48:1 # **config vlan default ipaddress 10.128.103.1 255.255.0.0**
* Summit 48:2 # **config iproute add default 10.128.0.1**

Note:

In the case of a clustered management server, the default address should be set to `10.128.101.1` (the IP address of the first node of the management server cluster).

4. Run the following command to allow you to use a Web interface to manage your switch:
* Summit 48:3 # **enable web**
You are reminded that you must reboot the switch before this command will take effect. However, you do not need to reboot the switch at this point (you will reboot in step 12).
5. Set the time zone (where +0 specifies the time difference in minutes from Greenwich Mean Time (GMT); for example, specify +60 for Paris, or -300 for New York), as follows:
* Summit 48:4 # **config timezone +0**
6. Set the site-specific admin password, as follows:
* Summit 48:5 # **config account admin**
password:
Reenter password:
Record the admin password in a safe place — you are prompted for this password each time you log into the Summit switch.

7. Configure each port that is attached to the management network interface adapter of each node, and the port that is attached to the management network interface adapter of the management server (if used), by running one of the following commands:
 - If the nodes use an Ethernet card from the DE600 family, run the following command:

```
* Summit 48:6 # config ports 1-48 auto on
```
 - If the nodes use an Ethernet card from the DE500 family, run the following command:

```
* Summit 48:6 # config ports 1-48 auto off duplex full
```

If you set all ports to full duplex, as in the above example, ensure that you reset the appropriate ports to half duplex; for example, the terminal server port (see step 8).

Note:

The DE600 family Ethernet cards cannot communicate correctly with ports that are configured for `duplex full`. The DE500 family Ethernet cards cannot communicate correctly with ports that are configured for `auto on`. If you have problems establishing connectivity with either of these types of cards, please review your settings.

8. Configure the port that is connected to the HP AlphaServer SC Interconnect Control Processor — set this port to 100Mbps full duplex. You should have already recorded this port number in Appendix D (see step 2 in Section 3.4.1 above). For example, if the HP AlphaServer SC Interconnect Control Processor is connected to port 46, use the following command:

```
* Summit 48:7 # config ports 46 auto on
```
9. Configure the port that is connected to the terminal server — set this port to 10Mbps half duplex. You should have already recorded this port number in Appendix D (see step 3 in Section 3.4.1 above). For example, if the terminal server is connected to port 47, use the following command:

```
* Summit 48:8 # config ports 47 auto off speed 10 duplex half
```
10. Configure the port that is connected to the Fibre Channel switch. You should already have recorded this number in Appendix D (see step 4 in Section 3.4.1 above). The command will differ depending on the model of the fibre channel switch. Newer models support a 100MB Ethernet connection. If the switch is connected to port number 48, one of the following commands will apply:

If the switch is a newer 2GB SAN switch model, the command is:

```
* Summit 48:9 # config ports 48 auto on
```

If the switch is an older 1GB SAN switch model, the command is:

```
* Summit 48:9 # config ports 48 auto off speed 10 duplex half
```

Connect the Management Network

11. Save this configuration, as follows:
* Summit 48:10 # **save**
12. Reboot the switch to apply your changes, as follows:
* Summit 48:11 # **reboot**

Note:

Do not enable spanning tree protocol (STP) on the Summit switch.

If your HP AlphaServer SC system has more than 48 nodes, repeat the above steps to configure as many Summit switches as necessary. If you have more than one Summit switch, you must connect the Summit switches to each other using fibre connectors; that is, connect ports 50 and 50R on one Summit switch to ports 49 and 49R respectively on the next Summit switch.

For more information, see the *Summit Hardware Installation Guide* and the *ExtremeWare Software User Guide*.

3.4.3 Configure the ProCurve Switch

To configure a ProCurve switch, perform the following tasks:

1. Connect a terminal to the ProCurve switch.
2. Power on the ProCurve switch. When the switch has powered up, the following prompt is displayed:
Waiting for Speed Sense. Press <Enter> twice to continue.

There is no default password. Press carriage return twice to display the following user prompt:
HP ProCurve Switch 2650#
3. At the user prompt, enter the `setup` command to receive the Switch Setup menu where the default values will be shown for each field. Within the Switch Setup menu, you can use the arrow keys to navigate through the fields on the menu and use the <space> key to toggle field choices.
HP ProCurve Switch 2650# **setup**

```
HP ProCurve Switch 2650                               1-Jan-1990   0:10:26
=====  CONSOLE - MANAGER MODE  =====
                               Switch Setup

System Name : HP ProCurve Switch 2650
System Contact :
Manager Password :                               Confirm Password :
Logon Default : CLI                               Time Zone [0] : 0
Community Name : public                           Spanning Tree Enabled [No] : No
```

Connect the Management Network

Default Gateway :
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled

IP Config [DHCP/Bootp] : DHCP/Bootp
IP Address :
Subnet Mask :

Actions-> Cancel Edit Save Help

4. For the HP AlphaServer SC, the recommended values within the Switch Setup menu are as follows:

Manager Password : **<Site specific choice>**
Confirm Password : **<As above>**
Time Zone : **0**
Spanning Tree Enabled : **No**
IP Config : **Manual**
IP Address : **10.128.103.1**
Subnet Mask : **255.255.0.0**
Default Gateway : **10.128.0.1**

(where 10.128.103.1 is the IP address of the Summit switch, and 10.128.0.1 is the IP address of the first node — see Table 2–1 on page 2–5 for more information about IP addresses).

Note:

Do not enable spanning tree protocol (STP) on the switch.

In the case of a clustered management server, the default gateway should be set to 10.128.101.1 (the IP address of the first node of the management server cluster).

When accessing the console once a password is configured, you will be prompted with a password prompt. If you enter no password at that prompt, you will be granted a minimal CLI with the following prompt "HP ProCurve Switch 2650>", however if you enter the correct password, then you will be placed into a fully functional CLI with prompt "HP ProCurve Switch 2650#"

Within the Switch Setup menu, you can use the arrow keys to navigate through the fields on the menu and use the <space> key to toggle field choices.

Once comfortable with the contents of the menu, you should press <enter> to move focus to the Actions line.

From the Actions line, you can elect to cancel the changes, edit the form again, save the form or access online help.

Connect the Management Network

5. For the HP AlphaServer SC, all nodes will have their management interfaces configured to a fixed configuration of 100BaseT Full Duplex. The ports on the Procurve unit are configured to auto-negotiation by default, and in compliance with IEEE 802.3, the link speed will be negotiated to 100BaseT half duplex. Left in this condition, this will result in high error rates and inefficient communications between the switch and the node.

One solution is to make sure that all nodes themselves are configured to auto — in which case, the link speeds will autonegotiate correctly as 100BaseT full duplex. The more predictable solution is to configure the ports to a fixed speed and mode. This is accomplished from the configuration context.

Enter the configuration context from the user prompt, as follows:

```
HP ProCurve Switch 2650# config
HP ProCurve Switch 2650(config)#
```

6. From the configure context, configure each port that is attached to the management network interface adapter of each node, and the port that is attached to the management network interface adapter of the management server (if used), by running the following command:

```
HP ProCurve Switch 2650(config)# interface ethernet 1-48 speed-duplex 100-full
```

If you set all ports to full duplex, as in the above example, ensure that you reset the appropriate ports to half duplex; for example, the terminal server port (see step 8).

7. From the configure context, configure the port that is connected to the HP AlphaServer SC Interconnect Control Processor — set this port to 100Mbps full duplex. You should have already recorded this port number in Appendix D (see step 2 in Section 3.4.1 above). For example, if the HP AlphaServer SC Interconnect Control Processor is connected to port 46, use the following command:

```
HP ProCurve Switch 2650(config)# interface ethernet 46 speed-duplex 100-full
```

8. From the configure context, configure the port that is connected to the terminal server — set this port to 10Mbps half duplex. You should have already recorded this port number in Appendix D (see step 3 in Section 3.4.1 above). For example, if the terminal server is connected to port 47, use the following command:

```
HP ProCurve Switch 2650(config)# interface ethernet 47 speed-duplex 10-half
```

9. From the configure context, configure the port that is connected to the Fibre Channel switch. You should already have recorded this number in Appendix D (see step 4 in Section 3.4.1 above). The command will differ depending on the model of the Fibre Channel switch. Newer models support a 100MB Ethernet connection. If the switch is connected to port number 48, one of the following commands will apply:

If the switch is a newer 2GB SAN switch model, the command is:

```
HP ProCurve Switch 2650(config)# interface ethernet 48 speed-duplex 100-full
```

If the switch is an older 1GB SAN switch model, the command is:

```
HP ProCurve Switch 2650(config)# interface ethernet 48 speed-duplex 10-half
```

Populate the HP AlphaServer SC PCI Slots

10. To leave the configuration context, the command line interface, and then to save the configuration changes, enter the following sequence:

```
HP ProCurve Switch 2650(config)# exit
HP ProCurve Switch 2650# logout
Do you want to log out [y/n]? y
Do you want to save current configuration [y/n]? y
```

If your HP AlphaServer SC system has more than 48 nodes, repeat the above steps to configure as many Procurve switches as necessary. If you have more than one Procurve switch, you must connect the Procurve switches to each other using fibre connectors; that is, connect ports 50 on one switch to ports 49 on the next switch.

3.5 Populate the HP AlphaServer SC PCI Slots

The following sections describe how to populate the HP AlphaServer ES40, HP AlphaServer ES45 and HP AlphaServer DS20L PCI slots.

3.5.1 HP AlphaServer ES40 PCI Slots

Note:

The six-slot HP AlphaServer ES40 Model 1 system is not supported.

Populate the HP AlphaServer ES40 PCI slots as described in Table 3–2.

Table 3–2 Populating the HP AlphaServer ES40 PCI Slots

Slot	Description	Part Number	Priority
1	Graphics Card	SN-PBXGF-AB SN-PBXGK-BB	1 1
2	DO NOT USE		
3	HP AlphaServer SC Elan #1	3X-CCNNA-AA	1
4	DO NOT USE		
5	Ethernet Card #1	3X-DE602-AA	1
6	SCSI Adapter #1	3X-KZPCA-AA	1
7	Fibre Channel #1	DS-KGPSA-CA	1
8	Fibre Channel #2	DS-KGPSA-CA	1

Populate the HP AlphaServer SC PCI Slots

Table 3–2 Populating the HP AlphaServer ES40 PCI Slots

Slot	Description	Part Number	Priority
9	Gigabit Ethernet #1	DEGPA-SA	1
	HIPPI Adapter #1	KZPHA-AX	1
	ATM Adapter (622M) #1	3X-DAPCA-FA	1
	ATM Adapter (155M) #1	3X-DAPBA-FA	1
	ATM Adapter UTP (155M) #1	3X-DAPBA-UA	1
	SCSI Adapter #2	3X-KZPCA-AA	1
10	AlphaServer SC Elan #2	3X-CCNNA-AA	1
	Gigabit Ethernet #1	DEGPA-SA	2
	HIPPI Adapter #1	KZPHA-AX	2
	ATM Adapter (622M) #1	3X-DAPCA-FA	2
	ATM Adapter (155M) #1	3X-DAPBA-FA	2
	ATM Adapter UTP (155M) #1	3X-DAPBA-UA	2
	SCSI Adapter #2	3X-KZPCA-AA	2

For more information on how to physically install a PCI card or other module, please refer to the *HP AlphaServer ES40 Service Guide*.

3.5.2 HP AlphaServer ES45 PCI Slots

In the ten-slot HP AlphaServer ES45 Model 2 system, populate the HP AlphaServer ES45 PCI slots as described in Table 3–3.

Table 3–3 Populating the HP AlphaServer ES45 PCI Slots

Slot	Description	Part Number	Priority
1	Fibre Channel #3	3X-KZPSA-CA	1
	SCSI Adapter #2	3X-KZPEA-DB	2
	Gigabit Ethernet #1	DEGPA-SA	3
	Fibre Channel (2Gb) #2	3X-KGPSA-DA	4
	ATM Adapter (644M) #1	3X-DAPCA-FA	4
	ATM Adapter (155M) #1	3X-DAPBA-FA	4
	ATM Adapter UTP (155M) #1	3X-DAPCA-UA	4
2	Fibre Channel #4	3X-KZPSA-CA	1
	SCSI Adapter #3	3X-KZPEA-DB	2
	Gigabit Ethernet #2	DEGPA-SA	3
	Fibre Channel (2Gb) #3	3X-KGPSA-DA	4
	ATM Adapter (644M) #2	3X-DAPCA-FA	4
	ATM Adapter (155M) #2	3X-DAPBA-FA	4
	ATM Adapter UTP (155M) #2	3X-DAPCA-UA	4

Populate the HP AlphaServer SC PCI Slots

Table 3–3 Populating the HP AlphaServer ES45 PCI Slots

Slot	Description	Part Number	Priority
3	HIPPI Adapter #1	KZPHA-AX	1
	Gigabit Ethernet #3	DEGPA-SA	2
	Graphics Card #2	SN-PBXGF-AB	3
	ATM Adapter (644M) #3	3X-DAPCA-FA	4
	ATM Adapter (155M) #3	3X-DAPBA-FA	4
	ATM Adapter UTP (155M) #3	3X-DAPCA-UA	4
4	AlphaServer SC Elan #2	3X-CCNNA-AA	1
	Fibre Channel (2Gb) #1	3X-KGPSA-DA	2
5	Fibre Channel (2Gb) #4	3X-KGPSA-DA	1
	Graphics Card #1	SN-PBXGF-AB	2
	Ethernet Card #2	3X-DE602-BB	3
6	Fibre Channel #2	3X-KZPSA-CA	1
	SCSI Adapter #5	3X-KZPEA-DB	2
7	HP AlphaServer SC Elan #1	3X-CCNNA-AA	1
8	Ethernet Card #1	3X-DE602-BB	1
9	Fibre Channel #1	3X-KZPSA-CA	1
	SCSI Adapter #4	3X-KZPEA-DB	2
10	SCSI Adapter #1	3X-KZPEA-DB	2

Note:

The DE602 ethernet cards can only reside in the slots indicated if they are of the DE602-BB (66MHz) variant. If the only available ethernet cards are of the DE602-AA (33MHz) variant, then these cards must be placed in slot 3, 6, 9, or 10. This is necessary so as not to restrict usage of the other slots which are 66MHz capable.

Table 3–3 supports the following mutually exclusive possibilities as well as combinations of these cards in an HP AlphaServer SC system:

- Up to one Graphics card
- Up to two Ethernet cards
- Up to two HP AlphaServer SC Elan adapter cards in a system
- Up to two High Performance (Gigabit Ethernet, ATM, HIPPI) cards

Configure Hardware for a Dual-Rail Configuration

- Up to four Fibre Channel (2Gb) cards (limit of three cards when dual rail)
- Up to four Fibre Channel (1Gb) cards
- Up to four SCSI cards

3.5.3 HP AlphaServer DS20L PCI Slots

In the two-slot HP AlphaServer DS20L system, populate the HP AlphaServer DS20L PCI slots as described in Table 3–4.

Table 3–4 Populating the HP AlphaServer DS20L PCI Slots

Slot	Description	Part Number	Priority
0	Fibre Channel #1	DS-KGPSA-CA	1
1	AlphaServer SC Elan #1	3X-CCNNA-AA	1

3.6 Configure Hardware for a Dual-Rail Configuration

The hardware must be configured as follows for a dual-rail configuration:

- Each node must have two HP AlphaServer SC Elan adapter cards.
- Each HP AlphaServer SC Elan adapter card must be on a different PCI bus.
- The HP AlphaServer SC Elan adapter cards must be in the same slot positions on each node.
- The HP AlphaServer SC Elan adapter cards in a given slot position must be connected to the same HP AlphaServer SC Interconnect switch.
- Both HP AlphaServer SC Elan adapter cards in a given node must be connected to the same port on the HP AlphaServer SC Interconnect switch (that is, the node must have the same network ID on each HP AlphaServer SC Interconnect switch).
- Connect the parallel port on the HP AlphaServer SC Interconnect as described in Section 3.7.
- The same type of HP AlphaServer SC Interconnect switch must be used for each rail.

3.7 Connect the HP AlphaServer SC Interconnect

Connect the HP AlphaServer SC Interconnect as follows:

1. Power down all nodes.
2. Power down the HP AlphaServer SC Interconnect switch.

3. Connect the nodes to the HP AlphaServer SC Interconnect switch. Connect Node 0 to port 0, Node 1 to port 1, and so on.
4. Connect the parallel port on the HP AlphaServer SC Interconnect switch as follows:
 - a. 16-Port Switch:
 - Connect the first switch to Node 0, the second switch to Node 1, and so on.
 - b. 128-Port Switch with traditional chassis (no HP AlphaServer SC Interconnect Control Card):
 - Connect the first switch to Node 0, the second switch to Node 1, and so on.
 - c. 128-Port switch with HP AlphaServer SC Interconnect Control Card:
 - Connect the HP AlphaServer SC Interconnect Control Card to the parallel port on the switch chassis
 - Connect the HP AlphaServer SC Interconnect Control Card to the Ethernet
 - Repeat these two steps for each 128-Port switch.

For more information, see the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

3.8 Connect the Node Console Port

Connect the appropriate cable to the console port (COM 1) of each node. For example, Node 0 must be connected to port 1 of the first terminal server, Node 1 must be connected to port 2, and so on.

The management server may be connected to any port on any terminal server, but nodes must be added to consecutive ports. Therefore, HP recommends that you connect the management server to the last terminal server port. This minimizes the possibility of having to move the management server connection to a different terminal server port when adding nodes. Record in Appendix D the port number of the terminal server port connected to the management server.

Configure the HP AlphaServer SC Interconnect Control Card with an IP Address

3.9 Configure the HP AlphaServer SC Interconnect Control Card with an IP Address

Each HP AlphaServer SC Interconnect control processor is connected to the management network *via* an Ethernet cable. To connect to the HP AlphaServer SC Interconnect control processors over the management network, you must first assign an IP address to each HP AlphaServer SC Interconnect control processor.

For instructions on how to configure the HP AlphaServer SC Interconnect Control Card with an IP address, please refer to Section 7.2.1 (Using the Serial Port to Assign a Static IP Address) of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

3.10 Configure the Terminal Servers with an IP Address

You must configure the terminal servers by performing the following steps on each terminal server:

1. Connect a terminal to port 1 of the terminal server (that is, disconnect the port from Node 0, and connect the port to a terminal instead), so that you can configure the terminal server with an IP address.

Note:

On an unconfigured terminal server, you can use any port to manage the terminal server itself. However, once the `sra` utility has configured any terminal server ports during installation (see either Section 5.2, step 10 on page 5–39 or Section 6.2, step 11 on page 6–26), the ports configured by the `sra` utility cannot be used to manage the terminal server — you must use an unconfigured port.

2. Press the Return key to display the `username` prompt. Enter any value at this prompt:
`Username: anything`
3. Enter privileged mode, as follows:
`LOCAL> SET PRIVILEGE`
The factory default password is `system`.
4. Set the subnet mask, as follows:
`LOCAL> DEFINE INTERNET SUBNET MASK 255.255.0.0`

Note:

Set the subnet mask before you set the IP address.

5. Set the IP address, as follows:

```
LOCAL> DEFINE INTERNET ADDRESS ip_address
```

where *ip_address* is one of the following:

- 10.128.100.1 (first terminal server)
- 10.128.100.2 (second terminal server)
- 10.128.100.3 (third terminal server)
- 10.128.100.4 (fourth terminal server)

6. Set the gateway, as follows:

```
LOCAL> DEFINE INTERNET GATEWAY gateway_ip_address
```

where *gateway_ip_address* is 10.128.0.1 (the IP address of the first node) or, in the case of a clustered management server, 10.128.101.1 (the IP address of the first node of the management server cluster).

7. Initialize the IP address, as follows:

```
LOCAL> INITIALIZE
```

8. Disconnect the terminal from port 1 of the terminal server, and reconnect the port to Node 0.

Once you have set the terminal server IP address as described in this section, and installed the `sra` utility (as described in either Section 5.1.9 on page 5–32 if using a management server, or Section 6.1.7 on page 6–20 if not using a management server), you can run the `sra setup` command to configure port characteristics (as described in either Section 5.2 on page 5–37 if using a management server, or Section 6.2 on page 6–24 if not using a management server).

3.11 Connect the Fibre Channel Switches

Cable your storage arrays and hosts to the fibre channel switches for multiple-bus failover mode — see the example configuration in Figure 3–9 on page 3–28 or the example HSV110 configuration in Figure 3–11 on page 3–45.

Verify Storage Component Revisions

3.12 Verify Storage Component Revisions

Table 3–5 shows the software, driver, and firmware revisions that are required for components in HP AlphaServer SC Version 2.6 (UK2) systems. Check that the components in your system meet these minimum requirements.

Table 3–5 Minimum System Driver and Firmware Versions

Product	Minimum Revision
SAN Appliance	V1.0C Build 20020108
SAN Switch 2/16	Kernel 5.3.1, Fabric OS V3.0.2a
Emx Driver	1.32a
KGPSA-BC	F/W Rev 3.03A1(1.31)
KGPSA-CA	F/W Rev 3.81A4(2.01A0)
KGPSA-DA	F/W Rev 3.81A4(1.01A0)
HSG80	ACS V8.7
HSV110	VCS V2.002
MSA1000	F/W Rev V3.30A

3.13 Configure the Fibre Channel Switch

The information in this section is organized as follows:

- Configure Fibre Channel Switch Network Addressing (see Section 3.13.1 on page 3–24)
- Switch Zoning (see Section 3.13.2 on page 3–26)

3.13.1 Configure Fibre Channel Switch Network Addressing

The following items are required to set network addressing.

- An IP address from your Network Administrator
- Fibre channels switch installed and connected to a power source
- Serial cable (supplied with the switch) for connecting the switch to the workstation
- A local workstation (desktop or notebook computer) with RS-232 serial communications software (VT100 emulation).

Configure the Fibre Channel Switch

- Ethernet cable for connecting the switch to the workstation or to a network containing the workstation
- Small form factor Gigabit interface convertors (SFF GBICs) and cables, as required to connect the switch to the fabric

Verify or change the fibre channel switch IP address, subnetmask, or gateway address as follows:

Note:

During first time setup, you must replace the factory IP, subnetmask and gateway addresses with addresses provided by your Network Administrator.

1. Remove the shipping plug from the fibre channel switch serial port.
2. Connect the serial cable to the fibre channel switch serial port
3. Connect the other end of the serial cable to an RS-232 serial port on the workstation.
4. Verify that the switch power is on and the power on self test (POST) is completed.
5. Power on the workstation and establish a connection to the switch using a terminal emulator application.
6. Configure the terminal emulation port settings as follows:
 - Bits per second: 9600
 - Databits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None
 - SHR-2497A

To configure port settings in a Tru64 UNIX environment, enter the following command:
tip /dev/ttyb -9600

7. Log on to the switch (with administrative privileges). The default administrative logon is *admin* and the default password is *password*.
 - a. Enter the following at the prompt:
ipAddrSet
 - b. Enter the following information at the corresponding prompts, listed below:
 - Ethernet IP Address [10.77.77.77]:Enter the new Ethernet IP address (see Table 2–1 on page 2–5).

Configure the Fibre Channel Switch

- Ethernet subnetmask [0.0.0.0]:
Enter the new Ethernet subnetmask (see Table 2–1 on page 2–5)
 - Fibre Channel IP Address [none]:
Press the Return key to select **none**.
 - Fibre Channel subnetmask [none]:
Press the Return key to select **none**.
 - Gateway Address [172.17.1.1]:
Enter the new gateway address (see Table 2–1). The gateway address will be 10.128.0.1 (the IP address of the first node) or in the case of a clustered management server, it will be 10.128.101.1 (the IP address of the first node of the clustered management server).
 - Set IP address now? [y = set now, n = next reboot]:
Enter **y** to set now.
- c. To verify that the IP address was entered correctly, enter:
ipAddrShow
- d. Once the IP address is verified as correct, remove the serial cable, and replace the shipping plug in the serial port.

Note:

The serial port is intended only for use during the initial setting of the IP address and for service purposes. Using the serial port during normal switch operation or for regular maintenance is not recommended.

8. Record the IP address on the label affixed to the front panel of the fibre channel switch.

3.13.2 Switch Zoning

If you have one RAID controller pair per domain, you can limit access to the RAID storagesets by using switch zoning. The following example shows how to create two zones on an 8-port switch:

1. Telnet to the fibre channel switch and log in as the admin user, as follows:

```
# telnet 10.128.104.1
user: admin
password: password
```
2. Create two zones on the switch, as follows:

```
admin> zoneCreate "zone_A", "1,0; 1,1; 1,4; 1,5"
admin> zoneCreate "zone_B", "1,2; 1,3; 1,6; 1,7"
admin> cfgCreate "atlas", "zone_A; zone_B"
```


Configure the System Storage on the HSG80

This command creates Zone A (`zone_A`) with ports 0, 1, 4, and 5, and Zone B (`zone_B`) with ports 2, 3, 6, and 7. Typically, the first domain (`atlasD0`) is connected to ports 0 and 1, and its RAID is connected to ports 4 and 5. Similarly, the second domain (`atlasD1`) is connected to ports 2 and 3, and its RAID is connected to ports 6 and 7.

3. Enable the zoning configuration, as follows:

```
admin> cfgEnable 'atlas'
```

4. Save the zoning configuration to flash memory, in case of power failure, as follows:

```
admin> cfgSave
```

For more information about switch zoning, see the *Fibre Channel Storage Switch User's Guide*.

3.14 Configure the System Storage on the HSG80

This section describes the tools for managing the physical components that make up the HSG80 storage subsystem (see Section 3.14.1 on page 3–29), and gives instructions for configuring system storage on the HSG80.

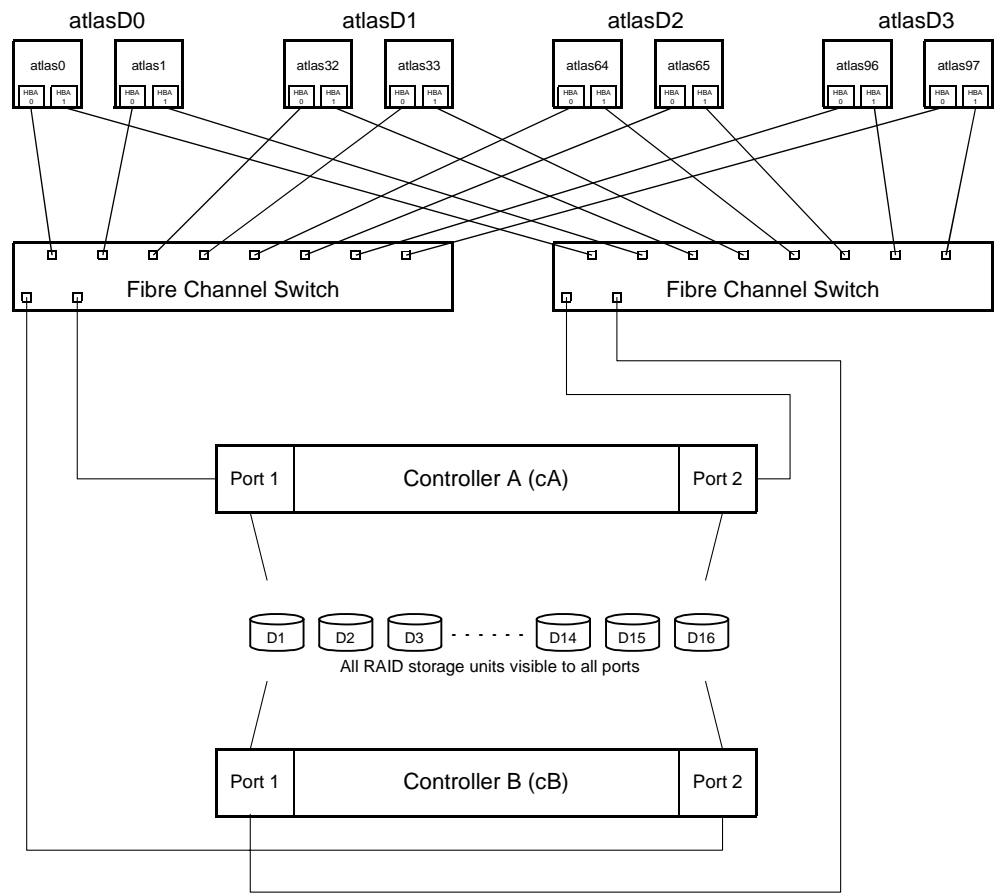
To configure the system storage on a Fibre Channel RAID subsystem, perform the following tasks on each domain:

- Configure the HSG80 RAID Controllers (see Section 3.14.2 on page 3–31)
- Configure the HSG80 RAID Storage (see Section 3.14.3 on page 3–36)

Before configuring the system storage on the HSG80, check that the system firmware on the HSG80 meets the requirements described in Section 3.12 on page 3–24

Configure the System Storage on the HSG80

Figure 3–9 shows the cabling in the example HSG80 system storage configured in this section.



HBA = Fibre Channel Host Bus Adapter

Figure 3–9 Example System Storage Configuration — Cabling

Configure the System Storage on the HSG80

Figure 3–10 shows the RAID storage units in the example system storage configured in this section.

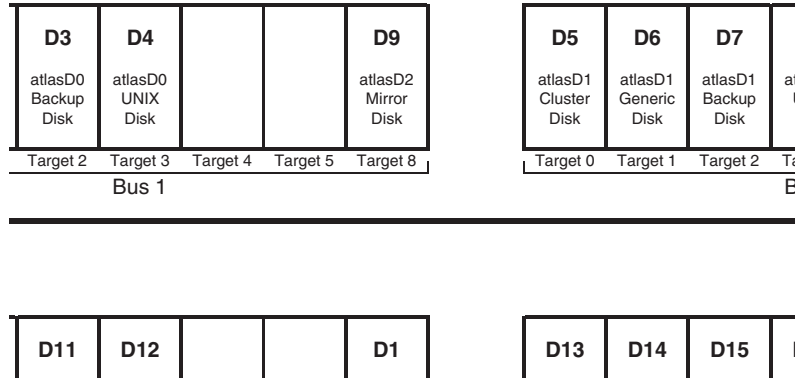


Figure 3–10 Example System Storage Configuration — RAID Storage Units

3.14.1 HSG80 Storage Management Tools

Several tools and methods are required to manage the physical components that make up the storage capabilities of the HSG80. This section describes the following tasks:

- Managing Storage Arrays and Associated HSG80 RAID Controllers (see Section 3.14.1.1 on page 3–29)
- Using SANworks Command Scriptor (see Section 3.14.1.2 on page 3–30)

3.14.1.1 Managing Storage Arrays and Associated HSG80 RAID Controllers

You can manage the storage arrays and HSG80 RAID controllers in four different ways:

- Connect a PC running the HP StorageWorks Command Console software, *via* a serial line, directly to the HSG80 RAID controllers. The StorageWorks command console client is an application that runs on Microsoft® Windows 95, Windows NT®, or Windows 2000. The client provides both a CLI and a Graphical User Interface (GUI).
- Attach a terminal emulator directly to the HSG80 RAID controllers.
In the example in Section 3.14 on page 3–27, this is the method used to configure the RAID controllers.
- Connect the HSG80 RAID controllers to a terminal server using a serial line.
- Use the HP SANworks Command Scriptor software product. This allows you to access the CLI of the HSG80 RAID controllers from any host that has a fibre channel HBA card connected by the fibre channel switch to the HSG80 RAID controllers. For more information on SANworks Command Scriptor, see Section 3.14.1.2 on page 3–30.

Configure the System Storage on the HSG80

Using a terminal server with telnet access provides the flexibility to manage many units from a PC or workstation anywhere on your network. Please see the documentation supplied with your hardware, and the StorageWorks Command Console software documentation, for details.

Storage that is configured as described in this guide is highly available. The system storage for each domain is located on a Mirrorset, and a spareset is created to automatically source a replacement disk in the case of a failed Mirrorset member. Storage created for user data can be created using any of the options supplied by your storage subsystem. Carefully plan the type of storage (RAIDset, Stripeset, Mirrorset — see Section 2.4.4 on page 2–13) to accomplish the correct combination of characteristics (for example, availability, I/O rate) for your environment and applications.

3.14.1.2 Using SANworks Command Scriptor

The SANworks Command Scriptor software product allows you to access the CLI of the HSG80 RAID controllers from any host that has a fibre channel HBA card connected by the fibre channel switch to the HSG80 RAID controllers.

SANworks Command Scriptor is a separate product that can be ordered. To install SANworks Command Scriptor, insert and mount the product CD-ROM on any node in the domain and follow the installation instructions in the accompanying documentation. Alternatively, copy the contents of the UNIX directory from the CD-ROM to a convenient directory, and install from there.

Once installed, use the SANworks Command Scriptor as described in this section.

SANworks Command Scriptor interacts with the HSG80 RAID controller through the fibre channel switch. To specify which HSG80 RAID controller you are interacting with, you must specify the name of the device (disk name) that the operating system has assigned.

To identify the device, use the `hwmgr -v dev` command as shown in the following example:

```
atlas0# hwmgr -v dev
```

HWID:	Device Name	Mfg	Model	Location
4:	/dev/kevm			
51:	/dev/disk/floppy0c		3.5in floppy	fdi0-unit-0
56:	/dev/disk/dsk0c	HP	BD018635C4	bus-0-targ-0-lun-0
57:	/dev/disk/dsk1c	HP	BD018122C9	bus-0-targ-1-lun-0
58:	/dev/disk/dsk2c	HP	BD018122C9	bus-0-targ-2-lun-0
59:	/dev/disk/dsk3c	DEC	HSG80	bus-1-targ-0-lun-1
60:	/dev/disk/dsk4c	DEC	HSG80	bus-1-targ-0-lun-2
61:	/dev/disk/dsk5c	DEC	HSG80	bus-1-targ-0-lun-3
62:	/dev/disk/dsk6c	DEC	HSG80	bus-1-targ-0-lun-4
63:	/dev/disk/cdrom0c	HP	CDR-8435	bus-3-targ-0-lun-0

In this example, there is only one HSG80 RAID controller (bus-1-targ-0). You can use `dsk3`, `dsk4`, `dsk5`, or `dsk6` to access this controller. It does not matter which of these devices you use — in each case, the commands go to the same controller.

Configure the System Storage on the HSG80

To execute CLI commands, you can either execute one command at a time, or you can place your commands in a file and execute the file.

The following example shows how to execute a single command:

```
atlas0# cmdscript -f dsk6c 'show units'
```

LUN	Uses	Used by
D1	MIRRORA	(partition)
D2	DISK10100	(partition)
D3	DISK10200	(partition)
D4	RAID1	(partition)

The following example shows how to execute commands in a file:

```
atlas0# cmdscript -f dsk6c < cmdfile.txt
```

3.14.2 Configure the HSG80 RAID Controllers

The HSG80 storage subsystem has two HSG80 controllers. The controller type is indicated by the CLI prompt (HSG80>). Configure the HSG80 controllers as follows:

1. Connect a PC *via* serial cable to one of the HSG80 controllers (cA) and use a terminal emulator to issue commands to the HSG80 command line interpreter (CLI).
2. Each controller has a high-speed cache to improve performance. Each cache must be protected from power failure — this protection is normally provided by connecting the cache to an external battery backup. If you do not have a battery backup, you should connect the cache to an uninterruptible power supply (UPS).

If using UPS instead of battery backup, enter one of the following commands when configuring the controller — see the *HP StorageWorks HSG80 Array Controller CLI Reference Guide* for more information about these commands:

```
HSG80> SET THIS UPS=NODE_ONLY
```

or

```
HSG80> SET THIS UPS=DATACENTER_WIDE
```

Note:

Setting the UPS controller variable for one controller sets it for both controllers.

3. To ensure proper operation of the HSG80 with Tru64 UNIX and TruCluster Server, set the controller values as follows:
 - a. Remove any failover mode that may have been previously configured:

```
HSG80> SET NOFAILOVER
```
 - b. Prevent the command line interpreter (CLI) from reporting a misconfiguration error resulting from not having a failover mode set:

```
HSG80> CLEAR CLI
```

Configure the System Storage on the HSG80

- c. Set up mirrored cache, if desired, for the controller pair:
`HSG80> SET THIS MIRRORED_CACHE`
- d. Put the controller pair into multiple-bus failover mode. Ensure that you copy the configuration information from the controller known to have a good array configuration:
`HSG80> SET MULTIBUS COPY=THIS`
- e. When the command is entered to set multiple-bus failover and copy the configuration information to the other controller, the other controller will restart. The restart may set off the audible alarm (which can be silenced by pressing the button on the EMU). The CLI will display an event report, and continue reporting the condition until cleared with the `CLEAR CLI` command:
`HSG80> CLEAR CLI`
- f. Take the ports off line and reset the topology to prevent an error message when setting the port topology:
`HSG80> SET THIS PORT_1_TOPOLOGY=OFFLINE`
`HSG80> SET THIS PORT_2_TOPOLOGY=OFFLINE`
`HSG80> SET OTHER PORT_1_TOPOLOGY=OFFLINE`
`HSG80> SET OTHER PORT_2_TOPOLOGY=OFFLINE`
- g. Set fabric as the switch topology:
`HSG80> SET THIS PORT_1_TOPOLOGY=FABRIC`
`HSG80> SET THIS PORT_2_TOPOLOGY=FABRIC`
`HSG80> SET OTHER PORT_1_TOPOLOGY=FABRIC`
`HSG80> SET OTHER PORT_2_TOPOLOGY=FABRIC`
- h. Set the date and time on this controller. In a dual-redundant configuration, the command sets the time on both controllers. The value takes effect immediately:
`HSG80> SET THIS TIME=DD-MMM-YYYY:HH:MM:SS`
- i. Specify the host protocol to use — you can use either SCSI-2 or SCSI-3:
 - Setting the `SCSI_VERSION` to SCSI-2 allows a disk unit to be at LUN0, and specifies that the command console LUN (CCL) is not fixed at a particular location, but floats to the first available LUN.
 - If `SCSI_VERSION` is set to SCSI-3, the CCL is presented at LUN0 for all connection offsets. Do not assign unit 0 at any connection offset because the unit would be masked by the CCL at LUN0 and would not be available.
 - Setting `SCSI_VERSION` to SCSI-3 is preferred because the CCL is fixed and it is much easier to manage a fixed CCL than a CCL that can change:
`HSG80> SET THIS SCSI_VERSION=SCSI-3`
`HSG80> SET OTHER SCSI_VERSION=SCSI-3`
- j. The HSG80 prompts you to restart both controllers after you set the SCSI version:
`HSG80> RESTART OTHER`
`HSG80> RESTART THIS`

Configure the System Storage on the HSG80

4. Compare the `HOST_ID` (the World Wide Name — WWN) reported by the console `show device` command with the connections reported by the controller `SHOW CONNECTIONS` command, as shown in the example below.

In this example, `atlas0` has two fibre channel Host Bus Adapter cards (HBAs). To identify the `atlas0` connections to the HSG80 controllers, perform the following steps:

- a. Identify the WWN of each HBA (see Table 3–6), by running the `show device` command at the `atlas0` SRM console prompt:

```
P00>>> show device
pga0.0.0.2.1          PGA0          WWN 2000-0000-c921-26e7
pgb0.0.0.3.1          PGB0          WWN 2000-0000-c921-310D
```

Table 3–6 Identifying the WWN of the HBAs — HSG80

HBA	WWN
0	2000-0000-c921-26e7
1	2000-0000-c921-310D

- b. List all of the connections (from all HBAs in all nodes) to the HSG80 controllers, by running the `SHOW CONNECTIONS` command on one of the HSG80 controllers (for example, Controller A):

```
HSG80> SHOW CONNECTIONS
Connection
Name      Operating system      Unit
Name      Operating system      Controller      Port      Address      Status      Offset
...
!NEWCON23  TRU64_UNIX            THIS            2          210613        OL this      0
HOST_ID=2000-0000-C921-26e7  ADAPTER_ID=1000-0000-C921-26e7
...
!NEWCON34  TRU64_UNIX            OTHER            1          210613        OL other      0
HOST_ID=2000-0000-C921-26e7  ADAPTER_ID=1000-0000-C921-26e7
...
!NEWCON42  TRU64_UNIX            OTHER            2          398576        OL other      0
HOST_ID=2000-0000-C924-310D  ADAPTER_ID=1000-0000-C924-310D
...
!NEWCON57  TRU64_UNIX            THIS            1          398576        OL this      0
HOST_ID=2000-0000-C924-310D  ADAPTER_ID=1000-0000-C924-310D
...
```

Configure the System Storage on the HSG80

- c. Identify the connections from the HSG80 controllers to the `atlas0` HBAs, by finding the entries in which the `HOST_IDs` (step b) match the `WWNs` (step a), as shown in Table 3–7:

Table 3–7 Identifying the Connections from the HSG80 Controllers to `atlas0`

From <code>show devices</code> on <code>atlas0</code> (step a):		From <code>SHOW CONNECTIONS</code> on HSG80 (step b):		
HBA	WWN	HOST_ID	Connection Name	Controller
0	2000-0000-c921-26e7	2000-0000-c921-26e7	!NEWCON23	THIS
0	2000-0000-c921-26e7	2000-0000-c921-26e7	!NEWCON34	OTHER
1	2000-0000-c921-310D	2000-0000-c921-310D	!NEWCON42	OTHER
1	2000-0000-c921-310D	2000-0000-c921-310D	!NEWCON57	THIS

You have now identified the four connections from the HSG80 controllers to `atlas0`:

- !NEWCON23 is the connection from Controller A to HBA 0 on `atlas0`.
 - !NEWCON34 is the connection from Controller B to HBA 0 on `atlas0`.
 - !NEWCON42 is the connection from Controller B to HBA 1 on `atlas0`.
 - !NEWCON57 is the connection from Controller A to HBA 1 on `atlas0`.
5. Create a naming convention for the connections, to give them meaningful names. Ensure that your naming convention allows you to add multiple connections to a node at a later date, and to identify the HBA (in the node) to which the controllers are connected. Note that the maximum length of the connection name is nine characters.

The convention used in this example is `N<node#>-<adapter#>-<connection#>` where

- `node#` is the node number (four-digits, left-pad with zeros as necessary)
 - `adapter#` is **0** for HBA 0, and **1** for HBA 1
 - `connection#` is **0** for this controller, and **1** for other controller
6. Rename the connections using the `rename` command.

For example, to rename the connections identified in step 4, enter the following commands:

```
HSG80> RENAME !NEWCON23 N0000-0-0
HSG80> RENAME !NEWCON34 N0000-0-1
HSG80> RENAME !NEWCON42 N0000-1-1
HSG80> RENAME !NEWCON57 N0000-1-0
```


Configure the System Storage on the HSG80

Repeat the rename command for each node connected, as shown in Table 3–8.

Table 3–8 Renaming the Connections — HSG80

Domain	Node	HBA	Controller	Connection Name
atlasD0	0	0	THIS	N0000-0-0
	0	0	OTHER	N0000-0-1
	0	1	THIS	N0000-1-0
	0	1	OTHER	N0000-1-1
	1	0	THIS	N0001-0-0
	1	0	OTHER	N0001-0-1
	1	1	THIS	N0001-1-0
	1	1	OTHER	N0001-1-1
atlasD1	32	0	THIS	N0032-0-0
	32	0	OTHER	N0032-0-1
	32	1	THIS	N0032-1-0
	32	1	OTHER	N0032-1-1
	33	0	THIS	N0033-0-0
	33	0	OTHER	N0033-0-1
	33	1	THIS	N0033-1-0
	33	1	OTHER	N0033-1-1
atlasD2	64	0	THIS	N0064-0-0
	64	0	OTHER	N0064-0-1
	64	1	THIS	N0064-1-0
	64	1	OTHER	N0064-1-1
	65	0	THIS	N0065-0-0
	65	0	OTHER	N0065-0-1
	65	1	THIS	N0065-1-0
	65	1	OTHER	N0065-1-1
atlasD3	96	0	THIS	N0096-0-0
	96	0	OTHER	N0096-0-1
	96	1	THIS	N0096-1-0
	96	1	OTHER	N0096-1-1
	97	0	THIS	N0097-0-0
	97	0	OTHER	N0097-0-1
	97	1	THIS	N0097-1-0
	97	1	OTHER	N0097-1-1

- For each connection to your domain, verify that the operating system is TRU64_UNIX and the unit offset is 0. Search the SHOW CONNECTION display for the WWN of each of the KGPSA adapters in your domain member systems. If the operating system and offsets are incorrect, set them, and then restart both controllers as follows:

Configure the System Storage on the HSG80

- a. Set the relative offset for LUN numbering to 0. You can set the `unit_offset` to nonzero values, but use caution. You may not be able to access storage units if you set the `unit_offset` improperly:

```
HSG80> SET N0000-0-0 UNIT_OFFSET=0
```

Repeat this command for each connection listed in Table 3–8.

- b. Specify that the host environment that is connected to the Fibre Channel port is TRU64_UNIX, as follows:

```
HSG80> SET N0000-0-0 OPERATING_SYSTEM=TRU64_UNIX
```

You must change each connection to TRU64_UNIX. This is very important.

Caution:

Failure to set this value to TRU64_UNIX will prevent your system from booting correctly, recovering from run-time errors, or from booting at all. The default operating system is Windows NT, which uses a different SCSI dialect to talk to the HSG80 controller.

Repeat this command for each connection listed in Table 3–8.

- c. Restart both controllers to cause all changes to take effect:

```
HSG80> RESTART OTHER
```

```
HSG80> RESTART THIS
```
- d. Enter the `SHOW CONNECTIONS` command once more and verify that all connections have the offsets set to 0 and the operating system set to TRU64_UNIX:

```
HSG80> SHOW CONNECTIONS
```

You have now configured both RAID controllers. For a complete description of all of the capabilities of HSG80 RAID controllers, see the *HP StorageWorks HSG80 Array Controller CLI Reference Guide*.

3.14.3 Configure the HSG80 RAID Storage

In the example in this section:

- Each domain has a Mirrorset and two individual disks.
- There is an additional disk, which is added as a spareset.

Configure the System Storage on the HSG80

Configure the HSG80 RAID disks as follows:

1. Using the CLI, run the configuration script to find the new disks in the array and add them, as follows:

```
HSG80> RUN CONFIG
```

The RUN CONFIG command produces the following output:

```
Config Local Program Invoked
```

```
Config is building its tables and determining what devices exist  
on the subsystem. Please be patient.
```

```
ADD DISK DISK10000 1 0 0  
ADD DISK DISK10100 1 1 0  
ADD DISK DISK10200 1 2 0  
ADD DISK DISK10300 1 3 0 [SC20 Product]  
ADD DISK DISK10800 1 8 0  
ADD DISK DISK20000 2 0 0  
ADD DISK DISK20100 2 1 0  
ADD DISK DISK20200 2 2 0  
ADD DISK DISK20300 2 3 0 [SC20 Product]  
ADD DISK DISK20800 2 8 0  
ADD DISK DISK30000 3 0 0  
ADD DISK DISK30100 3 1 0  
ADD DISK DISK30200 3 2 0  
ADD DISK DISK30300 3 3 0 [SC20 Product]  
ADD DISK DISK30800 3 8 0  
ADD DISK DISK40000 4 0 0  
ADD DISK DISK40100 4 1 0  
ADD DISK DISK40200 4 2 0  
ADD DISK DISK40300 4 3 0 [SC20 Product]  
ADD DISK DISK40500 4 5 0  
ADD DISK DISK40800 4 8 0
```

The format of the ADD DISK command is as follows:

```
ADD DISK DISKx0y0z x y z
```

where

- *x* is the SCSI device port number, from 1 to 6, on which the disk resides.
- *y* is the SCSI target ID, from 0 to 15 (excluding 6 and 7), of the disk on the port (values 6 and 7 are reserved for use by controllers cA and cB respectively).
- *z* is the LUN of the disk drive, and is always zero.

Note:

You can also add disks individually by entering the appropriate ADD DISK commands at the CLI prompt.

Configure the System Storage on the HSG80

2. Create the Mirrorset-type storagesets, as shown in Table 3–9.

Table 3–9 Creating the Mirrorsets

Domain	Command
atlasD0	HSG80> ADD MIRROR MIRRORD0 DISK10000 DISK30800
atlasD1	HSG80> ADD MIRROR MIRRORD1 DISK20000 DISK40800
atlasD2	HSG80> ADD MIRROR MIRRORD2 DISK30000 DISK10800
atlasD3	HSG80> ADD MIRROR MIRRORD3 DISK40000 DISK20800

3. Create a spareset for the HSG80 RAID subsystems, in case any of the other disks fails:
HSG80> **ADD SPARESET DISK40500**
4. Initialize the storagesets (Mirrorsets, JBOD, and spareset), as shown in Table 3–10.

Table 3–10 Initializing the Storagesets — HSG80

Domain	Command
atlasD0	HSG80> INITIALIZE MIRRORD0
	HSG80> INITIALIZE DISK10100
	HSG80> INITIALIZE DISK10200
<i>[SC20 Only]</i>	<i>HSG80> INITIALIZE DISK10300</i>
atlasD1	HSG80> INITIALIZE MIRRORD1
	HSG80> INITIALIZE DISK20100
	HSG80> INITIALIZE DISK20200
<i>[SC20 Only]</i>	<i>HSG80> INITIALIZE DISK20300</i>
atlasD2	HSG80> INITIALIZE MIRRORD2
	HSG80> INITIALIZE DISK30100
	HSG80> INITIALIZE DISK30200
<i>[SC20 Only]</i>	<i>HSG80> INITIALIZE DISK30300</i>
atlasD3	HSG80> INITIALIZE MIRRORD3
	HSG80> INITIALIZE DISK40100
	HSG80> INITIALIZE DISK40200
<i>[SC20 Only]</i>	<i>HSG80> INITIALIZE DISK40300</i>

Configure the System Storage on the HSG80

You have now created four Mirrorset storagesets, each with a usable capacity equivalent to an individual disk.

5. Create the virtual disks and introduce them as storage units, as shown in Table 3–11.

Table 3–11 Introducing the Storage Units (Virtual Disks) — HSG80

Domain	Command	UNIX Partition
atlasD0	HSG80> ADD UNIT D1 MIRRORD0 PREFERRED_PATH=THIS	/, /usr, /var
	HSG80> ADD UNIT D2 DISK10100 PREFERRED_PATH=THIS	generic boot
	HSG80> ADD UNIT D3 DISK10200 PREFERRED_PATH=THIS	backup/upgrade
	<i>[SC20 Only]</i> HSG80> ADD UNIT D4 DISK10300 PREFERRED_PATH=THIS	<i>Tru64</i>
atlasD1	HSG80> ADD UNIT D5 MIRRORD1 PREFERRED_PATH=OTHER	/, /usr, /var
	HSG80> ADD UNIT D6 DISK20100 PREFERRED_PATH=OTHER	generic boot
	HSG80> ADD UNIT D7 DISK20200 PREFERRED_PATH=OTHER	backup/upgrade
	<i>[SC20 Only]</i> HSG80> ADD UNIT D8 DISK20300 PREFERRED_PATH=OTHER	<i>Tru64</i>
atlasD2	HSG80> ADD UNIT D9 MIRRORD2 PREFERRED_PATH=THIS	/, /usr, /var
	HSG80> ADD UNIT D10 DISK30100 PREFERRED_PATH=THIS	generic boot
	HSG80> ADD UNIT D11 DISK30200 PREFERRED_PATH=THIS	backup/upgrade
	<i>[SC20 Only]</i> HSG80> ADD UNIT D12 DISK30300 PREFERRED_PATH=THIS	<i>Tru64</i>
atlasD3	HSG80> ADD UNIT D13 MIRRORD3 PREFERRED_PATH=OTHER	/, /usr, /var
	HSG80> ADD UNIT D14 DISK40100 PREFERRED_PATH=OTHER	generic boot
	HSG80> ADD UNIT D15 DISK40200 PREFERRED_PATH=OTHER	backup/upgrade
	<i>[SC20 Only]</i> HSG80> ADD UNIT D16 DISK40300 PREFERRED_PATH=OTHER	<i>Tru64</i>

Configure the System Storage on the HSG80

- 6. Set the IDENTIFIER label on the storage units, as shown in Table 3–12. This label simplifies the task of identifying disks when using the `hwmgr -v -d` command.

Table 3–12 Setting the Identifiers — HSG80

Domain	Command
atlasD0	HSG80> SET D1 IDENTIFIER=1
	HSG80> SET D2 IDENTIFIER=2
	HSG80> SET D3 IDENTIFIER=3
[SC20 Only]	HSG80> SET D4 IDENTIFIER=4
atlasD1	HSG80> SET D5 IDENTIFIER=1
	HSG80> SET D6 IDENTIFIER=2
	HSG80> SET D7 IDENTIFIER=3
[SC20 Only]	HSG80> SET D8 IDENTIFIER=4
atlasD2	HSG80> SET D9 IDENTIFIER=1
	HSG80> SET D10 IDENTIFIER=2
	HSG80> SET D11 IDENTIFIER=3
[SC20 Only]	HSG80> SET D12 IDENTIFIER=4
atlasD3	HSG80> SET D13 IDENTIFIER=1
	HSG80> SET D14 IDENTIFIER=2
	HSG80> SET D15 IDENTIFIER=3
[SC20 Only]	HSG80> SET D16 IDENTIFIER=4

Configure the System Storage on the HSG80

7. You must ensure that disks intended for use by a domain are visible only to that domain. Therefore, you should limit access to the storage units, as follows:
 - a. Disable access to all units, as shown in Table 3–13.

Table 3–13 Disabling Access to the Units

Domain	Command
atlasD0	HSG80> SET D1 DISABLE_ACCESS_PATH=ALL HSG80> SET D2 DISABLE_ACCESS_PATH=ALL HSG80> SET D3 DISABLE_ACCESS_PATH=ALL
[SC20 Only]	HSG80> SET D4 DISABLE_ACCESS_PATH=ALL
atlasD1	HSG80> SET D5 DISABLE_ACCESS_PATH=ALL HSG80> SET D6 DISABLE_ACCESS_PATH=ALL HSG80> SET D7 DISABLE_ACCESS_PATH=ALL
[SC20 Only]	HSG80> SET D8 DISABLE_ACCESS_PATH=ALL
atlasD2	HSG80> SET D9 DISABLE_ACCESS_PATH=ALL HSG80> SET D10 DISABLE_ACCESS_PATH=ALL HSG80> SET D11 DISABLE_ACCESS_PATH=ALL
[SC20 Only]	HSG80> SET D12 DISABLE_ACCESS_PATH=ALL
atlasD3	HSG80> SET D13 DISABLE_ACCESS_PATH=ALL HSG80> SET D14 DISABLE_ACCESS_PATH=ALL HSG80> SET D15 DISABLE_ACCESS_PATH=ALL
[SC20 Only]	HSG80> SET D16 DISABLE_ACCESS_PATH=ALL

Configure the System Storage on the HSG80

- b. Enable access for nodes that are connected to the storage, as shown in Table 3–14.

Table 3–14 Enabling Access to the Units for Connected Nodes — HSG80

Domain	Command
atlasD0	<pre>HSG80> SET D1 ENABLE_ACCESS_PATH=(N0000-0-0,N0000-0-1,N0000-1-0,N0000-1-1) HSG80> SET D1 ENABLE_ACCESS_PATH=(N0001-0-0,N0001-0-1,N0001-1-0,N0001-1-1) HSG80> SET D2 ENABLE_ACCESS_PATH=(N0000-0-0,N0000-0-1,N0000-1-0,N0000-1-1) HSG80> SET D2 ENABLE_ACCESS_PATH=(N0001-0-0,N0001-0-1,N0001-1-0,N0001-1-1) HSG80> SET D3 ENABLE_ACCESS_PATH=(N0000-0-0,N0000-0-1,N0000-1-0,N0000-1-1) HSG80> SET D3 ENABLE_ACCESS_PATH=(N0001-0-0,N0001-0-1,N0001-1-0,N0001-1-1)</pre>
[sc20] ¹	<pre>HSG80> SET D4 ENABLE_ACCESS_PATH=(N0000-0-0,N0000-0-1,N0000-1-0,N0000-1-1)</pre>
atlasD1	<pre>HSG80> SET D5 ENABLE_ACCESS_PATH=(N0032-0-0,N0032-0-1,N0032-1-0,N0032-1-1) HSG80> SET D5 ENABLE_ACCESS_PATH=(N0033-0-0,N0033-0-1,N0033-1-0,N0033-1-1) HSG80> SET D6 ENABLE_ACCESS_PATH=(N0032-0-0,N0032-0-1,N0032-1-0,N0032-1-1) HSG80> SET D6 ENABLE_ACCESS_PATH=(N0033-0-0,N0033-0-1,N0033-1-0,N0033-1-1) HSG80> SET D7 ENABLE_ACCESS_PATH=(N0032-0-0,N0032-0-1,N0032-1-0,N0032-1-1) HSG80> SET D7 ENABLE_ACCESS_PATH=(N0033-0-0,N0033-0-1,N0033-1-0,N0033-1-1)</pre>
[sc20] ¹	<pre>HSG80> SET D8 ENABLE_ACCESS_PATH=(N0032-0-0,N0032-0-1,N0032-1-0,N0032-1-1)</pre>
atlasD2	<pre>HSG80> SET D9 ENABLE_ACCESS_PATH=(N0064-0-0,N0064-0-1,N0064-1-0,N0064-1-1) HSG80> SET D9 ENABLE_ACCESS_PATH=(N0065-0-0,N0065-0-1,N0065-1-0,N0065-1-1) HSG80> SET D10 ENABLE_ACCESS_PATH=(N0064-0-0,N0064-0-1,N0064-1-0,N0064-1-1) HSG80> SET D10 ENABLE_ACCESS_PATH=(N0065-0-0,N0065-0-1,N0065-1-0,N0065-1-1) HSG80> SET D11 ENABLE_ACCESS_PATH=(N0064-0-0,N0064-0-1,N0064-1-0,N0064-1-1) HSG80> SET D11 ENABLE_ACCESS_PATH=(N0065-0-0,N0065-0-1,N0065-1-0,N0065-1-1)</pre>
[sc20] ¹	<pre>HSG80> SET D12 ENABLE_ACCESS_PATH=(N0064-0-0,N0064-0-1,N0064-1-0,N0064-1-1)</pre>
atlasD3	<pre>HSG80> SET D13 ENABLE_ACCESS_PATH=(N0096-0-0,N0096-0-1,N0096-1-0,N0096-1-1) HSG80> SET D13 ENABLE_ACCESS_PATH=(N0097-0-0,N0097-0-1,N0097-1-0,N0097-1-1) HSG80> SET D14 ENABLE_ACCESS_PATH=(N0096-0-0,N0096-0-1,N0096-1-0,N0096-1-1) HSG80> SET D14 ENABLE_ACCESS_PATH=(N0097-0-0,N0097-0-1,N0097-1-0,N0097-1-1) HSG80> SET D15 ENABLE_ACCESS_PATH=(N0096-0-0,N0096-0-1,N0096-1-0,N0096-1-1) HSG80> SET D15 ENABLE_ACCESS_PATH=(N0097-0-0,N0097-0-1,N0097-1-0,N0097-1-1)</pre>
[sc20] ¹	<pre>HSG80> SET D16 ENABLE_ACCESS_PATH=(N0096-0-0,N0096-0-1,N0096-1-0,N0096-1-1)</pre>

¹Note that only the first node of the domain has an access path to Tru64UNIX.

Configure the System Storage on the HSG80

Note:

If you are replacing a component (for example, a host adapter) or reseating a cable, you must update the access path as new connections will be created.

Table 3–15 on page 3–43 describes how these virtual disks are used.

Table 3–15 Allocating RAID Storagesets — HSG80

Virtual Disk Identifier	UNIX Name ¹	Partition	Recommended Size (GB) ²	Domain File System
IDENTIFIER=1	dsk3b	b	2	/
IDENTIFIER=1	dsk3g	g	18	/usr
IDENTIFIER=1	dsk3h	h	16	/var
IDENTIFIER=2	dsk4c	c	36	generic boot ³
IDENTIFIER=3	dsk5b	b	2	backup / ⁴
IDENTIFIER=3	dsk5g	g	18	backup /usr ⁴
IDENTIFIER=3	dsk5h	h	16	backup /var ⁴

¹ This UNIX naming convention applies if the ES40 or ES45 nodes are configured as specified; that is, three internal disks on the first node and two internal disks on the second node. For the DS20L product, the naming convention will be different

² The size has been rounded to the nearest gigabyte.

³ /cluster/admin/generic_boot_partition — used when adding domain members.

⁴ The backup disks are used during the upgrade process.

Note:

The identifier (IDENTIFIER=1) rather than the UNIX disk device name (dsk3) is used by the automatic installation process to identify the disks on the RAID (cluster disk, generic boot disk, and backup cluster disk).

For HP AlphaServer DS20L nodes, the usage of disk IDENTIFIER=4 is documented in Table 6–3 on page 6–6.

3.15 Configure the System Storage on the HSV110

This section describes how to configure system storage on the HSV110 storage subsystem and contains the following sections:

- HSV110 Storage Management (see Section 3.15.1 on page 3–46)
- Configure the HSV110 RAID Controllers (see Section 3.15.2 on page 3–47)

Note:

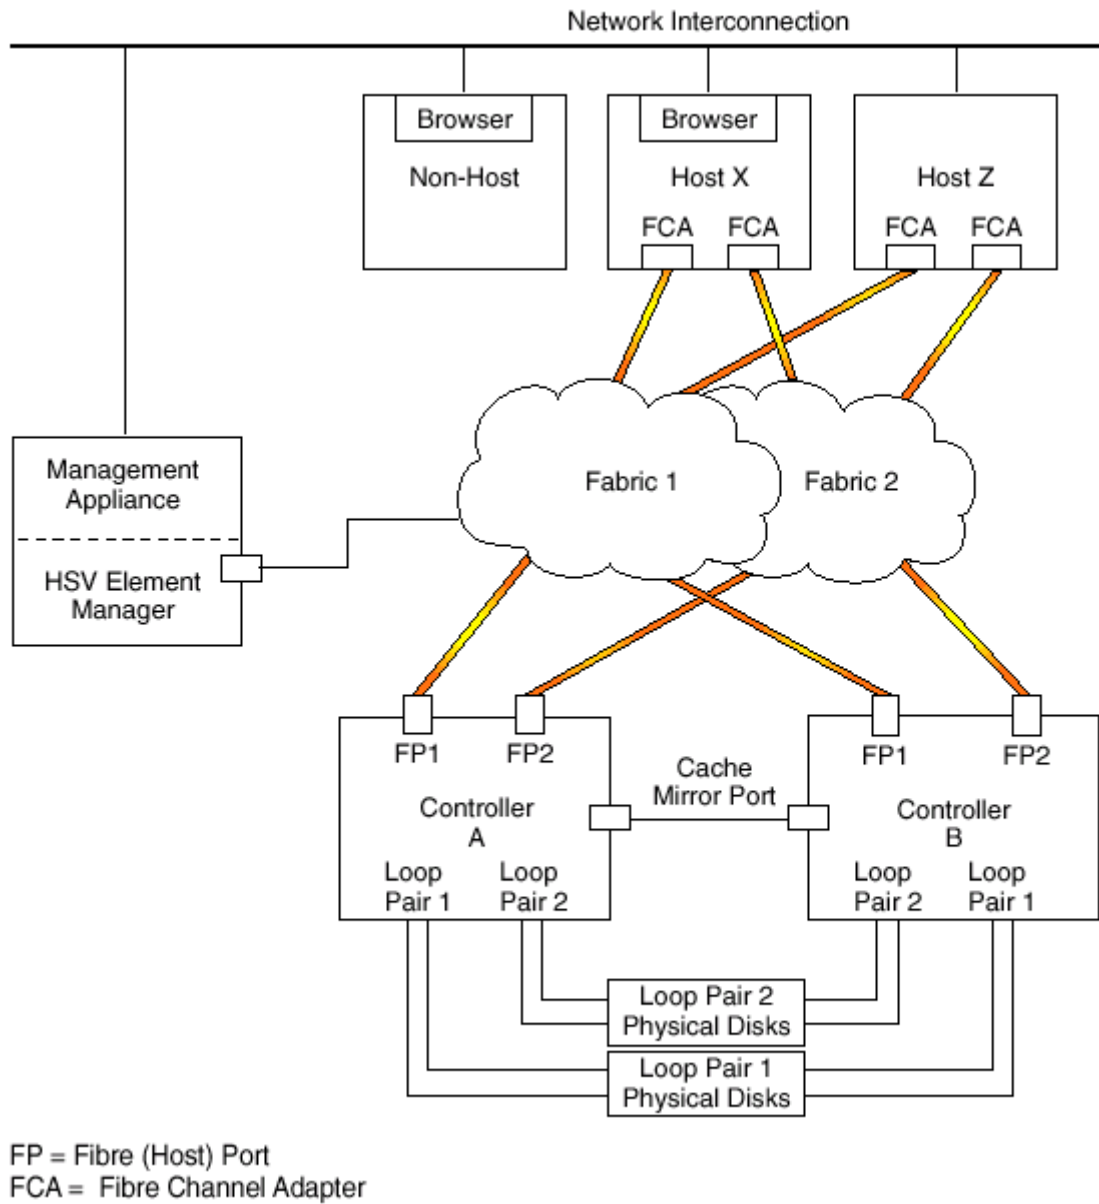
For instructions on how to configure fibre channel switch network addressing, see Section 3.13.1 on page 3–24; for instructions on how to configure fibre channel switch zoning, see Section 3.13.2 on page 3–26.

Before configuring the system storage on the HSV110, check that the system firmware on the HSV110 meets the requirements described in Section 3.12 on page 3–24.

Figure 3–11 shows a block diagram of how the HSV110 storage subsystem works, as follows:

- The HSV controller pair connects to two Fibre Channel fabrics, to which the hosts also connect.
- The HSV Element Manager is the software that controls the HSV110 storage subsystem. It resides on the SANworks Management Appliance. The SANworks Management Appliance connects into the fibre channel fabric.
- The controller pair connects to the physical disk array through fibre channel arbitrated loops. There are two separate loop pairs: loop pair 1 and loop pair 2. Each loop pair consists of 2 loops, each of which runs independently, but which can take over for the other loop in case of failure.

Configure the System Storage on the HSV110



CX07603A

Figure 3–11 Block Diagram of HSV110 Component Connections

Configure the System Storage on the HSV110

3.15.1 HSV110 Storage Management

The information in this section is organized as follows:

- Licenses (see Section 3.15.1.1 on page 3–46)
- The HSV Element Manager (see Section 3.15.1.2 on page 3–46)
- SANWorks Scripting Utility (see Section 3.15.1.3 on page 3–47)

For more detail on the HSV110 Array Controller, refer to the Enterprise Virtual Array Documentation, which is available at the following location:

`<http://h18006.www1.hp.com/products/storageworks/enterprise/documentation.html>`

3.15.1.1 Licenses

The HSV110 Controllers use Compaq SANWorks Virtual Controller Software (VCS) Package V2.002 and Compaq SANworks VCS Snapshot V2.002. These are licensed products. The VCS software kits contain a license authorization key which is used in conjunction with the controllers' WWN to generate a license key.

The license keys enable the use of VCS software and the value-added Snapshot functionality on a dual-controller basis. Snapshot for VCS requires, as a prerequisite, VCS Software and is licensed by the capacity attached to the dual controllers. After the storage configuration is initialized, the license is associated with the WWN and will persist if the controllers are replaced.

To enable licensing, the License Authorization Key from the VCS Software Kit and the VCS Snapshot Software Kit, and the WWN from the Enterprise Storage System WWN label are required. The customer may request a license key by visiting the licensing fulfillment website or using the manual methods outlined on the License Authorization Key Instruction Sheet.

Using the manual method, you will receive your license key in one to 48 hours. The key will be returned to you *via* email.

3.15.1.2 The HSV Element Manager

The HSV110 storage subsystem is monitored and managed centrally through the HSV Element Manager, a user-friendly graphical user interface (GUI). The element manager is installed on an HP SANworks Management Appliance that is attached to the fabric on which the hosts and the storage system reside. The element manager's client is a standard web browser. The element manager uses the paradigm of folders to organize the various storage system components.

3.15.1.3 SANWorks Scripting Utility

The SANWorks scripting utility is available from the Enterprise Platform Kit. Verify that the utility is installed in `/usr/sbin` and ensure that it has execute permissions.

The Scripting Utility V1.0 for Enterprise Virtual Array is a command line application that allows you to configure and control HSV controllers.

The Scripting Utility is a client of a client/server application. The Scripting Utility's server accepts commands from a client—the Scripting Utility. These commands are sent across a network to the HSV Element Manager. The Scripting Utility is also known as the Storage System Scripting Utility (SSSU).

Scripting utilities and SSSU script examples are also provided in the *HP AlphaServer SC System Software CD-ROM*. See Section 3.15.2.4 on page 3–55 and Section 3.15.2.5 on page 3–61 for information on editing and running these scripts.

3.15.2 Configure the HSV110 RAID Controllers

To configure the HSV110 RAID controllers, perform the following tasks:

- Locate the Example EVA (HSV110) Configuration Scripts (see Section 3.15.2.1 on page 3–48)
- Calculate the Minimum Storage Requirement (see Section 3.15.2.2 on page 3–48)
- Create Virtual Disk Units for the Clustered Management Server (see Section 3.15.2.3 on page 3–52)
- Edit the Configuration Script (see Section 3.15.2.4 on page 3–55)
- Run the Configuration Script (see Section 3.15.2.5 on page 3–61)

If you intend to install a clustered management server and you do not have any other Tru64 UNIX system available in the same network as your HSV110 storage subsystem, you must initialize the HSV110 through the SAN Appliance Element Manager browser and create the minimum entities needed to allow you to perform the initial Tru64 installation for the clustered management server. This is described in Section 3.15.2.3.

Configure the System Storage on the HSV110

3.15.2.1 Locate the Example EVA (HSV110) Configuration Scripts

Note:

The HSV110 belongs to the Enterprise Virtual Array (EVA) family of storage products. The scripts used to configure the HSV110 are called “EVA” configuration scripts.

Perform the following steps to get EVA configuration script examples for HP AlphaServer SC:

1. Insert the *HP AlphaServer SC System Software* CD-ROM in the disk drive.
2. Create a mount point for the CD-ROM, by running the following command:

```
# mkdir /cdrom
```
3. Mount the CD-ROM as follows:

```
# mount -r /dev/disk/cdrom0c /cdrom
```
4. Change to the directory where EVA scripting tools for Tru64 UNIX are located:

```
# cd /cdrom/Examples/eva_tools
```
5. List the files as follows:

```
# ls -l
```

```
1domain-nomgt.txt  
1domain-mgt.txt  
4domains-nomgt.txt  
4domains-mgt.txt
```

Note:

To configure the HSV110 through the Tru64 SSSU, you need to have completed the management server installation to the point where you have patched the management server. Otherwise, you can elect to use another external Tru64 UNIX system available in the same network as your HSV110 storage subsystem.

3.15.2.2 Calculate the Minimum Storage Requirement

As explained in Section 2.4.4.2 on page 2–14, there may be situations where it is advisable to separate virtual disks in different disk groups.

In such cases, use the following guidelines:

- Use the default disk group for holding system storage

Configure the System Storage on the HSV110

- Use the information provided in Table 3–16 on page 3–49, Table 3–17 on page 3–49, Table 3–18 on page 3–50, and Table 3–19 on page 3–50 to determine the minimum storage allocation needs for the domains in your system.
- Match your storage space needs with the information provided in Table 3–20 to determine the size of the default disk group.
- See Section 2.4.4.2 on page 2–14 for additional information about disk group considerations.

Table 3–16 Storage Requirements for a Single HP AlphaServer SC domain — HSV110

Disk	Required Space [GB]
CFS disk	16
Backup / Upgrade disk	16
Generic Boot	1
<i>Tru64 installation [SC20 only]</i>	<i>16</i>
Total SC45	33
<i>Total SC20</i>	<i>49</i>

Table 3–17 displays the storage management requirement for a clustered management server.

Table 3–17 Storage Requirement for the Clustered Management Server — HSV110

Disk	Required Space [GB]
CFS disk	16
Quorum disk	2
Tru64 installation disk	16
Node 0 boot disk	2
Node 1 boot disk	2
Total	38

Configure the System Storage on the HSV110

Table 3–18 displays the SC45 minimum storage requirements.

Table 3–18 SC45 Minimum Storage Requirements — HSV110

Number of domains	Without management server required space [GB]	With management server required space [GB]
1	33	71
2	66	104
3	99	137
4	132	170
5	165	203

Table 3–19 displays the SC20 minimal storage requirements.

Table 3–19 SC20 Minimum Storage Requirements — HSV110

Number of domains	Without management server required space [GB]	With management server required space [GB]
1	49	87
2	98	136
3	147	185
4	196	234
5	245	283

Table 3–20 displays the disk group virtual RAID1 space per number of disks.

Table 3–20 Disk Group Virtual RAID1 Available Space per Number of Disks — HSV110

Number of Disks [36 GB]	Available VR1 Space with Sparing Level 1 [GB]	Available VR1 Space with Sparing Level 2 [GB]
8	101	67
10	135	101
12	168	134
14	202	168

Configure the System Storage on the HSV110

Table 3–20 Disk Group Virtual RAID1 Available Space per Number of Disks — HSV110

Number of Disks [36 GB]	Available VR1 Space with Sparing Level 1 [GB]	Available VR1 Space with Sparing Level 2 [GB]
16	236	203
18	269	236
20	302	269
22	337	303

Table 3–21 summarizes the allocation of disks on the HSV110.

Table 3–21 Summary Allocation of Disks on the HSV110

Number of Domains	Total Number of 36GB disks with Sparing Level 2	Total Number of 36GB disks with Sparing Level 1	Total Number of 72GB disks with Sparing Level 1 or Level 2	Total Number of 146GB disks with Sparing Level 1 or Level 2
1	10	8	8	8
2	12	10	8	8
3	14	12	8	8
4	16	14	8	8
5	16	16	8	8

Configure the System Storage on the HSV110

3.15.2.3 Create Virtual Disk Units for the Clustered Management Server

In a clustered management server, all the disks are based on the HSV110 storage subsystem. This has the following consequences:

- The initial Tru64 UNIX system for the clustered management server cannot be installed before configuring the HSV110 storage system
- If no other Tru64 UNIX systems are available in the same network, you must perform a minimal HSV110 configuration through the SAN Appliance Element Manager browser. This minimal configuration is required to perform the initial Tru64 UNIX system installation for the clustered management server

To complete the HSV110 configuration, run the SSSU scripts as described in Section 3.15.2.4 and Section 3.15.2.5.

Section 3.15.2.3.1 describes minimal steps to be performed to configure the HSV110 through the SAN Appliance Element Manager browser.

Section 3.15.2.3.2 gives an SSSU script example to complete the storage configuration for the clustered management server.

3.15.2.3.1 Initialize the HSV110 Storage Subsystem through the Element Manager

Perform the following steps in order to configure the minimum storage needs for the initial Tru64 UNIX installation on the clustered management server:

Cell Initialization

To initialize the cell, perform the following tasks:

1. Log in the SAN Management appliance.
2. In the Navigation pane, select the Element Manager, HSV Element Manager option.
3. Click on the Launch button.
4. In the Navigation pane, select the Uninitialized Storage System option.
5. Click on the Initialize button.

Enter the cell name (`atlas` in the example) and the number of disks composing the default disk group

6. Click on the Advanced Options button.

The next steps allow you to optionally set the date format, the disk failure protection level (0, single or double), and enter a comment.

7. When completed, click on the Finish button and wait until completed.

The `atlas` cell is now displayed in the Navigation pane.

Configure the System Storage on the HSV110

Clustered Management Server Host Definition

To define the clustered management server host, perform the following tasks:

1. In the Navigation pane, select the `atlas` cell.
2. In the Navigation pane, select the `hosts` option.
3. Click on the Create Folder button.
4. Enter the folder name (SC in the example) and click on the Finish button.
5. Click on the Create Host button.
6. Enter the host name for the clustered management server (`atlasms` in the example).

Note:

The clustered management server host represents the full cluster, not a single node.

7. Click on the Next Step button.

The Add Host window displays a list of adapter port WWNs.

8. Select a WWN from the list, by comparing it with the SRM output from node 0 and node 1 of the clustered management server:

```
P00 >>> show device pg*
```

```
pga0.0.0.2.1PGA0
```

```
WWN 2000-0000-c92c-1c32
```

```
pgb0.0.0.3.1PGB0
```

```
WWN 2000-0000-c92c-1b05
```

Note:

When matching SRM output and Element Manager WWID information, please note that the SRM output displays the world wide node name (WWNN), where the leftmost character is 2. The Element Manager displays a menu of world wide port names (WWPN), where the leftmost character is 1.

9. Select the Tru64 UNIX operating system type for the host.
10. Click on the Finish button.
11. In the Navigation pane, select the created host, and click on the Add Port button.
12. Add the remaining WWN to complete the host definition.

Configure the System Storage on the HSV110

Tru64 Disk Unit Creation

To create the Tru64 UNIX disk unit, perform the following tasks:

1. On the Navigation pane, click on Virtual Disks.
2. Click on the Create Folder button.
3. Enter the folder name (SC in the example) and click on the Finish button.
4. Click on Create Virtual Disk Family button.
5. Create a virtual disk with the following characteristics:
 - Name: cfsms_Tru64
 - Redundancy: VRAID1
 - Write cache policy: mirrored write-back (default)
 - Read cache policy: on
 - OS unit ID: 3 (for this example)
 - Present to host: atlasms
 - Preferred path mode: No preference
6. Click on the Finish button.

The HSV110 storage subsystem is ready for the initial Tru64 UNIX system installation.

When the initial Tru64 UNIX installation is done, the SSSU utility can be used to complete the configuration of the HSV110 storage subsystem, as described in Section 3.15.2.4 on page 3–55 and Section 3.15.2.5 on page 3–61.

3.15.2.3.2 SSSU Script for the Clustered Management Server

For systems with a clustered management server with shared storage on an HSV110, Example 3–1 shows a suitable configuration script for creating the remaining virtual disk units for the clustered management server. This should be included in the script that will be created in Section 3.15.2.4. The amalgamated configuration script will be used to complete the configuration of the HSV110 in Section 3.15.2.5.

Example 3–1 Configuration Script Commands for the Clustered Management Server

```
!ADD HOST "\\Hosts\\SC\\atlasms"    WORLD_WIDE_NAME="1000-0000-c92c-1c32"
OPERATING_SYSTEM=TRU64
!SET HOST "\\Hosts\\SC\\atlasms"    ADD_WORLD_WIDE_NAME="1000-0000-c92c-1b05"
ADD STORAGE "\\Virtual Disks\\SC\\cfsms"    SIZE=16 REDUNDANCY=VRAID1 GROUP="\\Disk
Groups\\Default Disk Group" OS_UNIT_ID=1 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfsms_quorum"    SIZE=2 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=2 NOPREFERRED_PATH
```

Configure the System Storage on the HSV110

```
!ADD STORAGE "\Virtual Disks\SC\cfms Tru64"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\Disk Groups\Default Disk Group" OS_UNIT_ID=3 NOPREFERRED_PATH
ADD STORAGE "\Virtual Disks\SC\cfms_bootdisk1"  SIZE=2 REDUNDANCY=VRAID1
GROUP="\Disk Groups\Default Disk Group" OS_UNIT_ID=4 NOPREFERRED_PATH
ADD STORAGE "\Virtual Disks\SC\cfms_bootdisk2"  SIZE=2 REDUNDANCY=VRAID1
GROUP="\Disk Groups\Default Disk Group" OS_UNIT_ID=5 NOPREFERRED_PATH
  ADD LUN 91    HOST="\Hosts\SC\atlasms"        STORAGE="\Virtual
Disks\SC\cfms\ACTIVE"
  ADD LUN 92    HOST="\Hosts\SC\atlasms"        STORAGE="\Virtual
Disks\SC\cfms_quorum\ACTIVE"
  ADD LUN 93    HOST="\Hosts\SC\atlasms"        STORAGE="\Virtual
Disks\SC\cfms_backup\ACTIVE"
  ADD LUN 94    HOST="\Hosts\SC\atlasms"        STORAGE="\Virtual
Disks\SC\cfms_bootdisk1\ACTIVE"
  ADD LUN 95    HOST="\Hosts\SC\atlasms"        STORAGE="\Virtual
Disks\SC\cfms_bootdisk2\ACTIVE"
```

3.15.2.4 Edit the Configuration Script

Note:

If you intend to configure a clustered management server, Section 3.15.2.3 on page 3–52 describes the additional lines you will need to insert in the configuration script for allocating storage for the clustered management server.

Throughout this section, reference is made to the example SSSU configuration script in Example 3–2 on page 3–59. (Further example scripts are available on the *HP AlphaServer SC System Software* CD-ROM.)

Edit the following lines of the configuration script:

1. SELECT MANAGER command:

- Replace the IP address with the IP address of the SAN Appliance
- Check the username and password.

2. SELECT CELL "Uninitialized Storage system" command:

This command selects an uninitialized storage system. If you have already initialized the storage subsystem using the SAN Appliance Element Manager browser, you do not need to include the SELECT CELL command.

3. ADD CELL command:

If you have already initialized the storage subsystem using the SAN Appliance Element Manager browser, you do not need to include the ADD CELL command. Otherwise, modify the command as follows:

- Modify the cell name (atlas in the example) to customize the cell name.

Configure the System Storage on the HSV110

- b. Customize the `DEVICE_COUNT` according to the chosen number of disks for the default disk group; in this example, `DEVICE_COUNT` is equal to 28, which is the total number of disks for a fully equipped configuration of two controllers and two disk shelves (2C2D).
- c. Customize the sparing level.

Note:

In the example script (Example 3–2 on page 3–59), the HSV110 is configured with a unique disk group (the default disk group) including the totality of the disks of a 2C2D configuration (2 shelves, 28 disks). Section 3.15.2.2 provides information for calculating the minimum disk group size requirements if you wish to create a separate disk group for system storage.

4. `SELECT CELL cell name`

Change the *cell name* to the name of the storage cell as created with the `ADD CELL` command or the SAN Appliance Element Manager browser.

5. `ADD HOST` command

When modifying this command, note the following:

- A host is a domain, not a single node
- The `ADD HOST` command create a host, and there should be one `ADD HOST` command for each domain. The `WORLD_WIDE_NAME` option in the `ADD HOST` command specifies the WWN for the first port in the domain. (The `SET HOST` commands (see step 6) specify the WWN of the remaining HBAs in the domain.)

You should identify the WWN of each HBA, as follows:

```
P00 >>> show device pg
```

```
pga0.0.0.2.1      PGA0      WWN 2000-0000-c92c-1c5d
pgb0.0.0.3.1      PGB0      WWN 2000-0000-c92c-1b2e
```

Caution:

The HSV110 detects the world wide port name (WWPN), while SRM shows the world wide node name (WWNN) node name; for ES45 and DS20L the leftmost character of the WWPN is 1. When editing the script, use the WWPNs for the `WORLD_WIDE_OPTION` of the `ADD HOST` and `SET HOST` commands.

Configure the System Storage on the HSV110

Table 3–22 displays the world wide node and port names.

Table 3–22 World Wide Node and Port Names — HSV110

World Wide Node Name	World Wide Port Name
2000-0000-c92c-1c5d	1000-0000-c92c-1c5d

You can double-check the WWN of the HBA on the switch, as follows:

```
s1switch1:admin> switchshow
switchName:      s1switch1
switchType:      9.2
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     2
switchId:        fffc02
switchWwn:       10:00:00:60:69:51:35:9f
switchBeacon:    OFF
Zoning:          OFF
port  0: id N2 Online      F-Port 10:00:00:00:c9:2c:1c:5d
port  1: id N2 Online      F-Port 10:00:00:00:c9:2c:1b:2e
port  2: id N2 Online      F-Port 10:00:00:00:c9:2c:1b:05
port  3: -- N2 No_Module
port  4: id N2 Online      F-Port 10:00:00:00:c9:2c:1c:45
port  5: -- N2 No_Module
port  6: id N2 Online      F-Port 10:00:00:00:c9:2c:1b:ff
:
:
port 12: -- N2 No_Module
port 13: id N2 Online      F-Port 50:00:1f:e1:00:13:a3:cc
port 14: id N2 Online      F-Port 50:00:1f:e1:00:13:a3:c9
port 15: id N1 Online      F-Port 10:00:00:00:c9:23:2e:f9
s1switch1:admin>
```

For each domain, change the `WORLD_WIDE_NAME` option of an `ADD HOST` command to the WWN of the HBA in the first node of the domain.

Note:

If you have a host name that is a prefix for another host name (example `atlas`, `atlasms`), insert the command for creating the prefix (`add host "\Hosts\SC\atlas"`) **BEFORE** the command to create the host name (`"\Hosts\SC\atlasms"`). If you change the order, the second `ADD HOST` command will fail.

Configure the System Storage on the HSV110

6. SET HOST commands

When modifying these commands, note the following:

- The SET HOST command adds a port to the host. This port can be either the second HBA of the first node of the domain or one HBA of the second node of the same domain.
- There is one SET HOST command for each HBA in the domain that is additional to the HBA defined in the ADD HOST command (the first port in the domain). In Example 3–2 on page 3–59, there are a total of 16 HBAs in the system, with four HBAs in each of four domains. This results in four ADD HOST commands — one for each domain — and 12 SET HOST commands — three for each domain.)

For each domain, edit one SET HOST command for each remaining HBA in the domain. Change the WORLD_WIDE_NAME option of the SET HOST command to the WWN of the HBA.

7. ADD STORAGE commands

Edit these commands to allocate the storage you want to configure for your system. In Example 3–2 on page 3–59, the following allocations are made:

- 16 GB VRAID1 for the CFS disk
- 1 GB VRAID1 for the generic boot disk
- 16 GB VRAID1 for the backup/upgrade disk

Do not modify the OS_UNIT_ID options: they are coherent with sra setup settings.

For SC20 configurations, uncomment the ADD STORAGE command for creating the Tru64 disk.

8. ADD LUN commands:

Edit these commands to present the disk units to the hosts in your system.

For SC20 configurations, uncomment the ADD LUN commands to present the Tru64 disks.

Configure the System Storage on the HSV110

Example 3–2 shows the sample SSSU configuration script.

Example 3–2 SSSU Configuration Script

```
! Script to initialize and config atlas via 16.189.121.99
! Using SSSU (SANscript)
! Hewlett Packard
! Date: 10/janv/2003 15:34
!
SELECT MANAGER 16.189.121.99 USERNAME=administrator PASSWORD=administrator
!
! Stop if errors do occur

SET OPTIONS ON_ERROR=HALT_ON_ERROR
! Create the Default Disk Group if requested
SET OPTIONS COMMAND_DELAY=60
SELECT CELL "Uninitialized Storage System1"
ADD CELL atlas DEVICE_COUNT=28 SPARE_POLICY=DOUBLE
SELECT CELL atlas
SET OPTIONS COMMAND_DELAY=10

! Create one host with all adapters WWN per domain
ADD FOLDER "\\Hosts\\SC"
!
ADD HOST "\\Hosts\\SC\\atlasD0" WORLD_WIDE_NAME="1000-0000-c92c-1c5d"
OPERATING_SYSTEM=TRU64
SET HOST "\\Hosts\\SC\\atlasD0" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1b2e"
SET HOST "\\Hosts\\SC\\atlasD0" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1b2f"
SET HOST "\\Hosts\\SC\\atlasD0" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1b19"
!
ADD HOST "\\Hosts\\SC\\atlasD1" WORLD_WIDE_NAME="1000-0000-c92a-490b"
OPERATING_SYSTEM=TRU64
SET HOST "\\Hosts\\SC\\atlasD1" ADD_WORLD_WIDE_NAME="1000-0000-c92c-2cc8"
SET HOST "\\Hosts\\SC\\atlasD1" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1bc7"
SET HOST "\\Hosts\\SC\\atlasD1" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1bc8"
!
ADD HOST "\\Hosts\\SC\\atlasD2" WORLD_WIDE_NAME="1000-0000-c92c-1c63"
OPERATING_SYSTEM=TRU64
SET HOST "\\Hosts\\SC\\atlasD2" ADD_WORLD_WIDE_NAME="1000-0000-c92c-1b05"
SET HOST "\\Hosts\\SC\\atlasD2" ADD_WORLD_WIDE_NAME="1000-0000-c927-c9b3"
SET HOST "\\Hosts\\SC\\atlasD2" ADD_WORLD_WIDE_NAME="1000-0000-c927-cb14"

ADD HOST "\\Hosts\\SC\\atlasD3" WORLD_WIDE_NAME="1000-0000-c92c-1c45"
OPERATING_SYSTEM=TRU64
SET HOST "\\Hosts\\SC\\atlasD3" ADD_WORLD_WIDE_NAME="1000-0000-c92c-2bee"
SET HOST "\\Hosts\\SC\\atlasD3" ADD_WORLD_WIDE_NAME="1000-0000-c927-ccee"
SET HOST "\\Hosts\\SC\\atlasD3" ADD_WORLD_WIDE_NAME="1000-0000-c925-6ad9"

!
! Create Virtual disks
ADD FOLDER "\\Virtual Disks\\SC"

!
```

Configure the System Storage on the HSV110

```
ADD STORAGE "\\Virtual Disks\\SC\\cfs0"    SIZE=16 REDUNDANCY=VRAID1 GROUP="\\Disk
Groups\\Default Disk Group" OS_UNIT_ID=1 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs0_generic"    SIZE=1 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=2 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs0_backup"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=3 NOPREFERRED_PATH
! Uncomment next line for SC20 configurations
! ADD STORAGE "\\Virtual Disks\\SC\\cfs0_Tr64"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=4 NOPREFERRED_PATH

ADD STORAGE "\\Virtual Disks\\SC\\cfs1"    SIZE=16 REDUNDANCY=VRAID1 GROUP="\\Disk
Groups\\Default Disk Group" OS_UNIT_ID=1 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs1_generic"    SIZE=1 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=2 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs1_backup"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=3 NOPREFERRED_PATH
! Uncomment next line for SC20 configurations
! ADD STORAGE "\\Virtual Disks\\SC\\cfs1_Tr64"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=4 NOPREFERRED_PATH

ADD STORAGE "\\Virtual Disks\\SC\\cfs2"    SIZE=16 REDUNDANCY=VRAID1 GROUP="\\Disk
Groups\\Default Disk Group" OS_UNIT_ID=1 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs2_generic"    SIZE=1 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=2 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs2_backup"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=3 NOPREFERRED_PATH
! Uncomment next line for SC20 configurations
! ADD STORAGE "\\Virtual Disks\\SC\\cfs2_Tr64"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=4 NOPREFERRED_PATH

ADD STORAGE "\\Virtual Disks\\SC\\cfs3"    SIZE=16 REDUNDANCY=VRAID1 GROUP="\\Disk
Groups\\Default Disk Group" OS_UNIT_ID=1 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs3_generic"    SIZE=1 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=2 NOPREFERRED_PATH
ADD STORAGE "\\Virtual Disks\\SC\\cfs3_backup"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=3 NOPREFERRED_PATH
! Uncomment next line for SC20 configurations
! ADD STORAGE "\\Virtual Disks\\SC\\cfs3_Tr64"    SIZE=16 REDUNDANCY=VRAID1
GROUP="\\Disk Groups\\Default Disk Group" OS_UNIT_ID=4 NOPREFERRED_PATH

! Add unit numbers to the SC's and activate them

ADD LUN 1    HOST="\\Hosts\\SC\\atlasD0"    STORAGE="\\Virtual
Disks\\SC\\cfs0\\ACTIVE"
ADD LUN 2    HOST="\\Hosts\\SC\\atlasD0"    STORAGE="\\Virtual
Disks\\SC\\cfs0_generic\\ACTIVE"
ADD LUN 3    HOST="\\Hosts\\SC\\atlasD0"    STORAGE="\\Virtual
Disks\\SC\\cfs0_backup\\ACTIVE"
! Uncomment next line for SC20 configurations
!ADD LUN 4    HOST="\\Hosts\\SC\\atlasD0"    STORAGE="\\Virtual
Disks\\SC\\cfs0_Tr64\\ACTIVE"
```

Configure the System Storage on the HSV110

```
ADD LUN 11    HOST="\Hosts\SC\atlasD1"    STORAGE="\Virtual
Disks\SC\cfs1\ACTIVE"
ADD LUN 12    HOST="\Hosts\SC\atlasD1"    STORAGE="\Virtual
Disks\SC\cfs1_generic\ACTIVE"
ADD LUN 13    HOST="\Hosts\SC\atlasD1"    STORAGE="\Virtual
Disks\SC\cfs1_backup\ACTIVE"
! Uncomment next line for SC20 configurations
!ADD LUN 14    HOST="\Hosts\SC\atlasD1"    STORAGE="\Virtual
Disks\SC\cfs1_Tru64\ACTIVE"

ADD LUN 21    HOST="\Hosts\SC\atlasD2"    STORAGE="\Virtual
Disks\SC\cfs2\ACTIVE"
ADD LUN 22    HOST="\Hosts\SC\atlasD2"    STORAGE="\Virtual
Disks\SC\cfs2_generic\ACTIVE"
ADD LUN 23    HOST="\Hosts\SC\atlasD2"    STORAGE="\Virtual
Disks\SC\cfs2_backup\ACTIVE"
! Uncomment next line for SC20 configurations
!ADD LUN 24    HOST="\Hosts\SC\atlasD2"    STORAGE="\Virtual
Disks\SC\cfs2_Tru64\ACTIVE"

ADD LUN 31    HOST="\Hosts\SC\atlasD3"    STORAGE="\Virtual
Disks\SC\cfs3\ACTIVE"
ADD LUN 32    HOST="\Hosts\SC\atlasD3"    STORAGE="\Virtual
Disks\SC\cfs3_generic\ACTIVE"
ADD LUN 33    HOST="\Hosts\SC\atlasD3"    STORAGE="\Virtual
Disks\SC\cfs3_backup\ACTIVE"
! Uncomment next line for SC20 configurations
!ADD LUN 34    HOST="\Hosts\SC\atlasD3"    STORAGE="\Virtual
Disks\SC\cfs3_Tru64\ACTIVE"
!
```

3.15.2.5 Run the Configuration Script

When you have completed edits to the EVA configuration script, run the scripting utility either from the HP AlphaServer SC management server, or from a PC connected to the same IP network as your HSV110 storage subsystem.

Note:

In order to run SSSU from a PC you need the EVA Platform kit for Windows NT.

The following steps apply to running the scripting utility from the HP AlphaServer SC management server:

1. Run the SSSU scripting utility from anywhere on the same IP network as your HSV110 storage subsystem, as follows:

```
atlasms# ./sssu
```

The NoCellSelected prompt will be displayed.

Configure the System Storage on the MSA1000

2. Run the configuration script (in the example below, a script for 4 domains without clustered management server):

```
NoCellSelected > file 4domains-nomgt
```

The script may take several minutes to execute due to the number of scripted commands.

3. Verify the configuration on the SAN Appliance PC using the Graphical User Interface.

Note:

Script execution may take a long time. It may be quicker to manually execute script lines one at a time.

3.16 Configure the System Storage on the MSA1000

Note:

The HP StorageWorks Modular SAN Array 1000 (MSA1000) is the entry-level storage subsystem for HP AlphaServer SC systems and is only supported for systems that have one domain. If you have more than one domain, it is more cost-effective to use an alternative type of storage subsystem.

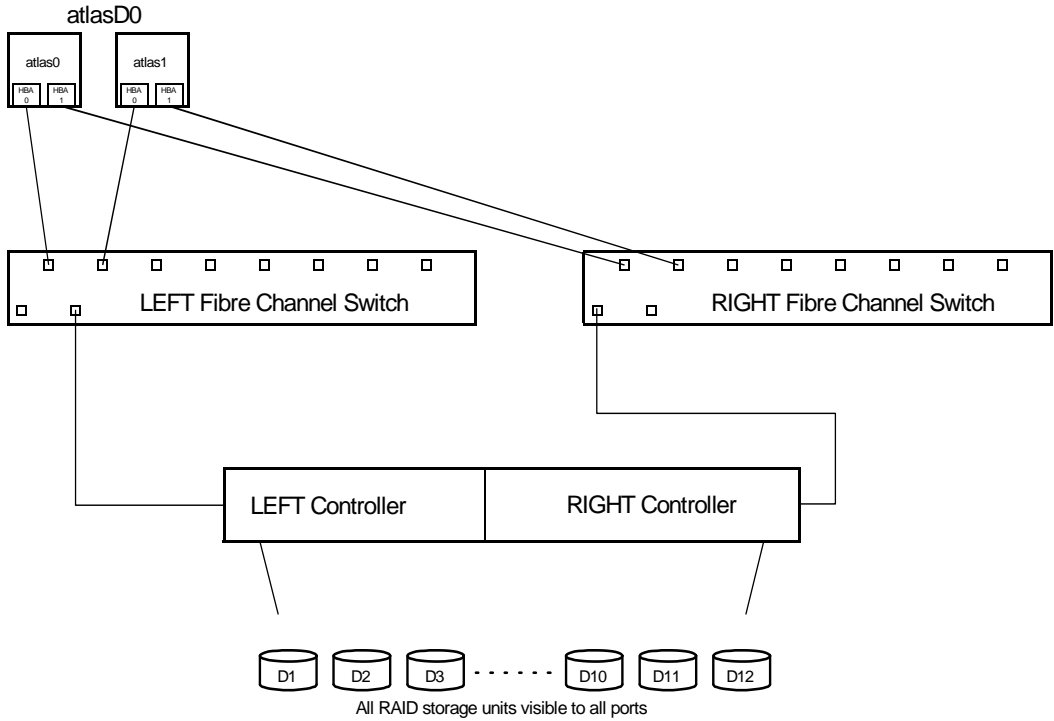
In an MSA1000 storage subsystem, two RAID controllers are used, as shown in Figure 3–12 on page 3–63, where the two controllers are labelled LEFT and RIGHT. You can configure the RAID subsystem using a command line interpreter (CLI) on one of the controllers (slot 1, on the RIGHT controller in Figure 3–12).

To configure the system storage on an MSA1000 storage subsystem, perform the following tasks on each domain:

- Check the MSA1000 Controller Firmware (see Section 3.16.1 on page 3–64)
- Configure the MSA1000 RAID Storage (see Section 3.16.2 on page 3–64)
- Rename the Connections (see Section 3.16.3 on page 3–66)

Configure the System Storage on the MSA1000

Figure 3–12 shows the cabling in the example MSA1000 storage subsystem configured in this section.



HBA = Fibre Channel Host Bus Adapter

Figure 3–12 Example System Storage Configuration — MSA1000

Configure the System Storage on the MSA1000

3.16.1 Check the MSA1000 Controller Firmware

Before configuring the system storage on the MSA1000, check that the system firmware on the MSA1000 meets the requirements described in Section 3.12 on page 3–24.

If necessary, see Appendix J.1. for information on upgrading the MSA1000 Controller Firmware.

3.16.2 Configure the MSA1000 RAID Storage

Configure the MSA1000 RAID controllers as follows:

1. Connect a terminal *via* serial cable to slot 1 of one of the MSA1000 controllers and use a terminal emulator to issue commands to the CLI. In the example in this section:
 - The MSA1000 is equipped with an additional dual-bus enclosure, for a total of 4 SCSI buses.
 - System storage uses two 36.4 GB disks.

2. Using the CLI, view the installed physical disks as follows:

```
CLI> show disks
```

Disk List:	(box,bay)	(bus,ID)	Size	Units
Disk101	(1,01)	(0,00)	36.4GB	none
Disk102	(1,02)	(0,01)	36.4GB	none
Disk103	(1,03)	(0,02)	36.4GB	none
Disk104	(1,04)	(0,03)	36.4GB	none
Disk108	(1,08)	(1,00)	36.4GB	none
Disk109	(1,09)	(1,01)	36.4GB	none
Disk110	(1,10)	(1,02)	36.4GB	none
Disk111	(1,11)	(1,03)	36.4GB	none
Disk112	(1,13)	(1,04)	36.4GB	none
Disk201	(2,01)	(2,00)	36.4GB	none
Disk202	(2,02)	(2,02)	36.4GB	none
Disk203	(2,03)	(2,03)	36.4GB	none
Disk204	(2,05)	(2,04)	36.4GB	none
Disk208	(2,08)	(3,00)	36.4GB	none
Disk209	(2,09)	(3,01)	36.4GB	none
Disk210	(2,10)	(3,02)	36.4GB	none
Disk211	(2,11)	(3,03)	36.4GB	none

3. Create the virtual disks and introduce them as storage units, as shown in Table 3–23.

For a clustered management server, create the following additional storage units:

- Unit 4: CFS disk
- Unit 5: Quorum disk
- Unit 6: Tru64 disk
- Unit 7: Node 0 boot
- Unit 8: Node 1 boot

Configure the System Storage on the MSA1000

Table 3–23 Introducing the Storage Units (Virtual Disks) — MSA1000

Domain	Command	UNIX Partition
atlasD0	CLI> ADD UNIT 1 RAID_LEVEL=1 DATA="disk101 disk201" SIZE=16GB SPARE="disk112"	/, /usr, /var
	CLI> ADD UNIT 2 RAID_LEVEL=1 DATA="disk102 disk202" SIZE=1500MB SPARE="disk112"	generic boot
	CLI> ADD UNIT 3 RAID_LEVEL=1 DATA="disk103 disk203" SIZE=16GB SPARE="disk112"	backup/upgrade
atlasms	CLI> ADD UNIT 5 RAID_LEVEL=1 DATA="disk104 disk204" size=16GB SPARE="disk112"	CFS disk
	CLI> ADD UNIT 8 RAID_LEVEL=1 DATA="disk108 disk208" size=2G SPARE="disk112"	quorum
	CLI> ADD UNIT 4 RAID_LEVEL=1 DATA="disk109 disk209" size=12GB SPARE="disk112"	Tru64
	CLI> ADD UNIT 6 RAID_LEVEL=1 DATA="disk110 disk210" size=2GB SPARE="disk112"	node 0 boot
	CLI> ADD UNIT 7 RAID_LEVEL=1 DATA="disk111 disk211" size=2GB SPARE="disk112"	node 1 boot ¹

¹For sizes shown in this example, when creating the last unit you may have to adjust the size to fit the remaining available space.

- Set the IDENTIFIER label on the storage units, as shown in Table 3–24. This label simplifies the task of identifying disks when using the `hwmgr -v -d` command.

Table 3–24 Setting the Identifiers — MSA1000

Domain	Command
atlasD0	CLI> SET UNIT_ID 1 1
	CLI> SET UNIT_ID 2 2
	CLI> SET UNIT_ID 3 3
atlasms	CLI> SET UNIT_ID 4 1
	CLI> SET UNIT_ID 5 2

Configure the System Storage on the MSA1000

Table 3–24 Setting the Identifiers — MSA1000

Domain	Command
	<i>CLI> SET UNIT_ID 6 3</i>
	<i>CLI> SET UNIT_ID 7 4</i>
	<i>CLI> SET UNIT_ID 8 5</i>

3.16.3 Rename the Connections

Before you begin to modify the connections, check that the system firmware on the MSA1000 meets the requirements described in Section 3.12 on page 3–24. Rename the connections as follows:

1. Compare the HOST_ID (WWN) reported by the console `show device` command with the connections reported by the controller `SHOW CONNECTIONS` command.

In this example, `atlas0` and `atlas1` each have two fibre channel Host Bus Adapter cards (HBAs). To identify the `atlas0` connections to the LEFT and RIGHT controllers, perform the following steps:

- a. Identify the WWN of each HBA (see Table 3–25), by running the `show device` command at the `atlas0` SRM console prompt:

```
P00>>> show devices pg
pga0.0.0.2.1      PGA0      WWN 1000-0000-c92c-1bba
pgb0.0.0.3.1      PGB0      WWN 1000-0000-c92c-1bff
```

Table 3–25 Identifying the WWN of the HBAs — MSA1000

Domain	HBA	WWN
atlas0	0	2000-0000-c92c-1bba
	1	2000-0000-c92c-1bff
atlas1	0	2000-0000-c92c-1c72
	1	2000-0000-c92c-1c12

Configure the System Storage on the MSA1000

For a clustered management server, the WWNs in Table 3–26 may be identified:

Table 3–26 Identifying the WWN of the HBAs — MSA1000 — Management Server

Domain	HBA	WWN
<i>atlasms0</i>	<i>0</i>	<i>2000-0000-c92c-1c32</i>
	<i>1</i>	<i>2000-0000-c92c-1c29</i>
<i>atlasms1</i>	<i>0</i>	<i>2000-0000-c92c-1c42</i>
	<i>1</i>	<i>2000-0000-c92c-1c51</i>

- b. Create a naming convention for the connections, to give them meaningful names. The convention used in this example is N<node#>-<adapter#>

where:

- *node#* is the node number (four-digits, left-pad with zeros as necessary)
- *adapter#* is **0** for HBA 0, and **1** for HBA 1

Note:

These two fields are sufficient to fully identify the connection, as each MSA1000 controller has only one port.

- c. List all of the connections (from all HBAs in all nodes) to the MSA1000 controllers, by running the `SHOW CONNECTIONS` command:

```
CLI> SHOW CONNECTIONS
Connection Name: <Unknown>
  Host WWNN = 20000000-C92C1C72
  Host WWPN = 10000000-C92C1C72
  Profile Name = Default
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 2 Port 1 Status = Online

Connection Name: <Unknown>
  Host WWNN = 20000000-C92C1BBA
  Host WWPN = 10000000-C92C1BBA
  Profile Name = Tru64
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 1 Port 1 Status = Online
```

Configure the System Storage on the MSA1000

```
Connection Name: <Unknown>
  Host WWNN = 20000000-C92C1BFF
  Host WWPN = 10000000-C92C1BFF
  Profile Name = Tru64
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 2 Port 1 Status = Online
```

```
Connection Name: <Unknown>
  Host WWNN = 20000000-C92C1C12
  Host WWPN = 10000000-C92C1C12
  Profile Name = Tru64
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 2 Port 1 Status = Online
```

...

If you have more domains, or a clustered management server connected to the same fabric, information similar to the output above is displayed for all nodes equipped with HBAs and connected to the same fabric.

- d. Rename the connections using the ADD CONNECTION command:

```
CLI> ADD CONNECTION N0000-0 WWPN = 10000000-C92C1BBA
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION N0000-1 WWPN = 10000000-C92C1BFF
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION N0001-0 WWPN = 10000000-C92C1C72
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION N0001-1 WWPN = 10000000-C92C1C12
  PROFILE=Tru64 OFFSET=0
```

For a clustered management server, rename the connections as follows:

```
CLI> ADD CONNECTION MGT0000-0 WWPN = 10000000-C92C1C32
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION MGT0000-1 WWPN = 10000000-C92C1C29
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION MGT0001-0 WWPN = 10000000-C92C1C42
  PROFILE=Tru64 OFFSET=0
CLI> ADD CONNECTION MGT0001-1 WWPN = 10000000-C92C1C51
  PROFILE=Tru64 OFFSET=0
```

Configure the System Storage on the MSA1000

Repeat the `ADD CONNECTION` command for each node connected in the same fabric, as shown in Table 3–27.

Table 3–27 Renaming the Connections — MSA1000

Domain	Node	HBA	Controller	Connection Name
atlasD0	0	0	LEFT	N0000-0-0
	0	1	RIGHT	N0000-0-1
	1	0	LEFT	N0000-1-0
	1	1	RIGHT	N0000-1-1
atlasms0		0	LEFT	MGT0000-0
		1	RIGHT	MGT0000-1
atlasms1		0	LEFT	MGT0001-0
		1	RIGHT	MGT0001-1

- You must ensure that disks intended for use by a domain are visible only to that domain. To do this, set the Access Control List (ACL) for the devices, as shown in Table 3–28.

Table 3–28 Enabling Access to the Units for Connected Nodes — MSA1000

Domain	Command
atlasD0	<pre> CLI> ADD ACL CONNECTION=N0000-0 UNIT=1-3 CLI> ADD ACL CONNECTION=N0000-1 UNIT=1-3 CLI> ADD ACL CONNECTION=N0001-0 UNIT=1-3 CLI> ADD ACL CONNECTION=N0001-1 UNIT=1-3 </pre>
atlasms	<pre> CLI> ADD ACL CONNECTION=MGT0000-0 UNIT=4-8 CLI> ADD ACL CONNECTION=MGT0000-1 UNIT=4-8 CLI> ADD ACL CONNECTION=MGT0001-0 UNIT=4-8 CLI> ADD ACL CONNECTION=MGT0001-1 UNIT=4-8 </pre>

Note:

If you are replacing a component (for example, a host adapter) or reseating a cable, you must update the access path as new connections will be created.

Table 3–15 on page 3–43 describes how these virtual disks are used.

For sample flow and output when configuring the MSA1000, refer to Appendix J.3.

Upgrade Installation Procedure

This chapter describes how to upgrade to HP AlphaServer SC Version 2.6 (UK2) from HP AlphaServer SC Version 2.6 (UK1).

If your system has a management server, read Section 4.4 on page 4–10 (Upgrade with a Management Server). If your system has a clustered management server, read Section 4.5 on page 4–21 (Upgrade with a Clustered Management Server). If your system does not have a management server, read Section 4.6 on page 4–32 (Upgrade without a Management Server).

The information in this chapter is structured as follows:

- Understanding the Upgrade Process (see Section 4.1 on page 4–2)
- Pre-Upgrade Audit (see Section 4.2 on page 4–8)
- Preparing RIS for C2 Security (see Section 4.3 on page 4–10)
- Upgrade with a Management Server (see Section 4.4 on page 4–10)
- Upgrade with a Clustered Management Server (see Section 4.5 on page 4–21)
- Upgrade without a Management Server (see Section 4.6 on page 4–32)
- Upgrading the Domains (see Section 4.7 on page 4–41)
- Recover from Failures during an Upgrade (see Section 4.8 on page 4–45)
- Post-Upgrade Tasks (see Section 4.9 on page 4–49)

Understanding the Upgrade Process

4.1 Understanding the Upgrade Process

The information in this section is organized as follows:

- Upgrade Versus Installation (see Section 4.1.1 on page 4–2)
- Upgrade Procedure Overview (see Section 4.1.2 on page 4–3)
- Upgrade Restart (see Section 4.1.3 on page 4–4)
- Upgrade States (see Section 4.1.4 on page 4–4)
- Preserved and Unpreserved Files (see Section 4.1.5 on page 4–5)
- Special Upgrade Mechanisms on Domains (see Section 4.1.6 on page 4–6)

4.1.1 Upgrade Versus Installation

There are two types of HP AlphaServer SC installation:

- **Full installation:** A full installation must be performed when a new system is to be created. See Chapter 5 or Chapter 6 as appropriate for more information.
- **Upgrade installation:** An upgrade is performed when a previously installed HP AlphaServer SC system is to be upgraded to a new version of the HP AlphaServer SC software.

Note:

The only earlier installation that is supported for upgrade to HP AlphaServer SC Version 2.6 (UK2) is HP AlphaServer SC Version 2.6 (UK1).

If you have a system with an older version of HP AlphaServer SC, and wish to upgrade to HP AlphaServer SC Version 2.6 (UK2), please contact your local HP support representative.

For the upgrade from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), if your system has a management server (or clustered management server) you need not upgrade the Tru64 UNIX operating system on the management server.

If your system does not have a management server, the steps for upgrading the first domain are different. The steps for upgrading the first domain are detailed in Section 4.6.2 on page 4–32.

During an upgrade, the `sra upgrade` command is used to upgrade the domains.

The `sra upgrade` command performs a series of tasks, as follows:

- The Tru64 UNIX patch kits are applied.
- The HP AlphaServer SC software is updated.
- The members are not deleted from the domain, and the Tru64 UNIX and TruCluster Server kits are not updated on any of the domain members.

The tasks to be performed by the `sra upgrade` command are automatically determined.

4.1.2 Upgrade Procedure Overview

The upgrade procedure is lengthy. Detailed instructions are provided later in this chapter, and these instructions should be followed very carefully.

The upgrade procedure can be summarized as follows:

1. Perform a pre-upgrade audit to identify any problems before the system is taken out of production. This audit can be carried out during production and helps to minimize the time that the system is out of production.
2. Back up the management server, or the clustered management server, or the first domain where no management server exists.
3. Complete the pre-upgrade steps.

From this point until the end of the upgrade, the full system will be out of production.

4. Upgrade the management server, or the clustered management server, or the first domain where no management server exists.

At this point, the management server will be upgraded, but the domains will remain out of production.

5. Back up the domains.
6. Upgrade the domains.
7. Complete the post-upgrade tasks.

At this point, the full system will be back in production.

If your system has a management server, follow the instructions in Section 4.4 on page 4–10. If your system has a clustered management server, follow the instructions in Section 4.5 on page 4–21. If your system does not have a management server, follow the instructions in Section 4.6 on page 4–32.

Understanding the Upgrade Process

4.1.3 Upgrade Restart

A state-based mechanism that controls the upgrade process allows you to restart the upgrade at various points. However, if certain upgrade steps fail or are interrupted, the affected domains will need to be restored as described in Section 4.8 on page 4–45, before the upgrade can be restarted.

4.1.4 Upgrade States

The upgrade state applies to the domain and not to the node. During an upgrade, a domain will move through several upgrade states from `Pre_Upgrade` to `Upgraded` as described in Table 4–1.

As each upgrade step is successfully performed, the current upgrade state for the domain (the `current_upg_state` in the `sc_domain` table) is updated with the name of the state just completed.

If an upgrade step fails and the upgrade is restarted, the `current_upg_state`, which contains the last successful upgrade step, is used to determine the next upgrade step to perform.

The upgrade process is controlled by the `sra upgrade` command and not by the `sra daemon` (`srad`). For this reason, do not interrupt the `sra upgrade` command while the upgrade is in progress.

Table 4–1 Upgrade States

Upgrade State	Description
<code>Pre_Upgrade</code>	This is the initial state of the domain.
<code>Upg_Installed</code>	The upgrade subset has been installed on the domain.
<code>Checked</code>	The check phase has been completed on the domain.
<code>Setup</code>	The setup phase to prepare the domain for upgrade has been completed.
<code>Installed</code>	The Tru64 UNIX patch kit and the HP AlphaServer SC software have been installed on the domain.
<code>Upgraded</code>	The cleanup has been performed and the system upgrade has been completed.

Note:

For more information about the typical actions that happen during each upgrade state, and where you can access log files to determine the cause of an error, see Section 11.24 on page 11–40.

4.1.4.1 Check the Upgrade State

To check the upgrade state of the domain, use the `rmsquery` command, as follows:

```
atlasms# rmsquery -v
sql> select name,current_upg_state,desired_upg_state from sc_domain
name      current_upg_state      desired_upg_state
-----
atlasD0 Upgraded                Upgraded
atlasD1 Upgraded                Upgraded
atlasD2 Pre_Upgrade            Upgraded
atlasD3 Pre_Upgrade            Upgraded
sql> quit
atlasms#
```

In the above example, atlasD0 and atlasD1 are upgraded while atlasD2 and atlasD3 are in the pre-upgrade state.

4.1.5 Preserved and Unpreserved Files

Local files and customizations on the management server or clustered management server will be preserved and/or migrated in line with the standard mechanisms of HP Tru64 UNIX Version 5.1B-3 (also known as HP Tru64 UNIX Version 5.1B Patch Kit 5). Please refer to the *Tru64 UNIX Release Notes for Version 5.1B-3* for further information.

During the upgrade process, certain member-specific files on domain members are automatically backed up and later restored, including the following files:

- `/etc/rc.config`
- `/etc/member_fstab`
- `/etc/securettys`
- `/etc/clu_alias.config`

The configuration of SCFS and PFS file systems, and data on these file systems, are also preserved during the upgrade. However, HP recommends that standard site backup policy be employed to ensure the safety of the configuration and data on these file systems in the event of a catastrophic failure.

Understanding the Upgrade Process

Site-specific customizations to `/etc/inetd.conf` may get changed during the upgrade to HP AlphaServer Version 2.6 (UK2). This is a side-effect of installing Tru64 UNIX Version 5.1B-3 during the upgrade. Where customizations have been made to `/etc/inetd.conf`, HP recommends that you back up the file and check the contents manually when the upgrade has completed.

If you want to preserve other site-specific modifications, you must back up the files containing the modifications before the upgrade and restore them manually when the upgrade is finished.

Files not mentioned here are not preserved by the upgrade process. For example, site-specific changes to the following areas on each domain member are lost during the upgrade:

- All `crontab` settings
- Special member swap settings
- Contents of member `/tmp`, `/tmp1`, `/local`, `/local1`
- Additional volumes added to member `/tmp`, `/tmp1`, `/local`, `/local1`

The files in this list and any other **site-specific files and customizations must be backed up manually, and then restored manually when the upgrade has finished**. If you are in any doubt as to whether files or modifications are automatically backed up, please take appropriate action to archive site-specific files and customizations.

4.1.6 Special Upgrade Mechanisms on Domains

In order to explain the special upgrade mechanism on the domains, it is necessary first to consider the full installation procedure for a domain.

When a domain is being fully installed, the patching of the system to Version 5.1B-3 happens when the domain is standalone (not part of the system) and running from the UNIX disk. Therefore, the `dupatch` process only applies the OSFBASE patches.

This reflects the fact that the TruCluster subsets within the HP AlphaServer SC product are re-mastered from later source code — where the source code actually includes the changes normally included in Version 5.1B-3. Therefore, the TruCluster binaries in the HP AlphaServer SC product are already patched and do not need to be patched again.

Once the HP AlphaServer SC kits are installed on the UNIX disk, the domain can be successfully created, and so on.

When the full installation is complete, the patch lists in the `setld` database on each domain indicate that only the Tru64 patches from Version 5.1B-3 are installed. However, through the re-mastering process, we know that the TruCluster subsets already include the Version 5.1B-3 functional changes.

However, in the case of a system upgrade, each domain is already built and a rolling upgrade is not supported on HP AlphaServer SC systems. In addition, the TruCluster subsets (in order to attempt to arrive at a similar configuration to the full installation scenario) cannot be removed because the TruCluster kernel/daemons are currently using files from the original TruCluster kits.

Therefore, the process is to first remove the original HP AlphaServer SC subsets from the domain, with the exception of the TruCluster kits that came with the original HP AlphaServer SC subsets. The original SRAOSFPATCH and SRATCRBASE subsets kits are then uninstalled because of the dependencies that they place in the lock files for the Tru64 and TruCluster base subsets.

The system will already be running Tru64 V5.1B and the `installupdate` command is not necessary, and therefore the non-lead members of each domain are not deleted. Instead, the upgrade will proceed to patch the cluster to Version 5.1B-3 using `sc_dupatch`.

The upgrade process simply unpacks the re-mastered TruCluster subset into the temporary area (and then patches the system to Version 5.1B-3 as described above). As the patch process completes, the contents of the temporary area are copied back onto the normal `/usr/opt/TruCluster` area, thereby reinstating the re-mastered TruCluster subset files and leaving the contents effectively unchanged by the general patching process.

Following `sc_dupatch`, the kernel rebuild and reboot stages at the end of the patch process are deliberately skipped. This provides the opportunity to install the new HP AlphaServer SC Version 2.6 (UK2) subsets — but excluding the TruCluster subset from within the HP AlphaServer SC Version 2.6 (UK2) product.

Once the HP AlphaServer SC kits are installed, the new SRAOSFPATCH320 and SRATCRBASE320 subsets are in place, and it is these subsets that reintroduce the special Tru64 and TruCluster hooks to support the HP AlphaServer SC product. At this point, the kernel can be built and the single node cluster rebooted.

When the domain upgrade process is complete, the patch lists in the `setld` database on each domain indicate that both Tru64 and TruCluster patches from Version 5.1B-3 are installed.

4.1.7 Approximate Timing Guideline

The amount of time an upgrade to HP AlphaServer SC Version 2.6 (UK2) takes is dependent on the number of nodes in the system, the amount of memory, and the node types. In addition, the time taken to upgrade the management server or clustered management server needs to be factored into the calculations.

Pre-Upgrade Audit

Table 4–2 displays the amount of time typically taken by the most time-consuming steps in an upgrade process.

Table 4–2 Upgrade Time Estimates

Activity	Time Estimate
Management Server or Management Cluster or First Domain	
Back up management server	Data dependent - allow 2 hours
Back up management cluster	Data dependent - allow 2 hours
Back up first domain	root/usr/var/boot only - approx 30 minutes
Patch management server to Version 5.1B-3	Approx 1.5 hours
Patch management cluster to Version 5.1B-3	Approx 2.5 hours
Patch first domain to Tru64 UNIX Version 5.1B-3	Approx 3 hours
Conclude upgrade tasks on management server	Approx 1.5 hours
Concluding upgrade tasks on management cluster	Approx 1.5 hours
Conclude upgrade tasks on first domain	Approx 1.5 hours
Each Subsequent Domain (Upgraded in Parallel)	
Back up domain	root/usr/var/boot only - approx 30 minutes
Check domain	Approx 10 minutes
Upgrade and patch domain to Tru64 UNIX Version 5.1B-3	Approx 5 hours
	Approx 5 hours
Add members back into domain	
Post-Upgrade Tasks	
Upgrade each PFS	Dependent on number of user-files and the number of SCFS components per PFS
Restoration in the Event of Failure	
Restore management server	Data dependent - allow 2 hours
Restore management cluster	Data dependent - allow 2 hours
Restore first domain	root/usr/var/boot only - approx 30 minutes
Restore other domains	root/usr/var/boot only - approx 30 minutes

4.2 Pre-Upgrade Audit

While your system is still in production mode, it is possible to perform an audit of the system in order to identify any problems that might subsequently affect the upgrade process.

The `upgrade_check` audit script will check that the basic requirements of the upgrade procedure are met by the system. The audit script will not fix any problems identified, but instead will suggest how the problem might be corrected. Refer to Appendix E.5 for a list of audits performed.

The audit script is non-invasive and can be run while the system is still in production mode. The audit script can be run repeatedly when the system is in production mode so that the problems identified can be resolved prior to the planned downtime window.

The same audit script is run again during the upgrade procedure once the system is removed from production. This is necessary in order to validate that certain invasive steps, such as marking file systems offline, have been performed correctly before proceeding with the upgrade.

If you run the audit script when a system is in production mode, you can expect to see warning messages regarding file systems still being online. Therefore, each reported audit failure should be considered on its own merits before continuing to plan the system downtime for the actual upgrade.

At this point, you should read the current version of the *HP AlphaServer SC Release Notes*, particularly any information about upgrade installations. You should also check with your account representative so that you are aware of any *HP AlphaServer SC Support Bulletins* relating to the upgrade procedure.

To run the `upgrade_check` script on the management server (or Node 0, if you do not have a management server), perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```
4. Run the following command:

```
atlasms# ./upgrade_check /tmp/upgrade_check.log
```

Preparing RIS for C2 Security

If `upgrade_check` finds conditions that would cause the upgrade to fail, it will report it to standard output (`stdout`) with a suggestion on how to correct the problem. Refer to Appendix E.5 for a list of tests generated from the `upgrade_check` command.

A detailed log of what actions `upgrade_check` performs to test the system is recorded in `/var/sra/adm/log/InstallSC/upgrade_check.log`, if no logfile is specified on the command line.

4.3 Preparing RIS for C2 Security

If your RIS server will have C2 security enabled, the RIS user file must be changed to ensure that the `ris` password does not expire and deny client access.

Perform the following steps on the RIS server as superuser to modify the RIS user file if you are going to use RIS with C2 security enabled:

1. Edit the file `/tcb/files/auth.db`. This requires you to use the `edauth` utility.

Use the `edauth` command line utility or the `dxaccounts` graphical user interface to modify the RIS account. See `edauth(8)` or `dxaccounts(8)` for more information.

Each field is delimited by a colon (`:`)

2. Set the current password field `u_pwd` to an asterisk (`*`).
3. Set the `u_succhg` value to any non-zero value. This value is a `time_t` type printed with `%ld`.
4. Set the `u_life` and `u_exp` fields to zero.

The following is an example of a modified `/tcb/files/auth/r/ris` user file:

```
ris:u_name=ris:u_id#11:\
    u_oldcrypt#0:\
    u_pwd=*\
    u_exp#0:u_life#0:\
    u_succhg#79598399:\
    u_suclog#79598399:\
    u_lock@:chkent:
```

After you make these changes, the RIS password should not expire and cause a denial of service to clients.

4.4 Upgrade with a Management Server

This section describes how to upgrade an HP AlphaServer SC system that has an unclustered management server. To upgrade with a management server, perform the following steps:

1. Back Up the Management Server (see Section 4.4.1 on page 4–11)
2. Upgrade the Management Server (see Section 4.4.2 on page 4–11)

4.4.1 Back Up the Management Server

Perform the following tasks on the management server:

1. Back up the management server root (/), /usr, and /var file domains and filesets. Guidance on one possible method for backing up these filesets on a management server is described in Back Up and Restore the Management Server Root (/), /usr, and /var File Systems (see Section 10.2.2 on page 10–4).
2. HP also recommends that you back up user data before beginning the upgrade installation. This includes data held locally on the management server and data held locally on the domains.

For further information on backup strategies in general, refer to Chapter 9 of the *Tru64 UNIX System Administration Guide*.

You are now ready to upgrade the management server, as described in Section 4.4.2.

4.4.2 Upgrade the Management Server

To upgrade the management server, perform the following tasks:

- Prepare for Upgrade (see Section 4.4.2.1 on page 4–11).
- Perform the Pre-Upgrade Check (see Section 4.4.2.2 on page 4–13).
- Back Up the SC Database (see Section 4.4.2.3 on page 4–13).
- Remove the HP AlphaServer SC Software (see Section 4.4.2.4 on page 4–14).
- Install the Operating System Patch Software (see Section 4.4.2.5 on page 4–14).
- Build the New Kernel and Reboot (see Section 4.4.2.6 on page 4–17).
- Configure the RIS Server (see Section 4.4.2.7 on page 4–17).
- Install the HP AlphaServer SC Software (see Section 4.4.2.8 on page 4–19).
- Restore the SC Database (see Section 4.4.2.9 on page 4–19).
- Migrate the SC Database (see Section 4.4.2.10 on page 4–20).
- Register RIS Clients for the New RIS Environment (see Section 4.4.2.11 on page 4–21).
- Build the New Kernel and Reboot (see Section 4.4.2.12 on page 4–21).

4.4.2.1 Prepare for Upgrade

With the system removed from production mode, the following steps should be performed:

1. Disable the SCFS and PFS file systems as follows:
`atlasms# pfsmgr offline all`

Upgrade with a Management Server

Once this has completed successfully, and the PFS file systems are offline on all nodes (check that systems are offline on all nodes by using the command `pfsmgr show`), configure out `pfs.mod` as follows:

```
atlasms# scrn -n all /sbin/sysconfig -u pfs
```

There is no need to build or deploy kernels or reboot any nodes at this time.

Then, mark the SCFS filesystems offline as follows:

```
atlasms# scfsmgr offline all
```

The SCFS and PFS filesystems will now be automatically unmounted. You should use `scfsmgr show` repeatedly until all filesystems are shown as unmounted.

2. If present, de-install the PFS subset as follows:

```
atlasms# setld -i |grep PFSMOD  
PFSMOD300 installed Quadrics PFS File System  
atlasms# setld -d PFSMOD300
```

3. If present de-install the OTABASE subset (part of Compaq Fortran kit) from the management server as follows:

```
atlasms# setld -i |grep OTABASE  
OTABASE219 installed Compaq Compiled Code Support Library #219  
atlasms# setld -d OTABASE219
```

4. If present, de-configure and de-install the SCIP software from the management server and from the domains. Please refer to the *HP AlphaServer SC Installing and Configuring SCIP* user manual for details on how to de-install SCIP.
5. Tru64 UNIX Version 5.1B-3 includes Ladebug patches that can only be applied to the default version of Ladebug that shipped with Tru64 UNIX Version 5.1B.

If an updated Ladebug subset has been installed, remove it now, and replace it with the default version of Ladebug (which is available on the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*), as follows:

```
atlasms# setld -i |grep LDB  
LDBBASE467 installed Ladebug Debugger Version 467  
LDBDOC467 installed Ladebug Debugger Version 467 Documentation  
OSFLDBBASE540 not installed Ladebug Debugger Version 467 (Software Development)  
OSFLDBDOC540 not installed Ladebug Debugger Version 467 Documentation (Software Development)  
atlasms# setld -d LDBBASE467 LDBDOC467
```

Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.

```
atlasms# mkdir /cdrom; mount -r /dev/disk/cdrom0c /cdrom  
atlasms# setld -l /cdrom/ALPHA/BASE OSFLDBBASE540 OSFLDBDOC540
```

You are now ready to perform the Pre-Upgrade Check, as described in Section 4.4.2.2.

4.4.2.2 Perform the Pre-Upgrade Check

Check that the system is now ready to be upgraded using the following commands:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```

4. Run the following command:

```
atlasms# ./upgrade_check /tmp/upgrade_check.log
```

If `upgrade_check` finds conditions that would cause the upgrade to fail, it will report it to standard output (`stdout`) with a suggestion on how to correct the problem. Refer to Appendix E.5 for a list of tests generated from the `upgrade_check` command.

Do not proceed with the upgrade procedure until all tests within this audit are successful. Make whatever corrections are necessary to resolve any issues identified.

A detailed log of what actions `upgrade_check` performs to test the system is recorded in `/var/sra/adm/log/InstallSC/upgrade_check.log`, if no logfile is specified on the command line.

Once all tests are successful, you are ready to Back Up the SC Database, as described in Section 4.4.2.3.

4.4.2.3 Back Up the SC Database

In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. In order to preserve the existing data, it is necessary to backup the current SC database to a file and restore from this file later in the upgrade process.

To back up the SC database on the management server, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```

Upgrade with a Management Server

- b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```
4. Run the following command:

```
atlasms# ./sc_pre_upgrade
```

You are now ready to remove the HP AlphaServer SC software, as described in Section 4.4.2.4.

4.4.2.4 Remove the HP AlphaServer SC Software

Before upgrading the operating system on the management server, you must first de-install the HP AlphaServer SC software as follows:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM 1 as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```
4. De-install the previous version of the HP AlphaServer SC System Software:

```
atlasms# ./InstallSC -remove -ms
```

You can now proceed to install the Tru64 UNIX patch software, as described in Section 4.4.2.5.

4.4.2.5 Install the Operating System Patch Software

Install the operating system patch software as follows:

1. Unpack the Tru64 UNIX Version 5.1B-3 patch kit tar file in the `/patches` directory. The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:

```
<http://www.itrc.hp.com/>
```


or from your local HP support representative.

Note:

For HP AlphaServer SC Version 2.6 (UK2), you should load Tru64 UNIX Version 5.1B-3 only. Do NOT run the standard `dupatch` script to patch the management server at this time. In HP AlphaServer SC Version 2.6 (UK2), use the alternative patching script, called `sc_dupatch`, that is available on the *HP AlphaServer SC System Software CD-ROM*.

2. Shut down the system to single-user mode as follows:

```
atlasms# shutdown now
```

3. Mount the local file systems:

```
atlasms# /sbin/bcheckrc
```

4. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.

5. Mount the CD-ROM as the root user, as follows:

- a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

6. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```

7. Perform the patch precheck phase to see if there are any potential problems with the patch kit on your management server. The check can take 20-30 minutes, depending on the patches already installed on the management server. Run the supplied HP AlphaServer SC patch script to check the patches for the management server as follows:

```
atlasms# ./sc_dupatch -ms_upgrade -kit path \  
-product Tru64_UNIX_V5.1B -precheck_only
```

where:

-ms_upgrade indicates that you are patching the management server

-kit path specifies the directory where you unpacked the patch kit

-product Tru64_UNIX_V5.1B specifies the product to which the patch applies

-precheck_only indicates that you only want to do the precheck phase of the patch process at this time.

Upgrade with a Management Server

Note:

When upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), the following patches may fail the prerequisite/file applicability check:

- Patch 26035.00 - SP05 OSFHWBIN540
- Patch 26010.00 - SP05 OSFBIN540 (SSRT4891 SSRT4743 SSRT4696 ...)
- Patch 26001.00 - SP05 TCRBASE540 (SSRT2265)

This is normal behavior and will not prevent the upgrade from continuing. However, if any unexpected problems are encountered installing patches, do not continue the installation process. Stop the patch install, resolve the problem, and try the patch install process again.

During patch installations, warning messages may be displayed about files that have unknown origins, if you have previously installed manual patches. To resolve this problem, add a new `-deps_only` flag to the install command as shown in the next step. This flag will notify the install command to ignore any problems with files that have unknown origins.

8. Run the HP AlphaServer SC patch script in full install mode as follows, (approx 90 minutes duration):

```
atlasms# ./sc_dupatch -ms_upgrade -kit path \  
-product Tru64_UNIX_V5.1B -deps_only -proceed
```

using the same *path* as that used in step 6 above.

where:

`-ms_upgrade` indicates that you are patching the management server

`-kit path` specifies the directory where you unpacked the patch kit

`-product Tru64_UNIX_V5.1B` specifies the product to which the patch applies

`-deps_only` specifies that patch dependency checks should be performed and that inventory checks should be skipped. This check is useful in a system where files have been replaced by hand (for example, where new binary is copied into a directory in place of the binary recorded in the `setld` inventory files).

`-proceed` specifies that the installation should proceed with the patches that pass the dependency checks.

9. When the HP AlphaServer SC patch script is finished, the management server kernel will NOT be rebuilt and the system will NOT be rebooted.

You are now ready to build the new kernel and reboot, as described in Section 4.4.2.6.

4.4.2.6 Build the New Kernel and Reboot

To build the kernel on a management server, perform the following steps:

1. Save a copy of the existing `/vmunix` file. If possible, save the file in the root (`/`) directory, as follows:

```
atlasms# cp /vmunix /vmunix.save
```

If there are disk space constraints, you can save the kernel file in a file system other than root. For example:

```
atlasms# cp /vmunix /usr/vmunix.save
```
2. Run the `/usr/sbin/doconfig` program specifying the name of the target configuration file with the `-c` option. For example, on a system named `atlasms`, enter the following command:

```
atlasms# /usr/sbin/doconfig -c ATLASMS
```

```
*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***
```

```
Saving /sys/conf/ATLASMS as /sys/conf/ATLASMS.bck
```
3. You are prompted to indicate whether or not you want to edit the configuration file:

```
Do you want to edit the configuration file? (y/n) [n]:
```

Accept the default to indicate that you do not want to edit the configuration file; the `/usr/sbin/doconfig` program builds a new kernel.
4. When the kernel configuration and build are completed without errors, copy the new `vmunix` file to `/vmunix`. On a system named `atlasms`, enter the following command:

```
atlasms# cp /sys/ATLASMS/vmunix /vmunix
```
5. Reboot the system as follows:

```
atlasms# /usr/sbin/shutdown -r now
```

You are now ready to configure the RIS server, as described in Section 4.4.2.7.

4.4.2.7 Configure the RIS Server

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, the management server is expected to have a properly-configured RIS server.

If you have previously configured RIS server (for example, during a previous upgrade or during the installation), please ensure that it is working properly.

Otherwise, configure the RIS server as described in this section.

Upgrade with a Management Server

If you have deliberately disabled a previous RIS configuration, for example, for security reasons, it will need to be re-enabled now. To ensure that the RIS `joind` daemon is started correctly, it is suggested that you enable RIS again, and then reboot your management server so that `dhcp/joind` starts automatically. If you elect to start `dhcp/joind` manually, there are occasions where it may exit unexpectedly during the upgrade process and cause the upgrade to fail.

Load the Tru64 UNIX V5.1B operating system into the RIS environment as follows:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:
`atlasms# mkdir /cdrom`
 - b. Mount the CD-ROM as follows:
`atlasms# mount -r /dev/disk/cdrom0c /cdrom`
3. Run the `ris` command as the root user, as follows:
`atlasms# ris`
4. Choose the `Install software products` option by entering `i` at the prompt:
Enter your choice: `i`
5. The RIS Installation menu displays the installation options. Choose option 1, the `Install software into a new area` option.
6. Enter the full pathname for the distribution media, as follows:
Enter the device special file name or the path of the directory where the software is located
(for example, `/mnt/ALPHA/BASE`): `/cdrom/ALPHA/BASE`
7. Choose the standard boot method.
8. Choose to extract the software from the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*.
9. Choose to install all mandatory and all optional subsets. You will need to go through a number of pages of options before selecting this option.
10. Enter `y` to confirm that the subset list is correct. The subset extraction process begins.

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to install the HP AlphaServer SC software, as described in Section 4.4.2.8.

4.4.2.8 Install the HP AlphaServer SC Software

Once the operating system on the management server has been upgraded, install the HP AlphaServer SC software as follows:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```
4. Install the HP AlphaServer SC System Software as follows:

```
atlasms# ./InstallSC -install -ms
```
5. If PFS was previously installed, then install the new PFS subset on the management server as follows:

```
atlasms# cd PFS  
atlasms# setld -l . PFSMOD320
```

You are now ready to restore the SC database, as described in Section 4.4.2.9.

4.4.2.9 Restore the SC Database

In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. Earlier in Section 4.4.2.3 the original SC database was backed up to a file. It is now necessary to restore the SC database from this file. To restore the SC database on the management server, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits
```

Upgrade with a Management Server

4. Run the following command:

```
atlasms# ./sc_post_upgrade
```

Note:

The `sc_post_upgrade` command will simply restore the original database. You may see the following warning:

```
rmstbladm: Warning: table sc_procurve is not present in backup file
```

The migration of the SC database as described in Section 4.4.2.10 will introduce this new table and other new database settings required by the HP AlphaServer SC software release.

The `sc_post_upgrade` command will display the following error messages, which can be ignored:

```
Setting RMS attributes...
rcontrol: Error: attribute where name='cleanup-timeout'
already exists in database
rcontrol: Error: attribute where name='ignore-cleanup-failure'
already exists in database
```

You are now ready to migrate the SC Database, as described in Section 4.4.2.10.

4.4.2.10 Migrate the SC Database

New versions of the HP AlphaServer SC software may include new tables, or new fields within existing tables in the SC Database. It is necessary to upgrade your database to include these new tables and fields and to populate them with appropriate data.

To migrate the database, run the following command:

```
atlasms# sc_upgrade_db
```

When prompted, you should supply a base address for the Preferred Server Cluster Aliases. Refer to chapter 19 of the *HP AlphaServer SC System Administration Guide* for further information.

You are now ready to register RIS clients for the New RIS Environment, as described in Section 4.4.2.11.

Upgrade with a Clustered Management Server

4.4.2.11 Register RIS Clients for the New RIS Environment

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, all domain members should be registered as clients of Tru64 UNIX Version 5.1B.

If you have not previously registered them with the RIS server (during a previous upgrade or installation), please follow the instructions in this section to do this.

Register all domain members as clients of Tru64 UNIX Version 5.1B, as follows:

```
atlasms# sra edit
sra> sys
sys> update ris all
```

Note:

In response to the `sra` command above and in response to subsequent `sra` commands, you may receive ongoing warnings to upgrade the SC Database. These warnings should be ignored, and once all domains have been upgraded, the procedure as described in Section 4.9.2 will be used to upgrade the database revision.

You are now ready to build the new kernel and reboot, as described in Section 4.4.2.12.

4.4.2.12 Build the New Kernel and Reboot

You now need to once again build the kernel and reboot. Follow the steps as described in Build the New Kernel and Reboot (see Section 4.4.2.6 on page 4–17).

You are now ready to upgrade the domains, as described in Section 4.7.

4.5 Upgrade with a Clustered Management Server

This section describes how to upgrade a HP AlphaServer SC system that has a clustered management server. To upgrade with a clustered management server, take the following steps:

1. Back Up the Clustered Management Server (see Section 4.5.1 on page 4–22)
2. Upgrade the Clustered Management Server (see Section 4.5.2 on page 4–22)

Upgrade with a Clustered Management Server

4.5.1 Back Up the Clustered Management Server

Perform the following tasks on the clustered management server:

1. Back up the clustered management server root (/), /usr, and /var file domains and filesets. Guidance on one possible method for backing up these filesets on a management server is described in Back Up and Restore the Management Server Root (/), /usr, and /var File Systems (see Section 10.2.2 on page 10–4). You should also backup up the member-specific boot partitions on the clustered management server.
2. HP also recommends that you back up user data before beginning the upgrade installation. This includes data held locally on the clustered management server and data held locally on the domains.

For further information on backup strategies in general, refer to Chapter 9 of the *Tru64 UNIX System Administration Guide*.

You are now ready to upgrade the clustered management server, as described in Section 4.5.2.

4.5.2 Upgrade the Clustered Management Server

To upgrade the clustered management server, perform the following tasks:

- Prepare for Upgrade (see Section 4.5.2.1 on page 4–22)
- Perform the Pre-Upgrade Check (see Section 4.5.2.2 on page 4–24)
- Back Up the SC Database (see Section 4.5.2.3 on page 4–24)
- Remove the HP AlphaServer SC Software (see Section 4.5.2.4 on page 4–25)
- Install the Operating System Patch Software (see Section 4.5.2.5 on page 4–25)
- Configure the RIS Server (see Section 4.5.2.6 on page 4–28)
- Install the HP AlphaServer SC Software (see Section 4.5.2.7 on page 4–29)
- Restore the SC Database (see Section 4.5.2.8 on page 4–29)
- Migrate the SC Database (see Section 4.5.2.9 on page 4–30)
- Register RIS Clients for the New RIS Environment (see Section 4.5.2.10 on page 4–31)
- Build the New Kernel and Reboot (see Section 4.5.2.11 on page 4–31)

4.5.2.1 Prepare for Upgrade

With the system removed from production mode, the following steps should be performed:

1. Disable the SCFS and PFS file systems as follows:

```
atlasms0# pfsmgr offline all
```

Upgrade with a Clustered Management Server

Once this has completed successfully, and the PFS file systems are offline on all nodes (check that systems are offline on all nodes by using the command `pfsmgr show`), configure out `pfs.mod` as follows:

```
atlasms0# scrunch -n all /sbin/sysconfig -u pfs
```

There is no need to build or deploy kernels or reboot any nodes at this time.

Then, mark the SCFS filesystems offline as follows:

```
atlasms0# scfsmgr offline all
```

The SCFS and PFS filesystems will now be automatically unmounted. You should use `scfsmgr` repeatedly until all filesystems are shown as unmounted.

2. If present, de-install the PFS subset as follows:

```
atlasms0# setld -i |grep PFSMOD  
PFSMOD300 installed Quadrics PFS File System  
atlasms0# setld -d PFSMOD300
```

3. If present de-install the OTABASE subset (part of Compaq Fortran kit) from the management server as follows:

```
atlasms0# setld -i |grep OTABASE  
OTABASE219 installed Compaq Compiled Code Support Library #219  
atlasms0# setld -d OTABASE219
```

4. If present, de-configure and de-install the SCIP software from the management server and from the domains. Please refer to the *HP AlphaServer SC Installing and Configuring SCIP* user manual for details on how to de-install SCIP.

5. Tru64 UNIX Version 5.1B-3 includes Ladebug patches that can only be applied to the default version of Ladebug that shipped with Tru64 UNIX Version 5.1B.

If an updated Ladebug subset has been installed, remove it now, and replace it with the default version of Ladebug (which is available on the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*), as follows:

```
atlasms0# setld -i |grep LDB  
LDBBASE467 installed Ladebug Debugger Version 467  
LDBDOC467 installed Ladebug Debugger Version 467 Documentation  
OSFLDBBASE540 not installed Ladebug Debugger Version 467 (Software Development)  
OSFLDBDOC540 not installed Ladebug Debugger Version 467 Documentation (Software Development)  
atlasms0# setld -d LDBBASE467 LDBDOC467
```

Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.

```
atlasms0# mkdir /cdrom; mount -r /dev/disk/cdrom0c /cdrom  
atlasms0# setld -l /cdrom/ALPHA/BASE OSFLDBBASE540 OSFLDBDOC540
```

You are now ready to perform the Pre-Upgrade Check, as described in Section 4.5.2.2.

Upgrade with a Clustered Management Server

4.5.2.2 Perform the Pre-Upgrade Check

Check that the system is now ready to be upgraded using the following commands:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms0# cd /cdrom/kits
```
4. Run the following command:

```
atlasms0# ./upgrade_check /tmp/upgrade_check.log
```

If `upgrade_check` finds conditions that would cause the upgrade to fail, it will report it to standard output (`stdout`) with a suggestion on how to correct the problem. Refer to Appendix E.5 for a list of tests generated from the `upgrade_check` command.

Do not proceed with the upgrade procedure until all tests within this audit are successful. Make whatever corrections are necessary to resolve any issues identified.

A detailed log of what actions `upgrade_check` performs to test the system is recorded in `/var/sra/adm/log/InstallSC/upgrade_check.log`, if no logfile is specified on the command line.

Once all tests are successful, you are ready to Back Up the SC Database, as described in Section 4.5.2.3.

4.5.2.3 Back Up the SC Database

In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. In order to preserve the existing data, it is necessary to backup the current SC database to a file and restore from this file later in the upgrade process.

To back up the SC database on the management server, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```

Upgrade with a Clustered Management Server

3. Change to the directory in which the kits are stored, as follows:

```
atlasms0# cd /cdrom/kits
```

4. Run the following command:

```
atlasms0# ./sc_pre_upgrade
```

You are now ready to remove the HP AlphaServer SC software, as described in Section 4.5.2.4.

4.5.2.4 Remove the HP AlphaServer SC Software

Before upgrading the operating system on the management server, you must first de-install the HP AlphaServer SC software as follows:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms0# mkdir /cdrom
```

- b. Mount the CD-ROM 1 as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```

3. Change to the directory in which the kits are stored, as follows:

```
atlasms0# cd /cdrom/kits
```

4. Deinstall the previous version of the HP AlphaServer SC System Software:

```
atlasms0# ./InstallSC -remove -ms
```

After the previous version of the HP AlphaServer SC System software is removed, the CAA applications for MSQ, RMS, and SRA will remain registered. This will cause CAA warnings to be displayed on the console. These warnings can be ignored and the CAA services will be restored to normal once the new HP AlphaServer SC system software is installed.

You are now ready to install the Tru64 UNIX patch software, as described in Section 4.5.2.5.

4.5.2.5 Install the Operating System Patch Software

Note:

The following instructions apply for installing Tru64 UNIX patches in a no-roll mode. Using the no-roll patch method requires a reboot of the entire management server cluster. If you require the management server cluster to always be up and

Upgrade with a Clustered Management Server

available, then you can use the rolling-patch method to install Tru64 UNIX patches. Refer to the patch documentation provided with the patch kit for more information on performing rolling patches.

Install the operating system patch software as follows:

1. Unpack the Tru64 UNIX Version 5.1B-3 patch kit tar file in the `/patches` directory. The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:
`<http://www.itrc.hp.com/>`
or from your local HP support representative.
2. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
3. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:
`atlasms0# mkdir /cdrom`
 - b. Mount the CD-ROM, as follows:
`atlasms0# mount -r /dev/disk/cdrom0c /cdrom`
4. Change to the directory in which the kits are stored, as follows:
`atlasms0# cd /cdrom/kits`
5. Perform the patch precheck phase to see if there are any potential problems with the patch kit on your management server. The check can take 20-30 minutes, depending on the patches already installed on the management server. Run the supplied HP AlphaServer SC patch script to check the patches for the management server, as follows:
`atlasms0# ./sc_dupatch -ms_upgrade -kit path -product all -precheck_only`
where:
 - ms_upgrade indicates that you are patching the management server
 - kit path specifies the directory where you unpacked the patch kit
 - product all specifies that both Tru64_UNIX_V5.1B and TruCluster_V5.1B product patches will be applied
 - precheck_only indicates that you only want to do the precheck phase of the patch process at this time.

Note:

When upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), the following patches may fail the prerequisite/file applicability check:

Upgrade with a Clustered Management Server

- Patch 26035.00 - SP05 OSFHWBIN540
- Patch 26010.00 - SP05 OSFBIN540 (SSRT4891 SSRT4743 SSRT4696 ...)
- Patch 26001.00 - SP05 TCRBASE540 (SSRT2265)

This is normal behavior and will not prevent the upgrade from continuing. However, if any unexpected problems are encountered when installing patches, do not continue the installation process. Stop the patch installation, resolve the problem, and try the patch install process again.

During patch installations, warning messages may be displayed about files that have unknown origins, if you have previously installed manual patches. To resolve this problem, add a new `-deps_only` flag to the install command as shown in the next step. This flag will notify the install command to ignore any problems with files that have unknown origins.

6. Run the HP AlphaServer SC patch script in full install mode as follows (approximately 90 minutes duration):

```
atlasms0# ./sc_dupatch -ms_upgrade -kit path \  
-product all -deps_only -proceed
```

using the same path as that used in step 6 above,

where:

`-ms_upgrade` indicates that you are patching the management server

`-kit path` specifies the directory where you unpacked the patch kit

`-product all` specifies that both Tru64_UNIX_V5.1B and TruCluster_V5.1B product patches will be applied

`-deps_only` specifies that patch dependency checks should be performed and that inventory checks should be skipped. This check is useful in a system where files have been replaced by hand (for example, where new binary is copied into a directory in place of the binary recorded in the `setld` inventory files).

`-proceed` specifies that the installation should proceed with the patches that pass the dependency checks.

7. When the HP AlphaServer SC patch script is finished, the clustered management server kernels will NOT be rebuilt and the system will NOT be rebooted.

You are now ready to configure the RIS server, as described in Section 4.5.2.6.

Upgrade with a Clustered Management Server

4.5.2.6 Configure the RIS Server

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, the management server is expected to have a properly-configured RIS server.

If you have previously configured the RIS server (for example, during a previous upgrade or during the installation), please ensure that it is working properly.

Otherwise, configure the RIS server as described in this section.

The upgrade of the Tru64 UNIX operating system on each domain will be performed using a Remote Installation Services (RIS) server. Load the Tru64 UNIX V5.1B operating system into the RIS environment as follows:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```
3. Run the `ris` command as the root user, as follows:

```
atlasms0# ris
```
4. Choose the `Install software products` option by entering `i` at the prompt:

```
Enter your choice: i
```
5. The RIS Installation menu displays the installation options. Choose option 1, the `Install software into a new area` option.
6. Enter the full pathname for the distribution media, as follows:

```
Enter the device special file name or the path of the directory  
where the software is located  
(for example, /mnt/ALPHA/BASE): /cdrom/ALPHA/BASE
```
7. Choose the standard boot method.
8. Choose to extract the software from the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*.
9. Choose to install all mandatory and all optional subsets. You will need to go through a number of pages of options before selecting this option.

10. Enter `y` to confirm that the subset list is correct. The subset extraction process begins.

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to install the HP AlphaServer SC software, as described in Section 4.5.2.7.

4.5.2.7 Install the HP AlphaServer SC Software

Once the operating system on the management server has been upgraded, install the HP AlphaServer SC software as follows:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlasms0# cd /cdrom/kits
```
4. Install the HP AlphaServer SC System Software:

```
atlasms0# ./InstallSC -install -ms
```
5. If PFS was previously installed, then install the new PFS subset on the management server as follows:

```
atlasms0# cd PFS  
atlasms0# setld -l . PFSMOD320
```

You are now ready to restore the SC database, as described in Section 4.5.2.8.

4.5.2.8 Restore the SC Database

In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. Earlier in Section 4.5.2.3 the original SC database was backed up to a file. It is now necessary to restore the SC database from this file. To restore the SC database on the management server, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

Upgrade with a Clustered Management Server

```
atlasms0# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlasms0# mount -r /dev/disk/cdrom0c /cdrom
```

3. Change to the directory in which the kits are stored, as follows:

```
atlasms0# cd /cdrom/kits
```

4. Run the following command:

```
atlasms0# caa_relocate SC05msql -c atlasms0
atlasms0# ./sc_post_upgrade
```

Note:

The `sc_post_upgrade` command will simply restore the original database. You may see the following warning message:

```
rmstbladm: Warning: table sc_procurve is not present in backup
file
```

The migration of the SC database as described in Section 4.5.2.9 will introduce this new table and other new database settings required by the HP AlphaServer SC software release.

The `sc_post_upgrade` command will display the following error messages, which can be ignored:

```
Setting RMS attributes...
rcontrol: Error: attribute where name='cleanup-timeout'
already exists in database
rcontrol: Error: attribute where name='ignore-cleanup-failure'
already exists in database
```

You are now ready to migrate the SC Database, as described in Section 4.5.2.9.

4.5.2.9 Migrate the SC Database

New versions of the HP AlphaServer SC software may include new tables, or new fields within existing tables in the SC Database. It is necessary to upgrade your database to include these new tables and fields and to populate them with appropriate data.

To migrate the database, run the following command:

```
atlasms0# sc_upgrade_db
```

When prompted, you should supply a base address for the Preferred Server Cluster Aliases. Refer to Chapter 19 of the *HP AlphaServer SC System Administration Guide* for further information.

Upgrade with a Clustered Management Server

You are now ready to register RIS clients for the New RIS Environment, as described in Section 4.5.2.10.

4.5.2.10 Register RIS Clients for the New RIS Environment

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, all domain members should be registered as clients of Tru64 UNIX Version 5.1B.

If you have not previously registered them with the RIS server (during a previous upgrade or installation), please follow the instructions in this section to do this.

Register all domain members as clients of Tru64 UNIX V5.1B, as follows:

```
atlasms0# sra edit
sra> sys
sys> update ris all
```

Note:

In response to the `sra` command above and in response to subsequent `sra` commands, you may receive ongoing warnings to upgrade the SC Database. These warnings should be ignored, and once all domains have been upgraded, the procedure as described in Section 4.9.2 will be used to upgrade the database revision.

You are now ready to build the new kernel and reboot, as described in Section 4.5.2.11.

4.5.2.11 Build the New Kernel and Reboot

Build the kernel on a clustered management server, as follows:

```
atlasms0# BuildKernels
atlasms0# DeployKernels
```

Once the kernels are built and deployed, the system should be shut down as follows:

```
atlasms0# shutdown -ch now
```

Once the nodes are shut down, boot them again from the consoles.

You are now ready to upgrade the domains, as described in Section 4.7.

Upgrade without a Management Server

4.6 Upgrade without a Management Server

This section describes how to upgrade an HP AlphaServer SC system that does not have a management server. The upgrade of the first domain requires the execution of a number of manual commands, that is, the user is responsible for driving the process. Once the first domain is upgraded, the automatic upgrade facility (`sra upgrade`) can be used to upgrade the other domains.

The steps to upgrade a system without a management server are as follows:

1. Back Up the First Domain (see Section 4.6.1 on page 4–32)
2. Upgrade the First Domain (see Section 4.6.2 on page 4–32)

4.6.1 Back Up the First Domain

Perform the following tasks on the first node of the first domain:

1. Back up the member-specific boot partitions and the cluster root (`/`), `/usr`, and `/var` file domains and filesets of the first domain as follows:

```
atlas0# /usr/sra/bin/sra_cluster_backup -force -disk dsk5 backup
```

where `dsk5` corresponds to the disk selected as the backup disk for domain 0.
2. HP also recommends that you back up user data before beginning the upgrade installation. This includes data held locally on the management server and data held locally on the domains.

For further information on backup strategies in general, refer to Chapter 9 of the *Tru64 UNIX System Administration Guide*.

You are now ready to upgrade the first domain, as described in Section 4.6.2.

4.6.2 Upgrade the First Domain

To upgrade the first domain, perform the following tasks:

- Prepare for Upgrade (see Section 4.6.2.1 on page 4–33).
- Perform the Pre-Upgrade Check (see Section 4.6.2.2 on page 4–34).
- Back Up the SC Database (see Section 4.6.2.3 on page 4–34).
- Install the Upgrade Subset on the First Domain (see Section 4.6.2.4 on page 4–35).
- Disable Cookies (see Section 4.6.2.5 on page 4–35).
- Configure the RIS Server (see Section 4.6.2.6 on page 4–36).
- Upgrade the First Domain (see Section 4.6.2.8 on page 4–38).

4.6.2.1 Prepare for Upgrade

With the system removed from production mode, the following steps should be performed:

1. Disable the SCFS and PFS file systems as follows:

```
atlas0# pfsmgr offline all
```

Once this has completed successfully, and the PFS file systems are offline on all nodes (check that systems are offline on all nodes by using the command `pfsmgr show`), configure out `pfs.mod` as follows:

```
atlas0# sscr -n all /sbin/sysconfig -u pfs
```

There is no need to build or deploy kernels or reboot any nodes at this time.

Then, mark the SCFS filesystems offline as follows:

```
atlas0# scfsmgr offline all
```

The SCFS and PFS filesystems will now be automatically unmounted. You should use `scfsmgr` repeatedly until all filesystems are shown as unmounted.

2. If present de-install the OTABASE subset (part of the HP Fortran Run-Time Library) as follows:

```
atlas0# setld -i |grep OTABASE
OTABASE219 installed Compaq Compiled Code Support Library #219
atlas0# setld -d OTABASE219
```

3. If present, de-configure and de-install the SCIP software. Please refer to the *HP AlphaServer SC Installing and Configuring SCIP* user manual for details on how to de-install SCIP.

4. Tru64 UNIX Version 5.1B-3 includes Ladebug patches that can only be applied to the default version of Ladebug that shipped with Tru64 UNIX Version 5.1B.

If an updated Ladebug subset has been installed, remove it now, and replace it with the default version of Ladebug (which is available on the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*), as follows:

```
atlas0# setld -i |grep LDB
LDBBASE467 installed Ladebug Debugger Version 467
LDBDOC467 installed Ladebug Debugger Version 467 Documentation
OSFLDBBASE540 not installed Ladebug Debugger Version 467 (Software Development)
OSFLDBDOC540 not installed Ladebug Debugger Version 467 Documentation (Software Development)
atlas0# setld -d LDBBASE467 LDBDOC467
```

Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.

```
atlas0# mkdir /cdrom; mount -r /dev/disk/cdrom0c /cdrom
atlas0# setld -l /cdrom/ALPHA/BASE OSFLDBBASE540 OSFLDBDOC540
```

You are now ready to perform the Pre-Upgrade Check, as described in Section 4.6.2.2.

Upgrade without a Management Server

4.6.2.2 Perform the Pre-Upgrade Check

Check that the system is now ready to be upgraded using the following commands:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlas0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlas0# mount -r /dev/disk/cdrom0c /cdrom
```
3. Change to the directory in which the kits are stored, as follows:

```
atlas0# cd /cdrom/kits
```
4. Run the following command:

```
atlas0# ./upgrade_check /tmp/upgrade_check.log
```

If `upgrade_check` finds conditions that would cause the upgrade to fail, it will report it to standard output (`stdout`) with a suggestion on how to correct the problem. Refer to Appendix E.5 for a list of tests generated from the `upgrade_check` command.

Do not proceed with the upgrade procedure until all tests within this audit are successful. Make whatever corrections are necessary to resolve any issues identified.

A detailed log of what actions `upgrade_check` performs to test the system is recorded in `/var/sra/adm/log/InstallSC/upgrade_check.log`, if no logfile is specified on the command line.

Once all tests are successful, you are ready to back up the SC Database, as described in Section 4.6.2.3.

4.6.2.3 Back Up the SC Database

In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. In order to preserve the existing data, it is necessary to backup the current SC database to a file and restore from this file later in the upgrade process.

To back up the SC database, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlas0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlas0# mount -r /dev/disk/cdrom0c /cdrom
```

3. Change to the directory in which the kits are stored, as follows:

```
atlas0# cd /cdrom/kits
```

4. Run the following command:

```
atlas0# ./sc_pre_upgrade
```

You are now ready to install the upgrade subset on the first domain, as described in Section 4.6.2.4.

4.6.2.4 Install the Upgrade Subset on the First Domain

To install the HP AlphaServer SC Upgrade subset on the first domain, perform the following steps:

1. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlas0# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlas0# mount -r /dev/disk/cdrom0c /cdrom
```

3. Change to the directory in which the kits are stored, as follows:

```
atlas0# cd /cdrom/kits
```

4. Delete the existing HP AlphaServer SC Upgrade subset as follows:

```
atlas0# setld -i |grep SRAUPG
SRAUPG300 installed AlphaServer SC Upgrade Utilities
atlas0# setld -d SRAUPG300
```

5. Install the new HP AlphaServer SC Upgrade subset, as follows:

```
atlas0# cd SRA
atlas0# setld -l . SRAUPG320
```

You are now ready to disable cookies, as described in Section 4.6.2.5.

4.6.2.5 Disable Cookies

Before upgrading the first domain, you should disable the `msql` cookie mechanism by running the following command on `rmshost`:

```
atlas0# sra cookie -enable no
```

To check if cookies are disabled, run the following command:

```
atlas0# sra cookie
```

You are now ready to configure the RIS server, as described in Section 4.6.2.6.

Upgrade without a Management Server

4.6.2.6 Configure the RIS Server

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, the first domain is expected to have a properly-configured RIS server.

If you have previously configured RIS server (for example, during a previous upgrade or during the installation), please ensure that it is working properly.

Otherwise, configure the RIS server as described in this section.

The upgrade of the Tru64 UNIX operating system on each domain will be performed using a Remote Installation Services (RIS) server. Load the Tru64 UNIX V5.1B operating system into the RIS environment as follows:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.
2. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```
3. Run the `ris` command as the root user, as follows:

```
atlasms# ris
```
4. Choose the `Install software products` option by entering `i` at the prompt:

```
Enter your choice: i
```
5. The RIS Installation menu displays the installation options. Choose option 1, the `Install software into a new area` option.
6. Enter the full pathname for the distribution media, as follows:

```
Enter the device special file name or the path of the directory  
where the software is located  
(for example, /mnt/ALPHA/BASE): /cdrom/ALPHA/BASE
```
7. Choose the standard boot method.
8. Choose to extract the software from the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*.
9. Choose to install all mandatory and all optional subsets. You will need to go through a number of pages of options before selecting this option.

Upgrade without a Management Server

10. Enter `y` to confirm that the subset list is correct. The subset extraction process begins.

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to register RIS clients for the new RIS environment, as described in Section 4.6.2.7.

4.6.2.7 Register RIS Clients for the New RIS Environment

Note:

When you are upgrading from HP AlphaServer SC Version 2.6 (UK1) to HP AlphaServer SC Version 2.6 (UK2), you are not required to perform an operating system upgrade from a RIS server.

However, for the upgrade to work properly, all domain members should be registered as clients of Tru64 UNIX Version 5.1B.

If you have not previously registered them with the RIS server (during a previous upgrade or installation), please follow the instructions in this section to do this.

Register all domain members as clients of Tru64 UNIX V5.1B, as follows:

```
atlas0# sra edit
sra> sys
sys> update ris all
```

Note:

In response to the `sra` command above and in response to subsequent `sra` commands, you may receive ongoing warnings to upgrade the SC Database. These warnings should be ignored, and once all domains have been upgraded, the procedure as described in Section 4.7.2 will be used to upgrade the database revision.

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to upgrade the remaining domains, as described in Section 4.7.

Upgrade without a Management Server

4.6.2.8 Upgrade the First Domain

Upgrade the first domain by performing the following steps:

1. Attach to the console of the first node of the first domain.
2. Unpack the Tru64 UNIX Version 5.1B-3 patch kit tar file in the `/patches` directory.

The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:

`<http://www.itrc.hp.com/>`

or from your local HP support representative.

3. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.
4. Mount the CD-ROM as the root user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:
`atlas0# mkdir /cdrom`
 - b. Mount the CD-ROM as follows:
`atlas0# mount -r /dev/disk/cdrom0c /cdrom`
 - c. Create a directory for the HP AlphaServer SC subsets as follows:
`atlas0# mkdir /usr/sckit`
 - d. Change to directory `cdrom` as follows:
`atlas0# cd /cdrom`
 - e. Copy the files to that location as follows:
`atlas0# tar cf - . | (cd /usr/sckit; tar xvf -)`
 - f. Unmount the CD-ROM as follows:
`atlas0# cd /; umount /cdrom`
5. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* into the disk drive.
6. Mount the CD-ROM as the root user, as follows:
`atlas0# mount -r /dev/disk/cdrom0c /cdrom`
7. Perform an upgrade check on the first domain as follows (approx 2 minutes duration):
`atlas0# sc_upgrade check`

If any errors are reported, take the appropriate steps to correct the situation. If any changes are required, run the command once again.

8. Begin the upgrade of the first domain as follows (approx 10 minutes duration):

```
atlas0# sc_upgrade setup -sckit /usr/sckit/kits \  
-unixpatch /patches/patch_kit -patchonly -cd /cdrom
```

Upgrade without a Management Server

This command will prompt for the root password for the domain. During this process, all other nodes will be halted. The `-patchonly` flag indicates that the operating system does not need to be updated.

9. Halt Node 0, first ensuring you are connected to the console of Node 0, as follows:

```
atlas0# shutdown -h now
```

10. Boot Node 0 to single-user mode, as follows:

```
P00>>> boot -fl s
```

11. Mount Local Filesystems:

```
atlas0# bcheckrc
```

12. Continue the upgrade process, as follows (approx 2.5 hours duration):

```
atlas0# sc_upgrade install -sckit /usr/sckit/kits \
        -unixpatch /patches/patch_kit -patchonly -cd /cdrom
```

During this step, the old HP AlphaServer SC software will be removed, the Tru64 UNIX Version 5.1B-3 patch kit will be loaded, and the new HP AlphaServer SC software will be installed onto the system.

13. When the install process is complete, reboot Node 0, as follows:

```
atlas0# shutdown -r now
```

Node 0 will now begin the process of configuring all the new kits and building a new kernel. It will automatically reboot following the kernel build. Once the system reaches multi-user mode, this node will be running the new version of the operating system.

14. In HP AlphaServer SC Version 2.6 (UK2), the `msql` daemon underlying the SC database has been updated. Earlier, in Section 4.6.2.3, the original SC database was backed up to a file. To restore the SC database from file, run the following command:

```
atlas0# /usr/sckit/kits/sc_post_upgrade
```

Note:

The `sc_post_upgrade` command will simply restore the existing database. You may see the following warning message:

```
rmstbladm: Warning: table sc_procurve is not present in backup
file
```

The migration of the SC database will introduce this new table and other new database settings required by the HP AlphaServer SC software release.

The `sc_post_upgrade` command will display the following error messages, which can be ignored:

Upgrade without a Management Server

```
Setting RMS attributes...
rcontrol: Error: attribute where name='cleanup-timeout'
already exists in database
rcontrol: Error: attribute where name='ignore-cleanup-failure'
already exists in database
```

15. New versions of the HP AlphaServer SC software may include new tables, or new fields within existing tables in the SC Database. It is necessary to upgrade your database to include these new tables and fields and to populate them with appropriate data. To migrate the database, log back into the system, and run the following command:

```
atlas0# sc_upgrade_db
```

16. It is now necessary to update the new database to reflect the current status of the upgrade procedure. This update is necessary because the commands run in step 12 made changes to the first domain that are not reflected in the database.

Update the database as follows:

```
atlas0# rmsquery "update sc_domain \
        set current_upg_state='Post_Installed' \
        where name='atlasD0' "
atlas0# rmsquery "update sc_domain set desired_upg_state ='Upgraded' \
        where name='atlasD0' "
```

17. Continue the upgrade process, as follows (approx 10 minutes duration):

```
atlas0# sc_upgrade post_install
```

At this stage member 1 will be upgraded, but it will be running in a degraded mode. The script will automatically begin the version switch process, and reboot the lead-node.

18. Configure the preferred server cluster aliases, as follows (approx 2 minutes duration):

```
atlas0# sc_upgrade member_config_aliases
atlas0# sra boot -nodes 'atlas[1-31]'
```

19. Boot all non-lead members as follows:

```
atlas0# sra boot -nodes 'atlas[1-31]'
```

20. Complete the upgrade process, as follows (approx 2 minutes duration):

```
atlas0# sc_upgrade clean
atlas0# rmsquery "update sc_domain \
        set current_upg_state='Upgraded' \
        where name='atlasD0' "
```

This removes all files that are no longer required and updates the current upgrade state in the database. At this point, the first domain is upgraded.

4.7 Upgrading the Domains

If you have not already upgraded the management server, clustered management server, or first domain (for systems with no management server), please do so now. See Section 4.4, Section 4.5, and Section 4.6 for more information.

Once the management server, clustered management server, or first domain (for systems with no management server) is updated, the steps in this section can be followed to back up and subsequently upgrade the remaining domains.

In this section, the term management server is used frequently. For a clustered management server, you can run the upgrade commands from either node of the clustered management server. In the case where there is no management server, then the first node of the first domain should be substituted for each occurrence.

It is recommended that you perform all of the following steps on the console of the management server. This will ensure that the commands are not interrupted by network events, and the screen output of the commands will be preserved for subsequent review if required.

At this point in the upgrade, the management server is running the latest software. However, some features of the software will only be fully functional when all of the domains are upgraded to the same level of software (for example, RMS and scrun). Do not use these software features until you have upgraded all of the domains. The commands in this section do not rely on these features.

The steps to upgrade the domains are as follows:

1. Back Up the Domains (see Section 4.7.1 on page 4–41)
2. Upgrade all of the Domains (see Section 4.7.2 on page 4–42)

4.7.1 Back Up the Domains

Before proceeding with the upgrade, you need to perform a backup of all the domains. When performing a backup of all the domains to the backup cluster disk, ensure that the backup cluster disk is the same size as the cluster disk to avoid any backup storage problems. For more information on disk layout for backup storage, see Section 2.4.1.

The backup process will archive one kernel image per domain as opposed to one image per domain member. It will archive the `genvmunix` kernel.

To ensure that the generic kernel image is up to date on each domain, run the following command:

```
atlasms# sra command -domains all -member 1 -command "BuildKernels"
```

To back up all the domains, run the following command:

Upgrading the Domains

```
atlasms# sra upgrade -domains all -backupdev dsk5
```

where `dsk5` corresponds to the disk selected as the backup disk within each domain. Where a different disk is used for different domains, then the individual commands can be run per domain as appropriate.

This command can take approximately 20 minutes, depending on the size of the cluster, /usr and /var domains.

You should check the console logs of the lead nodes of all domains after completing the `sra upgrade` command for any erroneous messages. Some errors may have occurred that are not visible to the controlling `sra upgrade` command.

You are now ready to upgrade all of the domains, as described in Section 4.7.2.

4.7.2 Upgrade all of the Domains

To upgrade the domains, perform the following steps on the upgraded management server:

1. If the management server has any NFS mounts from the domain, or if any domains are NFS mounting from any other domains, please unmount them now as follows:

```
atlasms# mount | grep atlas
atlasms# umount mountpoint
atlasms# sra command -domains all -member 1 -command 'mount | grep atlas'
atlasms# sra command -domains all -member 1 -command 'umount mountpoint'
```

2. Log into each domain, and check that all members in each domain are operational as follows:

```
atlasms# sra command -domains all -member 1 \
    -command 'hwmgr -v c | grep DOWN'
```

Boot any members that are identified as not operational. If you are unable to boot members that are not operational, remove the nodes from the domain by running the following command (while logged into the offending domain):

```
atlas0# sra delete_member -nodes atlas5
```

Note:

The upgrade will fail if all members are not booted when the upgrade begins.

3. If present, de-install the OTABASE subset (part of the Compaq Fortran kit) from each domain as follows:

```
atlas0# setld -i |grep OTABASE
OTABASE219 installed Compaq Compiled Code Support
Library #219
atlas0# setld -d OTABASE219
```

4. If present de-install the SYSCHECK utility from each domain as follows:

```
atlas0# setld -i |grep SYSCHECK
SYSCHECK128          installed          Syscheck utility
atlas0# setld -d SYSCHECK128
```

5. Tru64 UNIX Version 5.1B-3 includes Ladebug patches that can only be applied to the default version of Ladebug that shipped with Tru64 UNIX Version 5.1B.

If an updated Ladebug subset has been installed, remove it now, and replace it with the default version of Ladebug (which is available on the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*), as follows:

```
atlas0# setld -i |grep LDB
LDBBASE467 installed Ladebug Debugger Version 467
LDBDOC467 installed Ladebug Debugger Version 467 Documentation
OSFLDBBASE540 not installed Ladebug Debugger Version 467 (Software Development)
OSFLDBDOC540 not installed Ladebug Debugger Version 467 Documentation (Software Development)
atlas0# setld -d LDBBASE467 LDBDOC467
```

Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.

```
atlas0# mkdir /cdrom; mount -r /dev/disk/cdrom0c /cdrom
atlas0# setld -l /cdrom/ALPHA/BASE OSFLDBBASE540 OSFLDBDOC540
```

6. Unpack the Tru64 UNIX Version 5.1B-3 patch kit tar file in the /patches directory.

The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:

<<http://www.itrc.hp.com/>>

or from your local HP support representative.

7. Insert the *HP AlphaServer SC System Software CD-ROM* in the disk drive.

8. Mount the CD-ROM as the root user, as follows:

- a. Create a mount point for the CD-ROM, by running the following command:

```
atlas0# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlas0# mount -r /dev/disk/cdrom0c /cdrom
```

- c. Create a directory for HP AlphaServer SC subsets as follows:

```
atlas0# mkdir /usr/sckit
```

- d. Change to directory cdrom as follows:

```
atlas0# cd /cdrom
```

- e. Copy the files to that location as follows:

```
atlas0# tar cf - . | (cd /usr/sckit; tar xvf -)
```

- f. Unmount the CD-ROM as follows:

Upgrading the Domains

```
atlas0# cd /; umount /cdrom
```

9. Perform an upgrade check on each domain before starting the upgrade (approx 10 minutes duration):

```
atlasms# sra upgrade -domains list -unixpatch /patches/patch_kit \  
-sckit /usr/sckit/kits -endstate Checked -redo Pre_Upgrade
```

where:

`-domains list` specifies that the software should be upgraded on the domains. Replace `list` with `atlasD [0-3]` if you have a management server, or replace `list` with `atlasD [1-3]` where there is no management server.

`-unixpatch /patches/patch_kit` specifies the directory in which you had unpacked the Tru64 UNIX patch kit software (see step 5 above).

`-sckit /usr/sckit/kits` specifies the directory containing the HP AlphaServer SC software (see step 7 above).

`-endstate Checked` specifies an upgrade check is performed. The check installs the SRAUPG subset, if needed, on each domain and then performs the upgrade checks. If a domain fails the check then refer to the console log of member 1.

`-redo Pre_Upgrade` specifies that the upgrade check will start the process from the `Pre_Upgrade` state. Refer to Table 4–1 on page 4–4 for details on the upgrade states.

10. Upgrade each domain as follows (approx 10 hours duration consisting of 5 hours for member1 upgrade, and 5 hours for re-adding remaining members):

```
atlasms# sra upgrade -domains list -unixpatch /patches/patch_kit \  
-sckit /usr/sckit/kits
```

where:

`-domains list` specifies that the software should be upgraded on domains. Replace `list` with `atlasD [0-3]` if you have a management server, or replace `list` with `atlasD [1-3]` where there is no management server.

`-unixpatch /patches/patch_kit` specifies the directory in which you had unpacked the Tru64 UNIX patch kit software (see step 5 above).

`-sckit /usr/sckit/kits` specifies the directory containing the HP AlphaServer SC software (see step 7 above).

The script starts an automatic upgrade. When the script completes, the domain is upgraded.

Note:

Check the console logs of the domain lead nodes after completing the `sra upgrade` command for any erroneous messages. Some errors may have occurred that are not visible to the controlling `sra upgrade` command.

To check on the upgrade status of the domain, use the following command:

```
atlasms# sra upgrade_info -domains all
```

If you encounter problems during the upgrade, please refer to the section Recover from Failures during an Upgrade (see Section 4.8 on page 4–45).

You are now ready to perform any post-upgrade tasks as described in Post-Upgrade Tasks (see Section 4.9 on page 4–49).

4.8 Recover from Failures during an Upgrade

Upgrade may fail for a variety of reasons, and the failures can normally be divided into four categories: Minor, Medium, Serious, and Severe. For each of these categories, nodes, whole domains, or multiple domains may fail to upgrade. In all cases, an attempt should be made to find the underlying reason for the failure before continuing with the recovery steps outlined below. Checking of console logs and domain/node activity at the time of failures can be a good starting point to finding these reasons and resolving the underlying problems. Once you are satisfied that the problems have been identified and fixed, proceed as in the following sections.

4.8.1 Minor Failures: Re-add Members that Failed to Upgrade

If one or two members failed to upgrade, but the domain reports a `current_upg_state` of `Upgraded` in the `sc_domain` table, you should first check if the members just failed to boot into the newly upgraded domain. This can be checked in the console log of the failing member(s).

If this is the case, you can simply boot the member(s) into the domains as follows:

```
atlasms# sra boot -nodes list of failed nodes
```

If you can determine that the members failed to be re-added to the domain for another reason, but the domain in question is reporting the `current_upg_state` of `Upgraded` in the `sc_domain` table, re-add the failed members as follows:

```
atlasms# sra install -nodes list of failed nodes
```

Recover from Failures during an Upgrade

4.8.2 Medium Failures: Re-run the sra upgrade Command

If the `sra upgrade` command finishes but reports an error on one or more domains, check the `current_upg_state` in the `sc_domain` table for the domain in question.

Based on the `current_upg_state`, and with the help of the lead node console logs, you should try and determine what caused the upgrade to stop. If possible, the problem should be resolved before continuing. Once the problem is identified and resolved, run the upgrade command again as in the following example:

```
atlasms# rmsquery -v
sql> select name,startnode,current_upg_state,desired_upg_state from sc_domain
name          startnode current_upg_state desired_upg_state
-----
atlasD0        0             Null           Null
atlasD1        2             Null           Null
atlasD2        4             Setup          Upgraded
atlasD3        6             Upgraded        Upgraded
```

In this example, domain `atlasD2` is still at the `Setup` stage, while domain `atlasD3` has completed the upgrade. To restart the upgrade for the `atlasD2` domain, run the following command:

```
atlasms# sra upgrade -domains atlasD2 -unixpatch /patches/patch_kit \
-sckit /cdrom /kits
```

Note:

If you are not using the management server as the RIS server, you need to specify the RIS server using the optional `-rishost name-of-ris-server` option in the `sra upgrade` command.

This command will pick up the current upgrade state for `atlasD2` from the `sc_domain` table as described above, and continue the upgrade process from that point.

4.8.3 Serious Failures: Recover from Backup

If you encounter severe problems during an upgrade that cannot be overcome by the instructions in Section 4.8.1 or Section 4.8.2 above, you may have to restore the domain from the backup created in Section 4.6.1 or Section 4.7.1.

Unlike the other recovery methods mentioned above, you will need to restore from backup before fixing the underlying problems and repeating the upgrade process.

The following severe failure may occur:

Recover from Failures during an Upgrade

- If the lead member is rebooted during the setup phase of upgrade (for example, power outage, or accidental shutdown), then subsequent attempts to boot the lead member again will result in the panic:

```
ics_elan: elan device does not appear to be configured
```

This happens because the old HP AlphaServer SC software is already removed from the system and therefore the `elan.mod` kernel module is gone from the system. This kernel module is required by TruCluster when the lead member boots, and this is why TruCluster causes the panic.

The recovery action at this point is to restore the cluster from backup and start the upgrade again.

For further potential problems, you should read the current version of the *HP AlphaServer SC Release Notes*, particularly any information about upgrade installations. You should also check with your account representative so that you are aware of any *HP AlphaServer SC Support Bulletins* relating to the upgrade procedure.

To recover from backup, perform the following steps for each failed domain:

1. Halt the domain by running the following command:

```
atlasms# sra reset -dom name-of-domain
```
2. Connect to the console of the lead node
3. Boot the lead node using the Tru64 UNIX disk (standalone):

```
P00>>> boot dka200
```
4. When the prompt is displayed, log in to the lead node as `root`.
5. Restore the previous version of the operating system from backup as follows:

To restore using `/dev/disk/dsk5` as the restore device:

```
# /usr/opt/sra/bin/sra_cluster_backup -disk dsk5 restore
```

This will restore the `cluster /`, `/usr`, and `/var` filesystems, restore the boot partition of the lead node and the CNX partition.

Note:

The `sra_cluster_backup` script will RIS boot the other nodes and restore their boot and CNX partition. Following this, these nodes will be halted.

When restoring a domain from backup, if you encounter any errors reporting that the backup `cluster /`, `/usr`, and `/var` filesystems do not exist, refer to the troubleshooting item *Restoring from Backup Can Sometimes Fail* (see Section 11.23.7 on page 11–39).

6. Shut down the lead node as follows:

Recover from Failures during an Upgrade

```
# shutdown -h now
```

7. Disconnect from the console of the lead node.
8. Update the `sc_nodes` table to ensure that all nodes in the domain are set to bootable as follows:

```
atlasms# rmsquery "update sc_nodes set bootable='1' \  
                  where domain_name='name-of-domain'"
```

9. Query the `sc_domain` table to ensure that the `current_upg_state` of each domain is set to `Pre_Upgrade` as follows:

```
atlasms# rmsquery "update sc_domain set current_upg_state='Pre_Upgrade' \  
                  where name='name-of-domain'"
```

10. Boot the domain as follows:

```
atlasms# sra boot -domain name-of-domain
```

At this stage, the domain will have been restored to the backups created earlier.

If you intend to repeat the upgrade process on the domain, you should fix the problems with the upgrade, and then restart the upgrade procedure on this domain. When restarting the upgrade process, it is not necessary to repeat the backup of the domain just restored. Instead, you should restart the upgrade process from the step immediately after the domain backup process, specifically Section 4.6.2 or Section 4.7.2 as appropriate.

4.8.4 Severe Failures: Full System Restoration

If at any point you decide to defer the entire system upgrade until another time, and return the system to production using the original software image, then you should proceed as follows:

1. Restore all domains from their backups, as described on Section 4.8.3.
2. The RMS startup script should be restored on each domain, the new customized kernels should be built, and the domains should be rebooted as follows:

```
atlasms# sra command -domains all -member 1 \  
          -command "mv /sbin/init.d/rms.disabled /sbin/init.d/rms"  
atlasms# sra command -domains all -member 1 \  
          -command "BuildKernels"  
atlasms# sra command -domains all -member 1 \  
          -command "DeployKernels -g"  
atlasms# sra shutdown -domains all
```

3. For systems with a management server or clustered management server, once the above steps have been completed on the domains, you must restore the management server to its original image using the backups created in Section 4.4.1 or Section 4.5.1.
4. Once the management server is running the old software, you should boot all domains as follows:

```
atlasms# sra boot -domains all
```

At this point, the system will be running the original software on both the management server and on the domains. To conclude the system restoration, you should mount the SCFS and PFS file systems and you should configure in the nodes into the RMS partitions and start the RMS partitions.

4.9 Post-Upgrade Tasks

Once the upgrade has been completed, the following post-upgrade tasks should be performed:

1. Create Alternate Boot Partitions (see Section 4.9.1 on page 4–49)
2. Upgrade the Database Revision (see Section 4.9.2 on page 4–49)
3. Re-Add Members Deleted prior to Upgrade (see Section 4.9.3 on page 4–50)
4. Re-enable MSQl Cookie Mechanism (see Section 4.9.4 on page 4–50)
5. Add the cmf Archive entry into crontab (see Section 4.9.5 on page 4–50)
6. Add the rms Archive entry into crontab (see Section 4.9.6 on page 4–50)
7. Verify Network Configuration (see Section 4.9.7 on page 4–50)
8. Check that Members 2 and 3 Have a Vote (see Section 4.9.9 on page 4–51)
9. Configure Nodes into RMS Partitions (see Section 4.9.10 on page 4–51)
10. Setting Up the SC Monitor System (see Section 4.9.11 on page 4–51)
11. Assigning Cabinets in the SC Database (see Section 4.9.12 on page 4–51)
12. Reinstall Uninstalled Kits (see Section 4.9.13 on page 4–51)
13. Build and Deploy Generic Kernels (see Section 4.9.14 on page 4–52)

4.9.1 Create Alternate Boot Partitions

If alternate boot disks are in use, then update the alternate boot partitions on all nodes in the system, by running the following command from the management server (or rishost):

```
atlasms# sra copy_boot_disk -nodes all
```

For more information on this command, refer to the *HP AlphaServer SC System Administration Guide*.

4.9.2 Upgrade the Database Revision

The database migration performed in Section 4.4.2.10 does not upgrade the database revision.

Post-Upgrade Tasks

To upgrade the database revision, run the following command on the management server (or rishost):

```
atlasms# sc_upgrade_db -rev
```

4.9.3 Re-Add Members Deleted prior to Upgrade

Before the upgrade, some members may have been deleted. Add those members back in now by running the following command:

```
atlasms# sra install -nodes list of nodes
```

4.9.4 Re-enable MSQL Cookie Mechanism

To check if the cookie mechanism needs to be re-enabled, run the following command from the management server (or rishost):

```
atlasms# sra cookie
MSQL cookies are currently disabled
```

If the message above is displayed, run the following command:

```
atlasms# sra cookie -enable yes
```

For more information on this command, refer to Chapter 3 of the *HP AlphaServer SC System Administration Guide*.

4.9.5 Add the cmf Archive entry into crontab

The console log files can be archived automatically by adding the following crontab entry:

```
0 20 13,27 * * /sbin/init.d/cmf rotate
```

This will result in the console log files being archived on the 13th and 27th day of the month.

4.9.6 Add the rms Archive entry into crontab

The RMS database tables can be archived automatically by adding the following crontab entry:

```
5 2 * * * /usr/bin/rmsbackup
```

This will result in the RMS database tables being archived and backed up at 2:05am every day.

4.9.7 Verify Network Configuration

During the post-upgrade process, you should verify your network configuration. For more information on default route configuration, refer to Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

For more information on preferred server aliases, refer to Chapter 19, of the *HP AlphaServer SC System Administration Guide*.

4.9.8 Verifying Local Files and Customizations

The upgrade process endeavors to preserve many local files and customizations. However, as advised in Section 4.1.5 on page 4–5, some customizations may be affected by the upgrade process. You should now check that your original customizations and settings have been preserved correctly during the upgrade.

This is especially important if the upgrade failed and was restarted from a checkpoint, or if the upgrade needed to be restarted after having restored the domain from backup.

4.9.9 Check that Members 2 and 3 Have a Vote

At the start of the upgrade process all members votes, except member1, are removed. The process re-adds the member votes in the upgrade "clean" phase. However, this step can fail for a number of reasons. Check that members 2 and 3 have a vote using the `clu_quorum` command.

4.9.10 Configure Nodes into RMS Partitions

During the upgrade, all nodes in the domains will be automatically configured out of the RMS partitions. These nodes should now be configured back in using the command:

```
atlasms# rcontrol configure in nodes "atlas[0-1023]"
```

4.9.11 Setting Up the SC Monitor System

For systems with a management server, please complete the instructions documented in Set Up the SC Monitor System (see Section 5.6 on page 5–50).

For systems without a management server, please complete the instructions documented in Set Up the SC Monitor System (see Section 6.5 on page 6–35).

4.9.12 Assigning Cabinets in the SC Database

For systems with a management server, please complete the instructions documented in Assign Cabinets in the SC Database (see Section 5.7 on page 5–52).

For systems without a management server, please complete the instructions documented in Assign Cabinets in the SC Database (see Section 6.6 on page 6–38).

4.9.13 Reinstall Uninstalled Kits

During the upgrade, you may have needed to uninstall kits such as PFSMOD, OTABASE, SYSCHECK, and System Event Analyzer/WEBES from your management server (if your system has a management server) and your domains. Re-install any uninstalled kits now.

Post-Upgrade Tasks

4.9.14 Build and Deploy Generic Kernels

Once you are satisfied that the upgrade was successful and your applications are running smoothly for approximately 24 hours, it is recommended that you deploy the new generic kernels globally. If you do not perform this action, then a `delete member` followed by re-install of that member will attempt to boot initially on a non-compatible generic kernel.

From the management server (or Node 0 if you do not have a management server), run the following commands:

```
atlasms# scrun -d all "BuildKernels"  
atlasms# scrun -d all "DeployKernels -g"
```

This will build and deploy the generic kernels to all members of all domains.

You have now completed the Upgrade Installation Procedure.

Installing: When the System Has a Management Server

This chapter describes how to install an HP AlphaServer SC system that has a management server.

Note:

If your system does not have a management server at installation time, do not use this chapter — refer to Chapter 6 (*Installing: When the System Does Not Have a Management Server*) instead.

If you wish to add a management server to an HP AlphaServer SC system later — that is, after domain creation — use the checklist provided in Appendix C to ensure that you complete all installation tasks in the correct order.

Note:

If your system has a clustered management server, refer to Section 5.1.7.

When installing software on an HP AlphaServer SC system that has a management server, install the software on the management server first (see Chapter 1 of the *HP AlphaServer SC System Administration Guide* for more information about management servers).

For information on helpful tips and guidelines that may assist you when performing an HP AlphaServer SC system installation, see Section 11.1.

The information in this chapter is structured as follows:

- Set Up the Management Server (see Section 5.1 on page 5–2)
- Set Up the SC Database (see Section 5.2 on page 5–37)
- Configure Out All Nodes During Installation (see Section 5.3 on page 5–45)
- Check All Nodes in the HP AlphaServer SC System (see Section 5.4 on page 5–45)

Set Up the Management Server

- Configure and Diagnose the HP AlphaServer SC Interconnect (see Section 5.5 on page 5–47)
- Set Up the SC Monitor System (see Section 5.6 on page 5–50)
- Assign Cabinets in the SC Database (see Section 5.7 on page 5–52)

Note:

Use the checklist provided in Appendix A, to ensure that you complete all installation tasks in the correct order.

5.1 Set Up the Management Server

Setting up the management server involves the following tasks:

- Set the Console Variables (see Section 5.1.1 on page 5–3)
- Check the System Firmware (see Section 5.1.2 on page 5–4)
- Creating Bootable Devices on EMA/EVA Storage (see Section 5.1.3 on page 5–4)
- Install the Tru64 UNIX Operating System (see Section 5.1.4 on page 5–6)
- Customize the System Configuration (see Section 5.1.5 on page 5–10)
- Install the Operating System Patch Software (see Section 5.1.6 on page 5–22)
- Install and Configure the Clustered Management Server (see Section 5.1.7 on page 5–22)
- Configure the RIS Server (see Section 5.1.8 on page 5–31)
- Install the HP AlphaServer SC System Software (see Section 5.1.9 on page 5–32)
- Install the HP Fortran Run-Time Libraries (see Section 5.1.10 on page 5–33)
- Install Layered Products (Optional) (see Section 5.1.11 on page 5–33)
- Install the SANworks Storage System Scripting Utility (see Section 5.1.12 on page 5–33)
- Define the RMS Master Node (rmshost) (see Section 5.1.13 on page 5–34)
- Build the New Kernel and Reboot (see Section 5.1.14 on page 5–35)

5.1.1 Set the Console Variables

Before installing the Tru64 UNIX operating system on the management server, you must configure the system console. Display the SRM console prompt on the management server as follows:

- If your system has a factory-installed software (FIS) kernel, it will automatically start to boot the Tru64 UNIX operating system at power on. When prompted to continue this boot, enter **No** to return to the SRM console prompt.
- If your system does not have a FIS kernel, the system will display the SRM console prompt at power on.

To set the console variables, enter (at the SRM console prompt) the commands in Table 5–1.

Table 5–1 Setting the Console Variables

```
P00>>> set auto_action HALT
P00>>> set eia0_mode FastFD
P00>>> set eib0_mode1 FastFD2
P00>>> set boot_osflags A
P00>>> set boot_reset off
P00>>> set os_type UNIX
P00>>> set console serial
P00>>> set sys_com1_rmc off
P00>>> set ocp_text nodename3
P00>>> set bootdef_dev ''
P00>>> set pci_parity on
```

¹You need only set the `eib0_mode` variable on nodes that have an external network interface. You should set the `eia0_mode` variable on all nodes, for the management network.

² When setting the `eib0_mode` variable, specify the appropriate network speed.

³You may not need to set the `ocp_text` variable on the management server.

The remaining console variables for the management server, and all console variables for the remaining nodes, are automatically set during the installation process (see Section 5.2, step 20 on page 5–42).

For more information about the SRM console, see Chapter 2 of the *HP AlphaServer ES40 Owner's Guide*.

You are now ready to check the system firmware, as described in Section 5.1.2.

Set Up the Management Server

5.1.2 Check the System Firmware

Table 5–2 lists the minimum firmware versions supported by HP AlphaServer SC Version 2.6 (UK2) for management servers. If the management server is HP AlphaServer ES40 or HP AlphaServer ES45 based, then the minimum firmware revisions are given in Table 5–11.

Table 5–2 Minimum System Firmware Versions

Firmware Version	HP AlphaServer DS20E
SRM Console	6.2-1
OpenVMS PALcode	1.96-77
Tru64 UNIX PALcode	1.90-72
Serial ROM	1.13-44

Check that the management server meets these minimum system firmware requirements, by entering the `show config` command at the SRM prompt.

If necessary, update the system firmware, as described in Chapter 21 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to create bootable devices on EMA/EVA storage, as described in Section 5.1.3.

5.1.3 Creating Bootable Devices on EMA/EVA Storage

If your target management server is a standalone one, then you should skip this section and proceed to Section 5.1.4.

If your target management server is a cluster without a local disk, the target disk for the initial Tru64 installation is located on the EMA or EVA storage system.

Note:

It is assumed that you have already performed the steps in Configure the EMA/EVA Storage System in Section 3.14 on page 3–27 and Section 3.15 on page 3–44.

Before installing Tru64 UNIX from the installation CD-ROM, perform the following steps to create the SRM boot variable for the target disk based on the EMA/EVA storage system:

1. Initialize the hardware as follows:

```
P00>>> init
```

- On node 0 of the management cluster, locate the virtual disks presented to this system by the EVA storage system as follows:

```
P00>>> wwidmgr -show wwid
[0] UDID: 1 WWID:01000010:6005-08b4-0001-0130-0001-a000-014f-0000 (ev:none)
[1] UDID: 2 WWID:01000010:6005-08b4-0001-0130-0001-a000-0154-0000 (ev:none)
[2] UDID: 3 WWID:01000010:6005-08b4-0001-0130-0001-a000-0159-0000 (ev:none)
[3] UDID: 4 WWID:01000010:6005-08b4-0001-0130-0001-a000-015e-0000 (ev:none)
[4] UDID: 5 WWID:01000010:6005-08b4-0001-0130-0001-a000-0163-0000 (ev:none)
```

- Check the UDID numbers against the OS_UNIT_ID parameters used in the Chapter 3 for EMA/EVA.

For the example given in Create Virtual Disk Units for the Clustered Management Server (see Section 3.15.2.3 on page 3–52), the UDID numbers are assigned as follows:

```
UDID:1 Cluster File System (/ , /usr/ , /var)
UDID:2 quorum disk
UDID:3 Initial Tru64 disk
UDID:4 node0 boot disk
UDID:5 node1 boot disk
```

- Assign a wwid environment variable to the target disk for the initial Tru64 UNIX installation as follows:

```
P00>>> wwidmgr quickset -udid 3
```

Disk assignment and reachability after next initialization:

```
6005-08b4-0001-0130-0001-a000-015e-0000
      via adapter:          via fc nport:      connected:
dga3.1001.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3c8      No
dga3.1002.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3cd      Yes
P00>>>
```

- Assign a wwid environment variable to the target boot disk, to be used by the clu_create command, as follows:

```
P00>>> wwidmgr -quickset -udid 4
```

Disk assignment and reachability after next initialization:

```
6005-08b4-0001-0130-0001-a000-015e-0000
      via adapter:          via fc nport:      connected:
dga3.1001.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3c8      No
dga3.1002.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3cd      Yes

6005-08b4-0001-0130-0001-a000-014f-0000
      via adapter:          via fc nport:      connected:
dga4.1001.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3c8      No
dga4.1002.0.9.0             pga0.0.0.9.0          5000-1fe1-0013-a3cd      Yes
P00>>>
```

- Re-initialize the hardware as follows:

```
P00>>> init
```

The system is now ready for Tru64 UNIX installation.

Set Up the Management Server

Caution:

At Tru64 UNIX installation time, select the disk with IDENTIFIER=3. Remember that IDENTIFIER=4 is the target boot disk for the clusterized node. The target installation has to be labelled. Refer to Table 5-3 for the recommended partition layout. On node 1, create the SRM boot device for the target boot disk for UDID 5, as described above.

You are now ready to install the Tru64 UNIX operating system, as described in Section 5.1.4.

5.1.4 Install the Tru64 UNIX Operating System

Install the Tru64 UNIX operating system on the management server, from the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*, as described here and in the *HP Tru64 UNIX Installation Guide*.

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* into the disk drive.
2. This step only applies to clustered management servers. For clustered management servers, shared storage will be used for management server disks. The actual quantity of disks required is shown in Figure 3-2 on page 3-3. Before proceeding to the next step of the Tru64 UNIX operating system installation, you should follow your appropriate hardware documentation to make these shared disks visible as bootable devices to this management server. For example, the use of `wwidmgr` may be required for shared Fiber Channel storage as described in Section 5.1.3. Once the disks are visible from the SRM `show dev` command, you can proceed to the next step.
3. Initialize the hardware, as follows:
`P00>>> init`
4. Identify the CD-ROM device name by running the following SRM console command:
`P00>>> show device`

The `show device` command produces the following output:

dka0.0.0.1.1	DKA0	COMPAQ BD018122C9 B016
dka100.1.0.1.1	DKA100	COMPAQ BD018122C9 B016
dka200.2.0.1.1	DKA200	COMPAQ BD018122C9 B016
dqa0.0.0.15.0	DQA0	COMPAQ CDR-8435 0013
dva0.0.0.1000.0	DVA0	
eia0.0.0.2004.1	EIA0	00-50-8B-CF-46-CC
eib0.0.0.2005.1	EIB0	00-50-8B-CF-46-CD
pga0.0.0.3.1	PGA0	WWN-1000-0000-C922-391A
pka0.7.0.1.1	PKA0	SCSI Bus ID 7

In this example, the CD-ROM device name is `dqa0`.

5. Enter the following command at the SRM console prompt, to display the installation user interface:

```
P00>>> boot dga0
```

The system boot process can take several minutes. Several hardware-specific messages are displayed. The more complex the system (many peripheral devices, and so on), the longer the boot process takes.

Upon successful system boot, the installation user interface appears. The type of user interface presented during installation depends on the hardware configuration:

- Systems equipped with graphics consoles present a graphical user interface.
- Systems with consoles that do not have graphics capabilities present a text-based, menu-driven user interface.

The information you supply is the same regardless of the type of user interface, but the order in which it is requested may be different. This guide documents the graphical user interface.

Follow these guidelines:

1. When prompted to select a language in which to view the user interface, choose `United States English`.
2. The Installation Welcome dialog box appears. Click on the Next button.
3. The Host Information dialog box appears. You must set several values on this screen, as follows:
 - a. Set the host name. The host name should be the same as the network interface for the management network. For example, if the system name is `atlas`, the host name should be `atlasms`.
 - b. Set the area.
 - c. Set the location.
 - d. Set the date.
 - e. Set the time.

Click on the Next button.

4. The Set root Password dialog box appears. Set the root password for the system. The same root password is used for all nodes. Click on the Next button.
5. The Software Selection dialog box appears. You must set the value on this screen as follows:
 - a. Choose `All Software` when prompted for the software that you wish to install.
 - b. Click on the Next button.

Set Up the Management Server

- 6. The Kernel Options dialog box appears, prompting you to choose the type of kernel components to build into the kernel. Select `Customize` (you will choose the individual kernel options in step 13). Click on the `Next` button.
- 7. The Select File System Layout dialog box appears. Select `Customize File System Layout` and click on the `Next` button.
- 8. The Custom File System Layout dialog box appears. You must set several values on this screen, as follows:
 - a. Set the `Use LSM` option to `No`.
You must not use LSM at this stage; instead, configure LSM after all nodes have been added to the domain — see Section 8.11 on page 8–18.
 - b. Configure the recommended partition layout for the system disk.

Table 5–3 shows the recommended partition layout for an 18GB system disk, and for a 36GB disk, on a management server.

Table 5–3 Recommended Partition Layout for Management Server System Disk

File System	Disk	Partition	Size for an 18GB Disk	Size for a 36GB Disk	Type
root	dsk0	a	384MB	384MB	AdvFS
/usr	dsk0	d	5.4GB	11.1GB	AdvFS
swap1	dsk0	f	5.4GB	11.1GB	swap
/var	dsk0	e	5.4GB	11.1GB	AdvFS

See Chapter 6 of the *HP Tru64 UNIX Installation Guide* for more information about customizing the file system layout.

- 9. When you have entered all of the required information, click on the `Next` button.
- 10. The Installation Summary dialog box appears. Review the information on this screen. To change any values, click on the `Reset` button. When the information is correct, click on the `Finish` button.
- 11. The Ready to Begin Installation dialog box appears. Click on the `OK` button. The system saves the configuration, creates the file systems, and loads the software.
- 12. The system then automatically reboots from the system disk.

13. Software configuration occurs automatically after your system reboots. The Kernel Option Selection dialog box appears. You must select at least the options listed in Table 5–4.

Table 5–4 Minimum Kernel Options

Selection	Kernel Option
2	NTP V3 Kernel Phase Lock Loop (NTP_TIME)
3	Kernel Breakpoint Debugger (KDEBUG)
4	Packetfilter driver (PACKETFILTER)
12	ISO 9660 Compact Disc File System (CDFS)

If you need to install a kernel component after installation is complete, use the `doconfig(8)` command. See the *HP Tru64 UNIX Installation Guide* for more information.

14. When prompted to edit the configuration file, select `NO`.
15. The kernel build procedure automatically begins after software configuration.
16. After the kernel build, the system reboots automatically.
17. The final step is to log into the newly installed system as the `root` user. (You may see warnings about missing license PAKs for OSF-BASE — you can ignore these until Section 5.1.5.1 on page 5–10.)

When you log in for the first time, the Tru64 UNIX System Setup dialog box appears. Click on the Custom Setup icon. The Tru64 UNIX Custom Setup menu appears. Customize the system configuration, as described in Section 5.1.5.

Note:

If you choose Quick Setup, your configuration options are reduced — you will have to rerun the `sysman` command to complete the configuration.

Set Up the Management Server

5.1.5 Customize the System Configuration

The Tru64 UNIX Custom Setup menu offers a number of options, as illustrated in Figure 5–1.

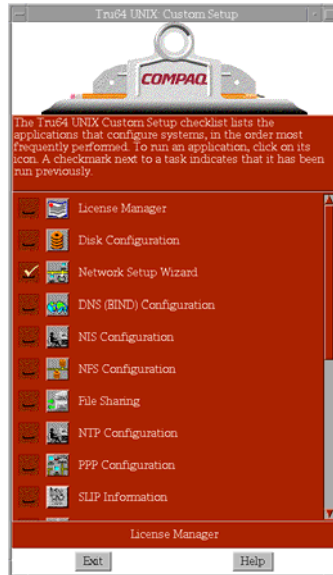


Figure 5–1 Tru64 UNIX Custom Setup Menu

Not all of these options are mandatory. This guide documents only the systems that you must configure. For more information about the Custom Setup Menu options, see the *HP Tru64 UNIX Installation Guide*.

First register the licenses, as described in Section 5.1.5.1.

5.1.5.1 Register Licenses (PAKs)

A license is a contract with terms and conditions. Each license has an associated Product Authorization Key (PAK), which is a set of characters.

When building the management server, you may choose to install the TCS-UA PAK specifically for the management server node type. This is not necessary and you should install the bundled PAKs, for all nodes in the system, onto the management server. Later on in the installation process, the licences database from the management server will automatically be reused when installing the domains and adding members.

Set Up the Management Server

You must use the license manager to enter the license information on each PAK into the license database; this process is called registering a license. Each time a user attempts to run a licensed product, the product calls the license-checking functions to be sure that the license allows the user to use the product.

For more information about licenses, see the *Tru64 UNIX Software License Management* manual.

Note:

The licenses listed below are those required to use software that is discussed elsewhere in this document. However:

- Not all of these licenses are required.
 - There are other licenses that may be useful. These licenses would be installed in a similar way to the licenses below, but are not covered in this document.
-

Click on the License Manager icon (on the Custom Setup menu) to register the following licenses:

- HP AlphaServer SC System Software License (TCS-UA). This license is required before you can create a cluster.
- HP AlphaServer SC System Software License (ASC). This license is required; you need this license to run the console logger daemon (`cmfd`), and to use the `sra setup` command. See below for more information.
- Remote Installation Services License (OSF-SVR). This license is required; it is essential for system operation.
- Operating System Base License (OSF-BASE). This license supports up to two interactive users.
- Symmetric Multiprocessing (SMP) Extension to Base License (OSF-BASE). This license is optional; it is only needed for each additional CPU on the system.
- Concurrent Use License (OSF-USR) or Unlimited Interactive User License (OSF-USR). These licenses are optional; they are only needed if you wish to support interactive users.
- Advanced File System Utilities License (ADVFS-UTILITIES). This license is required; it is needed if you wish to configure additional storage in an AdvFS file domain.

Set Up the Management Server

- HP AlphaServer SC Development Software License (OSF-DEV). This license is optional; it is only needed if you wish to use the Ladebug debugger (see below) or other products such as the C, C++, and Fortran compilers, as well as other performance and debugging tools.

HP AlphaServer SC System Software

The HP AlphaServer SC System Software License is sold in bundles of 1-, 16-, 32-, 64-, and 128-node licenses. You can combine any of these bundles so that the total number of licenses is equal to, or more than, the number of nodes in the HP AlphaServer SC system. For example, you can combine one 16-node license with one 32-node license for a 48-node HP AlphaServer SC system. The management server (if present) is not counted as a node — although the licences must be installed on the management server.

The *HP AlphaServer SC40 QuickSpecs* provide ordering information for the 1-, 16-, 32-, 64-, and 128-node licenses.

The license file for the HP AlphaServer SC System Software License contains two **Product Authorization Keys (PAKs)**: ASC and TCS-UA. Both of these PAKs must be installed on the system. A PAK contains a **units** field that identifies the capability of the PAK — the ASC PAK shown in Example 5–1 contains 6400 units.

Example 5–1 Sample ASC PAK

```
Issuer: DEC
Authorization Number: QS-SYS-16-NODE
Product Name: ASC
Producer: DEC
Number of units: 6400
Key Termination Date: 19-MAY-2001
Activity Table Code: CONSTANT=100
Checksum: 1-ABCD-GHDN-BOCJ-CLFH
```

Depending on the PAK, the units are used differently, as follows:

- ASC

In LMF terms, the ASC PAK is a **concurrent-use** PAK. The number of units determines the number of nodes that are licensed — 400 units are required to license a node. When you register and load the ASC PAKs, the LMF system combines the units of all of the ASC PAKs together. You will see messages to this effect when you install the second and subsequent ASC PAKs.

- TCS-UA

In LMF terms, the TCS-UA PAK is a **capacity-based** PAK. The number of units corresponds to the model of the AlphaServer — for example, a HP AlphaServer ES40 requires 1050 units. Unlike the ASC PAK, LMF does not combine the units of several TCS-UA PAKs. When you attempt to install a second or subsequent TCS-UA PAK,

Set Up the Management Server

LMF will print messages saying that the PAKs were not combined. This does not matter as you only require one TCS-UA PAK for the system to operate correctly. It does not matter if several TCS-UA PAKs are registered.

The ASC and TCS-UA PAKs are delivered as a shell script containing the LMF commands to register and install the PAKs. If you purchased several HP AlphaServer SC System Software Licenses (for example, a 16-node license and a 32-node license) at the same time, you will receive a single shell script containing the appropriate PAKs. To install the PAKs, copy the script to the system and execute the script. If you later buy additional licenses, you will receive a second shell script containing the appropriate additional PAKs — the new script does not contain PAKs for the original licenses. To ensure that all PAKs are installed, you should execute each shell script. You should also keep a copy of each license shell script, in case you ever need to do a complete system reinstall.

As mentioned earlier, when you run the second and subsequent shell scripts, you will see messages saying that the ASC PAKs were combined but that the TCS-UA PAKs were not. This is normal.

If you do not install the ASC and TCS-UA PAKs, or if the number of ASC units is not appropriate for the number of nodes in the system, the system will not operate correctly as follows:

- The `sra setup` command will print a warning when you first specify the number of nodes. You can ignore this warning and attempt to continue. However, you may be unable to operate parts of the system at a later stage.
- The console management daemon (`cmfd`) will not start. You will be unable to access node consoles. A message will be printed to the `/var/sra/adm/log/cmfd/cmfd_<hostname>_<port>.log` file.

The shell script automatically registers and loads the ASC and TCS-UA PAKs. If you inadvertently unload the PAKs, the system will not operate — that is, a PAK must be both registered and loaded for it to work. You can tell how many units are loaded as follows:

```
atlasms# lmf list full cache for ASC
```

You can load previously registered ASC and TCS-UA PAKs as follows:

```
atlasms# lmf load 0 ASC
atlasms# lmf load 0 TCS-UA
```

The Ladebug Debugger

The Ladebug debugger, distributed with Tru64 UNIX Version 5.1B, is a symbolic source-level debugger that supports debugging of ADA, C/C++, Fortran, and Fortran 90 applications. RMS uses the Ladebug debugger to print a back trace when an application core dumps.

Set Up the Management Server

To use the Ladebug debugger, you need the OSF-DEV license. You can obtain this license by purchasing, for example, HP AlphaServer SC Development Software, or Developer's Toolkit for Tru64 UNIX.

Note:

If you are not licensed to use the Ladebug debugger, RMS will not print a back trace.

When you have completed the license PAK registration, click on the Exit button to return to the Custom Setup Menu.

You are now ready to configure the networks, as described in Section 5.1.5.2.

5.1.5.2 Set Up Networks

Click on the Network Setup Wizard icon (on the Custom Setup menu) to configure the management and external networks. This displays a submenu as shown in Figure 5–2.

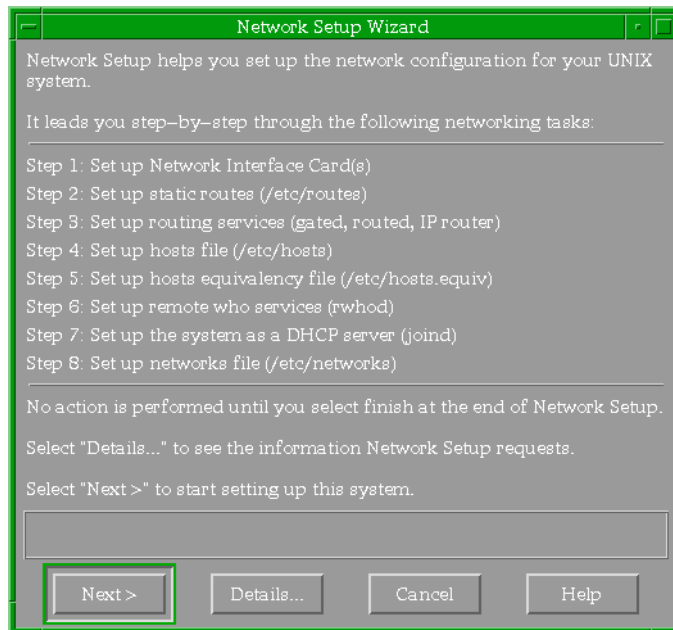


Figure 5–2 Tru64 UNIX Network Setup Wizard Menu

Configure the networks as follows:

1. Click on the Next button to display the Set up Network Interface Card(s) screen.

Set Up the Management Server

Set up the network interface cards using the settings provided in Table 5–5, where `atlas` is an example system name, and `x.x.x.x` and `y.y.y.y` are site-dependent values (recorded in Appendix D).

Table 5–5 Network Interface Cards on a Management Server

Network Type	Device Name	Host Name	IP Address	Network Mask
Management	ee0	atlasms	10.128.101.1	255.255.0.0
External	ee1	atlasms-ext1	x.x.x.x	y.y.y.y

2. Click on the Next button to display the Set Up Static Routes screen.

Set up a default static route, as follows:

- a. Click on the Add button to display the Network Setup: Set Up Static Routes: Add/Modify screen.
 - b. Set the Destination Type to Default Gateway.
 - c. In the Gateway field, enter the external gateway IP address (see Appendix C).
 - d. Set the Route Via option to Gateway, and click on the OK button.
3. Click on the Next button to display the Set Up Routing Services screen.

When configuring a management server, setting up routing services is an optional task.

Note:

For a clustered management server, it is mandatory to set up routing services using gated.

- If you do not wish to set up a routing service for the system, select None when asked Do you want to set up a routing service for this system.
- If you wish to set up routing services for the management server, set the routing services using the settings provided in Table 5–6:

Table 5–6 Routing Services

Question	Answer
Do you want to set up a routing service for this system:	Yes (use gated)
Do you want to run this system as an IP Router:	Yes

4. Click on the Next button to display the Set Up Hosts File screen.

Set Up the Management Server

The hosts file will automatically contain the entries listed in Table 5–7. The following notes apply to Table 5–7:

- atlas is an example system name.
- x.x.x.x represents a site-specific value (recorded in Appendix D).
- DNS servers and NIS servers will be added later (see Section 5.1.5.3 on page 5–17 and Section 5.1.5.6 on page 5–19 respectively).

Note:

The `sra setup` command will update the `/etc/hosts` file to add an entry for all of the hosts associated with the domain (for example, `atlas0`, `atlas1`, ..., `atlas127`) as well as each cluster alias (see Section 5.2, step 9 on page 5–39).

Table 5–7 Hosts File When Configuring a Management Server

IP Address	Host Name	Aliases/Comments
127.0.0.1	localhost	
10.128.101.1	atlasms ¹	
x.x.x.x	atlasms-ext1 ²	

¹This value is not displayed, but if you try to add it, you will be told that it is already in the hosts file.

²The `/etc/hosts` file may not contain the entries as shown in Table 5–7. If you are using the Network Setup Wizard, these entries will exist after you have completed all the Network Setup Wizard steps and select finish.

5. Click on the Next button to display the Set Up Hosts Equivalency File screen. Do not set up a hosts equivalency file.
6. Click on the Next button to display the Set Up Remote Who Services screen. Select No.
7. Click on the Next button to display the Set Up the System as a DHCP Server screen. Select No. Accept the default setting of Log No Messages.
8. Click on the Next button to display the Set Up Networks File screen. Click on the Next button to skip this step.
9. Click on the Finish button to return to the Custom Setup menu. When prompted to restart services, select Yes.

You are now ready to configure Domain Name Service / Berkeley Internet Name Daemon (DNS/BIND), as described in Section 5.1.5.3.

5.1.5.3 Configure DNS (BIND)

Note:

HP recommends configuring the HP AlphaServer SC Version 2.6 (UK2) system as a DNS Client as indicated below. If you wish to configure the system as a DNS Server, please refer to Appendix I.

Note:

If you have more than one DNS server, you need to add your site specific DNS servers.

If you wish to configure DNS, perform the following steps:

1. Click on the DNS (BIND) Configuration icon (on the Custom Setup menu) to display the DNS (BIND) Configuration screen.
2. Double click on the Configure System as a DNS Client entry.
3. When prompted to add the DNS server to the hosts file, enter the local domain and click Add.
4. When prompted to update the host name, select No (see Note below).

Note:

If you select Yes, the host name is changed. If the host name is changed, RMS will not work (because the `rmshost` attribute in the RMS database is wrong). If you change the host names (by selecting Yes), log onto each node in turn and change the host name back to the original value, by removing the domain name.

5. Click on the Exit button to return to the Custom Setup menu.

For more information about configuring DNS (BIND), see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network Time Protocol (NTP), as described in Section 5.1.5.4.

Set Up the Management Server

5.1.5.4 Configure NTP

Click on the NTP Configuration icon (on the Custom Setup menu) to configure the Network Time Protocol (NTP). You must know the name of an NTP server that is accessible through the external network interface on the first node (Nodes 0, 32, 64, or 96). Perform the following steps as the `root` user:

1. Double click on the Configure System as an NTP Client entry.
2. Click on the Add button, and enter the hostname of the NTP server on the external network.

Note:

The hostname of the NTP server does not need to be fully qualified. After network configuration, you can check to see that external network connection is valid by trying to `ping` one of the NTP servers. Otherwise the NTP restart will hang when it tries to start the service.

Note:

You can add as many NTP servers as your site requires.

3. Set the Mode to Server. Accept the default settings for Version and Key Number.
4. Click on the OK button to return to the Configure System as an NTP Client screen.
5. Click on the OK button to return to the Custom Setup Menu. When prompted to start `xntpd`, click on the Yes button.

For more information about configuring NTP, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network File System (NFS), as described in Section 5.1.5.5.

5.1.5.5 Configure NFS

Click on the NFS Configuration icon (on the Custom Setup menu) to display the NFS Configuration screen, and perform the following steps:

1. Double click on the Configure System as an NFS Client entry and accept the default settings.
2. Double click on the Configure System as an NFS Server entry and change the settings as follows:

Default Settings

Number of TCP Server Threads*: 8
Number of UDP Server Threads*: 8

New Settings:

Number of TCP Server Threads*: 32
Number of UDP Server Threads*: 96

where 128, as a sum of TCP and UDP, is the maximum allowed.

This will improve the RIS performance when completing the `sra install` step on all nodes as described in Chapter 7.

Ensure that the Enable Locking check box is selected (this is the default setting) and select OK to commit changes.

3. When prompted to restart the NFS daemons, select Yes.
4. Click on the Exit button to return to the Custom Setup menu.

For more information about configuring NFS, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network Information System (NIS), as described in Section 5.1.5.6.

5.1.5.6 Configure NIS

If you wish to configure the Network Information System (NIS), perform the following steps:

Note:

The installation of the HP AlphaServer SC RMS subset automatically adds a UNIX user named "rms" and a UNIX group named "rms" with specific identifiers.

However, the RMS subset installation will fail if any of the following already exist in the NIS database on the NIS master server:

- any group or user named "rms"
- any user with uid=15
- any group with gid=200

Such entries in the NIS database must be removed before proceeding. If you cannot remove such entries from the NIS database, then do not configure NIS at this time. NIS may be configured after the entire system is configured. However, this can be a security issue as NIS users with uid=15 or gid=200 will have access to RMS files on the HP AlphaServer SC system.

Set Up the Management Server

1. Click on the NIS Configuration icon (on the Custom Setup menu) to display the NIS Configuration screen.
2. Enter and confirm your system's NIS domain name.
3. If you already have a NIS master server, press Return to accept the default option to indicate that you are configuring a NIS client.
If you do not have an external NIS master server, but you wish to use NIS within the HP AlphaServer SC system, configure the management server as a master server.
4. When configuring a NIS client, complete the following sub-steps:
 - a. When prompted to use `ypbind` secure mode (option `-s`), select Yes.
 - b. When prompted to use `ypbind` domain locks (option `-S`), select Yes.
 - c. When prompted, enter the NIS MASTER server in the list of authorized NIS servers.
 - d. If prompted to add the NIS MASTER server to the hosts file, select Yes and then enter the relevant details.

When configuring a NIS master server, complete the following sub-steps:

- a. When prompted regarding enhanced security, select No.
 - b. When prompted regarding maintaining maps as "btree" files, select No.
 - c. When prompted for the names of the SLAVE servers, complete the list as desired.
 - d. When prompted to use `ypbind` secure mode (option `-s`), select Yes.
 - e. When prompted to use `ypbind` domain locks (option `-S`), select Yes.
 - f. When prompted, enter other NIS servers as appropriate in the list of authorized NIS servers.
 - g. If prompted to add these NIS servers to the hosts file, select Yes and enter the relevant details.
5. When prompted, choose option 3 to disallow all `ypset` requests.
 6. When prompted to configure the system to use all of the NIS databases, select Yes.
 7. When prompted to start the NIS daemons now, select Yes.
 8. Click on the Finish button to return to the Custom Setup menu.

For more information about configuring NIS, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure mail, as described in Section 5.1.5.7.

5.1.5.7 Configure Mail

Click on the Mail Configuration icon (on the Custom Setup menu) to configure mail.

Note:

You must configure mail.

Perform the following steps as the `root` user:

1. Double-click on the Configure Mail as a Client entry.
If DNS has not been set up, a warning is displayed.
2. Enter the mail server, and click the Commit button.
3. Select the Yes option to restart sendmail.
4. Select the Close option to return to the Custom Setup menu.

For more information about configuring mail, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure printers, as described in Section 5.1.5.8.

5.1.5.8 Configure Printers

If you wish to configure printers, click on the Printer Configuration icon (on the Custom Setup menu) and configure printers as described in Chapter 6 of the *HP Tru64 UNIX Network Administration* manual.

Note:

HP AlphaServer SC Version 2.6 (UK2) supports remote printing only — do not attach a printer directly to the HP AlphaServer SC system.

Click on the Finish button to return to the Custom Setup menu. Click on the Exit button to return to the operating system prompt.

For more information about the other configuration options, such as security, see the *HP Tru64 UNIX Installation Guide*.

You are now ready to install the latest operating system patch, as described in Section 5.1.6.

Set Up the Management Server

5.1.6 Install the Operating System Patch Software

Install the Tru64 UNIX patch software on the management server.

Note:

For HP AlphaServer SC Version 2.6 (UK2), you should load Tru64 UNIX Version 5.1B-3 (also known as Tru64 UNIX Version 5.1B Patch Kit 5) only.

The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:

<<http://www.itrc.hp.com/>>

or from your local HP support representative.

Download the operating system patch software kit, untar it, and then run `dupatch` in the `patch_kit` directory.

If you have a clustered management server, you are now ready to install and configure the clustered management server, as described in Section 5.1.7. If not, you are ready to configure the remote installation services (RIS) server, as described in Section 5.1.8.

5.1.7 Install and Configure the Clustered Management Server

Note:

This section is optional. Only perform the steps in this section if you are installing a clustered management server. If you are not installing a clustered management server, you can configure the RIS server, as described in Section 5.1.8.

Note:

To install and configure a clustered management server, use only the standard TruCluster Server software. This software, which is supplied on the *Tru64 UNIX Associated Products Volume 2 Version 5.1B CD-ROM*, operates over a Gigabit Ethernet Interconnect.

Never install the HP AlphaServer SC-specific TruCluster Server software (which is supplied on the HP AlphaServer SC System Software CD-ROM) on a management server. Occasionally, patches may be provided for HP

AlphaServer SC cluster software - do not apply such patches to a management server, unless specifically asked to do so. Always strictly adhere to all instructions supplied with patch kits.

To install and configure a clustered management server, perform the following tasks:

- Register the TruCluster Server Software License (see Section 5.1.7.1 on page 5–23)
- Load the TruCluster Server Subsets (see Section 5.1.7.2 on page 5–24)
- Install the TruCluster Server Patch Software (see Section 5.1.7.3 on page 5–24)
- Configure the EVA Storage System Disks for the Management Cluster (see Section 5.1.7.4 on page 5–25)
- Run the `clu_create` Command (see Section 5.1.7.5 on page 5–25)
- Back Up Important Configuration Files on Member 1 (see Section 5.1.7.6 on page 5–27)
- Prevent Other Members Booting During Installation (see Section 5.1.7.7 on page 5–27)
- Run the `clu_add_member` Command (see Section 5.1.7.8 on page 5–28)
- Boot the New Member (see Section 5.1.7.9 on page 5–29)
- Back Up Important Configuration Files on New Members (see Section 5.1.7.10 on page 5–30)
- Shut Down the Non-Lead Member of the Clustered Management Server (see Section 5.1.7.11 on page 5–31)

5.1.7.1 Register the TruCluster Server Software License

The TruCluster Server kit includes a license Product Authorization Key (PAK). Use this PAK when registering a TruCluster Server license. If you do not have a PAK, contact your HP Customer Support representative.

For information on installing a license PAK, refer to the *Tru64 UNIX Software License Management* manual.

Set Up the Management Server

5.1.7.2 Load the TruCluster Server Subsets

Note:

For a clustered management server, use only the standard TruCluster Server software (which is supplied on the *Tru64 UNIX Associated Products Volume 2 Version 5.1B CD-ROM*) operating over a Gigabit Ethernet Interconnect. **Never install the AlphaServer SC-specific TruCluster Server software (which is supplied on the HP AlphaServer SC System Software CD-ROM) on a management server.**

To load the TruCluster Server kit, perform the following steps:

1. Insert the *TruCluster Server Software Version 5.1B CD-ROM* in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 5.1.9, step 2 on page 5–32).
3. Install the TruCluster Server software, as follows:

```
atlasms# cd /cdrom/TruCluster/kit
atlasms# setid -l .
```

Select the option to install all mandatory and all optional subsets.

After you select an option, the installation procedure verifies that there is sufficient file system space before copying the subsets onto your system.

Note:

Patch kits include fixes for both the base operating system and for cluster software. If you are installing a patch kit, patch the system after loading the TruCluster Server subsets but before running `clu_create` to create a single-member cluster.

5.1.7.3 Install the TruCluster Server Patch Software

In Section 5.1.6, the Tru64 UNIX patch software was installed on the management server. It is now necessary to install the matching TruCluster Server patch software on the management server.

Note:

For HP AlphaServer SC Version 2.6 (UK2), you should load Tru64 UNIX Version 5.1B-3 (also known as Tru64 UNIX Version 5.1B Patch Kit 5) only.

The operating system patch software kit (`T64V51BB26AS0005-20050502.tar`) is available from the following location:

<<http://www.itrc.hp.com/>>

or from your local HP support representative.

When following the system patch instructions below, please select only the TruCluster Server patches. At the end of the patch process, you will be informed that the system needs to be rebooted. You should select "Do nothing at this time". You will also be asked if you wish to overwrite the existing pre-patched kernel. You should select "y" and allow the overwrite operation. The clu_create phase will handle further required operations.

5.1.7.4 Configure the EVA Storage System Disks for the Management Cluster

In the example below, the aim is to be able to understand which UNIX special device corresponds to each shared-storage target disk. The screen shot is taken from an HSV110 installation but closely resembles an HSG80 installation.

To locate the target disks for the CFS disk, the quorum disk, and the boot disks, enter the following command:

```
atlasms0# hwmgr -view devices
```

HWID: Device Name	Mfg	Model	Location
4: /dev/dmapi/dmapi			
5: /dev/scp_scsi			
6: /dev/kevm			
42: /dev/disk/floppy0c		3.5in floppy	fdi0-unit-0
54: /dev/disk/cdrom0c	DEC	RRD46 (C) DEC	bus-0-targ-5-lun-0
55: /dev/disk/dsk0c	COMPAQ	HSV110 (C) COMPAQ	IDENTIFIER=1
56: /dev/disk/dsk1c	COMPAQ	HSV110 (C) COMPAQ	IDENTIFIER=2
57: /dev/disk/dsk2c	COMPAQ	HSV110 (C) COMPAQ	IDENTIFIER=3
58: /dev/disk/dsk3c	COMPAQ	HSV110 (C) COMPAQ	IDENTIFIER=4
59: /dev/disk/dsk4c	COMPAQ	HSV110 (C) COMPAQ	IDENTIFIER=5
60: /dev/cport/scp0		HSV110 (C) COMPAQ	bus-3-targ-0-lun-0

To correctly assign disks, check the disk identifiers against the OS_UNIT_ID/IDENTIFIER parameters used in the EVA/EMA configuration script (Management Cluster section). For the example in Create Virtual Disk Units for the Clustered Management Server (see Section 3.15.2.3 on page 3–52), assign disks as follows:

```
IDENTIFIER=1 CFS disk
IDENTIFIER=2 quorum disk
IDENTIFIER=3 Tru64 disk
IDENTIFIER=4 node 0 boot disk
IDENTIFIER=5 node 1 boot disk
```

5.1.7.5 Run the clu_create Command

To create the first member of the cluster from the Tru64 UNIX system, perform the following steps:

Set Up the Management Server

1. For shared Fibre Channel storage, please ensure that steps were already taken in Section 5.1.3 to make the various shared cluster disks bootable.
2. If there are NFS filesystems mounted by the management server, where the NFS server is currently down, please take action to unmount such filesystems now.
3. Run the following command:
`atlasms# /usr/sbin/clu_create`
The `/usr/sbin/clu_create` command prompts for the information needed to create a single-member cluster. Respond using the guidelines below ensuring that the cluster name is unqualified.

Sample answers to `clu_create` command:

Cluster name:	atlasms
Cluster alias IP Address:	16.209.133.169
Clusterwide root partition:	dsk0b
Clusterwide usr partition:	dsk0g
Clusterwide var partition:	dsk0h
Clusterwide il8n partition:	Not-Applicable
Quorum disk device:	dsk1
Number of votes assigned to the quorum disk:	1
First member's member ID:	1
Number of votes assigned to this member:	1
First member's boot disk:	dsk3
First member's virtual cluster interconnect device name:	ics0
First member's virtual cluster interconnect IP name:	atlasms0-ics0
First member's virtual cluster interconnect IP address:	10.32.0.1
First member's physical cluster interconnect devices	alt0
First member's NetRAIN device name	Not-Applicable
First member's physical cluster interconnect IP address	10.33.0.1

Once all the questions are answered, the `clu_create` process will perform the following tasks:

- Sets up the clusterwide `root (/)`, `/usr` and `/var` file systems, and the first member's boot disk.
- Configures a quorum disk (optional).
- Builds a kernel with cluster components.
- If the kernel build succeeds, `clu_create` copies the new kernel to the first member's boot disk. If the kernel build fails, `clu_create` displays warning messages but continues creating this first member. (You can boot the cluster `genvmunix` from the boot disk and attempt to build a kernel on the single-member cluster.)
- Sets boot-related console variables: `bootdef_dev` and `boot_reset`
- Creates and sets console variable `boot_dev`

The `clu_create` command writes a log file of the installation to `/cluster/admin/clu_create.log`. This log file contains all installation prompts, responses, and messages and should be examined for errors if problems arise during the `clu_create` phase.

4. When `clu_create` completes the kernel build phase, it will prompt to reboot now. You should answer **Yes**. The management server will now reboot as a single node cluster.

5.1.7.6 Back Up Important Configuration Files on Member 1

Because cluster members rely on the information in the following files, we recommend that, after booting the first member of the cluster, you make on-disk copies of these files in case of inadvertent modification. For member-specific files, the following examples assume that the member ID of the first member is 1 (`memberid=1`):

- `/etc/sysconfigtab.cluster`
`cp /etc/sysconfigtab.cluster /etc/sysconfigtab.cluster.sav`
- `/etc/rc.config.common`
`cp /etc/rc.config.common /etc/rc.config.common.sav`
- `/etc/sysconfigtab`

This file is a CDSL whose target is:

```
../cluster/members/{memb}/boot_partition/etc/sysconfigtab
```

To make a backup copy, change directory to the first member's `boot_partition/etc` directory and make a copy of its `sysconfigtab` file.

For example:

```
# cd /cluster/members/member1/boot_partition/etc
# cp sysconfigtab sysconfigtab.sav
```

- `/etc/rc.config`

This file is a CDSL whose target is:

```
../cluster/members/{memb}/etc/rc.config
```

To make a backup copy, change directory to the first member's `etc` directory and make a copy of its `rc.config` file. For example:

```
# cd /cluster/members/member1/etc
# cp rc.config rc.config.sav
```

In addition, we recommend that you perform a full backup of the single-member cluster.

5.1.7.7 Prevent Other Members Booting During Installation

Halt or turn off the system that will become the new member. If halting the system:

1. Set the `auto_action` console variable to `halt`

```
>>> set auto_action halt
```

Set Up the Management Server

2. Set the `bootdef_dev` console variable to an empty string:

```
>>> set bootdef_dev ""
```

The reason for these precautions is to make sure that the system cannot boot from the disk that `clu_add_member` will configure as the new member's boot disk.

5.1.7.8 Run the `clu_add_member` Command

New members can now be added to the cluster using the following command:

```
atlasms# /usr/sbin/clu_add_member
```

The `/usr/sbin/clu_add_member` command prompts for the information needed to add new members. Respond using the guidelines below ensuring that the member name is unqualified.

Sample answers to `clu_add_member` command

Member's hostname:	atlasms1
Member's ID:	2
Number of votes assigned to this member:	1
Member's boot disk:	dsk4
Member's virtual cluster interconnect devices:	ics0
Member's virtual cluster interconnect IP name:	atlasms1-ics0
Member's virtual cluster interconnect IP address:	10.32.0.2
Member's physical cluster interconnect devices:	alt0
Member's NetRAIN device name:	Not-Applicable
Member's physical cluster interconnect IP address:	10.33.0.2
Member's cluster license:	Not Entered

The `clu_add_member` command configures the new member's boot disk, adds and modifies files in the clusterwide file systems, and gives you the option of loading the TruCluster Server license PAK.

Note:

You can boot a system that does not have a TruCluster Server license. The system joins the cluster and boots to multi-user mode, but only root can log in (with a maximum of two users). The cluster application availability (CAA) daemon, `caad` is not started. The system displays a license error message reminding you to load the license. This policy enforces license checks while making it possible to boot and repair a system during an emergency.

Load only the TruCluster Server license at this time. Do not load the Tru64 UNIX license PAK. (You will load the Tru64 UNIX license PAK and any other license PAKs you need after you boot the new member for the first time.)

The `clu_add_member` command writes a log file of the installation to `/cluster/admin/clu_add_member.log`. If problems arise, please examine this log file for errors before continuing.

5.1.7.9 Boot the New Member

After running `clu_add_member` go to the console of the newly installed member and perform the following steps:

1. At the console of the new member, set the console variable `boot_osflags` to `A` so the system will boot to multi-user mode:

```
>>> set boot_osflags A
```
2. At the console of the new member, boot `genvmunix` from the new member's boot disk:

```
>>> boot -file genvmunix new_member_boot_disk
```

Remember to specify the correct SRM device name for the boot disk at the console as opposed to the Tru64 UNIX special file name you supplied to `clu_add_member`. For shared Fibre Channel storage, refer to the steps taken in Section 5.1.4 where the various shared cluster disks were made bootable. For example, if the console device name for the new member's boot disk is `dka300`:

```
>>> boot -file genvmunix dka300
```

3. During its first boot, the new member automatically performs the following tasks:
 - a. Configures all loaded subsets.
 - b. Attempts to build a customized kernel.
 - If the kernel build succeeds, copies the new kernel to the member's boot partition.
 - If the build does not succeed, when the system reaches multi-user mode, you can run `doconfig` to build a kernel.

Copy (`cp`) the new kernel from `/sys/HOSTNAME/vmunix` to `/vmunix` (If you move (`mv`) the kernel to `/vmunix` you will overwrite the `/vmunix CDSL`.)

4. When prompted to configure network interfaces, please select `configure` and follow the dialog to configure both the management network interface (`ee0`) and also the external network interface (`ee1`). The interfaces should be configured as shown in Table 5–8, where `atlas` is an example system name, and `x.x.x.x` and `y.y.y.y` are site-dependent values.

Table 5–8 Configuration of New Member Network Interfaces

Network Type	Device Name	Host Name	IP Address	Network Mask
Management	ee0	atlasms1	10.128.101.2	255.255.0.0
External	ee1	atlasms1-ext1	x.x.x.x	y.y.y.y

The system will continue to set the `boot_reset` and `bootdef_dev` variables, and creates and sets the `boot_dev` console variable. The system will then boot to multi-user mode.

Set Up the Management Server

5. When the system finishes booting to multi-user mode, register the Tru64 UNIX license and any other required application licenses. If you did not register the TruCluster Server license while running `clu_add_member`, you must register it now.
6. Because this member is still running `genvmunix`, reboot the system so it is using its custom kernel.

```
# shutdown -r now
```

This reboot is a mandatory step when adding a member to the cluster.

During the first boot, the system runs the `clu_check_config` command to examine the configuration of several important cluster subsystems. Look at the `clu_check_config` log files in the `/cluster/admin` directory to verify that these subsystems are configured properly and operating correctly.

You can then run the command in verbose mode to display more information about why a subsystem failed the initial test. See the *Tru64 UNIX System Administration Guide* and the *TruCluster Server Cluster Administration Guide* for information on configuring subsystems.

5.1.7.10 Back Up Important Configuration Files on New Members

Because cluster members rely on the information in the following files, we recommend that, after booting each additional member of the cluster, you make on-disk copies of these files in case of inadvertent modification.

For member-specific files, the examples use member 2 (`memberid=2`). Substitute the correct member ID for your new member when making backup copies of files. For example:

- `/etc/sysconfigtab.cluster`

```
# cp /etc/sysconfigtab.cluster /etc/sysconfigtab.cluster.sav
```

Because quorum vote information is updated each time a member is added, make a backup copy of this file after adding each member.

- `/etc/sysconfigtab`

This file is a CDSL whose target is:

```
../cluster/members/{memb}/boot_partition/etc/sysconfigtab
```

To make a backup copy, change directory to the new member's `boot_partition/etc` directory and make a copy of its `sysconfigtab` file.

For example:

```
# cd /cluster/members/member2/boot_partition/etc
# cp sysconfigtab sysconfigtab.sav
```

- `/etc/rc.config`

This file is a CDSL whose target is:

```
../cluster/members/{memb}/etc/rc.config
```

To make a backup copy, change directory to the new member's `etc` directory and make a copy of its `rc.config` file. For example:

```
# cd /cluster/members/member2/etc
# cp rc.config rc.config.sav
```

5.1.7.11 Shut Down the Non-Lead Member of the Clustered Management Server

The non-lead member of the clustered management server should be shut down, so that the RIS CAA service can be guaranteed to reside on member1, while the `sra install` command installs all of the HP AlphaServer SC software on the specified nodes as described in Section 7.4 on page 7–23. For information about booting the non-lead member of the clustered management server, see Section 8.1 on page 8–2.

5.1.8 Configure the RIS Server

Installation of the Tru64 UNIX operating system on the lead member and the remaining nodes of each domain will be performed using a Remote Installation Services (RIS) server. The RIS server allows each new member to boot a network `vmunix`.

The RIS server is also necessary for performing particular diagnostics during the installation phase. It is important to configure the RIS server at this stage and before you actually install the HP AlphaServer SC software as described in Section 5.1.9.

Configure the management server as a RIS Server, as follows:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 5.1.9, step 2 on page 5–32).
3. Run the `ris` command as the `root` user, as follows:
`atlasms# ris`
4. Choose the `Install software products` option by entering `i` at the prompt:
Enter your choice: `i`
5. The RIS Installation menu displays the installation options. Choose option 1, the `Install software into a new area` option.
6. Enter the full pathname for the distribution media, as follows:
Enter the device special file name or the path of the directory where the software is located
(for example, `/mnt/ALPHA/BASE`): `/cdrom/ALPHA/BASE`
7. Choose the standard boot method.
8. Choose to extract the software from the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM*.

Set Up the Management Server

Note:

If you choose to link to the CD-ROM instead of extracting the software, the installation process is considerably slower.

9. Choose to install all mandatory and all optional subsets.
10. Enter `y` to confirm that the subset list is correct. The subset extraction process begins.

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to install the HP AlphaServer SC installation utility (SRA) software, as described in Section 5.1.9.

5.1.9 Install the HP AlphaServer SC System Software

To install the system software on the management server, perform the following steps as the `root` user:

1. Insert the *HP AlphaServer SC System Software* CD-ROM in the disk drive.
2. Mount the CD-ROM as the `root` user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```
 - b. Mount the CD-ROM¹ as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/kits/
```
4. Install the HP AlphaServer SC software. From the kits directory above, run the following command:

```
atlasms# ./InstallSC -install -ms
```

1. See the *HP Tru64 UNIX Installation Guide* for more information on how to identify the CD-ROM device name. The CD-ROM device name in this example is `/dev/disk/cdrom0c`.

Note:

On clustered management servers, there may be some warnings at the end of the `InstallSC` sequence regarding `srad` and `rmshost`. These warnings should be ignored.

You are now ready to install the HP Fortran Run-Time Libraries, as described in Section 5.1.10.

5.1.10 Install the HP Fortran Run-Time Libraries

To install the HP Fortran Run-Time Libraries, perform the following steps:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1 CD-ROM* in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 5.1.9, step 2 on page 5–32).
3. Install the HP Fortran Run-Time Library, as follows:

```
atlasms# cd /cdrom/DEC_Fortran_RTL/kit
atlasms# setld -l .
```

You are now ready to install layered products, as described in Section 5.1.11.

5.1.11 Install Layered Products (Optional)

This step is optional. If you wish to install layered products (for example, CXML Compaq Extended Math Library), do so at this point, if possible.

If the product has a license (PAK), install the license now.

You are now ready to install the SANworks Storage System Scripting Utility, as described in Section 5.1.12.

5.1.12 Install the SANworks Storage System Scripting Utility

The SANworks Storage System Scripting Utility (SSSU) is required so that the SC Monitor system can monitor HP SANworks Management Appliance and HSV110 RAID System devices. If your system does not have a SANworks Management Appliance or HSV110 RAID System, you can skip this step.

To install the SANworks Storage System Scripting Utility, follow these steps:

1. Order the HP SANworks Tru64 UNIX Kit for Enterprise Virtual Array.
2. This kit contains a CD-ROM. Mount the CD-ROM as the `root` user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```

Set Up the Management Server

- b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom
```

4. Install the software, as follows:

```
atlasms# setld -l .
```

```
*** Enter subset selections ***
```

The following subsets are mandatory and will be installed automatically unless you choose to exit without installing any subsets:

```
* Enterprise v2 For Tru64 Unix
```

You may choose one of the following options:

1) ALL of the above

2) CANCEL selections and redisplay menus

3) EXIT without installing any subsets

Estimated free diskspace(MB) in root:646.3 usr:

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 1

You are installing the following mandatory subsets:

Enterprise v2 For Tru64 Unix

You are installing the following optional subsets:

Estimated free diskspace(MB) in root:646.3 usr:1716.8

Is this correct? (y/n):y

5. When the `setld` command has completed, create a link to `/usr/bin/sss`, as follows:

```
atlasms# ln -fs /usr/opt/ENTP002/sbin/sss /usr/bin/sss
```

This step is important, because the monitoring scripts expect the `sss` binary to be located in the `/usr/bin` directory.

You are now ready to define the RMS master node (`rmshost`), as described in Section 5.1.13.

5.1.13 Define the RMS Master Node (`rmshost`)

Define the management server to be the RMS master node; that is, the `rmshost` system.

Note:

The `rmshost` is the cluster alias of the management server, in cases where you have a clustered management server.

To do this, update the `/etc/hosts` file on the management server, to define `rmshost` as a host alias for the management server, as shown in the following example:

```
10.128.101.1    atlasms    rmshost
```

On a clustered management server, the following steps are necessary and mandatory. On a standalone management server, these CAA steps do not apply.

1. Register the `SC05msql` resource profile with CAA as follows:
`/usr/sbin/caa_register SC05msql`
2. Start the CAA applications, as follows:
`/usr/sbin/caa_start SC05msql`

For further information about configuring these CAA applications, refer to Section 8.7.

You are now ready to build the new kernel and reboot, as described in Section 5.1.14.

5.1.14 Build the New Kernel and Reboot

To build the kernel on a management server, perform the following steps:

Note:

If you do not have a clustered management server, proceed directly to step 1.

Note:

On clustered management servers, use the following command sequence instead of the instructions in steps 2 to 5 below:

```
atlasms# BuildKernels  
atlasms# DeployKernels  
atlasms# shutdown -ch now
```

Once the nodes are shut down, boot them again from the consoles.

Note:

Do not use the `mv` command to copy the kernels. In particular, do not use the `mv` command to copy the kernels on a clustered management server. On a non-clustered management server, you can safely move the new `vmunix` file to `/vmunix`.

1. Save a copy of the existing `/vmunix` file. If possible, save the file in the root (`/`) directory, as follows:
atlasms# `cp /vmunix /vmunix.save`

Set Up the Management Server

If there are disk space constraints, you can save the kernel file in a file system other than root. For example:

```
atlasms# cp /vmunix /usr/vmunix.save
```

2. Run the `/usr/sbin/doconfig` program specifying the name of the target configuration file with the `-c` option. For example, on a system named `atlasms`, enter the following command:

```
atlasms# /usr/sbin/doconfig -c ATLASMS
*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***
Saving /sys/conf/ATLASMS as /sys/conf/ATLASMS.bck
```

3. You are prompted to indicate whether or not you want to edit the configuration file:
Do you want to edit the configuration file? (y/n) [n]:
Accept the default to indicate that you do not want to edit the configuration file; the `/usr/sbin/doconfig` program builds a new kernel.
4. When the kernel configuration and build are completed without errors, copy the new `vmunix` file to `/vmunix`. On a system named `atlasms`, enter the following command:

```
atlasms# cp /sys/ATLASMS/vmunix /vmunix
```

Note:

On a management server, you can safely move the new `vmunix` file to `/vmunix`.

However, in a domain, you must always copy the new `vmunix` file to `/vmunix`. This is because, in an HP AlphaServer SC system, `/vmunix` is a context-dependent symbolic link (CDSL):

```
/vmunix -> cluster/members/{memb}/boot_partition/vmunix
```

You should treat a CDSL as you would any other symbolic link: remember that copying a file follows the link, but moving a file replaces the link. If you were to move (instead of copy) a kernel to `/vmunix`, you would replace the symbolic link with the actual file.

It is good practice to always copy the new `vmunix` file to `/vmunix`, even on a management server.

5. Reboot the system as follows:

```
atlasms# /usr/sbin/shutdown -r now
```

You are now ready to run the `sra setup` command, as described in Section 5.2.

5.2 Set Up the SC Database

Note:

The SC database is an SQL-based database and holds the information that was previously stored in the SRA database (the flat file `/var/sra/sra-database.dat`) and the RMS database (also a SQL-based database).

Set up the SC database on the management server, by running the following commands as the root user:

1. If you have a second rail in each HP AlphaServer ES40 system, but the rail is either currently unconnected or powered off, then you must first remove the second (expansion) elan card from the nodes before performing the `sra install` step. When prompted during the `sra setup` dialog, please answer that the system will have one rail. Later, once the domains are built and all members added, the second rail can be added as a post-installation step by following the instructions in Section 8.13. Please refer to Section 3.5 for details on PCI slot selection rules.
2. Plan the domain attributes (see step 5) and record them in Appendix D.
3. Ensure that all nodes are halted; that is, that they are powered up at the SRM console prompt.

Note:

For some system installations that utilize CAA for the RMS service, there is a possibility that `rmsbuild` will report a warning during `sra setup`. If this warning arises, it can be ignored and the user should answer "yes" when prompted. The warning is being ignored as it will automatically be handled/fixed later by the CAA startup script for RMS.

4. Run the `sra setup` command to create and populate the SC database, as follows:
`atlasms# sra setup`

This information is needed by other `sra` commands.

5. The `sra setup` command prompts you for the following information:
 - System name¹
 - Number of nodes
 - Number associated with first domain and node²

Set Up the SC Database

- Management Server or not?
- Management Server Name¹
- Hardware type of the system
- Number of Cluster Interconnect Rails used in the system
- Number of domains
- Management Network IP address for Node 0
- Cluster Interconnect IP address for Node 0²
- System Interconnect IP address for Node 0³
- Terminal server model
- Number of ports on the terminal server
(if you do not accept the default terminal server model)
- IP address of first terminal server
- First port on first terminal server
- Interconnect Switch IP address at Node Level and Top Level for each Rail
- IP address for the Preferred Server cluster alias base address

-
1. The HP AlphaServer SC installation process uses the system name (which cannot end with a digit) to derive both the cluster alias names and the host names of each member in the domain. For example, in a 64-node system with system name `atlas`, the cluster aliases will be `atlasD0` and `atlasD1`, while the node host names will be `atlas0`, `atlas1`, ..., `atlas63`.

Note that if the number of nodes is less than or equal to 32, the cluster alias name will be the same as the system name — in the above example, the cluster alias name would be `atlas`. If you later increase the number of nodes to 33 or more, you must create a second domain (`atlasD1`), and rename the original cluster alias (from `atlas` to `atlasD0`).

2. Domain and node names in a system typically start at 0, for example `atlasD0`, `atlas0`. However, it is possible to start at a different number, for example `atlasD32`, `atlas1024`.
1. For a clustered management server, the management server name is the cluster alias for the clustered management server.
2. The Cluster Interconnect network is an IP network provided by TruCluster Server. This network only spans a domain. This network is an artifact of the TruCluster Server software and is not intended for end-user use. This network is denoted by the `ics` suffix.
3. The System Interconnect network is layered on the HP AlphaServer SC Interconnect, and spans all nodes connected to the HP AlphaServer SC Interconnect. This network can be configured with externally routable IP addresses (instead of the default 10.* network). For example, this allows you to perform high-speed FTP by routing through high-performance interfaces on externally connected nodes. Note that such use will impact bandwidth availability to parallel jobs. This network is denoted by the `eip` suffix.

6. The `sra setup` command displays these settings and asks you to confirm that they are correct. If any settings are incorrect, enter **No** and repeat step 5. When all settings are correct, enter **Yes**.
7. The `sra setup` command prompts you whether you would like to use optional automatic cluster configuration.

Note:

If you have allocated external node IP addresses and cluster alias IP addresses according to the scheme mentioned in Section 2.2, and if all clusters in the system (except possibly the first and last) are to be composed of the same number of nodes, enter **Yes** when the `sra setup` command asks you whether to use automatic cluster configuration. If you do, `sra setup` will prompt you for the cluster sizes and two base IP addresses. `sra setup` will check the IP addresses, and then automatically enter configuration information for each cluster into the database, saving you from having to answer several questions for each cluster

8. The `sra setup` command prompts you for the following information for each domain:
 - Number of Nodes in the First Cluster
 - Number of Node in Subsequent Clusters
 - External network IP address for First Node
 - IP address for cluster alias

The `sra setup` command automatically assigns the IP addresses and prompts you to add the nodes to the database.

9. The `sra setup` command asks if you would like to update the `/etc/hosts` file. Press Return to accept the default answer of **Yes**. The `sra setup` command adds an entry to the `/etc/hosts` file for each IP address specified in steps 5 and 8 above.
10. The `sra setup` command asks which host should run console monitor (the `cmf` utility). Press Return to accept the default value (the current node). The `sra setup` command starts `cmf`, and informs you of the location of the `cmf` log file.

If you require `cmf` to operate as a CAA application, you should wait and do so as described in Section 8.8.
11. The `sra setup` command asks if you would like to configure the terminal server(s). Press Return to accept the default answer of **Yes**; the `sra setup` command configures the terminal server ports.
12. The `sra setup` command asks if you would like to configure an alternate boot device. Enter **yes**. On SC20 systems, enter **no** for alternate boot device as there is no alternate boot device on these systems.

Set Up the SC Database

- 13. The `sra setup` command asks if you would like to use an alternate boot device. Enter **yes**.
- 14. The `sra setup` command prompts you for the information needed to configure the boot device(s). You can accept the default values for all settings except the SRM device name. To enter a value for the SRM device name, perform the following steps:
 - a. Enter the SRM device name for the primary boot disk at the `SRM device name for primary boot disk?` prompt. If you do not know the SRM device name, enter **probe**.
 - b. If prompted, enter the host name of a generic node that is currently at the SRM prompt. This is only necessary if you entered `probe` in step 14a. above. If you enter the host name of a node that is not currently at the SRM prompt, the hardware probe will fail.
 - c. The `sra setup` command displays the settings for the primary boot disk, including the SRM device name. To accept these settings, enter **y**. To change any setting, enter **n**, enter the item number of the setting to be changed, and enter the new value.

You can set the alternate boot device in a similar way.

Table 5–9 shows the recommended boot disk partition configuration when the memory size is 4GB and each boot disk is a 36GB disk.

Table 5–9 Recommended Boot Disk Partition Configuration

Partition	Description	Recommended Size (% and GB)	
		Boot Disk Only	Boot Disk and Alternate Boot
b	swap	30% = 10.8GB	15% = 5.4GB on each disk
d	/local	35% = 12.6GB	43% = 15.2GB on each disk
e	/tmp	35% = 12.6GB	42% = 15.4GB on each disk

For a 36GB disk, the swap space is calculated as 2.5 times the physical memory size. However, for an 18GB disk, the swap space is calculated as 1.25 times the physical memory size, which you may consider to be too low. Think carefully before accepting the default partition values.

Note:

If you configure an alternate boot disk and have chosen to use the alternate boot disk, its swap space is added to the `sysconfigtab` file; that is, you spread the swap space across the two disks. Also, the `tmp` and `local` partitions on the alternate boot disk are mounted on `/tmp1` and `/local1` respectively. You should take this into account when deciding the partition percentages, as follows:

- If configuring only one boot disk, use the values specified in the third column of Table 5–9.

- If configuring a boot disk and an alternate boot disk, use the values specified in the fourth column of Table 5–9.
-
15. The `sra setup` command asks for the information needed to configure the primary boot disk.
Press Return to accept all the default answers. The SRM device name is not required to complete the installation.
 16. The `sra setup` command asks for the information needed to configure the `Cluster /usr` and `/var` disk.
Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation.

The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit` (see Section 7.2).
 17. The `sra setup` command asks for the information needed to configure the backup cluster disk.
Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation. The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit`.
 18. The `sra setup` command asks for the information needed to configure the Generic boot disk.
Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation.

The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit`.
 19. The `sra setup` command prompts you for the SRM device name and the UNIX device name for the management and external LAN adapters. You can accept the default value, or enter a new value, or use **probe** as described in step 14.

Note:

For an external LAN adapter, if the external LAN device name is unknown, you can accept the default SRM device name (for example, `eia0`) but you must enter the appropriate UNIX device name. Page xxix lists the SRM and UNIX device names for the supported network adapters.

Set Up the SC Database

For HP AlphaServer DS20L systems, the SRM and UNIX names of the LAN devices are as follows:

	SRM	UNIX
management lan	eib0	ee1
external lan	eia0	ee0

20. The `sra setup` command asks if you would like to probe the nodes for the hardware ethernet (MAC) address.

Note:

The `probe` command will return the MAC address of the first network interface only. The `sra edit` command may be used to set the MAC address manually (see Section C.1).

Enter **yes** — the `sra setup` command asks two additional questions:

- a. The script asks which nodes you would like to probe.
You can specify a single node (for example, `atlas3`), several nodes (for example, `atlas3,atlas5`), a range of nodes (for example, `atlas[1-8]`), or all nodes (`all`).
Enter **all**.
- b. The script also asks if you would like to initialize hardware.
This runs the SRM `init` command on each node. During the probe, SRM console variables are set, as detailed in Table 5–10.
Enter **yes**.

Table 5–10 SRM Console Variables

Variable	Value
<code>auto_action</code>	<code>halt</code>
<code>boot_reset</code>	<code>off</code>
<code>boot_osflags</code>	<code>A</code>
<code>os_type</code>	<code>UNIX</code>
<code>ocp_text</code>	<code>nodename</code>
<code>eia0_mode</code> or <code>eib0_mode</code>	(<code>eia0_mode</code> if the management network interface is the first network adapter on the node; otherwise, the device name could be <code>eib0_mode</code>) <code>FastFD</code>
<code>console</code>	<code>serial</code>
<code>sys_com1_rmc</code>	<code>off</code>

Table 5–10 SRM Console Variables

Variable	Value
¹ sysvar	12
¹ ac_action	STANDBY
¹ wol_enable	YES

¹For use in HP AlphaServer DS20L systems only.

If any nodes fail the hardware probe, a list of failed nodes is written to the `/tmp/sra_hosts.failed` file. A node may fail the hardware probe because it is not at the SRM prompt, or because `cmf` is misconfigured — see Section 11.8 on page 11–26 for more details.

The `sra setup` command asks if you would like to reset the failed nodes. Enter **Yes**.
The `sra setup` command asks if you would like to retry the hardware probe. Enter **Yes**.
The `sra setup` command asks if you would like to initialize the hardware during the probe. Enter **Yes**.

Note:

For the DS20L product, on the first node of each domain, the `sra setup` command automatically runs `init` and `wwidmgr` to set the path to the HSG disks. The `sra setup` command will take a little longer to run in this case.

- 21. The `sra setup` command asks if you would like to add nodes to the RIS database. Press Return to accept the default answer of **Yes**. The `sra setup` command configures each node as a RIS client. During the Configure RIS section of `sra setup`, you are asked if you would like to add domains to RIS. Press Return to accept the default answer of **Yes**.
- 22. The `sra setup` command will now prompt to add two root crontab entries (one crontab for archiving the `cmf` logs and one crontab for backing up the RMS database). Press Return to accept the default answer of **Yes** to both questions. The `cmf` archives will be archived every two weeks and the RMS database backup will take place nightly.
- 23. The `sra setup` command incrementally saves all information to the database, and then exits.

Note:

After the SC database has been created, the `sra setup` command completes the installation of the SRACFENGINE subset by running the `cfconfig` script on the management server.

Set Up the SC Database

The `cfconfig` command automatically runs on each cluster node when the node is first booted. The `cfconfig` command examines the local host to populate the appropriate files with default values. The `cfengine` daemon (`cfengine`) starts for the first time after the `cfconfig` command has run.

If DNS has not been configured on the system, the `cfconfig` command will print an error message when it cannot find `/etc/resolv.conf`. This is normal. `cfengine` cannot operate unless DNS has been configured. You can run `cfconfig` at a later stage once DNS is configured.

If you later wish to change any of the data in the SC database, use the `sra edit` command (see Chapter 16 of the *HP AlphaServer SC System Administration Guide*) — it is not necessary to rerun the `sra setup` command.

See Appendix E for an example `sra setup` log file.

Note:

Now that you have set up the SC database, you can use the `sra -c|cl|m|ml` command on the management server to connect to the console of any system in the cluster. See Chapter 16 of the *HP AlphaServer SC System Administration Guide* for more information about the `sra -c|cl|m|ml` command.

24. You will now need to start RMS on the management server. If you have a standalone management server, this step is not necessary, so you can proceed to Section 5.3.

On a clustered management server, the following steps are necessary and mandatory. On a standalone management server, these CAA steps do not apply.

- a. Register the `SC20rms` resource profile with CAA as follows:
`/usr/sbin/caa_register SC20rms`
- b. Start the CAA application, as follows:
`/usr/sbin/caa_start SC20rms`

For further information about configuring this CAA application, refer to Section 8.7.

You are now ready to configure out all nodes in the system, as described in Section 5.3.

5.3 Configure Out All Nodes During Installation

While the installation is progressing, it is necessary to configure out all nodes. If the nodes are left configured in, then there will be many unwanted events reporting the fact that the nodes are not responding. Configuring out the nodes will reduce the load on the `mSQL` daemon by the RMS `mmanager` and `eventmgr`. You should configure out the nodes as follows:

```
atlasms# rcontrol configure out nodes "atlas[0-1023]"
```

In Configure the RMS Database (see Section 8.6 on page 8–10) you will have the opportunity to create the desired partition configuration and to configure in the nodes again.

You are now ready to check all nodes in the system, as described in Section 5.4.

5.4 Check All Nodes in the HP AlphaServer SC System

Check all of the nodes in the system by performing the following steps:

- Check the State of the Nodes (see Section 5.4.1 on page 5–45)
- Check the System Firmware (see Section 5.4.2 on page 5–46)

5.4.1 Check the State of the Nodes

The `sra diag` command examines the HP AlphaServer ES40 node using various SRM and RMC commands, and gathers as much data as possible about the state of the specified node(s).

Run the `sra diag` command on all nodes, as follows:

```
atlasms# sra diag -nodes 'atlas[0-1023]' -analyze no
```

Review the `/var/sra/diag/nodename.sra_diag_report` files to check for errors or warnings.

You are now ready to check the system firmware, as described in Section 5.4.2.

Check All Nodes in the HP AlphaServer SC System

5.4.2 Check the System Firmware

Table 5–11 lists the minimum firmware versions supported by HP AlphaServer SC Version 2.6 (UK2) for an HP AlphaServer ES40, an HP AlphaServer ES45, and an HP AlphaServer DS20L system.

Table 5–11 Minimum System Firmware Versions

Firmware	HP AlphaServer ES40	HP AlphaServer ES45	HP AlphaServer DS20L
SRM Console	6.2-1	6.2-8	6.3-1
ARC Console	5.71	Not Displayed	Not Displayed
OpenVMS PALcode	1.96-103	1.96-39	1.90.71
Tru64 UNIX PALcode	1.90-104	1.90-30	1.86.68
Serial ROM	2.12-F	2.20-F	Not Displayed
RMC ROM	1.0	1.0	Not Displayed
RMC Flash ROM	2.5	1.9	Not Displayed

Check that all nodes in your system meet these minimum system firmware requirements, by performing the following tasks:

1. Run the `show config` command on all nodes and save the output to file, as follows:
`atlasms# sra command -nodes all -command 'show config' |tee /tmp/fw.txt`
(press Return when prompted for the root password)
2. View the SRM console firmware version:
`atlasms# grep -i 'SRM Console' /tmp/fw.txt`
3. View the ARC console firmware version:
`atlasms# grep -i 'ARC Console' /tmp/fw.txt`
4. View the UNIX PALcode firmware version:
`atlasms# grep -i 'PALcode' /tmp/fw.txt`
5. View the Serial ROM firmware version:
`atlasms# grep -i 'Serial Rom' /tmp/fw.txt`
6. View the RMC ROM firmware version:
`atlasms# grep -i 'RMC Rom' /tmp/fw.txt`
7. View the RMC Flash ROM firmware version:
`atlasms# grep -i 'RMC Flash Rom' /tmp/fw.txt`
8. Delete the output file:

Configure and Diagnose the HP AlphaServer SC Interconnect

```
atlasms# rm /tmp/fw.txt
```

If necessary, update the system firmware, as described in Chapter 21 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure and diagnose the HP AlphaServer SC Interconnect, as described in Section 5.5.

5.5 Configure and Diagnose the HP AlphaServer SC Interconnect

It is **vital** that these sections be performed at exactly this point in the installation.

If you neglect to perform these steps, then the remainder of the installation will have problems and there will be no way to diagnose the root cause. Rather than duplicate the text from the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*, this section contains a number of references to the relevant sections of that guide. Please be careful not to miss any steps as you refer to the sections. All steps are mandatory for successful operation of the diagnostic tools.

Note:

As you follow the steps in the instructions in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*, please be aware that the instructions within that manual assume that all nodes are already installed and booted. However, from the perspective of this point in this document, this is not actually the case, and instead only the management server is installed.

Examples are situations where you are prompted to restart daemons on all nodes, or where you are asked to start swmserver daemons on nodes 0/1 for directly-connected switches. As such, you should follow the steps in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual* in so far as is possible at this point in the installation sequence, and skip over the steps that obviously require further nodes to be installed and booted. Later, when the nodes in question are installed and booted, the diagnostics will operate correctly, and you can elect to run the diagnostics again as a confidence-building measure.

Configure and Diagnose the HP AlphaServer SC Interconnect

5.5.1 Upgrading the HP AlphaServer SC Interconnect Control Processor Software

Note:

This section does not apply to HP AlphaServer SC 16-port switches or old-type HP AlphaServer SC 128-way switches. Such switches are known as directly-connected switches, and do not have any firmware. This section only applies to the new-type HP AlphaServer SC 128-way switches containing an HP AlphaServer SC Interconnect control processor.

In HP AlphaServer SC Version 2.6 (UK2), the software (also known as firmware) running on the HP AlphaServer SC Interconnect Control Processors has changed, and must be upgraded. The new software image is packaged in the SWM subset, which is part of the HP AlphaServer SC Version 2.6 (UK2) System Software.

You should follow the steps in Section 7.2, Upgrading the HP AlphaServer SC Interconnect Control Processor Software and all subsections, of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

5.5.2 Creating an Interconnect Configuration Using SC Viewer

To receive accurate status and diagnostic information about the HP AlphaServer SC Interconnect, it is mandatory to create an Interconnect Configuration using SC Viewer. This will allow the switch managers to report anomalies in the network, and it will allow the command-line diagnostic tools to make pass/fail decisions based on the configured size of the network compared with the detected size of the network. Once configured, SC Viewer will also allow the user to view the interconnect configuration and events, which is very important in identifying problem locations.

You should follow Section 7.3, Creating an Interconnect Configuration Using SC Viewer and all subsections, of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

Note:

For systems with directly connected switches, you can simply skip over those instructions in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual* that assume that nodes 0/1 are already installed and booted (for example, as described in step 3 of Section 7.3.2.2, Switch Has an HP AlphaServer SC Interconnect Control Processor).

5.5.3 Confirming the Operation of the HP AlphaServer SC Interconnect

At this stage, the database and nodes will be in suitable condition to perform an initial sweep of diagnostics.

The following diagnostics should be performed:

- Assign a common clock source per rail (see Section 7.4.1 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Verify the clock frequency on each switch module (see Section 7.4.2 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `elanenvtest` to verify the switch modules (see Section 7.4.3 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*).

Note:

For directly-connected switches, please be aware that `elanenvtest` will only succeed once nodes 0/1 are installed and booted with their `swmsserver` daemons running.

- Show links in reset (see Section 7.4.4 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

Note:

For directly-connected switches, please be aware that the `show links in reset` examination will only succeed once nodes 0/1 are installed and booted with their `swmsserver` daemons running.

- Run `elanpcitest` on all nodes (see Section 7.4.5 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Run `elanlinktest` on all nodes (see Section 7.4.6 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Run `riscabletest` on each node-level switch (see Section 7.4.7 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

The remaining diagnostics will be performed once the system nodes are installed (see Chapter 8).

You are now ready to set up the SC Monitor System, as described in Section 5.6.

Set Up the SC Monitor System

5.6 Set Up the SC Monitor System

After completing the installation steps in Section 5.2, the SC Database is populated with default entries. As there are many existing configurations, it is possible that the database may not reflect entirely the system configuration. For example, it is possible that your system may have more terminal servers than are added by default, or more switches.

Some devices, like HSG and SAN Appliance, are not added automatically in the database as monitored devices at install time, so it is necessary to add these devices if you want to monitor the device status.

Setting Up the SC Monitor system involves the following tasks:

- Add a SAN Appliance as Monitored Devices (see Section 5.6.1 on page 5–50)
- Change Terminal Servers Monitoring Distribution (see Section 5.6.2 on page 5–50)
- Change Extreme Switches Monitoring Distribution (see Section 5.6.3 on page 5–51)
- Activate the Changes (see Section 5.6.4 on page 5–52)

5.6.1 Add a SAN Appliance as Monitored Devices

The `sra setup` command does not add any entry for SAN Appliance devices. To monitor this type of device, add the device manually as a monitored device by using the `scmonmgr add` command. The command syntax is as follows:

```
scmonmgr add -c appliance -o object -i ip-addr -s server [-r rack -u unit]
```

For example:

```
atlasms# scmonmgr add -c appliance -o sanapp0 -i 16.21.24.13 -s atlasms
```

This command adds the SAN Appliance with IP 16.21.24.13 as a monitored device with name `sanapp0`. The node that will perform the monitoring functions will be `atlasms`.

5.6.2 Change Terminal Servers Monitoring Distribution

Immediately after installation (when `sra setup` has completed), the terminal servers specified are added as monitored devices. All terminal servers are initially monitored by the management server, if there is a management server, or by the first domain otherwise.

To see the initial distribution for terminal server monitoring functions, run the following command:

```
# scmonmgr distribution -c tserver
```

Example (server `atlasms` monitors all 8 terminal servers from the system):

```
atlasms# scmonmgr distribution -c tserver
Class: tserver
Server atlasms monitors: atlas-tc[1-8]
```

atlasms#

To add additional terminal servers to be monitored by the system, use the `scmonmgr add` command.

Example:

```
atlasms# scmonmgr add -c tserver -o atlas-tc9 -s atlasms
```

Note:

A terminal server cannot be added if there is no entry for it in the `/etc/hosts` file. The system needs to know the IP address of the terminal server.

To change the distribution of the monitoring for terminal servers, use the `scmonmgr move` command as follows:

```
atlasms# scmonmgr move -o atlas-tc2 -s atlas10
```

After this `atlas10` will perform the monitoring for `atlas-tc2`.

5.6.3 Change Extreme Switches Monitoring Distribution

The `sra setup` process creates default entries for Extreme switches in the SC database using the following rules:

- If the number of nodes in the system is less than or equal to 16, a single Extreme switch of type Summit24 will be added as a monitored device.
- If the number of nodes is between 16 and 128 (including 128), the `setup` process adds three Summit48 switches for monitoring.
- If the number of nodes is greater than 128, the number of switches added for monitoring is calculated using the following expression: $\{number\ of\ nodes\ from\ system / 32\}$. So for a 255 nodes system, we will have $255/32 = 7$ Summit48 switches added by default as monitored devices.

This model was established based on the standard configuration of the systems, but in reality, it is possible that a system may not respect these rules and a post-install customization will be necessary.

If necessary, you can add/remove Extreme switches from the monitoring system by using `scmonmgr` command.

For example:

1. To add an Extreme switch as a monitored device

```
atlasms# scmonmgr add -c extreme -o extreme10 -i 10.128.103.10 -s atlasms
```

Assign Cabinets in the SC Database

2. To remove an Extreme switch as a monitored device

```
atlasms# scmonmgr remove -o extreme10
```

The monitoring of all Extreme switches is initially set to be performed by the management server if the system has one, or by the first domain otherwise. But this can be changed using the `scmonmgr move` command.

5.6.4 Activate the Changes

After you have finished all the tasks required to change/customize the monitoring and distribution of the monitoring for a system, it is necessary to activate these changes by sending a reload command to the monitoring daemons running on all nodes of the system.

Before doing this, you can analyze/display the distribution of the monitoring for the entire system by using the following command:

```
atlasms# scmonmgr distribution
```

If you are satisfied, send the `reload` command to the daemons on each management server and domain. The monitoring daemon has a reload mechanism that works on domain level. If you send a `reload` command to a daemon (node) from a domain, this will trigger the reload for all daemons from that domain.

If no domains have been installed, then you can skip the reload process for all domains.

However, if some domains are already installed, then you can start the reload process for all the daemons on all domains in the system by running the following command:

```
atlasms# scrun -d all '/sbin/init.d/scmon reload'
```

This will send the reload signal (`SIGHUP`) to one node from each domain, and it will trigger the reload procedure on all members from that domain.

After this, it is necessary to run the reload command on the management server:

```
atlasms# /sbin/init.d/scmon reload
```

For more information on SC Monitor management, refer to Chapter 27 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to assign cabinets in the SC Database, as described in Section 5.7.

5.7 Assign Cabinets in the SC Database

The *Physical* view in the SC Viewer needs to know about how your HP AlphaServer SC system is packaged, which cabinets hold nodes, terminal servers, and other pieces of hardware. The sysman `sc_cabinet` menu is used to load this information into the SC Database.

Note:

You can skip this step now and perform it later. However, until you complete this step, the `scviewer` program will be unable to display the physical relationships between hardware components.

Also, if you are already running `scviewer`, please close it down before proceeding, as the *Physical* view will not correctly reflect the cabinet changes you are making until `scviewer` is restarted.

5.7.1 Load the Physical Relationships

To load the physical relationships, you should follow these steps:

1. Create a description of cabinets in the SC Database.

Each cabinet (or rack) has the following data associated with it:

- Number. Each cabinet has a unique number (starting at 1).
 - Label. In this release, the label can be blank or the same as the number. In a later release, you can assign a meaningful label to each cabinet.
 - Area. This allows you to group cabinets into different rooms or different areas within a room. A small system would probably have one area—a large system might have several. The area concept is used by the `scviewer` when displaying cabinets, so it helps if an area is smaller than 20 rows of 30 cabinets (these numbers are guidelines, you can have larger numbers).
 - Position within area. The position of a cabinet within an area is specified by a row and column number. If you drew a plan of your cabinets (as `scviewer` does), row 1, column 1 is on the top left. Row numbers increase as you move downwards and column numbers increase as you move to the right.
2. You can use the `sysman sc_cabinet` menu to create the cabinets in an area as follows:
 - Start the `sysman sc_cabinet` menu (as root):
`atlasms# sysman sc_cabinet`
 - Select the Create block of cabinets... option.
 - Specify the Area Name, First cabinet number, the label style and the number of rows and columns you want.
 - Click on Next, and confirm your choices by clicking on OK. A block of cabinets in this area is now created in the SC Database.
 3. You can repeat these steps several times for different areas or for different blocks of cabinets.

Assign Cabinets in the SC Database

5.7.2 Populate Cabinets with Nodes

You can populate a block of cabinets with a range of nodes as follows:

1. Start the sysman *sc_cabinet* menu (as root):
atlasms# **sysman sc_cabinet**
2. Select the Assign nodes... option.
3. Enter the range of nodes that are to be placed into the cabinets.

Note:

The first node number should be 0 if the first node name is atlas0.

4. Enter the range of cabinet numbers where these nodes are located.
 - Nodes can be placed top to bottom (the first node is on top, last on bottom) or vice versa.
 - Nodes can be placed into the cabinet range first to last (the first node is placed into the first cabinet, the last node is placed into the last cabinet) or vice versa.
5. Click OK to confirm you choices. The SC Database is updated to reflect this node placement.
6. You can repeat these steps several times for different ranges of nodes.

5.7.3 Populate Cabinets with Other Hardware

You can populate cabinets with other hardware (terminal servers, and so on) as follows:

Note:

If you wish to modify the properties of one or more existing cabinets (for example, to contain additional hardware) select the cabinet, and click the Modify option, to edit the contents as described in Step 5.

1. Start the sysman *sc_cabinet* menu (as root):
atlasms# **sysman sc_cabinet**
2. Select the Cabinets... option.
This shows a menu where all existing nodes are listed.
3. Select Add.. to add a new cabinet.
4. Specify a unique number, label, area, row and column. Click OK to create the cabinet in the SC Database.

Assign Cabinets in the SC Database

5. The Edit a cabinet dialog box now appears. This allows you to place hardware components into the cabinet. A list on the right shows the Unassigned Objects—that is the hardware components that have no record of their location in the SC Database. The list does not show nodes—instead it shows other types of hardware components (terminal server, Extreme Switch, and so on).
6. To place a hardware component in the cabinet, select the component and click on the Move button. This places the component into the cabinet—so it appears in the Cabinet Contents list on the left.
7. You can reorder components in the cabinet by selecting the component and clicking on the Up and Down buttons.
8. To confirm your changes, click OK. The SC Database is updated.

The Cabinets... option also allows you to modify or delete existing cabinets.

When you have finished modifying the cabinet data in the SC Database, you can review the data using the Make report... option. This menu allows you to specify a file where a textual description of the data is saved. The data is saved in a format that allows subsequent machine processing rather than being human readable.

You are now ready to build the domains, as described in Chapter 7 (*Building the Domains*).

Installing: When the System Does Not Have a Management Server

This chapter describes how to install an HP AlphaServer SC system that does not have a management server.

Note:

For HP AlphaServer DS20L systems, you will have a management server, so please refer to **Chapter 5 (Installing: When the System Has a Management Server)** for installation instructions.

Note:

If your system has a management server, do not use this chapter — refer to Chapter 5 (Installing: When the System Has a Management Server) instead.

If you wish to add a management server to an HP AlphaServer SC system later — that is, after domain creation — use the checklist provided in Appendix C to ensure that you complete all installation tasks in the correct order.

When installing software on an HP AlphaServer SC system that does not have a management server, install the software on Node 0 first.

For information on helpful tips and guidelines that may assist you when performing an HP AlphaServer SC system installation, see Section 11.1.

The information in this chapter is structured as follows:

- Set Up Node 0 (see Section 6.1 on page 6–2)
- Set Up the SC Database (see Section 6.2 on page 6–24)
- Check All Nodes in the HP AlphaServer SC System, Except Node 0 (see Section 6.3 on page 6–31)

Set Up Node 0

- Configure and Diagnose the HP AlphaServer SC Interconnect (see Section 6.4 on page 6–33)
- Set Up the SC Monitor System (see Section 6.5 on page 6–35)
- Assign Cabinets in the SC Database (see Section 6.6 on page 6–38)
- Review the SC Database Disk Settings (see Section 6.7 on page 6–41)
- Transform Node 0 into a Single Node Domain (see Section 6.8 on page 6–43)
- Configure Out All Nodes During Installation (see Section 6.9 on page 6–44)
- Run the HP AlphaServer SC Interconnect Tests on Node 0 (see Section 6.10 on page 6–44)

Chapter 7 describes how to build domains.

Note:

Use the checklist provided in Appendix B, to ensure that you complete all installation tasks in the correct order.

6.1 Set Up Node 0

Setting up Node 0 involves the following tasks:

- Set the Console Variables (see Section 6.1.1 on page 6–3)
- Check the System Firmware (see Section 6.1.2 on page 6–4)
- Install the Tru64 UNIX Operating System (see Section 6.1.3 on page 6–4)
- Customize the System Configuration (see Section 6.1.4 on page 6–8)
- Install the Latest Operating System Patch Software (see Section 6.1.5 on page 6–18)
- Configure the RIS Server (see Section 6.1.6 on page 6–19)
- Install the HP AlphaServer SC System Software (see Section 6.1.7 on page 6–20)
- Install the HP Fortran Run-Time Libraries (see Section 6.1.8 on page 6–21)
- Install Layered Products (Optional) (see Section 6.1.9 on page 6–21)
- Install the SANworks Storage System Scripting Utility (see Section 6.1.10 on page 6–21)
- Add sysconfigtab Parameters (see Section 6.1.11 on page 6–22)
- Define the RMS Master Node (rms host) (see Section 6.1.12 on page 6–23)

6.1.1 Set the Console Variables

Before installing the Tru64 UNIX operating system on Node 0, you must configure the system console.

Display the SRM console prompt on Node 0 as follows:

- If your system has a factory-installed software (FIS) kernel, it will automatically start to boot the Tru64 UNIX operating system at power on. When prompted to continue this boot, enter **No** to return to the SRM console prompt.
- If your system does not have a FIS kernel, the system will display the SRM console prompt at power on.

To set the console variables, enter (at the SRM console prompt) the commands specified in Table 6–1.

Table 6–1 Setting the Console Variables

```
P00>>> set auto_action HALT

P00>>> set eia0_mode FastFD

P00>>> set eib0_mode1 FastFD2

P00>>> set boot_osflags A

P00>>> set boot_reset off

P00>>> set os_type UNIX

P00>>> set console serial

P00>>> set sys_com1_rmc off

P00>>> set ocp_text nodename

P00>>> set bootdef_dev ''

P00>>> set pci_parity on
```

¹You need only set the `eib0_mode` variable on nodes that have an external network interface. You should set the `eia0_mode` variable on all nodes, for the management network.

²When setting the `eib0_mode` variable, specify the appropriate network speed.

The remaining console variables for Node 0, and all console variables for the remaining nodes, are automatically set during the installation process (see Section 6.2, step 22 on page 6–30).

For more information about the SRM console, see Chapter 2 of the *HP AlphaServer ES40 Owner's Guide*.

Set Up Node 0

You are now ready to check the system firmware, as described in Section 6.1.2.

6.1.2 Check the System Firmware

Table 6–2 lists the minimum firmware versions supported by HP AlphaServer SC Version 2.6 (UK2) for an HP AlphaServer ES40 and an HP AlphaServer ES45 system.

Table 6–2 Minimum System Firmware Versions

Firmware	HP AlphaServer ES40	HP AlphaServer ES45
SRM Console	6.2-1	6.2-8
ARC Console	5.71	Not Displayed
OpenVMS PALcode	1.96-103	1.96-39
Tru64 UNIX PALcode	1.90-104	1.90-30
Serial ROM	2.12-F	2.20-F
RMC ROM	1.0	1.0
RMC Flash ROM	2.5	1.9

Check that Node 0 meets these minimum system firmware requirements, by entering the `show config` command at the SRM prompt.

If necessary, update the system firmware, as described in Chapter 21 of the *HP AlphaServer SC System Administration Guide*.

You will check the system firmware on the other nodes later (see Section 6.3.2 on page 6–31).

You are now ready to install the Tru64 UNIX operating system, as described in Section 6.1.3.

6.1.3 Install the Tru64 UNIX Operating System

Install the Tru64 UNIX operating system on Node 0, from the *Tru64 UNIX Version 5.1B Operating System Volume 1* CD-ROM, as described here and in the *HP Tru64 UNIX Installation Guide*.

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1* CD-ROM into the disk drive.
2. Initialize the hardware, as follows:

```
P00>>> init
```
3. Identify the CD-ROM device name by running the following SRM console command:

```
P00>>> show device
```

The *show device* command produces the following output:

dka0.0.0.1.1	DKA0	COMPAQ BD018122C9 B016
dka100.1.0.1.1	DKA100	COMPAQ BD018122C9 B016
dka200.2.0.1.1	DKA200	COMPAQ BD018122C9 B016
dqa0.0.0.15.0	DQA0	COMPAQ CDR-8435 0013
dva0.0.0.1000.0	DVA0	
eia0.0.0.2004.1	EIA0	00-50-8B-CF-46-CC
eib0.0.0.2005.1	EIB0	00-50-8B-CF-46-CD
pga0.0.0.3.1	PGA0	WWN-1000-0000-C922-391A
pka0.7.0.1.1	PKA0	SCSI Bus ID 7

In this example, the CD-ROM device name is *dqa0*.

4. Enter the following command at the SRM console prompt, to display the installation user interface:

```
P00>>> boot dqa0
```

The system boot process can take several minutes. Several hardware-specific messages are displayed. The more complex the system (many peripheral devices, and so on), the longer the boot process takes.

Upon successful system boot, the installation user interface appears. The type of user interface presented during installation depends on the hardware configuration:

- Systems equipped with graphics consoles present a graphical user interface.
- Systems with consoles that do not have graphics capabilities present a text-based, menu-driven user interface.

The information you supply is the same regardless of the type of user interface, but the order in which it is requested may be different. The following steps describe the steps when using the graphical user interface:

Follow these guidelines:

1. When prompted to select a language in which to view the user interface, choose *United States English*.
2. The *Installation Welcome* dialog box appears. Click on the *Next* button.
3. The *Host Information* dialog box appears. You must set several values on this screen, as follows:
 - a. Set the *host name*. The host name should be the same as the network interface for the management network. For example, if the system name is *atlas*, the host name should be *atlas0*.
 - b. Set the *area*.
 - c. Set the *location*.
 - d. Set the *date*.
 - e. Set the *time*.

Set Up Node 0

Click on the Next button.

4. The Set root Password dialog box appears. Set the root password for the system. The same root password is used for all nodes. Click on the Next button.
5. The Software Selection dialog box appears. You must set the values on this screen as follows:
 - a. Choose `All Software` when prompted for the software that you wish to install.
 - b. Click on the Next button.
6. The Kernel Options dialog box appears, prompting you to choose the type of kernel components to build into the kernel. Select `Customize` (you will choose the individual kernel options in step 13). Click on the Next button.
7. The Select File System Layout dialog box appears. Select `Customize File System Layout` and click on the Next button.
8. The Custom File System Layout dialog box appears. You must set several values on this screen, as follows:
 - a. Set the Use LSM option to `No`.

You must not use LSM at this stage; instead, configure LSM after all nodes have been added to the domain — see Section 8.11 on page 8–18.
 - b. Configure the recommended partition layout for the system disk.

Table 6–3 shows the recommended partition layout for an 18GB system disk, and for a 36GB disk, on Node 0.

Table 6–3 Recommended Partition Layout for Node 0 System Disk

File System	Disk	Partition	Size for an 18GB Disk	Size for a 36GB Disk	Type
root	dsk2	a	384MB	384MB	AdvFS
/usr	dsk2	d	5.4GB	11.1GB	AdvFS
swap1	dsk2	f	5.4GB	11.1GB	swap
/var	dsk2	e	5.4GB	11.1GB	AdvFS

See Chapter 6 of the *HP Tru64 UNIX Installation Guide* for more information about customizing the file system layout.

9. When you have entered all of the required information, click on the Next button.
10. The Installation Summary dialog box appears. Review the information on this screen. To change any values, click on the Reset button. When the information is correct, click on the Finish button.

- 11. The Ready to Begin Installation dialog box appears. Click on the OK button. The system saves the configuration, creates the file systems, and loads the software.
- 12. The system then automatically reboots from the system disk.
- 13. Software configuration occurs automatically after your system reboots. The Kernel Option Selection dialog box appears. You must select at least the options listed in Table 6–4.

Table 6–4 Minimum Kernel Options

Selection	Kernel Option
2	NTP V3 Kernel Phase Lock Loop (NTP_TIME)
3	Kernel Breakpoint Debugger (KDEBUG)
4	Packetfilter driver (PACKETFILTER)
12	ISO 9660 Compact Disc File System (CDFS)

If you need to install a kernel component after installation is complete, use the `doconfig(8)` command. See the *HP Tru64 UNIX Installation Guide* for more information.

- 14. When prompted to edit the configuration file, select `No`.
- 15. The kernel build procedure automatically begins after software configuration.
- 16. After the kernel build, the system reboots automatically.
- 17. The final step is to log into the newly installed system as the `root` user. (You may see warnings about missing license PAKs for OSF-BASE — you can ignore these until Section 6.1.4.1 on page 6–8.)

When you log in for the first time, the Tru64 UNIX System Setup dialog box appears. Click on the Custom Setup icon. The Tru64 UNIX Custom Setup menu appears. Customize the system configuration, as described in Section 6.1.4.

Note:

If you choose Quick Setup, your configuration options are reduced — you will have to rerun the `sysman` command to complete the configuration.

Set Up Node 0

6.1.4 Customize the System Configuration

The Tru64 UNIX Custom Setup menu offers a number of options, as illustrated in Figure 6–1.

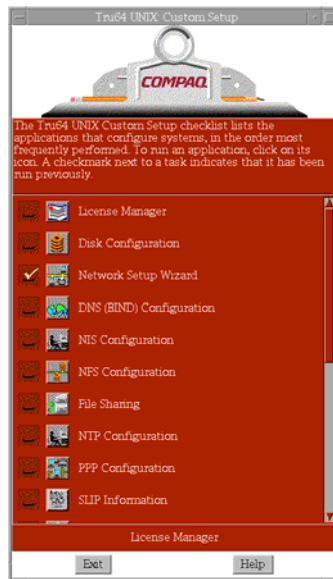


Figure 6–1 Tru64 UNIX Custom Setup Menu

Not all of these options are mandatory. This section describes only the systems that you must configure. For more information about the Custom Setup Menu options, see the *HP Tru64 UNIX Installation Guide*.

First register the licenses, as described in Section 6.1.4.1.

6.1.4.1 Register Licenses (PAKs)

A license is a contract with terms and conditions. Each license has an associated Product Authorization Key (PAK), which is a set of characters.

You must use the license manager to enter the license information on each PAK into the license database; this process is called registering a license. Each time a user attempts to run a licensed product, the product calls the license-checking functions to be sure that the license allows the user to use the product.

For more information about licenses, see the *Tru64 UNIX Software License Management* manual.

Note:

The licenses listed below are those required to use software that is discussed elsewhere in this document. However:

- Not all of these licenses are required.
 - There are other licenses that may be useful. These licenses would be installed in a similar way to the licenses below, but are not covered in this document.
-

Click on the License Manager icon (on the Custom Setup menu) to register the following licenses:

- HP AlphaServer SC System Software License (TCS-UA and ASC). This license is required; you need this license to run the console logger daemon (`cmfd`), and to use the `sra setup` command. See below for more information.
- Remote Installation Services License (OSF-SVR). This license is required; it is essential for system operation.
- Operating System Base License (OSF-BASE). This license supports up to two interactive users.
- Symmetric Multiprocessing (SMP) Extension to Base License (OSF-BASE). This license is optional; it is only needed for each additional CPU on the system.
- Concurrent Use License (OSF-USER) or Unlimited Interactive User License (OSF-USER). These licenses are optional; they are only needed if you wish to support interactive users.
- Advanced File System Utilities License (ADVFS-UTILITIES). This license is required; it is needed if you wish to configure additional storage in an AdvFS file domain.
- HP AlphaServer SC Development Software License (OSF-DEV). This license is optional; it is only needed if you wish to use the Ladebug debugger (see below) or other products such as the C, C++, and Fortran compilers, as well as other performance and debugging tools.

HP AlphaServer SC System Software

The HP AlphaServer SC System Software License is sold in bundles of 1-, 16-, 32-, 64-, and 128-node licenses. You can combine any of these bundles so that the total number of licenses is equal to, or more than, the number of nodes in the HP AlphaServer SC system. For example, you can combine one 16-node license with one 32-node license for a 48-node HP AlphaServer SC system. The management server (if present) is not counted as a node — although the licences must be installed on the management server.

Set Up Node 0

The *HP AlphaServer SC40 QuickSpecs* provide ordering information for the 1-, 16-, 32-, 64-, and 128-node licenses.

The license file for the HP AlphaServer SC System Software License contains two **Product Authorization Keys (PAKs)**: ASC and TCS-UA. Both of these PAKs must be installed on the system. A PAK contains a **units** field that identifies the capability of the PAK — the ASC PAK shown in Example 6–1 contains 6400 units.

Example 6–1 Sample ASC PAK

```
Issuer: DEC
Authorization Number: QS-SYS-16-NODE
Product Name: ASC
Producer: DEC
Number of units: 6400
Key Termination Date: 19-MAY-2002
Activity Table Code: CONSTANT=100
Checksum: 1-ABCD-GHDN-BOCJ-CLFH
```

Depending on the PAK, the units are used differently, as follows:

- **ASC**

In LMF terms, the ASC PAK is a **concurrent-use** PAK. The number of units determines the number of nodes that are licensed — 400 units are required to license a node. When you register and load the ASC PAKs, the LMF system combines the units of all of the ASC PAKs together. You will see messages to this effect when you install the second and subsequent ASC PAKs.

- **TCS-UA**

In LMF terms, the TCS-UA PAK is a **capacity-based** PAK. The number of units corresponds to the model of the AlphaServer — for example, an HP AlphaServer ES40 requires 1050 units. Unlike the ASC PAK, LMF does not combine the units of several TCS-UA PAKs. When you attempt to install a second or subsequent TCS-UA PAK, LMF will print messages saying that the PAKs were not combined. This does not matter as you only require one TCS-UA PAK for the system to operate correctly. It does not matter if several TCS-UA PAKs are registered.

The ASC and TCS-UA PAKs are delivered as a shell script containing the LMF commands to register and install the PAKs. If you purchased several HP AlphaServer SC System Software Licenses (for example, a 16-node license and a 32-node license) at the same time, you will receive a single shell script containing the appropriate PAKs. To install the PAKs, simply copy the script to the system and execute the script. If you later buy additional licenses, you will receive a second shell script containing the appropriate additional PAKs — the new script does not contain PAKs for the original licenses. To ensure that all PAKs are installed, you should execute each shell script. You should also keep a copy of each license shell script, in case you ever need to do a complete system reinstall.

As mentioned earlier, when you run the second and subsequent shell scripts, you will see messages saying that the ASC PAKs were combined but that the TCS-UA PAKs were not. This is normal.

If you do not install the ASC and TCS-UA PAKs, or if the number of ASC units is not appropriate for the number of nodes in the system, the system will not operate correctly as follows:

- The `sra setup` command will print a warning when you first specify the number of nodes. You can ignore this warning and attempt to continue. However, you may be unable to operate parts of the system at a later stage.
- The console management daemon (`cmfd`) will not start. You will be unable to access node consoles. A message will be printed to the `/var/sra/adm/log/cmfd/cmfd_<hostname>_<port>.log` file.

The shell script automatically registers and loads the ASC and TCS-UA PAKs. If you inadvertently unload the PAKs, the system will not operate — that is, a PAK must be both registered and loaded for it to work. You can tell how many units are loaded as follows:

```
# lmf list full cache for ASC
```

You can load previously registered ASC and TCS-UA PAKs as follows:

```
# lmf load 0 ASC
# lmf load 0 TCS-UA
```

The Ladebug Debugger

The Ladebug debugger, distributed with Tru64 UNIX Version 5.1B, is a symbolic source-level debugger that supports debugging of ADA, C/C++, Fortran, and Fortran 90 applications. RMS uses the Ladebug debugger to print a back trace when an application core dumps.

To use the Ladebug debugger, you need the OSF-DEV license. You can obtain this license by purchasing, for example, HP AlphaServer SC Development Software, or Developer's Toolkit for Tru64 UNIX.

Note:

If you are not licensed to use the Ladebug debugger, RMS will not print a back trace.

When you have completed the license PAK registration, click on the Exit button to return to the Custom Setup Menu.

You are now ready to configure the networks, as described in Section 6.1.4.2.

Set Up Node 0

6.1.4.2 Set Up Networks

Click on the Network Setup Wizard icon (on the Custom Setup menu) to configure the management and external networks. This displays a submenu as shown in Figure 6–2.

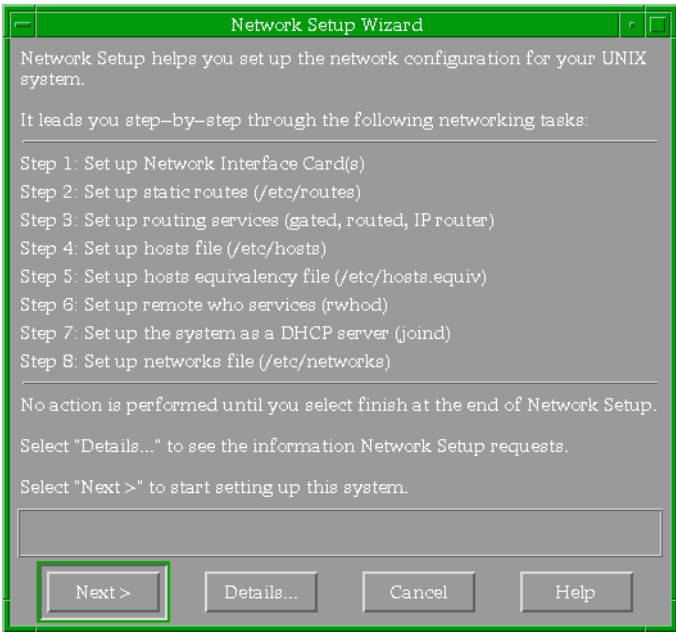


Figure 6–2 Tru64 UNIX Network Setup Wizard Menu

Configure the networks as follows:

1. Click on the Next button to display the Set up Network Interface Card(s) screen.

Set up the network interface cards using the settings provided in Table 6–5, where *atlas* is an example system name, and *x.x.x.x* and *y.y.y.y* are site-dependent values (recorded in Appendix D).

Table 6–5 Network Interface Cards on Node 0

Network Type	Device Name	Host Name	IP Address	Network Mask
Management	ee0	atlas0	10.128.0.1	255.255.0.0
External	ee1	atlas0-ext1	x.x.x.x	y.y.y.y

2. Click on the Next button to display the Set Up Static Routes screen.

Set up a default static route, as follows:

- a. Click on the Add button to display the Network Setup: Set Up Static Routes: Add/Modify screen.
- b. Set the Destination Type to Default Gateway.
- c. In the Gateway field, enter the external gateway IP address (see Appendix C).
- d. Set the Route Via option to Gateway, and click on the OK button.

Note:

The default route for each node added to a domain is automatically set to the IP address of the first node in the domain.

3. Click on the Next button to display the Set Up Routing Services screen.
- You must set the routing services using the settings provided in Table 6–6:

Table 6–6 Routing Services

Question	Answer
Do you want to set up a routing service for this system:	Yes (use gated)
Do you want to run this system as an IP Router:	Yes

4. Click on the Next button to display the Set Up Hosts File screen.
- The hosts file will automatically contain the entries listed in Table 6–7. The following notes apply to Table 6–7:
- atlas is an example system name.
 - x.x.x.x represents a site-specific value (recorded in Appendix D).
 - DNS servers and NIS servers will be added later (see Section 6.1.4.3 on page 6–14 and Section 6.1.4.6 on page 6–16 respectively).

Note:

The `sra setup` command will update the `/etc/hosts` file to add an entry for all of the hosts associated with the cluster (for example, `atlas1`, `atlas2`, ..., `atlas127`) as well as each cluster alias (see Section 6.2, step 10 on page 6–26).

Table 6–7 Hosts File When Configuring Node 0

IP Address	Host Name	Aliases/Comments
127.0.0.1	localhost	

Set Up Node 0

Table 6–7 Hosts File When Configuring Node 0

IP Address	Host Name	Aliases/Comments
10.128.0.1	atlas0 ¹	
x.x.x.x	atlas0-ext1 ²	

¹This value is not displayed, but if you try to add it, you will be told that it is already in the hosts file.

²The `/etc/hosts` file may not contain the entries as shown in Table 6–7. If you are using the Network Setup Wizard, these entries will exist after you have completed all the Network Setup Wizard steps and select finish.

- 5. Click on the Next button to display the Set Up Hosts Equivalency File screen. Do not set up a hosts equivalency file.
- 6. Click on the Next button to display the Set Up Remote Who Services screen. Select No.
- 7. Click on the Next button to display the Set Up the System as a DHCP Server screen. Select No. Accept the default setting of Log No Messages.
- 8. Click on the Next button to display the Set Up Networks File screen. Click on the Next button to skip this step.
- 9. Click on the Finish button to return to the Custom Setup menu. When prompted to restart services, select Yes.

You are now ready to configure Domain Name Service / Berkeley Internet Name Daemon (DNS/BIND), as described in Section 6.1.4.3.

6.1.4.3 Configure DNS (BIND)

Note:

HP AlphaServer SC Version 2.6 (UK2) supports configuring the system as a DNS client only — do not configure the system as a DNS server.

If you wish to configure DNS, perform the following steps:

- 1. Click on the DNS (BIND) Configuration icon (on the Custom Setup menu) to display the DNS (BIND) Configuration screen.
- 2. Double click on the Configure System as a DNS Client entry.
- 3. When prompted to add the DNS server to the hosts file, enter the local domain and click Add.
- 4. When prompted to update the host name, select No (see Note below).

Note:

If you select Yes, the host name is changed. If the host name is changed, RMS will not work (because the `rmshost` attribute in the RMS database is wrong). If you change the host names (by selecting Yes), log onto each node in turn and change the host name back to the original value, by removing the domain name.

5. Click on the Exit button to return to the Custom Setup menu.

For more information about configuring DNS (BIND), see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network Time Protocol (NTP), as described in Section 6.1.4.4.

6.1.4.4 Configure NTP

Click on the NTP Configuration icon (on the Custom Setup menu) to configure the Network Time Protocol (NTP). You must know the name of an NTP server that is accessible through the external network interface on the first node (Nodes 0, 32, 64, or 96). Perform the following steps as the `root` user:

1. Double click on the Configure System as an NTP Client entry.
2. Click on the Add button, and enter the hostname of the NTP server on the external network.

Note:

The hostname of the NTP server does not need to be fully qualified. After network configuration, you can check to see that external network connection is valid by trying to `ping` one of the NTP servers. Otherwise the NTP restart will hang when it tries to start the service.

3. Set the Mode to Server. Accept the default settings for Version and Key Number.
4. Click on the OK button to return to the Configure System as an NTP Client screen.
5. Click on the OK button to return to the Custom Setup Menu. When prompted to start `xntpd`, click on the Yes button.

For more information about configuring NTP, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network File System (NFS), as described in Section 6.1.4.5.

Set Up Node 0

6.1.4.5 Configure NFS

Click on the NFS Configuration icon (on the Custom Setup menu) to display the NFS Configuration screen, and perform the following steps:

1. Double click on the Configure System as an NFS Client entry and accept the default settings.
2. Double click on the Configure System as an NFS Server entry and accept the default settings.

Ensure that the Enable Locking check box is selected (this is the default setting).

3. When prompted to restart the NFS daemons, select Yes.
4. Click on the Exit button to return to the Custom Setup menu.

For more information about configuring NFS, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure the Network Information System (NIS), as described in Section 6.1.4.6.

6.1.4.6 Configure NIS

If you wish to configure the Network Information System (NIS), perform the following steps:

Note:

The installation of the HP AlphaServer SC RMS subset automatically adds a UNIX user named "rms" and a UNIX group named "rms" with specific identifiers.

However, the RMS subset installation will fail if any of the following already exist in the NIS database on the NIS master server:

- any group or user named "rms"
- any user with uid=15
- any group with gid=200

Such entries in the NIS database must be removed before proceeding. If you cannot remove such entries from the NIS database, then do not configure NIS at this time. NIS may be configured after the entire system is configured, however, this can be a security issue as NIS users with uid=15 or gid=200 will have access to RMS files on the HP AlphaServer SC system.

1. Click on the NIS Configuration icon (on the Custom Setup menu) to display the NIS Configuration screen.

2. Enter and confirm your system's NIS domain name.
3. Choose option 2 to indicate that you are configuring a slave server.

Note:

HP AlphaServer SC Version 2.6 (UK2) supports configuring Node 0 as a NIS slave server only. Do not configure the system as a NIS master.

4. When prompted to copy the current maps from the master server, select Yes.
5. When prompted, please provide the NIS MASTER server for the domain.
6. If prompted to add the NIS server to the hosts file, select Yes and subsequently enter the relevant details.
7. When prompted regarding enhanced security, select No.
8. When prompted regarding maintaining maps as "btree" files, select No.
9. When prompted to use `ypbind` secure mode (option -s), select Yes.
10. When prompted to use `ypbind` domain locks (option -S), select Yes.
11. When prompted, enter the NIS MASTER in addition to the slave server itself in the list of authorized NIS servers.
12. When prompted, choose option 3 to disallow all `ypset` requests.
13. When prompted to configure the system to use all of the NIS databases, select Yes.
14. When prompted to start the NIS daemons now, select Yes.
15. Click on the Finish button to return to the Custom Setup menu.

For more information about configuring NIS, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure mail, as described in Section 6.1.4.7.

6.1.4.7 Configure Mail

Click on the Mail Configuration icon (on the Custom Setup menu) to configure mail.

Note:

You must configure mail.

Perform the following steps as the `root` user:

1. Double click on the Configure Mail as a Client entry.

Set Up Node 0

If DNS has not been set up, a warning is displayed.

2. Enter the mail server, and click the Commit button.
3. Select the Yes option to restart sendmail.
4. Select the Close option to return to the Custom Setup menu.

For more information about configuring mail, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure printers, as described in Section 6.1.4.8.

6.1.4.8 Configure Printers

If you wish to configure printers, click on the Printer Configuration icon (on the Custom Setup menu) and configure printers as described in Chapter 6 of the *Tru64 UNIX Network Administration* manual.

Note:

HP AlphaServer SC Version 2.6 (UK2) supports remote printing only — do not attach a printer directly to the HP AlphaServer SC system.

Click on the Finish button to return to the Custom Setup menu. Click on the Exit button to return to the operating system prompt.

For more information about the other configuration options, such as security, see the *HP Tru64 UNIX Installation Guide*.

You are now ready to install the latest operating system patch, as described in Section 6.1.5.

6.1.5 Install the Latest Operating System Patch Software

Install the Tru64 UNIX patch software on the management server.

Note:

For HP AlphaServer SC Version 2.6 (UK2), you should load Tru64 UNIX Version 5.1B-3 (also known as Tru64 UNIX Version 5.1B Patch Kit 5) only.

The operating system patch software kit (T64V51BB26AS0005-20050502.tar) is available from the following location:

<<http://www.itrc.hp.com/>>

or from your local HP support representative.

You are now ready to configure the remote installation services (RIS) server, as described in Section 6.1.6.

6.1.6 Configure the RIS Server

Installation of the Tru64 UNIX operating system on the remaining nodes of each domain will be performed using a Remote Installation Services (RIS) server. The RIS server allows each new member to boot a network `vmunix`.

The RIS server is also necessary for performing particular diagnostics during the installation phase. It is important to configure the RIS server at this stage and before you install the HP AlphaServer SC software in section Install the HP AlphaServer SC System Software (see Section 6.1.7 on page 6–20).

Configure Node 0 as a RIS Server, as follows:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1* CD-ROM in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 6.1.7, step 2 on page 6–20).
3. Run the `ris` command as the `root` user, as follows:
`atlas0# ris`
4. Choose the `Install software products` option by entering `i` at the prompt:
`Enter your choice: i`
5. The RIS Installation menu displays the installation options. Choose option 1, the `Install software into a new area` option.
6. Enter the full pathname for the distribution media, as follows:
`Enter the device special file name or the path of the directory
 where the software is located
 (for example, /mnt/ALPHA/BASE): /cdrom/ALPHA/BASE`
7. Choose the standard boot method.
8. Choose to extract the software from the *Tru64 UNIX Version 5.1B Operating System Volume 1* CD-ROM.

Note:

If you choose to link to the CD-ROM instead of extracting the software, the installation process is considerably slower.

9. Choose to install all mandatory and all optional subsets (option 72).
10. Enter `y` to confirm that the subset list is correct. The subset extraction process begins.

Set Up Node 0

Do not configure any RIS clients at this time — configure RIS clients later using the `sra setup` command (see Section 6.2, step 22 on page 6–30).

Note:

For information on RIS security recommendations, see Section 10.2.4, page 10–10.

You are now ready to install the HP AlphaServer SC installation utility (SRA) software, as described in Section 6.1.7.

Note:

You can rebuild the kernel up to the point where the HP AlphaServer SC software is installed.

6.1.7 Install the HP AlphaServer SC System Software

To install the HP AlphaServer SC software on Node 0, perform the following steps as the `root` user:

1. Insert the *HP AlphaServer SC System Software* CD-ROM in the disk drive.
2. Mount the CD-ROM as the `root` user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:
`atlas0# mkdir /cdrom`
 - b. Mount the CD-ROM¹ as follows:
`atlas0# mount -r /dev/disk/cdrom0c /cdrom`

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:
`atlas0# cd /cdrom/kits`
4. Install the HP AlphaServer SC software. From the kits directory above, if you do not have a management server, run the following command:
`atlas0# ./InstallSC -install -noms`

You are now ready to install the HP Fortran Run-Time Libraries, as described in Section 6.1.8.

1. See the *HP Tru64 UNIX Installation Guide* for more information on how to identify the CD-ROM device name. The CD-ROM device name in this example is `/dev/disk/cdrom0c`.

6.1.8 Install the HP Fortran Run-Time Libraries

To install the HP Fortran Run-Time Libraries, perform the following steps:

1. Insert the *Tru64 UNIX Version 5.1B Operating System Volume 1* CD-ROM in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 6.1.7, step 2 on page 6–20).
3. Install the HP Fortran Run-Time Library, as follows:


```
atlas0# cd /cdrom/DEC_Fortran_RTL/kit
atlas0# setld -l .
```

You are now ready to install layered products, as described in Section 6.1.9.

6.1.9 Install Layered Products (Optional)

This step is optional. If you wish to install layered products (for example, CXML Compaq Extended Math Library), do so at this point, if possible.

If the product has a license (PAK), install the license now. Licenses added to the first domain member are propagated to new members added to the domain.

You are now ready to install the SANworks Storage System Scripting Utility, as described in Section 6.1.10.

6.1.10 Install the SANworks Storage System Scripting Utility

The SANworks Storage System Scripting Utility (SSSU) is required so that the SC Monitor system can monitor HP SANworks Management Appliance and HSV110 RAID System devices. If your system does not have a SANworks Management Appliance or HSV110 RAID System, you can skip this step.

To install the SANworks Storage System Scripting Utility, follow these steps:

1. Order HP SANworks Tru64 UNIX Kit for Enterprise Virtual Array.
2. This kit contains a CD-ROM. Mount the CD-ROM as the `root` user, as follows:
 - a. Create a mount point for the CD-ROM, by running the following command:


```
atlas0# mkdir /cdrom
```
 - b. Mount the CD-ROM as follows:


```
atlas0# mount -r /dev/disk/cdrom0c /cdrom
```

For more information about mounting a CD-ROM, see Appendix B of the *HP Tru64 UNIX Installation Guide*.

3. Change to the directory in which the kits are stored, as follows:


```
atlas0# cd /cdrom
```
4. Install the software, as follows:

Set Up Node 0

```
atlas0# setld -l .
*** Enter subset selections ***
The following subsets are mandatory and will be installed automatically
unless you choose to exit without installing any subsets:
* Enterprise v2 For Tru64 Unix
```

You may choose one of the following options:

- 1) ALL of the above
- 2) CANCEL selections and redisplay menus
- 3) EXIT without installing any subsets

Estimated free disk space (MB) in root: 646.3 usr:

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 1

You are installing the following mandatory subsets:

Enterprise v2 For Tru64 Unix

You are installing the following optional subsets:

Estimated free disk space (MB) in root: 646.3 usr: 1716.8

Is this correct? (y/n): y

5. When the `setld` command has completed, create a link to `/usr/bin/sss`, as follows:

```
atlas0# ln -fs /usr/opt/ENTP002/sbin/sss /usr/bin/sss
```

This step is important, because the monitoring scripts expect the `sss` binary to be located in the `/usr/bin` directory.

You are now ready to add the `sysconfigtab` parameters, as described in Section 6.1.11.

6.1.11 Add `sysconfigtab` Parameters

The Tru64 UNIX operating system includes various subsystems that are used to define or extend the kernel. Kernel variables control the behavior of these subsystems, or track subsystem statistics since boot time.

Kernel variables are assigned default values at boot time. For certain configurations and workloads, especially memory- or network-intensive systems, the default values of some attributes may not be appropriate, so you must modify these values to provide optimal performance. You can modify the values of kernel variable attributes by adding `sysconfigtab` parameters to the `/etc/sysconfigtab` file.

The recommended `sysconfigtab` parameters are provided in the file `/Examples/sysconfigtab` on the *HP AlphaServer SC System Software* CD-ROM. This sample `sysconfigtab` file assumes a memory size of 4GB.

Note:

The `sysconfigdb` commands below will apply the values from this sample `sysconfigtab` file. Before issuing the `sysconfigdb` commands, review the `/cdrom/Examples/sysconfigtab` file.

If necessary, copy the file to a temporary area on your system and modify the values — then specify the location of the modified file in the `sysconfigdb` commands.

Note that the default value for the `shm_max` attribute is 4GB.

To add the recommended `sysconfigtab` parameters on Node 0, perform the following steps:

1. Insert the *HP AlphaServer SC System Software* CD-ROM in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 6.1.7, step 2 on page 6–20).
3. Add the parameters to the `/etc/sysconfigtab` file by running the `sysconfigdb(8)` command as the `root` user, as follows:
`atlas0# sysconfigdb -m -f /cdrom/Examples/sysconfigtab`
4. When creating a new domain member (see Section 7.4 on page 7–23), the `clu_add_member` command uses the `/etc/.proto..sysconfigtab` file instead of the `/etc/sysconfigtab` file. Therefore, you must update the `/etc/.proto..sysconfigtab` file as follows:
`atlas0# sysconfigdb -m -t /etc/.proto..sysconfigtab -f /cdrom/Examples/sysconfigtab`

For more information about kernel variables and `sysconfigtab` parameters, see the *Tru64 UNIX System Configuration and Tuning* manual.

You are now ready to run the `sra setup` command, as described in Section 6.2.

6.1.12 Define the RMS Master Node (rmshost)

Define Node 0 to be the RMS master node; that is, the `rmshost` system.

To do this, update the `/etc/hosts` file on the management server, to define `rmshost` as a host alias for Node 0, as shown in the following example:

```
10.128.0.1      atlas0      rmshost
```

If using a C shell, run the `rehash` command, as follows:

```
atlas0# rehash
```

Start the `mSQL` daemon, as follows:

```
atlas0# /sbin/init.d/msqld start
```

You may see an error indicating that the `mSQL` daemon was already running. Ignore this error.

You are now ready to run the `sra setup` command, as described in Section 6.2.

6.2 Set Up the SC Database

Note:

The SC database supersedes both the SRA database (the flat file `/var/sra/sra-database.dat`) and the RMS database (also a SQL-based database).

Set up the SC database on Node 0 by running the following commands as the `root` user:

1. If you have a second rail in each HP AlphaServer ES40 system, but the rail is either currently unconnected or powered off, then you must first remove the second (expansion) elan card from the nodes before performing the `sra install` step. When prompted during the `sra setup` dialog, please answer that the system will have one rail. Later, once the domains are built and all members added, the second rail can be added as a post-installation step by following the instructions in Section 8.13. Please refer to Section 3.5 for details on PCI slot selection rules.
2. Plan the domain attributes (see step 5) and record them in Appendix D.
3. Ensure that all nodes are halted; that is, that they are powered up at the SRM console prompt.

Note:

For some system installations that use CAA for the RMS service, there is a possibility that `rmsbuild` will report a warning during `sra setup`. If this warning occurs, it can be ignored and you should answer "yes" when prompted. The warning is being ignored as it will automatically be handled later by the CAA startup script for RMS.

4. Run the `sra setup` command to create and populate the SC database, as follows:
`atlas0# sra setup`

This information is needed by other `sra` commands.

5. The `sra setup` command prompts you for the following information:
 - System name¹
 - Number of nodes
 - Number associated with first domain and node²
 - Management Server name
 - Hardware type of the system

- Number of Cluster Interconnect Rails used in the system
 - Number of domains in the system
 - Management Network IP address
 - Cluster Interconnect IP address for Node 0¹
 - System Interconnect IP address for Node 0²
 - Terminal server model
 - Number of ports on the terminal server
(if you do not accept the default terminal server model)
 - IP address of first terminal server
 - First port on first terminal server
 - Interconnect Switch IP address at Node Level and Top Level for each Rail
 - IP address for the Preferred Server cluster alias base address
6. The `sra setup` command displays these settings and asks you to confirm that they are correct. If any settings are incorrect, enter **No** and repeat step 5. When all settings are correct, enter **Yes**.
 7. The `sra setup` command prompts you to indicate whether you would like automatic cluster configuration.

-
1. The HP AlphaServer SC installation process uses the system name (which cannot end with a digit) to derive both the cluster alias names and the host names of each member in the domain. For example, in a 64-node system with system name `atlas`, the cluster aliases will be `atlasD0` and `atlasD1`, while the node host names will be `atlas0`, `atlas1`, ..., `atlas63`.
Note that if the number of nodes is less than or equal to 32, the cluster alias name will be the same as the system name — in the above example, the cluster alias name would be `atlas`. If you later increase the number of nodes to 33 or more, you must create a second domain (`atlasD1`), and rename the original cluster alias (from `atlas` to `atlasD0`).
 2. Domain and node names in a system typically start at 0, for example `atlasD0`, `atlas0`. However, it is possible to start at a different number, for example `atlasD32`, `atlas1024`.
 1. The Cluster Interconnect network is an IP network provided by TruCluster Server. This network only spans a domain. This network is an artifact of the TruCluster Server software and is not intended for end-user use. This network is denoted by the `ics` suffix.
 2. The System Interconnect network is layered on the HP AlphaServer SC Interconnect, and spans all nodes connected to the HP AlphaServer SC Interconnect. This network can be configured with externally routable IP addresses (instead of the default 10.* network). For example, this allows you to perform high-speed FTP by routing through high-performance interfaces on externally connected nodes. Note that such use will impact bandwidth availability to parallel jobs. This network is denoted by the `eip` suffix.

Set Up the SC Database

Note:

If you have allocated external node IP addresses and cluster alias IP addresses according to the scheme mentioned in Section 2.2, and if all domains in the system (except possibly the first and last) are composed of the same number of nodes, enter **Yes** when the `sra setup` command prompts you to indicate whether to use automatic cluster configuration. If you enter **Yes**, `sra setup` will prompt you to enter the cluster sizes and the two base IP addresses. `sra setup` will check the IP addresses, and then it will automatically enter configuration information for each cluster into the database, saving you from having to answer several questions for each cluster.

8. The `sra setup` command prompts you for the following information for each domain:
 - Number of Nodes in the First Cluster
 - Number of Node in Subsequent Clusters
 - External network IP address for First Node
 - IP address for cluster alias
9. The `sra setup` command automatically assigns the IP addresses and prompts you to add the nodes to the database.
10. The `sra setup` command asks if you would like to update the `/etc/hosts` file. Press Return to accept the default answer of **Yes**. The `sra setup` command adds an entry to the `/etc/hosts` file for each IP address specified in steps 5 and 7 above.
11. The `sra setup` command asks which host should run console monitor (the `cmf` utility). Press Return to accept the default value (the current node). The `sra setup` command starts `cmf`, and informs you of the location of the `cmf` log file.
12. The `sra setup` command asks if you would like to configure the terminal server(s). Press Return to accept the default answer of **Yes**; the `sra setup` command configures the terminal server ports.
13. The `sra setup` command asks if you would like to configure an alternate boot device. Enter **yes**.
14. The `sra setup` command asks if you would like to use an alternate boot device. Enter **yes**.
15. The `sra setup` command prompts you for the information needed to configure the boot device(s). You can accept the default values for all settings except the SRM device name. To enter a value for the SRM device name, perform the following steps:
 - a. Enter the SRM device name for the primary boot disk at the SRM device name for primary boot disk? prompt. If you do not know the SRM device name, enter **probe**.

- b. If prompted, enter the host name of a generic node that is currently at the SRM prompt. This is only necessary if you entered `probe` in step 15a. above.
If you enter the host name of a node that is not currently at the SRM prompt, the hardware probe will fail.
- c. The `sra setup` command displays the settings for the primary boot disk, including the SRM device name. To accept these settings, enter `y`. To change any setting, enter `n`, enter the item number of the setting to be changed, and enter the new value.

You can set the alternate boot device in a similar way.

Table 6–8 shows the recommended boot disk partition configuration when the memory size is 4GB and each boot disk is a 36GB disk.

Table 6–8 Recommended Boot Disk Partition Configuration

Partition	Description	Recommended Size (% and GB)	
		Boot Disk Only	Boot Disk and Alternate Boot
b	swap	30% = 10.8GB	15% = 5.4GB on each disk
d	/local	35% = 12.6GB	42% = 15.2GB on each disk
e	/tmp	35% = 12.6GB	43% = 15.4GB on each disk

For a 36GB disk, the swap space is calculated as 2.5 times the physical memory size. However, for an 18GB disk, the swap space is calculated as 1.25 times the physical memory size, which you may consider to be too low. Think carefully before accepting the default partition values.

Note:

If you configure an alternate boot disk and have chosen to use the alternate boot disk, its swap space is added to the `sysconfigtab` file; that is, you spread the swap space across the two disks. Also, the `tmp` and `local` partitions on the alternate boot disk are mounted on `/tmp1` and `/local1` respectively. You should take this into account when deciding the partition percentages, as follows:

- If configuring only one boot disk, use the values specified in the third column of Table 6–8.
- If configuring a boot disk and an alternate boot disk, use the values specified in the fourth column of Table 6–8.

16. The `sra setup` command asks for the information needed to configure the primary boot disk.

Set Up the SC Database

Press Return to accept the default answer. The SRM device name is not required to complete the installation.

17. The `sra setup` command asks for the information needed to configure the Cluster /`/usr` and/`/var` disk.

Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation.

The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit` (see Section 7.2).

18. The `sra setup` command asks for the information needed to configure the backup cluster disk.

Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation.

The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit`.

19. The `sra setup` command asks for the information needed to configure the generic boot disk.

Press Return to accept the default answer for all settings. The SRM device name is not required to complete the installation.

The Disk Location Identifier is required but at this point you can accept the default value. After running `sra setup` the identifiers can be updated using `sra edit`.

20. The `sra setup` command prompts you for the SRM device name and the UNIX device name for the management and external LAN adapters. You can accept the default value, or enter a new value, or use **probe** as described in step 15.

Note:

For an external LAN adapter, if the external LAN device name is unknown, you can accept the default SRM device name (for example, `eia0`) but you must enter the appropriate UNIX device name. Page xxix lists the SRM and UNIX device names for the supported network adapters.

21. The `sra setup` command asks if you would like to probe each node in the cluster for its hardware ethernet (MAC) address.

Note:

The `probe` command will return the MAC address of the first network interface only. The `sra edit` command may be used to set the MAC address manually (see Section C.1, step 4 on page C-3).

Enter **yes** — the `sra setup` command asks two additional questions:

- a. The script asks which nodes you would like to probe.
You can specify a single node (for example, `atlas3`), several nodes (for example, `atlas3,atlas5`), a range of nodes (for example, `atlas[1-8]`), or all nodes (`all`).
Enter **atlas[1-127]**.
- b. The script also asks if you would like to initialize hardware.
This runs the `SRM init` command on each node. During the probe, SRM console variables are set, as detailed in Table 6-9.
Enter **yes**.

Table 6-9 SRM Console Variables

Variable		Value
<code>auto_action</code>		<code>halt</code>
<code>boot_reset</code>		<code>off</code>
<code>boot_osflags</code>		<code>A</code>
<code>os_type</code>		<code>UNIX</code>
<code>ocp_text</code>		<code>nodename</code>
<code>eia0_mode</code> or <code>eib0_mode</code>	(<code>eia0_mode</code> if the management network interface is the first network adapter on the node; otherwise, the device name could be <code>eib0_mode</code>)	<code>FastFD</code>
<code>console</code>		<code>serial</code>
<code>sys_com1_rmc</code>		<code>off</code>

If any nodes fail the hardware probe, a list of failed nodes is written to the `/tmp/sra_hosts.failed` file. A node may fail the hardware probe because it is not at the SRM prompt, or because `cmf` is misconfigured — see Section 11.8 on page 11-26 for more details.

Set Up the SC Database

Note:

If you specified **a11** in step a on page 6–29, you will get a message that Node 0 has failed the hardware probe — ignore this message. Node 0 fails the hardware probe because it is not at the SRM prompt. If you wish, you can set MAC address manually using the `sra edit` command (see Section C.1, step 4 on page C–3).

The `sra setup` command asks if you would like to reset the failed nodes. Enter **Yes**.

The `sra setup` command asks if you would like to retry the hardware probe. Enter **Yes**.

The `sra setup` command asks if you would like to initialize the hardware during the probe. Enter **Yes**.

22. The `sra setup` command asks if you would like to add nodes to the RIS database. Press Return to accept the default answer of **Yes**. The `sra setup` command configures each node as a RIS client. During the Configure RIS section of `sra setup`, you are asked if you would like to add domains to RIS. Press Return to accept the default answer of **Yes**.
23. The `sra setup` command will now prompt to add two root crontab entries (one crontab for archiving the `cmf` logs and one crontab for backing up the RMS database). Press Return to accept the default answer of **Yes** to both questions. The `cmf` archives will be archived every two weeks and the RMS database backup will take place nightly.
24. The `sra setup` command saves all information to the database, and then exits.

Note:

After the SC database has been created, the `sra setup` command completes the installation of the SRACFENGINE subset by running the `cfconfig` script on Node 0.

The `cfconfig` command automatically runs on each cluster node when the node is first booted. The `cfconfig` command examines the local host to populate the appropriate files with default values. The `cfengine` daemon (`cfengine`) starts for the first time after the `cfconfig` command has run.

If DNS has not been configured on the system, the `cfconfig` command will print an error message when it cannot find `/etc/resolve.conf`. This is normal.

`cfengine` cannot operate unless DNS has been configured. You can run `cfconfig` at a later stage once DNS is configured.

If you later wish to change any of the data in the SC database, use the `sra edit` command (see Chapter 16 of the *HP AlphaServer SC System Administration Guide*) — it is not necessary to rerun the `sra setup` command.

Check All Nodes in the HP AlphaServer SC System, Except Node 0

See Appendix E for an example `sra setup` log file.

Note:

Now that you have set up the SC database, you can use the `sra -c|cl|m|ml` command on Node 0 to connect to the console of any system in the cluster. See Chapter 16 of the *HP AlphaServer SC System Administration Guide* for more information about the `sra -c|cl|m|ml` command.

You are now ready to check all nodes (except node 0), as described in Section 6.3.

6.3 Check All Nodes in the HP AlphaServer SC System, Except Node 0

Check all of the nodes in the system by performing the following steps:

- Check the State of the Nodes (see Section 6.3.1 on page 6–31)
- Check the System Firmware (see Section 6.3.2 on page 6–31)

6.3.1 Check the State of the Nodes

The `sra diag` command examines the HP AlphaServer ES40 node using various SRM and RMC commands, and gathers as much data as possible about the state of the specified node(s).

Run the `sra diag` command on all nodes, as follows:

```
atlas0# sra diag -nodes 'atlas[1-1023]' -analyze no
```

Review the `/var/sra/diag/nodename.sra_diag_report` files to check for errors or warnings.

You are now ready to check the system firmware, as described in Section 6.3.2.

6.3.2 Check the System Firmware

Table 6–10 lists the minimum firmware versions supported by HP AlphaServer SC Version 2.6 (UK2) for an HP AlphaServer ES40, an HP AlphaServer ES45, and an HP AlphaServer DS20L system.

Table 6–10 Minimum System Firmware Versions

Firmware	HP AlphaServer ES40	HP AlphaServer ES45	HP AlphaServer DS20L
SRM Console	6.2-1	6.2-8	6.3-1

Check All Nodes in the HP AlphaServer SC System, Except Node 0

Table 6–10 Minimum System Firmware Versions

Firmware	HP AlphaServer ES40	HP AlphaServer ES45	HP AlphaServer DS20L
ARC Console	5.71	Not Displayed	Not Displayed
OpenVMS PALcode	1.96-103	1.96-39	1.90.71
Tru64 UNIX PALcode	1.90-104	1.90-30	1.86.68
Serial ROM	2.12-F	2.20-F	Not Displayed
RMC ROM	1.0	1.0	Not Displayed
RMC Flash ROM	2.5	1.9	Not Displayed

Check that all nodes in your system meet these minimum system firmware requirements, by performing the following tasks:

1. Run the `show config` command on all nodes and save the output to file, as follows:

```
atlas0# sra command -nodes all -command 'show config' | tee /tmp/fw.txt
```


(press Return when prompted for the root password)
2. View the SRM console firmware version:

```
atlas0# grep -i 'SRM Console' /tmp/fw.txt
```
3. View the ARC console firmware version:

```
atlas0# grep -i 'ARC Console' /tmp/fw.txt
```
4. View the UNIX PALcode firmware version:

```
atlas0# grep -i 'PALcode' /tmp/fw.txt
```
5. View the Serial ROM firmware version:

```
atlas0# grep -i 'Serial Rom' /tmp/fw.txt
```
6. View the RMC ROM firmware version:

```
atlas0# grep -i 'RMC Rom' /tmp/fw.txt
```
7. View the RMC Flash ROM firmware version:

```
atlas0# grep -i 'RMC Flash Rom' /tmp/fw.txt
```
8. Delete the output file:

```
atlas0# rm /tmp/fw.txt
```

If necessary, update the system firmware, as described in Chapter 21 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to configure and diagnose the HP AlphaServer SC Interconnect, as described in Section 6.4.

6.4 Configure and Diagnose the HP AlphaServer SC Interconnect

It is **vital** that these sections be performed at exactly this point in the installation.

If you neglect to perform these steps, then the remainder of the installation will have problems and there will be no way to diagnose the root cause. Rather than duplicate the text from the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*, this section contains a number of references to the relevant sections of that guide. Please be careful not to miss any steps as you refer to the sections. All steps are mandatory for successful operation of the diagnostic tools.

Note:

As you follow the steps in the instructions in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*, please be aware that the instructions within that manual assume that all nodes are already installed and booted. However, from the perspective of this point in this document, this is not actually the case, and instead only node 0 is installed.

Examples are situations where you are prompted to restart daemons on all nodes, or where you are asked to start swmsvr daemons on nodes 0/1 for directly-connected switches. As such, you should follow the steps in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual* in so far as is possible at this point in the installation sequence, and skip over the steps that obviously require further nodes to be installed and booted. Later, when the nodes in question are installed and booted, the diagnostics will operate correctly, and you can elect to run the diagnostics again as a confidence-building measure.

6.4.1 Upgrading the HP AlphaServer SC Interconnect Control Processor Software

Note:

This section does not apply to HP AlphaServer SC 16-port switches or old-type HP AlphaServer SC 128-way switches. Such switches are known as directly-connected switches, and do not have any firmware. This section only applies to the new-type HP AlphaServer SC 128-way switches containing an HP AlphaServer SC Interconnect control processor.

Configure and Diagnose the HP AlphaServer SC Interconnect

In HP AlphaServer SC Version 2.6 (UK2), the software (also known as firmware) running on the HP AlphaServer SC Interconnect Control Processors has changed, and must be upgraded. The new software image is packaged in the SWM subset, which is part of the HP AlphaServer SC Version 2.6 (UK2) System Software.

You should follow the steps in Section 7.2, Upgrading the HP AlphaServer SC Interconnect Control Processor Software and all subsections, of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

6.4.2 Creating an Interconnect Configuration Using SC Viewer

To receive accurate status and diagnostic information about the HP AlphaServer SC Interconnect, it is mandatory to create an Interconnect Configuration using SC Viewer. This will allow the switch managers to report anomalies in the network and it will allow the command line diagnostic tools to make pass/fail decisions based on the configured size of the network versus the detected size of the network. Once configured, SC Viewer will also allow the user to view the interconnect configuration and events which is very important in identifying problem locations.

You should follow Section 7.3, Creating an Interconnect Configuration Using SC Viewer and all subsections, of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*.

Note:

For systems with directly connected switches, you can simply skip over those instructions in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual* that assume that nodes 0/1 are already installed and booted (for example, as described in step 3 of Section 7.3.2.2, Switch Has an HP AlphaServer SC Interconnect Control Processor).

6.4.3 Confirming the Operation of the HP AlphaServer SC Interconnect

At this stage, the database and nodes will be in suitable condition to perform an initial sweep of diagnostics.

The following diagnostics should be performed:

- Assign a common clock source per rail (see Section 7.4.1 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Verify the clock frequency on each switch module (see Section 7.4.2 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

- Use `elanenvtest` to verify the switch modules (see Section 7.4.3 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

Note:

For directly-connected switches, please be aware that `elanenvtest` will only succeed once nodes 0/1 are installed and booted with their `swmserver` daemons running.

- Show links in reset (see Section 7.4.4 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

Note:

For directly-connected switches, please be aware that the show links in reset examination will only succeed once nodes 0/1 are installed and booted with their `swmserver` daemons running.

- Run `elanpcitest` on all nodes (see Section 7.4.5 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Run `elanlinktest` on all nodes (see Section 7.4.6 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Run `riscabletest` on each node-level switch (see Section 7.4.7 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

The remaining diagnostics will be performed once the system nodes are installed (see Chapter 8).

You are now ready to set up the SC Monitor System, as described in Section 6.5.

6.5 Set Up the SC Monitor System

After completing the installation steps in Section 6.2, the SC Database is populated with default entries. As there are many existing configurations, it is possible that the database may not reflect entirely the system configuration. For example, it is possible that your system may have more terminal servers than are added by default, or more switches.

Some devices, like HSG and SAN Appliance, are not added automatically in the database as monitored devices at install time, so it is necessary to add these devices if you want to monitor the device status.

Set Up the SC Monitor System

Setting Up the SC Monitor system involves the following tasks:

- Add a SAN Appliance as Monitored Devices (see Section 6.5.1 on page 6–36)
- Change Terminal Servers Monitoring Distribution (see Section 6.5.2 on page 6–36)
- Change Extreme Switches Monitoring Distribution (see Section 6.5.3 on page 6–37)
- Activate the Changes (see Section 6.5.4 on page 6–38)

6.5.1 Add a SAN Appliance as Monitored Devices

The `sra setup` command does not add any entry for SAN Appliance devices. To monitor this type of device, add the device manually as a monitored device by using the `scmonmgr add` command. The command syntax is as follows:

```
scmonmgr add -c appliance -o object -i ip-addr -s server [-r rack -u unit]
```

For example:

```
atlas0# scmonmgr add -c appliance -o sanapp0 -i 16.21.24.13 -s atlas0
```

This command adds the SAN Appliance with IP 16.21.24.13 as a monitored device with name `sanapp0`. The node that will perform the monitoring functions will be `atlasms`.

6.5.2 Change Terminal Servers Monitoring Distribution

Immediately after installation (when `sra setup` has completed), the terminal servers specified are added as monitored devices. All terminal servers are initially monitored by the management server, if there is a management server, or by the first domain otherwise.

To see the initial distribution for terminal server monitoring functions, run the following command:

```
atlas0# scmonmgr distribution -c tserver
```

Example (server `atlasms` monitors all 8 terminal servers from the system):

```
atlas0# scmonmgr distribution -c tserver
Class: tserver
Server atlasms monitors: atlas-tc[1-8]
atlas0#
```

To add additional terminal servers to be monitored by the system, use the `scmonmgr add` command.

Example:

```
atlas0# scmonmgr add -c tserver -o atlas-tc9 -s atlas0
```

Note:

A terminal server cannot be added if there is no entry for it in the `/etc/hosts` file. The system needs to know the IP address of the terminal server.

To change the distribution of the monitoring for terminal servers, use the `scmonmgr move` command as follows:

```
atlas0# scmonmgr move -o atlas-tc2 -s atlas10
```

After this `atlas10` will perform the monitoring for `atlas-tc2`.

6.5.3 Change Extreme Switches Monitoring Distribution

The `sra setup` process creates default entries for Extreme switches in the SC database using the following rules:

- If the number of nodes in the system is less than or equal to 16, a single Extreme switch of type Summit24 will be added as a monitored device.
- If the number of nodes is between 16 and 128 (including 128), the setup process adds three Summit48 switches for monitoring.
- If the number of nodes is greater than 128, the number of switches added for monitoring is calculated using the following expression: $\{number\ of\ nodes\ from\ system / 32\}$. So for a 255 nodes system, we will have $255/32 = 7$ Summit48 switches added by default as monitored devices.

This model was established based on the standard configuration of the systems, but in real life it is possible that a system may not respect these rules and a post-install customization will be necessary.

If necessary, you can add/remove Extreme switches from the monitoring system by using `scmonmgr` command.

For example:

1. To add an Extreme switch as a monitored device

```
atlas0# scmonmgr add -c extreme -o extreme10 -i 10.128.103.10 -s atlas0
```
2. To remove an Extreme switch as a monitored device

```
atlas0# scmonmgr remove -o extreme10
```

The monitoring of all Extreme switches is initially set to be performed by the management server if the system has one, or by the first domain otherwise. But this can be changed using the `scmonmgr move` command.

Assign Cabinets in the SC Database

6.5.4 Activate the Changes

After you have finished all the tasks required to change/customize the monitoring and distribution of the monitoring for a system, it is necessary to activate these changes by sending a `reload` command to the monitoring daemons running on all nodes of the system.

Before doing this, you can analyze and display the distribution of the monitoring for the entire system by using the following command:

```
atlas0# scmonmgr distribution
```

If you are satisfied, send the `reload` command to the daemons on each domain. The monitoring daemon has a reload mechanism that works on domain level. If you send a `reload` command to a daemon (node) from a domain this will trigger the reload for all daemons from that domain.

If no domains have been installed, then you can skip the reload process for all domains. However, if some domains are already installed, then you can start the reload process for all the daemons on all domains in the system by running the following command:

```
atlas0# scrun -d all '/sbin/init.d/scmon reload'
```

This will send the reload signal (`SIGHUP`) to one node from each domain, and will trigger the reload procedure on all members from that domain.

For more information on SC Monitor management, refer to Chapter 27 of the *HP AlphaServer SC System Administration Guide*.

You are now ready to assign cabinets in the SC Database, as described in Section 6.6.

6.6 Assign Cabinets in the SC Database

The *Physical* view in the `scviewer` needs to know about how your HP AlphaServer SC system is packaged, which cabinets hold nodes, terminal servers, and other pieces of hardware. The sysman `sc_cabinet` menu is used to load this information into the SC Database.

Note:

You can skip this step now and perform it later. However, until you complete this step, the `scviewer` program will be unable to display the physical relationships between hardware components.

Also, if you are already running `scviewer`, please close it down before proceeding, as the *Physical* view will not correctly reflect the cabinet changes you are making until `scviewer` is restarted.

6.6.1 Load the Physical Relationships

To load the physical relationships, you should follow these steps:

1. Create a description of cabinets in the SC Database.

Each cabinet (or rack) has the following data associated with it:

- Number. Each cabinet has a unique number (starting at 1).
 - Label. In this release, the label can be blank or the same as the number. In a later release, you can assign a meaningful label to each cabinet.
 - Area. This allows you to group cabinets into different rooms or different areas within a room. A small system would probably have one area - a large system might have several. The area concept is used by the scviewer when displaying cabinets, so it helps if an area is smaller than 20 rows of 30 cabinets (these numbers are guidelines, you can have larger numbers).
 - Position within area. The position of a cabinet within an area is specified by a row and column number. If you drew a plan of your cabinets (as scviewer does), row 1, column 1 is on the top left. Row numbers increase as you move downwards and column numbers increase as you move to the right.
2. You can use the sysman *sc_cabinet* menu to create the cabinets in an area as follows:
 - Start the sysman *sc_cabinet* menu (as root):
`atlas0# sysman sc_cabinet`
 - Select the Create block of cabinets... option
 - Specify the Area Name, First cabinet number, the label style and the number of rows and columns you want.
 - Click on Next, and confirm your choices by clicking on OK. A block of cabinets in this area is now created in the SC Database.
 3. You can repeat these steps several times for different areas or for different blocks of cabinets.

6.6.2 Populate Cabinets with Nodes

You can populate a block of cabinets with a range of nodes as follows:

1. Start the sysman *sc_cabinet* menu (as root):
`atlas0# sysman sc_cabinet`
2. Select the Assign nodes... option
3. Enter the range of nodes that are to be placed into the cabinets.

Assign Cabinets in the SC Database

Note:

The first node number should be 0 if the first node name is `atlas0`.

4. Enter the range of cabinet numbers where these nodes are located
 - Nodes can be placed top to bottom (the first node is on top, last on bottom) or vice versa.
 - Nodes can be placed into the cabinet range first to last (the first node is placed into the first cabinet, the last node is placed into the last cabinet) or vice versa.
5. Click OK to confirm your choices. The SC Database is updated to reflect this node placement.
6. You can repeat these steps several times for different ranges of nodes.

6.6.3 Populate Cabinets with Other Hardware

You can populate cabinets with other hardware (terminal servers, and so on) as follows:

Note:

If you wish to modify the properties of one or more existing cabinets (for example, to contain additional hardware) select the cabinet, and click the Modify option, to edit the contents as described in Step 5.

1. Start the sysman `sc_cabinet` menu (as `root`):
`atlas0# sysman sc_cabinet`
2. Select the Cabinets... option.
This shows a menu where all existing nodes are listed.
3. Select Add.. to add a new cabinet.
4. Specify a unique number, label, area, row and column. Click OK to create the cabinet in the SC Database.
5. The Edit a cabinet dialog box is displayed. This allows you to place hardware components into the cabinet. A list on the right shows the Unassigned Objects—that is the hardware components that have no record of their location in the SC Database. The list does not show nodes—instead it shows other types of hardware component (terminal server, Extreme Switch, and so on).
6. To place a hardware component in the cabinet, select the component and click on the Move button. This places the component into the cabinet—so it appears in the Cabinet Contents list on the left.

7. You can reorder components in the cabinet by selecting the component and clicking on the Up and Down buttons.
8. To confirm your changes, click OK. The SC Database is updated.

The Cabinets... option also allows you to modify or delete existing cabinets.

When you have finished modifying the cabinet data in the SC Database, you can review the data using the Make report... option. This menu allows you to specify a file where a textual description of the data is saved. The data is saved in a format that allows subsequent machine processing rather than being human readable.

You are now ready to review the SC Database Settings, as described in Section 6.7.

6.7 Review the SC Database Disk Settings

The `sra setup` command gathers information on disk usage. You should review this information now, in particular the Disk Location Identifiers assigned to the cluster and generic boot disks. When using RAID, the UNIX device name assigned to each of the RAID disks is not guaranteed to be in ascending order. For example `disk3` (assuming 3 local disks), typically associated with the cluster disk, will not necessarily be the RAID disk configured to be the cluster disk.

Note:

For SC20, the location ID for the UNIX disk should also be reviewed and set appropriately.

To map a UNIX device name to a RAID disk, use the RAID disk identification label assigned in Table 3–12.

For example the cluster disk for `atlasD0` has an identification label 'IDENTIFIER=1' and the generic boot disk for `atlasD0` has an identification label 'IDENTIFIER=2'. The Disk Location Identifiers in the SC database should be updated to reflect the identifiers used. Use the `sra edit` command to update the identifiers assigned to each cluster as follows:

```
atlas0# sra edit
sra> sys
sys> edit cluster atlasD0
```

Id	Description	Value

[0]	Cluster name	atlasD0
[1]	Cluster alias IP address	site-specific
[2]	Domain Type	fs
[3]	First node in the cluster	0
[4]	I18n partition device name	

Review the SC Database Disk Settings

```
[5 ] SRA Daemon Port Number          6600
[6 ] File Serving Partition           0
[7 ] Number of Cluster IC Rails       1
[8 ] Image Role                       cluster
[9 ] Image name                       first
[10 ] UNIX device name                dsk6
[11 ] SRM device name
[12 ] Disk Location (Identifier)       IDENTIFIER=5
[14 ] root partition size (%)         5
[15 ] root partition                  b
[16 ] usr partition size (%)          50
[17 ] usr partition                   g
[18 ] var partition size (%)          45
[19 ] var partition                   h
[27 ] Image Role                      gen_boot
[28 ] Image name                      first
[29 ] UNIX device name                dsk5
[30 ] SRM device name
[31 ] Disk Location (Identifier)       IDENTIFIER=6
[33 ] root partition size (%)
[34 ] root partition
[35 ] usr partition size (%)
[36 ] usr partition
[37 ] var partition size (%)
[38 ] var partition
[46 ] Image Role                      unix
[47 ] Image name                      first
[48 ] UNIX device name                dsk7
[49 ] SRM device name
[50 ] Disk Location (Identifier)
[52 ] root partition size (%)         10
[53 ] root partition                  a
[54 ] usr partition size (%)          35
[55 ] usr partition                   g
[56 ] var partition size (%)          35
[57 ] var partition                   h
```

Select attributes to edit, q to quit
eg. 1-5 10 15

```
edit? 12 31
Disk Location (Identifier)             [IDENTIFIER=5]
new value? IDENTIFIER=1
Disk Location (Identifier)             [IDENTIFIER=6]
new value? IDENTIFIER=2

Disk Location (Identifier)             [IDENTIFIER=1]
Disk Location (Identifier)             [IDENTIFIER=2]
Correct? [y|n] y
sys>
```

Transform Node 0 into a Single Node Domain

When using `sra edit` for larger systems, it is recommended you take the following steps:

1. Create a file in the following format:

```
domain_name disk image identifier
```

where:

`domain_name` is the default cluster alias e.g. `atlasD0`

`disk` is typically either `cluster` or `gen_boot`

`image` is the image name: `first` or `second`

`identifier` is the disk label in the form `IDENTIFIER=1`

The following is a sample file:

```
atlasD0 cluster first IDENTIFIER=1
atlasD0 gen_boot first IDENTIFIER=2
atlasD1 cluster first IDENTIFIER=4
atlasD1 gen_boot first IDENTIFIER=5
atlasD2 cluster first IDENTIFIER=7
atlasD2 gen_boot first IDENTIFIER=8
atlasD3 cluster first IDENTIFIER=10
atlasD3 gen_boot first IDENTIFIER=11
```

2. Using `sra edit` load this data into the SC database as follows (assuming the file created is `/var/sra/disk_id`):

```
sra> sys
sys> update
update hosts
update cmf
update ris <nodes>
update ds <nodes>
update diskid <filename>
sys> update diskid /var/sra/disk_id
Disk Location Identifiers loaded successfully
sys>
```

6.8 Transform Node 0 into a Single Node Domain

In Section 6.1, you installed the HP AlphaServer SC system software on Node 0. This section describes how to transform Node 0 into a single node domain using the `sra install` command. You can later install, clusterize, and add members on each domain, as described in Chapter 7.

Create the first domain member by running the following command on Node 0 (as the `root` user where `atlas` is an example system name):

```
atlas0# sra install -nodes atlas0
```

This command performs the following tasks:

- a. Partitions the boot disk, and the alternate boot disk (if specified), using the `createbootlabel` command.

Configure Out All Nodes During Installation

- b. Partitions the cluster disk, and the backup cluster disk (if specified), using the `createclusterlabel` command.
- c. Creates the first domain member.

The `sra install` command performs the following tasks:

- Checking
- Creating disk labels
- Creating AdvFS domains
- Populating the cluster `/`, `/usr`, and `/var`¹
- Creating CDSLs
- Modifying configuration files
- Building the kernel

To display information on the `sra` commands, use the following command:

```
atlasms# sra command_info
```

Note:

After performing the above steps, the node automatically reboots as a single node domain.

6.9 Configure Out All Nodes During Installation

While the installation is progressing, it is necessary to configure out all nodes. If the nodes are left configured in, then there will be many unwanted events reporting the fact that the nodes are not responding. Configuring out the nodes will reduce the load on the `mSQL` daemon by the RMS `mmanager` and `eventmgr`. You should configure out the nodes as follows:

```
atlas0# rcontrol configure out nodes "atlas[0-1023]"
```

In Configure the RMS Database (see Section 8.6 on page 8–10) you will have the opportunity to create the desired partition configuration and to configure in the nodes again.

You are now ready to run the Interconnect tests on Node 0 as described in Section 6.10.

6.10 Run the HP AlphaServer SC Interconnect Tests on Node 0

In Section 6.4 on page 6–33, you ran the HP AlphaServer SC Interconnect tests on all nodes except Node 0. Now run the HP AlphaServer SC Interconnect tests on Node 0, as follows:

-
1. This step may take some time (about 20 minutes); the blinking cursor indicates that the command is running and has not hung.

Run the HP AlphaServer SC Interconnect Tests on Node 0

1. Change to the directory in which the diagnostic scripts are stored, as follows:
`atlas0# cd /usr/opt/qswdiags/bin`
2. Run the `elanpcitest` script, to check that Node 0 can access its HP AlphaServer SC Elan adapter card across its PCI bus, as follows:
`atlas0# ./elanpcitest`
 - If the `elanpcitest` script completes successfully, the output is as follows:
`elanpcitest: accessing Network Adapter memory, please wait`
`elanpcitest: test completed successfully`
 - If the `elanpcitest` script fails, the output is as follows:
`elanpcitest: accessing QM401 Network Adapter memory, please wait`
`elanpcitest: test failed`
Check that the HP AlphaServer SC Elan adapter card is correctly seated in Node 0.
3. Run the `elanlinktest` script, to check that the Node 0 HP AlphaServer SC Elan adapter card can access the HP AlphaServer SC Interconnect switch, as follows:
`atlas0# ./elanlinktest`
 - If the `elanlinktest` script completes successfully, the output is as follows:
`elanlinktest: testing connection from network adapter to network switch card`
`elanlinktest: test completed successfully`
 - If the `elanlinktest` script fails, the output is as follows:
`elanlinktest: testing connection from network adapter to network switch card`
`elanlinktest: test failed`
If the `elanlinktest` script fails, check the following components:
 - The cable connecting Node 0 to the HP AlphaServer SC Interconnect switch
 - The HP AlphaServer SC Elan adapter card in Node 0
4. Run the `cabletest` on Node 0 as follows:
`atlas0# cd /usr/opt/qswdiags/bin`
Run the test directly on the node as follows:
`atlas0# ./cabletest`

The HP AlphaServer SC system is now in the following state:

- The Tru64 UNIX operating system has been installed and configured on Node 0.
- The HP AlphaServer SC software has been installed on Node 0.
- The SC database has been created on Node 0.
- Node 0 has been clusterized and is running as a single-node domain.
- All other nodes are uninstalled, and at the SRM prompt. Check that all domain members are at the SRM prompt:
`atlas0# sra info -nodes 'atlas[1-1023]'`
Nodes not at the SRM prompt should be shut down or halted.

Run the HP AlphaServer SC Interconnect Tests on Node 0

You are now ready to build the domains, as described in Chapter 7 (*Building the Domains*).

Building the Domains

By following the instructions in earlier chapters, you have manually installed and configured one host: either the management server (see Chapter 5) or Node 0 (see Chapter 6). This chapter describes how to use this manually-installed host to build the domains.

Note:

Do not proceed with this chapter until you have completed the manual installation and configuration of the first host.

The information in this chapter is organized as follows:

- Understanding the Automated Installation Process (see Section 7.1 on page 7–1)
- Review the SC Database System Settings (see Section 7.2 on page 7–14)
- Add sysconfigtab Parameters (see Section 7.3 on page 7–23)
- Create the Domains (see Section 7.4 on page 7–23)
- Boot the System (see Section 7.5 on page 7–25)
- Complete the Setup of the Domains (see Section 7.6 on page 7–25)

7.1 Understanding the Automated Installation Process

The information in this section is organized as follows:

- SC Database Installation Tables (see Section 7.1.1 on page 7–2)
- The sra install Command (see Section 7.1.2 on page 7–2)
- The Installation Daemon (see Section 7.1.3 on page 7–2)
- Installation States (see Section 7.1.4 on page 7–3)
- Monitoring the Installation State (see Section 7.1.5 on page 7–4)

Understanding the Automated Installation Process

7.1.1 SC Database Installation Tables

The installation steps performed are determined by the interaction between data in two database tables:

- `sc_command` Table - contains the following domain specific information:
 - The command performed (for example `install`, `boot`, `shutdown`)
 - The command status and associated message (`failed`, `completed`)
 - The desired installation state of the domain (see Table 7–1)
- `sc_nodes` Table - contains an entry for each node, and each entry has a `current_state` and a `desired_state`.

The database entries and values assigned capture the state of the system as it transitions from Uninstalled to Member_Added.

7.1.2 The `sra install` Command

The `sra install` command automates the installation of domains. When you have manually installed one host (either the management server, if used, or Node 0), you can build all of the domains by running a single `sra install` command. You can also run the `sra install` command to complete the installation of a partially built domain. The `sra install` command determines the steps required to change each domain from its current installation state to the desired installation state, and then performs the identified steps.

7.1.3 The Installation Daemon

The `sra install` command works with the installation daemon (`srad`). One `srad` daemon runs on the management server (when the system has a management server), and one `srad` daemon runs on each domain.

The `srad` daemon on the management server performs the following installation steps on the first node of each domain:

1. RIS-installs the Tru64 UNIX operating system.
2. Configures the Tru64 UNIX operating system.
3. Installs the Tru64 UNIX patch software.
4. Installs the HP AlphaServer SC software.
5. Installs the HP AlphaServer SC patch software.
6. Installs the New Hardware Delivery (NHD) software.
7. Starts the `srad` daemon on the first node of each domain.

Understanding the Automated Installation Process

The `srad` daemon runs on the first node of each domain and completes the installation as follows:

1. Runs the `clu_create` command to clusterize the first node of the domain.
2. Performs the following steps on each domain member:
 - a. Runs the `clu_add_member` command to add the node to the domain.
 - b. RIS-boots the node and downloads the boot partitions from the first node of the domain.
 - c. Boots `genvmunix`
 - d. Shuts the member down to the SRM prompt.

Note:

The log files for the `srad` daemon is located in the directory: `/var/sra/adm/log/srad`.

7.1.4 Installation States

The current installation state of a domain will change during the installation process. During a complete installation, a domain will change through several installation states from Uninstalled to `Member_Added`, as described in Table 7–1.

Table 7–1 Installation States

Installation State	Description
Uninstalled	No Tru64 UNIX or HP AlphaServer SC software has been installed on the domain
UNIX_Installed	The Tru64 UNIX operating system has been installed on the first node of the domain
UNIX_Config	The Tru64 UNIX operating system has been configured on the first node of the domain
UNIX_Patched	The Tru64 UNIX operating system patch software has been installed on the first node of the domain
SC_Installed	The HP AlphaServer SC software has been installed on the first node of the domain
SC_Patched	The HP AlphaServer SC patch software has been installed on the first node of the domain

Understanding the Automated Installation Process

Table 7–1 Installation States

Installation State	Description
NHD_Installed	The New Hardware Delivery software has been installed on the first node of the domain
CLU_Create	The first node of the domain has been clusterized
CLU_Added	The remaining nodes have been added to the domain
Bootp_Loaded	The boot partitions have been downloaded to each node in the domain, from its first node
Member_Added	All members have been added (to the domain), booted, and shut down to the SRM prompt

The installation process is incremental. A domain cannot change to a particular installation state until it has passed through all of the previous installation states. If any step fails, you can rerun the `sra install` command and the `srad` daemon will restart the installation process from the last successful step. The `srad` daemon determines the installation steps required for each domain, based on the current installation state and desired installation state of the domain.

Note:

The installation process assumes that each domain may be in a different installation state.

To partially build a system, you can use the `sra install` command with the `-endstate` option. Table 7–1 lists the installation states that can be specified.

Note:

For more information about the typical actions that happen during each installation state, and where you can access log files to determine the cause of an error, see Section 11.3 on page 11–5.

7.1.5 Monitoring the Installation State

This section contains the following information:

- Displaying Information on `sra` Commands (see Section 7.1.5.1 on page 7–5)

Understanding the Automated Installation Process

- Querying Node Installation Status (see Section 7.1.5.2 on page 7–5)
- Starting the Monitor Utility (see Section 7.1.5.3 on page 7–5)
- Filtering Views of Installation Progress (see Section 7.1.5.4 on page 7–11)

7.1.5.1 Displaying Information on sra Commands

Use the `sra command_info` command to display information on `sra` commands. Use the `-states` flag to specify which command states to display. The possible states are Allocated, Unallocated, Success, Error, and Abort.

7.1.5.2 Querying Node Installation Status

Use the `sra install_info` command to monitor the status of installation on specific nodes. The command provides information similar to the `sramon` utility (see Section 7.1.5.3). The syntax is as follows:

```
atlasms# sra install_info
```

```
sra install_info {-nodes <nodes> | -domains <domains> | -members <members>}  
[...] [-display <yes|no>]
```

The following shows sample output from the `sra install_info` command on a specified domain:

```
atlasms# sra install_info -dom 0
```

```
=====
Node Name      Current State   Desired State   Commands Status
-----
atlas0         Member_Added    Member_Added    Finished
atlas1         Member_Added    Member_Added    System-Up
atlas2         Member_Added    Member_Added    System-Up
atlas3         Member_Added    Member_Added    System-Up
atlas4         CLU_Create      Member_Added    Error: failed to RIS boot
atlas5         Member_Added    Member_Added    System-Up
atlas6         Member_Added    Member_Added    System-Up
atlas7         Member_Added    Member_Added    System-Up
=====
CLU_Create: atlas4
Member Added: atlas[0-3,5-7]
=====
#
```

7.1.5.3 Starting the Monitor Utility

The `sramon` utility is used to monitor the progress of an installation on the system of domains.

Start the `sramon` utility by running the following command:

```
atlasms# sramon &
```

Understanding the Automated Installation Process

[1] 6899
atlasms #

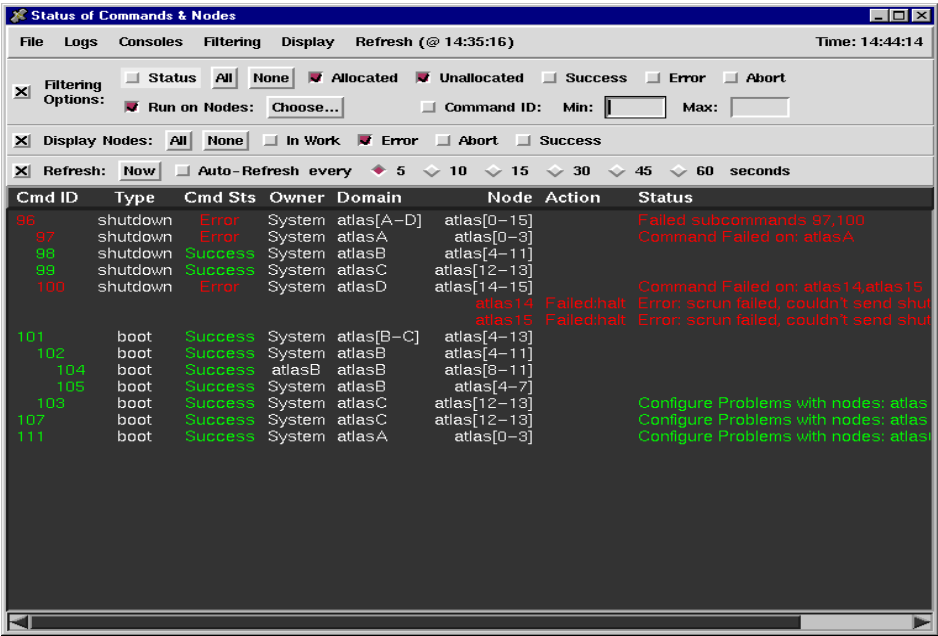


Figure 7–1 Status of Commands and Nodes

Table 7–2 describes the menu items:

Table 7–2 Status of Commands and Nodes Menu Items

Menu Item	Description
File - Exit	Used to exit the monitor
Logs	Used to view the <code>srtd</code> log for selected nodes
Consoles	Used to view console output for selected nodes
Filtering	Used to select filtering options
Display	Used to select nodes to display
Refresh	Used to specify auto-refresh settings or refresh current display

Note:

To exit an `sra` command from within the monitor, select the Dismiss option from the File menu.

The menu bar contains the following additional items (when the menu bar is expanded):

- Filtering Options allowing you to filter by:
 - Command Status
 - Run on Nodes (atlas0 to atlas7)
 - Command ID
- Display Nodes
 - All, None, In Work, Error, Abort, Success
- Refresh Options and Details allowing you to:
 - View current time and last refresh time
 - Specify auto-refresh or immediate refresh

Table 7–3 describes the information displayed in the main window (Figure 7–1):

Table 7–3 Status of Commands and Nodes Window

Column	Description
Cmd ID	The command ID / number.
Type	The command type (<code>install</code>).
Cmd Sts	The status of the command (allocated, unallocated, success, error, abort).
Owner	The owner of the command (<code>system</code>)
Domain	The domain where the activity is being monitored (<code>atlasD0</code>)
Node	The node in the domain where the activity is being monitored
Action	The action being performed on the node
Status	The current installation status on the node

Log Menu Items

Use the Log menu items to display log file details of installation on the selected domain.

Understanding the Automated Installation Process


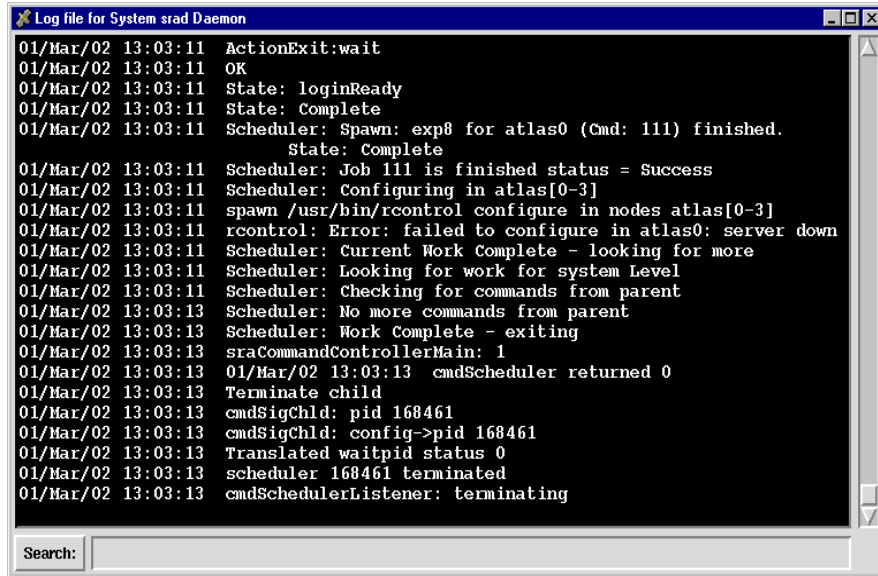


Figure 7–2 Log Menu Items

Select an option from the menu to open a new screen that displays the log file detail for the selected system or domain (atlasD0, atlasD1, atlasD2, atlasD3). Select the Systems option to display the `srad` log from either the management server or from Node 0.

Understanding the Automated Installation Process



```
01/Mar/02 13:03:11 ActionExit:wait
01/Mar/02 13:03:11 OK
01/Mar/02 13:03:11 State: loginReady
01/Mar/02 13:03:11 State: Complete
01/Mar/02 13:03:11 Scheduler: Spawn: exp8 for atlas0 (Cmd: 111) finished.
                        State: Complete
01/Mar/02 13:03:11 Scheduler: Job 111 is finished status = Success
01/Mar/02 13:03:11 Scheduler: Configuring in atlas[0-3]
01/Mar/02 13:03:11 spawn /usr/bin/rcontrol configure in nodes atlas[0-3]
01/Mar/02 13:03:11 rcontrol: Error: failed to configure in atlas0: server down
01/Mar/02 13:03:11 Scheduler: Current Work Complete - looking for more
01/Mar/02 13:03:11 Scheduler: Looking for work for system Level
01/Mar/02 13:03:11 Scheduler: Checking for commands from parent
01/Mar/02 13:03:13 Scheduler: No more commands from parent
01/Mar/02 13:03:13 Scheduler: Work Complete - exiting
01/Mar/02 13:03:13 sraCommandControllerMain: 1
01/Mar/02 13:03:13 01/Mar/02 13:03:13 cmdScheduler returned 0
01/Mar/02 13:03:13 Terminate child
01/Mar/02 13:03:13 cmdSigChld: pid 168461
01/Mar/02 13:03:13 cmdSigChld: config->pid 168461
01/Mar/02 13:03:13 Translated waitpid status 0
01/Mar/02 13:03:13 scheduler 168461 terminated
01/Mar/02 13:03:13 cmdSchedulerListener: terminating
```

Figure 7–3 Sample Log File

Console Menu Items

This section describes the Console menu. Use this option to view the console output for a node. The console logs are saved in `/var/sra/log/cmfd`. The log file contains the nodes console output. Figure 7–4 shows sample console output.

Understanding the Automated Installation Process




Figure 7-4 Console Menu Items

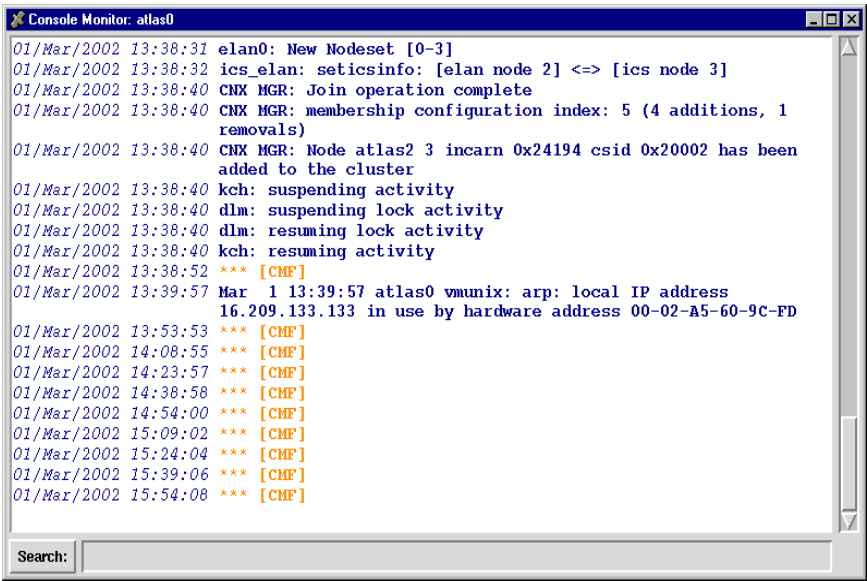


Figure 7-5 Console File Details

Understanding the Automated Installation Process

7.1.5.4 Filtering Views of Installation Progress

This section will describe the type of information displayed and how to change your view of installation progress.

Filtering Based On Status

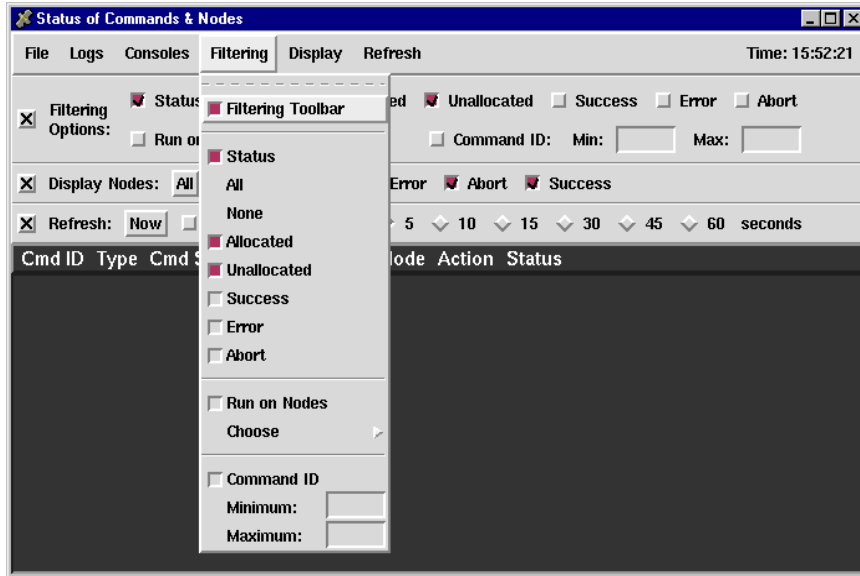


Figure 7–6 Filtering Options Based on Status

The following shows the detached external filtering toolbar.

Understanding the Automated Installation Process



Display Options

This section will describe how to change the type of information displayed.

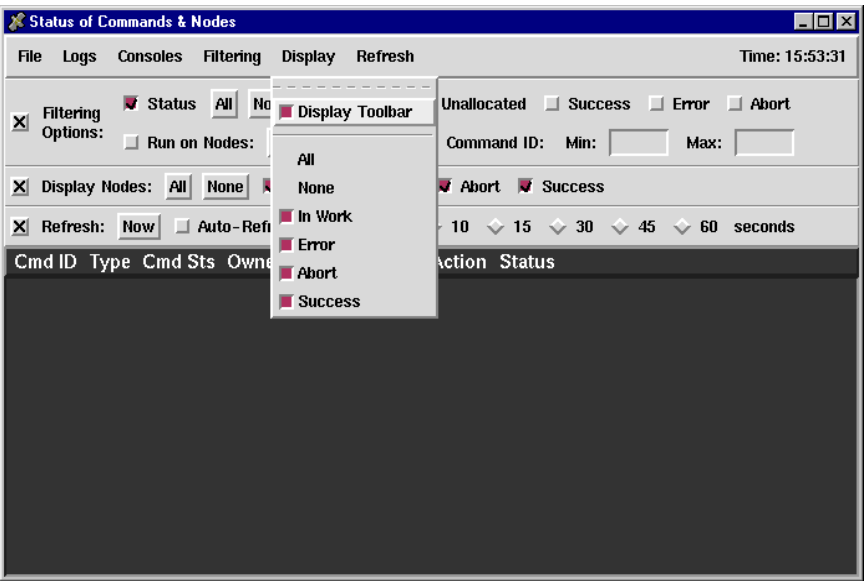
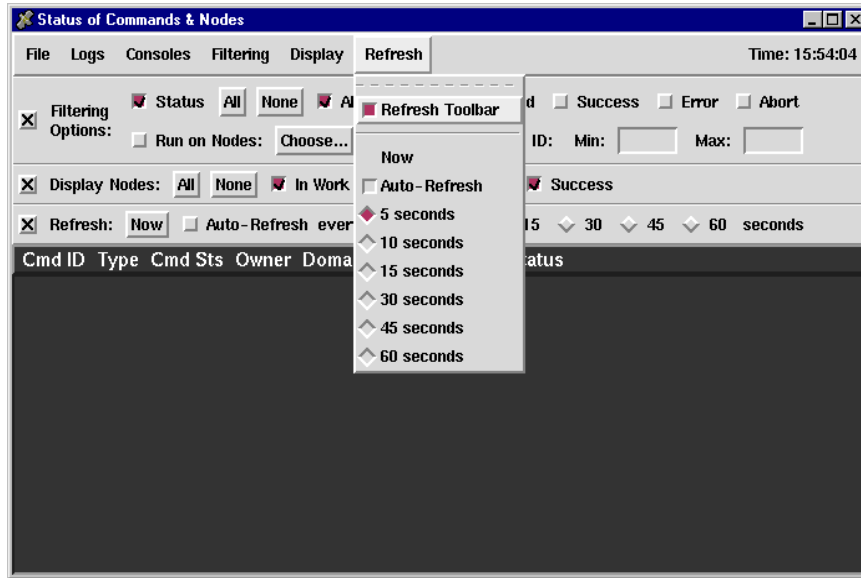


Figure 7-7 Display Options Based on Nodes

Understanding the Automated Installation Process

Refresh Options

This section describes the refresh options.



7.1.5.5 Text Based sramon Utility

The sramon utility can also be run from the command line. The following commands can be used:

```
atlasms# sramon cmd_id [-status]
```

For example:

```
atlasms# sramon 105
15:11:49 Command 105 (boot) atlasB: atlas[4-7] -- <Success>
15:11:49 Node atlas4 -- <Complete:wait> Finished
15:11:49 Node atlas5 -- <Complete:wait> Finished
15:11:49 Node atlas6 -- <Complete:wait> Finished
15:11:49 Node atlas7 -- <Complete:wait> Finished
```

Command has finished:

```
Command 105 (boot) atlasB : atlas[4-7] -- <Success>
```

```
*** Node States *** Completed: atlas[4-7]
```

```
# sramon 105 -status
Success
```

Review the SC Database System Settings

7.2 Review the SC Database System Settings

When the `sra setup` command runs for the first time, the installation process gathers information about the system and stores this information in the SC database. The installation process uses this information when installing the other nodes, so you should review the SC database system settings to ensure that the values are set correctly.

To display the SC database system settings, use the `sra edit` command. The following example shows the SC database system settings for a 128-node system called `atlas`:

```
atlas0# sra edit
sra> sys
sys> show system
```

Id	Description	Value
[0]	System name	atlas
[1]	SC database revision	2.6.0
[2]	Connect method	cmf
[3]	First DECserver IP address	site specific
[4]	First port on the terminal server	7
[5]	Hardware type	ES45
[6]	Default image	0
[7]	Number of nodes	8
[8]	Node running console logging daemon (cmfd)	atlasms
[9]	cmf home directory	/var/sra/
[10]	cmf port number	6500
[11]	cmf port number increment	2
[12]	cmf max nodes per daemon	256
[13]	cmf max daemons per host	4
[14]	Allow cmf connections from this subnet	255.255.0.0
[15]	cmf reconnect wait time (seconds)	60
[16]	cmf reconnect wait time (seconds) for failed ts	1800
[17]	Software selection	1
[18]	Software subsets	
[19]	Kernel selection	3
[20]	Kernel components	2 3 4 11
[21]	DNS Domain Name	site specific
[22]	DNS server IP list	site specific
[23]	DNS Domains Searched	
[24]	NIS server name list	site specific
[25]	NTP server name list	site specific
[26]	MAIL server name	site specific
[27]	Default Internet route IP address	site specific
[28]	Management Server name	atlasms
[29]	Management Server IP address	10.128.101.1
[30]	Use swap, tmp & local on alternate boot disk	no
[31]	SRA Daemon (srad) port number	6600
[32]	SRA Daemon Monitor host	
[33]	SRA Daemon Monitor port number	
[34]	SC Database setup and ready for use	1
[35]	IP address of First Top level switch (rail 0)	10.128.128.128
[36]	IP address of First Node level switch (rail 0)	10.128.128.1

Review the SC Database System Settings

```
[37 ] IP address of First Top level switch (rail 1)      10.128.129.128
[38 ] IP address of First Node level switch (rail 1)    10.128.129.1
[39 ] Port used to connect to the scmountd on MS        5555
[40 ] /etc/clua_default metric                          2
[41 ] Preferred Server cluster alias base address      10.128.106.1
```

sys>

Note:

The UNIX disk is no longer displayed in the system output.

To change any of these settings, use the `sra edit` command, as shown in the following sections:

- Changing the Tru64 UNIX Disk
- Configuring the Cluster Disk
- Changing the External Network Netmask
- Changing the External IP Address on Additional Nodes

Changing the Tru64 UNIX Disk

The Tru64 UNIX disk is set to `dsk2` by default. If you wish to change this value, use the `sra edit` command as shown in the following example:

```
atlasms# sra edit
sra> sys
sys> show image
valid images are [unix-first cluster-first cluster-second boot-first boot-second
gen_boot-first]1
sys> edit image unix-first
```

Id	Description	Value
[0]	Image Role	unix
[1]	Image name	first
[2]	UNIX device name	dsk2
[3]	SRM device name	dkb200
[4]	root partition size (%)	15
[5]	root partition	a
[6]	usr partition size (%)	35

-
1. `unix-first` = Tru64 UNIX disk (no backup); `cluster-first` = cluster (`/`, `/usr`, `/var`) disk; `cluster-second` = backup cluster (`/`, `/usr`, `/var`) disk; `boot-first` = primary boot disk; `boot-second` = backup boot disk; `gen_boot-first` = generic boot disk (no backup).

Review the SC Database System Settings

```
[7 ] usr partition          g
[8 ] var partition size (%) 30
[9 ] var partition          h
[10] swap partition size (%) 20
[11] swap partition          b
```

Select attributes to edit, q to quit
eg. 1-5 10 15

probe = probe for value

edit? **2**

UNIX device name [dsk2]

new value? **dsk3**

UNIX device name [dsk3]

Correct? [y|n] **y**

sys>

Note:

Do not spread the base operating system across more than one disk. Do not change the disk partition selection — use b, g, and h for swap, /usr, and /var respectively.

Configuring the Cluster Disk

The `sra setup` command gathers information on disk usage. You should review this information now, in particular the Disk Location Identifiers assigned to the cluster and generic boot disks. When using RAID, the UNIX device name assigned to each of the RAID disks is not guaranteed to be in ascending order. For example `dsk3` (assuming 3 local disks), typically associated with the cluster disk, will not necessarily be the RAID disk configured to be the cluster disk.

To map a UNIX device name to a RAID disk, use the RAID disk identification label assigned in Table 3–12.

For example, the cluster disk for `atlasD0, D1`, has an identification label 'IDENTIFIER=1' and the generic boot disk for `atlasD0, D1`, has an identification label 'IDENTIFIER=2'.

The Disk Location Identifiers in the SC database should be updated to reflect the identifiers used. Use the `sra edit` command to update the identifiers assigned to each cluster as follows:

```
atlasms# sra edit
sra> sys
sys> edit cluster atlasD0
```

Review the SC Database System Settings

Id	Description	Value
[0]	Cluster name	atlasD0
[1]	Cluster alias IP address	<i>site-specific</i>
[2]	Domain Type	fs
[3]	First node in the cluster	0
[4]	I18n partition device name	
[5]	SRA Daemon Port Number	6600
[6]	File Serving Partition	0
[7]	Number of Cluster IC Rails	2
[8]	Current Upgrade State	Pre_Upgrade
[9]	Desired Upgrade State	Pre_Upgrade
[10]	Image Role	cluster
[11]	Image name	first
[12]	UNIX device name	dsk3
[13]	SRM device name	
[14]	Disk Location (Identifier)	IDENTIFIER=5
[15]	root partition size (%)	5
[16]	root partition	b
[17]	usr partition size (%)	50
[18]	usr partition	g
[19]	var partition size (%)	45
[20]	var partition	h
[21]	Image Role	cluster
[22]	Image name	second
[23]	UNIX device name	dsk5
[24]	SRM device name	
[25]	Disk Location (Identifier)	IDENTIFIER=3
[26]	root partition size (%)	5
[27]	root partition	b
[28]	usr partition size (%)	50
[29]	usr partition	g
[30]	var partition size (%)	45
[31]	var partition	h
[32]	Image Role	gen_boot
[33]	Image name	first
[34]	UNIX device name	dsk4
[35]	SRM device name	
[36]	Disk Location (Identifier)	IDENTIFIER=6
[37]	default or not	
[38]	swap partition size (%)	30
[39]	tmp partition size (%)	35
[40]	local partition size (%)	35
[41]	Image Role	unix
[42]	Image name	first
[43]	UNIX device name	dsk2
[44]	SRM device name	
[45]	Disk Location (Identifier)	
[46]	root partition size (%)	10
[47]	root partition	a
[48]	usr partition size (%)	35
[49]	usr partition	g
[50]	var partition size (%)	35

Review the SC Database System Settings

```
[51 ] var partition                h
[52 ] swap partition size (%)      20
[53 ] swap partition              b
-----
```

Select attributes to edit, q to quit
eg. 1-5 10 15
edit? **14 36**

Disk Location (Identifier) [IDENTIFIER=5]
new value? **IDENTIFIER=1**
Disk Location (Identifier) [IDENTIFIER=6]
new value? **IDENTIFIER=2**

Disk Location (Identifier) [IDENTIFIER=1]
Disk Location (Identifier) [IDENTIFIER=2]
Correct? [y|n] **y**
sys>

When using `sra edit` for larger systems, it is recommended you take the following additional steps:

1. Create a file in the following format:

domain_name disk image identifier

where:

domain_name - is the default cluster alias e.g. atlasD0

disk - is typically either cluster or gen_boot

image - is the image name: first or second

identifier - is the disk label in the form IDENTIFIER=1

The following is a sample file:

```
atlasD0 cluster first IDENTIFIER=1
atlasD0 gen_boot first IDENTIFIER=2
atlasD1 cluster first IDENTIFIER=4
atlasD1 gen_boot first IDENTIFIER=5
atlasD2 cluster first IDENTIFIER=7
atlasD2 gen_boot first IDENTIFIER=8
atlasD3 cluster first IDENTIFIER=10
atlasD3 gen_boot first IDENTIFIER=11
```

2. Using `sra edit` load this data into the SC database as follows (assuming the file created is `/var/sra/disk_id`):

```
sra> sys
sys> update
update hosts
update cmf
update ris <nodes>
update ds <nodes>
update diskid <filename>
sys> update diskid /var/sra/disk_id
Disk Location Identifiers loaded successfully
sys>
```


Changing the External Network Netmask

The netmask of an external network is set to 255.255.255.0 by default. If you wish to change this value, use the `sra edit` command as shown in the following example:

```
atlasms# sra edit
sra> sys
sys> show ip
valid ips are [man ext ics eip]1
sys> edit ip ext
```

Id	Description	Value
[0]	Interface name	ext
[1]	Hostname suffix	-ext1
[2]	Network address (IP)	
[3]	UNIX device name	ee1
[4]	SRM device name	eib0
[5]	Netmask	255.255.255.0
[6]	Cluster Alias Metric	

```
-----

Select attributes to edit, q to quit
eg. 1-5 10 15

edit? 5
Netmask [255.255.255.0]
new value? 255.255.0.0

Netmask [255.255.0.0]
correct? [y|n] y
sys>
```

Each node's external network netmask will be affected by this change *via* the rule set.

Changing the External IP Address on Additional Nodes

The `sra install` command will automatically configure the external interfaces as required on additional nodes.

The address and UNIX device name (if different not using DE602) should be defined as follows for each node with an external interface.

If desired, the external interfaces can be configured.

1. `man` = management network; `ext` = external network; `ics` = cluster interconnect;
`eip` = system interconnect

Review the SC Database System Settings

```
atlasms# sra edit
sra> node
node> edit atlas1
```

Id		Description	Value

[0]	Hostname		atlas1
...etc...			
[66]	ip01:Hostname suffix		atlas1-ext1 *
[67]	ip01:Network address (IP)		#
[68]	ip01:UNIX device name		ee1
...etc...			

Select attributes to edit, q to quit
eg. 1-5 10 15

edit? 67 68

enter a new value, probe or auto
auto = generate value from system
probe = probe hardware for value

ip01:Network address (IP)	[]	(no value)
new value? 16.209.133.226		
ip01:UNIX device name	[ee1]	(set)
new value? alt0		
ip01:Network address (IP)	[16.209.133.226]	(new)
ip01:UNIX device name	[alt0]	(new)
correct? [y n] y		
node>		

For more information about the `sra` command, see Chapter 16 of the *HP AlphaServer SC System Administration Guide*.

Changing the Default Route

The `sra setup` command gathers the default route setting from the management server, or node0, and stores it in the SC database. The installation process uses this setting when configuring the first node of each domain. If you wish to change this value, use the `sra edit` command as follows:

```
atlasms# sra edit
sra> sys
sys> edit system
```

Review the SC Database System Settings

Id	Description	Value
[0]	System name	atlas
.	.	.
[27]	Default Internet route IP address	site specific

.

.

Select attributes to edit, q to quit
eg. 1-5 10 15

edit? **27**

Default Internet route IP address [site specific]
new value? **uu.xx.yy.zz**

Default Internet route IP address [uu.xx.yy.zz]
correct? [y|n] **y**

sys>

Changing the Node Type

The `sra setup probe` should determine any differences in hardware type. Please verify that the hardware type for the non-dominant type is correct by using the `sra edit` command as follows:

```
atlasms# sra edit
sra> node
node> edit 1
```

Id	Description	Value
[0]	Hostname	atlas1
[1]	DECserver name	atlas-tcl
[2]	DECserver internal port	2
[3]	cmf host for this node	atlasms
[4]	cmf port number for this node	6500
[5]	TruCluster memberid	2
[6]	Cluster name	atlas
[7]	Hardware address (MAC)	00-02-56-00-0C-99
[8]	Number of votes	0
[9]	Node specific image_default	0
[10]	Elan Id	*
[11]	Bootable or not	0
[12]	Hardware type	DS20L
[13]	Current Installation State	Member_Added
[14]	Desired Installation State	Member_Added
[15]	Current Installation Action	Complete:wait
[16]	Command Identifier	21
[17]	Node Status	Finished
[19]	im00:Image Role	boot

Review the SC Database System Settings

```

[20 ] im00:Image name           first
[21 ] im00:UNIX device name     dsk2
[22 ] im00:SRM device name      dqa0
[23 ] im00:Disk Location (Identifier)
[24 ] im00:default or not       yes
[31 ] im00:swap partition size (%) 30
[33 ] im00:tmp partition size (%) 35
[35 ] im00:local partition size (%) 35
[38 ] ip03:Interface name       eip
[39 ] ip03:Hostname suffix      atlas1-eip0      *
[40 ] ip03:Network address (IP) 10.64.0.2      *
[41 ] ip03:UNIX device name     eip0
[42 ] ip03:SRM device name
[43 ] ip03:Netmask              255.255.0.0
[44 ] ip03:Cluster Alias Metric 16
[46 ] ip02:Interface name       ics
[47 ] ip02:Hostname suffix      atlas1-ics0      *
[48 ] ip02:Network address (IP) 10.0.0.2      *
[49 ] ip02:UNIX device name     ics0
[50 ] ip02:SRM device name
[51 ] ip02:Netmask              255.255.255.0
[52 ] ip02:Cluster Alias Metric
[54 ] ip01:Interface name       ext
[55 ] ip01:Hostname suffix      atlas1-ext1      *
[56 ] ip01:Network address (IP) #
[57 ] ip01:UNIX device name     ee0
[58 ] ip01:SRM device name      eia0
[59 ] ip01:Netmask              255.255.255.0
[60 ] ip01:Cluster Alias Metric
[62 ] ip00:Interface name       man
[63 ] ip00:Hostname suffix      atlas1      *
[64 ] ip00:Network address (IP) 10.128.0.2      *
[65 ] ip00:UNIX device name     ee1
[66 ] ip00:SRM device name      eib0
[67 ] ip00:Netmask              255.255.0.0
[68 ] ip00:Cluster Alias Metric

```

* = default generated from system

= no default value exists

 Select attributes to edit, q to quit
 eg. 1-5 10 15

edit? 12

enter a new value, probe or auto
 auto = generate value from system
 probe = probe hardware for value

Hardware type [DS20L] (set)
 new value? **ES45**

Hardware type [ES45] (new)

```
correct? [y|n] y
node>
```

7.3 Add sysconfigtab Parameters

The Tru64 UNIX operating system includes various subsystems that are used to define or extend the kernel. Kernel variables control the behavior of these subsystems, or track subsystem statistics since boot time.

The recommended `sysconfigtab` parameters are provided in the file `/Examples/sysconfigtab` on the *HP AlphaServer SC System Software* CD-ROM. This sample `sysconfigtab` file assumes a memory size of 4GB.

You should review the recommended values and if necessary copy the file to a temporary location on your system, and modify the values. Note that the default value for the `shm_max` attribute is 4GB.

The recommended `sysconfigtab` parameters will be added automatically by `sra install` (see Section 7.4), where you will specify the temporary location for the modified file in the `sra install` command line.

7.4 Create the Domains

Note:

For a clustered management server, ensure that the non-lead member of the clustered management server is shut down as described in Section 5.1.7.11.

To create the domains, perform the following steps on the management server, if used, or on Node 0 (if not using a management server).

1. Insert the *HP AlphaServer SC System Software* CD-ROM in the disk drive.
2. If it is not already mounted, mount the CD-ROM (see Section 5.1.9, step 2 on page 5–32).
3. Unpack the latest Tru64 UNIX patch kit tar file in the `/patches` directory.

Note:

For HP AlphaServer SC Version 2.6 (UK2), you should load Tru64 UNIX Version 5.1B-3 (also known as Tru64 UNIX Version 5.1B Patch Kit 5) only.

The operating system patch software kit (`T64V51BB26AS0005-20050502.tar`) is available from the following location:
<<http://www.itrc.hp.com/>>

Create the Domains

or from your local HP support representative.

4. Install, clusterize, and add members to each domain, as follows:

```
atlasms# sra install -domains all -unixpatch /patches/patch_kit -sckit /  
cdrom/kits -sysconfig /cdrom/Examples/sysconfigtab
```

where

- `-domains all` specifies that the software should be installed on all of the domains.
- `-unixpatch /patches/patch_kit` specifies the directory in which you had unpacked the Tru64 UNIX patch kit software (see step 3 above).
- `-sckit /cdrom/kits` specifies the CD-ROM mount point created earlier (see step 2 above) containing the HP AlphaServer SC software.
- `-sysconfig /cdrom/Examples/sysconfigtab` adds the specified `sysconfigtab` parameters to each node.

Note:

The `sra install` command installs all of the HP AlphaServer SC software on the specified nodes.

When this command has successfully completed, the first node in each domain is booted off the cluster disk, and all the members have been added, booted, and shut-down to the SRM prompt.

To display information on `sra` commands, use the following command:

```
atlasms# sra command_info
```

When running the `sra install` command to customize an installation, note the following:

- The `sra install` command options `-domains` and `-nodes` can be abbreviated, as in the following examples:

```
sra install -domains 0-2 -nodes 96-127  
sra install -do 0
```

- The `-nodes` option is not a qualifier for the `-domains` option.

By specifying the command: `sra install -domains atlasD0,atlasD1 -nodes atlas96` you install all nodes in domains `atlasD0` and `atlasD1`, that is nodes `atlas0-63`, and also node `atlas96`.

However, by specifying the command: `sra install -domains atlasD0 -nodes atlas0`, you install all of `atlasD0`. In this example, the option `-nodes atlas0` is redundant, and the command does not install only `atlas0`.

- The `-member` option can be used as a qualifier for the `-domains` flag. By specifying the command: `sra install -domains atlasD0,atlasD1 -member 1,2` you only install members 1,2 in domains atlasD0 and atlasD1 (that is nodes atlas0, atlas1, atlas32 and atlas33).

For more information on the `sra install` command, see Section 7.1.2 on page 7–2. For information on how to monitor the progress of the installation, see Section 7.1.5 on page 7–4.

7.5 Boot the System

After running the `sra install` command, the first member of each domain will be booted and all other members will be halted.

You can now boot the entire system with the following command:

```
atlasms# sra boot -domains all
```

Note

This command will ignore any nodes already booted.

7.6 Complete the Setup of the Domains

To complete the setup of the domains, perform the following steps:

1. Install the HP Fortran Run-Time Libraries on each domain. To do this, follow the instructions in Section 5.1.10 on each domain in the system.
2. Install any further layered products required on each domain.

Completing the Installation

When you have added all nodes to the HP AlphaServer SC system, perform the following tasks to complete the installation:

- Boot the Non-Lead Member of the Clustered Management Server (see Section 8.1 on page 8–2)
- Configure the External Network Interfaces (see Section 8.2 on page 8–2)
- Improve Cluster Availability (see Section 8.3 on page 8–3)
- Load File System Configuration Data in the SC Database (see Section 8.4 on page 8–7)
- Initial Setup of Monitoring for HSG80 RAID Systems (see Section 8.5 on page 8–8)
- Configure the RMS Database (see Section 8.6 on page 8–10)
- Provide RMS with CAA Failover Capability (see Section 8.7 on page 8–12)
- Enable CMF as a CAA Application (see Section 8.8 on page 8–15)
- Run the Example MPI Program (see Section 8.9 on page 8–17)
- Verify the HP AlphaServer SC Interconnect (see Section 8.10 on page 8–17)
- Configure LSM (see Section 8.11 on page 8–18)
- Verify Swap Mode (see Section 8.12 on page 8–18)
- Add a Second Rail to an HP AlphaServer SC System after Domain Creation (see Section 8.13 on page 8–18)

Note:

Use the checklist provided in Appendix A (if using a management server) or Appendix B (if not using a management server), to ensure that you complete all installation tasks in the correct order.

Boot the Non-Lead Member of the Clustered Management Server

If adding a management server to an HP AlphaServer SC system after domain creation, use the checklist provided in Appendix C to ensure that you complete all installation tasks in the correct order.

8.1 Boot the Non-Lead Member of the Clustered Management Server

For a clustered management server, ensure that the non-lead member of the server is booted. See Section 5.1.7.11 on page 5–31 for information about why the non-lead member of the clustered management server was previously shut down.

8.2 Configure the External Network Interfaces

Note:

In Section 7.2, you will have defined those nodes with external IP addresses and the `sra install` command will have automatically configured these addresses during the installation process. However, if you have further changes to make, proceed with the instructions below.

Network interfaces on additional nodes can be used to move or share the routing load to cluster aliases within a domain. They can also be used to make a specific node a client of an external NFS file system. When such a node mounts the external file system, NFS traffic between the node and the external file server will be *via* the node's external interface. This is more efficient than routing through the primary nodes.

To configure the external network interfaces, perform the following tasks:

- Set Up NFS to Allow Mounts from External Machines (see Section 8.2.1 on page 8–2)

For more information about configuring network devices, see Chapter 22 of the *HP AlphaServer SC System Administration Guide*.

8.2.1 Set Up NFS to Allow Mounts from External Machines

The `vi` command hangs when trying to open existing files on a file system that is NFS-mounted from a machine that is not part of the cluster management network. The `vi` command hangs because it is attempting to obtain an exclusive lock on the file. The hang persists for many minutes; it may not be possible to use Ctrl/C to return to the command prompt.

The workaround is to ensure that the system that is NFS-serving the file system to a domain can resolve the internal domain member names (for example, `atlas0`) of the domain members that mount the NFS file system. The usual way of doing this is to use the internal domain member names as aliases for the address of the external interface on those nodes (for example, create an alias called `atlas0` for the `atlas0-ext1` external interface).

For example, domains `atlasD0` and `atlasD1` both NFS-mount the `/data` file system from the NFS server `dataserv`. The `/data` file system is being mounted by domain members `atlas0` and `atlas32`. These nodes have external interfaces `atlas0-ext1` and `atlas32-ext1` respectively. To avoid the `vi` hang problem, ensure that `dataserv` can resolve `atlas0` to `atlas0-ext1` and `atlas32` to `atlas32-ext1`.

This section describes three common ways of ensuring that the internal domain names can be resolved:

- `/etc/hosts`
In the `/etc/hosts` file on `dataserv`, define `atlas0` as an alias for `atlas0-ext1`.
In the `/etc/hosts` file on `dataserv`, define `atlas32` as an alias for `atlas32-ext1`.
You must perform this action on every system that is NFS-serving file systems to a domain.
- NIS/YP
If NIS/YP is in use, and is distributing a hosts table, put the alias definitions for `atlas0` and `atlas32` into this table.
- DNS
If DNS is in use, and is distributing host address information, define `atlas0` and `atlas32` as aliases for their respective external interface entries.

Note:

If you choose either the NIS/YP option or the DNS option, ensure that `svc.conf` is configured so that hostname resolution checks locally (that is, `/etc/hosts`) before going to `bind` or `yp`. For more information, see the `svc.conf(4)` reference page.

8.3 Improve Cluster Availability

The availability of a cluster is determined by the following factors:

- Cluster quorum (see Section 8.3.1 on page 8–4)
- The number of voting nodes present — if necessary, you can add votes (see Section 8.3.2 on page 8–4)

Improve Cluster Availability

8.3.1 Cluster Quorum

Cluster quorum is defined as follows:

Cluster Quorum = Round Down $[(\text{The number of votes} + 2) / 2]$

At this stage in the installation, each domain in the system has one voting node, the first node, giving it a quorum of 1:

Cluster Quorum = Round Down $[(1 + 2) / 2] = \text{Round Down } [1.5] = 1$

If this first node fails, there are no voting nodes present — the domain loses quorum and becomes unavailable.

Giving a vote to a second node would give a quorum of 2:

Cluster Quorum = Round Down $[(2 + 2) / 2] = \text{Round Down } [2] = 2$

To make the domain more available, we assign votes to Node 1 and Node 2 (in addition to the existing vote assigned to Node 0) to give a quorum of 2:

Cluster Quorum = Round Down $[(3 + 2) / 2] = \text{Round Down } [2.5] = 2$

Therefore, when any two of the first three nodes are running, the domain is available. Since Node 0 and Node 1 can each serve all of the system file store, the domain can tolerate the loss of either node.

Note:

Node 2 does not have to be connected to the system file store. We assign a vote to Node 2 simply to satisfy the quorum algorithm, to ensure that any two of the three nodes will suffice.

See Section 8.3.2 for more information on how to add votes to a node.

8.3.2 Add Votes

The first node of each cluster already has one vote. To improve cluster availability, add votes to the next two nodes of each cluster.

Ensure that all nodes of each cluster are booted, before adding more votes; this ensures that the configuration files of all nodes are updated appropriately.

Perform the following steps as the `root` user (where `atlas` is an example system name):

Note:

You must run the `clu_quorum` command (step 1) on the first node of the relevant cluster. The `sra` commands (step 2) may be run on the management server or on the first node of the cluster.

1. Add one vote to the next two nodes of each cluster, as follows:

```
On Node 0:      atlas0# clu_quorum -m atlas1 1
                  atlas0# clu_quorum -m atlas2 1
On Node 32:     atlas32# clu_quorum -m atlas33 1
                  atlas32# clu_quorum -m atlas34 1
On Node 64:     atlas64# clu_quorum -m atlas66 1
                  atlas64# clu_quorum -m atlas65 1
.
.
.
On Node 992:    atlas992# clu_quorum -m atlas993 1
                  atlas992# clu_quorum -m atlas994 1
```

2. Activate the additional votes by shutting down and booting the updated nodes, as follows:

```
a. Shut down the nodes:
atlas0# sra shutdown -nodes 'atlas[1-2]'
atlas32# sra shutdown -nodes 'atlas[33-34]'
atlas64# sra shutdown -nodes 'atlas[65-66]'
.
.
.
atlas992# sra shutdown -nodes 'atlas[993-994]'

b. When the shutdown is complete, boot the nodes:
atlas0# sra boot -nodes 'atlas[1-2]'
atlas32# sra boot -nodes 'atlas[33-34]'
atlas64# sra boot -nodes 'atlas[65-66]'
.
.
.
atlas992# sra boot -nodes 'atlas[993-994]'
```

When the nodes have booted, each cluster will have three voting members and a quorum of 2. It will be resilient to the loss of any one of the first three nodes.

To view the quorum status, run the `clu_quorum` command with no arguments, as shown in the following example:

```
# clu_quorum
```

Improve Cluster Availability

Note:

After updating the quorum configuration of the cluster, you must update the SC Database to reflect the votes of each member. Use the `sra edit` command to change the Number of Votes for the nodes you have configured.

See Appendix E for sample output from the `clu_quorum` command.

8.3.3 Side Effects of Using Quorum

As already demonstrated, quorum provides increased resilience and availability. However, there are a number of quorum-related side effects that you should be aware of, especially if you have been using a system with a quorum value of 1.

8.3.3.1 Shutting Down a Domain

On a system, regardless of the number of votes or nodes in the domain, use the following command to shut down the domain:

```
atlasms# sra shutdown -domains all
```

Note:

The `shutdown` command can be issued from the management server, or from another domain.

If you shut down the first domain on a system that does not have a management server, the shutdown will perform successfully. When you subsequently boot the domain, the `shutdown` command will not be listed as having completed because the shutdown could not be monitored by the system.

The `sra shutdown` command determines the best action to take to shut down your system by running either of the following commands:

```
atlasms# sra shutdown -nodes 'atlas[0-31]'  
atlasms# sra shutdown -domains atlasD0
```

Therefore, if `sra` detects that an entire domain is to be shut down, `sra` will run the `shutdown -ch` command on the relevant domain.

Note:

When shutting down just some nodes in a domain, `sra` detects if the `shutdown` command will cause a loss of quorum. If this is the case, `sra` will cancel the `shutdown` command, unless the `-force` flag is specified.

Load File System Configuration Data in the SC Database

8.3.3.2 Booting a System

To boot an entire system, you must have access to the console network *via* the console manager.

For systems that are configured with a management server, you can use the `sra boot` command from the management server to boot the domains. To achieve the fastest boot time for the 1024-node example system documented in this guide, boot all domains in parallel, by running the following command on the management server:

```
atlasms# sra boot -domains all
```

For systems without a management server, the console manager runs on the first domain. The console manager will not start until the first domain is established. This requires the presence of at least two of the first three nodes. Boot the first two nodes — in parallel — by entering the following command on the graphics console on Node 0 and on Node 1:

```
P00>>> boot boot_device
```

Note:

If you do not boot Node 0 and Node 1 in parallel, both nodes will hang until they reach a common point in the boot process and attain quorum.

Once these nodes have booted, you can use the `sra boot` command to boot the remaining nodes. As each domain can boot in parallel, you can reduce the boot time by running the following command:

```
atlas0# sra boot -domains all
```

Attaining quorum during boot is not an issue for the remaining domains as the `sra boot` command will, by default, boot four nodes at a time.

8.4 Load File System Configuration Data in the SC Database

The `scfsmgr Sysman` menu needs information about the configuration and usage of all disks on all File Serving (FS) Domains. You load this data into the database using the `scfsmgr scan` command as follows:

```
# scfsmgr scan
```

On a large system, this command can take a long time to complete.

If all nodes on the FS Domain(s) are not running when you use `scfsmgr scan`, you should rerun `scfsmgr scan` when those nodes are next booted.

8.5 Initial Setup of Monitoring for HSG80 RAID Systems

The HSG80 devices are not automatically added as monitored device by the install process (`sra setup`). In a large system, it is possible to have hundreds of HSG80 devices, and adding these as monitored devices manually can be time consuming.

To set the initial distribution of HSG80 RAID systems, use the `scmonmgr detect` command. This command runs the detect process on each node from the domain, or on the management server on which it is run. To detect all HSG devices installed in a system, run the command on each management server and each domain.

The syntax of the command is as follows:

```
scmonmgr detect -c class [-d <0|1>]
```

Where:

`-c class` – the class of the devices that must be detected.

Note:

In the current release, only the HSG class is supported.

`-d <0|1>` - an optional parameter which sets the debugging mode OFF/ON. If the parameter is not specified, the default value is OFF. When the debugging mode is ON (`-d 1`) the detect process displays additional information for each phase. This mode is used for diagnosis and testing purposes.

If the debug is turned off, the command displays only the new detected devices and errors. The new detected devices are printed in the following form:

```
node_name: found hsgX (scpX) on node_name
```

and the errors are printed with the node name followed by the error explanation:

```
node_name: [error explanation]
```

Some Examples:

1. To run the command on all domains with debug option disabled (and so ensure that all the detected HSG devices are added in the database as monitored objects), do the following:

```
atlasms# scrun -d all '/usr/bin/scmonmgr detect -c hsg'
```

2. To detect the HSG devices on a management server with the debug disabled, do the following:

```
atlasms# scmonmgr detect -c hsg
```

```
This system is not configured in a cluster.
```

```
The detect process will be run only on atlasms.
```

```
Detect process finished.
```


Initial Setup of Monitoring for HSG80 RAID Systems

```
atlasms#
```

3. To detect the HSG devices on a management server with the debug enabled, do the following:

```
atlasms# scmonmgr detect -c hsg -d 1
This system is not configured in a cluster.
The detect process will be run only on atlasms.
INFO   : Starting the HSG detect process on 'atlasms'.
INFO   : There are no HSG devices attached to this node.
INFO   : Exit
Detect process finished.
atlasms#
```

4. To detect the HSG devices on a certain domain (for example, atlasD0) it is necessary to run the command on a node from this domain. This will detect all HSG80 devices attached to the domain nodes, as follows:

```
atlas0# scmonmgr detect -c hsg
Starting HSG detect process for cluster atlasD0
atlas0: found hsg1 (scp0) on atlas0
atlas1: found hsg2 (scp0) on atlas1
atlas4: error in the detect process
Detect process finished on cluster atlasD0
atlas0#
```

5. To detect the HSG devices on a certain domain, with only 2 nodes in domain atlasD0 and with debugging enabled, do the following:

```
atlas# scmonmgr detect -c hsg -d 1
Starting HSG detect process for cluster atlasD0
INFO   : Number of nodes in cluster: 2
INFO   : Detecting HSG devices on atlas0
atlas0 : INFO   : Starting the HSG detect process on 'atlas0'.
atlas0 : INFO   : Found HSG device scp0 using bus 1 targ 0 lun 0
atlas0 : INFO   : Successfully logged into the controller scp0 at bus 1
targ 0 lun 0
atlas0 : INFO   : Sending 'show this' command for scp0 at bus 1 targ 0 lun 0
atlas0 : INFO   : WWID for scp0 at bus 1 targ 0 lun 0 is 5000-1FE1-0009-5180
atlas0 : INFO   : New HSG name = hsg47 for scp0 (wwid=5000-1FE1-0009-5180)
atlas0 : INFO   : Sending 'quit' for scp0 at bus 1 targ 0 lun 0
atlas0 : INFO   : port_wwid=5000-1FE1-0009-5183 and serial=ZG03200669 is
Controller A for hsg47.
atlas0 : INFO   : Controller A for hsg47 added into database.
atlas0 : INFO   : port_wwid=5000-1FE1-0009-5181 and serial=ZG03200650 is
Controller B for hsg47.
atlas0 : INFO   : Controller B for hsg47 added into database.
atlas0 : INFO   : Adding hsg47 (wwid=5000-1FE1-0009-5180) as a monitored
device
atlas0: found hsg47 (scp0) on atlas0
atlas0: INFO   : Detect process finished.
atlas0: INFO   : Exit
INFO   : Detecting HSG devices on atlas1
atlas1: INFO   : Starting the HSG detect process on 'atlas1'.
atlas1: INFO   : Found HSG device scp0 using bus 1 targ 0 lun 0
atlas1: INFO   : Successfully logged into the controller scp0 at bus 1
```

Configure the RMS Database

```
targ 0 lun 0
atlas1: INFO      : Sending 'show this' command for scp0 at bus 1 targ 0 lun 0
atlas1: INFO      : WWID for scp0 at bus 1 targ 0 lun 0 is 5000-1FE1-0009-5180
atlas1: INFO      : Device 'scp0' with WWID=5000-1FE1-0009-5180 exists in the
database.
atlas1: INFO      : Detect process finished.
atlas1: INFO      : Exit
Detect process finished on cluster atlasD0
atlas0#
```

Note:

An HSG80 device can be monitored only from a node that is directly attached to it. The detect process sets this correctly. It is not recommended to set the monitoring to another node. If you want to do this, be sure that the HSG80 device is attached to the node (monitoring server) which you specify as the server in the `scmonmgr move` command.

8.6 Configure the RMS Database

To configure the RMS database, you use `sra setup` command to perform a number of automatic steps. You can configure the RMS database in the following ways:

- Set Up RMS Partitions (see Section 8.6.1 on page 8–10)
- Customize RMS Partitions (see Section 8.6.2 on page 8–11)

8.6.1 Set Up RMS Partitions

Partitions are used to organize nodes in the HP AlphaServer SC system into groups for organizational, management, and policy reasons. The nodes in the system can be organized into named partitions. The `rcontrol` command is used to define partitions. The `prun` and `allocate` commands use the interface to the RMS system to execute programs. The `prun` and `allocate` commands allocate resources from a named partition.

Log on to the `rmshost` node (management server) as the `root` user.

The following example shows how to organize the 1024-node `atlas` system into three partitions.

Note:

Partitions are site-dependent; the partitions created below are examples only. A node cannot be in more than one partition at a time.

1. Create each partition using the `rcontrol` command, as follows:

```
atlasms# rcontrol create partition=fileserv configuration=day nodes='atlas0'
```

```
atlasms# rcontrol create partition=interactive configuration=day nodes='atlas[1,2] '
atlasms# rcontrol create partition=parallel configuration=day nodes='atlas[3-1023] '
```

2. Start each partition using the `rcontrol` command, as follows:

```
atlasms# rcontrol start partition=fileserv
atlasms# rcontrol start partition=interactive
atlasms# rcontrol start partition=parallel
```

3. Run the `rinfo` command to check the status of the partitions, as follows:

```
atlasms# rinfo
```

The output should be similar to the following:

PARTITION	CPUS	STATUS	TIME	TIMELIMIT	NODES
root	4096				atlas[0-1023]
fileserv	0/4				atlas0
interactive	0/8				atlas[1-2]
parallel	0/4084				atlas[3-1023]

8.6.2 Customize RMS Partitions

To change a partition, you must first stop the partition, delete it, create it again, and then start it. For example, the following commands will change the size of the `interactive` and `parallel` partitions:

1. Make sure that no jobs are running on the partitions (use the `rinfo` command).

2. Stop the partitions as follows:

```
atlasms# rcontrol stop partition=interactive
atlasms# rcontrol stop partition=parallel
```

3. Delete the partitions, as follows:

```
atlasms# rcontrol remove partition=interactive configuration=day
atlasms# rcontrol remove partition=parallel configuration=day
```

4. Recreate and restart the partitions, as follows:

```
atlasms# rcontrol create partition=interactive configuration=day nodes='atlas[1-3] '
atlasms# rcontrol create partition=parallel configuration=day nodes='atlas[4-1023] '
atlasms# rcontrol start partition=interactive
atlasms# rcontrol start partition=parallel
```

Note:

All nodes are always in the `root` partition. You must not attempt to change the `root` partition in any way. You must not use the `root` partition in the `prun` command. If you wish to run a job on all nodes, create a single partition containing all nodes.

For more information about setting up RMS partitions, see Chapter 5 of the *HP AlphaServer SC System Administration Guide*.

8.7 Provide RMS with CAA Failover Capability

At this point in the installation, RMS can operate correctly. However, if the host defined as `rmshost` (see Section 5.2 on page 5–37, or Section 6.2 on page 6–24) fails, RMS services become unavailable. This section describes how to use CAA to allow the `rmshost` functions to fail over to another node.

To make RMS a CAA failover application, perform the following steps on the first domain (where `atlas` is an example system name, and `atlas0` is the current `rmshost`):

Note:

The following CAA steps for RMS and MySQL only apply in cases where there is no management server, or where the system has a clustered management server.

Note:

These steps assume that RMS/MySQL are currently not under CAA control and that if you need to update CAA profiles and registrations, then you should consult the CAA man page.

1. Stop the RMS daemon, as follows:

```
atlas0# rmsctl stop
```

2. Stop the MySQL daemon, as follows:

```
atlas0# /sbin/init.d/msqld stop
```

3. Create two CAA application resource profiles for RMS, as follows:

```
# caa_profile -create SC05mysql -t application -a mysql.scr -p restricted  
-h 'atlas0 atlas1' -o as=1,ci=300,st=500  
# caa_profile -create SC20rms -t application -a rms.scr -p restricted -h  
'atlas0 atlas1' -o as=1,ci=300,st=500
```

If you receive a warning that the profiles already exist and that creation is cancelled, then you should update the existing profiles, as follows:

```
# caa_profile -update SC05mysql -t application -a mysql.scr -p restricted  
-h 'atlas0 atlas1' -o as=1,ci=300,st=500  
# caa_profile -update SC20rms -t application -a rms.scr -p restricted -h  
'atlas0 atlas1' -o as=1,ci=300,st=500
```

Provide RMS with CAA Failover Capability

Note:

The `caa_profile` commands create or update the resource profiles `SC05mysql` and `SC20rms`: defining the resources, their dependencies, and how the resources are managed by CAA. The commands also create a restricted placement policy - the application can run on `atlas0` or `atlas1` only, where `atlas0` and `atlas1` are two of the nodes that will provide RMS services.

4. Register the `SC05mysql` and `SC20rms` resource profiles with CAA as follows:

```
# /usr/sbin/caa_register SC05mysql
# /usr/sbin/caa_register SC20rms
```

This allows CAA to monitor and manage the resources.

Note:

If the profiles have been deleted (using the `caa_profile -delete` command), the associated scripts will also have been deleted, causing the following messages to appear:

```
# /usr/sbin/caa_register SC05mysql
Action Script `/var/cluster/caa/script/mysql.scr` does not exist!
Could not register resource mysql.
# /usr/sbin/caa_register SC20rms
Action Script `/var/cluster/caa/script/rms.scr` does not exist!
Could not register resource rms.
```

If these messages appear, copy the standard profile scripts to the appropriate directory, as follows:

```
# cp /usr/opt/rms/examples/scripts/mysql.scr /var/cluster/caa/script
# cp /usr/opt/rms/examples/scripts/rms.scr /var/cluster/caa/script
and re-run the caa_register commands.
```

5. Edit the `/etc/hosts` file — on each domain and on the management server (if used) — so that `rmshost` is a host alias of the default cluster alias of the first domain (for example, `atlasD0`), as follows:

Before:

```
atlasD0
atlas0      rmshost
```

After:

```
atlasD0      rmshost
atlas0
```

6. Start RMS as a CAA application, as follows:

```
# /usr/sbin/caa_start SC05mysql
# /usr/sbin/caa_start SC20rms
```

The `caa_start SC05mysql` and `caa_start SC20rms` commands will be performed by CAA on subsequent reboots.

Provide RMS with CAA Failover Capability

Note:

The RMS startup script can recognize whether it is being run on a domain, and whether RMS is registered as a CAA application.

7. At this stage, MySQL and the RMS servers are running on one node only (typically `atlas0`). You can start RMS on the remaining nodes as follows:
`atlas0# rmsctl start`

Because RMS servers are already running on a node (typically `atlas0`), `rmsctl` will report the RMS is already running on that node. You can ignore this message.

RMS is now up and running as a failover application.

The `caa_start SC05mysql` command starts the MySQL daemon on one of the nodes specified by the `SC05mysql` application profile. The profile was created in step 3 above. In this example, it is running on either `atlas0` or `atlas1`. Once you have registered `SC05mysql` as a CAA application, the `/sbin/init.d/mysqld` script no longer starts or stops the MySQL daemon; instead, you must use the `caa_start` or `caa_stop` commands. If you manually kill the MySQL daemon, without first using `caa_stop`, the CAA system will automatically restart the daemon after a short interval.

To see which node should be running the MySQL daemon, use the `caa_stat SC05mysql` command.

The `caa_start SC20rms` command starts RMS servers on one of the nodes specified by the `SC10rms` application profile. The profile was created in step 3 above. As CAA is starting RMS servers on the node, the `rmshost` attribute is set to the name of that node. To see which node should be running the `SC20rms` service, use the `caa_stat SC20rms` command. While the `SC20rms` service is running, the `rmshost` attribute and `caa_stat` should agree (for example, if the `SC20rms` application is running on `atlas1`, the `rmshost` attribute should have a value of `atlas1`). You can see the value of the `rmshost` attribute as follows:

```
# rmsquery "select val from attributes where name='rmshost' "
```

If you stop RMS on the node currently running the `rms` CAA application, using either `rmsctl stop` or `/sbin/init.d/rms stop`, the CAA application will automatically restart it after a short interval. To stop RMS, you must run the following commands in the specified sequence:

```
# caa_stop SC20rms
# rmsctl
```

If you have stopped RMS and then use `rmsctl` to start RMS, the RMS system will start and appear to operate correctly. The last value in the `rmshost` attribute will determine which node runs the RMS servers. However, should that node fail, CAA will not start RMS on the backup node. This is because RMS is in the stopped state. If you inadvertently use `rmsctl`

start while the rms CAA state is stopped, you may safely use `caa_start SC20rms` to enable failover. There will be a brief interruption in RMS availability because the RMS servers are stopped and then restarted.

Once you have started the `SC05msql` and `SC20rms` applications using `caa_start`, the CAA system will automatically relocate the applications to another node if the original node fails. However, you can manually initiate this using the `caa_relocate` command. This can be useful if you plan to shut down the node. It is possible to independently relocate either the `SC05msql` or `SC20rms` application, as it is not mandatory that these both run on the same node.

During the relocation process, there will be a brief interruption of service. This is similar to the interruption that occurs when a node failure causes the relocation. During the period of relocation, RMS commands will report messages such as the following:

```
rinfo: Warning: Can't connect to MSOL server on rmshost: retrying ...
rcontrol: Warning: RMS server pmanager-parallel (rmshost) not responding
```

The CAA commands (`caa_start`, and so on) can be executed on any node in the domain running the application.

See Chapter 5 of the *HP AlphaServer SC System Administration Guide* for more information about RMS (and CAA failover).

8.8 Enable CMF as a CAA Application

On a system with no management server, or where using a clustered management server, the Console Logging Service should be provided with failover capability.

In the following, CMFHOST refers to the host running the `cmf` daemon(s) [`cmfd`], that is, `atlas0` (where there is no management server) or `atlasms[0|1]` (where there is a clustered management server).

After CMF is CAA enabled, CMFHOST will be a cluster alias: `atlasD0` (where there is no management server) or `atlasms` (where there is a clustered management server).

To enable CMF as a CAA application, perform the following steps:

1. Stop the console logging daemons by running the following command on CMFHOST:
`/sbin/init.d/cmf stop`
2. Use the `sra edit` command to set the CMF host in the SC database to be the cluster alias name of the domain hosting the CMF service, as follows:

```
# sra edit
sra> sys
sys> edit system
Id          Description                               Value
-----
:
```

Enable CMF as a CAA Application

```
[8 ] Node running console logging daemon (cmfd) atlas0
:
:
Select attributes to edit, q to quit
eg. 1-5 10 15

edit? 8
Node running console logging daemon (cmfd) [atlas0]
new value? atlasD0

Node running console logging daemon (cmfd) [atlasD0]
correct? [y|n] y
```

The `sra edit` command then asks if you would like to modify the SC database, as follows:

Modify SC database only (1), update daemons (2), restart daemons (3) [3]:

Enter 1 to modify the SC database only.

3. Check the CMF CAA profile, as follows:

```
atlas0# caa_stat -p SC10cmf
NAME=SC10cmf
TYPE=application
ACTION_SCRIPT=cmf.scr
ACTIVE_PLACEMENT=0
AUTO_START=1
CHECK_INTERVAL=60
DESCRIPTION=AlphaServer SC Console Management Facility
FAILOVER_DELAY=10
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=
OPTIONAL_RESOURCES=
PLACEMENT=balanced
REQUIRED_RESOURCES=
RESTART_ATTEMPTS=1
SCRIPT_TIMEOUT=300
```

When CMFHOST is the first domain (that is, when the HP AlphaServer SC system does not have a management server), the `HOSTING_MEMBERS` field should contain the hostnames of the first two nodes, and the `PLACEMENT` field should contain the text `restricted`, as shown in the following example:

```
HOSTING_MEMBERS=atlas0 atlas1
PLACEMENT=restricted
```

When CAA-enabled, the `cmfd` daemon will run on any node in the cluster (CMFHOST is the default cluster alias). However, it is preferable to use nodes that have a network interface on the subnet on which the cluster alias is defined — that is, the first two nodes in the default configuration — to avoid an extra routing hop.

If the output of the `caa_stat -p SC10cmf` command does not reflect the values specified above for the `HOSTING_MEMBERS` and the `PLACEMENT` fields, use a text editor to make the necessary changes to the `/var/cluster/caa/profile/SC10cmf.cap` file. Alternatively, use the `caa_profile` command to make these changes. For more information, see the *TruCluster Server Cluster Highly Available Applications* manual.

If CMFHOST is a clustered management server, the default values should be used for these fields, as follows:

```
HOSTING_MEMBERS=  
PLACEMENT=balanced
```

4. On the new CMFHOST (atlasD0), register CMF as a CAA application, as follows:

```
# caa_register SC10cmf
```
5. On the new CMFHOST (atlasD0), start the CAA service, as follows:

```
# caa_start SC10cmf
```

8.9 Run the Example MPI Program

When you have set up the RMS partitions as described in Section 8.6.1 on page 8–10, you can run the example MPI program provided on the *HP AlphaServer SC System Software* CD-ROM.

The installation process copies the example MPI program and associated README file into the `/usr/opt/mpi/examples/sra/mpi` directory. To run the example program, copy the contents into your own directory. The associated README file provides instructions on how to compile and run the program.

8.10 Verify the HP AlphaServer SC Interconnect

In Chapter 5 and Chapter 6, you will have completed various diagnostic tests while building the system. At this stage, you should complete the range of diagnostic tests that are available.

The diagnostic tests are documented in the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*. This section contains a number of references to the relevant sections of that guide. Please be careful not to miss any steps as you refer to the sections. It is advisable to complete all the diagnostic steps suggested below.

- Run a simple boot-time check (see Section 7.4.8 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `elan_level_test` to check all node cables (see Section 7.4.9 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `cabletest` to check all uplinks (see Section 7.4.10 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

Configure LSM

- Use `elan_level_test` to check each node-level switch (see Section 7.4.11 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `elan_level_test` to check all uplinks (see Section 7.4.12 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `elan_level_test` to check each top-level switch (see Section 7.4.13 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `elansoaktest` to test the entire HP AlphaServer SC Interconnect (see Section 7.4.14 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)
- Use `neterror` to identify any faulty components (see Section 7.4.15 of the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual*)

Once all tests are complete, you should prepare a maintenance schedule where you will elect to run at least `elan_level_test` and `neterror` at intervals - perhaps weekly. You should also monitor the events table with `sc_event` to identify any potential errors with the interconnect hardware (See Chapter 9 of the *HP AlphaServer SC System Administration Guide* for information on Managing Events).

8.11 Configure LSM

If you wish to configure LSM, do so after all nodes have been added to the domain. For more information about LSM, see Chapter 25 of the *HP AlphaServer SC System Administration Guide*.

8.12 Verify Swap Mode

The HP AlphaServer SC installation procedure configures all nodes with "eager" swap mode. Do not switch to "lazy" swap mode.

Note:

Lazy swap is not supported in HP AlphaServer SC Version 2.6 (UK2) — use eager swap only.

8.13 Add a Second Rail to an HP AlphaServer SC System after Domain Creation

If you add a second rail to an HP AlphaServer SC system after domain creation, you must perform some additional steps.

Add a Second Rail to an HP AlphaServer SC System after Domain Creation

Note:

The second rail should be populated before proceeding with the following steps.

The steps are different depending on the type of node used, as follows:

- HP AlphaServer SC System Composed of HP AlphaServer ES40 Nodes (page 8-19)
- HP AlphaServer SC System Composed of HP AlphaServer ES45 Nodes (page 8-21)

Note:

The information in this section is not applicable to HP AlphaServer DS20L systems.

8.13.1 HP AlphaServer SC System Composed of HP AlphaServer ES40 Nodes

This section applies only to an HP AlphaServer SC system composed of HP AlphaServer ES40 nodes. The additional steps are necessary because of the order in which the PCI buses are probed on an HP AlphaServer ES40. The name assigned to each rail (`elan0` or `elan1`) is related to the probe order, and whether one or two rails were present when the system was initially installed.

To ensure that the rails are consistently named, perform the following steps:

1. Ensure that all nodes are up.
2. Run the following command:
`# sra command -nodes all -command '/usr/sra/bin/sra_multirail -move'`
3. Shut down the entire system.
4. Boot the system.
5. Use the `sra edit` command to set the rail count for each domain, as follows:

```
# sra edit
sra> sys
sys> edit cluster atlasD0
Id Description Value
-----
[0 ] Cluster name atlasD0
[1 ] Cluster alias IP address www.xxx.yyy.zzz
[2 ] Cluster domain type fs
[3 ] First node in the cluster 0
[4 ] Primary partition list {root dsk3b 5} {usr dsk3g 50}
{var dsk3h 45}
[5 ] Backup partition list {root dsk5b 5} {usr dsk5g 50}
{var dsk5h 45}
[6 ] I18n partition device name
[7 ] Number of Cluster Interconnect Rails 1
```

Add a Second Rail to an HP AlphaServer SC System after Domain Creation

```
-----  
Select attributes to edit, q to quit  
  
eg.1-5 10 15  
edit? 7  
Number of Cluster Interconnect Rails [1]  
new value? 2  
Number of Cluster Interconnect Rails [2]  
Correct? [y|n] y  
sys>
```

6. Repeat step 5 for each domain in the system.

Add a Second Rail to an HP AlphaServer SC System after Domain Creation

8.13.2 HP AlphaServer SC System Composed of HP AlphaServer ES45 Nodes

This section applies only to an HP AlphaServer SC system composed of HP AlphaServer ES45 nodes. To add a second rail after domain creation, use the `sra edit` command to set the rail count for each domain, as follows:

```
# sra edit
sra> sys
sys> edit cluster atlasD0
Id      DescriptionValue
-----
.
.
.
[7 ]    Number of Cluster IC Rails1
.
.
.
-----
Select attributes to edit, q to quit
eg. 1-5 10 15
edit? 7
Number of Cluster IC Rails      [1]
new value? 2
Number of Cluster IC Rails      [2]
Correct? [y|n] y
sys>
```

Installing LSF for HP AlphaServer SC

This chapter describes how to install Platform Computing Corporation's LSF® software ("LSF") on an HP AlphaServer SC system. The chapter presents the following topics:

- LSF Overview (see Section 9.1 on page 9–2)
- Installing a New LSF for HP AlphaServer SC (see Section 9.2 on page 9–3)
- Configuring LSF for HP AlphaServer SC and Starting the LSF (see Section 9.3 on page 9–10)

9.1 LSF Overview

LSF for HP AlphaServer SC combines the strengths of LSF and HP AlphaServer SC software to provide a comprehensive Distributed Resource Management (DRM) solution.

LSF acts primarily as the workload scheduler, providing policy and topology-based scheduling. RMS acts as a parallel subsystem, and other HP AlphaServer SC software provides enhanced fault tolerance.

Nodes within HP AlphaServer SC systems are arranged into domains. LSF can be configured to treat each domain as a single virtual server host. Each virtual host has the same name as the domain.

For example, if two domains containing 64 HP AlphaServer ES45 systems are configured to appear as two 128 CPU virtual LSF server hosts, LSF policies are applied to the virtual hosts as if they were real hosts.

9.1.1 Preparing for LSF Installation

Before you install the LSF software, please note the following points:

- A domain should be configured as a single virtual host or a set of real hosts.
- You must upgrade all Platform LSF for HP AlphaServer SC hosts to Platform LSF/HP AlphaServer SC Version 2.6 (UK1). Platform LSF has not changed for HP AlphaServer SC Version 2.6 (UK2) so this **is** the correct LSF version.
- The LSF log directory `LSF_LOGDIR` in `lsf.conf` must be a local directory. Do not use an NFS-mounted directory as `LSF_LOGDIR`. The `LSF_LOGDIR` should be set to `/var/lsf_logs`. This mount point (`/var`) is on the domain (not on NFS) and it is shared among cluster members. To set this during the installation procedure, add the following line to the end of the `install.config` file:
`LSF_LOGDIR="/var/lsf_logs"`
- LSF requires an administrative user. This must not be the root user. Frequently, the username is `lsfadmin`. However, any non-root user name can be used as the LSF administrator. The LSF administrator has to be a known user on all LSF hosts.
- The LSF software and configuration files are stored in a filesystem that is mounted using NFS on all domains and management servers. You can use the management server, one of the domains, or some NFS server external to the HP AlphaServer SC System to serve this file system. The path (`LSF_TOP`) for the filesystem is usually `/usr/share/lsf`, although you can configure your system differently. Do not use `/`, `/var/lsf` or `/usr/opt/lsf` for `LSF_TOP`

- There should be no nodes on the HP AlphaServer SC system defined in the RMS database that do not physically exist, or that have never been booted. To check for such nodes on the system, run the following command:

```
rmsquery "select name from nodes where cpus=-2"
```


If these nodes do exist, they should be excluded from the RMS partition(s) that are under LSF control.

9.2 Installing a New LSF for HP AlphaServer SC

Installing the LSF software involves the following tasks:

- Before You Install (see Section 9.2.1 on page 9–3)
- Installing LSF for HP AlphaServer SC (see Section 9.2.2 on page 9–3)
- Setting up LSF hosts (hostsetup) (see Section 9.2.3 on page 9–8)
- Configuring the LSF hosts File for Multiple Network Interfaces (see Section 9.3.1 on page 9–10)

9.2.1 Before You Install

These installation instructions assume that you are installing a new LSF cluster on the HP AlphaServer SC system.

You can add an HP AlphaServer SC system to an existing LSF cluster, but you should back up your existing LSF configuration files before doing this.

9.2.2 Installing LSF for HP AlphaServer SC

To install LSF for HP AlphaServer SC, follow these steps:

1. Log on as `root` to the management node. If you do not have management nodes, log on to the first node of the first domain.
2. Change to the directory containing the LSF distribution kit. For example:

```
# cd /tmp
```
3. Use the `zcat` and `tar` commands to decompress and extract the `lsf6.0_lsfinstall.tar.Z` distribution file as follows:

```
# zcat lsf6.0_lsfinstall.tar.Z | tar xvf -
```


Do not extract `lsf6.0_alpha5-rms.tar.Z`.
4. Change to `lsf6.0_lsfinstall`:

```
# cd /tmp/lsf6.0_lsfinstall
```
5. Read `lsf6.0_lsfinstall/install.config` and decide which installation variables you need to set.

Installing a New LSF for HP AlphaServer SC

- 6. Edit `lsf6.0_lsfinstall/install.config` to specify the installation variables you want.
- 7. Run `lsfinstall` as root:
`# ./lsfinstall -f install.config`
- 8. NFS mount the `LSF_TOP` directory (`/usr/share/lsf`), from the node the `lsfinstall` script has been run, to each domain using the standard UNIX mount procedure on the head node of each domain. Add an entry for the `LSF_TOP` directory to the `/etc/member_fstab` file on the head node of each domain. For more information on how to use NFS mount, see the *HP AlphaServer SC System Administration Guide*.

Note:

Do not use autoFS to mount an LSF installation. LSF for HP AlphaServer SC Version 2.5 and higher needs to mount only the `LSF_TOP` directory. Both `indir` and `cmddir` are located under the `LSF_TOP` directory (in `$LST_TOP/work`), so they do not need to be mounted individually.

The `LSF_TOP` directory needs to be exported in order to support the LSF master failover capability. By default, each LSF host can potentially become the LSF master host, and it needs to be able to write LSF accounting data to `LSF_TOP/work`. LSF provides the functionality to limit potential LSF master hosts by listing candidate master hosts in the `LSF_MASTER_LIST` variable in `lsf.conf`. In this way, only the listed hosts can become LSF masters.

A possible solution is to use the `hostlist` option when exporting `LSF_TOP`, where `hostlist` would include only the nodes that can become LSF master hosts. In practice, this means that `hostlist` should include the `man` and `ext` aliases of the first two nodes on each domain that LSF is running on (this assumes that `LSF_TOP` is exported from the management server).

See the *HP AlphaServer SC Platform LSF® Reference* for information on the `install.config` file.

Table 9–1 shows the required `install.config` variables.

Table 9–1 Required install.config Variables

Variable	Example
LSF_TOP	<code>"/usr/share/lsf"</code>
LSF_ADMINS	<code>"user_name [user_name...]"</code>
LSF_CLUSTERNAME	<code>"cluster_name"</code>

Table 9–1 Required install.config Variables

Variable	Example
LSF_TARDIR	"path"
LSF_LICENSE	"<none>"
LSF_ADD_SERVERS	"atlasD0 atlasD1 atlasD2 atlasD3"

Table 9–2 shows the variables that require an absolute path.

Table 9–2 Variables that Require an Absolute Path

Variable	Path
LSF_TOP	"/usr/share/lsf"
LSF_LOGDIR	"/var/lsf_logs"
LSF_TARDIR	"path"

See the *HP AlphaServer SC Platform LSF® Reference* for more information about `lsfinstall` and the `install.config` file.

9.2.2.1 Example install.config

Example 9–1 displays a sample `install.config` file used for installing LSF without client nodes.

Example 9–1 Sample install.config file

```

*****
# Name:      install.config
#
# Purpose:   LSF installation options file template
#
# $Id: install.config,v 1.9 2002/12/09 22:12:51 pdetina Exp $
#
# Format:
#   o Each line is in the format of LSF_OPTION_NAME="VALUE"
#   o Options can only appear once in the file
#   o Blank lines and lines starting with a pound sign (#)
#     are ignored.
#
# Instructions:
#   1. Edit install.config to specify the options for
#      your cluster. Uncomment the options you want and
#      replace the example values with your own settings.
#   2. Run ./lsfinstall -f install.config

```

Installing a New LSF for HP AlphaServer SC

```
#
# See install.config in the HP AlphaServer SC Platform LSF Reference for details.
#*****

# -----
LSF_TOP="/usr/share/lsf"
# -----
# Required argument
# Full path to the top-level installation directory
# The path to LSF_TOP must be shared and accessible to all hosts
# in the cluster. Cannot be the root directory (/).
# File system containing LSF_TOP must have enough disk space for
# all host types (approximately 300 MB per host type).

# -----
LSF_ADMINS="lsfadmin user1 user2"
# -----
# Required argument
# List of LSF administrators
# The first user account name in the list is the primary LSF
# administrator. Cannot be the root user account.
# LSF administrator accounts must exist on all LSF hosts before
# installing LSF.
#
# The primary LSF administrator account is typically named lsfadmin.
# It owns the LSF configuration files and log files for job events.
# It also has permission to reconfigure LSF and to control batch
# jobs submitted by other users. It typically does not have
# authority to start LSF daemons. Usually, only root has
# permission to start LSF daemons.

# -----
LSF_CLUSTER_NAME="atlas"
# -----
# Required argument
# The name of the LSF cluster
# Must be 39 characters or less, and cannot contain any
# white spaces. Do not use an LSF host name.

# -----
LSF_TARDIR="/usr/share/lsf_distrib/6.0"
# -----
# Full path to the directory containing the LSF distribution
# tar files
# Default: parent directory of the current working directory
#           where lsfinstall is running.
#           For example, if lsfinstall is running in
#           /usr/share/lsf_distrib/lsf_lsfinstall,
#           the default is /usr/share/lsf_distrib/.

# -----
LSF_LICENSE="<none>"
# -----
```

Installing a New LSF for HP AlphaServer SC

```
# Full path to the LSF license.dat license file
# Default: license.dat in the parent directory of the current
#       working directory where lsfinstall is running.
#       For example, if lsfinstall is running in
#       /usr/share/lsf_distrib/lsf_lsfinstall,
#       the default is /usr/share/lsf_distrib/license.dat.

# -----
LSF_ADD_SERVERS="atlasms atlasD0 atlasD1 atlasD2 atlasD3"
# -----
# List of LSF server hosts to be added to the cluster
# Specify a list of host names two ways:
# o Host names separated by spaces
# o Name of a file containing a list of host names,
#   one host per line. For example:
#       LSF_ADD_SERVERS=:lsf_server_hosts
#       The file lsf_server_hosts contains:
#           hosta
#           hostb
#           hostc
#           hostd
# The first host listed is the LSF master host.
# Default: For initial installation, the local host where
# lsfinstall is running

# -----
# LSF_ADD_CLIENTS="hoste hostf"
# -----
# List of LSF client-only hosts to be added to the cluster
# Specify a list of host names two ways:
# o Host names separated by spaces
# o Name of a file containing a list of host names,
#   one host per line. For example:
#       LSF_ADD_CLIENTS=:lsf_client_hosts
#       The file lsf_server_hosts contains:
#           hoste
#           hostf
# Default: None

# -----
# UNIFORM_DIRECTORY_PATH="/usr/local/lsf"
# -----
# Local directory for the root of the path to the machine-dependent
# LSF files.
#
# Must be an absolute path to a local directory and not shared.
# Cannot be the root directory (/).
#
# Maintains existing uniform directory path for upgrade from
# LSF Version 4.x ONLY. DO NOT use for new installations.
# After lsfinstall is finished, you must run hostsetup on each
# host to set up UNIFORM_DIRECTORY_PATH.
# Default: None
```

Installing a New LSF for HP AlphaServer SC

```
# -----
# LSF_QUIET_INST="n"
# -----
# Do not display LSF installation messages
# Default: LSF_QUIET_INST="n"

*****
# The following options are useful for Dynamic addhost
# feature
*****
# -----
# LSF_MASTER_LIST="atlasms atlasD0 atlasD1 atlasD2 atlasD3"
# -----
# List of LSF server hosts to be master or master candidate in the cluster
# Specify a list of host names two ways:
# o Host names separated by spaces
# o Name of a file containing a list of host names,
#   one host per line. For example:
#       LSF_MASTER_LIST=:lsf_master_list
#   The file lsf_master_list contains:
#       hoste
#       hostf
#   The first host listed is the LSF master host.
# Default: None

#-----
LSF_LOGDIR="/var/lsf_logs"
#-----
# LSF log directory
```

Note:

LSF client nodes require a flexlm license, which has to be obtained from Platform Computing.

9.2.3 Setting up LSF hosts (hostsetup)

9.2.3.1 What hostsetup does

The `hostsetup` script performs the following:

1. Checks that LSF is available from the host where `hostsetup` is running, and that the host type matches one of the host types installed in `LSF_TOP`.
2. Creates a table named `lsfrids` in the RMS database. This table is used internally by LSF.

3. Creates `/usr/opt/lsf` and `/var/lsf` if they do not exist.
4. On a node of a virtual host, creates the appropriate cluster application availability (CAA) rules that register LSF as a CAA service; also modifies the `/etc/clua_services` file to register the ports for the LSF daemons to match the ports configured in `/usr/share/lsf/conf/lsf.conf`.
5. Adds a Boolean resource `rms` to `lsf.shared` and assigns this resource to all LSF hosts that run on an RMS partition in `lsf.cluster.cluster_name`.
6. Creates the following symbolic links:
 - `/usr/opt/lsf/bin` symbolic link to `LSF_BINDIR`
 - `/usr/opt/lsf/etc` symbolic link to `LSF_SERVERDIR`
 - `/usr/opt/lsf/lib` symbolic link to `LSF_LIBDIR`
 - `/var/lsf/conf` symbolic link to `LSF_ENVDIR`
 - `/var/lsf/work` symbolic link to `LSF_WORKDIR`
 - `/var/lsf/work` symbolic link to `LSB_SHAREDIR`
 - `/usr/opt/include` symbolic link to `LSF_INCLUDEDIR`
 - `/usr/opt/man` symbolic link to `LSF_MANDIR`
 - `/usr/opt/misc` symbolic link to `LSF_MISCDIR`
 - `/usr/opt/install` symbolic link to `LSF_TOP/LSF_VERSION/install`

9.2.3.2 Running `hostsetup`

To set up LSF hosts, run `hostsetup` as follows:

1. Log on to each LSF server host as `root`. Start with the LSF master host.
2. Run `hostsetup` on each LSF server host. For example:

```
# cd /usr/share/lsf/6.0/install/  
# ./hostsetup --top="/usr/share/lsf"
```

If you are setting up LSF on a real LSF host, management server or a stand-alone node, use the `--boot="y"` option on the `hostsetup` command to configure system scripts to automatically start and stop LSF daemons at system startup or shutdown. You must run `hostsetup` as `root` to use this option to modify the system scripts. The default is:
`--boot="n"`.

The `--boot` option is not needed to set up virtual hosts.

For information on `hostsetup` usage, enter `hostsetup -h`.

Configuring LSF for HP AlphaServer SC and Starting the LSF

9.2.4 Next Steps

You have now installed the software on all systems. However, not all configuration steps are complete. To complete the installation, proceed to the section Configuring LSF for HP AlphaServer SC and Starting the LSF (see Section 9.3 on page 9–10).

9.3 Configuring LSF for HP AlphaServer SC and Starting the LSF

This section describes the following:

- Configuring the LSF hosts File for Multiple Network Interfaces (see Section 9.3.1 on page 9–10)
- Starting LSF on the SC Cluster (see Section 9.3.2 on page 9–12)
- Verifying that the Configuration is Correct (see Section 9.3.3 on page 9–13)

9.3.1 Configuring the LSF hosts File for Multiple Network Interfaces

Your LSF server hosts can be configured to use multiple network interfaces; for example:

- Domain (for example, `atlasD0`)
- Management network interfaces (for example, `atlas0`)
- External network interfaces (for example, `atlas0-ext1`)

To manage the communication across multiple network interfaces, you must configure the `LSF_CONFDIR/hosts` file to associate the name of LSF server hosts (where the LSF daemons are running) with the external IP addresses and names of the nodes or domain names.

Note:

The `/usr/share/lsf/conf/hosts` file relates the virtual host IP address (that is, the cluster alias address) and the individual hosts in the domain. If there are any errors in this file, the Load Information Manager (LIM) on the master will be unable to communicate with the LIM on the domain.

The LSF hosts file has a compact format where each line in the hosts file represents one host (real or virtual) of the LSF cluster.

- For each real host (standalone hosts or management servers), the LSF hosts file should include the following line:
`<external host interface address> <host name> <external interface names>`
- For each virtual host, the LSF hosts file should include the following line:

Configuring LSF for HP AlphaServer SC and Starting the LSF

```
<cluster alias address> <cluster name> <member names> <member interface names>
```

Example 9–2 describes the format of the `LSF_CONFDIR/hosts` file.

Example 9–2 Format of LSF_CONFDIR/hosts file

```
16.21.24.15 atlasms0 atlasms0-ext1
16.21.24.16 atlasms1 atlasms1-ext1
16.21.24.60 atlasD0 atlas0 atlas0-ext1 atlas0-ext2 atlas0-eip0 atlas1 atlas1-ext1
atlas1-eip0 atlas2 atlas2-eip0 atlas3 atlas3-eip0
16.21.24.61 atlasD1 atlas4 atlas4-ext1 atlas4-ext2 atlas4-eip0 atlas5 atlas5-ext1
atlas5-eip0 atlas6 atlas6-eip0 atlas7 atlas7-eip0
```

Example 9-2 is explained as follows:

- System comprises two management servers and two domains configured as virtual hosts
- 16.21.24.15 is the external interface address of atlasms0 (management server)
- 16.21.24.16 is the external interface address of atlasms1 (management server)
- 16.21.24.60 is the cluster alias address for atlasD0 (virtual host atlasD0)
- 16.21.24.61 is the cluster alias address for atlasD1 (virtual host atlasD1)
- Each domain (atlasD0, atlasD1) contains 4 nodes
- The first two members of each domain have an external network interface (atlas0-ext1, atlas1-ext1, and so on).
- The first member of each domain has an additional external network interface (for example, this could be an ATM network). The extension -ext2 (for example, atlas4-ext2) has been chosen for this example, but another name or extension could have been used.

Note:

If the nodes contain additional external interfaces (for example, atlas4-ext2), the names of these interfaces should also be included in the LSF hosts file. This format ensures that the LSF CAA failover operates correctly and that LSF commands can be run from all nodes that run LSF daemons.

9.3.1.1 Configuring LSF_SERVER_HOSTS

To submit LSF commands from the nodes which do not run LSF daemons, define the following variable in the `lsf.conf` file:

```
LSF_SERVER_HOSTS="the list of all LSF hosts"
```

Configuring LSF for HP AlphaServer SC and Starting the LSF

In Example 9–2, the entry should be as follows:

```
LSF_SERVER_HOSTS="atlasms0 atlasms1 atlasD0 atlasD1"
```

For more information about the LSF hosts file, see the *HP AlphaServer SC Platform LSF® Reference*.

9.3.1.2 When to Update the Hosts File

The `LSF_CONFDIR/hosts` file should be updated in the following cases:

- When you add a domain as an LSF server host
- When you add other (external) interfaces to any node in a domain.
- When you add a node to a domain
- If you change any IP address listed in the file
- If you change the name of any network interface

9.3.2 Starting LSF on the SC Cluster

To start the cluster, do the following:

- If a domain is configured as a single virtual LSF host, log on to any node of that domain, and run the following command on each domain on the cluster configured as virtual LSF host:

```
# caa_start lsf
```

- On the management server, and on the cluster nodes configured as real LSF hosts, log on to the node and set your LSF environment as follows:

For `cs`h or `tc`sh:

```
% source /usr/share/lsf/conf/cshrc.lsf
```

For `sh`, `ksh`, or `bash`:

```
$ . /usr/share/lsf/conf/profile.lsf
```

Run the following command:

```
# lsf_daemons start
```

Repeat this procedure on each LSF host.

- Use the following command to reactivate all LSF queues after upgrading:

```
% badmin qact all
```
- To test your cluster, run some basic LSF commands (for example, `lsid`, `lshosts`, `bhosts`)

Configuring LSF for HP AlphaServer SC and Starting the LSF

Note:

After testing your cluster, be sure all LSF users include `LSF_CONFDIR/cshrc.lsf` or `LSF_CONFDIR/profile.lsf` in their `.cshrc` or `.profile`.

9.3.3 Verifying that the Configuration is Correct

To check that the configuration is correct, do the following:

1. Log on as root to any LSF server.
2. Verify the configuration of the `/etc/clua_services` file, which registers LSF as a CAA service.
 - The port numbers for the LSF daemons `lim` (`LSF_LIM_PORT`), `sbatchd` (`LSB_SBD_PORT`), `res` (`LSF_RES_PORT`), `mbatchd` (`LSB_MBD_PORT`), `rla` (`LSF_RLA_PORT`) must match the ports configured in `/usr/share/lsf/conf/lsf.conf`.
 - Virtual hosts are set up with `in_single`, `out_alias` options, single-node hosts are set up with `in_multi`, `in_noalias` options.
 - Each daemon should have two entries, one for TCP and one for UDP. Sample `/etc/clua_services` entries are as follows:

```
# LSF -- appended Tue Jul 2 14:46:45 EDT 2002
lim 6879/tcp in_single,out_alias
lim 6879/udp in_single,out_alias
res 6878/tcp in_single,out_alias
res 6878/udp in_single,out_alias
sbatchd 6882/tcp in_single,out_alias
sbatchd 6882/udp in_single,out_alias
rla 6883/tcp in_single,out_alias
rla 6883/udp in_single,out_alias
mbatchd 6881/tcp in_single,out_alias
mbatchd 6881/udp in_single,out_alias
```

If you change the LSF daemon ports in `/usr/share/lsf/conf/lsf.conf`, you must make corresponding changes to all `clua_services` files.

For more information about the `clua_services` file, see the `clua_services(5)` TruCluster 64 man page.

For more information about CAA, see the *HP TruCluster Server: High Availability Applications*.

3. Use the `lsload -l` and `bhosts` commands to display the load information for all domains in the cluster.

Configuring LSF for HP AlphaServer SC and Starting the LSF

- 4. After verifying that LSF is operating properly, make LSF available to your users by having them include `/usr/share/lsf/conf/cshrc.lsf` or `/usr/share/lsf/conf/profile.lsf` in their `.cshrc` or `.profile`.

9.3.3.1 Sample lsload -l output

The status for all domains should be `ok`. Hosts with the static resource `rms` defined do not report `io` or `it` load indices. The output should look similar to the following:

```
# lsload -l
HOST_NAME status  r15s r1m r15m  ut  pg   io  ls it tmp    swp   mem
atlasD0  ok          0.0 0.0 0.0    0% 0.0 -   0 - 7184M 5016M 2432M
atlasD1  ok          0.0 0.0 0.0    0% 0.0 -   1 - 7184M 5004M 2436M
atlas64  ok          0.0 0.0 0.8    0% 0.0 -   2 - 7184M 4756M 32M
atlas65  ok          1.6 1.6 33.2   1% 0.0 -   2 - 7184M 1296M 448M
atlas66  ok          5.2 5.2 5.7    3% 0.0 -   4 - 7184M 4656M 2498M
atlasms  ok          0.9 0.9 0.9   22% 0.0 3840 15 0  442M 5088M 2876M
```

9.3.3.2 Sample bhosts output

The status for all domains should be `ok`. The output should look similar to this:

```
% bhosts

% bhosts
HOST_NAME      status JL/U      MAX      NJOBS      RUN      SSUSP      USUS  PRSV
atlasD0        ok      -      128      0      0      0      0      0
atlasD1        ok      -      128      0      0      0      0      0
atlas64        ok      -      4      0      0      0      0      0
atlas65        ok      -      4      0      0      0      0      0
atlas66        ok      -      4      0      0      0      0      0
```

Post-Installation Tasks

This chapter provides an overview of the administrative tasks that are normally carried out on a recently installed system. It describes, in high-level detail, the kind of operations normally carried out prior to putting the system into general use. For more detail on how to perform these tasks, see the *HP AlphaServer SC System Administration Guide*.

This chapter presents the following topics:

- State of System Immediately After Installation (see Section 10.1 on page 10–2)
- General Post-Installation Tasks (see Section 10.2 on page 10–3)
- Console Network (see Section 10.3 on page 10–10)
- Storage and File Systems (see Section 10.4 on page 10–11)
- User Administration (see Section 10.5 on page 10–12)
- Cluster Aliases (see Section 10.6 on page 10–13)
- RMS (see Section 10.7 on page 10–13)

State of System Immediately After Installation

10.1 State of System Immediately After Installation

The typical state of your system immediately after installation is expected to be as follows:

- All nodes are booted and operational.
- The first two members of each domain have physical connections to the system storage array. If you have followed the recommended storage guidelines, four 36GB drives are configured for each domain: two as a Mirrorset, one as the generic boot disk, and one as the backup cluster disk. An additional disk is configured as a spareset — the spareset is shared by all domains.
- The Mirrorset stores the clusterwide system storage, specifically `/`, `/var`, and `/usr`. These file systems are seen by all nodes within the domain and contain the installed system software and any optional layered products installed during the installation process.
- Each node has two local disks configured as follows (see Section 2.4.1 on page 2–9):
 - First disk: boot partition containing bootable `vmunix`, `genvmunix` and startup files; node swap; node `/tmp`; and node `/local`
 - Second disk: alternate boot partition; additional swap, `/tmp1`, and `/local1`

Configuring a second disk allows you to boot from that disk if the primary boot disk fails. *Using* a second disk permits access to its `tmp` and `local` partitions, and its swap space.
- The console manager is running on Node 0, or on the management server (if configured). It is monitoring the console lines of all nodes.
- All nodes are connected to the management network.
- The first two nodes of each domain have external LAN interfaces configured.
- Each domain has been configured as a DNS client and as an NTP client. Each may also be configured as an NIS slave, if NIS is in use.
- Each domain has been configured as both an NFS server and NFS client.
- You may have demonstrator partitions — named `fileserv`, `interactive`, and `parallel` — configured and running. These are normally set up at the end of the installation process as part of the installation verification. The partitions as defined may not match your site requirements.

Before putting the system into general use, you typically need to perform a number of tasks to develop a final system configuration that meets your site requirements and policies. These include such tasks as configuring further network interfaces, storage, and file systems, as well as user administration. The remainder of this section provides a high-level description of the nature of these tasks, and how they differ from configuring a normal SMP system.

10.2 General Post-Installation Tasks

This section describes the following general post-installation tasks:

- Restrict Access to Login Nodes (see Section 10.2.1 on page 10–3)
- Back Up and Restore the Management Server Root (/), /usr, and /var File Systems (see Section 10.2.2 on page 10–4)
- Back Up the Cluster Root (/), /usr, and /var File Systems (see Section 10.2.3 on page 10–7)
- Implement Security Recommendations (see Section 10.2.4 on page 10–10)

10.2.1 Restrict Access to Login Nodes

You may wish to prevent users from logging on to nodes that are used to run parallel programs. This will ensure that the resources of the node are devoted to running the parallel program, and are not affected by interactive users. The procedure to restrict access to login nodes is as follows:

- Identify the nodes that are to be used for parallel programs.
- Generally, such nodes correspond to one or more RMS partitions. You should mark these partitions as being of type `parallel`. This means that `rlogin` (and `rsh` access) is disabled on these nodes. The process for marking RMS partitions as `parallel` is described in Chapter 5 of the *HP AlphaServer SC System Administration Guide*.
- Users use a cluster alias, generally the default cluster alias, to log onto the HP AlphaServer SC system. By default, all nodes in a domain are members of (or join) the default cluster alias. Also, the alias attributes are the same on all nodes. This means that when a user logs into the system through the default cluster alias, they may log into any of the nodes within the domain. You can configure the default cluster alias properties so that you control the nodes that the cluster alias mechanism will select.

The process to configure the default cluster alias is as follows:

1. Run the Cluster Alias Manager System Management Menu as follows:

```
# sysman clua_aliases
```
2. Choose DEFAULTALIAS and select Modify...
3. Modify the properties of each member as follows:
 - a. To disable login to a member (node): set the selection weight (`SELW`) to 0 (zero).
 - b. To enable login to a member (node): set the selection weight (`SELW`) to a positive number.

These changes do not take effect until you shut down and boot all nodes. To change alias properties without shutting down and booting, run `cluamgr(8)` on every node of a domain.

General Post-Installation Tasks

For example, to restrict users so that they can only login to Node 0 and Node 1 in a 16-node domain, use the following sample command:

```
# scrun -n 'atlas[0-1]' '/usr/sbin/cluamgr -a alias=DEFAULTTALIAS,selw=3'
# scrun -n 'atlas[2-15]' '/usr/sbin/cluamgr -a alias=DEFAULTTALIAS,selw=0'
```

10.2.2 Back Up and Restore the Management Server Root (/), /usr, and /var File Systems

This section describes how to backup and restore the management server root (/), /usr, and /var file systems.

Specific details for backing up a clustered management server are not included, however, the same basic commands could be extended for that purpose.

10.2.2.1 Backup the Management Server

To back up the management server, perform the following steps:

1. Clone the root, usr, and var file sets as follows:

```
atlasms# clonefsset root_domain root rootclone
atlasms# clonefsset usr_domain usr usrclone
atlasms# clonefsset var_domain var varclone
```
2. Create mount points for the cloned file systems as follows:

```
atlasms# mkdir /rootclonemnt
atlasms# mkdir /usrclonemnt
atlasms# mkdir /varclonemnt
```
3. Mount the cloned file systems as follows:

```
atlasms# mount root_domain#rootclone /rootclonemnt
atlasms# mount usr_domain#usrclone /usrclonemnt
atlasms# mount var_domain#varclone /varclonemnt
```
4. Using a spare disk (assumed to be dsk1c) of equal size to the original boot disk, create a backup file domain as follows:

```
atlasms# mkfdmn /dev/disk/dsk1c ms_backup_domain
```
5. Create backup root, usr, and var file sets as follows:

```
atlasms# mkfset ms_backup_domain root_bkup
atlasms# mkfset ms_backup_domain usr_bkup
atlasms# mkfset ms_backup_domain var_bkup
```
6. Create mount points for the backup file systems as follows:

```
atlasms# mkdir /root_bkup
atlasms# mkdir /usr_bkup
atlasms# mkdir /var_bkup
```
7. Mount the backup file systems as follows:

```
atlasms# mount ms_backup_domain#root_bkup /root_bkup
atlasms# mount ms_backup_domain#usr_bkup /usr_bkup
atlasms# mount ms_backup_domain#var_bkup /var_bkup
```


8. Dump the original data to the backup file systems as follows:
atlasms# **vdump -0f - /rootclonemnt | vrestore -xf - -D /root_bkup**
atlasms# **vdump -0f - /usrclonemnt | vrestore -xf - -D /usr_bkup**
atlasms# **vdump -0f - /varclonemnt | vrestore -xf - -D /var_bkup**
9. Unmount the backup file systems as follows:
atlasms# **umount /root_bkup**
atlasms# **umount /usr_bkup**
atlasms# **umount /var_bkup**
10. Remove the backup mount points as follows:
atlasms# **rmdir /root_bkup**
atlasms# **rmdir /usr_bkup**
atlasms# **rmdir /var_bkup**
11. Unmount the cloned file systems as follows:
atlasms# **umount /rootclonemnt**
atlasms# **umount /usrclonemnt**
atlasms# **umount /varclonemnt**
12. Remove the clone mount points as follows:
atlasms# **rmdir /rootclonemnt**
atlasms# **rmdir /usrclonemnt**
atlasms# **rmdir /varclonemnt**
13. Remove the clone file sets as follows:
atlasms# **rmfset root_domain rootclone**
atlasms# **rmfset usr_domain usrclone**
atlasms# **rmfset var_domain varclone**

After the above commands have been run, the layout of the file systems (assuming that disk0 is the original boot disk) is as follows:

disk partition=dsk0a	domain=root_domain	filesystem=root
disk partition=dsk0d	domain=usr_domain	filesystem=usr
disk partition=dsk0e	domain=var_domain	filesystem=var
disk partition=dsk1c	domain=ms_backup_domain	filesystem=root_bkup filesystem=usr_bkup filesystem=var_bkup

General Post-Installation Tasks

10.2.2.2 Restore the Management Server

To restore the management server, perform the following steps:

1. Connect to the console of the management server and ensure that it is at the SRM prompt.
2. Boot the management server from the CD-ROM (or from RIS if available) as follows:

```
P00>>> set bootdef_dev ""
P00>>> boot dqa0
```
3. Exit to a UNIX shell prompt by selecting option 3 from the installation menu.
4. Create the symbolic link necessary for the backup domain (where it is assumed that `dsk1` is the backup disk used earlier):

```
atlasms# mkdir /etc/fdmns/ms_backup_domain
atlasms# cd /etc/fdmns/ms_backup_domain
atlasms# ln -s /dev/disk/dsk1c
```
5. Create mount points for the backup file systems as follows:

```
atlasms# mkdir /var/tmp/root_bkup
atlasms# mkdir /var/tmp/usr_bkup
atlasms# mkdir /var/tmp/var_bkup
```
6. Mount the backup file systems as follows:

```
atlasms# mount ms_backup_domain#root_bkup /var/tmp/root_bkup
atlasms# mount ms_backup_domain#usr_bkup /var/tmp/usr_bkup
atlasms# mount ms_backup_domain#var_bkup /var/tmp/var_bkup
```
7. Create the symbolic link necessary for the original root domain (where it is assumed that `dsk0` is the original boot disk):

```
atlasms# mkdir /etc/fdmns/root_domain
atlasms# cd /etc/fdmns/root_domain
atlasms# ln -s /dev/disk/dsk0a
```
8. Create the symbolic link necessary for the original usr domain (where it is assumed that `dsk0` is the original boot disk):

```
atlasms# mkdir /etc/fdmns/usr_domain
atlasms# cd /etc/fdmns/usr_domain
atlasms# ln -s /dev/disk/dsk0d
```
9. Create the symbolic link necessary for the original var domain (where it is assumed that `dsk0` is the original boot disk):

```
atlasms# mkdir /etc/fdmns/var_domain
atlasms# cd /etc/fdmns/var_domain
atlasms# ln -s /dev/disk/dsk0e
```
10. Create mount points for the original file systems as follows:

```
atlasms# mkdir /var/tmp/root
atlasms# mkdir /var/tmp/usr
atlasms# mkdir /var/tmp/var
```
11. Remove the original file systems as follows:

```
atlasms# rmfset root_domain root
```

- ```

atlasms# rmfset usr_domain usr
atlasms# rmfset var_domain var

```
12. Recreate the original file systems as follows:
 

```

atlasms# mkset root_domain root
atlasms# mkfset usr_domain usr
atlasms# mkfset var_domain usr

```
  13. Mount the original file systems as follows:
 

```

atlasms# mount root_domain#root /var/tmp/root
atlasms# mount usr_domain#usr /var/tmp/usr
atlasms# mount var_domain#var /var/tmp/var

```
  14. Restore the backed up data to the original file systems as follows:
 

```

atlasms# vdump -0f - /var/tmp/root_bkup | vrestore -xf - -D /var/tmp/root
atlasms# vdump -0f - /var/tmp/usr_bkup | vrestore -xf - -D /var/tmp/usr
atlasms# vdump -0f - /var/tmp/var_bkup | vrestore -xf - -D /var/tmp/var

```
  15. Shut down the management server as follows:
 

```

atlasms# shutdown -h now

```
  16. Boot the management server from the original disk as follows:
 

```

P00>>> set bootdef_dev dka0
P00>>> boot dka0

```

### 10.2.3 Back Up the Cluster Root (/), /usr, and /var File Systems

When creating the cluster (using `sra clu_create clustername`), the `sra` command prompts for, and partitions, the backup cluster disk. However, it does not copy the contents of the cluster disk to the backup cluster disk. You must back up the cluster disk contents.

---

#### Note:

Before creating the cluster, the `/usr` file system and the `/var` file system may be in separate file domains, or they may be in a single file domain. After creating the cluster (using `sra clu_create clustername`), the `/usr` file system and the `/var` file system will be in two separate file domains, even if they were in a single file domain before creation of the cluster.

---

This section describes how to back up the cluster disk. Chapter 2 of the *HP AlphaServer SC System Administration Guide* describes how to boot the cluster using the backup cluster disk created in this section.

To back up the cluster disk, perform the following steps (on the first node of each domain):

1. Make the backup cluster root (/), `/usr`, and `/var` file domains and filesets. For example:

```

atlas0# mkfdmn /dev/disk/dsk5b cluster_root_bkup
atlas0# mkfset cluster_root_bkup root

```

## General Post-Installation Tasks

```
atlas0# mkfdmn /dev/disk/dsk5g cluster_usr_bkup
atlas0# mkfset cluster_usr_bkup usr
atlas0# mkfdmn /dev/disk/dsk5h cluster_var_bkup
atlas0# mkfset cluster_var_bkup var
```

The primary cluster disk now contains the following:

- The `cluster_root` file domain points to the primary cluster disk (`dsk3`).
- The `cluster_root_bkup` file domain represents storage (`dsk5`) that will be used to back up the contents of the `cluster_root` disk.

The primary cluster disk contains similar file domains for `cluster /usr` and `cluster /var`.

2. Make backup mount points for cluster root (`/`), `/usr`, and `/var`. For example:

```
atlas0# mkdir /root_bkup
atlas0# mkdir /usr_bkup
atlas0# mkdir /var_bkup
```

3. Mount the backup cluster root (`/`), `/usr`, and `/var` file systems. For example:

```
atlas0# mount cluster_root_bkup#root /root_bkup
atlas0# mount cluster_usr_bkup#usr /usr_bkup
atlas0# mount cluster_var_bkup#var /var_bkup
```

4. Create a clone fileset for cluster root (`/`), `/usr`, and `/var` file systems. For example:

```
atlas0# clonefset cluster_root root rootclone
atlas0# clonefset cluster_usr usr usrclone
atlas0# clonefset cluster_var var varclone
```

5. Make mount points for the clone filesets:

```
atlas0# mkdir /rootclonemnt
atlas0# mkdir /usrclonemnt
atlas0# mkdir /varclonemnt
```

6. Mount the clone filesets:

```
atlas0# mount cluster_root#rootclone /rootclonemnt
atlas0# mount cluster_usr#usrclone /usrclonemnt
atlas0# mount cluster_var#varclone /varclonemnt
```

7. Back up the cluster root (`/`), `/usr`, and `/var` file systems. For example:

```
atlas0# vdump -0f - /rootclonemnt | vrestore -xf - -D /root_bkup
atlas0# vdump -0f - /usrclonemnt | vrestore -xf - -D /usr_bkup
atlas0# vdump -0f - /varclonemnt | vrestore -xf - -D /var_bkup
```

8. Step 7 performs an exact copy of the contents of the primary cluster disk. Therefore, the `cluster_root`, `cluster_usr`, and `cluster_var` file domains — on the primary cluster disk and on the backup cluster disk — point to the primary cluster disk (`dsk3`).

Change these file domains on the backup cluster disk, so that they point to the backup cluster disk (`dsk5`) instead. For example:

```
atlas0# cd /root_bkup/etc/fdmns/cluster_root
atlas0# rm *
atlas0# ln -s /dev/disk/dsk5b
atlas0# cd /root_bkup/etc/fdmns/cluster_usr
```

- ```
atlas0# rm *
```
- ```
atlas0# ln -s /dev/disk/dsk5g
```
- ```
atlas0# cd /root_bkup/etc/fdmns/cluster_var
```
- ```
atlas0# rm *
```
- ```
atlas0# ln -s /dev/disk/dsk5h
```
9. Since the `cluster_root` and `cluster_root_bkup` file domains on the backup boot disk now point to the same disk (`dsk5`), you must remove the `cluster_root_bkup` file domain on the backup cluster disk. Similarly, you must remove the `cluster_usr_bkup` and `cluster_var_bkup` file domains on the backup cluster disk.
To remove the backup file domains from the backup cluster disk, enter the following commands:

```
atlas0# rm -rf /root_bkup/etc/fdmns/cluster_root_bkup
```

```
atlas0# rm -rf /root_bkup/etc/fdmns/cluster_usr_bkup
```

```
atlas0# rm -rf /root_bkup/etc/fdmns/cluster_var_bkup
```
 10. Unmount the backup cluster root (`/`), `/usr`, and `/var` file systems. For example:

```
atlas0# umount /root_bkup
```

```
atlas0# umount /usr_bkup
```

```
atlas0# umount /var_bkup
```
 11. Remove the mount points for the backup cluster root (`/`), `/usr` and `/var` file systems. For example:

```
atlas0# rmdir /root_bkup
```

```
atlas0# rmdir /usr_bkup
```

```
atlas0# rmdir /var_bkup
```
 12. Unmount the clone filesets for cluster root (`/`), `/usr` and `/var` file systems. For example:

```
atlas0# umount /rootclonemnt
```

```
atlas0# umount /usrclonemnt
```

```
atlas0# umount /varclonemnt
```
 13. Remove the mount points for the clone file systems. For example

```
atlas0# rmdir /rootclonemnt
```

```
atlas0# rmdir /usrclonemnt
```

```
atlas0# rmdir /varclonemnt
```
 14. Use the `file` command to identify the major and minor numbers of the new `cluster_root` device. For example:

```
atlas0# file /dev/disk/dsk5b
```

```
/dev/disk/dsk5b: block special (19/221)
```


Record these numbers; you will need them to boot using the backup cluster disk, if the primary cluster disk fails.

Console Network

10.2.4 Implement Security Recommendations

When configuring the RIS Server, note that RIS will add `/ris/*` and `/var/adm/ris/*` to the `/etc/exports` file. While this has not proved to be a security problem previously, for highly-secure clusters, you may want to remove these exports after the install/upgrade procedure has completed. However, if you remove these exports, they must be added back before you add any additional nodes.

Following the installation of the AlphaServer SC V2.6 software, the management server (or the first cluster when there is no management server) will NFS export the `/var/sra/diag/quadrics` directory to the other domains, and the first node of the other domains will mount this directory using their `/etc/member_fstab` file.

This directory is used during the Interconnect Diagnostics, specifically the `elan_level_test` and the `elansoaktest` diagnostics (refer to the *HP AlphaServer SC Interconnect Installation and Diagnostics Manual* for more information). These diagnostic utilities are run on each node of the system, and their outputs are recorded beneath this directory for subsequent parsing and analysis by the management server.

By default, this exported filesystem can be mounted by systems other than the HP AlphaServer SC domains. While this has not proved to be a security problem previously, security conscious sites may want to lock down these exports, when the installation is complete. For example, the line in the `/etc/exports` file could be modified to contain a space-separated list of aliases for the lead members of each domain as follows:

```
/var/sra/diag/quadrics -root=0 atlas0 atlas0-ext1 atlas1 atlas1-ext1 atlas32...
```

Such a change will mean that the filesystem can only be mounted by those servers, and this will mean that the diagnostic utility will continue to function normally.

10.3 Console Network

The console manager allows you to monitor and connect to the console of each node within the system. It also produces a time-stamped log of each console's output.

The console manager is automatically configured at installation time, and the information is stored in the `sc_cmf` table in the SC Database. This information in the table defines which ports of the terminal server each node is connected to.

Most sites normally have other devices (for example, RAID controllers) that use a serial port for monitoring and control. It may be convenient to use the console manager facilities to do this. Local extensions to the configuration can be stored in the `/var/sra/cmf.conf.local` file. This file has the same format as the standard configuration file:

```
<name> <terminal server hostname> 20<port number>
```

If the console manager detects the existence of this file at startup, it will manage the ports defined therein.

For more details on the console network, see Chapter 14 of the *HP AlphaServer SC System Administration Guide*.

10.4 Storage and File Systems

When installed, the system's mandatory file systems (`/`, `/usr`, and `/var`) are served by the first node in each domain. The second node in each domain is also connected to the system storage, for failover.

Each node has its own local storage configured as a mixture of `tmp`, `local`, and `swap`.

Note:

These local disks are now mounted as “server only”, by using the `-o server_only` flag. Specifying this mount option means that if a node panics or is reset, the local file systems are unmounted. If you do not specify this mount option, the file systems will remain mounted if a node panics or is reset, which can make it difficult to delete a member, or to switch to the alternate boot disk.

However, mounting these disks as “server only” also means that these file systems will not be accessible from other members in the cluster. If you wish to remove the `server_only` mount option, run the following command:

```
atlasms# scrun -d atlasD0 '/usr/sbin/rcmgr delete SC_MOUNT_OPTIONS'
```

You may wish to make the following alterations to your storage system:

1. Populate the storage controller with additional spindles and create file systems on the new storage.
2. Connect additional nodes to the storage array.
3. Add further local storage to individual nodes.
4. Make decisions regarding which node serves which file system.
5. Serve file systems to other domains within the HP AlphaServer SC system using SC File System (SCFS). See Chapter 7 of the *HP AlphaServer SC System Administration Guide* for more information about SCFS.
6. Set up a Parallel File System (PFS). See Chapter 8 of the *HP AlphaServer SC System Administration Guide* for more information about PFS.
7. Serve file systems to external clients.
8. Mount file systems served by external servers.

User Administration

Note:

Each node within a domain is *theoretically* a server to all other nodes within the domain. It is automatically a client of all file systems within the domain. We use the term *theoretically* here because, under typical operation, a small subset of nodes *actually* serves file systems to other nodes.

For example, as configured, the first node of the domain serves the `/`, `/usr`, and `/var` file systems to all other nodes.

10.5 User Administration

Immediately after the installation, no users (with the exception of compulsory system users) have been added to the system.

If your site is using NIS for user administration, then it is not necessary to set up user accounts. NIS, if required, will have been configured at installation time.

If you do not have an external NIS server, then consistent user administration data must be available on each domain.

Note:

User administration data will be consistent within a domain.

To ensure that user administration data is consistent on each domain, perform either of the following tasks:

- On the first domain, use the system administration tools for user administration and replicate the relevant files (`/etc/passwd`, `/etc/group`) to the other domains. If enhanced security is configured, use the `convuser` and `convauth` utilities to extract data from one domain and transfer it to another.
- Nominate one domain as an internal NIS master, and configure the other domains as NIS slave servers.

As well as adding users to the UNIX configuration files, you must register new users with the RMS system. RMS users may operate as standalone users or as members of projects. Therefore, registering new users with RMS may result in updating three RMS database tables — the `users` table, the `projects` table, and the `access_controls` table — as follows:

- The `users` table lists all registered RMS users. It also records project membership for each user. Every user is automatically a member of the default projects.

- The `projects` table lists all registered projects. Users can submit jobs as themselves or as a member of a specific project.
- The `access_controls` table specifies resource limits that apply to a user or project.

The tables can be manipulated by using SQL commands or by using the `sra_user` graphical utility.

SQL commands can be useful when used in site-specific tools to convert batches of users. For more details about RMS administration, see Section 10.7 on page 10–13 of this manual, and Chapter 5 of the *HP AlphaServer SC System Administration Guide*.

10.6 Cluster Aliases

Each domain within the HP AlphaServer SC system can be addressed by a single IP alias. This is called the default alias. For example, in a 64-node system called `atlas`, the default alias of the first domain is `atlasD0`; the default alias of the second domain is `atlasD1`.

The alias `atlasD0` is a special IP address that can refer to any of the first 32 nodes. You may wish to create further aliases; for example, you may wish to use the first four nodes for login/interactive and program development work.

To do this, create an alias (for example, `atlas-dev`), and make the first four nodes members of this alias. The cluster alias feature is quite flexible, you can assign different weights and priorities to different nodes. This can be used to preferentially select specific nodes, or to only select some nodes when others are available.

The cluster alias facility is discussed in more detail in Chapter 19 of the *HP AlphaServer SC System Administration Guide*.

10.7 RMS

For a recently installed system, there are two types of RMS system administration tasks:

- Mandatory RMS Administration Tasks (see Section 10.7.1 on page 10–13)
- Optional RMS Administration Tasks (see Section 10.7.2 on page 10–14)

10.7.1 Mandatory RMS Administration Tasks

The following RMS administration tasks are mandatory:

1. Define and set up partition(s).
2. Start the partition(s).

3. Configure the system to notify the system administrators when significant events occur. For example, an notification might be sent for an RMS partition state transition. Refer to the *HP AlphaServer SC System Administration Guide* for instructions on configuring event notifications.
4. Back up the RMS database.

The installation process creates one project, named `default`. When a user uses `prun` to run a job without specifying a partition, the job is automatically run as a member of the `default` project. Without making any changes, any UNIX user of the system can run parallel jobs on the system. See Chapter 5 of the *HP AlphaServer SC System Administration Guide* for further details.

10.7.2 Optional RMS Administration Tasks

You may wish to configure various RMS tables to reflect site policy, by performing some or all of the following tasks:

1. Add users to the RMS database.
You can use the graphical utility `sra_user` to import users from the `/etc/passwd` file.
2. Define projects in the RMS database.
As mentioned earlier, the installation process creates one project, named `default`. You can create additional projects.
3. Set up access controls for users and projects in the RMS database.
You can limit access to certain partitions. Within a partition, you can specify CPU and memory limits.
4. Define partition types.
Partition types specify the way in which the partition should be used (for example: interactive development, parallel programs only, both).
5. Define partition attributes.
Partition attributes control resource usage and apply usage limits to users of the partition.
6. Update the RMS database to specify which users should be notified when the following error conditions are found:
 - Node status has changed (set the `email-node-status` attribute)
 - Partition is blocked (set the `email-partition-blocked` attribute)
 - Temperature has changed by more than 2°C (set the `email-module-tempwarn` attribute)
 - Event handler did not complete in time (set the `email-event-escalate` attribute)
7. Create a policy for cleaning old entries from the database.

You can use the graphical utility `sra_user` to manipulate RMS database tables.

Alternatively, you can use SQL commands to manipulate the individual tables directly. See Chapter 5 of the *HP AlphaServer SC System Administration Guide* for further details.

Troubleshooting

This chapter describes some of the common problems found when installing an HP AlphaServer SC system, and provides solutions to these problems. See also the "Known Problems" section of the *HP AlphaServer SC Release Notes* and the Troubleshooting chapter of the *HP AlphaServer SC System Administration Guide*.

The information in this chapter is structured as follows:

- Tips for Installing an HP AlphaServer SC System (see Section 11.1 on page 11–2)
- InstallSC Errors (see Section 11.2 on page 11–5)
- Interpreting Problems During Software Installation (see Section 11.3 on page 11–5)
- Boot Errors (see Section 11.4 on page 11–17)
- Adding RIS Client With No Default Route (see Section 11.5 on page 11–25)
- Corrupt .member.list File Causes Core Dumps (see Section 11.6 on page 11–25)
- Adding a New Member Fails (see Section 11.7 on page 11–26)
- sra setup Errors (see Section 11.8 on page 11–26)
- Terminal Server Errors (see Section 11.9 on page 11–27)
- rinfo Command Displays UID or Wrong User Name (see Section 11.10 on page 11–28)
- How to Drop and Rebuild the RMS Database (see Section 11.11 on page 11–29)
- rcontrol Reports Errors During Node Boot (see Section 11.12 on page 11–29)
- rcontrol Reports Error When Starting Partitions (see Section 11.13 on page 11–30)
- clu_get_info Prints CONFIGURATION_ERROR (see Section 11.14 on page 11–30)
- How to Powercycle an HP AlphaServer SC System (see Section 11.15 on page 11–30)
- Error When Installing the Elan Subset on a Management Server (see Section 11.16 on page 11–32)

Tips for Installing an HP AlphaServer SC System

- Database Access Denied Errors (see Section 11.17 on page 11–32)
- Database Access Denied Error on Some Domains (see Section 11.18 on page 11–33)
- Diagnosing Federated Network Routing Problems (see Section 11.19 on page 11–34)
- Wakeup on LAN (WOL) Problem (see Section 11.20 on page 11–34)
- scfsmgr/pfsmgr report Could Not Open Socket (see Section 11.21 on page 11–35)
- Problem with HSG Devices in Installation Process (see Section 11.22 on page 11–35)
- Upgrade Errors (see Section 11.23 on page 11–37)
- Interpreting Problems During Software Upgrade (see Section 11.24 on page 11–40)
- Increasing the Number of ptys (see Section 11.25 on page 11–45)
- Increasing Socket Listen Queue Limits (see Section 11.26 on page 11–46)

11.1 Tips for Installing an HP AlphaServer SC System

This section contains various helpful tips and guidelines that may assist you when performing the reinstallation of an HP AlphaServer SC system. The steps in this section are designed to interleave with the usual steps of Chapter 5 and Chapter 6 as opposed to being a replacement for those chapters. If you are uncertain about any steps, you should consider Chapter 5 and Chapter 6 to be the authority.

1. Before taking the existing system out of production, perform a database backup using the `rmsbackup` utility. Archive the database backup file (the newest file in `/var/rms/backup`) offline, so that it is not affected by the management server re-installation.
2. If you have a clustered management server, then once the SC software has been installed or upgraded, and the management server is running from the rebuilt kernels, you should actually shut down the member2 node while you build the sra database and install or upgrade the nodes. For a clustered management server or a cluster with no management server, then after installing the first domain, you should shut down all nodes, except the lead node in the domain, before installing or upgrading the other domains. This action will eliminate any risk of complex routing problems occurring with bootp packets, where the dhcp CAA service is accidentally located on member2 while the gateway entry in the `bootptab` file will indicate member1.
3. Immediately following the `sra setup` step, you should back up the new sra database with the `rmsbackup` utility. Perform the following checks:
 - a. Perform a line-by-line comparison of the earlier database backup with the latest backup, with a focus on the following tables: `sc_nodes`, `sc_system`, `sc_networks`, `sc_system_devices`, `attributes`, and `sc_domain`.

Tips for Installing an HP AlphaServer SC System

- b. Perform a line-by-line comparison of the earlier database backup with the latest backup, with a focus on the following tables: `sc_scfs`, `sc_pfs`, and their associated tables.
4. Review the SRA database as follows:
 - a. Check the generic *swap*, *tmp*, and *local* settings in the sra database using the *sys* menu option within the `sra edit` command, and review the boot-first and boot-second image settings.
 - b. Check the IP addresses for the corporate Ethernet connections on any nodes with a corporate connection.
 - c. Check that the router, netmask, mail server, NIS server, DNS server settings as defined by *show sys* in the *sys* menu option within the `sra edit` command are accurate and can be pinged.
 - d. Check the seed addresses for the preferred server aliases.
5. In order to avoid problems with the probe ordering of local SCSI disks versus disks presented *via* Fibre Channel, some additional steps can be taken. While the compute nodes are at SRM (that is, before dispatching the `sra install` commands on the management server), perform the following changes:
 - a. For SC40/SC45 systems (not SC20 systems), run the following commands in the order shown to ensure that there are no bootable HSG/HSV devices:

```
# sra command -domains all -member 1 -command init -limit no
# sra command -domains all -member 1 -command wwidmgr -clear all \
-limit no
# sra command -domains all -member 1 -command init -limit no
```
 - b. Disconnect the FC cables (or disable LUN access) to the member1 nodes (lead members) to the Cluster System storage.
 - c. Disconnect the FC cables (or disable LUN access) to the member2 nodes to the Cluster System storage.
 - d. Disconnect the FC cables (or disable LUN access) to all SCFS server nodes from the relevant EMA/EVA storage.
 - e. If you have performed any type of SCSI disk movements, then run the command:

```
# sra reset -nodes all
```

Note:

Alternatively, you could have allowed the `sra setup` command to initialize the nodes when it prompted you to do this.

Tips for Installing an HP AlphaServer SC System

6. On the lead members, before they are installed with SC software, check that the UNIX devices as seen from the RIS environment have the expected UNIX special device names. This check determines if `dsk0` is the UNIX special device name for the primary boot disk, and if `dsk1` is the UNIX special device name for the alternative boot disk, and to see if `dsk2` exists in readiness for the standalone UNIX image. To do this, perform the following steps:
 - a. Run the following commands (the first command will ensure that an old `hwmgr` database is not used — this variable is typically cleared before a member is installed):

```
# sra script -domains all -members 1 -command set bootdef_dev
# sra script -domains all -members 1 -script rishwmgrviewdevices
```
 - b. When finished, check the list of disks reported by the script, and, if necessary, you can search the console logs of the lead members for the outputs to the resulting `hwmgr` command.
 - c. Where you see anomalies, the hardware layout should be changed to ensure a more correct probe order, or alternatively, the generic entry (or node-specific entries) in the `sra edit` command should be changed to reflect the reported probe order.
7. On the non-lead members (members2 through 32), before they are installed with SC software, check that the UNIX devices as seen from the RIS environment have the expected UNIX special device names. You can use the same SRA script as described in the previous step. This check determines if `dsk0` is the UNIX special device name for the primary boot disk, and if `dsk1` is the UNIX special device name for the alternative boot disk (note that `dsk2` is not used for SC purposes on the non-lead members). As mentioned in the previous step, where you see anomalies, the hardware layout should be changed to ensure a more correct probe order, or alternatively, the generic entry (or node-specific entries) in the `sra edit` command should be changed to reflect the reported probe order.
8. When running the initial `sra install` command for the domains, include the parameter `-end_state UNIX_Config` with the `sra install` command, so that the automation will stop after the UNIX configuration is complete.
9. Later, when the installation command arrives at the `end_state` of `UNIX_Config` on all domains (after approximately 1 hour), reconnect the FC cables (or enable LUN access) to the member1 nodes (lead members) to the Cluster System storage.
10. Run the initial `sra install` command again and exclude the `end_state` parameter (ensuring that the `redo` parameter is not included). The command will automatically continue from the `UNIX_Config` state, and proceed to apply the patch kits onto the domains.
11. When the `sra install` automation command has finished (approximately 8 hours) and all members are added to the domains, perform the following tasks:

- a. Reconnect the FC cables (or enable LUN access) to the member2 nodes to the Cluster System storage.
 - b. Reconnect the FC cables (or enable LUN access) to all SCFS server nodes from the relevant EMA/EVA storage.
12. Continue remaining steps as described in the installation chapters of this guide, (and for systems with clustered management servers, you should now boot member2 of the clustered management server).

11.2 InstallSC Errors

If any errors are noted when running `InstallSC` while upgrading the management server in Chapter 4, while building the Management Server in Chapter 5, or while building Node 0 in Chapter 6, then the errors must be addressed before continuing with the installation.

Possible errors that may occur when running the `InstallSC` command are as follows:

- Failure to create the correct RIS area prior to running `InstallSC`. The correct RIS area should be created as instructed in the appropriate chapter. Once the RIS area is created, the `InstallSC` command should be run again. By running the command again, you will quickly step through completed steps to the point where it had failed earlier.
- Failure of the RMS subset installation to add the rms UNIX user or UNIX group. In this case, the `InstallSC` command should be run again with the `-remove` option. The ground rules for NIS should be checked in Configure NIS (see Section 5.1.5.6 on page 5–19) or Configure NIS (see Section 6.1.4.6 on page 6–16). Once the problem is solved, then the `InstallSC` command can be run again with the `-install` option.

The same `InstallSC` utility is used when updating the HP AlphaServer SC software on the domains. During this process, there may be a message to indicate that the RIS area does not exist. This is a benign warning on a domain because the only RIS area required is the one on the management server.

11.3 Interpreting Problems During Software Installation

The installation of a management server is effectively a manual process and the steps required are fully described in Chapter 5. The steps are clearly explained and if problems occur, the diagnosis should be straightforward.

Once the management server is installed, then the subsequent installation of the domains is an automated process that is controlled and monitored by various sra scripts. There is a state machine that controls this automated process. A basic description of the installation state mechanism is provided in Understanding the Automated Installation Process (see Section 7.1 on page 7–1).

Interpreting Problems During Software Installation

In the following subsections, further information is provided that may assist you in resolving problems that may arise during each state of the installation process.

The installation status of each node can be reported using the command:

```
sra install_info -node atlas[0-31]
```

This command will report the node's status as being *one* of the following:

- UNIX_Installed
- UNIX_Config
- UNIX_Patched
- SC_Installed
- SC_Patched
- NHD_Installed
- CLU_Create
- CLU_Added
- Bootp_Loaded
- Member_Added

Having initially established that a node is in a particular state, it is then necessary to understand if any command is currently in progress or queued for a given node. This command information is visible from the `sra install_info` output command above, and the command information can also be reported using the following command:

```
sra command_info
```

The advantage of the above command is that it will list all commands queued and in progress without needing to specify what nodes might be involved.

Where a command has recently stopped as a result of an error, then an error message will be indicated in the status column of either reports above. It is anticipated that the status message will be descriptive enough to indicate the cause of the error. However, for more complex problems, it may be necessary to understand the typical actions that happen during each installation state, and to be aware of the log file locations in order to determine the cause of the error. Examples of such information are provided in the following subsections.

A node will remain in the current state until all actions required to progress to the next state are successfully completed. In the case of some states, these actions can be repeated after command errors until such time that the automation process succeeds and it can move to the next state. However, depending on the error, and depending on the current state, it might not be possible to repeat the actions. Examples of such scenarios are outlined in the following subsections.

11.3.1 Uninstalled State

The current state is that no Tru64 UNIX or HP AlphaServer SC software has been installed on the domain and all nodes the domain should be at the SRM prompt.

Normal actions in this state are to RIS-boot each lead member and install the Tru64 UNIX base operating system on the UNIX disk. The UNIX disk is usually the third local SCSI disk in each lead-member and the SRM device name and UNIX special device name of this disk will have been specified during the `sra setup` step.

Typical problems during this process are that the RIS boot fails to happen at all or fails to complete correctly. Typical causes of problems are as follows:

1. Console logger daemon not running.
2. Incorrect management Ethernet SRM device name defined in the SC database for this node.
3. Ethernet cable problems on the management network.
4. Incorrect Ethernet MAC address registration of the nodes by `sra setup` in the `/etc/bootptab` file on the management server.
5. Joind running on member2 on clustered management servers.
6. Incorrect `sa` or `gw` entries in the `/etc/bootptab` file on clustered management servers.
7. Missing `/etc/exports` entries for the RIS areas.
8. Incorrect UNIX disk specification for a particular domain.
9. Probe order causes HSV/HSG-based LUNs to be seen earlier in probe order than local SCSI disks.
10. Timeouts during subset installation because of problems with speed or duplex settings on management network.

The typical log file locations that you can access to diagnose problems in this state are as follows:

- Console logger daemon logfile in `/var/sra/adm/log/cmfd` on the management server.
- Lead member console log in `/var/sra/logs` on the management server.
- Logs for the system-level `sra` daemon in `/var/sra/adm/log/srad/srad.log` on the management server.

Interpreting Problems During Software Installation

As a debugging aid, it may be useful to RIS-boot the offending node manually from the SRM prompt using the command `boot -p bootp eia0`. It might also be useful to run the `joind` daemon in the foreground with debugging enabled, using the command `joind -f -d5` on the management server.

A further debugging aid would be to run the `sra ethercheck` diagnostic on the relevant node. Once the cause of the problem is understood and resolved, then the lead member can be returned to the SRM prompt, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.2 UNIX_Installed State

The current state is that the Tru64 UNIX operating system has been installed on the first node of the domain and the lead member of the domain should be booted from the UNIX disk. Normal actions in this state are to configure UNIX services and certain other parameters on the lead member of each domain. This work involves completing the following steps in the order shown:

1. Log into the lead member while booted from the UNIX disk.
2. Configure the Ethernet interfaces.
3. Populate the `/etc/routes` file.
4. Set the GATED parameter in the `/etc/rc.config` file.
5. Enable the NFS server and client services.
6. Populate the `/etc/fstab` file.
7. Populate the `/etc/hosts` file.
8. Configure DNS.
9. Configure NIS.
10. Configure NTP.
11. Configure Mail.
12. Configure LMF licence management.
13. Append the `/var/sra/diag/quadrics` NFS mount to the `/etc/fstab` file.
14. Set the `SC_MS` parameter in the `/etc/rc.config.common` file.
15. Populate the `/var/cookies/root` cookie file.

Interpreting Problems During Software Installation

If there are problems with actions in this state, then the cause can usually be determined by checking the lead member console log in `/var/sra/logs` on the management server, or by checking logs for the `sra` daemon in `/var/sra/adm/log/srad/srad.log` on the management server.

The parameters used when configuring the services listed above are replicated from the same services on the management server. During `sra setup`, the parameter values currently in use on the management server are recorded in the `sc_system` table.

Problems in this state are often caused by these parameters being incorrect, perhaps, because the real parameters on the management server have changed since `sra setup` established a baseline set of parameters in the `sc_system` table.

If required, then these parameters can be edited using the `sys` menu within `sra edit`, as described in Review the SC Database System Settings (see Section 7.2 on page 7–14) prior to configuring the domains.

Also, this will be the first time that the external Ethernet interface of the lead member of the domain will have been used. It is important for the NIS and Mail steps that connectivity is available from the lead member, through the router and onwards to the NIS server, and the Mail relay. In particular, for the Mail setup, it is important that connectivity to the DNS server is also available from the lead member.

Typically, problems in this state will result in timeouts during the automation and the lead member will be reset and revert to the SRM prompt. However, once the cause of the problem is understood and resolved, then the node can be booted manually to multi-user mode from the UNIX disk, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.3 UNIX_Config State

The current state is that the Tru64 UNIX operating system has been configured on the first node of the domain, and the lead member of the domain should be booted from the UNIX disk. The normal action for this state is to install the UNIX patch software. This work involves completing the following steps in the order shown:

1. Log in the lead member while booted from the UNIX disk.
2. Shut the lead member down to the SRM prompt.
3. Boot the node to single-user mode using the UNIX disk.
4. Start the Internet Services on the node.
5. Check for pre-existing mounted patch directory.
6. Mount the patch directory exported from the management server.

Interpreting Problems During Software Installation

7. Run the `dupatch` script to install the patch kit.
8. Reboot the node.
9. Remove the mount point used for the patch directory.

If there are problems with actions in this state, then the cause can usually be determined by checking in any of the following locations:

- The lead member console log in `/var/sra/logs` on the management server.
- Logs for the `sra` daemon in `/var/sra/adm/log/srad/srad.log` on the management server.
- The `/var/adm/patch/log/session.log` file and the `/sys/<system_name>/errs` files on the lead member.

Typical problems in this state are associated with Ethernet connectivity to the management server. The management network is used when the management server is unclustered, and the external network is used when the management server is clustered.

One occasional problem seen in the past occurred due to failure of the Ethernet card and Ethernet switch port to maintain negotiation. In the case of 100Mbps-Full Duplex connections, then this problem is remedied by setting the switch to use this speed instead of auto-negotiate, and adding the following line to the `/etc/inet.local` file:

```
/sbin/ifconfig ee0 down speed 200 up
```

When you add this line, the Internet services should be restarted with the `rcinet restart` command.

If this state had problems during the `dupatch` session itself, and after the stage where the CAUTION message is displayed warning about interruptions to the process, then the lead member will likely be damaged to the point where this domain will need to be reinstalled using the `-redo Uninstalled` parameter of the `sra install` command. If this is the case, then the node should first be returned to the SRM prompt manually.

For other problems in this state, once the cause of the problem is understood and resolved, then the node can be booted manually to multi-user mode from the UNIX disk, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.4 UNIX_Patched State

The current state is that the Tru64 UNIX operating system patch software has been installed on the first node of the domain. The normal action for this state is to install the HP AlphaServer SC software. This work involves completing the following steps in the order shown:

Interpreting Problems During Software Installation

1. Log into the lead member while booted from the UNIX disk.
2. Invoke the switch from SVC to NSS configuration files for the DNS service.
3. Stop the `sendmail` daemon to avoid interruption of the automation.
4. Check for the pre-existing mounted `sckit` directory.
5. Mount the `sckit` directory exported from the management server.
6. Install the individual SC subsets in turn, by first checking that they are not already installed, then performing the installation, and finally checking that they are correctly installed.
7. Unmount and remove the mount point used for the `sckit` directory.
8. Start the `sra` daemon.
9. Restart the `sendmail` daemon.

If there are problems with actions in this state, then the cause can usually be determined by checking in any of the following locations:

- The lead member console log in `/var/sra/logs` on the management server.
- Logs for the system-level `sra` daemon in the `/var/sra/adm/log/srad/srad.log` file on the management server.

The most frequent problem experienced during this stage of the process is with access to the `mysql` daemon on the management server, which will affect the ability to start the `sra` daemon successfully. Possible causes include problems with Ethernet connectivity between the lead member and the management server on both external and management networks, or an incompatible `mysql` cookie file `/var/cookies/root` on this node with that of the management server.

Another possible problem at this stage of the process is related to NIS, and the requirement that user `id 15` and group `id 200` be unused, and that they are available for configuration as `local users/groups` by the RMS component.

Once the cause of the problem is understood and resolved, then the node can be booted manually to multi-user mode from the UNIX disk, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

Interpreting Problems During Software Installation

11.3.5 SC_Installed State

The current state is that the HP AlphaServer SC software has been installed on the first node of the domain. There are currently no actions defined for this state. This state was a placeholder for functionality that was not implemented. Therefore, the automation will immediately proceed to the next state.

11.3.6 SC_Patched State

The current state is that the HP AlphaServer SC patch software has been installed on the first node of the domain. The normal action for this state is to install the New Hardware Delivery (NHD) kit. This kit was required in previous HP AlphaServer SC software releases and specifically for the SC20 product. However, more recent HP AlphaServer SC software releases are based on the Tru64 V5.1B kit, which includes all base software required for all supported products. This means that the installation of the NHD kit is no longer required.

Therefore, assuming that the `-nhdkit` parameter to the `sra install` command was included, then the automation will immediately proceed to the next state.

11.3.7 NHD_Installed State

The current state is that the NHD software has been installed on the first node of the domain. Although the NHD kit is no longer required for any HP AlphaServer SC products, this state of the state machine will still be used briefly during all installations.

The actions during this state are to create a single node cluster from the lead member of each domain. This work involves completing the following steps in the order shown:

1. Dynamically load the Elan3 kernel module on the lead member.
2. Use the `/usr/sys/conf/.product.list` file to populate the member-specific list file.
3. Create a disk label on the lead member's primary boot disk.
4. Create a disk label on the lead member's alternative boot disk.
5. Run the `/usr/sbin/clu_create` command on the lead member, which in turn creates the cluster `root`, `/usr` and `/var` file domains on the cluster disk, and populates these domains with data from the UNIX disk.
6. Reboot the node from its primary boot disk as a single node cluster.
7. Perform the first kernel build on the lead member.
8. Activate the various AlphaServer SC special routing mechanisms as described in chapters 19 and 22 of the *HP AlphaServer SC System Administration Guide*.
9. Activate the HP AlphaServer SC binary error logging feature.

Interpreting Problems During Software Installation

10. Continue booting until a UNIX prompt is displayed.

Typical problems associated with this state are as follows:

- Cabling problems or PCI card problems with the Elan3 interconnect card.
- Probe order causes HSV/HSB RAIDed LUNs to be seen earlier in probe order than local SCSI disks, therefore affecting the selection of primary or alternative boot disks.
- Default swap, local, or tmp settings for primary or alternative boot disks in the SC database are unsuitable.
- Probe order causes HSV/HSB RAIDed LUNs to be seen earlier in probe order than local SCSI disks, therefore affecting the selection of the cluster disk.
- Persistent reservations on the HSV/HSB-based LUNs causing `/usr/sra/bin/createclusterlabel` to fail to create a disk label because of i/o errors.
- Incorrect network settings: subnet mask, router, or hosts file.

If there are problems with actions in this state, then the cause can usually be determined by checking in the lead member console log in the `/var/sra/logs` file on the management server, or by checking the `/var/sra/adm/log/srad/srad.log` file and the `/cluster/admin/clu_create.log` files on the lead member.

All actions associated with this state are controlled by the `srad` daemon on the lead member. The work of that daemon does not complete until the node reaches the UNIX prompt following clusterization. At that point, the current state is updated to the next state.

However, network configuration problems and/or cluster aliases can prevent the `sra` daemon updating the state of the domain in the SC database. This leaves the current `sra` automation command unallocated, yet, it will appear that the domain has been clusterized. Such networking issues need to be resolved, the lead-member needs to be booted manually from the UNIX disk, and automation run again using the `-redo NHD_Installed` parameter of the `sra install` command.

A further debugging aid for cabling problems or PCI problems is to bring the respective lead member to the SRM prompt, and then run the `sra elancheck` diagnostic on the relevant node from the management server.

For other problems in this state, once the cause of the problem is understood and resolved, then the node can be booted manually to multi-user mode from the UNIX disk, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

Interpreting Problems During Software Installation

11.3.8 CLU_Create State

This is a state associated with a non-lead member. The current state is that the lead member of the domain has been clusterized and that the respective non-lead member has not yet been added to the domain.

The actions associated with this state are to add further non-lead members to the single node cluster. The lead member is already part of the cluster, and as such, the automation will immediately move the lead member from this state to the `Member_Added` state.

For the non-lead members, perform the following steps:

1. Create a disk label on the generic boot disk on the lead member.
2. Run the `/usr/sbin/clu_add_member` command on the lead member (targeting the non-lead member), which in turn uses the generic boot disk as a temporary member boot disk that allows the script to create a temporary member specific root domain. Once the member root domain is complete, the `clu_add_member` command continues to populate the cluster file systems with the member-specific files for the new member.
3. Create a `generic_boot_partition` for the new member.
4. Populate the `/etc/clu_alias.config` file with data for the new member.

If there are problems with actions in this state, then the cause can usually be determined by checking in the non-lead member console log in `/var/sra/logs` on the management server, or by checking the `/var/sra/adm/log/srad/srad.log` file and the `/cluster/admin/clu_add_member.log` files on the lead member.

Typically, this step is only affected by problems in accessing the generic boot disk such as basic connectivity, size, or persistent reservations. However, the UNIX disk is used temporarily during this state, and basic connectivity problems or an inappropriate disk label on the UNIX disk will cause the process to fail.

Also, if there are any non-responding NFS mounts on the lead member of the target domain, then this can cause the member add process to stall until the issue with non-responding mounts is resolved.

Once the cause of the problem is understood and resolved, then the non-lead member should be returned to the SRM prompt, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.9 CLU_Added State

This is a state associated with a non-lead member. The current state is that the lead member of the domain has been clusterized and that the respective node has been added to the domain. During this state, the non-lead member is RIS-booted so that the local SCSI primary and alternative boot disk can be populated with the root partition information.

For the non-lead members, perform the following steps:

1. Create a `generic_boot_partition` generic boot disk on the lead member.
2. RIS-boot the non-lead member over the management network.
3. Mount the `/cluster/admin` directory from the lead member.
4. Run the `buildbootdisk` command on the non-lead member to automatically provide the pertinent information about this member, which in turn partitions and populates the primary and alternative boot disks.
5. Shut the non-lead member down to the SRM prompt.

Typical problems during this process are that the RIS boot fails to happen or fails to complete correctly. Typical causes of problems are as follows:

- Console logger daemon not running.
- Incorrect management Ethernet SRM device name defined in the SC database for this node.
- Ethernet cable problems on the management network.
- Incorrect Ethernet MAC address registration of the nodes by running the `sra setup` command in the `/etc/bootptab` file on the management server.
- Joind running on member2 on clustered management servers.
- Incorrect `sa` or `gw` entries in the `/etc/bootptab` file on clustered management servers.
- Missing `/etc/exports` entries for the RIS areas.
- Incorrect generic disk specification for a particular domain.

The typical log file locations that you can access to diagnose problems in this state are as follows:

- Console logger daemon logfile in `/var/sra/adm/log/cmfd` on the management server.
- Non-lead member console log in `/var/sra/logs` on the management server.
- Logs for the domain-level `sra` daemon in the `/var/sra/adm/log/srad/srad.log` file on the relevant domain.

Interpreting Problems During Software Installation

As a debugging aid, it may be useful to RIS-boot the offending node manually from the SRM prompt using the command `boot -p bootp eia0`. It might also be useful to run the `joind` daemon in the foreground with debugging enabled using the command `joind -f -d5` on the management server. A further debugging aid would be to run the `sra ethercheck` diagnostic on the relevant node.

If you wish to repeat the step in question, you can elect to delete the failing member at this point using the `sra delete` command, and attempt the process again using the `sra install` command.

Once the cause of the problem is understood and resolved, then the non-lead member should be returned to the SRM prompt, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.10 Bootp_Loaded State

This is a state associated with a non-lead member. The current state is that the lead member of the domain has been clusterized and that the respective node has been RIS-booted, and the boot partitions have been downloaded to each node in the domain.

During this state, the non-lead members are booted in turn using the `genvmunix` kernel from their newly created member boot disks.

Typical problems during this process are that the RIS boot fails to happen, or it fails to complete correctly. Typical causes of problems are as follows:

- Failure to boot from the primary boot disk because of incorrect SRM value defined in the SC database for this node.
- Boot disk built on incorrect disk because of incorrect UNIX special device name defined in the SC database. This may in turn have been caused by the probe order of HSV/HSG-based LUNs compared with the probe order of local SCSI disks.
- Failure to establish position on the interconnect due to faulty Elan3 PCI card, or bad connectivity of Elan3 cable.
- Kernel panic during boot from `genvmunix` caused by failure to deploy generic kernels following any recent point patches (this can happen when adding a member later in the lifetime of a system).

The typical log file locations that you can access to diagnose problems in this state are as follows:

- Console logger daemon logfile in `/var/sra/adm/log/cmfd` on the management server.
- Non-lead member console log in `/var/sra/logs` on the management server.

- Logs for the domain-level sra daemon in the `/var/sra/adm/log/srad/srad.log` file on the relevant domain.

As a debugging aid for problems relating to disk probe order, it may be useful to RIS-boot the offending node manually from the SRM prompt using the command `boot -p bootp eia0` and exit from the RIS session to the UNIX shell. Then at the UNIX shell, use the `/sbin/hwmgr -view devices` command to establish the correct UNIX special device names for the disks. Later, at the SRM prompt, you can use the `show dev` command to identify the corresponding SRM names for these same disks, always remembering the same `bus/target/lun` values. A further debugging aid for Elan3-related problems, would be to run the `sra elancheck` diagnostic on the relevant node.

For most of the typical problems experienced in this state, it may be necessary to delete the failing member at this point using the `sra delete` command, resolve the underlying problem, and later, attempt the process again using the `sra install` command.

Once the cause of the problem is understood and resolved, then the non-lead member can be returned to the SRM prompt, and the `sra install` command can simply be run again on the management server. The automation process will continue from the current state.

11.3.11 Member_Added State

All members have been added (to the domain), booted, and shut down to the SRM prompt. This is the final state of every member during the installation process, and there are no further actions.

11.4 Boot Errors

This section describes the following boot errors:

- RIS Boot Failures (see Section 11.4.1 on page 11–18)
- RIS Boot Reports a Bootstrap Failure (see Section 11.4.2 on page 11–19)
- RIS Boot Failure: Cannot Determine RIS Home Directory (see Section 11.4.3 on page 11–19)
- Failed to RIS Boot a Node When Attempting to Add It to the Cluster (see Section 11.4.4 on page 11–20)
- Cannot Communicate with the CAA Daemon (see Section 11.4.5 on page 11–21)
- New Member Fails to RIS Boot (see Section 11.4.6 on page 11–21)
- New Member Fails After RIS Boot (see Section 11.4.7 on page 11–21)
- RIS: Boot Error (see Section 11.4.8 on page 11–22)

Boot Errors

- sra boot Has Timeout Error (see Section 11.4.9 on page 11–22)
- Elan Error During Node Boot (see Section 11.4.10 on page 11–23)
- Error During Node Boot (see Section 11.4.11 on page 11–23)
- Node Hang During RIS Boot (see Section 11.4.12 on page 11–24)

11.4.1 RIS Boot Failures

The `Bootstrap-fail` error indicates that the system has failed to boot from its network interface. Generally this is because the system has received no response from its BOOTP requests.

Solution:

Perform the following checks:

- a. Check that the network cables are connected correctly.
- b. Check that the network interface is set to the correct speed (use the `show eia0_mode SRM Console` command).
- c. Check that the corresponding port in the FastEthernet switch is configured correctly.
- d. Check that spanning trees are disabled on the FastEthernet switch.
- e. Check that the `joind` daemon is running on the RIS server (the management server, if used, or the first node of the cluster). If not, restart as follows:
`# /sbin/init.d/dhcp start`
- f. For a clustered management server, ensure that the non-lead member of the clustered management server is shut down as described in Section 5.1.7.11.
- g. Ensure that clients have been added to the RIS database.

Note:

If you abort an install command (`sra abort -command cmd-id`) while the install is in progress, the abort procedure will attempt to return the node, which may be booting from RIS, to the correct state.

- h. Check that all disks are in their original disk drive.

Note:

If you move disks, even on the same node, the system may not boot. This is because the UNIX disk number is associated with the disk, and not with the disk drive. When you move a disk, the file domain expects the disk number of the original disk; as this number is now different, the system fails to boot. Therefore, do not move disks in an HP AlphaServer SC system.

11.4.2 RIS Boot Reports a Bootstrap Failure

The RIS boot may report a bootstrap failure, as follows:

```
16/Mar/2001 00:09:41 Broadcasting BOOTP Request...
16/Mar/2001 00:09:41 Received BOOTP Packet File Name is: /ris/r0k1
16/Mar/2001 00:09:41 local inet address: 10.128.0.47
16/Mar/2001 00:09:42 remote inet address: 10.128.101.1
16/Mar/2001 00:09:42 TFTP Read File Name: /ris/r0k1
16/Mar/2001 00:09:42 netmask = 255.255.0.0
16/Mar/2001 00:09:42 Server is on same subnet as client.
16/Mar/2001 00:10:16 ..bootstrap failure
```

Ensure that the `tftp` entry in `/etc/inetd.conf` is not commented out — the `tftp` entry should be similar to the following:

```
tftp  dgram  udp      wait    root      /usr/sbin/tftpd      tftpd  /tmp
/var/adm/ris /ris
```

If the problem persists, add a debug flag (`-d`) to the `tftpd` entry, as follows:

```
tftp  dgram  udp      wait    root      /usr/sbin/tftpd -d      tftpd  /tmp
/var/adm/ris /ris
```

Debug information is sent to `syslogd`.

11.4.3 RIS Boot Failure: Cannot Determine RIS Home Directory

The RIS boot may fail with the following output:

```
Cannot determine RIS home directory on iamms0
Unable to save existing hardware configuration. New configuration will be used.
*** Performing RIS Installation from iamms0
Loading installation process and scanning system hardware.
```

The base operating system software is not listed in the table of contents. This probably means you have not registered this client for an Operating System product on the RIS server. It could also mean that this `ris` area is corrupt.

The installation cannot proceed.

This failure can occur if the RIS server is running C2 enhanced security and the RIS user password expired. Try the following command to check this:

Boot Errors

```
# rsh -l ris atlasms0 pwd
Your password has expired.
#
```

If your RIS server will have C2 security enabled, the RIS user file must be changed to ensure that the RIS password does not expire and deny client access.

Perform the following steps on the RIS server as superuser to modify the RIS user file if you are going to use RIS with C2 security enabled:

1. Edit the file `/tcb/files/auth.db`. This requires you to use the `edauth` utility.

Use the `edauth` command line utility or the `dxaccounts` graphical user interface to modify the RIS account. See `edauth(8)` or `dxaccounts(8)` for more information.

Each field is delimited by a colon (:).

2. Set the current password field `u_pwd` to an asterisk (*).
3. Set the `u_succhg` value to any non-zero value. This value is a `time_t` type printed with `%ld`.
4. Set the `u_life` and `u_exp` fields to zero.

The following is an example of a modified `/tcb/files/auth/r/ris` user file:

```
ris:u_name=ris:u_id#11:\
    u_oldcrypt#0:\
    u_pwd=*\:\
    u_exp#0:u_life#0:\
    u_succhg#79598399:\
    u_suclog#79598399:\
    u_lock@:chkent:
```

After you make these changes, the RIS password should not expire and cause a denial of service to clients.

11.4.4 Failed to RIS Boot a Node When Attempting to Add It to the Cluster

This error occurs if `joind` is not running. This may happen if, during the boot process, `caa` failed to register the `dhcp` service. The boot output is as follows:

```
CAA daemon started
Starting CAA application registration
Cannot communicate with the CAA daemon!
Cannot communicate with the CAA daemon!
Cannot communicate with the CAA daemon!
```

If the `dhcp` service is not registered, `joind` fails to start.

Solution:

Re-register the `dhcp` service:

```
# caa_register dhcp
```



```
Restart joind:
# /sbin/init.d/dhcp start
```

11.4.5 Cannot Communicate with the CAA Daemon

See Section 11.4.4.

11.4.6 New Member Fails to RIS Boot

The RIS boot command may fail with either of the following errors:

- `status:failed RIS-Boot-fail - Could not ris boot`
- `panic (cpu 0): vfs_mountroot: cannot mount root`

If the `status:failed RIS-Boot-fail` message is displayed on your HP AlphaServer SC system, perform the checks described in Section 11.4.1 on page 11–18.

If the `panic (cpu 0): vfs_mountroot` message is displayed on your HP AlphaServer SC system, perform the following steps:

1. Ensure that the RIS server is configured as an NFS server.
2. Ensure that the `/etc/bootptab` file on the RIS server has an entry similar to the following:

```
.ris0.alpha:tc=.ris.dec:bf=/ris/r0k1:sa=www.xxx.yyy.zzz:rp="rrr:/ris/r0p1":
```

where the `sa` and `rp` fields are populated depending on the type of RIS server in use as follows:

- When a standalone management server is used as the RIS server, then the `sa` field contains the management LAN IP address (in the above example, `www.xxx.yyy.zzz`), and the `rp` field contains the hostname (in the above example, `rrr`).
 - When a cluster is used as the RIS server, then the `sa` field contains the default cluster alias IP address (in the above example, `www.xxx.yyy.zzz`), and the `rp` field contains the default cluster alias name (in the above example, `rrr`).
3. Try restarting the `nfs` daemons on the RIS server as shown in Section 11.4.7.

11.4.7 New Member Fails After RIS Boot

Adding a member to a domain may fail with the following error just after RIS boot:

```
02/Mar/2000 19:36:58 # mount 10.128.0.29:/cluster/admin /mnt
02/Mar/2000 19:37:08 Cannot talk to mount server at 10.128.0.29: RPC: Timed out
```

If this message is displayed on your HP AlphaServer SC system, perform the following steps on the first node of the domain:

1. Stop NFS as follows:
atlas0# `/sbin/init.d/nfs stop`

Boot Errors

2. Kill the portmap daemon.
3. Restart NFS as follows:
`atlas0# /sbin/init.d/nfs start`

11.4.8 RIS: Boot Error

The RIS boot command may continuously output a message similar to the following:

```
..Block FFFFFFF2 is not in any zone
```

If this message is displayed on your HP AlphaServer SC system, perform the following steps:

1. Start RIS, and check that the first RIS product is Tru64 UNIX.
2. Check that the booting node is a client of the Tru64 UNIX product.
3. Retry the RIS boot.

11.4.9 sra boot Has Timeout Error

This error occurs because the node fails to display the `login:` prompt after being booted. The boot process appears to have hung. The cause of the problem depends on the state of the node, as follows:

- The node was previously working normally. The problem may be caused by a HP AlphaServer ES40 fault that prevents successful reboot. Try the `sra` command described below to powercycle the node.
- The node is still being configured. There are many possible causes. Try the solutions described below.

Solution:

1. Manually connect to the console using the `sra console` command. Press Ctrl/C. If the boot process now continues, the boot process was probably hung while trying to NFS-mount remote file systems. Check that the network cables are attached, and that the first node of each domain is booted and is operating correctly. If you correct an error, remember to shut down and then boot the node so that file systems are correctly mounted.
2. If the node is unresponsive, shut it down, then boot it as follows, so that you can observe the message printed to the console.
 - a. To halt the node, close your console connection and use the `sra shutdown` command.
 - b. Reconnect to the console using the `sra console` command.

- c. Boot the system from the appropriate disk. Stay connected so that you can observe the complete boot sequence. You may spot errors that explain the failure to boot correctly.
3. Finally, powercycle the node using the `sra` command, as follows:
 - a. Make sure that you are not connected to the console.
 - b. On the first node of the domain, issue the following command (where `atlas` is an example system name):


```
# sra power_off -nodes atlas7
```
 - c. Wait for 30 seconds.
 - d. On the first node of the domain, issue the following command (where `atlas` is an example system name):


```
# sra power_on -nodes atlas7
```

Note:

If you manually power down a node, you must always wait 30 seconds before powering the node up again.

11.4.10 Elan Error During Node Boot

A node boot may fail with the following error:

```
panic (cpu 0): elan0: cannot determine network position - assuming disconnected
```

If this message is displayed on your HP AlphaServer SC system, ensure that the cable connecting the node to the HP AlphaServer SC Interconnect switch is plugged in, or reseal this cable.

11.4.11 Error During Node Boot

A node boot may fail with the following error:

```
01/Feb/2000 10:10:26 Waiting for cluster mount to complete
01/Feb/2000 10:10:26 panic (cpu 0): cfs_mountroot: cfs_lock_devts failed for boot partition
01/Feb/2000 10:10:27 syncing disks...
01/Feb/2000 10:10:28 trap: invalid memory write access from kernel mode
01/Feb/2000 10:10:28
01/Feb/2000 10:10:28      faulting virtual address:      0x0000000000000008
01/Feb/2000 10:10:28      pc of faulting instruction:  0xfffffc00002c80bc
01/Feb/2000 10:10:28      ra contents at time of fault: 0xfffffc000066a97c
01/Feb/2000 10:10:28      sp contents at time of fault: 0xfffffe05449e7650
01/Feb/2000 10:10:28
01/Feb/2000 10:10:28 DUMP: Will attempt to compress 257818624 bytes of dump
01/Feb/2000 10:10:28      : into 3923763184 bytes of memory.
01/Feb/2000 10:10:28 DUMP: Dump to 0x200005: ..halted CPU 1
01/Feb/2000 10:10:33 halted CPU 2
01/Feb/2000 10:10:33 halted CPU 3
```

Boot Errors

```
01/Feb/2000 10:10:33
01/Feb/2000 10:10:34 halted CPU 0
01/Feb/2000 10:10:34
```

If this message is displayed on your HP AlphaServer SC system, you may be trying to boot a node into a cluster of which it is no longer a member.

11.4.12 Node Hang During RIS Boot

A node boot may hang if it is RIS-booted and there is a speed mismatch; for example, if the Summit switch port for that node is configured for one speed, and the Ethernet card on the node is configured for a different speed. Both should be configured for 100Mbps.

11.4.13 RIS: File Open Failure for BOOTP

When bootp packets cannot solicit a response and fail with the following message:

```
P00>>>boot -p bootp eib0
(boot eib0.0.0.2005.1 -flags A)
Trying BOOTP boot.
Broadcasting BOOTP Request...
..file open failed for bootp/eib0.0.0.2005.1
```

Enter the following command:

```
P00>>>net -stop eia0
P00>>>net -start eia0
```

Then, enter the bootp command once again, as follows:

```
P00>>>boot -p bootp eia0
(boot eia0.0.0.2004.1 -flags A)
Trying BOOTP boot.
Broadcasting BOOTP Request...
..Received BOOTP Packet File Name is: /ris/r0k1
```

11.4.14 RIS: Boot Failure Access Violation

The RIS boot may fail with the following message:

```
P00>>>boot -protocol bootp eia0
(boot eia0.0.0.2004.1 -file vmunix -flags A)
Trying BOOTP boot.
Broadcasting BOOTP Request...
Received BOOTP Packet File Name is: /ris/r0k1
local inet address: 10.128.0.1
remote inet address: 10.128.101.2
TFTP Read File Name: vmunix
netmask = 255.255.0.0
Server is on same subnet as client.
..Tftp error 2: Access violation.
bootstrap failure
```

Check the nodes boot_file SRM variable, and ensure that it is blank.

11.5 Adding RIS Client With No Default Route

Adding a client to RIS results in the following warning if there is no default route:

```
usage: ifconfig interface
      [ af [ address[/bitmask] [ dest_addr ] ] [ up ] [ down ] [ netmask mask ] ]
      [ broadcast address ]
      [ alias address[/bitmask | netmask mask ] ]
      [ -alias address ]
      [ aliaslist address-list[/bitmask | netmask mask ] ]
      [ -aliaslist address-list ]
      [ abort ]
      [ delete [ address ] ]
      [ metric n ]
      [ trailers | -trailers ]
      [ promisc | -promisc ]
      [ allmulti | -allmulti ]
      [ filter | -filter ]
      [ arp | -arp ]
      [ ipmtu mtu ]
      [ speed value ]
      [ debug | -debug ]
      [ trustgrp group ]
      [ add [interface-list] ]
      [ switch ]
      [ remove ]
      [ nrtimers value-list ]
      [ autofail [value] | -autofail ]
```

This is just a warning; the client is successfully added. However, to avoid this warning, you should set a default route.

11.6 Corrupt .member.list File Causes Core Dumps

If the `sra add_member` command fails, the ASCII file `/cluster/admin/.member.list` may become corrupt. This causes certain commands — such as `clu_get_info`, `dsfmgr`, and `hwmgr` — to core dump.

The workaround is to re-create the `/cluster/admin/.member.list` file, using any text editor, as shown in the following example:

```
# This is the cluster member list file.
# This file is created and modified by the cluster administration commands.
# The format of this file is:
#   o Lines starting with a # are comments
#   o The 1st non-comment line is the number of members entry, which begins
#     with the word 'members' followed by a space and an integer whose value
#     is equal to the number of cluster members.
#   o One line for each member in the form member<id>, where:
#     <id> is a an integer in the range 1-63 inclusive.
# DO NOT edit this file!
#####
```

Adding a New Member Fails

```
members 2
member1
member2
```

11.7 Adding a New Member Fails

When adding a new member to a cluster, you may see the following error:

```
15:44:28 atlas1      Info:      Adding Member atlas1 (id:1) to cluster
15:45:03 atlas1      Info:      Creating Disk
15:45:07 atlas1      Info:      Creating AdvFS
15:45:08 atlas1      Info:      Creating member-specific files
15:45:08 atlas1      Info:      Creating new members root member-specific files
15:45:27 atlas1      Info:      Error: Tar returned an error when coping member
specific files to: /cluster/members/member2/
```

The reason for the failure is that the root file system (/) is full. Free some space in the root file system, and add the member again.

11.8 sra setup Errors

If any nodes fail the hardware probe during `sra setup` (see Section 5.2 on page 5–37 or Section 6.2 on page 6–24), a list of failed nodes is written to the `/tmp/sra_hosts.failed` file.

The hardware probe will fail in the following situations:

- The Node is Not at the SRM Prompt (see Section 11.8.1 on page 11–26)
- The Console Logger is Misconfigured (see Section 11.8.2 on page 11–27)

11.8.1 The Node is Not at the SRM Prompt

To find out whether a failed node is at the SRM prompt, log in to its console by running the following command in another window:

```
# sra -cl nodename
```

Note:

The node may have started the Tru64 UNIX Factory-Installed Software startup procedure. If so, the console prompts you to answer the following question:

Would you like to continue? (y/n):

Enter **n** to return the node to the SRM prompt; then enter Ctrl/G to exit the console.

The `sra setup` command asks if you would like to reset the failed nodes. If you have returned the node to the SRM prompt as described in the Note above, enter **No**. Otherwise, enter **Yes** to reset the failed nodes and return them to the SRM prompt.

The `sra setup` command asks if you would like to try the hardware probe again. Enter **Yes**.

The `sra setup` command asks if you would like to initialize the hardware during the probe. If the node had previously started the Tru64 UNIX Factory-Installed Software startup procedure as described in the Note above, enter **No**. Otherwise, enter **Yes**.

11.8.2 The Console Logger is Misconfigured

The console manager is automatically configured at installation time, and the information is stored in the `sc_cmf` table in the SC Database. This error indicates that the table contains incorrect information. To fix this problem, perform the following steps:

1. Exit the `sra setup` command.
2. Stop the console logger, as follows:
`# /sbin/init.d/cmf stop`
3. Rerun the `sra setup` command.

Note:

Take care when entering information at the `sra setup` prompts — if you enter incorrect values for certain information (for example, terminal server IP address), the `cmfd` daemon will not start.

The `sra setup` command will recreate the `sc_cmf` table with correct information from the SC database, and will restart the `cmfd` daemon.

11.9 Terminal Server Errors

This section describes the following terminal server errors:

- Terminal Server Configuration Has Changed (see Section 11.9.1 on page 11–27)
- Terminal Server Refuses a Connection (see Section 11.9.2 on page 11–27)

11.9.1 Terminal Server Configuration Has Changed

If you press the Reset button on the terminal server, it loses its configuration information and is reset to the factory configuration. Reconfigure the terminal server as described in Chapter 14 of the *HP AlphaServer SC System Administration Guide*.

11.9.2 Terminal Server Refuses a Connection

The terminal server allows only one connection to a node's console port to exist at any point in time. If you attempt to connect to a console port that is already in use, the following messages are reported in the Xterm window:

rinfo Command Displays UID or Wrong User Name

Console-Busy port already in use
Press return to continue

Press Return to close the connection.

If you need to connect to the console port, follow these steps:

1. Check to see if anyone is connected to the console port on that node. To find the user process that is connected to the console, run the following command:

```
# ps -a | grep sra | grep nodename
```

Search for output similar to one of the following:

```
18288 pts/2 I +      0:00.09 /usr/bin/sra console -cl <nodename>  
18288 pts/2 I +      0:00.09 /usr/bin/sra -cl <nodename>
```

where node is the name of the node to which you wish to connect.

If you can identify the person who is connected to the node's console, ask them to close their connection.

If you cannot identify the person who is connected to the node's console, or if you fail to find an sra console process, go to step 2.

2. The sra ds_logout has three possible actions which you can use to recover console access to a node.

- a. Enter the following command as the root user:

```
# sra ds_logout -nodes nodename
```

This will end a user's connection, or a ghost connection caused by the terminal server failing to close an sra console session.

- b. Enter the following command as the root user:

```
# sra ds_logout -ts yes -nodes nodename
```

This will cause CMF to drop its connection to the terminal server. It will reconnect after a short delay. This option is useful if the problem had been caused by, for example, the physical cable being disconnected and reconnected, causing the CMF connection to "hang".

- c. Enter the following command as the root user:

```
# sra ds_logout -force yes -nodes nodename
```

This will issue a logout command for the node, directly to the terminal server.

11.10 rinfo Command Displays UID or Wrong User Name

The rinfo command may display the UID instead of the user name, or may display a different username than the name of the user who allocated the resource. This indicates that the node on which the allocate was performed has a different set of users than the set of

How to Drop and Rebuild the RMS Database

users on the `rmshost` system (the management server). Typically, this happens if the `/etc/passwd` files differ, or if the same NIS server is not being used by all nodes (including the management server).

To correct this problem, ensure that the UID and username — of all RMS users — are the same on all nodes in the system.

11.11 How to Drop and Rebuild the RMS Database

To drop and recreate the RMS database, perform the following steps on the RMS master node (the management server, if used, or Node 0):

1. Halt all other nodes.
2. Stop the RMS daemons by running the `rmctl stop` command.
3. Drop the old database using the `msqladmin drop` command, as follows (where `atlas` is an example system name):

```
# msqladmin drop rms_atlas
```

Caution:

If the RMS database is dropped by using the command: `# msqladmin drop rms_atlas`, the *entire* SC database is dropped. The SC database needs to be recreated by using the `sra setup` command, and not the `rmsbuild` command.

4. Create the new database using the `sra setup` command (if the management server is used, see Section 5.2 on page 5–37, or if Node 0 is used, see Section 6.2 on page 6–24).
5. Start RMS on the remaining nodes by running the `rmctl start` command.

11.12 rcontrol Reports Errors During Node Boot

The `rcontrol` command will report a failure during node boot if the parallel cable is not connected to the switch. The failure message can take several forms — two of these are shown below:

```
timestamp { 03/12/01 18:06:47 } atlas1 Info:    rcontrol: Error: can't find
machine containing node atlas0
18:40:22 atlas3      Info:    rcontrol: Error: swmgr failed to connect atlas3:
check control cable is plugged
18:40:22 atlas3      Info:                                in
```

You can ignore these messages.

rcontrol Reports Error When Starting Partitions

11.13 rcontrol Reports Error When Starting Partitions

The following message may be reported when you attempt to start an RMS partition for the very first time:

```
# rcontrol start partition parallel
rcontrol: Error: can't start partition parallel: bad Elan Id
ordering
```

This message implies that the nodes are not connected in sequence to the switch ports (for example, two cables may have been swapped in error).

To solve this problem, halt all suspected nodes, and run the command:

```
atlasms# sra elancheck -nodes "atlas[0-1023]"
```

This will report whether all nodes are connected to the correct elan ports in ascending order, and inform you which nodes are offending if appropriate.

11.14 clu_get_info Prints CONFIGURATION_ERROR

If `/etc/rc.config` has been corrupted, `clu_get_info` displays a warning similar to the following:

```
Cluster memberid = 2
Hostname = CONFIGURATION_ERROR
Cluster interconnect IP name = CONFIGURATION_ERROR
Member state = UP
```

A possible cause for this is that the root (`/`) file system is full and someone has attempted to edit the `/etc/rc.config` file.

The general solution is as follows:

1. Restore the `/etc/rc.config` file (copy from backup or from another location).
2. Shut down the problem node.
3. When the shutdown has finished, boot the node.

11.15 How to Powercycle an HP AlphaServer SC System

This section describes the order in which you should perform the following activities:

- Shutting Down (see Section 11.15.1 on page 11–31)
- Powering Off (see Section 11.15.2 on page 11–31)
- Powering On (see Section 11.15.3 on page 11–31)
- Starting Up (see Section 11.15.4 on page 11–32)

How to Powercycle an HP AlphaServer SC System

Chapter 2 of the *HP AlphaServer SC System Administration Guide* provides more details about how to shut down and start up an HP AlphaServer SC system.

11.15.1 Shutting Down

When shutting down an HP AlphaServer SC system, do so in the following order:

1. Shut down all nodes.
2. If using a management server, shut it down.
3. Shut down the RAID controllers.

11.15.2 Powering Off

When powering off an HP AlphaServer SC system, do so in the following order:

1. Shut the system down, as described in Section 11.15.1.
2. If using a management server, power off all nodes. If not using a management server, power off all nodes except the first node in the first domain.
3. If using a management server, power it off. If not using a management server, power off the first node in the first domain.
4. For non-federated installations, power off the HP AlphaServer SC Interconnect switch(es). For federated installations, power off the HP AlphaServer SC Interconnect in the following order:
 - a. node-level switches
 - b. top-level switches
 - c. clock-distribution boxes.
5. Power off the terminal servers and the ethernet switches.
6. Power off the Fibre Channel switches.
7. Power off the RAID controllers.

11.15.3 Powering On

When powering on an HP AlphaServer SC system, do so in the following order:

1. Power on the RAID controllers.
2. Power on the Fibre Channel switches.
3. Power on the ethernet switches and the terminal servers.
4. For non-federated installations, power on the HP AlphaServer SC Interconnect switch(es). For federated installations, power on the HP AlphaServer SC Interconnect in the following order:
 - a. clock-distribution boxes

Error When Installing the Elan Subset on a Management Server

- b. top-level switches
- c. node-level switches
5. If using a management server, power it on. If not using a management server, power on the first node in the first domain.
6. Power on all remaining nodes.
7. Start the system up, as described in Section 11.15.4.

11.15.4 Starting Up

When starting up an HP AlphaServer SC system, do so in the following order:

1. If using a management server, boot it up.
2. Boot the minimum number of nodes to allow each domain to attain quorum.
3. Boot all remaining nodes.

All of the other HP AlphaServer SC components start automatically when powered up.

11.16 Error When Installing the Elan Subset on a Management Server

The following error may be reported when installing the Elan subset on a management server:

```
Configuring "Elan Device Driver TS2.5-BL5-0021_EFT4" (ELNMOD320)
eip0: cannot create kernel comms (rail 0) rcvr service for 65536 byte svc(2)
eip: failure to attach to kernel comms - check if an elan is present
The dynamically loading Elan Device Driver has been installed on your system.
Installation of the Elan Device Driver (ELNMOD320) subset is complete.
```

You can ignore this message.

11.17 Database Access Denied Errors

To make modifications to the database you must be the root user and a member of the RMS group. So it is normal for non-root users to get "access denied" errors if they attempt to modify the database. However, if the root user gets access denied errors you should take the following actions to identify the problem:

- Check that the root user is a member of the RMS group. On an HP AlphaServer SC, the RMS group should be present in `/etc/group` (and not in NIS). The entry for the RMS group should be present and identical in `/etc/group` on all domains and management servers.
- Check that there is an RMS user in `/etc/passwd`. The entry for the RMS user should be present and identical in `/etc/passwd` on all domains and management servers.

Database Access Denied Error on Some Domains

- Check that the `mSQL` daemon was started by the root user.
- If the access denied error occurs when connecting to the database, i.e., before you attempt to perform a database modification, this might indicate a problem with the cookie security mechanism. You can determine whether the access denied error occurs on connect or on a modify by simply running `rmsquery` as root. If the error occurs before `rmsquery` returns a prompt, then the error has occurred on connect. However, if you can perform a query operation, but not an update operation, then the connect has succeeded.

You can disable the cookie mechanism by running the following command on `rmshost` as follows:

```
# sra cookie -enable no
```

Once you do this, you are operating in a less secure mode. Once you have determined and corrected the source of the problem (see Section 11.18), you can re-enable the cookie mechanism as follows:

```
# sra cookie -enable yes
```

11.18 Database Access Denied Error on Some Domains

If the root user or an SC daemon is getting access denied errors on some domains, but not on others or is getting errors on all domains but not on the management server, you should review the information in Section 11.17 about the RMS user and group.

However, if the RMS user and group information appear correct on every domain, you should disable cookies by running the following command on `rmshost`:

```
# sra cookie -enable no
```

If you no longer get access denied errors, this indicates a failure in the cookie security mechanism. You can diagnose potential causes of this failure as follows:

- The cookie mechanism is managed by the daemons associated with the `scrun` command. These are `gxmgmtd`, `gxclusterd` and `gxnoded`. You should check that:
 - `gxmgmtd` is running on the management server (or member 1 of the first domain if there is no management server)
 - `gxclusterd` is running on all nodes in the system

If not, you should start the daemons as follows:

```
# /sbin/init.d/gxdaemons start
```

You must run this command on each node or management server where the daemons are not running.

Diagnosing Federated Network Routing Problems

- Check that the file `/var/cookies/root` is identical on all management servers on each domain. The file must be owned by `root` and group `rms` and have permissions `owner:rw, group:r`. If one or more domains have a `/var/cookies/root` file that is different than the host where the database is located (`rmshost`), you can copy the file from `rmshost` to the domains.

Once, the `gx` daemons are working normally or you have manually copied the `/var/cookies/root` file, you can re-enable the cookie security mechanism by running the following command on `rmshost`:

```
# sra cookie -enable yes
```

You can check if cookies are enabled or not by running the following command on `rmshost`:

```
# sra cookie
```

11.19 Diagnosing Federated Network Routing Problems

When a federated network is wired correctly, a boot sequence for any node should display the following message:

```
28/Feb/2002 10:58:07 elan0: nodeid=127 level=5 numnodes=512
(adaptive routing ok)
28/Feb/2002 10:58:07 elan0: New Nodeset [127]
```

When a federated network is not wired correctly, a boot sequence for any node may display the following message:

```
27/Feb/2002 13:13:01 elan0: nodeid=127 level=5 numnodes=512
(adaptive routing disabled 0x14)
27/Feb/2002 13:13:01 elan0: New Nodeset [127]
```

To identify the cause of the error, run the SC Viewer application to check for the following potential problems:

- QM410 controller card not working
- QM410 controller card with incorrect firmware revision
- Badly populated combination of high level cards in node-level chassis
- Incorrect cabling of uplinks from node-level cards to top level chassis

11.20 Wakeup on LAN (WOL) Problem

The WOL feature may report some problems if the DLI kernel subsystem is unconfigured on the management server (or running node). This problem may occur even when the node is at the SRM prompt.

To resolve this problem, configure the DLI kernel subsystem on the management server by running the following command:

```
atlasms# sysconfig -c dli
```

11.21 scfsmgr/pfsmgr report Could Not Open Socket

If `scfsmgr` or `pfsmgr` cannot connect to the `scmountd`, they report an error such as:
Could not open socket to the daemon: couldn't open socket: connection refused

To correct this, use one of the following procedures:

- If you have a stand alone management server:
 - a. Log onto the management server
 - b. Stop `scmountd` as follows:

```
# /sbin/init.d/scmountd stop
```
 - c. Wait two minutes and restart it as follows:

```
# /sbin/init.d/scmountd start
```
- If you have a dual-redundant management server or have no management server:
 - a. Log into the management server or domain 0
 - b. Stop `scmountd` as follows:

```
# caa_stop SC30scmountd
```
 - c. Start `scmountd` as follows:

```
# caa_start SC30scmountd
```

11.22 Problem with HSG Devices in Installation Process

The install process gets as far as the `clucreate` stage, which requires access to the HSG storages area, before failing with the following message:

```
20:41:56      Node                      eve12      -- <Failed:cluCreate> Error:
createclusterlabel (primary) failed - /usr/sra/bin/createclusterlabel[311]:
rootDiskSize*rootPartitionSize/100: bad number: <ABORT code completed>
```

Command has finished:

```
Command 165 (install) eveC : eve[12-15] -- <Error> Command Failed on:
eve12
*** Node States *** Errored: eve12
```

The explanation is that the HSG controllers create "reserves" on the disks belonging to a cluster when the cluster is created, that is, a locking mechanism that allows access to the HSG disks from a particular cluster only.

When the cluster is down "persistent reserves" appear beside each disk on the HSG controller. The problem may appear as a hardware issue but is caused by this access problem.

Problem with HSG Devices in Installation Process

The disks can be seen but not accessed, and display an I/O error or a more detailed hardware error. This can be seen in booted nodes, or RIS-booted nodes as follows:

```
# hwmngr -v d
HWID: Device Name           Mfg      Model      Location
-----
 6: /dev/dmapi/dmapi
 7: /dev/scp_scsi
 8: /dev/kevm
37: /dev/disk/floppy0c      3.5in floppy  fdi0-unit-0
73: /dev/disk/dsk0c        COMPAQ      BD018635C4   bus-0-targ-0-lun-0
74: /dev/disk/dsk1c        COMPAQ      BD018635C4   bus-0-targ-1-lun-0
75: /dev/disk/dsk2c        DEC         HSG80        IDENTIFIER=2
76: /dev/disk/dsk3c        DEC         HSG80        IDENTIFIER=3
77: /dev/disk/dsk4c        DEC         HSG80        IDENTIFIER=1
78: /dev/disk/dsk5c        DEC         HSG80        IDENTIFIER=4
79: /dev/disk/cdrom0c      COMPAQ      CRD-8402B    bus-3-targ-0-lun-0
80: /dev/cport/scp0        HSG80CCL    bus-2-targ-0-lun-0

# disklabel -p dsk4
disklabel: dsk4: I/O error

# disklabel -p dsk3
cam_logger: SCSI event packet
cam_logger: hardware_id=-86 bus 2 target 0 lun 12
cdisk_op_spin
Device Not Ready
Hard Error Detected
Hardware ID = -86
DEC      HSG80      V86F
Active CCB at time of error
CCB request completed with an error
disklabel: dsk3: I/O error
#
```

However, on the HSG controller, the "persistent reserved" can be seen as follows:

```
D10                                DISK60000    (partition)
LUN ID:      6000-1FE1-0009-5160-0009-0320-0632-01F3
IDENTIFIER = 1
Switches:
RUN              NOWRITE_PROTECT      READ_CACHE
READAHEAD_CACHE  WRITEBACK_CACHE
MAX_READ_CACHED_TRANSFER_SIZE = 32
MAX_WRITE_CACHED_TRANSFER_SIZE = 32
Access:
EVE12,    EVE12B,    EVE12C,    EVE13,    EVE13B,    EVE13C
State:
ONLINE to this controller
Persistent reserved
NOPREFERRED_PATH
Size:      35557306 blocks
Geometry (C/H/S): ( 7000 / 20 / 254 )
```


To clear these reserves so that the disks are accessible again, run the following command from a node in the relevant cluster, from Single User Mode:

```
/usr/sbin/cleanPR clean
```

To show any existing reserves, before cleaning them, and to check that the script has run properly, run the command:

```
/usr/sbin/cleanPR show
```

11.23 Upgrade Errors

This section describes the following upgrade errors:

- Upgrade Setup: RIS Host Appears to Be Invalid (see Section 11.23.1 on page 11–37)
- Problem with dupatch Failure when Upgrading an HP AlphaServer SC System (see Section 11.23.2 on page 11–37)
- Restarting a Failed Upgrade (see Section 11.23.3 on page 11–38)
- Problem with Dependency on First Cluster (see Section 11.23.4 on page 11–38)
- Problems when clu_quorum does not Complete (see Section 11.23.5 on page 11–38)
- Upgrade Backup: node failed with cfs_find_drv_handle panic (see Section 11.23.6 on page 11–39)
- Restoring from Backup Can Sometimes Fail (see Section 11.23.7 on page 11–39)

11.23.1 Upgrade Setup: RIS Host Appears to Be Invalid

An upgrade setup may fail with the following error:

```
Permission denied.
```

```
SC_UPGRADE: <name-of-ris-server> appears to be invalid!
```

This error may appear if the RIS user accounts `.rhosts` file (`/var/adm/ris/.rhosts`) is incomplete. Make sure that the default cluster alias of the cluster being upgraded is a RIS client.

11.23.2 Problem with dupatch Failure when Upgrading an HP AlphaServer SC System

The `dupatch` script may fail while upgrading an HP AlphaServer SC system and return the following error message:

```
*** You have selected 1 patches ***
/native_threads//usr/opt/compaq/svctools/common/jre/bin/./bin/alpha/
native_threads/java is
/usr/opt/compaq/svctools/common/jre/bin/./bin/alpha/native_threads/
/usr/opt/compaq/svctools/common/jre/bin/./bin/alpha/native_threads/java
is /usr
```

Upgrade Errors

```
/opt/compaq/svc: no space
```

To solve this problem, kill the java process and re-run the `sra upgrade` command.

11.23.3 Restarting a Failed Upgrade

If the upgrade fails and for some reason will not restart the `sra upgrade`, you may need to check that the `clusters/nodes` in question are in the correct state to restart the upgrade. In order to successfully restart the upgrade, here are a few simple tips:

1. Check the current state of the upgrade by querying the `sc_domain` table.
2. Ensure that all nodes are in the `SYSTEMUP` state.

If the current upgrade state of your failed cluster is `Setup`, you should ensure that all members of that cluster are up and running before trying to re-issue the `sra upgrade` command.

11.23.4 Problem with Dependency on First Cluster

In a system with multiple clusters, there is a dependency during upgrade that the first cluster is available. Due to the current design of the `/etc/bootptab` file, there is a dependency on the first cluster to act as a gateway.

This is particularly evident if using a clustered management server.

If you are having difficulty upgrading the first cluster, or making it available during upgrade, you can edit the `/etc/bootptab` file on the first cluster and change the gateway settings to point to the lead nodes of that cluster.

For example, currently the `atlas32` entry in `/etc/bootptab` is as follows:

```
atlas32:tc=.ris0.alpha:ht=ethernet:gw=10.128.0.1:ha=00508BDFAAAE:ip=10.128.0.33:
```

You need to change the gateway setting to the IP of the cluster `atlasD1` to look like the following:

```
atlas32:tc=.ris0.alpha:ht=ethernet:gw=10.128.0.33:ha=00508BDFAAAE:ip=10.128.0.33:
```

11.23.5 Problems when `clu_quorum` does not Complete

In two places during an upgrade, the `clu_quorum` command is issued on each cluster to change the number of votes a cluster member has. In some rare cases, this command may not complete for some reason and the upgrade will hang waiting for it to finish. If this happens, you will see the `clu_quorum` process running on the cluster in question.

If you can verify that the `clu_quorum` command has changed the vote successfully, then it is safe to kill the `clu_quorum` process (and associated logging processes). The upgrade should continue from that point.

11.23.6 Upgrade Backup: node failed with `cfs_find_drv_handle` panic

When backing up a system prior to upgrade, you might see a node panic with the following message:

```
panic (cpu 1): cfs_find_drv_handle: no dev on list
```

This is typically caused by the member's boot file domain being incorrectly set in the `/etc/fdmns` directory.

For information on how to correct this problem, please refer to Pre-Upgrade Audit (see Section 4.2 on page 4–8).

11.23.7 Restoring from Backup Can Sometimes Fail

When restoring a domain using the `sra_cluster_backup` utility as described in Section 4.8.3, you may see the following error messages:

```
# /usr/opt/sra/bin/sra_cluster_backup -disk dsk8 restore
sc_cluster_backup#cluster_root DOES NOT EXIST!
sc_cluster_backup#cluster_usr DOES NOT EXIST!
sc_cluster_backup#cluster_var DOES NOT EXIST!
sc_cluster_backup#boot_partitions DOES NOT EXIST!
```

The possible causes are as follows:

- Incorrect special device number for the backup disk
- Persistent reservations in the HSG for the backup disk

The solution to each possible cause is described below.

11.23.7.1 Incorrect Special Device Number for the Backup Disk

In order to restore a domain from backup when an upgrade fails with a severe error, the lead member will be booted from the UNIX disk.

When the domain was being backed up, the backup process itself will have created a directory `/etc/fdmns/sc_cluster_backup` on the root partition of the UNIX disk. Within this directory, a symbolic link will have been created to the special device file of the backup disk, for example, `/dev/disk/dsk5c`.

However, the special device name assigned within the domain might be different to the special device name assigned to the same disk when the lead member is booted from the UNIX disk.

To solve this issue, first confirm the actual special device name for the backup disk using the following command when booted from the UNIX disk:

```
# hwmgr -v d
```

Interpreting Problems During Software Upgrade

Once the actual special device is known, then you can correct the special device for the `sc_cluster_backup` file domain using the following commands:

```
# cd /etc/fdms/sc_cluster_backup
# unlink dsk8c
# ln -s /dev/disk/dsk5c
```

Where:

- the partition used for the backups is always the C: partition
- `dsk8c` is the special device name assigned within the domain
- `dsk5c` is the special device name assigned to the same disk when booted from the UNIX disk

Once the `sc_cluster_backup` file domain is corrected, you can proceed with the restore operation as described in Section 4.8.3.

11.23.7.2 Persistent Reservations in the HSG for the backup disk

The explanation is that the HSG controllers create "reserves" on the disks belonging to a domain when the domain is created, that is, a locking mechanism that allows access to the HSG disks from a particular domain only.

This issue is described in more detail in Problem with HSG Devices in Installation Process (see Section 11.22 on page 11–35).

To clear these reserves so that the disks are accessible again, run the following command while the lead member of the domain is booted from the UNIX disk.

```
# /usr/sbin/cleanPR clean
```

To show any existing reserves, before cleaning them, and to check that the script has run properly, run the command:

```
# /usr/sbin/cleanPR show
```

Once the reservations are cleaned, you can proceed with the restore operation as described in Section 4.8.3.

11.24 Interpreting Problems During Software Upgrade

The upgrade of a management server is effectively a manual process and the steps required are fully described in Chapter 4. The steps are clearly explained and if problems occur, the diagnosis should be straightforward.

Once the management server is upgraded, then the subsequent upgrade of the domains is an automated process that is controlled and monitored by various `sra` scripts. There is a state machine that controls this automated process. A basic description of the upgrade state mechanism is provided in Upgrade States (see Section 4.1.4 on page 4–4).

Interpreting Problems During Software Upgrade

In the following subsections, further information is provided that may assist you in resolving problems that may arise during each state of the upgrade process.

The upgrade status of each domain can be determined using the command:

```
sra upgrade_info -domain atlasD[0-3]
```

This command will report the domain's status as being *one* of the following:

- Pre_Upgrade
- Upg_Installed
- Checked
- Setup
- Installed
- Upgraded

Most work performed by the upgrade automation is controlled by a script spawned in the controlling shell where the `sra upgrade` command is entered by the user. As such, it is important to keep open the controlling shell/window where the `sra upgrade` command was entered. If this shell/window is closed, then the upgrade process will terminate when the current action completes, and the process will not move beyond the current state.

The `sra upgrade` script will attach to the consoles of the lead-members and perform various tasks on the lead-members and subsequently on the non-lead members. As the process continues, the `sra upgrade` script will summarize the progress and report the events to the command line. This is necessary in order to allow the upgrade process to scale to multiple domains at once. However, if you require more information on what actions are happening in the background on each domain, then you may wish to monitor the console output of the lead members in another window.

A node will remain in the current state until all actions required to progress to the next state are successfully completed. In the case of some states, these actions can be repeated after errors until such time that the automation process succeeds, and it can move to the next state. However, depending on the error, and depending on the current state, it might not be possible to repeat the actions. Examples of such scenarios are outlined in the following subsections.

11.24.1 Pre_Upgrade State

This is the initial state of all domains. In this state, the first action is to install the upgrade software subset (SRAUPG) on the target domain(s). Possible problems with this operation may include the following:

1. Password problems for the serial console of lead member of the domain.

Interpreting Problems During Software Upgrade

2. Inability of `setld` to delete the existing upgrade software subset and install the new subset.

In order to diagnose problems in this state, typically, you can look in the console log for the lead member in `/var/sra/logs` on the management server. You should ensure that all nodes of the target domain are booted to multi-user mode and that the `rsh` service is enabled (if this service was previously disabled for security concerns).

Once the cause of the problem is understood and resolved, then the target domain can be booted to multi-user mode, and the `sra upgrade` command can simply be run again on the management server. The automation process will continue from the current state.

11.24.2 Upg_Installed State

The current state is that the upgrade subset has been installed on the domain. The next action in this state is for the automation to perform various pre-upgrade checks on the target domain. These checks include the following (however, these checks are subject to change in the future):

1. Check the operating system version of the system, in order to decide between shortcut upgrade and operating system upgrade.
2. Check that the cluster `evm` event manager is functional, in order to ensure that patch installation is successful later.
3. Check that the SCFS and PFS file systems are not mounted.
4. Stop the `scmon` daemon on all nodes, and rename the `scmon` startup script.
5. Stop the `rms` daemon on all nodes, and rename the `rms` startup script.
6. Run the `/usr/sbin/clu_upgrade` script from the operating system installation update suite.
7. Check that there is adequate disk space in `/usr` and `/var`; the upgrade process recommends 50% disk space free in these file systems.
8. Check for certain subsets that should not be installed as described in the HP AlphaServer SC documentation set. For example, see step 2 in Section 4.6.2.1 on page 4–33.
9. Check that the generic boot disk is available, which is necessary for adding members again during full upgrades.
10. Check that the UNIX boot disk is free, which is necessary for backup, restore, and full upgrade procedures.
11. Check that the `/etc/inetd.conf` file has not changed in a manner that will affect the upgrade.
12. Check that the cluster aliases are correct.

13. Check that kernel lockmode is less than the value 4.
14. Check that the SC database is set up and accessible.
15. Check that the domain is a registered RIS client, which is necessary for the operating system upgrade.
16. Check that `clu_quorum` is functional, which implicitly checks the integrity of intra-cluster communication.

In order to diagnose problems in this state, you should look in the following locations:

1. The console log for the lead member in `/var/sra/logs` on the management server.
2. The upgrade log file in `/var/adm/smlogs/sc_upgrade.log` on the target domain.

Once the cause of the problem is understood and resolved, then the target domain can be booted to multi-user mode, and the `sra upgrade` command can simply be run again on the management server. The automation process will continue from the current state.

11.24.3 Checked State

The current state is that the check phase has been completed on the domain. In this state, the following actions are performed:

1. Copy the new HP AlphaServer SC kits from the management server to `/var/adm/update/Sierra` on the target domain.
2. Populate `/var/adm/update/TruCluster` on the target domain with a symbolic link to the HP AlphaServer SC TruCluster subset.
3. Copy the new Tru64 UNIX patch kit from the management server to `/var/adm/update/Patch` on the target domain.
4. Remove the quorum votes of the non-lead members.

In order to diagnose problems in this state, you should look in the following locations:

1. The console log for the lead member in `/var/sra/logs` on the management server.
2. The contents of `/var/adm/update` on the target domain.

The typical causes of problems in this state are connectivity problems on the management network, problems with NFS mounts to the management server, and incorrect kit path names specified on the `sra upgrade` command line.

In the past, one problem with this state has related to lingering content in `/var/adm/update` for previously successful and failed upgrades. However, these issues have been addressed, and are unlikely to cause problems again.

Interpreting Problems During Software Upgrade

Once the cause of the problem is understood and resolved, then the target domain can be booted to multi-user mode, and the `sra upgrade` command can simply be run again on the management server. The automation process will continue from the current state.

11.24.4 Setup State

The current state is that the setup phase to prepare the domain for upgrade has been completed. In this state, the following actions are performed:

1. Use `rcmgr` to set the correct values for `SC_MS`, `SC_CLUSTER`, `SC_MOUNT_OPTIONS`, and `ALIASD_NONIFF` in the `/etc/rc.config.common` file
2. Ensure that `/var/sra` is not mounted from the management server.
3. For multi-rail ES40, configure the `elan0` and `elan1` rails appropriately, in order that the rail on `pci0` is used for kernel communications.
4. Enable any daemons disabled in previous kits, for example, `smsd`, `advfsd`, and `clu_mibs`.
5. Delete the current HP AlphaServer SC kits.
6. Patch the domain to the correct patch version using the `scdupatch` utility.
7. Install the latest HP AlphaServer SC kits.
8. Copy the generic kernel from the HP AlphaServer SC subset to the primary boot disk.
9. Reboot the domain members, so that the new software subsets are configured and a new kernel is built.

In order to diagnose problems in this state, you should look in the following locations:

- The console log for the lead member in `/var/sra/logs` on the management server.
- The contents of `/var/adm/update` on the target domain.

For problems experienced while in this upgrade state, it may be safest to restore the target domain from backups as described in *Serious Failures: Recover from Backup* (see Section 4.8.3 on page 4–46). In particular, if the upgrade process causes errors at any point between the time that the HP AlphaServer SC kits are deleted and the time that the new kernel is built, then the target domain must be restored from the backup images created earlier, and the upgrade will need to be started from the `pre_upgrade` state as described in *Serious Failures: Recover from Backup* (see Section 4.8.3 on page 4–46).

11.24.5 Installed State

The current state is that the Tru64 UNIX V5.1B software and patch kit and the HP AlphaServer SC software have been installed on the domain.

The actions in this state are associated with cleaning up after the operating system and the SC kit upgrade actions in the earlier step. The cluster version switch is completed, daemons that are to be disabled in the new HP AlphaServer SC release have their startup scripts renamed, and the lead member is rebooted from the primary boot disk.

Typically, there are no problems seen with actions in this state, and the automation process will perform the actions and continue to the next state.

11.24.6 Upgraded State

The cleanup has been performed and the target domain upgrade has been completed. However, the upgrade process as a whole is not fully complete, in that, there are various post-upgrade steps that must be completed as described in Post-Upgrade Tasks (see Section 4.9 on page 4–49).

11.25 Increasing the Number of ptys

When a management server is being built, then the `sra setup` phase will automatically increase the default number of ptys to a level suited to the number of nodes in the system.

For the scenario of a clustered management server, or for the scenario where further standalone management servers are being added to a system, then `sra setup` might not be rerun on each individual server. In such circumstances, then the default number of ptys will be 255 and during normal maintenance on systems with large node counts, system administrators may find that certain `sra` commands issued from these servers will fail with the message:

```
<Failed:boot> Error: The system has no more ptys. Ask your system
administrator to create more.
```

In these cases, the solution is to increase the number of ptys manually in the `pts` stanza of the `sysconfigtab` file using the `sysconfigdb` command.

```
pts:
    nptys=1024
```

Table 11–1 should be used as a guide when manually adjusting the number of ptys:

Table 11–1 Manually Adjusting ptys

Node Range	nptys
0-127	256
128-255	512
256-383	768

Increasing Socket Listen Queue Limits

Table 11–1 Manually Adjusting ptys

Node Range	nptys
384-511	1024
512-768	1280

For an example of how to change `sysconfigtab` values using the `sysconfigdb` command, please refer to Add `sysconfigtab` Parameters (see Section 6.1.11 on page 6–22).

11.26 Increasing Socket Listen Queue Limits

For systems with very large node counts (for example, 1024 nodes), then the default Socket Listen Queue Limits will need to be increased on the management server.

Failing to increase the number of sockets may cause delays when configuring in nodes and starting the partitions.

The limits can be reconfigured dynamically using the command below, once satisfied with the new values, and then the `socket` stanza in the `/etc/sysconfigtab` file should be updated with the new values using the `sysconfigdb` command.

```
atlasms# /sbin/sysconfig -r socket somaxconn=65535
atlasms# /sbin/sysconfig -r socket sominconn=65535
```

For more information on tuning please refer to the Tru64 System Tuning Guide.

Installation Overview and Checklist: When the System Has a Management Server

This appendix provides the following information to help you to install a system that has a management server:

- Installation Overview (page A-2)

Installation Overview

A.1 Installation Overview

Table A–1 provides an overview of the installation process for a system that has a management server.

Table A–1 Installation Process: When the System Has a Management Server

Perform This Installation Task...	On...	Scope of Task
Set Up the Management Server:		
1. Set the Console Variables	Mgt server	Mgt server
2. Check the System Firmware	Mgt server	Mgt server
3. Install the Tru64 UNIX Operating System	Mgt server	Mgt server
4. Customize the System Configuration	Mgt server	Mgt server
5. Install the Operating System Patch Software	Mgt server	Mgt server
6. Install and Configure the Clustered Management Server	Mgt server	Mgt server
7. Configure the RIS Server	Mgt server	Mgt server
8. Install the HP AlphaServer SC System Software	Mgt server	Mgt server
9. Install the HP Fortran Run-Time Libraries	Mgt server	Mgt server
10. Install Layered Products (Optional)	Mgt server	Mgt server
11. Install the SANworks Storage System Scripting Utility	Mgt server	Mgt server
12. Define the RMS Master Node (rmshost)	Mgt server	Mgt server
13. Build the New Kernel and Reboot	Mgt server	Mgt server
Set Up the SC Database		
14. Set Up the SC Database	Mgt server	Mgt server
Check All Nodes in the HP AlphaServer SC System:		
15. Check the State of the Nodes	Mgt server	All Nodes
16. Check the System Firmware	Mgt server	All Nodes
Configure and Diagnose the HP AlphaServer SC Interconnect		
17. Upgrading the HP AlphaServer SC Interconnect Control Processor Software	Mgt server	Mgt server
18. Creating an Interconnect Configuration Using SC Viewer	Mgt server	Mgt server
19. Confirming the Operation of the HP AlphaServer SC Interconnect	Mgt server	Mgt server
Set Up the SC Monitor System		
20. Add a SAN Appliance as Monitored Devices	Mgt server	Mgt server

Table A–1 Installation Process: When the System Has a Management Server

Perform This Installation Task...	On...	Scope of Task
21. Change Terminal Servers Monitoring Distribution	Mgt server	Mgt server
22. Change Extreme Switches Monitoring Distribution	Mgt server	Mgt server
23. Activate the Changes	Mgt server	Mgt server
Assign Cabinets in the SC Database		
24. Load the Physical Relationships	Mgt server	Mgt server
25. Populate Cabinets with Nodes	Mgt server	Mgt server
26. Populate Cabinets with Other Hardware	Mgt server	Mgt server
Building the Domains:		
27. Review the SC Database System Settings	Mgt server First node of each domain	Systemwide
28. Add sysconfigtab Parameters	Mgt server First node of each domain	Systemwide
29. Create the Domains	Mgt server First node of each domain	Systemwide
30. Boot the System	Mgt server First node of each domain	Systemwide
31. Complete the Setup of the Domains	Mgt server	Systemwide
Completing the Installation:		
32. Configure the External Network Interfaces	<i>Specified nodes</i>	<i>Specified nodes</i>
33. Improve Cluster Availability	<i>Specified nodes</i>	<i>Specified nodes</i>
34. Load File System Configuration Data in the SC Database	Mgt server	Systemwide
35. Initial Setup of Monitoring for HSG80 RAID Systems	Mgt server	Systemwide
36. Configure the RMS Database	Mgt server	Systemwide
37. Provide RMS with CAA Failover Capability	Mgt server and Node 0	Systemwide
38. Enable CMF as a CAA Application	Mgt server and Node 0	Systemwide
39. Run the Example MPI Program	<i>Any node</i>	<i>Specified nodes</i>

Installation Overview

Table A–1 Installation Process: When the System Has a Management Server

Perform This Installation Task...	On...	Scope of Task
40. Verify the HP AlphaServer SC Interconnect	Mgt server	Systemwide
41. Configure LSM	First three nodes of each domain	First three nodes of each domain
42. Verify Swap Mode	N/A	N/A
43. Add a Second Rail to an HP AlphaServer SC System after Domain Creation	Mgt server	Systemwide

Installation Overview and Checklist: When the System Does Not Have a Management Server

This appendix provides the following information to help you to install a system that does not have a management server:

- Installation Overview (page B-2)

Installation Overview

B.1 Installation Overview

Table B–1 provides an overview of the installation process for a system that does not have a management server.

Table B–1 Installation Process: When the System Does Not Have a Management Server

Perform This Installation Task...	On...	Scope of Task
Set Up Node 0:		
1. Set the Console Variables	Node 0	Node 0
2. Check the System Firmware	Node 0	Node 0
3. Install the Tru64 UNIX Operating System	Node 0	Node 0
4. Customize the System Configuration	Node 0	Node 0
5. Install the Latest Operating System Patch Software	Node 0	Node 0
6. Configure the RIS Server	Node 0	Node 0
7. Install the HP AlphaServer SC System Software	Node 0	Node 0
8. Install the HP Fortran Run-Time Libraries	Node 0	Node 0
9. Install Layered Products (Optional)	Node 0	Node 0
10. Install the SANworks Storage System Scripting Utility	Node 0	Node 0
11. Add sysconfigtab Parameters	Node 0	Node 0
12. Define the RMS Master Node (rmshost)	Node 0	Systemwide
Set Up the SC Database:		
13. Set Up the SC Database	Node 0	Node 0 ¹
Check All Nodes in the HP AlphaServer SC System, Except Node 0:		
14. Check the State of the Nodes	Node 0	All remaining nodes
15. Check the System Firmware	Node 0	All remaining nodes
Configure and Diagnose the HP AlphaServer SC Interconnect		
16. Upgrading the HP AlphaServer SC Interconnect Control Processor Software	Node 0	Node 0
17. Creating an Interconnect Configuration Using SC Viewer	Node 0	Node 0
18. Confirming the Operation of the HP AlphaServer SC Interconnect	Node 0	Node 0
Set Up the SC Monitor System		
19. Add a SAN Appliance as Monitored Devices	Node 0	Node 0
20. Change Terminal Servers Monitoring Distribution	Node 0	Node 0
21. Change Extreme Switches Monitoring Distribution	Node 0	Node 0
22. Activate the Changes	Node 0	Node 0

Table B–1 Installation Process: When the System Does Not Have a Management Server

Perform This Installation Task...	On...	Scope of Task
Assign Cabinets in the SC Database		
23. Load the Physical Relationships	Node 0	Node 0
24. Populate Cabinets with Nodes	Node 0	Node 0
25. Populate Cabinets with Other Hardware	Node 0	Node 0
Review the SC Database Settings:		
26. Review the SC Database Disk Settings	Node 0	Node 0
Transform Node 0 into a Single Node Domain:		
27. Transform Node 0 into a Single Node Domain	Node 0	Node 0
28. Run the HP AlphaServer SC Interconnect Tests on Node 0	Node 0	Node 0
Building the Domains:		
29. Review the SC Database System Settings	Node 0	Systemwide
30. Add sysconfigtab Parameters	Node 0	Systemwide
31. Create the Domains	Node 0	Systemwide
32. Boot the System	Node 0	Systemwide
33. Complete the Setup of the Domains	Node 0	Systemwide
Completing the Installation:		
34. Configure the External Network Interfaces	<i>Specified nodes</i>	<i>Specified nodes</i>
35. Improve Cluster Availability	<i>Specified nodes</i>	<i>Specified nodes</i>
36. Load File System Configuration Data in the SC Database	Node 0	Systemwide
37. Initial Setup of Monitoring for HSG80 RAID Systems	Node 0	Systemwide
38. Configure the RMS Database	Node 0	Systemwide
39. Provide RMS with CAA Failover Capability	Node 0	Systemwide
40. Enable CMF as a CAA Application	Node 0	Systemwide
41. Run the Example MPI Program	<i>Any node</i>	<i>Specified nodes</i>
42. Verify the HP AlphaServer SC Interconnect	Node 0	Systemwide
43. Configure LSM	First three nodes of each domain	First three nodes of each domain
44. Verify Swap Mode	N/A	N/A
45. Add a Second Rail to an HP AlphaServer SC System after Domain Creation	Node 0	Systemwide

¹The `rmshost` alias is manually defined on Node 0 and automatically propagated to each domain.

Checklist: Adding a Management Server to a Cluster

Use this checklist to ensure that you complete all installation tasks in the correct order, when adding a management server after cluster creation:

Table C–1 Installation Checklist

Installation Task	Page	Task Complete?
Pre-Installation Planning:		
1. Review the Release Notes	2–2	
2. Assign the External Network IP Addresses	2–2	
3. Assign the System Name and Default Cluster Aliases	2–3	
4. Plan the Local and Global Storage	2–7	
5. Choose the Root Password	2–16	
6. Record the External Gateway IP Address	2–16	
Physical Installation:		
7. Connect the Management Network	3–11	
8. Populate the HP AlphaServer SC PCI Slots	3–17	
9. Configure Hardware for a Dual-Rail Configuration	3–20	
10. Connect the HP AlphaServer SC Interconnect	3–20	
11. Connect the Node Console Port	3–21	
12. Configure the HP AlphaServer SC Interconnect Control Card with an IP Address	3–22	
13. Configure the Terminal Servers with an IP Address	3–22	
14. Connect the Fibre Channel Switches	3–23	
15. Configure the System Storage on the HSG80	3–27	
16. Edit the Configuration Script	3–55	
Set Up the Management Server:		
17. Set the Console Variables	5–3	
18. Check the System Firmware	5–4	

Table C–1 Installation Checklist

Installation Task	Page	Task Complete?
19. Install the Tru64 UNIX Operating System	5–6	
20. Register Licenses (PAKs)	5–10	
21. Set Up Networks	5–14	
22. Configure DNS (BIND)	5–17	
23. Configure NTP	5–18	
24. Configure NFS	5–18	
25. Configure NIS	5–19	
26. Configure Mail	5–21	
27. Configure Printers	5–21	
28. Install the Operating System Patch Software	5–22	
29. Install and Configure the Clustered Management Server	5–22	
30. Install the HP AlphaServer SC System Software	5–32	
31. Install the HP Fortran Run-Time Libraries	5–33	
32. Install Layered Products (Optional)	5–33	
33. Disable RIS on Node 0	C–6	
34. Build the New Kernel and Reboot	5–35	
35. Update SRA	C–3	
36. Check the System Firmware	5–46	
37. Set Up RMS Partitions	8–10	

C.1 Update SRA

Follows these steps to move the SC database from the existing node (usually Node 0) to the management server, update the `/etc/hosts` file, and configure RIS:

1. On a system without a management server, the `cmf` host is normally Node 0. Identify the `cmf` host, as follows:
`# sra dbget cmf.host`
2. Stop the console logger (`cmf`) on the `cmf` host, as follows:
`atlas0# /sbin/init.d/cmf stop`
3. Move the RMS database from `atlas0` to `atlasms` using the instructions provided in Section , You are now ready to upgrade the management server, as described in Section 4.4.2. (page 4-11).
4. On the management server, edit the SC database to enter the MAC address for Node 0. This value was not entered earlier by the `sra setup` command (see Section 6.2, step 22 on page 6–30), because Node 0 was not then at the SRM prompt. To enter the MAC address for Node 0, run the `sra edit` command, as shown below:

```
atlasms# sra edit
sra> node
node> edit atlas0
Id      Description                               Value
-----
...
[6 ] Cluster name                               atlasD0
[7 ] Hardware address (MAC)                     #
[8 ] Number of votes                             1
...
# = no default value exists
-----

Select attributes to edit, q to quit
eg. 1-5 10 15

edit? 7

enter a new value, probe or auto
auto = generate value from system
probe = probe hardware for value

ip00:hardware address (MAC)                     [] (set)
new value? probe

info Connected through cmf
info Connected through cmf

ip00:hardware address (MAC)                     [00-00-F8-1B-2E-BA] (probed)

correct? [y|n] y

Remote Installation Services (ris) should be updated
Update ris ? n
```

Update SRA

```
node> quit
sra>
```

5. On the management server, update the `/etc/hosts` file and configure RIS by running the `sra edit` command, as shown below.

Note:

You must use the `sra edit` command — the `sra setup` command will not update the `cmf` host when this value is already set in the database.

The management server is already at the `sra>` prompt. Enter the following commands:

```
sra> sys
sys> edit system
```

Id	Description	Value
...		
[8]	Node running console logging daemon (cmfd)	atlas0
...		
[28]	Management Server name	atlas0
...		

Select attributes to edit, q to quit
eg. 1-5 10 15

```
edit? 8 28
Node running console logging daemon (cmfd) [atlas0]
new value? atlasms
Management Server name [atlas0]
new value? atlasms
```

```
Node running console logging daemon (cmfd) [atlasms]
Management Server name [atlasms]
correct? [y|n] y
```

You have modified fields which effect the console logging system. The SC database will be updated. In addition you may chose to update (ping) the daemons to reload from the modified database, or restart the daemons.

```
Modify SC database only (1), update daemons (2), restart daemons (3) [3]:3
Finished adding nodes to CMF table
Finished updating nodes in CMF table
CMF reconfigure: success
sys> update hosts
Updating /etc/hosts...
sys> update ris all
Gateway for subnet 10 is 10.128.101.1
```

```

Setup RIS for host atlas0
.
.
Setup RIS for host atlas1023
sys> quit
sra> quit
#

```

6. Add an entry to the RMS database for the management server, as follows:
rcontrol create node name=atlasms
7. Start the RMS daemons on the management server, as follows:
atlasms# **/sbin/init.d/rms start**
8. If you have changed the factory default password (that is, **system**) for the terminal server(s), perform the following steps (where **atlas-tcl** is an example terminal server name):
 - a. Connect to the terminal server, as follows:
sra -c atlas-tcl
 - b. Change the terminal server password to the factory default password, as follows:
access
Network Access SW V2.4 BL50 for DS732
(c) Copyright 2000, Digital Networks - All Rights Reserved
Please type HELP if you need assistance
Enter username> **system**
Local> **set priv**
Password> **site_specific_password**¹
Local> **change server privileged password**
Password> **system**¹
Verification> **system**¹
Local>
 - c. Run the **sra ds_passwd** command to change the terminal server password back to the site-specific value, as follows:
sra ds_passwd -server atlas-tcl
This command will set the password on the named terminal server (atlas-tcl).
Confirm change password for server atlas-tcl [yes]:
Enter new password for atlas-tcl: **site_specific_password**¹
Please re-enter new password for atlas-tcl: **site_specific_password**¹
Info: connecting to terminal server atlas-tcl (10.128.100.1)
Info: Connected through cmf

This command sets the password on the terminal server, and updates the entry in the SC database.

1. The value that you enter is not echoed on screen.

Disable RIS on Node 0

9. The value of SC_MS in /etc/rc.config.common needs to be changed on the management server and on all clusters as follows:

```
atlasms# rcmgr -c set SC_MS atlasms
atlasms# sra command -domains all -m 1 -command "rcmgr -c set SC_MS atlasms"
```
10. In order for the HP AlphaServer SC system management daemons on the cluster nodes to correctly recognize and sign on with the new management server, it is advisable to shut down and reboot the cluster nodes at this point.

C.2 Disable RIS on Node 0

When adding a management server to an HP AlphaServer SC system, you typically configure the management server as the RIS server for the system. However, Node 0 was previously configured as the RIS server. You should remove this functionality from Node 0 before you configure the management server as the RIS server, as described in this section.

To disable RIS, some tasks are mandatory and others are optional.

C.2.1 Mandatory Tasks

You must perform the following tasks — on any node in the first domain:

1. In the /etc/rc.config.common file, change the setting of the JOIND entry from "yes" to "no":
Before: JOIND="yes"
After: JOIND="no"
2. Stop the joind daemon, as follows:
/sbin/init.d/dhcp stop

C.2.2 Optional Tasks

You may also choose to remove the Tru64 UNIX software product from the RIS environment on Node 0 — this is an optional task. To do this, perform the following steps:

1. Start RIS on Node 0, as follows:
ris
2. Choose the *DELETE software products* option by entering **d** at the prompt.
3. Enter the product number that represents the product to be removed; typically, the product number for Tru64 UNIX is **1**.
4. When prompted to remove all clients associated with the product, enter **yes**.

C.3 Stopping the RMS System and mSQL

To stop the RMS system, perform the following steps:

1. Ensure that there are no allocated resources. One way to do this is to stop each partition using the `kill` option, as shown in the following example:

```
# rcontrol stop partition=big option kill
```

- 2.

Note:

If the `rms` CAA application has not been enabled, skip this step.

If the `rms` CAA application has been enabled and is running, stop the `rms` application.

You can determine the current status of the `rms` application by running the `caa_stat` command on the first domain in the system (that is, `atlasD0`, where `atlas` is an example system name) as follows:

```
# caa_stat SC20rms
```

To stop the `rms` application, use the `caa_stop` command, as follows:

```
# caa_stop SC20rms
```

3. Stop the RMS daemons on every node, by running the following command once on any node:

```
# rmsctl stop
```

Note:

If the `rms` CAA application has been enabled and you did not stop the `rms` application as described in step 2, then you will not be able to stop the RMS daemons in this step — CAA will automatically restart RMS daemons on the node where the `rms` application was last located.

4. Stop the `mSQL` daemon in one of the following ways, depending on whether you have registered `mSQL` as a CAA service:

- Case 1: `mSQL` is registered with CAA

If the `mSQL` CAA application has been enabled and is running, stop the `mSQL` application.

You can determine the current status of the `mSQL` application by running the `caa_stat` command on the first domain in the system (that is, `atlasD0`, where `atlas` is an example system name), as follows:

```
# caa_stat SC05mSQL
```

Stopping the RMS System and mSQL

To stop the `msql` application, use the `caa_stop` command, as follows:

```
# caa_stop SC05msql
```

- Case 2: `msql` is not registered with CAA

If the `msql` CAA application has not been enabled, stop the `mSQL` daemon by running the following command on the RMS master node (`rmshost` — usually Node 0):

```
# /sbin/init.d/msqld stop
```

This process stops the RMS system.

At this stage, any attempt to use an RMS command will result in an error similar to the following:

```
rinfo: Warning: Can't connect to mSQL server on rmshost: retrying ...
```

This is because the `mSQL` daemon was stopped in step 4 above. If you skip step 4 (so that the `mSQL` daemon is running, but RMS daemons are stopped), you will be able to access the database but unable to execute commands that require the RMS daemons. Different commands need different RMS daemons, so the resulting error messages will differ. A typical message is similar to the following:

```
rcontrol: Warning: RMS server pmanager-parallel (rmshost) not responding
```

Note:

If you did not perform step 1 above, this process will not stop any jobs that are running when the RMS system is stopped.

C.3.1 Manually Starting RMS

If you stopped RMS as described in Section C.3 on page C–7, you can restart RMS by performing the following steps:

1. Start the `mSQL` daemon in one of the following ways, depending on whether you have registered `msql` as a CAA service:

- Case 1: `msql` is registered with CAA

If the `msql` CAA application has been enabled and is stopped, start the `msql` application.

You can determine the current status of the `msql` application by running the `caa_stat` command on the first domain in the system (that is, `atlasD0`, where `atlas` is an example system name) as follows:

```
# caa_stat SC05msql
```

To start the `msql` application, use the `caa_start` command as follows:

```
# caa_start SC05msql
```

Stopping the RMS System and mSQL

- Case 2: msql is not registered with CAA

If the msql CAA application has not been enabled, start the mSQL daemon by running the following command on the RMS master node (rms host — usually Node 0):

```
# /sbin/init.d/msqld start
```

2. **Note:**

If the rms CAA application has not been enabled, skip this step.

If the rms CAA application has been enabled and is stopped, start the rms application.

You can determine the current status of the rms application by running the `caa_stat` command on the first domain in the system (that is, `atlasD0`, where `atlas` is an example system name) as follows:

```
# caa_stat SC20rms
```

To start the rms application, use the `caa_start` command as follows:

```
# caa_start SC20rms
```

3. Start the RMS daemons on the remaining nodes, by running the following command once on any node:

```
# rmsctl start
```

For more information about RMS, see Chapter 5 of the *HP AlphaServer SC System Administration Guide*.

Information Checklists

Table D–1 Summit Switch and Terminal Server Port Numbers

Attribute	S1 ¹	S2 ¹	T/Server ¹	Page
Port Number of Summit switch port connected to the terminal server			n/a	3–12
Port Number of the Summit switch port connected to the Fibre Channel switch			n/a	3–12
Port Number of the Summit switch port connected to the AlphaServer SC Interconnect Control Processor				3–12
Port Number of terminal server port connected to the management server	n/a	n/a		3–21

¹S1 = first Summit switch; S2 = second Summit switch, and so on; T/Server = terminal server

Table D–2 Tru64 UNIX System Attributes

Attribute	MS ¹	Domains				Page
		D0	D1	D2	D3	
AlphaServer SC system name (for example, atlas)						2–3
Host name						5–7 / 6–5
Host name IP address						5–14 / 6–12
External gateway						5–14 / 6–12
Tru64 UNIX /, /usr, and /var device from console (for example, dka200)						2–9
Tru64 UNIX / partition (for example, dsk2a)						5–8 / 6–6
Tru64 UNIX /usr partition (for example, dsk2g)						5–8 / 6–6
Tru64 UNIX /var partition (for example, dsk2h)						5–8 / 6–6

¹MS = Management Server

Table D–3 Domain Attributes

Attribute	Domains				Page
	D0	D1	D2	D3	
Number of nodes					5–37 / 6–24
Hardware type of the system (for example, ES40)					5–37 / 6–24
Number of domains					5–37 / 6–24
First node of this domain					5–37 / 6–24
domain name					2–4, 5–37 / 6–24
Default cluster alias IP address					2–4, 5–37 / 6–24
Terminal server model (for example, terminal server)					5–37 / 6–24
Number of ports on the terminal server					5–37 / 6–24
Terminal server IP address (first terminal server)					5–37 / 6–24
First port on terminal server					5–37 / 6–24
Domain /, /usr, and /var device from console (for example, DKC101)	primary:				6–43
	backup:				
Domain / partition (for example, dsk3b)	primary:				6–43
	backup:				
Domain /usr partition (for example, dsk3g)	primary:				6–43
	backup:				
Domain /var partition (for example, dsk3h)	primary:				6–43
	backup:				
Domain / disk size (percentage of total size)	primary:				6–43
	backup:				
Domain /usr disk size (percentage of total)	primary:				6–43
	backup:				
Domain /var disk size (percentage of total)	primary:				6–43
	backup:				
Generic boot disk (for example, dsk4)					6–30

Table D–4 Member Attributes of Domain D0

Attribute	Node 0: Member 1	Node 1: Member 2	Node 2: Member 3	Node <i>n</i>: Member <<i>n</i>+1>
Member host name ¹	atlas0	atlas1	atlas2	atlasn
Host name IP address ¹	10.128.0.1	10.128.0.2	10.128.0.3	10.128.0.<n+1>
Member ID ² (memberid: 1–32)	1	2	3	n+1
Number of votes assigned to this member	1	1	1	0
Boot disk ³ (for example, dsk0)	dsk0	dsk6	dsk8	dsk<2m+2>
Boot device from console (for example, dka0)	dka0	dka0	dka0	dka0
Cluster interconnect IP name ⁴	atlas0-ics0	atlas1-ics0	atlas2-ics0	atlasn-ics0
Cluster interconnect IP address ⁴	10.0.0.1	10.0.0.2	10.0.0.3	10.0.0.<n+1>
System interconnect IP name ⁵	atlas0-eip0	atlas1-eip0	atlas2-eip0	atlasn-eip0
System interconnect IP address ⁶	10.64.0.1	10.64.0.2	10.64.0.3	10.64.0.<n+1>
External Network IP name ⁶	atlas0-ext1	atlas1-ext1	atlas2-ext1	atlasn-ext1
External Network IP address	site-dependent	site-dependent	site-dependent	site-dependent
External Network subnet mask	site-dependent	site-dependent	site-dependent	site-dependent

¹The first member inherits its host name and IP address from the Tru64 UNIX operating system.

²The default member ID for the first member is 1, and increments by 1 for each additional member.

³This naming convention is based on the configuration described in this guide: 3 local disks on the first node, 4 disks shared on the RAID subsystem, and 2 local disks on all other nodes.

⁴By default, the installation programs offer IP addresses on the 10.0 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-ics0`.

⁵By default, the installation programs offer IP addresses on the 10.64 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-eip0`.

⁶The suffix is `ext1` for the first external network, `ext2` for the second, and so on.

Table D–5 Member Attributes of Domain D1

Attribute	Node 32: Member 1	Node 33: Member 2	Node 34: Member 3	Node <i>n</i>: Member <i>m</i>
Member host name ¹	atlas32	atlas33	atlas34	atlasn
Host name IP address ¹	10.128.0.33	10.128.0.34	10.128.0.35	10.128.0.<n+1>
Member ID ² (memberid: 1–32)	1	2	3	m
Number of votes assigned to this member	1	1	1	0
Boot disk ³ (for example, dsk0)	dsk0	dsk6	dsk8	dsk<2m+2>
Boot device from console (for example, dka0)	dka0	dka0	dka0	dka0
Cluster interconnect IP name ⁴	atlas32-ics0	atlas33-ics0	atlas34-ics0	atlasn-ics0
Cluster interconnect IP address ⁴	10.0.0.33	10.0.0.34	10.0.0.35	10.0.0.<n+1>
System interconnect IP name ⁵	atlas32-eip0	atlas33-eip0	atlas34-eip0	atlasn-eip0
System interconnect IP address ⁶	10.64.0.33	10.64.0.34	10.64.0.35	10.64.0.<n+1>
External Network IP name ⁶	atlas32-ext1	atlas33-ext1	atlas34-ext1	atlasn-ext1
External Network IP address	site-dependent	site-dependent	site-dependent	site-dependent
External Network subnet mask	site-dependent	site-dependent	site-dependent	site-dependent

¹The first member inherits its host name and IP address from the Tru64 UNIX operating system.

²The default member ID for the first member is 1, and increments by 1 for each additional member.

³This naming convention is based on the configuration described in this guide: 3 local disks on the first node, 4 disks shared on the RAID subsystem, and 2 local disks on all other nodes.

⁴By default, the installation programs offer IP addresses on the 10.0 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-ics0`.

⁵By default, the installation programs offer IP addresses on the 10.64.0 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-eip0`.

⁶The suffix is `ext1` for the first external network, `ext2` for the second, and so on.

Table D–6 Member Attributes of Domain D2

Attribute	Node 64: Member 1	Node 65: Member 2	Node 66: Member 3	Node <i>n</i>: Member <i>m</i>
Member host name ¹	atlas64	atlas65	atlas66	atlas <i>n</i>
Host name IP address ¹	10.128.0.65	10.128.0.66	10.128.0.67	10.128.0.< <i>n</i> +1>
Member ID ² (memberid: 1–32)	1	2	3	<i>m</i>
Number of votes assigned to this member	1	1	1	0
Boot disk ³ (for example, dsk0)	dsk0	dsk6	dsk8	dsk<2 <i>m</i> +2>
Boot device from console (for example, dka0)	dka0	dka0	dka0	dka0
Cluster interconnect IP name ⁴	atlas64-ics0	atlas65-ics0	atlas66-ics0	atlas <i>n</i> -ics0
Cluster interconnect IP address ⁴	10.0.0.65	10.0.0.66	10.0.0.67	10.0.0.< <i>n</i> +1>
System interconnect IP name ⁵	atlas64-eip0	atlas65-eip0	atlas66-eip0	atlas <i>n</i> -eip0
System interconnect IP address ⁵	10.64.0.65	10.64.0.66	10.64.0.67	10.64.0.< <i>n</i> +1>
External Network IP name ⁶	atlas64-ext1	atlas65-ext1	atlas66-ext1	atlas <i>n</i> -ext1
External Network IP address	site-dependent	site-dependent	site-dependent	site-dependent
External Network subnet mask	site-dependent	site-dependent	site-dependent	site-dependent

¹The first member inherits its host name and IP address from the Tru64 UNIX operating system.

²The default member ID for the first member is 1, and increments by 1 for each additional member.

³This naming convention is based on the configuration described in this guide: 3 local disks on the first node, 4 disks shared on the RAID subsystem, and 2 local disks on all other nodes.

⁴By default, the installation programs offer IP addresses on the 10.0 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-ics0`.

⁵By default, the installation programs offer IP addresses on the 10.64 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-eip0`.

⁶The suffix is `ext1` for the first external network, `ext2` for the second, and so on.

Table D–7 Member Attributes of Domain D32

Attribute	Node 992: Member 1	Node 993: Member 2	Node 994: Member 3	Node <i>n</i>: Member <i>m</i>
Member host name ¹	atlas992	atlas993	atlas994	atlasn
Host name IP address ¹	10.128.7.97	10.128.7.98	10.128.7.99	10.128.(<i>n</i> +1)/ 128.(<i>n</i> +1)%128
Member ID ² (memberid: 1–32)	1	2	3	<i>m</i>
Number of votes assigned to this member	1	1	1	0
Boot disk ³ (for example, dsk0)	dsk0	dsk6	dsk8	dsk<2 <i>m</i> +2>
Boot device from console (for example, dka0)	dka0	dka0	dka0	dka0
Cluster interconnect IP name ⁴	atlas992- ics0	atlas993- ics0	atlas994- ics0	atlasn-ics0
Cluster interconnect IP address ⁴	10.0.7.97	10.0.7.98	10.0.7.99	10.0.0.< <i>n</i> +1>
System interconnect IP name ⁵	atlas992- eip0	atlas993- eip0	atlas994- eip0	atlasn-eip0
System interconnect IP address ⁴	10.64.7.97	10.64.7.98	10.64.7.99	10.64.7.< <i>n</i> +1>
External Network IP name ⁶	atlas992- ext1	atlas993- ext1	atlas994- ext1	atlasn-ext1
External Network IP address	site-dependent	site-dependent	site-dependent	site-dependent
External Network subnet mask	site-dependent	site-dependent	site-dependent	site-dependent

¹The first member inherits its host name and IP address from the Tru64 UNIX operating system.

²The default member ID for the first member is 1, and increments by 1 for each additional member.

³This naming convention is based on the configuration described in this guide: 3 local disks on the first node, 4 disks shared on the RAID subsystem, and 2 local disks on all other nodes.

⁴By default, the installation programs offer IP addresses on the 10.0 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-ics0`.

⁵By default, the installation programs offer IP addresses on the 10.64 subnet, with the host portion of the address set to the member ID and the IP name set to the member's host name followed by `-eip0`.

⁶The suffix is `ext1` for the first external network, `ext2` for the second, and so on.

Example Installation Output

This appendix provides samples of installation output from the following commands:

- `sra setup` (see Section E.1 on page E-2)
- `clu_create` (see Section E.2 on page E-20)
- `clu_add_member` (see Section E.3 on page E-23)
- `clu_quorum` (see Section E.4 on page E-25)
- `upgrade_check` (see Section E.5 on page E-27)

E.1 sra setup

This section provides the following example output:

- Sample sra setup Output When Using a Management Server (Example E–1 on page E–2)
- Sample sra setup Output When Not Using a Management Server (Example E–2 on page E–11)

Example E–1 Sample sra setup Output When Using a Management Server

```
atlasms# sra setup
```

We will need some configuration information in order to build the system.

The system name is used to derive both the cluster alias names and the hostname of each node in the system.
For example, in a 64-node, 2-cluster system with system name "atlas", the cluster aliases would be atlasD0 and atlasD1, while the node hostnames would be atlas0, atlas1, ..., atlas63.

The system name length is limited to 10 characters.

However, a deprecated naming scheme is also available for cases where the number of clusters is less than 27. Using this scheme the clusters in the above example would be named atlasA and atlasB.

Note: if the system has only one cluster, the cluster alias is the same as the system name. In the above example the cluster alias would be "atlas".

Enter the system name: **atlas**

Enter the number of nodes in the system: **1024**

Domain names in the system typically start at 0, for example atlasD0.
Or when using the old naming scheme they would start at A, for example atlasA.
However, it is possible to start at a different number.
Enter the domain number (or letter) of the first domain in the system [0]: **0**

Node names in the system typically start at 0, for example atlas0.
However, it is possible to start at a different number.
Enter the node number of the first node in the system [0]: **0**

Is this system a Management Server? [yes]: Enter the name of the Management Server
[atlasms]: **atlasms**

We need to know the hardware type of the system
Supported hardware types are ES40, ES45 and DS20L
If your system is to contain clusters of mixed hardware type,
enter the type which composes the majority of the system.

Enter the hardware type [ES40]: **ES45**

We need to know the number of Cluster Interconnect Rails used in the system. The supported number of rails is 1 or 2

Enter the number of Cluster Interconnect Rails [1]: **1**

Building the SC database (-m atlas -N "atlas[0-1023]" -t ES45 -I atlasms)

Enter the number of clusters in the system [1]: **32**

Do you wish to use the old alphabetic cluster naming scheme? [no]: **no**

The Internet protocol (IP) address associated with the Management Network for each node in the cluster is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 Management Network [10.128.4.1]: **10.128.4.1**

The address associated with the Compaq Alphaserver SC Cluster Interconnect for each node in the system is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 Cluster Interconnect [10.0.0.1]: **10.0.0.1**

The address associated with the Compaq Alphaserver SC System (Eip) Interconnect for each node in the system is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 System Interconnect [10.64.0.1]: **10.64.0.1**

If you know the model of the terminal server used to connect the Console Network, enter it here. Otherwise just accept the default or set it to "unknown"

Enter the terminal server model [DECserver900]: **DECserver900**

The IP address of each terminal server in the system is determined by incrementing from that of the first

Enter the IP address for the first DECserver900 [10.128.100.1]: **10.128.100.1**

In normal practice the terminal server ports are assigned sequentially starting from 1 for node 0

Enter the first port on the DECserver900 (in the range 1-32) [1]: **1**

The file /etc/clua_default needs to be populated. When populated, This file allows the selection of a preferred network for routes to Cluster Alias IP addresses. The contents of the /etc/clua_default file affect how gated behaves. SRA has selected the Management LAN as the network to provide a default route and the value 2 as the metric. If you wish to change his value use the SRA edit command. If you change the metric value in the database you will also need to change the value in /etc/clua_default file on each domain.

Enter the IP address for the Preferred Server cluster alias base address [10.128.106.1]: **10.128.106.1**

System name

atlas

sra setup

Number of nodes 1024
Domain number of first domain 0
Node number of first node 0
Management Server atlasms
IP address for node 0 Management Network 10.128.0.1
IP address for node 0 Cluster Interconnect 10.0.0.1
IP address for node 0 System Interconnect 10.64.0.1
Terminal server model DECserver900
IP address for the first DECserver900 10.128.100.1
First port on the DECserver900 1
Hardware type ES45
Number of Cluster Interconnect Rails 1
Number of clusters 32
Cluster naming scheme Alphanumeric
Preferred server cluster alias base address 10.128.106.1

Is this correct? [yes]: **yes**
Setting up clusters

If you choose manual cluster configuration, then for each cluster you will be asked to enter the external IP address for member 1, the cluster alias IP address, the cluster type (file server or compute server), and the start node.

If you choose automatic cluster configuration, the system will assign default values for the cluster type. The IP addresses for members 1 and 2 of each cluster, and the cluster aliases, will be assigned in sequence starting with the base IPs which you enter.

Configurations where all clusters (except possibly the first and last) are the same size can be automatically generated. Other configurations must be manually entered.

Would you like to use automatic cluster configuration? [yes]: **yes**
Number of nodes in the first cluster [32]: **32**
Number of nodes in subsequent clusters [32]: **32**
Enter the External Network IP address for node 0 []: **16.209.133.30**
Enter the IP for cluster 0 cluster alias []: **16.209.133.28**

Cluster Name	Start Node	Cluster Alias IP Address	Member 1 IP Address	Member 2 IP Address	Cluster Type
atlasD0	0	16.209.133.28	16.209.133.30	16.209.133.31	fs
atlasD1	32	16.209.133.29	16.209.133.32	16.209.133.33	cs
atlasD2	64	16.209.133.30	16.209.133.34	16.209.133.35	cs
atlasD3	96	16.209.133.31	16.209.133.36	16.209.133.37	cs
atlasD4	128	16.209.133.32	16.209.133.38	16.209.133.39	cs
atlasD5	160	16.209.133.33	16.209.133.40	16.209.133.41	cs
atlasD6	192	16.209.133.34	16.209.133.42	16.209.133.43	cs
atlasD7	224	16.209.133.35	16.209.133.44	16.209.133.45	cs
atlasD8	256	16.209.133.36	16.209.133.46	16.209.133.47	cs
atlasD9	288	16.209.133.37	16.209.133.48	16.209.133.49	cs

atlasD10	320	16.209.133.38	16.209.133.50	16.209.133.51	cs
atlasD11	352	16.209.133.39	16.209.133.52	16.209.133.53	cs
atlasD12	384	16.209.133.40	16.209.133.54	16.209.133.55	cs
atlasD13	416	16.209.133.41	16.209.133.56	16.209.133.57	cs
atlasD14	448	16.209.133.42	16.209.133.58	16.209.133.59	cs
atlasD15	480	16.209.133.43	16.209.133.60	16.209.133.61	cs
atlasD16	512	16.209.133.44	16.209.133.62	16.209.133.63	cs
atlasD17	544	16.209.133.45	16.209.133.64	16.209.133.65	cs
atlasD18	576	16.209.133.46	16.209.133.66	16.209.133.67	cs
atlasD19	608	16.209.133.47	16.209.133.68	16.209.133.69	cs
atlasD20	640	16.209.133.48	16.209.133.70	16.209.133.71	cs
atlasD21	672	16.209.133.49	16.209.133.72	16.209.133.73	cs
atlasD22	704	16.209.133.50	16.209.133.74	16.209.133.75	cs
atlasD23	736	16.209.133.51	16.209.133.76	16.209.133.77	cs
atlasD24	768	16.209.133.52	16.209.133.78	16.209.133.79	cs
atlasD25	800	16.209.133.53	16.209.133.80	16.209.133.81	cs
atlasD26	832	16.209.133.54	16.209.133.82	16.209.133.83	cs
atlasD27	864	16.209.133.55	16.209.133.84	16.209.133.85	cs
atlasD28	896	16.209.133.56	16.209.133.86	16.209.133.87	cs
atlasD29	928	16.209.133.57	16.209.133.88	16.209.133.89	cs
atlasD30	960	16.209.133.58	16.209.133.90	16.209.133.91	cs
atlasD31	992	16.209.133.59	16.209.133.92	16.209.133.93	cs

Is this correct? [yes]: **yes**

Adding nodes to the database

Finished adding nodes to database

Populating the SC Monitor database tables

add Extreme switch: extremel/10.128.103.1

...etc...

add Extreme switch: extreme32/10.128.103.32

Configuring cfengine

Update /etc/hosts? [yes]: Updating /etc/hosts...

Info - No SRA section in /etc/hosts file

CMF is a utility used to monitor and connect to the Console Network. It must be correctly configured before the cluster setup can continue

Which host should run the CMF daemon (cmfd)? [atlasms]: **atlasms**

Finished adding nodes to CMF table

Finished updating nodes in CMF table

If the terminal servers have never been configured, cmf will fail.

Configure terminal servers? [yes]: no

The CMF host is set as atlasms in the database, but is not responding

Attempting to restart CMF on this node

Restarting the CMF daemon, wait...ok

Checking we can talk to the daemon...ok

sra setup

The CMF daemon was successfully restarted on atlasms

Each member has its own boot disk, which has an associated clusterized UNIX device name; for example dsk7
An alternate boot disk may be configured -- this is optional.

Would you like to configure an alternate boot disk? [yes]: **yes**

If you USE an alternate boot disk, the swap space from the alternate boot disk is added to the swap space from the primary boot disk, thus doubling the available swap space. Also, the tmp and local partitions on the alternate boot disk are mounted on /tmp1 and /local1 respectively.

Would you like to use the alternate boot disk? [yes]: **yes**
Setting up primary boot disk

We need to know the SRM device name for the primary boot disk
If you know the SRM device name (eg dka0) enter it now
otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for primary boot disk? [probe]: **dka0**
Settings for the primary boot disk:

```
-----
UNIX device name           dsk0
SRM device name            dka0
swap partition size (%)    15
tmp partition size (%)     42
local partition size (%)   43
-----
```

Are these settings correct? [yes]: **yes**

Setting up secondary boot disk

We need to know the SRM device name for the secondary boot disk
If you know the SRM device name (eg dka100) enter it now
otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for secondary boot disk? [probe]: **dka100**
Settings for the secondary boot disk:

```
-----
UNIX device name           dsk1
SRM device name            dka100
swap partition size (%)    15
tmp partition size (%)     42
local partition size (%)   43
-----
```

Are these settings correct? [yes]: **yes**

Setting up Tru64 UNIX disk

We need to know the UNIX device name for the Tru64 UNIX disk
If you know the UNIX device name (eg dsk2) enter it now
otherwise accept the default

UNIX device name for the Tru64 UNIX disk? [dsk2]: **dsk2**

Settings for the Tru64 UNIX disk:

Id	Description	Value

[0]	Image Role	unix
[1]	Image name	first
[2]	UNIX device name	dsk2
[3]	SRM device name	#
[4]	Disk Location (Identifier)	
[5]	root partition size (%)	10
[6]	root partition	a
[7]	usr partition size (%)	35
[8]	usr partition	g
[9]	var partition size (%)	35
[10]	var partition	h
[11]	swap partition size (%)	20
[12]	swap partition	b

Are these settings correct? [yes]: **yes**

Setting up Cluster /, /usr & /var disk

We need to know the UNIX device name for the Cluster /, /usr & /var disk
If you know the UNIX device name (eg dsk3) enter it now
otherwise accept the default

UNIX device name for the Cluster /, /usr & /var disk? [dsk3]: **dsk3**

Settings for the Cluster /, /usr & /var disk:

Id	Description	Value

[0]	Image Role	cluster
[1]	Image name	first
[2]	UNIX device name	dsk3
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=1
[5]	root partition size (%)	5
[6]	root partition	b
[7]	usr partition size (%)	50
[8]	usr partition	g
[9]	var partition size (%)	45

sra setup

```
[10 ] var partition                               h
```

Are these settings correct? [yes]: **yes**

Would you like to configure a backup cluster disk? [no]: **yes**

Setting up backup Cluster /, /usr & /var disk

We need to know the UNIX device name for the backup Cluster /, /usr & /var disk

If you know the UNIX device name (eg dsk5) enter it now

otherwise accept the default

UNIX device name for the backup Cluster /, /usr & /var disk? [dsk5]: dsk5

Settings for the backup Cluster /, /usr & /var disk:

Id	Description	Value
[0]	Image Role	cluster
[1]	Image name	second
[2]	UNIX device name	dsk5
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=3
[5]	root partition size (%)	5
[6]	root partition	b
[7]	usr partition size (%)	50
[8]	usr partition	g
[9]	var partition size (%)	45
[10]	var partition	h

Are these settings correct? [yes]: **yes**

Setting up Generic boot disk

We need to know the UNIX device name for the Generic boot disk

If you know the UNIX device name (eg dsk4) enter it now

otherwise accept the default

UNIX device name for the Generic boot disk? [dsk4]: **dsk4**

Settings for the Generic boot disk:

Id	Description	Value
[0]	Image Role	gen_boot
[1]	Image name	first
[2]	UNIX device name	dsk4
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=2
[5]	default or not	#
[6]	swap partition size (%)	30
[7]	tmp partition size (%)	35
[8]	local partition size (%)	35

Are these settings correct? [yes]: **yes**

We need to know the SRM device name for the Management LAN adapter and the External LAN adapter.

If you know the SRM device name (eg eia0) enter it now otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for Management LAN adapter? [eia0]: **eia0**

Please enter the UNIX device name for the Management LAN adapter, eg "ee0". Accept the default if you are not sure

UNIX device name for Management LAN adapter? [ee0]: **ee0**

Settings for the Management LAN adapter

```
-----
UNIX device name                      ee0
SRM device name                      eia0
-----
```

Are these settings correct? [yes]: **yes**

SRM device name for External LAN adapter? [eib0]: **eib0**

Please enter the UNIX device name for the External LAN adapter, eg "ee1". Accept the default if you are not sure

UNIX device name for External LAN adapter? [ee1]: **ee1**

Settings for the External LAN adapter

```
-----
UNIX device name                      ee1
SRM device name                      eib0
-----
```

Are these settings correct? [yes]: **yes**

We will now do a hardware probe in order to collect the hardware ethernet address of each node.

Probe hardware? [yes]: **yes**

Nodes to probe [all]: **all**

Initialize node hardware during probe? [yes]: **yes**

Log file is /var/sra/sra.logd/sra.log.0

Display not loaded

02:38:46 atlas0 Phase: Initializing

02:59:09 atlas1023 status:success

sra setup

* Summary * Success: atlas[0-1023]

All nodes were successfully probed

In order to add members to the cluster we need to
add nodes to the Remote Installation Services (RIS) database

Add nodes to the RIS database? [yes]: **yes**

Gateway for subnet 10 is 10.128.0.1

Setup RIS for host atlas0

Setting up tftp in inetd.conf...

JOIN Server Release 4.1.0b for Compaq Tru64 UNIX

Copyright 1992-1998 Competitive Automation, Inc. All Rights Reserved.

DHCP daemon started

Setup RIS for host atlas1

...etc...

Setup RIS for host atlas1023

In order for future Operating System upgrades to function correctly we need to
add domains to the Remote Installation Services (RIS) database.

Add domains to the RIS database? [yes]: **yes**

Gateway for subnet 10 is 10.128.0.1

DHCP daemon started

Setup RIS for domain atlasD0

...etc...

Setup RIS for domain atlasD31

Enabling consolidated binary.errlog on rms host

If you need to make any changes to the database
use "sra edit"

To complete the installation, run "sra install" to install
the cluster

The following crontab entry will rotate cmf logs every two weeks:

"20 0 17,3 * * /sbin/init.d/cmf rotate"

Add entry to root crontab? [yes]: yes

The following crontab entry will archive and backup the rms database at 2:05 every
day:

"5 2 * * * /usr/bin/rmsbackup"

Add entry to root crontab? [yes]: **yes**

sra setup completed successfully.

#

Example E-2 Sample sra setup Output When Not Using a Management Server

```
atlas0# sra setup
```

We will need some configuration information in order to build the system.

The system name is used to derive both the cluster alias names and the hostname of each node in the system.

For example, in a 64-node, 2-cluster system with system name "atlas", the cluster aliases would be atlasD0 and atlasD1, while the node hostnames would be atlas0, atlas1, ..., atlas63.

The system name length is limited to 10 characters.

However, a deprecated naming scheme is also available for cases where the number of clusters is less than 27. Using this scheme the clusters in the above example would be named atlasA and atlasB.

Note: if the system has only one cluster, the cluster alias is the same as the system name. In the above example the cluster alias would be "atlas".

```
Enter the system name: atlas
```

```
Enter the number of nodes in the system: 1024
```

Domain names in the system typically start at 0, for example atlasD0. Or when using the old naming scheme they would start at A, for example atlasA. However, it is possible to start at a different number.

```
Enter the domain number (or letter) of the first domain in the system [0]: 0
```

Node names in the system typically start at 0, for example atlas0.

However, it is possible to start at a different number.

```
Enter the node number of the first node in the system [0]: 0
```

```
Is this system a Management Server? [yes]: no
```

We need to know the hardware type of the system

Supported hardware types are ES40, ES45 and DS20L

If your system is to contain clusters of mixed hardware type, enter the type which composes the majority of the system.

```
Enter the hardware type [ES40]: ES45
```

We need to know the number of Cluster Interconnect Rails used in the system. The supported number of rails is 1 or 2

```
Enter the number of Cluster Interconnect Rails [1]: 1
```

```
Building the SC database (-m atlas -N "atlas[0-3]" -t ES45 )
```

```
Enter the number of clusters in the system [1]: 32
```

```
Do you wish to use the old alphabetic cluster naming scheme? [no]: no
```

sra setup

The Internet protocol (IP) address associated with the Management Network for each node in the cluster is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 Management Network [10.128.0.1]: **10.128.0.1**

The address associated with the Compaq Alphaserver SC Cluster Interconnect for each node in the system is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 Cluster Interconnect [10.0.0.1]: **10.0.0.1**

The address associated with the Compaq Alphaserver SC System (Eip) Interconnect for each node in the system is determined by incrementing from the first node (atlas0)

Enter the IP address for the node 0 System Interconnect [10.64.0.11]: **10.64.0.11**

If you know the model of the terminal server used to connect the Console Network, enter it here. Otherwise just accept the default or set it to "unknown"

Enter the terminal server model [DECserver900]: **DECserver900**

The IP address of each terminal server in the system is determined by incrementing from that of the first

Enter the IP address for the first DECserver900 [10.128.100.1]: **10.128.100.1**

In normal practice the terminal server ports are assigned sequentially starting from 1 for node 0

Enter the first port on the DECserver900 (in the range 1-32) [1]: **1**

The file /etc/clua_default needs to be populated. When populated, This file allows the selection of a preferred network for routes to Cluster Alias IP addresses. The contents of the /etc/clua_default file affect how gated behaves. SRA has selected the Management LAN as the network to provide a default route and the value 2 as the metric. If you wish to change his value use the SRA edit command. If you change the metric value in the database you will also need to change the value in /etc/clua_default file on each domain.

Enter the IP address for the Preferred Server cluster alias base address [10.128.106.1]: **10.128.106.1**

```
-----
System name                      atlas
Number of nodes                  1024
Domain number of first domain    0
Node number of first node        0
Management Server
IP address for node 0 Management Network  10.128.0.1
IP address for node 0 Cluster Interconnect 10.0.0.1
IP address for node 0 System Interconnect  10.64.0.1
Terminal server model            DECserver900
IP address for the first DECserver900     10.128.100.1
First port on the DECserver900           1
```

```

Hardware type                               ES45
Number of Cluster Interconnect Rails        1
Number of clusters                          32
Cluster naming scheme                       Alphanumeric
Preferred server cluster alias base address 10.128.106.1
-----

```

Is this correct? [yes]: **yes**

Setting up clusters

If you choose manual cluster configuration, then for each cluster you will be asked to enter the external IP address for member 1, the cluster alias IP address, the cluster type (file server or compute server), and the start node.

If you choose automatic cluster configuration, the system will assign default values for the cluster type. The IP addresses for members 1 and 2 of each cluster, and the cluster aliases, will be assigned in sequence starting with the base IPs which you enter.

Configurations where all clusters (except possibly the first and last) are the same size can be automatically generated. Other configurations must be manually entered.

Would you like to use automatic cluster configuration? [yes]: **no**

Enter the first node in cluster atlasD0 [atlas0]: **atlas0**

Does node atlas0 have an External Network? [yes]: **yes**

Enter the External Network IP address for node atlas0 []: **10.196.3.137**

Enter the IP address for cluster alias atlasD0 []: **16.196.1.137**

Enter the Domain type for cluster atlasD0 [fs]: **fs**

Enter the first node in cluster atlasD1 [atlas32]: **atlas32**

...etc...

Enter the Domain type for cluster atlasD31 [cs]: **cs**

Adding nodes to the database

Adding node atlas0

...etc...

Adding node atlas1023

Populating the SC Monitor database tables

add Extreme switch: extreme1/10.128.103.1

...etc...

add Extreme switch: extreme32/10.128.103.32

Configuring cfengine

Update /etc/hosts? [yes]: **yes**

Updating /etc/hosts...

Info - No SRA section in /etc/hosts file

CMF is a utility used to monitor and connect

sra setup

to the Console Network. It must be correctly configured before the cluster setup can continue

Which host should run the CMF daemon (cmfd)? [atlas0]: **atlas0**

Adding node atlas0 to CMF table

...etc...

Adding node atlas1023 to CMF table

Finished updating nodes in CMF table

If the terminal servers have never been configured, cmf will fail.

Configure terminal servers? [yes]: yes

The CMF host is set as atlas0 in the database, but is not responding

Attempting to restart CMF on this node

Restarting the CMF daemon, wait...ok

Checking we can talk to the daemon...ok

The CMF daemon was sucessfully restarted on atlas0

Each member has its own boot disk, which has an associated clusterized UNIX device name; for example dsk7

An alternate boot disk may be configured -- this is optional.

Would you like to configure an alternate boot disk? [yes]: **yes**

If you USE an alternate boot disk, the swap space from the alternate boot disk is added to the swap space from the primary boot disk, thus doubling the available swap space. Also, the tmp and local partitions on the alternate boot disk are mounted on /tmp1 and /local1 respectively.

Would you like to use the alternate boot disk? [yes]: **yes**

Setting up primary boot disk

We need to know the SRM device name for the primary boot disk

If you know the SRM device name (eg dka0) enter it now otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for primary boot disk? [probe]: **dka0**

Settings for the primary boot disk:

```
-----
UNIX device name                      dsk0
SRM device name                       dka0
swap partition size (%)                15
tmp partition size (%)                 42
local partition size (%)               43
-----
```

Are these settings correct? [yes]: **yes**

Setting up secondary boot disk

We need to know the SRM device name for the secondary boot disk
 If you know the SRM device name (eg dka100) enter it now
 otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for secondary boot disk? [probe]: **dka100**
 Settings for the secondary boot disk:

```
-----
UNIX device name                dsk1
SRM device name                 dka100
swap partition size (%)         15
tmp partition size (%)          42
local partition size (%)        43
-----
```

Are these settings correct? [yes]: **yes**

Setting up Tru64 UNIX disk

We need to know the UNIX device name for the Tru64 UNIX disk
 If you know the UNIX device name (eg dsk2) enter it now
 otherwise accept the default

UNIX device name for the Tru64 UNIX disk? [dsk2]: **dsk2**
 Settings for the Tru64 UNIX disk:

Id	Description	Value
[0]	Image Role	unix
[1]	Image name	first
[2]	UNIX device name	dsk2
[3]	SRM device name	#
[4]	Disk Location (Identifier)	
[5]	root partition size (%)	10
[6]	root partition	a
[7]	usr partition size (%)	35
[8]	usr partition	g
[9]	var partition size (%)	35
[10]	var partition	h
[11]	swap partition size (%)	20
[12]	swap partition	b

Are these settings correct? [yes]: **yes**

Setting up Cluster /, /usr & /var disk

We need to know the UNIX device name for the Cluster /, /usr & /var disk
 If you know the UNIX device name (eg dsk3) enter it now
 otherwise accept the default

UNIX device name for the Cluster /, /usr & /var disk? [dsk3]: **dsk3**
 Settings for the Cluster /, /usr & /var disk:

sra setup

Id	Description	Value

[0]	Image Role	cluster
[1]	Image name	first
[2]	UNIX device name	dsk3
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=1
[5]	root partition size (%)	5
[6]	root partition	b
[7]	usr partition size (%)	50
[8]	usr partition	g
[9]	var partition size (%)	45
[10]	var partition	h

Are these settings correct? [yes]: **yes**

Would you like to configure a backup cluster disk? [no]: **yes**
Setting up backup Cluster /, /usr & /var disk

We need to know the UNIX device name for the backup Cluster /, /usr & /var disk
If you know the UNIX device name (eg dsk5) enter it now
otherwise accept the default

UNIX device name for the backup Cluster /, /usr & /var disk? [dsk5]: **dsk5**
Settings for the backup Cluster /, /usr & /var disk:

Id	Description	Value

[0]	Image Role	cluster
[1]	Image name	second
[2]	UNIX device name	dsk5
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=3
[5]	root partition size (%)	5
[6]	root partition	b
[7]	usr partition size (%)	50
[8]	usr partition	g
[9]	var partition size (%)	45
[10]	var partition	h

Are these settings correct? [yes]: **yes**

Setting up Generic boot disk

We need to know the UNIX device name for the Generic boot disk
If you know the UNIX device name (eg dsk4) enter it now
otherwise accept the default

UNIX device name for the Generic boot disk? [dsk4]: **dsk4**
Settings for the Generic boot disk:

Id	Description	Value

[0]	Image Role	gen_boot
[1]	Image name	first
[2]	UNIX device name	dsk4
[3]	SRM device name	#
[4]	Disk Location (Identifier)	IDENTIFIER=2
[5]	default or not	#
[6]	swap partition size (%)	30
[7]	tmp partition size (%)	35
[8]	local partition size (%)	35

Are these settings correct? [yes]: **yes**

We need to know the SRM device name for the Management LAN adapter and the External LAN adapter.

If you know the SRM device name (eg eia0) enter it now otherwise accept the default

"probe" will probe a node at SRM for the device name

SRM device name for Management LAN adapter? [eia0]: **eia0**

Please enter the UNIX device name for the Management LAN adapter, eg "ee0". Accept the default if you are not sure

UNIX device name for Management LAN adapter? [ee0]: **ee0**

Settings for the Management LAN adapter

```
-----
UNIX device name                ee0
SRM device name                 eia0
-----
```

Are these settings correct? [yes]:

SRM device name for External LAN adapter? [eib0]: **eib0**

Please enter the UNIX device name for the External LAN adapter, eg "ee1". Accept the default if you are not sure

UNIX device name for External LAN adapter? [ee1]: **ee1**

Settings for the External LAN adapter

```
-----
UNIX device name                ee1
SRM device name                 eib0
-----
```

Are these settings correct? [yes]: **yes**

We will now do a hardware probe in order to collect the hardware ethernet address of each node.

sra setup

```
Probe hardware? [yes]: yes
Nodes to probe [all]: all
Initialize node hardware during probe? [yes]: yes
Log file is /var/sra/sra.logd/sra.log.0
```

```
Display not loaded
18:38:50 atlas1      Phase: Initializing
...etc...
18:59:13 atlas1023   status:success
* Summary * Success: atlas[1-1023]
```

All nodes were successfully probed

In order to add members to the cluster we need to
add nodes to the Remote Installation Services (RIS) database

```
Add nodes to the RIS database? [yes]: yes
Gateway for subnet 10 is 10.128.0.1
Setup RIS for host atlas1
Setting up tftp in inetd.conf...
JOIN Server Release 4.1.0b for Compaq Tru64 UNIX
Copyright 1992-1998 Competitive Automation, Inc. All Rights Reserved.
```

```
DHCP daemon started
Setup RIS for host atlas2
...
...
Setup RIS for host atlas1023
```

In order for future Operating System upgrades to function correctly we need to
add domains to the Remote Installation Services (RIS) database.

```
Add domains to the RIS database? [yes]: yes
Gateway for subnet 10 is 10.128.0.1
Setup RIS for domain atlasD0
...etc...
Setup RIS for domain atlasD31
```

Enabling consolidated binary.errlog on rms host

If you need to make any changes to the database
use "sra edit"

To complete the installation, run "sra install" to install
the cluster

```
The following crontab entry will rotate cmf logs every two weeks:
"20 0 16,2 * * /sbin/init.d/cmf rotate"
Add entry to root crontab? [yes]: yes
```

```
The following crontab entry will archive and backup the rms database at 2:05 every
day:
"5 2 * * * /usr/bin/rmsbackup"
```

Add entry to root crontab? [yes]: **yes**

sra setup completed successfully.

#

E.2 clu_create

One of the tasks performed by the `sra clu_create` command is to run the `clu_create` command. Each time you run `clu_create`, it writes log messages to the file `/cluster/admin/clu_create.log`. Example E-3 shows a sample `clu_create` log file.

Example E-3 Sample clu_create Log File

```
#####
clu_create on 'atlas0' begin logging at Tue Nov 18 09:25:20 EST 2003
-----
Cluster interconnect present

*** Error ***
This system has only Tru64 UNIX patches installed.
Please install the latest TruCluster Server patches on your system.
You can obtain the most recent patch kit from:
    http://www.support.compaq.com/patches/

*** Info ***
clu_create: Using configuration file '/tmp/sra_tmp.27624'.
Checking cluster name: atlasD0.
Checking cluster alias IP address: 16.21.24.64.
Checking the cluster root partition: dsk3b.
Checking the cluster usr partition: dsk3g.
Checking the cluster var partition: dsk3h.
Checking cluster member ID: 1.
Checking number of votes for this member: 1.
Checking the member boot disk: dsk0.
Checking virtual cluster interconnect device: ics0.
Checking virtual cluster interconnect IP name: atlas0-ics0.
Checking virtual cluster interconnect IP address: 10.0.1.1.
Checking physical cluster interconnect interface device name(s): mc0.
clu_common: Cluster interconnect present

Creating required disk labels.
    Creating disk label on member disk: dsk0.
    Initializing cnx partition on member disk: dsk0h.

Creating AdvFS domains:
    Creating AdvFS domain 'root1_domain#root' on partition '/dev/disk/dsk0a'.
    Creating AdvFS domain 'cluster_root#root' on partition '/dev/disk/dsk3b'.
    Creating AdvFS domain 'cluster_usr#usr' on partition '/dev/disk/dsk3g'.
    Creating AdvFS domain 'cluster_var#var' on partition '/dev/disk/dsk3h'.

Populating clusterwide root, usr, and var file systems:
    Copying root file system to 'cluster_root#root'.
.
    Copying usr file system to 'cluster_usr#usr'.
.....
    Copying var file system to 'cluster_var#var'.
..
```

Creating Context Dependent Symbolic Links (CDSLs) for file systems:

- Creating CDSLs in root file system.
- Creating CDSLs in usr file system.
- Creating CDSLs in var file system.
- Creating links between clusterwide file systems.

Populating member's root file system.

Modifying configuration files required for cluster operation:

- Creating /etc/fstab file.
- Configuring cluster alias.
- Updating member-specific /etc/inittab file with 'cms' entry.
- Updating /etc/rc.config file.
- Creating gated.conf file.
- Updating /etc/sysconfigtab file.

clu_create: type is mct!

- Retrieving cluster_root major and minor device numbers.
- Creating cluster device file CDSLs.
- Updating /.rhosts - adding hostname 'atlasD0'.
- Updating /etc/hosts.equiv - adding hostname 'atlasD0'.
- Updating /.rhosts - adding hostname 'atlas0-ics0'.
- Updating /etc/hosts.equiv - adding hostname 'atlas0-ics0'.

Configuring /.shosts file for: atlasD0

- Updating /etc/ifaccess.conf - adding deny entry for 'ee0'.
- Updating /etc/ifaccess.conf - adding deny entry for 'ee1'.
- Updating /etc/ifaccess.conf - adding deny entry for 'sl0'.
- Updating /etc/ifaccess.conf - adding deny entry for 'tun0'.
- Updating /etc/ifaccess.conf - adding deny entry for 'tun1'.
- Updating /etc/cfgmgr.auth - adding hostname 'atlas0'.
- Finished updating member1-specific area.

Building a kernel for this member.

- Saving kernel build configuration.
- The kernel will now be configured using the doconfig program.

*** Warning ***

File in /usr/sys/BINARY found as a file, expected symlink: GENERIC.mod.

*** Warning ***

File in /usr/sys/BINARY found as a file, expected symlink: GENERIC_EXTRAS.mod.

*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***

Saving /sys/conf/ATLAS0 as /sys/conf/ATLAS0.bck

*** PERFORMING KERNEL BUILD ***

Working....Tue Nov 18 09:31:26 EST 2003

The new kernel is /sys/ATLAS0/vmunix

Finished running the doconfig program.

clu_create

```
The kernel build was successful and the new kernel
has been copied to this member's boot disk.
Restoring kernel build configuration.
```

```
Updating console variables.
```

```
Setting console variable 'bootdef_dev' to dsk0.
Setting console variable 'boot_dev' to dsk0.
Setting console variable 'boot_reset' to ON.
Saving console variables to non-volatile storage.
```

```
clu_create: Cluster created successfully.
```

```
To run as a single member cluster, this system must be rebooted.
To reboot this system as a cluster member, do the following:
    o shutdown -r now
```

```
-----
clu_create on 'Tue Nov 18 09:32:10 EST 2003' end logging at atlas0
```


E.3 clu_add_member

One of the tasks performed by the `sra add_member` command is to run the `clu_add_member` command. Each time you run `clu_add_member`, it writes log messages to the file `/cluster/admin/clu_add_member.log`.

Example E-4 shows a sample `clu_add_member` log file.

Example E-4 Sample clu_add_member Log File

```
#####
clu_add_member on 'atlas0' begin logging at Tue Nov 18 09:37:43 EST 2003
-----

*** Info ***
clu_add_member: Using configuration file '/tmp/sra_tmp.527123'.
Checking member's hostname: atlas1.
Checking cluster member ID: 2.
Checking number of votes for this member: 0.
Checking the member boot disk: dsk6.
Checking virtual cluster interconnect device: ics0.
Checking virtual cluster interconnect IP name: atlas1-ics0.
Checking virtual cluster interconnect IP address: 10.0.0.2.
Checking physical cluster interconnect interface device name(s): mc0.

Creating required disk labels.
  Creating disk label on member disk: dsk6.
  Initializing cnx partition on member disk: dsk6h.

Creating AdvFS domains:
  Creating AdvFS domain 'root1_domain#root' on partition 'dsk6a'.

Creating cluster member-specific files:
  Creating new member's root member-specific files.
  Creating new member's usr  member-specific files.
  Creating new member's var  member-specific files.
  Creating new member's boot member-specific files.

Modifying configuration files required for new member operation:
  Updating /etc/rc.config.
  Updating /etc/sysconfigtab.
  Updating member-specific /etc/inittab file with 'cms' entry.
  Updating /.rhosts - adding hostname 'atlas1-ics0'.
  Updating /etc/hosts.equiv - adding hostname 'atlas1-ics0'.
  Updating /etc/cfgmgr.auth - adding hostname 'atlas1'.
  Configuring cluster alias.
  Configuring Network Time Protocol for new member.
  Adding interface 'atlas0-ics0' as an NTP peer to member 'atlas1'.
  Adding interface 'atlas1-ics0' as an NTP peer to member 'atlas0'.

Configuring automatic subset configuration and kernel build.
```

clu_add_member

clu_add_member: Initial member 2 configuration completed successfully.
From the newly added member's console, perform the following steps to
complete the newly added member's configuration:

1. Set the console variable 'boot_osflags' to 'A'.
2. Identify the console name of the newly added member's boots device:

```
>>>show device
```

The newly added member's boot device has the following properties:

```
Manufacturer: DEC
Model: HSG80
Target: IDENTIFIER=14
Lun: UNKNOWN
Serial Number: SCSI-WWID:01000010:6000-1fe1-000d-6480-0009-0440-4188-00ef
```

Note: The SCSI bus number may differ when viewed from different members.

3. Boot the newly added member using genvmunix:

```
>>>boot -file genvmunix <new-member-boot-device>
```

During this initial boot the newly added member will:

- o Configure each installed subset.
- o Attempt to build and install a new kernel. If the system cannot build a kernel, it starts a shell where you can attempt to build a kernel manually. If the build succeeds, copy the new kernel to /vmunix. When you are finished, exit the shell using ^D or 'exit'.
- o The newly added member will attempt to set boot related console variables and continue to boot to multi-user mode.
- o After the newly added member boots, you should setup your system default network interface using the appropriate system management command.

```
-----
clu_add_member on 'Tue Nov 18 09:40:48 EST 2003' end logging at atlas0
```

E.4 clu_quorum

Example E–5 shows sample output from the `clu_quorum` command for a five-node cluster. This example shows the output generated by running `clu_quorum` immediately after adding votes (see Section 8.3.2 on page 8–4) but before booting the updated nodes:

- The Running Value of `cluster_expected_votes` for updated nodes is 3.
- The Running Value of `cluster_expected_votes` for the remaining nodes is 1.

After booting the nodes, the Running Value of `cluster_expected_votes` is 3 for all nodes.

Example E–5 Sample `clu_quorum` Output

```
atlas0 # clu_quorum
Cluster Quorum Data for: atlasD0 as of Fri Nov 21 11:04:52 EST 2003

Cluster Common Quorum Data
Quorum disk:   Not Configured
File:          /etc/sysconfigtab.cluster

Attribute                                     File Value
expected votes                                     3

Member 1 Quorum Data
Host name:    atlas0                               Status:      UP
File:         /cluster/members/member1/boot_partition/etc/sysconfigtab

Attribute      Running Value      File Value
current votes      3                N/A
quorum votes      2                N/A
expected votes      3                3
node votes         1                1
qdisk votes        0                0
qdisk major        0                0
qdisk minor        0                0

Member 2 Quorum Data
Host name:    atlas1                               Status:      UP
File:         /cluster/members/member2/boot_partition/etc/sysconfigtab

Attribute      Running Value      File Value
current votes      3                N/A
quorum votes      2                N/A
expected votes      3                3
node votes         1                1
qdisk votes        0                0
qdisk major        0                0
qdisk minor        0                0
```

clu_quorum

```
Member 3 Quorum Data
Host name:      atlas2                      Status:                                UP
File:           /cluster/members/member3/boot_partition/etc/sysconfigtab

Attribute      Running Value      File Value
current votes  3                          N/A
quorum votes   2                          N/A
expected votes 3                          3
node votes     1                          1
qdisk votes    0                          0
qdisk major    0                          0
qdisk minor    0                          0

Member 4 Quorum Data
Host name:      atlas3                      Status:                                UP
File:           /cluster/members/member4/boot_partition/etc/sysconfigtab

Attribute      Running Value      File Value
current votes  3                          N/A
quorum votes   2                          N/A
expected votes 3                          3
node votes     0                          0
qdisk votes    0                          0
qdisk major    0                          0
qdisk minor    0                          0

<with similar output for each additional member>
atlas0 #
```

E.5 upgrade_check

Table E-1 summarizes the output from sample tests run using the `upgrade_check` command.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.accounting.in_place</code>	The SC domains listed under <i>Additional Information</i> that have HP AlphaServer SC-specific UNIX accounting in place.	Follow the instructions for removing HP AlphaServer SC-specific UNIX Accounting in Chapter 21 of the <i>HP AlphaServer SC Administration Guide</i> .
<code>pre-inst.adapter_mode.cant.get.srm_name</code>	Could not get the srm name of the generic management adapter from the <code>sc_networks</code> table of the database.	Use the <code>sra edit</code> command to update the database.
<code>pre-inst.adapter_mode.bad_adapter_mode</code>	The SC domains listed under <i>Additional Information</i> have the mode for their management adapter set to a value other than <code>FastFD</code> . If a node was down or <code>consvar</code> returned an error, then the test could not be performed and the node is marked with ? (for example, <code>atlas1?</code>)	Check the adapter modes use <code>scrunch -n all consvar -g (srm_name)_mode</code> where <code>srm_name</code> is the SRM name of the management adapter (for example, <code>eia0</code>).
<code>preinst.advfs_domain.root_domain</code>	The SC domains listed under <i>Additional Information</i> do not have an <code>advfs root_domain</code> domain defined for the UNIX disk.	Without this <code>advfs</code> domain, you will not be able to perform the pre-upgrade backup. You may proceed at your own risk, or follow the instructions in the <i>Tru64 UNIX Installation Guide</i> for building a UNIX disk.
<code>preinst.advfs_domain.unix_disk</code>	The SC domains listed under <i>Additional Information</i> have <code>advfs</code> domains referencing the UNIX disk. If a domain is down, the test could not be performed and its entry is marked with a ? (for example, <code>atlasD1?</code>)	Replace the link in <code>/etc/fdms/*_domain</code> with a link to the member's <code>bootdef_dev</code> .

Table E–1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.advfs_domain.usr_domain</code>	The SC domains listed under <i>Additional Information</i> do not have an <code>advfs usr_domain</code> domain defined for the UNIX disk.	Without this <code>advfs</code> domain, you will not be able to perform the pre-upgrade backup. You may proceed at your own risk, or follow the instructions in the <i>Tru64 UNIX Installation Guide</i> for building a UNIX disk.
<code>preinst.advfs_domain.var_domain</code>	The SC domains listed under <i>Additional Information</i> do not have an <code>advfs var_domain</code> domain defined for the UNIX disk.	Without this <code>advfs</code> domain, you will not be able to perform the pre-upgrade backup. You may proceed at your own risk, or follow the instructions in the <i>Tru64 UNIX Installation Guide</i> for building a UNIX disk.
<code>preinst.clu_upgrade.error</code>	Running <code>clu_upgrade check setup</code> on the SC domains listed under <i>Additional Information</i> produced errors. If a domain was not available, the test could not be run and its entry is marked with a ? (for example, <code>atlasD1?</code>)	Run <code>clu_upgrade check setup</code> manually to determine the cause of the errors.
<code>preinst.cluvote.expected_file_mismatch</code>	The expected votes (file value) is not equal to the sum of each member's node votes (file value) for the SC nodes listed under <i>Additional Information</i> .	Use <code>clu_quorum</code> to reset the members voting rights, or the cluster's expected votes.
<code>preinst.cluvote.expected_running_mismatch</code>	The expected votes (running value) is not equal to the sum of each member's node votes (running value) for the SC nodes listed under <i>Additional Information</i> .	Use <code>clu_quorum</code> to reset the members voting rights, or the cluster's expected votes.
<code>preinst.cluvote.not_run</code>	The SC domains listed under <i>Additional Information</i> were unavailable, or the <code>clu_quorum</code> command could not be run when checking for cluster member voting rights.	Bring the cluster up or remove it from the HP AlphaServer SC system. If the cluster was already up, then run the <code>clu_quorum</code> command to determine the cause of the error.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.cluvote.no_vote</code>	The SC domains listed under <i>Additional Information</i> have lead nodes without voting rights.	Use <code>clu_quorum</code> to reset the voting rights of the lead nodes.
<code>preinst.cluvote.node_vote_mismatch</code>	The SC nodes listed under <i>Additional Information</i> have a different vote value at runtime (running value) than that displayed in <code>/etc/sysconfig-tab</code> (file value).	Use <code>clu_quorum</code> to reset the members voting rights, or the cluster's expected votes.
<code>preinst.cluvote.total_file_mismatch</code>	The total node votes (file value) for the SC domains listed under <i>Additional Information</i> is different than the cluster common expected votes value.	Use <code>clu_quorum</code> to reset the members voting rights, or the cluster's expected votes.
<code>preinst.cluvote.total_running_mismatch</code>	The total node votes (running value) for the SC domains listed under <i>Additional Information</i> is different than the cluster common expected votes value.	Use <code>clu_quorum</code> to reset the members voting rights, or the cluster's expected votes.
<code>preinst.database.not_setup</code>	The database is not set up on the SC (lead) nodes listed under <i>Additional Information</i> . If a node is not available, its entry is marked with a ? (for example, <code>atlasD1?</code>)	Ensure that the HP AlphaServer SC database is set up and <code>rmshost</code> is set correctly.
<code>preinst.dbrev.incompatible</code>	The upgrade is incompatible with the current database version.	Return the database to the version originally installed with your HP AlphaServer SC system.
<code>preinst.dir.not_directory</code>	On the SC domains listed under <i>Additional Information</i> , the listed file should be a directory, but it is not. If a domain was down, then the test could not be run and its entry is marked with a ? (for example, <code>atlasD1?</code>)	If the file does not exist, then create it using <code>mkdir</code> . If a link exists in its place, then remove the link and copy the directory it points to.

Table E–1 Upgrade Check Output

Test Name	Description	Solution
preinst.disk_space.fs_capacity	The SC domains listed under <i>Additional Information</i> have important file systems over 50% full. If a domain is down, the test could not be performed and its entry is marked with a ? (for example, atlasD1?)	If you are sure there is enough space, perform the upgrade, then run <code>sc_upgrade</code> with the <code>-diskcap</code> option to override the 50% limit. Otherwise, remove unimportant files from the indicated file systems.
preinst.domain_members.unavailable	The SC nodes listed under <i>Additional Information</i> are not in multi-user mode.	Bring the nodes to multi-user state with <code>sra boot</code> or delete them from their domains with <code>sra delete_member</code> .
preinst.evm.evm_errors	There are errors in EVM on the SC domains listed under <i>Additional Information</i> . If a domain is down, the test could not be performed and its entry is marked with a ? (for example, atlasD1?)	Run the command: <code>evmget -A -f '[name *.evm] & [age < 5m]'</code> on the domains and look for errors.
preinst.ext_adapter.bad_ext_adapter	The SC nodes listed under <i>Additional Information</i> have external network adapters that are not correctly described in the SC database. If a node is not available, its entry is marked with a ? (for example, atlas1?)	Use the <code>sra edit</code> command to update the database to reflect these network adapters.
preinst.file_system.pfs_mod_loaded	The SC nodes listed under <i>Additional Information</i> have the <code>pfs.mod</code> module configured into the kernel. If a node is down, the test could not be performed and its entry is marked with a ? (for example, atlas1?)	Unload the kernel module on all nodes using the command: <code>scrunch -n all /sbin/sysconfig -u pfs</code>
preinst.file_system.pfs_online	The PFS file systems listed under <i>Additional Information</i> are online.	Use the command: <code>pfsmgr offline all</code> to set the file systems in the offline state. Then use the command: <code>scrunch -n all /sbin/sysconfig -u pfs</code> to unload the kernel module on all nodes.

Table E–1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.file_system.scfs_online</code>	The SCFS file systems listed under <i>Additional Information</i> are online.	Set the file systems in the offline state using the command: <code>scfsmgr offline all</code>
<code>preinst.generic_boot_disk.used</code>	The SC domains listed under <i>Additional Information</i> have advfs file domains that use the generic boot disk. If a domain was unavailable, the test could not be performed and its entry is marked with a ? (for example, atlasD1?)	Remove the generic boot disk from any advfs domains.
<code>preinst.general.scrun_error</code>	There was a problem running <code>scrun</code> . The current test could not be run.	Correct the problem with <code>scrun</code> and run this program again.

Table E–1 Upgrade Check Output

Test Name	Description	Solution
preinst.inetd.atypical	The SC domains listed under <i>Additional Information</i> have changed one or more of the typical entries in /etc/inetd.conf (ftp, telnet, shell, login, exec, cfmgr, and only on the management server: tftp).	Replace the modified /etc/inetd.conf with the original or replace the appropriate entry with one of the following: ftp stream tcp nowait root /usr/sbin/ftpd ftpd telnet stream tcp nowait root /usr/sbin/ telnetd telnetd shell stream tcp nowait root /usr/sbin/rshd rshd login stream tcp nowait root /usr/sbin/rlogind rlogind exec stream tcp nowait root /usr/sbin/rexecd rexecd cfmgr stream tcp nowait root /sbin/cfg- mgr cfmgr tftp dgram udp wait root /usr/sbin/tftpd tftpd /tmp /var/adm/ris /ris

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.inetd.missing	The SC domains listed under <i>Additional Information</i> are missing one or more of the typical entries in /etc/inetd.conf (ftp, telnet, shell, login, exec, cfgmgr, and only on the management server: tftp).	Replace the modified /etc/inetd.conf with the original or add one of the following as appropriate: ftp stream tcp nowait root /usr/sbin/ftpd ftpd telnet stream tcp nowait root /usr/sbin/ telnetd telnetd shell stream tcp nowait root /usr/sbin/rshd rshd login stream tcp nowait root /usr/sbin/rlogind rlogind exec stream tcp nowait root /usr/sbin/rexecd rexecd cfgmgr stream tcp nowait root /sbin/cfg- mgr cfgmgr tftp dgram udp wait root /usr/sbin/tftpd tftpd /tmp /var/adm/ris /ris
preinst.inetd.notrun	The test for typical entries in /etc/inetd.conf could not be run on the SC domains listed under <i>Additional Information</i> .	Bring the domains back online and run this program again.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.license.advfs-utilities	The lead node of each of the SC domains listed under <i>Additional Information</i> did not have the ADVFS-UTILITIES license installed. If the lead node of a domain is down, the test could not be performed and its entry is marked with a ? (for example, atlasD1?).	Install the ADVFS-UTILITIES on the listed domains.
preinst.mac_addr.bad_mac_addr	The SC nodes listed under <i>Additional Information</i> are missing from /etc/bootptab, or have MAC addresses that are different from what hwmgr returned, or the MAC address from /etc/bootptab does not agree with the address in the sc_nodes table of the database. If a node was down or hwmgr returned an error, then the test could not be performed and the node is marked with a ? (for example, atlas1?).	Check the actual MAC addresses using hwmgr and update /etc/bootptab to have the correct values. Check the database using the command: rmsquery select name, mac from sc_nodes.
preinst.osrev.wrong_rev	The SC domains listed under <i>Additional Information</i> are running a version of the operating system that is not compatible with the upgrade. If a domain was not available or the uname command produced an error, then the test could not be performed and its entry is marked with a ? (for example, atlasD1?).	Bring the domains to the recommended operating system level before upgrading HP AlphaServer SC.
preinst.ris.not_installed	No RIS server detected on the management server.	If you have not installed RIS elsewhere, then install it on the management server.

Table E–1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.rc.config_CLUA_PASSIVE.set</code>	The SC nodes listed under <i>Additional Information</i> have the <code>CLUA_PASSIVE</code> variable set in their <code>rc.config</code> file. This setting may conflict with the setting required by the V2.6 network routing. Please remove this setting before performing the upgrade. If a node was down or <code>consvar</code> returned an error, then the test could not be performed and the node is marked with a ? (for example, <code>atlas1?</code>).	Use the <code>rcmgr delete</code> command to unset the variable.
<code>preinst.rc.config_CLUAMGR_ROUTE_ARGS.set</code>	The SC nodes listed under <i>Additional Information</i> have the <code>CLUAMGR_ROUTE_ARGS</code> variable set in their <code>rc.config</code> file. This setting may conflict with the setting required by the V2.6 network routing. Please remove this setting before performing the upgrade. If a node was down or <code>consvar</code> returned an error, then the test could not be performed and the node is marked with a ? (for example, <code>atlas1?</code>).	Use the <code>rcmgr delete</code> command to unset the variable.
<code>preinst.rc.config_GATED_FLAGS.set</code>	The SC nodes listed under <i>Additional Information</i> have the <code>GATED_FLAGS</code> variable set in their <code>rc.config</code> file. This setting may conflict with the setting required by the V2.6 network routing. Please remove this setting before performing the upgrade. If a node was down or <code>consvar</code> returned an error, then the test could not be performed and the node is marked with a ? (for example, <code>atlas1?</code>).	Use the <code>rcmgr delete</code> command to unset the variable.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.rc.config_GATED.set	The SC nodes listed under <i>Additional Information</i> have the GATED variable set in their rc.config file. This setting may conflict with the setting required by the V2.6 network routing. Please remove this setting before performing the upgrade. If a node was down or consvar returned an error, then the test could not be performed and the node is marked with a ? (for example, atlas1?).	Use the rcmgr delete command to unset the variable.
preinst.ris.not_installed	No RIS server detected on the management server.	If you have not installed RIS elsewhere, then install it on the management server.
preinst.ris_rsh.bad_number	The SC nodes listed under <i>Additional Information</i> must update /etc/sysconfigtab before performing the upgrade. If a node is down, the test could not be run and its entry is marked with a ? (for example, atlas1?).	<p>Update the sysconfigtab file on the management server and lead nodes of each domain as follows:</p> <ol style="list-style-type: none">1. Create a temp file (for example, /tmp/risfix) with the following entries: inet: tcp_sendspace = 61440 tcp_recvspace = 614402. Run the command: sysconfigdb -m -f /tmp/risfix3. Run the command: sysconfig -r inet tcp_sendspace=614404. Run the command: sysconfig -r inet tcp_recvspace=614405. Reboot. <p>Repeat these steps on the management server and all lead nodes.</p>

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.routes.no_default	The SC nodes listed under <i>Additional Information</i> do not have a correct <code>/etc/routes</code> . Only nodes with external interfaces should have a default route. All other nodes should have an empty routes file. If a node was unavailable, the test could not be performed and its entry is marked with a ? (for example, atlas1?)	If the nodes require a default route, add a default route by running the following command for the appropriate nodes: <pre>sra command -node 'atlas[2-31]' -command /usr/sra/bin/SetDefaultRoute -m 1</pre> (substitute the listed nodes for atlas[2-31])
preinst.sia.grep_error	When checking <code>/var/adm/cdsl_admin.inv</code> for <code>/etc/sia</code> , either the domain was down, or an error occurred with <code>grep</code> (most likely, the file was missing).	Boot the domain(s), if necessary, and check that the file <code>/var/adm/cdsl_admin.inv</code> exists. If not, check the log for the exact error that <code>grep</code> is returning. After you have corrected the problem, run this program again.
preinst.sia.incorrect_file	The file <code>/etc/sia</code> was found to be a link, or it was not found.	Replace <code>/etc/sia</code> with a plain file.
preinst.sia.in_inventory	An entry was found for <code>/etc/sia</code> in <code>/var/adm/cdsl_admin.inv</code> .	Remove the <code>/etc/sia</code> entry in <code>/var/adm/cdsl_admin.inv</code> .
preinst.superuser.bad_shell	The SC domains listed under <i>Additional Information</i> have a default root shell other than <code>/bin/sh</code> . The upgrade and backup tools require a default shell of <code>/bin/sh</code> . Note: <code>/bin/ksh</code> may work but this is not guaranteed.	Set the root default shell in the <code>/etc/passwd</code> file on each of the listed domains to <code>/bin/sh</code> . This can be done either by using NIS or by editing the <code>/etc/passwd</code> file.
preinst.superuser.password_different	The root password for some domains may differ from others as the encrypted passwords are different. This can occur if passwords are actually different or if they have been set individually on each domain rather than sharing/replicating password files for all domains.	Check the root password on one domain in each set. If the sets have different root passwords, change the password on the appropriate domains (using NIS or by copying the <code>/etc/passwd</code> file) so that all domains are in agreement.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.superuser.password_notrun	The test for root passwords could not be run on the listed domains.	Boot the domains and run this program again.
preinst.unix_disk.bad_disk	The SC domains listed under <i>Additional Information</i> may not have a UNIX disk, or the disk may not have a /vmunix. If a domain was unavailable or if there were problems during test execution, then the test could not be performed and its entry will be marked with a ? (for example, atlasD1?).	Check that the UNIX disk is available so that the system can boot from it.

Table E-1 Upgrade Check Output

Test Name	Description	Solution
preinst.unix_disk.explanation	<p>During upgrade, it is very important to initially perform a backup. The selection of a backup device on a domain is straightforward - it would usually be the device with IDENTIFIER=3 because this would have been the disk created specifically for the role of cluster backup as described in Chapter 3 of this guide. But, it can be any disk that you choose.</p> <p>Test 20 is aimed to highlight the possibility that you might encounter problems identifying this backup device once you are booted from the UNIX disk to do a restore.</p> <p>The test mounts the root partition of the UNIX disk and reviews the /dev/disk contents. It compares this with the <code>hwmgr -v d</code> command output from the lead member of the cluster. If there is a mismatch, then the test fails indicating that there is a disk displayed, which states that the cluster which was not visible/probed the previous time, was booted from the UNIX disk.</p> <p>While the new disk will become visible when you next boot/probe from the UNIX disk, it's important to understand that the UNIX special device assigned to the new disk during that probe might be different from when the full domain probed it.</p>	

Table E-1 Upgrade Check Output

Test Name	Description	Solution
	<p>If the new disk in question is being used as the backup device, then the parameter of the <code>restore</code> command will possibly be different to the parameter for the <code>backup</code> command even though they identify the same spindle.</p> <p>During the lifetime of a domain, it is normal for a lead member of a domain to receive new SCSI disks from those on the lead member when the domain was created (while it had been booted from the UNIX disk). As such, the test is basic. However, it does serve its purpose and it highlights the fact that when you attempt a domain restore procedure, you may have to perform some research to determine the special device name of the disk containing the domain backup files.</p> <p>The fact that the test fails is not catastrophic; it means that if you need to restore, then you need to be careful. For example, if you know that you made your backups on <code>disk55</code>, then you know that this was a certain bus/target/lun combination. Later, when you boot from UNIX, the same disk might be known by a different special device, and you should check your disk selection by looking for the matching bus/target/lun combination in the <code>hwmgx -v d</code> output when booted from the UNIX disk.</p>	

Table E–1 Upgrade Check Output

Test Name	Description	Solution
<code>preinst.unix_disk.mounted</code>	The SC domains listed under <i>Additional Information</i> have mounted root_domain#root.	Unmount root_domain#root on the indicated domains.
<code>preinst.unix_disk.no_comparison_made</code>	Unable to compare the disks known to hwmgr to those listed in the devices/disk directory for the SC domains listed under <i>Additional Information</i> .	This usually means that the domains were unavailable. Reboot the domains and run this script again.
<code>preinst.unix_disk.not_unmounted</code>	Unable to unmount root_domain#root on the SC domains listed under <i>Additional Information</i> .	Unmount root_domain#root on the indicated domains.
<code>preinst.unix_disk.not_in_hwmgr</code>	The SC domains listed under <i>Additional Information</i> have disks that are listed in devices/disk on the unix_disk, but which are not known to hwmgr.	See Chapter 11: Troubleshooting of this guide for information on correcting this problem.
<code>preinst.unix_disk.not_on_unixdisk</code>	The SC domains listed under <i>Additional Information</i> have disks known to hwmgr, but they do not appear in the devices/disk directory on the unix_disk.	See Chapter 11: Troubleshooting of this guide for information on correcting this problem.
<code>preinst.volumes.multiple</code>	The listed advfs domains on the named SC domain listed under <i>Additional Information</i> are composed of multiple volumes. After the upgrade, the advfs domain will have only one volume defined.	Recreate the missing symbolic link(s) after the upgrade.
<code>preinst.volumes.unavailable</code>	The SC domains listed under <i>Additional Information</i> were unavailable when checking for multiple volumes in the advfs domains.	Reboot the domains and run this program again.
<code>preinst.webes.installed</code>	The SC domains listed under <i>Additional Information</i> have the indicated WEBES subsets installed.	Use <code>setld -d</code> to remove the required subset from the node(s).

Table E–1 Upgrade Check Output

Test Name	Description	Solution
preinst.wwl_support.subsets_installed	The SC nodes listed under <i>Additional Information</i> have World Wide Language (WWL) Support subsets installed. If a node was down or consvar returned an error, then the test could not be performed and the node is marked with ? (for example, atlas1?).	Use the setld -i grep IOS grep -v 'not installed' command to obtain a list of the WWL Support subsets installed on each node and then use the setld -d command to remove them. They can be reinstalled after the upgrade. Remember to install any WWL patches after the subsets have been reinstalled.

Cluster-Related Messages in System Log Files

The following sections show excerpts from system log files in the `/var/adm/syslog.dated/date` directories:

- Startup Messages After Creating a Domain (see Section F.1 on page F–2)
- Startup Messages After Adding a Domain Member (see Section F.2 on page F–9)

These messages track normal cluster startup operations; therefore, in addition to providing some level of assurance that cluster formation and recovery operations are proceeding in an orderly fashion, they also provide a starting point for troubleshooting cluster-related problems.

Startup Messages After Creating a Domain

F.1 Startup Messages After Creating a Domain

Example F–1 shows a transcript of a portion of the startup messages displayed during a reboot of the first cluster member system after running `sra install -node atlas0`. This information is also sent to `/var/adm/syslog.dated/date/kern.log`.

Example F–1 Startup Messages After Creating a Domain

```
18/Nov/2003 09:34:11 resetting all I/O buses
18/Nov/2003 09:34:34 (boot dka0.0.0.2.1 -flags A)
18/Nov/2003 09:34:34 block 0 of dka0.0.0.2.1 is a valid boot block
18/Nov/2003 09:34:34 reading 19 blocks from dka0.0.0.2.1
18/Nov/2003 09:34:34 bootstrap code read in
18/Nov/2003 09:34:34 base = 200000, image_start = 0, image_bytes = 2600(9728)
18/Nov/2003 09:34:34 initializing HWRPB at 2000
18/Nov/2003 09:34:34 initializing page table at 3ff56000
18/Nov/2003 09:34:34 initializing machine state
18/Nov/2003 09:34:34 setting affinity to the primary CPU
18/Nov/2003 09:34:34 jumping to bootstrap code
18/Nov/2003 09:34:36
18/Nov/2003 09:34:36 UNIX boot - Wednesday October 16, 2002
18/Nov/2003 09:34:36
18/Nov/2003 09:34:36 Loading vmunix ...
18/Nov/2003 09:34:36 Loading at 0xfffffc0000230000
18/Nov/2003 09:34:36
18/Nov/2003 09:34:36 Sizes:
18/Nov/2003 09:34:36 text = 9418368
18/Nov/2003 09:34:40 data = 2387520
18/Nov/2003 09:34:40 bss = 2564224
18/Nov/2003 09:34:40 Starting at 0xfffffc00002439c0
18/Nov/2003 09:34:40
18/Nov/2003 09:34:42 Loading vmunix symbol table ... [2319928 bytes]
18/Nov/2003 09:34:43 Memory trolling not supported, cpu Major id 13, Minor id 3
18/Nov/2003 09:34:44 Alpha boot: available memory from 0x38d4000 to 0xffff4000
18/Nov/2003 09:34:44 Compaq Tru64 UNIX V5.1B (Rev. 2650); Tue Nov 18 09:31:54 EST
2003
18/Nov/2003 09:34:44 physical memory = 4096.00 megabytes.
18/Nov/2003 09:34:44 available memory = 3968.10 megabytes.
18/Nov/2003 09:34:44 using 15658 buffers containing 122.32 megabytes of memory
18/Nov/2003 09:34:45 Master cpu at slot 0
18/Nov/2003 09:34:45 Starting secondary cpu 1
18/Nov/2003 09:34:45 Starting secondary cpu 2
18/Nov/2003 09:34:45 Starting secondary cpu 3
18/Nov/2003 09:34:45 Firmware revision: 6.4-17
18/Nov/2003 09:34:45 PALcode: UNIX version 1.91-104
18/Nov/2003 09:34:45 AlphaServer ES40
18/Nov/2003 09:34:45 pci1 (primary bus:1) at nexus
18/Nov/2003 09:34:45 pci2 (primary bus:1 subordinate bus:2) at pci1 slot 1
18/Nov/2003 09:34:45 ee0 at pci2 slot 4
18/Nov/2003 09:34:45 ee0: COMPAQ Intel 82558 (10/100 Mbps) Ethernet Interface
18/Nov/2003 09:34:45 ee0: Driver Rev = V1.0.23, Chip Rev = 5, hardware address:
00-02-A5-6B-28-7C
```

Startup Messages After Creating a Domain

```
18/Nov/2003 09:34:45 ee1 at pci2 slot 5
18/Nov/2003 09:34:45 ee1: COMPAQ Intel 82558 (10/100 Mbps) Ethernet Interface
18/Nov/2003 09:34:45 ee1: Driver Rev = V1.0.23, Chip Rev = 5, hardware address:
00-02-A5-6B-28-7D
18/Nov/2003 09:34:46 itpsa0 at pci1 slot 2
18/Nov/2003 09:34:46 IntraServer ROM Version V2.0 (c)1998
18/Nov/2003 09:34:47 ee0: Autonegotiated, 100 Mbps full duplex
18/Nov/2003 09:34:47 ee1: Autonegotiated, 100 Mbps full duplex
18/Nov/2003 09:34:49 scsi0 at itpsa0 slot 0 rad 0
18/Nov/2003 09:34:49 emx0 at pci1 slot 3
18/Nov/2003 09:34:49 KGPSA-CA : Driver Rev 2.10 : F/W Rev 3.81A4(2.01A0) : wwn
1000-0000-c922-39c4
18/Nov/2003 09:34:49 emx0: Using console topology setting of : Fabric
18/Nov/2003 09:34:50 scsi1 at emx0 slot 0 rad 0
18/Nov/2003 09:34:50 emx2 at pci1 slot 4
18/Nov/2003 09:34:50 KGPSA-CA : Driver Rev 2.10 : F/W Rev 3.81A4(2.01A0) : wwn
1000-0000-c923-e69c
18/Nov/2003 09:34:50 emx2: Using console topology setting of : Fabric
18/Nov/2003 09:34:50 scsi2 at emx2 slot 0 rad 0
18/Nov/2003 09:34:51 pci0 (primary bus:0) at nexus
18/Nov/2003 09:34:51 elan0: Device revision = 1 (Rev B)
18/Nov/2003 09:34:51 elan0: QSW Elan3 PCI Network Adaptor
18/Nov/2003 09:34:51 elan0: Serial Number - 380
18/Nov/2003 09:34:51 elan0: memory bank 0 is 32768K
18/Nov/2003 09:34:56 elan0: memory bank 1 is 32768K
18/Nov/2003 09:35:05 elan30 at pci0 slot 3
18/Nov/2003 09:35:05 isa0 at pci0
18/Nov/2003 09:35:05 gpc0 at isa0
18/Nov/2003 09:35:05 gpc1 not probed
18/Nov/2003 09:35:05 ace0 at isa0
18/Nov/2003 09:35:05 ace1 at isa0
18/Nov/2003 09:35:05 jtag0 at isa0
18/Nov/2003 09:35:05 fdi0 at isa0
18/Nov/2003 09:35:05 fd0 at fdi0 unit 0
18/Nov/2003 09:35:05 ata0 at pci0 slot 15
18/Nov/2003 09:35:05 ata0: ACER M1543C
18/Nov/2003 09:35:07 scsi3 at ata0 slot 0 rad 0
18/Nov/2003 09:35:07 scsi4 at ata0 slot 1 rad 0
18/Nov/2003 09:35:07 usb0 at pci0 slot 19
18/Nov/2003 09:35:07 Created FRU table binary error log packet
18/Nov/2003 09:35:07 kernel console: ace0
18/Nov/2003 09:35:07 dli: configured
18/Nov/2003 09:35:07 NetRAIN configured.
18/Nov/2003 09:35:07 Random number generator configured.
18/Nov/2003 09:35:08 ATM Subsystem configured with 4 restart threads
18/Nov/2003 09:35:08 ATMUNI: configured
18/Nov/2003 09:35:08 ATMSIG: 3.x (module=uni3x) configured
18/Nov/2003 09:35:08 ILMI: 3.x (module=ilmi) configured
18/Nov/2003 09:35:08 ATM IP: configured
18/Nov/2003 09:35:08 ATM LANE: configured.
18/Nov/2003 09:35:08 ATM IFMP: configured
18/Nov/2003 09:35:08 elan0: nodeid=128 level=5 numnodes=512
```

Startup Messages After Creating a Domain

```
18/Nov/2003 09:35:08 TruCluster Server V5.1-SC-V2.6-BL5-3320-EAGLE (Rev. 13320);
11/16/03 01:44
18/Nov/2003 09:35:08 elan0: waiting for network position to be found
18/Nov/2003 09:35:08 elan0: nodeId=128 level=3 numnodes=128
18/Nov/2003 09:35:08 elan0: Online [0]
18/Nov/2003 09:35:08 elan0: Nodeset [0]
18/Nov/2003 09:35:08 elan0: network position found at nodeId 0
18/Nov/2003 09:35:08 eip: checksums disabled - all other nodes must also disable
checksums
18/Nov/2003 09:35:08 eip: Elan IP initialized
18/Nov/2003 09:35:08 TNC kproc_creator_daemon: Initialized and Ready
18/Nov/2003 09:35:08 clubase: configured
18/Nov/2003 09:35:08 icsnet: configured
18/Nov/2003 09:35:08 drd configured 0
18/Nov/2003 09:35:08 kch: configured
18/Nov/2003 09:35:08 dlm: configured
18/Nov/2003 09:35:08 Starting CFS daemons
18/Nov/2003 09:35:08 Registering CFS Services
18/Nov/2003 09:35:08 Initializing CFSREC ICS Service
18/Nov/2003 09:35:08 Registering CFSMSFS remote syscall interface
18/Nov/2003 09:35:09 Registering CMS Services
18/Nov/2003 09:35:09 ep_enable_cluster_protocols: no longer implemented
18/Nov/2003 09:35:09 ics_elan: seticsinfo: [elan node 0] <=> [ics node 1]
18/Nov/2003 09:35:19 CNX MGR: Cluster atlasD0 incarnation 0x8ae4 has been formed
18/Nov/2003 09:35:19 CNX MGR: Founding node id is 1 csid is 0x10001
18/Nov/2003 09:35:19 CNX MGR: membership configuration index: 1 (1 additions, 0
removals)
18/Nov/2003 09:35:19 CNX MGR: quorum (re)gained, (re)starting cluster operations.
18/Nov/2003 09:35:19 Joining versw kch set.
18/Nov/2003 09:35:19 CNX MGR: Node atlas0 1 incarn 0x8ae4 csid 0x10001 has been
added to the cluster
18/Nov/2003 09:35:19 dlm: resuming lock activity
18/Nov/2003 09:35:19 kch: resuming activity
18/Nov/2003 09:35:19 clsm: checking for peer configurations
18/Nov/2003 09:35:19 clsm: initialized
18/Nov/2003 09:35:19 Waiting for cluster mount to complete
18/Nov/2003 09:35:20 scsi0: SCSI Bus was reset
18/Nov/2003 09:35:20 cam_logger: SCSI event packet
18/Nov/2003 09:35:20 cam_logger: bus 0
18/Nov/2003 09:35:20 itpsa SCSI HBA
18/Nov/2003 09:35:20 SCSI Bus was reset
18/Nov/2003 09:35:20
18/Nov/2003 09:35:20 elan0: Online Ack from [1-3]
18/Nov/2003 09:35:21 SCFS: SCFS/Elan Services Initialised on node 0(0x0) railmask
0x1
18/Nov/2003 09:35:21 vm_swap_init: swap is set to eager allocation mode
18/Nov/2003 09:35:22 CMS: Joining deferred filesystem sets
18/Nov/2003 09:35:22 Checking device naming:
18/Nov/2003 09:35:22 Checking local filesystems
18/Nov/2003 09:35:22 Mounting / (root)
18/Nov/2003 09:35:23 user_cfg_pt: reconfigured
18/Nov/2003 09:35:23 root_mounted_rw: reconfigured
```


Startup Messages After Creating a Domain

```
18/Nov/2003 09:35:23 Mounting /cluster/members/member1/boot_partition (boot
filesystem)
18/Nov/2003 09:35:23 user_cfg_pt: reconfigured
18/Nov/2003 09:35:23 root_mounted_rw: reconfigured
18/Nov/2003 09:35:23 Device Naming: first boot initialization . . .
18/Nov/2003 09:35:26 elan0: Online Ack from [4-15]
18/Nov/2003 09:35:27 ptm
18/Nov/2003 09:35:27 user_cfg_pt: reconfigured
18/Nov/2003 09:35:28 Mounting local filesystems
18/Nov/2003 09:35:28 exec: /sbin/mount_advfs -F 0x14000 cluster_root#root /
18/Nov/2003 09:35:28 cluster_root#root on / type advfs (rw)
18/Nov/2003 09:35:28 exec: /sbin/mount_advfs -F 0x4000 cluster_usr#usr /usr
18/Nov/2003 09:35:28 cluster_usr#usr on /usr type advfs (rw)
18/Nov/2003 09:35:28 exec: /sbin/mount_advfs -F 0x4000 cluster_var#var /var
18/Nov/2003 09:35:29 cluster_var#var on /var type advfs (rw)
18/Nov/2003 09:35:29 /proc on /proc type procfs (rw)
18/Nov/2003 09:35:29 Nov 18 09:35:27 esmd: Essential Services Monitor daemon
started
18/Nov/2003 09:35:29 Subsystem hwautoconfig was successfully configured.
18/Nov/2003 09:35:29 Subsystem shmem was successfully configured.
18/Nov/2003 09:35:29 Subsystem pfs was successfully configured.
18/Nov/2003 09:35:30 Environmental Monitoring Subsystem Configured.
18/Nov/2003 09:35:30 Subsystem envmon was successfully configured.
18/Nov/2003 09:35:30 Nov 18 09:35:28 update: started
18/Nov/2003 09:35:31
18/Nov/2003 09:35:31 Checking for Installation Tasks...
18/Nov/2003 09:35:33
18/Nov/2003 09:35:33 Executing Installation Tasks...
18/Nov/2003 09:35:34 Clusterizing NIS server configuration
18/Nov/2003 09:35:35 Clusterizing mail...
18/Nov/2003 09:35:35 Removing /var/adm/sendmail/Makefile.cf.atlas0
18/Nov/2003 09:35:35 Renaming /var/adm/sendmail/atlas0.m4 to /var/adm/sendmail/
atlasD0.m4
18/Nov/2003 09:35:35 Renaming /var/adm/sendmail/atlas0.cf to /var/adm/sendmail/
atlasD0.cf
18/Nov/2003 09:35:35 Saving original /var/adm/sendmail/sendmail.cf as /var/adm/
sendmail/sendmail.cf.cluster.sav
18/Nov/2003 09:35:35 Saving original /var/adm/sendmail/atlasD0.m4 as /var/adm/
sendmail/atlasD0.m4.cluster.sav
18/Nov/2003 09:35:36 Changes to mail configuration complete
18/Nov/2003 09:35:37 set boot_reset = off
18/Nov/2003 09:35:38
18/Nov/2003 09:35:38 *** KERNEL CONFIGURATION AND BUILD PROCEDURE ***
18/Nov/2003 09:35:38
18/Nov/2003 09:35:38 Saving /sys/conf/GENERIC as /sys/conf/GENERIC.bck
18/Nov/2003 09:35:38
18/Nov/2003 09:35:38 *** PERFORMING KERNEL BUILD ***
18/Nov/2003 09:35:50 Working....Tue Nov 18 09:35:47 EST 2003
18/Nov/2003 09:36:37
18/Nov/2003 09:36:37 The new kernel is /sys/GENERIC/vmunix
18/Nov/2003 09:36:42 Primary Boot Disk: dsk0
18/Nov/2003 09:36:42 Alternate Boot Disk: dsk1
```

Startup Messages After Creating a Domain

```
18/Nov/2003 09:36:46 Using /var/adm/sc.binary.errlog as the common log Replacing
existing /etc/binlog.conf entries ... Removing existing /etc/binlog.conf ...
18/Nov/2003 09:36:47 The system is coming up. Please wait...
18/Nov/2003 09:36:47 Checking for crash dumps
18/Nov/2003 09:36:47 Initializing paging space
18/Nov/2003 09:36:47 Mounting Memory filesystems
18/Nov/2003 09:36:51 evmstart: Daemon started
18/Nov/2003 09:36:51 Nov 18 09:36:48 esmd: Started monitoring the EVM daemon
18/Nov/2003 09:36:51 security configuration set to default (BASE).
18/Nov/2003 09:36:51 File /etc/sia/matrix.conf updated successfully.
18/Nov/2003 09:36:51 Successful SIA initialization
18/Nov/2003 09:36:51
18/Nov/2003 09:36:52 /usr/sbin/autopush: Can't push requested modules on STREAM
for entry 36
18/Nov/2003 09:36:52 /usr/sbin/autopush: Device (6,-1) already configured
18/Nov/2003 09:36:52 Streams autopushes configured
18/Nov/2003 09:36:52 Initializing random number driver
18/Nov/2003 09:36:53 CSSM_ModuleLoad: CSSM error 4107
18/Nov/2003 09:36:55 NIFF daemon started
18/Nov/2003 09:36:55 Configuring network
18/Nov/2003 09:36:55 hostname: atlas0
18/Nov/2003 09:36:56 Mounted root1_tmp#tmp on /tmp
18/Nov/2003 09:36:57 Mounted root1_local#local on /local
18/Nov/2003 09:36:58 Mounting advfs member file systems
18/Nov/2003 09:36:58 Loading IMF licenses
18/Nov/2003 09:36:59 System error logger started
18/Nov/2003 09:36:59 add net default: gateway 11.0.24.253
18/Nov/2003 09:37:01 gateway daemon started
18/Nov/2003 09:37:01 Setting kernel timezone variable
18/Nov/2003 09:37:01 Setting the current time and date with ntpdate
18/Nov/2003 09:37:01 Tue Nov 18 09:37:00 EST 2003
18/Nov/2003 09:37:01 Ntpdate succeeded.
18/Nov/2003 09:37:01 starting cluster alias
18/Nov/2003 09:37:01 cluster alias subsystem enabled
18/Nov/2003 09:37:01 enable: reconfigured
18/Nov/2003 09:37:02 aliasd: setting up NIFF monitor for interface ee0
18/Nov/2003 09:37:02 aliasd: setting up NIFF monitor for interface ee1
18/Nov/2003 09:37:02 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:02 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:04 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:04 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:04 gateway daemon started
18/Nov/2003 09:37:07 ONC portmap service started
18/Nov/2003 09:37:07 /usr/sbin/cluster/caa_rollDatastore check called
18/Nov/2003 09:37:08 CAA daemon started
18/Nov/2003 09:37:08 Starting CAA application registration
18/Nov/2003 09:37:09 Global Exec: started gxmgmtd.
18/Nov/2003 09:37:10 Global Exec: started gxclusterd.
18/Nov/2003 09:37:10 Global Exec: started gxnoded.
18/Nov/2003 09:37:12 CMF: console logging daemon does not run on this host
18/Nov/2003 09:37:12 Starting CAA registration of SC15srad
18/Nov/2003 09:37:12 management netmask = 255.255.0.0
18/Nov/2003 09:37:12 management network device = ee0
```

Startup Messages After Creating a Domain

```
18/Nov/2003 09:37:13 timezone = EST
18/Nov/2003 09:37:13 domain = eng.lkg.dec.com
18/Nov/2003 09:37:14 nameserver = 11.0.0.10
18/Nov/2003 09:37:16 starting cfd
18/Nov/2003 09:37:23 gated seems to be hung; killing it
18/Nov/2003 09:37:25 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:25 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:27 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:27 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:29 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:29 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:31 aliasd: skipping NIFF monitor for interface eip0
18/Nov/2003 09:37:33 CAA Profile for srad has been updated
18/Nov/2003 09:37:33 Attempting to start `SC15srad` on member `atlas0`
18/Nov/2003 09:37:33 Start of `SC15srad` on member `atlas0` succeeded.
18/Nov/2003 09:37:34 NIS domain name set to sc
18/Nov/2003 09:37:35 ypserv daemon started
18/Nov/2003 09:37:35 ypbind daemon started
18/Nov/2003 09:37:35 ypbind: Secure mode sunos 3.x servers rejected.
18/Nov/2003 09:37:35 NFS mount daemon started
18/Nov/2003 09:37:35 NFS export service started
18/Nov/2003 09:37:35 Attempting to start `cluster_lockd` on member `atlas0`
18/Nov/2003 09:37:35 cluster NFS Locking:
18/Nov/2003 09:37:35         cluster rpc.statd started
18/Nov/2003 09:37:36         cluster rpc.lockd started
18/Nov/2003 09:37:39 Start of `cluster_lockd` on member `atlas0` succeeded.
18/Nov/2003 09:37:39 NFS IO service started
18/Nov/2003 09:37:39 NFS Locking:
18/Nov/2003 09:37:39     rpc.statd started
18/Nov/2003 09:37:39     rpc.lockd started
18/Nov/2003 09:37:40 Mounting nfs member file systems
18/Nov/2003 09:37:40 scfe0-ext1:/usr/users on /usr/users type nfs (rw,bg)
18/Nov/2003 09:37:40 scfe0-ext1:/kits on /kits type nfs (rw,bg)
18/Nov/2003 09:37:40 atlasms:/var/sra/diag/quadrics on /var/sra/diag/quadrics type
nfs (rw,soft,bg,intr)
18/Nov/2003 09:37:40 Mounting nfsv3 member file systems
18/Nov/2003 09:37:40 NOTE:
18/Nov/2003 09:37:40     Skipping NFS file system entries found in /etc/fstab.
18/Nov/2003 09:37:40
18/Nov/2003 09:37:40     Please migrate the NFS file system entries to the member
18/Nov/2003 09:37:41     specific /etc/member_fstab of the member(s) which should
be
18/Nov/2003 09:37:41     mounting these NFS file systems.
18/Nov/2003 09:37:41
18/Nov/2003 09:37:41 Preserving editor files
18/Nov/2003 09:37:41 Clearing temporary files
18/Nov/2003 09:37:41 Unlocking ptys
18/Nov/2003 09:37:46 Secure Shell daemon (sshd2) started.
18/Nov/2003 09:37:46 SMTP Mail Service started.
18/Nov/2003 09:37:47 Network Time Service started
18/Nov/2003 09:37:47 The SNMP trap to Event Manager interface is disabled.
18/Nov/2003 09:37:47 GS Platform View and Discovery V1.3 for Insight Manager is
only supported on Alpha GS series platforms.
```

Startup Messages After Creating a Domain

```
18/Nov/2003 09:37:47 Internet services provided.
18/Nov/2003 09:37:48 clustercron entry started
18/Nov/2003 09:37:48 Cron service started
18/Nov/2003 09:37:52 rms: RMS service started on atlas0
18/Nov/2003 09:37:52 Binary error logger started
18/Nov/2003 09:37:53 CFS load monitoring (cfsd) started
18/Nov/2003 09:37:53 CAA Applications now started
18/Nov/2003 09:37:53 cluster wall daemon started
18/Nov/2003 09:37:54 Starting Cluster Configuration Check...
18/Nov/2003 09:37:56
18/Nov/2003 09:37:56 The boottime cluster check found a potential problem.
18/Nov/2003 09:37:56 For details search for !!!!!ATTENTION!!!!!! in /cluster/admin/
clu_check_log_atlas0
18/Nov/2003 09:37:56 check_cdsl_config : Boot Mode : Running /usr/sbin/cdslinrchk
in the background
18/Nov/2003 09:37:56 check_cdsl_config : Results can be found in : /var/adm/
cdsl_check_list
18/Nov/2003 09:37:57 clu_check_config : detected one or more configuration errors
18/Nov/2003 09:37:59 Enabling SCFS Serving on node atlas0.
18/Nov/2003 09:37:59 scmon: daemon started
18/Nov/2003 09:37:59 scmountd: not necessary to run on this node
18/Nov/2003 09:37:59 The system is ready.
18/Nov/2003 09:38:00
18/Nov/2003 09:38:00
18/Nov/2003 09:38:00 Compaq Tru64 UNIX V5.1B (Rev. 2650) (atlas0) console
18/Nov/2003 09:38:00
18/Nov/2003 09:38:47 login:
```

F.2 Startup Messages After Adding a Domain Member

Example F–2 shows a transcript of a portion of the startup messages displayed during a boot of the second cluster member system after running `sra install -node atlas1`. This information is also sent to `/var/adm/syslog.dated/date/kern.log`.

Example F–2 Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:28:05 P00>>>boot dka0 -file vmunix
19/Nov/2003 13:28:05 (boot dka0.0.0.2.1 -file vmunix -flags A)
19/Nov/2003 13:28:05 block 0 of dka0.0.0.2.1 is a valid boot block
19/Nov/2003 13:28:05 reading 19 blocks from dka0.0.0.2.1
19/Nov/2003 13:28:05 bootstrap code read in
19/Nov/2003 13:28:05 base = 200000, image_start = 0, image_bytes = 2600(9728)
19/Nov/2003 13:28:05 initializing HWRPB at 2000
19/Nov/2003 13:28:05 initializing page table at 3ff56000
19/Nov/2003 13:28:05 initializing machine state
19/Nov/2003 13:28:05 setting affinity to the primary CPU
19/Nov/2003 13:28:05 jumping to bootstrap code
19/Nov/2003 13:28:07
19/Nov/2003 13:28:07 UNIX boot - Wednesday October 16, 2002
19/Nov/2003 13:28:07
19/Nov/2003 13:28:07 Loading vmunix ...
19/Nov/2003 13:28:08 Loading at 0xfffffc0000230000
19/Nov/2003 13:28:08
19/Nov/2003 13:28:08 Sizes:
19/Nov/2003 13:28:08 text = 9403584
19/Nov/2003 13:28:11 data = 2383232
19/Nov/2003 13:28:12 bss = 2559232
19/Nov/2003 13:28:12 Starting at 0xfffffc00002439c0
19/Nov/2003 13:28:12
19/Nov/2003 13:28:13 Loading vmunix symbol table ... [2317024 bytes]
19/Nov/2003 13:28:14 Memory trolling not supported, cpu Major id 13, Minor id 3
19/Nov/2003 13:28:22 Alpha boot: available memory from 0x38cc000 to 0xffff4000
19/Nov/2003 13:28:22 Compaq Tru64 UNIX V5.1B (Rev. 2650); Tue Nov 18 10:12:23 EST
2003
19/Nov/2003 13:28:22 physical memory = 4096.00 megabytes.
19/Nov/2003 13:28:22 available memory = 3968.13 megabytes.
19/Nov/2003 13:28:22 using 15658 buffers containing 122.32 megabytes of memory
19/Nov/2003 13:28:23 Master cpu at slot 0
19/Nov/2003 13:28:23 Starting secondary cpu 1
19/Nov/2003 13:28:23 Starting secondary cpu 2
19/Nov/2003 13:28:23 Starting secondary cpu 3
19/Nov/2003 13:28:23 Firmware revision: 6.4-17
19/Nov/2003 13:28:23 PALcode: UNIX version 1.91-104
19/Nov/2003 13:28:23 AlphaServer ES40
19/Nov/2003 13:28:23 pci1 (primary bus:1) at nexus
19/Nov/2003 13:28:23 pci2 (primary bus:1 subordinate bus:2) at pci1 slot 1
19/Nov/2003 13:28:23 ee0 at pci2 slot 4
19/Nov/2003 13:28:23 ee0: COMPAQ Intel 82558 (10/100 Mbps) Ethernet Interface
19/Nov/2003 13:28:23 ee0: Driver Rev = V1.0.23, Chip Rev = 5, hardware address:
00-02-A5-6B-22-AC
```

Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:28:23 ee1 at pci2 slot 5
19/Nov/2003 13:28:24 ee1: COMPAQ Intel 82558 (10/100 Mbps) Ethernet Interface
19/Nov/2003 13:28:24 ee1: Driver Rev = V1.0.23, Chip Rev = 5, hardware address:
00-02-A5-6B-22-AD
19/Nov/2003 13:28:24 itpsa0 at pci1 slot 2
19/Nov/2003 13:28:24 IntraServer ROM Version V2.0 (c)1998
19/Nov/2003 13:28:25 ee0: Autonegotiated, 100 Mbps full duplex
19/Nov/2003 13:28:25 ee1: Autonegotiated, 100 Mbps full duplex
19/Nov/2003 13:28:28 scsi0 at itpsa0 slot 0 rad 0
19/Nov/2003 13:28:28 emx0 at pci1 slot 3
19/Nov/2003 13:28:28 KGPSA-CA : Driver Rev 2.10 : F/W Rev 3.81A4(2.01A0) : wwn
1000-0000-c922-4ad5
19/Nov/2003 13:28:28 emx0: Using console topology setting of : Fabric
19/Nov/2003 13:28:28 scsi1 at emx0 slot 0 rad 0
19/Nov/2003 13:28:28 emx2 at pci1 slot 4
19/Nov/2003 13:28:28 KGPSA-CA : Driver Rev 2.10 : F/W Rev 3.81A4(2.01A0) : wwn
1000-0000-c922-4627
19/Nov/2003 13:28:28 emx2: Using console topology setting of : Fabric
19/Nov/2003 13:28:28 scsi2 at emx2 slot 0 rad 0
19/Nov/2003 13:28:29 pci0 (primary bus:0) at nexus
19/Nov/2003 13:28:29 elan0: Device revision = 1 (Rev B)
19/Nov/2003 13:28:29 elan0: QSW Elan3 PCI Network Adaptor
19/Nov/2003 13:28:29 elan0: Serial Number - 438
19/Nov/2003 13:28:29 elan0: memory bank 0 is 32768K
19/Nov/2003 13:28:34 elan0: memory bank 1 is 32768K
19/Nov/2003 13:28:43 elan30 at pci0 slot 3
19/Nov/2003 13:28:43 isa0 at pci0
19/Nov/2003 13:28:43 gpc0 at isa0
19/Nov/2003 13:28:43 gpc1 not probed
19/Nov/2003 13:28:43 ace0 at isa0
19/Nov/2003 13:28:43 ace1 at isa0
19/Nov/2003 13:28:43 jtag0 at isa0
19/Nov/2003 13:28:43 fdi0 at isa0
19/Nov/2003 13:28:43 fd0 at fdi0 unit 0
19/Nov/2003 13:28:43 ata0 at pci0 slot 15
19/Nov/2003 13:28:43 ata0: ACER M1543C
19/Nov/2003 13:28:45 scsi3 at ata0 slot 0 rad 0
19/Nov/2003 13:28:45 scsi4 at ata0 slot 1 rad 0
19/Nov/2003 13:28:45 usb0 at pci0 slot 19
19/Nov/2003 13:28:45 Created FRU table binary error log packet
19/Nov/2003 13:28:45 kernel console: ace0
19/Nov/2003 13:28:45 dli: configured
19/Nov/2003 13:28:45 NetRAIN configured.
19/Nov/2003 13:28:46 Random number generator configured.
19/Nov/2003 13:28:46 ATM Subsystem configured with 4 restart threads
19/Nov/2003 13:28:46 ATMUNI: configured
19/Nov/2003 13:28:46 ATMSIG: 3.x (module=uni3x) configured
19/Nov/2003 13:28:46 ILMI: 3.x (module=ilmi) configured
19/Nov/2003 13:28:46 ATM IP: configured
19/Nov/2003 13:28:46 ATM LANE: configured.
19/Nov/2003 13:28:46 ATM IFMP: configured
19/Nov/2003 13:28:46 elan0: nodeid=1 level=3 numnodes=128
```

Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:28:46 TruCluster Server V5.1-SC-V2.6-BL5-3320-EAGLE (Rev. 13320);
11/16/03 01:44
19/Nov/2003 13:28:46 elan0: waiting for network position to be found
19/Nov/2003 13:28:46 elan0: nodeid=1 level=3 numnodes=128
19/Nov/2003 13:28:46 elan0: Online [1]
19/Nov/2003 13:28:46 elan0: Nodeset [1]
19/Nov/2003 13:28:46 elan0: network position found at nodeid 1
19/Nov/2003 13:28:46 eip: checksums disabled - all other nodes must also disable
checksums
19/Nov/2003 13:28:46 eip: Elan IP initialized
19/Nov/2003 13:28:46 TNC kproc_creator_daemon: Initialized and Ready
19/Nov/2003 13:28:46 clubase: configured
19/Nov/2003 13:28:46 icsnet: configured
19/Nov/2003 13:28:46 drd configured 0
19/Nov/2003 13:28:46 kch: configured
19/Nov/2003 13:28:47 dlm: configured
19/Nov/2003 13:28:47 Starting CFS daemons
19/Nov/2003 13:28:47 Registering CFS Services
19/Nov/2003 13:28:47 Initializing CFSREC ICS Service
19/Nov/2003 13:28:47 Registering CFSMSFS remote syscall interface
19/Nov/2003 13:28:47 Registering CMS Services
19/Nov/2003 13:28:47 elan0: Restart Request from [0-0] [2-3]
19/Nov/2003 13:28:47 ep_enable_cluster_protocols: no longer implemented
19/Nov/2003 13:28:47 ics_elan: seticsinfo: [elan node 1] <=> [ics node 2]
21/Nov/2003 16:43:29 elan0: Online Ack from [0-0] [2-3]
21/Nov/2003 16:43:29 elan0: Online [0,2-3]
21/Nov/2003 16:43:29 elan0: Nodeset [0-1]
19/Nov/2003 13:28:49 ics_elan: seticsinfo: [elan node 128] <=> [ics node 1]
19/Nov/2003 13:28:49 CNX MGR: Join operation complete
19/Nov/2003 13:28:49 CNX MGR: membership configuration index: 64 (34 additions, 30
removals)
19/Nov/2003 13:28:49 CNX MGR: quorum (re)gained, (re)starting cluster operations.
19/Nov/2003 13:28:49 Joining versw kch set.
19/Nov/2003 13:28:49 CNX MGR: Node atlas0 1 incarn 0x8ae4 csid 0x10001 has been
added to the cluster
19/Nov/2003 13:28:50 CNX MGR: Node atlas1 2 incarn 0xf385e csid 0x80004 has been
added to the cluster
19/Nov/2003 13:28:50 dlm: resuming lock activity
19/Nov/2003 13:28:50 kch: resuming activity
19/Nov/2003 13:28:50 scsi0: SCSI Bus was reset
19/Nov/2003 13:28:50 cam_logger: SCSI event packet
19/Nov/2003 13:28:50 cam_logger: bus 0
19/Nov/2003 13:28:50 itpsa SCSI HBA
19/Nov/2003 13:28:50 SCSI Bus was reset
19/Nov/2003 13:28:50
19/Nov/2003 13:28:51 scsi0: SCSI Bus was reset
19/Nov/2003 13:28:51 cam_logger: SCSI event packet
19/Nov/2003 13:28:51 cam_logger: bus 0
19/Nov/2003 13:28:51 itpsa SCSI HBA
19/Nov/2003 13:28:51 SCSI Bus was reset
19/Nov/2003 13:28:51
19/Nov/2003 13:28:51 clsm: incoming CNX data: 'a'
19/Nov/2003 13:28:51 clsm: checking for peer configurations
```

Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:28:51 clsm: initialized
19/Nov/2003 13:28:52 Waiting for cluster mount to complete
19/Nov/2003 13:28:53 SCFS: SCFS/Elan Services Initialised on node 1(0x1) railmask
0x1
19/Nov/2003 13:28:54 vm_swap_init: swap is set to eager allocation mode
19/Nov/2003 13:28:55 CMS: Joining deferred filesystem sets
19/Nov/2003 13:28:55 Checking device naming:
19/Nov/2003 13:28:56     Passed.
19/Nov/2003 13:28:56 Checking local filesystems
19/Nov/2003 13:28:56 Mounting / (root)
19/Nov/2003 13:28:56 user_cfg_pt: reconfigured
19/Nov/2003 13:28:57 root_mounted_rw: reconfigured
19/Nov/2003 13:28:57 Mounting /cluster/members/member2/boot_partition (boot
filesystem)
19/Nov/2003 13:28:57 user_cfg_pt: reconfigured
19/Nov/2003 13:28:57 root_mounted_rw: reconfigured
19/Nov/2003 13:28:57 user_cfg_pt: reconfigured
19/Nov/2003 13:28:58 dsfmgr: NOTE: updating kernel basenames for system at /
19/Nov/2003 13:28:58     scp kevmm tty00 tty01 random urandom dsk4 dsk5 dsk6 scp0
scpl scp2 dsk26 dsk27 dsk28 floppy2 cdrom2
19/Nov/2003 13:28:59 Mounting local filesystems
19/Nov/2003 13:28:59 exec: /sbin/mount_advfs -F 0x14000 cluster_root#root /
19/Nov/2003 13:28:59 cluster_root#root on / type advfs (rw)
19/Nov/2003 13:28:59 exec: /sbin/mount_advfs -F 0x4000 cluster_usr#usr /usr
19/Nov/2003 13:28:59 cluster_usr#usr on /usr: Device busy
19/Nov/2003 13:28:59 exec: /sbin/mount_advfs -F 0x4000 cluster_var#var /var
19/Nov/2003 13:28:59 cluster_var#var on /var: Device busy
19/Nov/2003 13:28:59 /proc on /proc type procfs (rw)
19/Nov/2003 13:29:00 Nov 19 13:28:59 esmd: Essential Services Monitor daemon
started
19/Nov/2003 13:29:00 Subsystem hwautoconfig was successfully configured.
19/Nov/2003 13:29:00 Subsystem pfs was successfully configured.
19/Nov/2003 13:29:00 Subsystem shmem was successfully configured.
19/Nov/2003 13:29:00 Environmental Monitoring Subsystem Configured.
19/Nov/2003 13:29:00 Subsystem envmon was successfully configured.
19/Nov/2003 13:29:00 Nov 19 13:29:00 update: started
19/Nov/2003 13:29:02
19/Nov/2003 13:29:02 Checking for Installation Tasks...
19/Nov/2003 13:29:03
19/Nov/2003 13:29:03 Executing Installation Tasks...
19/Nov/2003 13:29:04 The system is coming up. Please wait...
19/Nov/2003 13:29:04 Checking for crash dumps
19/Nov/2003 13:29:04 Initializing paging space
19/Nov/2003 13:29:05 Mounting Memory filesystems
19/Nov/2003 13:29:08 evmstart: Daemon started
19/Nov/2003 13:29:09 Nov 19 13:29:08 esmd: Started monitoring the EVM daemon
19/Nov/2003 13:29:09 security configuration set to default (BASE).
19/Nov/2003 13:29:09 File /etc/sia/matrix.conf updated successfully.
19/Nov/2003 13:29:09 Successful SIA initialization
19/Nov/2003 13:29:09
19/Nov/2003 13:29:09 /usr/sbin/autopush: Can't push requested modules on STREAM
for entry 36
19/Nov/2003 13:29:09 /usr/sbin/autopush: Device (6,-1) already configured
```


Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:29:09 Streams autopushes configured
19/Nov/2003 13:29:10 CSSM_ModuleLoad: CSSM error 4107
19/Nov/2003 13:29:12 NIFF daemon started
19/Nov/2003 13:29:12 Configuring network
19/Nov/2003 13:29:12 hostname: atlas1
19/Nov/2003 13:29:16 Mounted root2_tmp#tmp on /tmp
19/Nov/2003 13:29:17 Mounted root2_local#local on /local
19/Nov/2003 13:29:18 Mounting advfs member file systems
19/Nov/2003 13:29:18 Loading LMF licenses
19/Nov/2003 13:29:20 System error logger started
19/Nov/2003 13:29:20 add net default: gateway 11.0.24.253
19/Nov/2003 13:29:20 gateway daemon started
19/Nov/2003 13:29:20 Setting kernel timezone variable
19/Nov/2003 13:29:20
19/Nov/2003 13:29:20 Setting the current time and date with ntpdate
19/Nov/2003 13:29:25 Wed Nov 19 13:29:25 EST 2003
19/Nov/2003 13:29:25 Ntpdate succeeded.
19/Nov/2003 13:29:25
19/Nov/2003 13:29:25 starting cluster alias
19/Nov/2003 13:29:25 cluster alias subsystem enabled
19/Nov/2003 13:29:25 enable: reconfigured
19/Nov/2003 13:29:27 aliasd: setting up NIFF monitor for interface ee0
19/Nov/2003 13:29:27 aliasd: setting up NIFF monitor for interface ee1
19/Nov/2003 13:29:27 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:27 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:29 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:29 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:29 gateway daemon started
19/Nov/2003 13:29:36
19/Nov/2003 13:29:36 ONC portmap service started
19/Nov/2003 13:29:36 /usr/sbin/cluster/caa_rollDatastore check called
19/Nov/2003 13:29:37 CAA daemon started
19/Nov/2003 13:29:38 Nov 19 13:29:38 atlas1 vmunix: arp: local IP address
10.128.106.2 in use by hardware address 00-02-A5-6B-28-7C
19/Nov/2003 13:29:38 Global Exec: started gxmgmtd.
19/Nov/2003 13:29:39 Global Exec: started gxclusterd.
19/Nov/2003 13:29:39 Global Exec: started gxnoded.
19/Nov/2003 13:29:42 CMF: console logging daemon does not run on this host
19/Nov/2003 13:29:48 gated seems to be hung; killing it
19/Nov/2003 13:29:50 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:50 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:52 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:53 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:54 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:54 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:29:56 aliasd: skipping NIFF monitor for interface eip0
19/Nov/2003 13:30:05
19/Nov/2003 13:30:35
19/Nov/2003 13:30:50 CAA Profile for srاد has been updated
19/Nov/2003 13:30:50 Resource SC15srاد is already running on member atlas0
19/Nov/2003 13:30:50 CAA cannot start a resource twice.
19/Nov/2003 13:30:51 NIS domain name set to sc
19/Nov/2003 13:30:52 ypserv daemon started
```

Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:30:52 ypbind: Secure mode sunos 3.x servers rejected.
19/Nov/2003 13:30:52 ypbind daemon started
19/Nov/2003 13:30:52 NFS mount daemon started
19/Nov/2003 13:30:52 NFS export service started
19/Nov/2003 13:30:53 Resource cluster_lockd is already running on member atlas0
19/Nov/2003 13:30:53 CAA cannot start a resource twice.
19/Nov/2003 13:30:53 NFS IO service started
19/Nov/2003 13:30:53 NFS Locking:
19/Nov/2003 13:30:53   rpc.statd started
19/Nov/2003 13:30:53   rpc.lockd started
19/Nov/2003 13:30:54 Mounting nfs member file systems
19/Nov/2003 13:30:54 Mounting nfsv3 member file systems
19/Nov/2003 13:30:54 NOTE:
19/Nov/2003 13:30:54     Skipping NFS file system entries found in /etc/fstab.
19/Nov/2003 13:30:54
19/Nov/2003 13:30:54     Please migrate the NFS file system entries to the member
19/Nov/2003 13:30:54     specific /etc/member_fstab of the member(s) which should
be
19/Nov/2003 13:30:54     mounting these NFS file systems.
19/Nov/2003 13:30:54
19/Nov/2003 13:30:54 Preserving editor files
19/Nov/2003 13:30:54 Clearing temporary files
19/Nov/2003 13:30:55 Unlocking ptys
19/Nov/2003 13:31:00 Secure Shell daemon (sshd2) started.
19/Nov/2003 13:31:00 SMTP Mail Service started.
19/Nov/2003 13:31:01 Network Time Service started
19/Nov/2003 13:31:01 The SNMP trap to Event Manager interface is disabled.
19/Nov/2003 13:31:01 GS Platform View and Discovery V1.3 for Insight Manager is
only supported on Alpha GS series platforms.
19/Nov/2003 13:31:02 Internet services provided.
19/Nov/2003 13:31:02 removing cron
19/Nov/2003 13:31:02 Cron service started
19/Nov/2003 13:31:06
19/Nov/2003 13:31:06 rms: RMS service started on atlas1
19/Nov/2003 13:31:07 Binary error logger started
19/Nov/2003 13:31:07 CFS load monitoring (cfsd) started
19/Nov/2003 13:31:07 CAA Applications now started
19/Nov/2003 13:31:07 cluster wall daemon started
19/Nov/2003 13:31:09 Starting Cluster Configuration Check...
19/Nov/2003 13:31:32
19/Nov/2003 13:31:32 The boottime cluster check found a potential problem.
19/Nov/2003 13:31:33 For details search for !!!!!ATTENTION!!!!!! in /cluster/admin/
clu_check_log_atlas1
19/Nov/2003 13:31:33 check_cdsl_config : Boot Mode : Running /usr/sbin/cdslinvchk
in the background
19/Nov/2003 13:31:33 check_cdsl_config : Results can be found in : /var/adm/
cdsl_check_list
19/Nov/2003 13:31:34 clu_check_config : detected one or more configuration errors
19/Nov/2003 13:31:35
19/Nov/2003 13:31:37 Enabling SCFS Serving on node atlas1.
19/Nov/2003 13:31:37 scmon: daemon started
19/Nov/2003 13:31:38 scmountd: not necessary to run on this node
19/Nov/2003 13:31:38 The system is ready.
```

Startup Messages After Adding a Domain Member

```
19/Nov/2003 13:31:38
19/Nov/2003 13:31:38
19/Nov/2003 13:31:38
19/Nov/2003 13:31:38 Compaq Tru64 UNIX V5.1B (Rev. 2650) (atlas1) console
19/Nov/2003 13:31:38
19/Nov/2003 13:32:01 login:
```

Configuring Networker for HP AlphaServer SC

The Networker software is a suite of storage management software that provides backup and recovery mechanisms for your data. One machine is designated to be your Networker Server which provides control and scheduling for Networker operations. All of the other machines are Network clients which provide recover and on-demand backup functionality.

This appendix describes how to install and configure the Networker software and contains the following sections:

- Domains as a Client of a Corporate Networker Server (see Section G.1 on page G–2)
- Domains as the Networker Server (see Section G.2 on page G–5)

Domains as a Client of a Corporate Networker Server

G.1 Domains as a Client of a Corporate Networker Server

G.1.1 Installing the Networker Client Software on the Cluster

1. Install the Networker client on `atlas` as follows:

Mount the CD-ROM as follows:

- a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/NetWorker/kit
```

Install the software, as follows:

```
atlasms# setld -l .
```

Note:

This step installs the software on every node within the cluster (including `atlas0`).

2. When installing the server software, accept all of the default questions:

- 1) Legato Networker Basic Client
- 2) Legato Networker Driver & Storage Node
- 3) Legato Networker Manpages
- 4) Legato Networker Server

Or you may choose one of the following options:

- 5) ALL of the above
- 6) CANCEL selections and redisplay menus
- 7) EXIT without installing any subsets

Estimated free disk space(MB) in root:332.7 usr:2501.1 var:1648.1

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 1

- Location of Networker database:

Enter the location for this Client's NetWorker home directory
(it must sit on a local disk!) [/var/nsr]:
/var/nsr

Domains as a Client of a Corporate Networker Server

- List of servers:

atlasD0

Note:

Ensure that each member of the cluster is added to ensure that failover to another node operates correctly.

Note:

When installing Networker, the software does not start up all the daemons on atlas0. It will only start the client daemons, and will do this on every node on the cluster.

3. After `setld` has completed, run the following script on every node in the cluster:

`/usr/opt/networker/bin/networker.cluster`

Note:

This script may fail and requires some manual updates as described in Step 4.

4. Edit the script as follows:

Before:

```
CLUINFO="\`setld -i | grep "TCRBASE.*installed" | awk '/TCRBASE/ { print $1 }'\`"
```

After:

```
CLUINFO="\`setld -i | grep '^TCRBASE.*installed' | awk '/TCRBASE/ { print $1 }'\`"
```

G.1.2 Adding Domains as a Client of Corporate Networker Server

In this example, the cluster alias is `atlasD0`, and the client of a corporate Networker Server (`gibbnew`). To configure the domain as a client of the corporate Networker Server:

1. Add `atlasD0` to `gibbnew` as a client
2. Back up the following partitions on `atlasD0`

```
/one_node_advfs  
/two_node_advfs  
/two_node_pfs
```

Domains as a Client of a Corporate Networker Server

3. Add each member of the `atlasD0` cluster to the remote-access list on the client in Networker. If this step is not performed, the backup will not be successful.

```
atlas0-ext1
```

4. Edit the `/etc/hosts` file on `atlas0` to include the Networker server, by adding the following:

```
gibbnew  
gibbnew.ilo.dec.com
```

Note:

Stopping and starting Networker takes a few minutes when Networker is installed on a cluster.

G.2 Domains as the Networker Server

In this example, the Networker Server is installed on domains. To configure the domain as the Networker Server, do the following:

1. Install the SCSI single-ended card into `atlas1`
2. Connect the SCSI tape drive to the cluster (for example, connect the drive to `atlas1`)
3. Reboot on `genvmunix`
4. Rebuild the kernel to boot the machine.
5. Install the Networker Server and client on `atlas` as follows:

Mount the CD-ROM as follows:

- a. Create a mount point for the CD-ROM, by running the following command:

```
atlasms# mkdir /cdrom
```

- b. Mount the CD-ROM as follows:

```
atlasms# mount -r /dev/disk/cdrom0c /cdrom
```

Change to the directory in which the kits are stored, as follows:

```
atlasms# cd /cdrom/NetWorker/kit
```

Install the software, as follows:

```
atlasms# setld -l .
```

Note:

This step installs the software on every node within the cluster (including `atlas0`).

6. When installing the server software, accept all of the default questions:

- 1) Legato Networker Basic Client
- 2) Legato Networker Driver & Storage Node
- 3) Legato Networker Manpages
- 4) Legato Networker Server

Or you may choose one of the following options:

- 5) ALL of the above
- 6) CANCEL selections and redisplay menus
- 7) EXIT without installing any subsets

```
Estimated free diskspace(MB) in root:332.7 usr:2501.1 var:1648.1
```

Enter your choices or press RETURN to redisplay menus.

```
Choices (for example, 1 2 4-6): 5
```

Domains as the Networker Server

- Location of Networker database:

Enter the location for this Client's NetWorker home directory
(it must sit on a local disk!) [/var/nsr]:
/var/nsr

- List of servers:

atlasD0

Note:

Ensure that each member of the cluster is added to ensure that failover to another node operates correctly.

Note:

When installing Networker, the software does not start up all the daemons on atlas0. It will only start the client daemons, and will do this on every node on the cluster.

7. After setld has completed, run the following script on every node in the cluster:

/usr/opt/networker/bin/networker.cluster

Note:

This script may fail and requires some manual updates as described in Step 8.

8. Edit the script as follows:

Before:

```
CLUINFO="`setld -i | grep "TCRBASE.*installed" | awk '/TCRBASE/ { print $1 }'`"
```

After:

```
CLUINFO="`setld -i | grep '^TCRBASE.*installed' | awk '/TCRBASE/ { print $1 }'`"
```

9. When this script has run to completion, register the Networker software within caa, by entering the following:

```
caa_register networker  
caa_start networker
```

If you prefer to specify the node on which to start the Networker software (for example, atlas0), enter the following command:

```
caa_start networker -c atlas0
```

10. Install the tape drive by running the following command:

```
/usr/opt/networker/bin/jbconfig
```

The options vary depending on the type of tape drive and also on whether it is SCSI or fibre. Ensure that you assign the following no-rewind tape device:

```
/dev/ntape/tape0
```

Note:

It is recommended to start `nwadmin` on the same node running the server. Initially, start `nwadmin` on the node running the daemon, and edit the server setup to change the authorized list of users to add `root@atlas0` (if the daemon is running on `atlas1`). An entry should be included for each node within the cluster.

11. Start `nwadmin` as follows:

```
/usr/opt/networker/bin/nwadmin
```

12. Select your backup preferences as follows:

- Full or Incremental
- Tape Name (pool names)

For more information, refer to the Networker documentation. It is recommended to create a default tape, a full tape, and an incremental tape.

13. Set up the following:

- Registration (license)
- Label_templates
- Groups
- Pools
- Server setup
- Client setup

Sample Output

G.3 Sample Output

```
atlas0 # cd /cdrom/NetWorker/kit
atlas0 # setld -i . l .
```

The subsets listed below are optional:

There may be more optional subsets than can be presented on a single screen. If this is the case, you can choose subsets screen by screen or all at once on the last screen. All of the choices you make will be collected for your confirmation before any subsets are installed.

- 1) Legato Networker Basic Client
- 2) Legato Networker Driver & Storage Node
- 3) Legato Networker Manpages
- 4) Legato Networker Server

Or you may choose one of the following options:

- 5) ALL of the above
- 6) CANCEL selections and redisplay menus
- 7) EXIT without installing any subsets

Estimated free disk space(MB) in root:332.7 usr:2501.1 var:1648.1

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 5

You are installing the following optional subsets:

```
Legato Networker Basic Client
Legato Networker Driver & Storage Node
Legato Networker Manpages
Legato Networker Server
```

Estimated free disk space(MB) in root:332.7 usr:2407.0 var:1648.1

Is this correct? (y/n): y

Checking file system space required to install selected subsets:

File system space checked OK.

4 subsets will be installed.

Loading subset 1 of 4 ...

This subset may take some time to complete.

```
Legato Networker Basic Client
Copying from . (disk)
```

```
Verifying
Loading subset 2 of 4 ...
This subset may take some time to complete.
Legato Networker Driver & Storage Node
  Copying from . (disk)
    Working....Tue Aug  6 10:33:58 IST 2002
  Verifying
Loading subset 3 of 4 ...
This subset may take some time to complete.
Legato Networker Manpages
  Copying from . (disk)
    Verifying
Loading subset 4 of 4 ...
This subset will take some time to complete.
Legato Networker Server
  Copying from . (disk)
    Working....Tue Aug  6 10:34:26 IST 2002
  Verifying
4 of 4 subsets installed successfully.
Configuring "Legato Networker Basic Client" (LGTOCLNT600) on member0
Legato Networker Basic Client
Copyright (c) 1990-2000, Legato Systems, Inc.
***** File Configuration on NetWorker Client *****
/nsr not found!
Enter the location for this Client's NetWorker home directory
(it must sit on a local disk!) [ /var/nsr ]:
The installation procedure adds entries to the /etc/rpc and
/etc/syslog.conf files on the NetWorker server; the original
files are renamed and saved. The installation also creates
the /sbin/init.d/NSRstartstop file.
Do you wish to continue? (y/n) [ y ]: y
Modifying /etc/rpc
Modifying /etc/syslog.conf
* * * Restarting syslog daemon * * *
Do you wish to remove the saved files? (y/n) [ n ]:
```

Sample Output

The modified files were saved and

renamed as follows:

File	Location of saved file
-----	-----
/etc/rpc	/etc/rpc_nsrsave
/etc/syslog.conf	/etc/syslog.conf_nsrsave

Creating /sbin/init.d/NSRstartstop

Starting nsrexecd...

The nsr/res/servers file will need to be updated with the list of servers that will back up this system as a client.

This is also needed if this machine is to be used as an HSM client.

/nsr/res/servers file does not exist..

Do you wish to create the file? (y/n):**y**

Creating a new /nsr/res/servers file..

Creation of /nsr/res/servers file done...

Enter one server name per prompt. eg: atom or atom.loc.xyz.com

Enter <CR> to terminate the list.

Please enter the name of a server to back up to (or <CR> to end):

atlasD0

Please enter the name of a server to back up to (or <CR> to end):

Do you want to stop and restart nsrexecd? (y/n): **y**

Starting nsrexecd...

LGTOCLNT600 software installed successfully

The Legato NetWorker Client version 600 binaries have been installed in /usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato Networker Driver & Storage Node" (LGTONODE600) on member0

Legato NetWorker Driver & Storage Node

Copyright (c) 1990-2000, Legato Systems, Inc.
Starting nsrexecd...

LGTONODE600 software installed successfully

The Legato NetWorker Storage Node version 600 binaries have been installed in /usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato NetWorker Manpages" (LGTOMAN600) on member0

Legato NetWorker Manpages
Copyright (c) 1990-2000, Legato Systems, Inc.
LGTOMAN600 software installed successfully

The Manpages are installed in the following locations:

/usr/opt/networker/man/man3
/usr/opt/networker/man/man5
/usr/opt/networker/man/man8

Please update the MANPATH environment variable to include the path
/usr/opt/networker/man.

Configuring "Legato NetWorker Server" (LGTOSERV600) on member0

Legato NetWorker Server

Copyright (c) 1990-2000, Legato Systems, Inc.

***** File Configuration on NetWorker Server *****

starting nsrd...
NetWorker icon has been successfully installed into CDE.

LGTOSERV600 software installed successfully

The Legato NetWorker Server version 600 binaries have been installed in /usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato NetWorker Basic Client" (LGTACLNT600) on member1

Sample Output

Legato NetWorker Basic Client

Copyright (c) 1990-2000, Legato Systems, Inc.
mkdir: cannot create /usr/opt/networker/bin/C.
/usr/opt/networker/bin/C: File exists

***** File Configuration on NetWorker Client *****

The installation procedure adds entries to the /etc/rpc and /etc/syslog.conf files on the NetWorker server; the original files are renamed and saved. The installation also creates the /sbin/init.d/NSRstartstop file.

Do you wish to continue? (y/n) [y]: Modifying /etc/rpc
/etc/rpc already modified for NetWorker
Modifying /etc/syslog.conf
/cluster/members/{memb}/etc/syslog.conf already modified for NetWorker
Do you wish to remove the saved files? (y/n) [n]:
The modified files were saved and

renamed as follows:

File	Location of saved file
----	-----
/etc/rpc	/etc/rpc_nsrsave
/etc/syslog.conf	/etc/syslog.conf_nsrsave

Creating /sbin/init.d/NSRstartstop

Starting nsrexecd...

The nsr/res/servers file will need to be updated with the list of servers that will back up this system as a client.

This is also needed if this machine is to be used as an HSM client.

/nsr/res/servers contains the following list of NetWorker servers:

atlasD0

Do you wish to add servers to the /nsr/res/servers file? (y/n):No changes will be made to the /nsr/res/servers file.

LGTOCLNT600 software installed successfully

The Legato NetWorker Client version 600 binaries have been installed in /usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato Networker Driver & Storage Node" (LGTONODE600) on member1

Legato Networker Driver & Storage Node

Copyright (c) 1990-2000, Legato Systems, Inc.
Starting nsrexecd...

LGTONODE600 software installed successfully

The Legato NetWorker Storage Node version 600 binaries have been installed
in /usr/opt/networker/bin. Please update the PATH environment
variable to include /usr/opt/networker/bin.

Configuring "Legato Networker Manpages" (LGTOMAN600) on member1

Legato Networker Manpages
Copyright (c) 1990-2000, Legato Systems, Inc.
LGTOMAN600 software installed successfully

The Manpages are installed in the following locations:

/usr/opt/networker/man/man3
/usr/opt/networker/man/man5
/usr/opt/networker/man/man8

Please update the MANPATH environment variable to include the path
/usr/opt/networker/man.

Configuring "Legato Networker Server" (LGTOSERV600) on member1

Legato Networker Server

Copyright (c) 1990-2000, Legato Systems, Inc.

***** File Configuration on NetWorker Server *****

starting nsrd...

LGTOSERV600 software installed successfully

The Legato NetWorker Server version 600 binaries have
been installed in /usr/opt/networker/bin.
Please update the PATH environment variable to include
/usr/opt/networker/bin.

Sample Output

Configuring "Legato Networker Basic Client" (LGTOCLNT600) on member2

Legato Networker Basic Client

Copyright (c) 1990-2000, Legato Systems, Inc.
mkdir: cannot create /usr/opt/networker/bin/C.
/usr/opt/networker/bin/C: File exists

***** File Configuration on NetWorker Client *****

The installation procedure adds entries to the /etc/rpc and /etc/syslog.conf files on the NetWorker server; the original files are renamed and saved. The installation also creates the /sbin/init.d/NSRstartstop file.

Do you wish to continue? (y/n) [y]: Modifying /etc/rpc
/etc/rpc already modified for NetWorker
Modifying /etc/syslog.conf
* * * Restarting syslog daemon * * *
Do you wish to remove the saved files? (y/n) [n]:
The modified files were saved and

renamed as follows:

File	Location of saved file
-----	-----
/etc/rpc	/etc/rpc_nsrsave
/etc/syslog.conf	/etc/syslog.conf_nsrsave

Creating /sbin/init.d/NSRstartstop

Starting nsrexecd...

The nsr/res/servers file will need to be updated with the list of servers that will back up this system as a client.
This is also needed if this machine is to be used as an HSM client.

/nsr/res/servers contains the following list of NetWorker servers:

atlasD0

Do you wish to add servers to the /nsr/res/servers file? (y/n):No changes will be made to the /nsr/res/servers file.

LGTOCLNT600 software installed successfully

The Legato NetWorker Client version 600 binaries have been installed in

/usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato Networker Driver & Storage Node" (LGTONODE600) on member2

Legato Networker Driver & Storage Node

Copyright (c) 1990-2000, Legato Systems, Inc.
Starting nsrexecd...

LGTONODE600 software installed successfully

The Legato NetWorker Storage Node version 600 binaries have been installed in /usr/opt/networker/bin. Please update the PATH environment variable to include /usr/opt/networker/bin.

Configuring "Legato Networker Manpages" (LGTOMAN600) on member2

Legato Networker Manpages
Copyright (c) 1990-2000, Legato Systems, Inc.
LGTOMAN600 software installed successfully

The Manpages are installed in the following locations:

/usr/opt/networker/man/man3
/usr/opt/networker/man/man5
/usr/opt/networker/man/man8

Please update the MANPATH environment variable to include the path /usr/opt/networker/man.

Configuring "Legato Networker Server" (LGTOSERV600) on member2

Legato Networker Server

Copyright (c) 1990-2000, Legato Systems, Inc.

***** File Configuration on NetWorker Server *****

starting nsrd...

LGTOSERV600 software installed successfully

The Legato NetWorker Server version 600 binaries have been installed in /usr/opt/networker/bin.

Sample Output

```
Please update the PATH environment variable to include
/usr/opt/networker/bin.
```

```
atlas0 # /usr/opt/networker/bin/networker.cluster
```

```
-----
TruCluster is a high availability product for Tru64 Unix.
It defines failover applications which may move from one node to
another, depending on the availability and health of the machine
it is running upon.
The Networker Client and Networker Server have the default cluster
identity.
-----
```

```
Do you wish to continue? [Yes]?
Shutting down NetWorker services...
Restarting syslog daemon...
/ccluster/members/{memb}/etc/syslog.conf already modified for NetWorker
Restarting syslog daemon...
```

```
NetWorker has been successfully cluster configured.
atlas0 # exit
```

```
Script done on Tue Aug 6 10:45:02 2002
```

Cluster-Aliases and External Networks in the 10.x.x.x Range

In cases where the addressing scheme on the external LAN needs to be within the 10.x.x.x range, you must change the addressing scheme of both the internal management LAN and the HP AlphaServer SC Interconnect to use the 172.x.x.x notation.

Table H–1 displays the network address convention used for these networks and attached devices. Use the address notation in Table H–1 during the `sysman` (Section 6.1.4.2) and `sra setup` (Section 6.2) stages of the HP AlphaServer SC installation.

Table H–1 Alternative HP AlphaServer SC IP Addresses

Component	IP Address Range
Net mask	255.255.0.0
Cluster Interconnect (IP suffix: <code>-ics0</code>) for node n , where n is 0-127	172.20.0.($n+1$)
System Interconnect (IP suffix: <code>-eip0</code>) for node n , where n is 0-127	172.24.0.($n+1$)
Management network interface card for node n , where n is 0-127	172.28.0.($n+1$)
Terminal server t , where t is 1–254	172.28.100. t
Management server m on management LAN, where m is 1–2	172.28.101. m
Management server m Cluster Interconnect (IP suffix: <code>-ics0</code>)	172.22.0. m
Management server m Gigabit Interconnect (IP suffix: <code>-icstcp0</code>)	172.23.0. m
HP AlphaServer SC Interconnect switch JTAG port j , where j is 1–48	172.28.102. j
Summit switch s , where s is 1–254	172.28.103. s
HP SANworks Management Appliance or Fibre Channel switch f , where f is 1-254	172.28.104. f
RAID array controller c , where c is 1, 2, and so on	172.28.105. c

Table H–1 Alternative HP AlphaServer SC IP Addresses

Component	IP Address Range
HP AlphaServer SC Interconnect Control Card for Node-Level switch <i>N</i> , where <i>r</i> is the rail number and <i>N</i> is 0-31	$172.28.(128+r).(N+1)$
HP AlphaServer SC Interconnect Control Card for Top-Level switch <i>T</i> , where <i>r</i> is the rail number and <i>T</i> is 0-15	$172.28.(128+r).(T+128)$
Preferred server cluster alias addresses (where <i>FS</i> is 0-3, thus accommodating a maximum of four FS domains, and <i>y</i> is the member ID within the FS domain)	$172.28.(106+FS).y$

Configuring DNS Servers

This appendix describes how to set up a management server as a DNS server and contains the following sections:

- Management Server as DNS Server and Cluster as Client (see Section I.1 on page I-2)
- Management Server as Master DNS Server and Cluster as Slave DNS Server (see Section I.2 on page I-5)
- Other Configurations (see Section I.3 on page I-8)

I.1 Management Server as DNS Server and Cluster as Client

I.1.1 On the Management Server

Note:

This step should be performed before you run `sra setup`, and so should replace the procedure described in Section 5.1.5.3 and Section 6.1.4.3.

To configure the management server as the DNS server, and the cluster as the client, perform the following steps:

1. Log on to the management server as `root`.
2. Run the following command:

```
atlasms# sysman dns
```

or within Sysman, select the option, *Networking, Additional Network Services, Domain Name Service (DNS/BIND)*.
3. Configure the system as a DNS server.
4. Enter a local domain name (`mydomain.net`).
5. Enter the *DNS server type* = `Master` (assume you are not using slave or any other types).
6. Enter the default Hostname Resolution order. Select the default entries: `Local Host file`, `DNS database`, `NIS`.
7. Select the *Additional/Advanced* options, and select the default input (for example, location of configuration files).
8. Select the *OK/Next screen* option.
9. Add a name server, and select the hostname (unqualified) of the management server.
10. Allow Sysman to resolve the IP address and confirm that it is correct (check that the address is the `10.128.xxx.xxx` address of the management server).
11. Select the *OK/Next screen* option.
12. When prompted to create the DNS database from the `/etc/hosts` file, enter **Yes** (or select the checkbox).
13. When prompted to change the hostname, enter **No** (to avoid confusing the `rmshost` settings).
14. When prompted, select the option to *Start/Restart* the `named` daemon.

Management Server as DNS Server and Cluster as Client

15. To check that the DNS is working, run `nslookup` to confirm the resolution of hostnames defined in the `/etc/hosts` file.
16. Confirm that the Name Server address and name is correct.
17. If you wish to add addresses to `nslookup`, do the following:
 - a. Copy the hosts file to the DNS source directory as follows:

```
atlasms# cp /etc/hosts /etc/namedb/src/
```

- b. Edit this copy of hosts: `/etc/namedb/src/hosts` and make the appropriate additions/deletions.
 - c. Run the makefile to update the DNS database as follows:

```
atlasms# /etc/namedb/make all
```

This command takes this copy of the hosts file, rebuilds the DNS database, and restarts the `named` daemon (DNS server).

This procedure assumes that clients only need to know about addresses in their own domain and are not concerned with external domains. Other configurations are described in Section I.1.2.

I.1.2 On the Cluster

Perform the following steps after the cluster has been successfully installed and booted:

Note:

In this case, the hosts file contains all entries for all nodes and their networks (for example, `atlas24-ext`, `atlas24-eip`, `atlas24-ics` and so on). The entries are added to the DNS database.

1. Log on to any node in the cluster (typically the master node `atlas0`) as `root`.
2. Run the following command:

```
atlas0# sysman dns
```

or within Sysman select the option, *Networking, Additional Network Services, Domain Name Service (DNS/BIND)*.
3. Configure the system as a DNS client.
4. On the local domain, select the domain name (`mydomain.net`).
5. Confirm that the IP address is correct.
6. Enter the Hostname Resolution Order. Select the default entries: Local Host file, DNS database, NIS.

Management Server as DNS Server and Cluster as Client

7. Select the *Additional/Advanced* options, and select the Domains searched. Select the default options (none).

8. Select the *OK/Next screen* option.

A message is displayed that DNS is configured on the machine.

9. To check that the DNS is working, run `nslookup` to confirm the resolution of hostnames defined in the `/etc/hosts` file.

10. Confirm that the Name Server address and name is correct.

I.2 Management Server as Master DNS Server and Cluster as Slave DNS Server

I.2.1 On the Management Server

Note:

This step should be performed before you run `sra setup`, and so should replace the procedure described in Section 5.1.5.3 and Section 6.1.4.3.

To configure the management server as a master DNS server, refer to the steps in Section I.1.1. To configure the cluster as a slave DNS server refer to the steps in Section I.2.2.

I.2.2 On the Cluster

1. Log on to any node in the cluster (typically the master node `atlas0`) as `root`.
2. Run the following command:

```
atlas0# sysman dns
```


or within Sysman select the option, *Networking, Additional Network Services, Domain Name Service (DNS/BIND)*.
3. Configure the system as a DNS server
4. Enter a local domain name (`mydomain.net`).
5. Enter the *DNS server type* = `Slave`.
6. Enter the default Hostname Resolution order. Select the default entries: `Local Host file`, `DNS database`, `NIS`.
7. Select the *Additional/Advanced* options, and select the default input (for example, location of configuration files).
8. Select the *OK/Next screen* option.
9. On the *Zones Served* screen, the first entry (`0.0.127.in-addr.arpa`) is displayed by default. Do not change or remove this entry. Perform the following steps:
 - a. Select the *Add* option, to display the next screen.
 - b. Select the *Slave* option.
 - c. Enter the domain identifier for this zone

Management Server as Master DNS Server and Cluster as Slave DNS Server

Note:

This is only different from the domain entered in step 4 above if you intend for different slaves to operate for different zones within the domain. In the simplest case, you will not have multiple zones or slaves, so this will be the local domain entered in step 4.

- d. Enter the file name `hosts.slave.db` (effectively the slave copy of the DNS database)
 - e. Enter the IP address of the Master DNS Server (in this example, the IP address of the Management Server, that was set up in Section I.1)
10. Repeat steps a to e above specifying the domain identifier as follows: `'nnn.nnn.nnn.inaddr.arpa'`, where `'nnn.nnn.nnn'` is the reverse IP address of the Master DNS Server.

Note

If the Master IP address is 16.123.456.789, this entry will be for `456.123.16.inaddr.arpa` thereby defining reverse lookups for all IP addresses in that range. The file name in this case will be `hosts.slave.rev` (that is, the reverse lookup file for the DNS database). The IP address will be the IP address of the Master DNS server.

11. Add any more zones being served by this slave (if appropriate) in the same manner.
12. Select the *OK/Next screen* option.
13. Add a name server. Specify the hostname and IP address of the Master DNS server (as setup in Section I.1)
14. Select the *OK/Next screen* option.
15. When prompted, select the option to *Start/Restart* the `named` daemon. Click the checkbox and select the *Next* option.
16. When prompted to change the hostname, enter **no** (to avoid confusing the `rmshost` settings).
17. Select the *Finish* option to complete the setup.

I.2.3 Test the Slave DNS Server

To test that the slave DNS server is working:

1. Verify that the `named` daemon has been started on the slave server. If this is a cluster (management server cluster or standard SC cluster), the `named` daemon is started as a `caa_application`, so use `caa_stat` to check the status.
2. Check that the `hosts.slave.db` and `hosts.slave.rev` database files have been transferred successfully over from the master DNS server. The files should reside in the location: `/etc/namedb/` on the slave server. On a cluster, log on to the cluster alias to check the file location.
3. Run `nslookup` on the slave server (in this case any node in the cluster) and check that all hosts defined in the master database can be resolved.
4. On the Master DNS server (in this case the management server), add another host entry into the `/etc/namedb/src/hosts` file, and rebuild the DNS database as follows:

```
atlasms# /etc/namedb/make all
```

Refer to step 17 in Section I.1.1 above for more explanation.

5. Stop and restart the `named` daemon on the master (use the `caa_stop` or `caa_start` commands if on a cluster).
6. On the slave server, run the `nslookup` command and resolve one of the new names added to the server's DNS database.

Other Configurations

I.3 Other Configurations

The management server can be also set up as a slave DNS server, with the cluster(s) as DNS clients, by following the steps in Section I.2.

Configuring MSA1000

This appendix provides information on configuring MSA1000 and contains the following sections:

- Preparing to Upgrade the MSA1000 Firmware (see Section J.1 on page J-2)
- Upgrading MSA1000 Firmware (see Section J.2 on page J-3)
- MSA1000 Sample Configuration (see Section J.3 on page J-7)

Preparing to Upgrade the MSA1000 Firmware

J.1 Preparing to Upgrade the MSA1000 Firmware

For qualification reasons, the MSA1000 controller firmware needs to be upgraded to at least the recommended revision before the storage is actually used by the Tru64 systems.

The process of upgrading the MSA1000 firmware requires a single Tru64 node with a fibre channel connection to the MSA1000. If the management server has such a fibre channel connection, then the upgrade can be performed using the management server while it is booted from CD-ROM.

If the management server does not have a fibre channel connection, then the first node of the first domain can be used to perform the upgrade. However, it will be necessary to install the domain(s) in two phases so that the MSA1000 controller firmware can be upgraded prior to the MSA1000 storage itself being used by the `clu_create` phase. The sequence in such a situation is as follows:

1. Follow the steps to install when the system has a management server as described in Chapter 5.
2. Begin the installation of the domains as described in Create the Domains (see Section 7.4 on page –23), but halt the installation process at the `NHD_Installed` state as follows:

```
atlasms# sra install -domains all -unixpatch /patches/patch_kit -sckit  
/cdrom/kits -sysconfig /cdrom/Examples/sysconfigtab -endstate  
NHD_Installed
```

3. Upgrade the MSA1000 controller firmware from the `atlas0` node as described in Upgrading MSA1000 Firmware (see Section J.2 on page J–3)
4. Configure the MSA1000 RAID Controllers as described in Configure the MSA1000 RAID Storage (see Section 3.16.2 on page –64)
5. At this point, the normal installation process should be resumed beginning with the following command:

```
atlasms# sra install -domains all -unixpatch /patches/patch_kit -sckit  
/cdrom/kits -sysconfig /cdrom/Examples/sysconfigtab
```

Where this command will continue from the endstate defined in the earlier command. The remainder of the system installation will now proceed as documented in Chapter 7.

J.2 Upgrading MSA1000 Firmware

Assuming that `atlas0` is booted to a UNIX prompt, then to upgrade the MSA1000 controller firmware, you should perform the following steps.

1. Verify the current MSA1000 controller firmware, by running the following command from the MSA1000 command line interface:

```
CLI> show this_controller
Controller:
MSA1000(c) Compaq P56350D9IOS070 Software 2.38 Build 122 Hardware 7
Controller Identifier:
NODE_ID = 500805F3-000751E0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy Mode: Active/Standby
Current Role: Active
Device Port SCSI address 6
Host Port_1:
REPORTED PORT_ID 500805F3-000751E1
PORT_1_TOPOLOGY = F_Port
Cache:
Unconfigured Version 2
Cache is GOOD, but Cache is NOT configured.
No unflushed data in cache
Battery:
Module #1 is fully charged and turned on.
CLI>
```

Table 3–5 lists the supported firmware versions. If the controller firmware version is lower than the recommended version then you should upgrade the MSA1000 firmware as follows:

2. Download the MSA1000 firmware zip file from the following location:

```
<http://h18006.www1.hp.com/products/storageworks/software/drivers/msa1000/index.html>
```

3. Create a directory on the management server as follows:

```
atlas0# mkdir /usr/kits/msa1000-fw
```

4. Extract to the new directory the following files: `MSAv330A.bin` and `Tru64InstallNotes.txt`
5. Follow the `Tru64InstallNotes.txt` file for the most up to date steps required to upgrade the firmware. However, an example of the steps required to upgrade the MSA1000 firmware is shown in Example J–1.

Upgrading MSA1000 Firmware

Example J-1 Upgrading MSA1000 Firmware - Sample Output

```
CLI> show this_controller
Controller:
MSA1000(c) Compaq P56350D9IOS070 Software 2.38 Build 122 Hardware 7
Controller Identifier:
NODE_ID = 500805F3-000751E0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy Mode: Active/Standby
Current Role: Active
Device Port SCSI address 6
Host Port_1:
REPORTED PORT_ID 500805F3-000751E1
PORT_1_TOPOLOGY = F_Port
Cache:
Unconfigured Version 2
Cache is GOOD, but Cache is NOT configured.
No unflushed data in cache
Battery:
Module #1 is fully charged and turned off.
CLI>
```

Identify the bus/target/lun identification for the MSA1000 to upgrade as follows:

```
atlas0# /sbin/hwmgrr view device | grep MSA1000 | grep -v VOLUME
81: /dev/cport/scp0 MSA1000 bus-4-targ-0-lun-0
```

```
atlas0# /sbin/hwmgrr show scsi -full -id 81
SCSI DEVICE DEVICE DRIVER NUM DEVICE FIRST
HWID: DEVICEID HOSTNAME TYPE SUBTYPE OWNER PATH FILE VALID PATH
-----
81: 4 scmsa0 raid none 0 2 scp0 [4/0/0]
WWID:02000008:5008-05f3-0007-51e0
BUS TARGET LUN PATH STATE
-----
4 0 0 valid
5 0 0 valid
```

Launch the SCSI CAM Utility program to upgrade the MSA controller firmware as follows:

```
atlas0# /sbin/scu
scu> sbtl 4 0 0

Device: MSA1000, Bus: 4, Target: 0, Lun: 0, Type: Array Controller
scu> tur
scu> download MSAv330A.bin save segment
Downloading File 'MSAv330A.bin' of 1048576 bytes in 8192 byte segments...
Download completed successfully, now saving the microcode...
Delaying for 120 seconds while firmware is saved, please be patient... 100
.
Delaying for 120 seconds while firmware is saved, please be patient... 85
.
scu> exit
```

Upgrading MSA1000 Firmware

atlas0#

At this point, the MSA1000 command line interface will display the following message:

```
CLI>
WRITE_BUFFER SUCCESS: Successfully Validated MSA1000 FW Image!
Flashing MSA1000 Firmware...done
Powercycle to activate new MSA1000 FW!
```

You should now perform the following additional steps to upgrade the left controller:

1. Powercycle the right controller and wait for the boot sequence to complete.
2. Plug in the left controller.
3. Clone the left controller firmware by using the front panel arrows on the MSA1000 left controller.
4. Verify the updated firmware MSA1000 controller firmware, by running the following commands from the MSA1000 command line interface:

```
CLI> show this_controller
Controller:
MSA1000(c) Compaq P56350D9IOS070 Software 3.30 Build 211 Hardware 7
Controller Identifier:
NODE_ID = 500805F3-000751E0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy Mode: Active/Standby
Current Role: Active
Device Port SCSI address 6
Host Port_1:
REPORTED PORT_ID 500805F3-000751E1
PORT_1_TOPOLOGY = F_Port
Cache:
Unconfigured Version 2
Cache is GOOD, but Cache is NOT configured.
No unflushed data in cache
Battery:
Module #1 is fully charged and turned off.
```

```
CLI> show other_controller
Controller:
MSA1000(c) Compaq P56350D9IOS0KX Software 3.30 Build 211 Hardware 7
Controller Identifier:
NODE_ID = 500805F3-000751E0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy Mode: Active/Standby
Current Role: Standby
Device Port SCSI address 7
Host Port_1:
REPORTED PORT_ID 500805F3-000751E9
PORT_1_TOPOLOGY = F_Port
```

Upgrading MSA1000 Firmware

```
Cache:  
Unconfigured Version 2  
Cache is GOOD, but Cache is NOT configured.  
No unflushed data in cache  
Battery:  
Module #1 is fully charged and turned off.  
CLI>
```

J.3 MSA1000 Sample Configuration

Example J-2 displays a sample flow and output when configuring MSA1000.

Example J-2 Sample MSA1000 Configuration

CLI> **show connections**

Global_ris doesn't exist. No named connections exist.

Connection Name: <Unknown>

Host WWNN = 20000000-C9311D90

Host WWPNN = 10000000-C9311D90

Profile Name = Default

Unit Offset = 0

Controller 2 Port 1 Status = Online

Connection Name: <Unknown>

Host WWNN = 20000000-C9311CD6

Host WWPNN = 10000000-C9311CD6

Profile Name = Default

Unit Offset = 0

Controller 1 Port 1 Status = Online

CLI>

The message Global_ris doesn't exist. No named connections exist shown above is displayed when there is no storage unit created. Creating storage units is the first action to perform on MSA1000 configuration, otherwise all commands would fail, as can be seen for the add connection command below:

CLI> **add connection N0000-0 Host WWPNN = 10000000-C9311CD6 profile=Tru64 offset=0**
Cannot save any connection information because global_ris is NULL.

CLI> **show disks**

Disk List:	(box,bay)	(bus,ID)	Size	Units
Disk101	(1,01)	(0,00)	36.4GB	none
Disk102	(1,02)	(0,01)	36.4GB	none
Disk103	(1,03)	(0,02)	36.4GB	none
Disk104	(1,04)	(0,03)	36.4GB	none
Disk105	(1,05)	(0,04)	36.4GB	none
Disk108	(1,08)	(1,00)	36.4GB	none
Disk109	(1,09)	(1,01)	36.4GB	none
Disk110	(1,10)	(1,02)	36.4GB	none
Disk111	(1,11)	(1,03)	36.4GB	none
Disk112	(1,12)	(1,04)	36.4GB	none
Disk114	(1,14)	(1,08)	36.4GB	none
Disk201	(2,01)	(2,00)	36.4GB	none
Disk202	(2,02)	(2,01)	36.4GB	none
Disk203	(2,03)	(2,02)	36.4GB	none
Disk204	(2,04)	(2,03)	36.4GB	none

MSA1000 Sample Configuration

Disk205	(2,05)	(2,04)	36.4GB	none
Disk208	(2,08)	(3,00)	36.4GB	none
Disk209	(2,09)	(3,01)	36.4GB	none
Disk210	(2,10)	(3,02)	36.4GB	none
Disk211	(2,11)	(3,03)	36.4GB	none
Disk212	(2,12)	(3,04)	36.4GB	none

```
CLI> add unit 1 RAID_level=1 data="disk101 disk201" size=16GB spare="disk114"
```

First volume to be configured on these drives.

The logical unit size has been adjusted by 6MB for optimal performance.

Logical Unit size = 16378 MB

RAID overhead = 16378 MB

Total space occupied by new unit = 32756 MB

Free space left on this volume: = 36696 MB

Unit 1 is created successfully.

```
CLI> add unit 2 RAID_level=1 data="disk101 disk201" size=1500MB spare="disk114"
```

The logical unit size has been adjusted by 2MB for optimal performance.

Logical Unit size = 1498 MB

RAID overhead = 1498 MB

Total space occupied by new unit = 2996 MB

Free space left on this volume: = 33698 MB

Unit 2 is created successfully.

```
CLI> add unit 3 RAID_level=1 data="disk101 disk201" size=16GB spare="disk114"
```

The logical unit size has been adjusted by 6MB for optimal performance.

Logical Unit size = 16378 MB

RAID overhead = 16378 MB

Total space occupied by new unit = 32756 MB

Free space left on this volume: = 942 MB

Unit 3 is created successfully.

```
CLI> show disks
```

Disk List:	(box,bay)	(bus,ID)	Size	Units
Disk101	(1,01)	(0,00)	36.4GB	1, 2, 3
Disk102	(1,02)	(0,01)	36.4GB	none
Disk103	(1,03)	(0,02)	36.4GB	none
Disk104	(1,04)	(0,03)	36.4GB	none
Disk105	(1,05)	(0,04)	36.4GB	none
Disk108	(1,08)	(1,00)	36.4GB	none
Disk109	(1,09)	(1,01)	36.4GB	none
Disk110	(1,10)	(1,02)	36.4GB	none
Disk111	(1,11)	(1,03)	36.4GB	none
Disk112	(1,12)	(1,04)	36.4GB	none
Disk114	(1,14)	(1,08)	36.4GB	1, 2, 3 (spare)
Disk201	(2,01)	(2,00)	36.4GB	1, 2, 3
Disk202	(2,02)	(2,01)	36.4GB	none
Disk203	(2,03)	(2,02)	36.4GB	none
Disk204	(2,04)	(2,03)	36.4GB	none
Disk205	(2,05)	(2,04)	36.4GB	none
Disk208	(2,08)	(3,00)	36.4GB	none
Disk209	(2,09)	(3,01)	36.4GB	none

MSA1000 Sample Configuration

Disk210	(2,10)	(3,02)	36.4GB	none
Disk211	(2,11)	(3,03)	36.4GB	none
Disk212	(2,12)	(3,04)	36.4GB	none

CLI> **show connections**

Connection Name: <Unknown>
Host WWNN = 20000000-C9311D90
Host WWPN = 10000000-C9311D90
Profile Name = Default
Unit Offset = 0
Controller 2 Port 1 Status = Online

Connection Name: <Unknown>
Host WWNN = 20000000-C9311CD6
Host WWPN = 10000000-C9311CD6
Profile Name = Default
Unit Offset = 0
Controller 1 Port 1 Status = Online

CLI>

CLI> **set unit_id 1 1**
Device identifier 1 created.

CLI> **set unit_id 2 2**
Device identifier 2 created.

CLI> **set unit_id 3 3**
Device identifier 3 created.

CLI>

CLI> **add connection N0000-0 WWPN=10000000-C9311CD6 profile=Tru64 offset=0**
Connection has been added successfully.
Profile Tru64 is set for the new connection.

CLI> **add connection N0000-1 WWPN=10000000-C9311D90 profile=Tru64 offset=0**
Connection has been added successfully.
Profile Tru64 is set for the new connection.

CLI> **show connections**

Connection Name: N0000-0
Host WWNN = 20000000-C9311CD6
Host WWPN = 10000000-C9311CD6
Profile Name = Tru64
Unit Offset = 0
Controller 1 Port 1 Status = Offline

Connection Name: N0000-1
Host WWNN = 20000000-C9311D90

MSA1000 Sample Configuration

Host WWPN = 100000000-C9311D90
Profile Name = Tru64
Unit Offset = 0
Controller 2 Port 1 Status = Offline

A

Abbreviations, xxv
AdvFS Domains, 6-44
AlphaServer SC Interconnect, 3-20
AlphaServer SC System Components, 3-4
Audit
 Pre-Upgrade, 4-8

B

Backing Up File Systems, 10-7
BIND
 See DNS/BIND
Boot Disks, 5-40, 6-27, 6-43
Boot Errors, 11-17
Bootstrap-fail Error, 11-18

C

CAA
 See Failover
Cabinets
 Assigning, 5-52, 6-38
Cables
 AlphaServer SC Interconnect, 3-20
 Console Network, 3-21
 Fibre Channel, 3-23, 3-28, 3-63
 Management Network, 3-11
CDSLs, 6-44

Checklist
 See Information Checklists
Cluster Alias
 Default, 2-3
 Overview, 10-13
Cluster Availability, 8-3
Cluster Disk, 6-44
Cluster Member
 Adding Votes, 8-4
 Restricting Access, 10-3
Cluster Naming Scheme, 2-6
Cluster Quorum
 See Cluster Availability
Clustered Management Server
 Installing and Configuring, 5-22
 Upgrading with, 4-21
CMF as CAA Application, 8-15
Commands
 caa_register, 11-20
 clu_add_member, E-23
 clu_create, E-20
 clu_quorum, 8-5, E-25
 dhcp, 11-18
 msqladm, 11-29
 rcontrol, 8-11
 rinfo, 8-11
 ris, 4-18, 4-28, 4-36, 5-31, 6-19
 setld, 5-24, 5-33, 6-21
 sra_add_member, F-9
 sra boot, 8-5, 8-7
 sra edit, 5-44, 6-30, 7-14, 7-15, 7-19
 sra install, 7-2, 7-24
 sra power_off, 11-23
 sra power_on, 11-23

- sra setup, 11-26, E-2
- sra shutdown, 8-5
- sramon, 7-13
- upgrade_check, E-27

Console Network, 3-21, 10-10

Console Variables
See SRM Console

CS Domain, 2-4

D

Daemons

- CAA, 11-21
- Console Logger (cmfd), 11-27
- Installation, 7-2
- joind, 11-18
- portmap, 11-22

Database, Dropping and Rebuilding, 11-29

Default Cluster Alias, 2-3

Disk

- Labels, 6-44
- Layout, 2-9

DNS (BIND), 5-17, 6-14

DNS Server

- Configuring, I-1, J-1

Documentation

- Conventions, xxvii
- Online, xxx

Domain, 2-4

Domains, xxiix

- Creating, 7-23
- Upgrade Mechanism, 4-6
- Upgrading, 4-41

Dual-Rail

- Configuring Hardware, 3-20

E

Elan Error, 11-23

Example MPI Program, 8-17

Examples

- Code, xxx

F

Failover

- RMS, 8-12

Fibre Channel, 3-23

Fibre Channel Switch, 2-16

Files

- Preserved and Unpreserved, 4-5

Fortran Run-Time Libraries, 5-33, 6-21

FS Domain, 2-4

G

Gateway

- Default, 5-15, 6-13
- External, 2-16

H

Hardware Ethernet (MAC) Address, 6-28

Host Name, 5-7, 6-5

Hosts File, 5-15, 6-13

I

Information Checklists, D-1

Installation Checklists, C-1

Installation Output Samples, E-1

Installation Overview, 1-2

Installation Problems, 11-5

Installation State

- Monitoring, 7-4

Installation Tips, 11-2

InstallSC Errors, 11-5

IP Addresses

- AlphaServer SC Interconnect, 5-38, 6-25
- Alternative Notation, H-1
- External Gateway, 2-16
- External Network, 2-2
- Management Network, 5-15, 5-16, 6-13
- Table of, H-1

Terminal Server, 3-22

K

Kernel

Build, 5-9, 6-7, 6-44
Variables, 6-22

L

Ladefug Debugger, 5-13, 6-11

Layered Products, 5-33, 6-21

License

PAKs, 5-10, 5-33, 6-8, 6-21
Requirements, 1-3

LSF

Configuring, 9-10
Installing in NFS System Environment, 9-3
Overview, 9-2

LSM, 5-8, 6-6, 8-18

M

Mail Configuration, 5-21, 6-17

Management Server

Clustered, 3-3
Set up the, 5-2
Standalone, 3-2
Upgrading with, 4-10
Upgrading without, 4-32

Manual Cluster Creation, 11-27

Monitoring setup for HSG80 RAID Systems, 8-8

Multiple-Bus Failover Mode, 3-32

N

Network

External, 2-2, 8-2
Setup, 5-14, 6-12, 8-2
Speed, 11-18

Network Adapter, xxix

Networker

Configuring, G-1

New and Changed Features, xxi

NFS (Network File System), 11-22

NIS (Network Information System), 5-19, 6-16

Node 0 Setup, 6-2

Node Boot Error, 11-23

Node Installation Status, 11-6

Nodes

Checking, 5-45, 6-31
Configuring out during Installation, 5-45

NTP (Network Time Protocol), 5-18, 6-15

P

PAK

See License

Post-Installation Tasks, 10-3

Powercycle, 11-23

Printer Configuration, 5-21, 6-18

Q

Quorum

See Cluster Availability

R

RAID, 3-31

Reset Button, 11-27

RIS (Remote Installation Services), 5-31, 6-19, 11-20, 11-21

RMS (Resource Management System)

Database, 8-10
Mandatory Administration Tasks, 10-13
Optional Administration Tasks, 10-14
Partitions, 8-10, 8-11
rmshost, 5-34, 6-23
Starting Manually, C-8

Root Password, 2-16, 5-7, 6-6

S

- SANworks Command Scripter, 3-30
- SC Database Disk Settings, 6-41
- SC Database Setup, 5-37, 6-24
- SC Database System Settings, 7-14
- SC Interconnect Verification, 8-17
- SC Monitor
 - Setting Up, 5-50, 6-35
- SRM Console, 3-11, 5-3, 5-6, 5-42, 6-3, 6-4, 6-29
- SRM Device Name, 5-41, 6-28
- States
 - Checking Upgrade, 4-5
 - Upgrade, 4-4
- Storage
 - Global, 10-11
 - Local, 10-11
 - System, 3-27, 10-11
- Subnet mask, 2-5
- Subsets, 1-4
- Summit Switch, 3-12
- Swap Mode, 8-18
- Swap Space, 2-11, 5-40, 6-27
- Switch Zoning, 3-26
- sysconfigtab Parameters, 6-22, 7-23
- System Name, 2-3
- System Storage
 - Configuring, 3-27

T

- Terminal Server, 3-22, 11-27
- Troubleshooting, 11-1
- Tru64 UNIX Operating System Software, 5-6, 6-4, 7-23

U

- Upgrade

- Files Preserved and Unpreserved, 4-5
 - Interpreting Problems, 11-40
 - Mechanisms, 4-6
 - Overview of Procedure, 4-3
 - Post Upgrade Tasks, 4-49
 - Pre-check, 4-8
 - Restarting, 4-4
 - Time Estimates, 4-7
- Upgrade Errors, 11-37
- Upgrade Failure
 - Recovery, 4-45
- Upgrade Paths, 4-2
- Upgrade Process
 - Description, 4-2
- Upgrade States, 4-4
- Upgrading
 - Domains, 4-41
- User Administration, 10-12

V

- Votes
 - See* Cluster Availability