N1 Provisioning Server 3.1, Blades Edition, System Administration Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817–4684–10
March, 2004

040219@7940

Contents

Preface

This section of the System Administration Guide briefly introduces the product, lists the contents of this guide, and how to use this guide. The following topics are covered in this section:

About the Product

N1 Provisioning Server software provides a comprehensive infrastructure automation solution that enhances the management, visibility, and control of computing environments. N1 Provisioning Server software is part of an I-Fabric, which may contain multiple farms serving multiple organizations and users within those organizations. The software enables the complete design configuration, deployment, and management of multiple independent, secure, logical server farms.

Who Should Use This Book?

This guide is intended for system administrators who are responsible for managing an overall Infrastructure Fabric (I-Fabric) computing environment. Administrators should have knowledge of or be familiar with the following:

- Text editors and software that employs command-line interfaces
- UNIX® operating environments, file management commands and operations
- Server blades
- Computer networking
- Database administration
- Client/server software and related management issues
- Traditional server farm management
- Running scripts
- Markup languages

How This Book Is Organized

This guide contains procedures for the system administrator to perform various tasks for running and managing an I-Fabric. This guide is divided into the following chapters.

- Chapter 1 provides a general overview of an I-Fabric, its components, and the N1 Provisioning Server software that controls how I-Fabric components interoperate.
- Chapter 2 describes day-to-day I-Fabric operation activities, how to add, delete, and replace I-Fabric components, and how to set up passwords. This chapter also discusses the `tspr.properties` configuration file, which modifies the behaviors of all N1 Provisionable Server software.
- Chapter 3 describes how the N1 Provisioning Server software enables the management of various software images in an I-Fabric. Managing software images includes creating images and updating, deleting, and listing existing images.
- Chapter 4 describes the monitoring and messaging system and the events being monitored.
- Chapter 5 describes how to back up and restore various I-Fabric components.
- Chapter 6 describes error messages, their meanings, and actions required to solve a problem. It also describes how to gather additional information before reporting a problem.
- Chapter 7 describes how to handle device failure and how to troubleshoot and diagnose symptoms. It explains problem-reporting mechanisms and troubleshooting concepts.
- Appendix A describes events generated by the Farm Manager that can be used to bill customers.
- Appendix B provides a list and brief description of all command-line tools described in detail in the manual pages.

Related Publications

You might consider the following the related publications helpful:

- *N1 Provisioning Server 3.1, Blades Edition, Installation Guide*
- *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide*
- *N1 Provisioning Server 3.1, Blades Edition, Release Notes*

Accessing Sun Documentation Online

The docs.sun.com Web site enables you to access Sun technical documentation online. At this site, you can browse the archives or search for a specific book title or subject. The URL is http://docs.sun.com.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see “Buy printed documentation at http://docs.sun.com.

Using UNIX Commands

This document does not contain information about basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

For more information about UNIX command and procedures, see the following documents:

- *Solaris Handbook for Sun Peripherals*
- Online documentation for the Solaris operating environment
- Other software documentation that you received with your system

Typographic Conventions

A number of typographic conventions are used throughout this document to help you recognize special terms and instructions. These conventions are summarized in the following table:

TABLE P–1 Typographic Conventions

Convention	Description
Aa123	This typeface is used for buttons on a screen that can be clicked. Examples: 1. Click the Deployed tab to display a list of active farms. 2. Click Submit to submit a completed farm design for approval.
Aa123	This typeface is used for the following: Key words, such as terms that are defined in the text, names of books, and URLs Examples: 1. For more information, refer to the *System Administration Guide*. 2. Go to the vendor's web site, *XXX.com*, to download the information.
`Aa123`	This typeface is used for the following: ■ Code or command-line output examples ■ Directory paths ■ File names ■ Command syntax and usage Examples: 1. `1 free disks found on device 000184500535 (device id: 22001)` 2. `/opt/terraspring/sbin/disk` 3. `const char* getName() const`

TABLE P–1 Typographic Conventions *(Continued)*

Convention	Description
Aa123	This typeface is used for commands, file names, and other text strings that you must type exactly as shown. Examples: 1. **/opt/terraspring/sbin/disk -a** 2. **/opt/terraspring/sbin/device -LFt cpu -M hp -O lpr -r pes**
Aa123	This typeface, is used for placeholders that you must replace with an actual value. Actual values might be commands, file names, parameters, or other values required to run a command, change file directories, or run a script. Examples: 1. **/opt/terraspring/sbin/device -sU** *device ID* 2. **disk -l** *disk ID*

Sun Welcomes Your Comments

Sun is interested in improving documentation and welcomes your comments and suggestions. Send email with your comments to us at `docfeedback@sun.com`.

Include the part number of your document in the subject line of your email.

CHAPTER 1

N1 Provisioning Server Introduction

This chapter provides an introduction of the Infrastructure Fabric (I-Fabric) components and the N1 Provisioning Server. Later chapters of this guide describe in detail how to operate and perform day-to-day administrative tasks on the software and the I-Fabric.

This chapter discusses the following topics:

I-Fabric Overview

N1 Provisioning Server software provides a comprehensive infrastructure automation solution that enhances the visibility, and control of complex computing environments. An I-Fabric consists of various hardware components, such as blade system chassis, server blades, and load balancer blades, and the N1 Provisioning Server software. N1 Provisioning Server software controls how I-Fabric components interoperate. An I-Fabric, controlled and managed by N1 Provisioning Server software, combines computing and networking resources into a contiguous automated fabric of infrastructure that you deploy and reallocate to meet changing requirements. Resources are dedicated to a farm until returned to the common resource pool. With root access to devices, you can deploy any software or application onto the servers. Secure partitions enforced by N1 Provisioning Server software and methodologies enable you to exercise independent administrative control over each logical server farm.

Resources within an I-Fabric are accessed through the Control Center, a web-browser-based graphical user interface (GUI). This drag-and-drop interface enables you to design and deploy logical server farms. Within a design, you can define numerous characteristics, including topology, monitors, and alerts.

When you create a logical server farm design and apply it within the Control Center, N1 Provisioning Server software creates a logical description of the server farm design. The logical description is captured in the Farm Markup Language (FML), an eXtensible Markup Language (XML) dialect developed for N1 Provisioning Server software. FML enables the abstraction of design and configuration data for the deployment of actual physical resources. Abstraction of monitoring processes is captured using the Monitoring Markup Language (MML). Abstraction of the wiring configuration is captured using the Wiring Markup Language (WML).

I-Fabric Components

An I-Fabric integrates individually managed discontinuous networks, servers, and infrastructure devices into a coordinated, automated fabric. This fabric enables easy management, deployment, and redeployment of logical server farms. An I-Fabric is made up of three functional areas:

- Control plane
- Fabric layer
- Resource pool

Control Plane

The control plane provides intelligence, management, and control of an I-Fabric. N1 Provisioning Server software, which provides the intelligence that enables an I-Fabric, resides within the control plane. The control plane consists of N1 Provisioning Server software, the Control Center server software, and the associated server hardware on which the software is deployed.

Fabric Layer

The fabric layer is a highly integrated Ethernet environment based on industry-standard networking technologies. The fabric layer contains the networking infrastructure and switching fabric that ties the control plane and resource pool together.

Resource Pool

The resource pool consists of infrastructure resources, such as Sun Fire™ B1600 Blade System Chassis, load balancer blades, and server blades that function as resource pool servers for deployment to server farms. The resource pool of an I-Fabric contains all resource pool servers capable of being provisioned to a farm.

N1 Provisioning Server Configuration

The N1 Provisioning Server is set up in the following configuration:

- All N1 Provisioning Server software runs on a single server in the control plane. The image library can either reside on a separate server or on the control plane server. Optionally, you can connect a terminal server. The Control Center GUI is accessed through a separate PC.
- The resource pool configuration consists of one-blade to twelve-blade system chassis with 16 server blades on which you can deploy farms. For details on how to configure I-Fabric components see the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide*.

Logical Server Farms

The control plane, fabric layer, and resource pool work together to dynamically create logical server farms within. Logical server farms are securely allocated from the resource pool and managed by N1 Provisioning Server software. N1 Provisioning Server software creates logical server farms from the resources available within the resource pool. Logical server farms are analogous to traditional, manually built, dedicated server farms except that you can create, grow, shrink, and delete them as data structures that reside within N1 Provisioning Server software.

Secure partitions enforced by N1 Provisioning Server software and methodologies enable you to exercise independent administrative control over each logical server farm. Although you can have administrative access on all devices within a logical server farm, you cannot view, access, or modify the devices or data associated with a different logical server farm unless you have the appropriate permission.

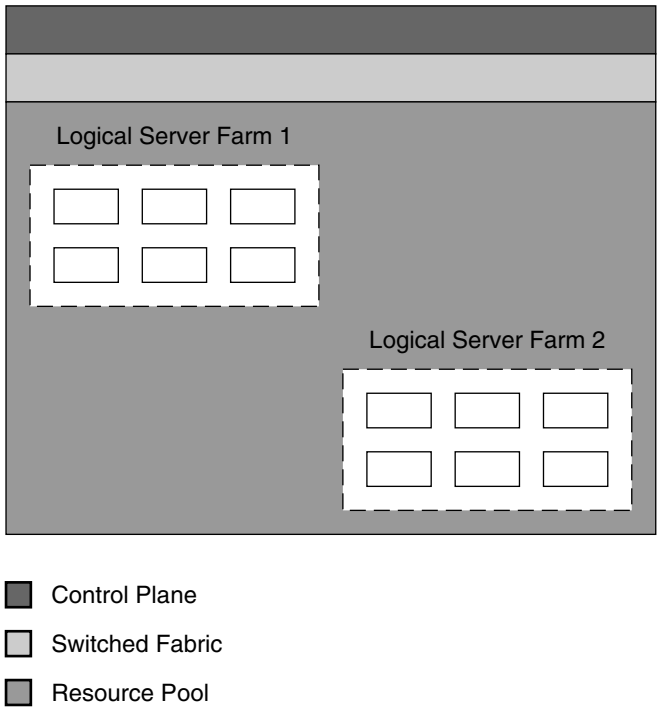

FIGURE 1–1 Logical Server Farms Provisioned Within the Resource Pool of an I-Fabric

N1 Provisioning Server Components

The N1 Provisioning Server provides a comprehensive set of infrastructure automation and management capabilities. The N1 Provisioning Server comprises two primary components: the Control Center and the control plane.

- The Control Center is the web browser-based GUI that enables design, configuration, deployment, and ongoing management of logical server farms. For more details about the Control Center, see the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide*.
- The control plane comprises the N1 Provisioning Server software, which provides the interface between the Control Center and the physical infrastructure resources, and any third-party components required to run the N1 Provisioning Server software.

Operating Environments

N1 Provisioning Server software runs on the Solaris™ 8 operating environment.

Control Center Software

The Control Center is the GUI to the N1 Provisioning Server software. You deploy and manage logical server farms through the Control Center. For more details about the Control Center, see the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide*.

N1 Provisioning Server Software

The N1 Provisioning Server software resides on the control plane server and provides the infrastructure automation services required to manage and deploy logical server farms within an I-Fabric. At a high level, the N1 Provisioning Server software manages the logical-to-physical mappings between a logical server farm and the physical resources assigned to it.

N1 Provisioning Server software comprises the following roles and software components:

- Segment manager, which controls and coordinates activities for an I-Fabric
- Farm manager, which monitors activities related to logical server farms
- Dynamic Host Configuration Protocol (DHCP) and Domain Name Server (DNS) services
- Monitoring software, which monitors the health and state of an I-Fabric and the logical server farms within it
- Storage manager client (STMC), an interface required by the farm manager to access storage functionality
- Control plane database (CPDB), which is a persistent, central repository of data
- The image server, which manages images. The image server can be any stand-alone server that supports network file server (NFS), or it can reside on the control plane server.

N1 Provisioning Server software provides six key areas of infrastructure automation services:

- Provisioning and configuration services
- Flexing services
- Software image management services
- Monitoring and messaging services
- Failure and recovery services

- Physical infrastructure management services

Provisioning and Configuration

The ability to automatically provision and configure resources within the resource pool of an I-Fabric is a core capability of the N1 Provisioning Server. The N1 Provisioning Server manages and automates the ongoing evolution of logical server farms and their initial activation. As resources are added or deleted on server farms, the N1 Provisioning Server continues to manage and automatically configure all virtual wiring as well as DHCP and DNS services.

Flexing

Flexing enables you to add or delete capacity on a farm. N1 Provisioning Server software automatically provisions and configures resources. N1 Provisioning Server software provides two types of flexing services:

- Adding and deleting individual servers within a logical server farm
- Adding and deleting server groups through a server group mechanism

Software Image Management

The N1 Provisioning Server manages software images and configurations of servers. The N1 Provisioning Server supports creation and management of two categories of images: global images and account images.

- Global images consist of the operating system (including patches and service packs), integrated N1 Provisioning Server agents, and certain customizations. Global images may also contain applications. At the very least, the customizations are modifications that make images compatible with an I-Fabric. You also may choose to customize images with account-specific software and data.
- Account images are for a particular account and consist of account-specific customizations of global images, blank disks, or application and data images.

I-Fabric Monitoring

The N1 Provisioning Server actively monitors the state and health of devices in an I-Fabric. Monitoring provides visibility of an I-Fabric and supports failover and recovery of devices or the restarting of failed processes.

Failure and Recovery of Devices in the Resource Pool

The N1 Provisioning Server automatically detects and then replaces failed devices in the resource pool, such as resource pool servers and load balancers. The failed device is replaced with a device of the same type from the available devices in the resource pool. The replacement device automatically takes on the network configuration of the failed device.

Load Balancer Failover Modes

Load balancers support the following failover modes:

- Path failover
- Device failover or high availability (HA)
- Single device failover (non-HA)

A path failover occurs when one interface to the router fails. In that case, the path on the failed interface is automatically restored on the live interface. In a high availability setup, one load balancer device is on standby as the secondary load balancer, while the primary, active load balancer device handles processing. If the primary load balancer fails, the secondary load balancer automatically takes over processing. A single device failover situation is resolved manually through the `replaceFailedDevice` tool.

Physical Infrastructure Management

As a part of the N1 Provisioning Server initialization process, the N1 Provisioning Server performs resource and wiring validation. This validation enables the CPDB to have a complete physical topology map of all resources within the I-Fabric. The wiring validation provides an automated way of confirming the physical wiring map of all devices within an I-Fabric. The integrity of the physical wiring of the resources within an I-Fabric enables the N1 Provisioning Server to successfully manage the virtual wiring of a logical server farm.

Description Languages

N1 Provisioning Server software uses description languages based on eXtensible Markup Language (XML) to create a digital blueprint of a farm's logical structure. This logical blueprint facilitates the automation of many manual tasks involved in constructing logical server farms. N1 Provisioning Server software employs the following description languages:

- Farm Markup Language (FML) - Represents the logical blueprint of a logical server farm and describes network and configuration data for physical resources associated with a logical server farm
- Monitoring Markup Language (MML) - Describes monitor deployments and configurations as you defined them within the Control Center
- Wiring Markup Language (WML) - Describes the physical wiring characteristics within an I-Fabric

I-Fabric Security

N1 Provisioning Server software and methodologies provide a multitenant infrastructure that can be securely partitioned and allocated. The benefits of multitenancy, such as resource optimization through reallocating and the ability to flex from a common resource pool, are achievable only if a rigorous security model and architecture are in place. An I-Fabric provides multiple types of security enforcement and ensures that security exists between logical server farms.

Within an I-Fabric, security is implemented at the following levels:

- Fabric layer — Ethernet and Virtual Local Area Network (VLAN) security
- Resource pool — Power cycling and memory scrubbing of resources
- Control plane — N1 Provisioning Server and Control Center security
- Password encryption at all levels

CHAPTER 2

I-Fabric Operation

This chapter describes day-to-day operational activities within an Infrastructure Fabric (I-Fabric) and how to add and remove various I-Fabric components. It also describes how to synchronize the physical and database environments. This information is useful in the event of hardware failure. Sections in this chapter include:

I-Fabric Operation Overview

You can add or remove resources (chassis and server blades) that were not part of your original I-Fabric design. Similarly, you can remove resources that are in your current I-Fabric design.

Applying Role-Based Access Control

Role-based access control (RBAC) is a security methodology used to assign root account and superuser account capabilities to specific accounts or roles based on specific task requirements. This section provides the procedure for RBAC security to the N1 Provisioning Server software and file systems.

This section does not cover general RBAC tasks, commands, and procedures. For more information on how to use RBAC see http://docs.eng.sun.com/db/doc/817-0365/6mg5vpmbd?q=Role-Based+Access+Control.

▼ To Apply RBAC Security

Steps

1. **Log in as root on the control plane server.**

2. **Back up the `/etc/security` attribute files.**

   Create a subdirectory and copy the attribute files to the subdirectory. Use the `cp` command option `-p` to preserve file permissions and ownership.

   Example:

```
cd /etc/security
ls -l *attr
-rw-r--r--   1 root     sys        42871 Nov  8  2002 audit_record_attr
-rw-r--r--   1 root     sys         5907 Nov  8  2002 auth_attr
-rw-r--r--   1 root     sys        12672 Nov  8  2002 exec_attr
-rw-r--r--   1 root     sys         4715 Nov  8  2002 prof_attr
mkdir attr-backup
cp -p *attr attr-backupls -l attr-backup
total 136
drwxr-xr-x   2 root     other        512 Jan  7 18:00 .
drwxr-xr-x   7 root     sys          512 Jan  7 18:00 ..
-rw-r--r--   1 root     sys        42871 Nov  8  2002 audit_record_attr
-rw-r--r--   1 root     sys         5907 Nov  8  2002 auth_attr
-rw-r--r--   1 root     sys        12672 Nov  8  2002 exec_attr
-rw-r--r--   1 root     sys         4715 Nov  8  2002 prof_attr
```

3. **Type the following commands to append the predefined N1 Provisioning Server security attributes to the `/etc/security` files.**

```
cd /opt/terraspring/lib/rbac/
cat n1_auth_attr >> /etc/security/auth_attr
cat n1_exec_attr >> /etc/security/exec_attr
cat n1_prof_attr >> /etc/security/prof_attr
```

4. **Create the `/n1admin` directory by typing:**

```
mkdir n1admin
```

5. **Add the `n1admin` role account.**

roleadd -d /n1admin -P"N1ADMIN" n1admin

For further information about the `roleadd` command, see the `roleadd` man page `roleadd(1M)`.

6. **Change the `n1admin` account password.**

   passwd n1admin

   You are prompted for the new password. Type the password and press Enter.

   You are prompted to reenter the new password.

   Retype the new password and press Enter.

7. **Assign user accounts to the `n1admin` security role.**

   Type the command **/opt/terraspring/lib/rbac]# usermod -R "n1admin"** *user01* where *user01*is a valid UNIX user account.

Note – The user account must exist. For further information about RBAC commands, see the http://docs.eng.sun.com/db/doc/817-0365/6mg5vpmbd?q=Role-Based+Access+Control.

---

8. **Log on to the Control Plane server using the user account.**

9. **Type `roles` to list the security roles.**

   Example:

   roles
   n1admin

10. **Verify that the user account cannot run N1 Provisioning Server commands.**

    Type **/opt/terraspring/sbin/image –l**. The response should be `image: cannot execute`. If a listing of available N1 Provisioning Server images appears, then `/opt/terraspring/sbin` file permissions have not been set correctly. Set and validate the file permissions as directed by step 2.

11. **Type `su n1admin` to log in as the `n1admin`superuser.**

12. **Verify that the user account can run N1 Provisioning Server administration commands using the `n1admin` role.**

Note – To run any administration command, you must either specify the full path for the command or be in the directory that contains the command and prefix the command with `./`.

You can do this using either of the following two methods:

- Log in as superuser `n1admin`, and type **`/opt/terraspring/sbin/image –l`** to list the available N1 Provisioning Server images. For example:

  ```
  su - n1admin
  /opt/terraspring/sbin/image –l
  ```

- Log in as superuser `n1admin` and use the `su` command `-c` option to run the `image –l` command. For example:

  ```
  su - n1admin -c "/opt/terraspring/sbin/image -l"
  ```

A list of N1 images is displayed. For example:

```
$ image -l
IMAGE_ID IMAGE_NAME               CUSTOMER         SIZE         OS       TYPE         STATE
1        solaris9u5-sun4ublade-flash __grid__      3004789248   solaris  flash        READY
LOCATION: nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
2        solaris9u5-sun4ublade-disk-image __grid__ 30000000000  solaris  disk_image   READY
 LOCATION:nfs://3001//images/master-images/solaris9u5-sun4ublade-disk-image
```

13. Verify the user ID (uid) and group ID (gid).

Type **`id`** to display the uid and gid. The uid should be `50004(n1admin)`, and the gid should be `1(other)`. For example:

```
id
uid=50004(n1admin) gid=1(other)
```

Understanding I-Fabric Component IDs

Addition and removal of I-Fabric components is based on a device ID scheme. Understanding the ID scheme will help you to effectively add and remove components within the I-Fabric. Execute the `device -l` command to obtain a typical listing of components and their IDs:

```
/opt/terraspring/sbin/device -l
```

Example output:

```
 DEVICE_ID  PARENT_ID STATUS   FARM_ID  TYPE
         2          - USED     -        rf:cisco-rf (cisco RouterFabric)
      3002          - USED     -        cpu:sun-svr-420R-idb (Sun 420R)
     50001          - USED     -        sf:stiletto-sf (Sun StilettoSwitchFabric)
     50100          - USED     103      cpu:sun-svr-blade (- -)
     50101          - USED     102      cpu:sun-svr-blade (- -)
     50102          - USED     101      cpu:sun-svr-blade (- -)
     50113          - NEW      -        cpu:sun-svr-blade (- -)
     50115          - FREE     -        cpu:sun-svr-blade (- -)
     50150      50001 USED     -        sw:stiletto-sw (- -)
     50160          - USED     -        pwr:stiletto-pwr (- -)
10 devices found.
```

The IDs of the chassis components define the chassis to which they belong. Chassis component IDs have five digits. The second and third digits represent the chassis ID. The fourth and fifth digits represent the chassis component's ID. For example, all devices with the ID 503xx belong to a chassis with the ID 50300.

For each chassis , the IDs 5yy00 through 5yy15 are reserved for server blades, IDs 5yy50 and 5yy51 are reserved for switches, and ID 5yy60 is reserved for the system and switch controller (SSC). For example, for the fifth chassis the IDs in the following table would be assigned:

ID	Device
50500 - 50515	Blades
50550 and 50551	Switch
50560	System and switch controller

Caution – IDs 50000 through 59999 are reserved for chassis components. Do not assign any other devices to those IDs.

Note – The chassis itself is not listed in the CPDB as a device. In that sense, a chassis has a "virtual" ID that is the same as the ID of the `S0` blade if one exists.

Changing the System Password

For security reasons, reset the default N1 Provisioning Server system password. The default password for access to the N1 Provisioning Server is

```
u8 (!xq] [~be4
```

To change the system password, run the `resetpasswd -s` command at the N1 Provisioning Server command line.

Configuring Load Balancers in High Availability Pairs

The N1 Provisioning Server software supports high availability (HA) load balancer pairs. Configuring load balancers for high availability mode involves two steps:

- Creating a new device type

- Creating a load balancer high availability pair

▼ How to create a new device type

Steps

1. **Type `devtype -l |grep lb` to get the ID of the device type.**
   Example

```
6       lb        -       -
                  0
7       halb                -       -
                  0
```

2. **To create a device type for a high availability load balancer pair run:**
   `devtype -c -n SFB10LB-ha-pair "SFB10LB HA Pair" -b -q none`
   Output example:

```
Added SFB10LB ha pair, id 25013
```

3. **Set the HAL class for the device type you just created by running**
   `devtype -s` *SFB10LB ha pair ID* **`-H com.terraspring.drivers.sun.HighAvailabilitySFB10LB`**

▼ How to create a high availability load balancer device

Steps

1. **To create an instance of a high availability load balancer device run**
   `device -c` *SFB10LB ha pair ID* **`-t SFB10LB ha pair`**

2. **To get a list of available load balancers run**
   `device -l | grep lb`
   Example output:

```
50103          - FREE   -  lb:sunfireb10n-299-blade
(sunw sunfireb10n)
       50115          - FREE   -  lb:sunfireb10n-299-blade
(sunw sunfireb10n)
```

3. **Set the parent for two available load balancer device IDs by running**
   `device -s -P` *SFB10LB ha pair ID SFB10LB device ID*
   Example:
   `device -s -P 50103`
   `device -s -P` *SFB10LB ha pair ID lSFB10LB ha pair device ID*
   Example:
   `device -s -P 50115`

4. **Set the the HA load balancer pair to a state of `FREE` by typing**
   `device -s -F` *SFB10LB ha pair device ID*

Configuring Sun Fire B1600 Blade Chassis Switch Connections

The blade system chassis switch may have physical connections to the following:

- External router or switch
- External L2 switch ports
- Switch ports of other chassis
- VLAN-aware NIC of provisioning servers

The presence or absence of these connections and the number of ports used depends on the network topology implemented. Ensure that the duplex and speed on both ends of the connection is properly auto-negotiated. Otherwise, network performance might be adversely affected. Also, if multiple ports are used for improving bandwidth between switches or switch—router connections, enable link aggregation on these ports.

Before setting VLAN rules to ports, VLANs need to exist in the switch database. To do this, execute the following sequence of commands:

```
Vty-0# configure
Vty-0(config)# vlan database
Vty-0(config-vlan)# vlan 1 name DefaultVlan media ethernet state active
Vty-0(config-vlan)# vlan 2 name MgtVlan media ethernet state active
Vty-0(config-vlan)# vlan 4 name IdleVlan media ethernet state active
Vty-0(config-vlan)# vlan 8 name ImageVlan media ethernet state active
Vty-0(config-vlan)# vlan 9 name ControlVlan media ethernet state active
Vty-0(config-vlan)# vlan 10 name VLAN10 media ethernet state active
Vty-0(config-vlan)# vlan 11 name VLAN11 media ethernet state active
```

When done creating all VLANs, press Ctrl Z or end to leave the configuration mode.

The following describes the configuration steps for ports involved in these connections.

▼ To Configure Chassis Switch Ports

Steps

1. **Connect ports to the external router.**

   Configure these ports as trunk ports that allow tagged packets. The blade system chassis switches require explicit enabling of VLANs allowed on trunk connections. Configure all VLANs except the switch management VLAN (factory default set to VLAN 2) to pass through.

   On the blade system chassis switch, the set of commands used to achieve this is as follows:

```
Vty-0(config)#interface ethernet NETP0
Vty-0(config-if)# description External RJ-45 connector NETP0
Vty-0(config-if)# switchport allowed vlan add 1 untagged
Vty-0(config-if)# switchport native vlan 1
Vty-0(config-if)# switchport mode trunk
Vty-0(config-if)# switchport allowed vlan add 10-1024 tagged
Vty-0(config-if)# switchport forbidden vlan add 2
Vty-0(config-if)#
```

In the above example VLANs 10-1024 are allowed on tagged port `NETP0`. All other VLANs are forbidden.

2. **Connect ports to the external switch ports.**

   These connections may be of two types: trunk connections to the external switch using the `NetP7` port and control plane connections using the `NetMgt` port.

   For trunk connections using the `NetP7` port, the configuration requirements and commands to do this are the same as in the previous step.

   For control plane connections using the `NetMgmt` port, the `NetMgmt` port must be in the management VLAN of the switch (by default, VLAN 2). This is a factory default setting, but verify this during the configuration phase.

   interface ethernet NETMGT

   switchport allowed vlan add 2 untagged

   switchport native vlan 2

   switchport allowed vlan remove 1

   switchport forbidden vlan add 1

3. **When done creating all VLANs, press Ctrl Z or end to leave the configuration mode.**

4. **For ports connected to switch ports of other chassis, the requirements and commands are the same as in step 1.**

5. **For ports connected to VLAN-aware NIC of the provisioning server, the requirements and commands are the same as in step 3.**

Configuring a Router for External Access to a Farm

After a farm is activated, you might want to configure an external router to enable external access to a farm and its components. Because the external router is considered external to the I-Fabric, you need to configure it manually.

After the farm is activated, execute the command `fmrouter -i` *farm_ID* to get the VLAN and IP address information for the farm. Three IP addresses are allocated for the external router configuration: one as the gateway IP address and the other two for up to two routers you can use. The VLAN for the farm on its external network is also displayed. For the farm to be accessible from an external network, enable a route to the gateway IP address on the specified VLAN.

Adding and Deleting Chassis

Depending on your computing needs, you might want to add and remove chassis.

▼ To Add a Chassis

This section describes how to add and remove chassis dynamically within an I-Fabric.

Steps

1. **Type `/opt/terraspring/sbin/discoversfb1600`**

   Running this tool launches a command-line UI that prompts you for chassis information.

2. **When prompted for the IP address of the SSC of the chassis to be discovered, type each address on a separate r line.**

3. **When prompted to accept the list of IP addresses, type `Yes` to accept or `No` to edit the values.**

4. **At the prompt, type your system controller username.**

5. **At the prompt, type your system controller password.**

   You are then asked whether your system controller username and password are the same as your switch username and password.

6. **Indicate whether the user names and passwords are the same.**

   - **If these user names and passwords are different, type `no` and provide the switch login information.**
     - **If the user names and passwords are the same, type `yes`.**

7. **Continue to provide the information as prompted by the UI.**

   The `discoversfb1600` command then attempts to discover the specified chassis. The`discoversfb1600` command automatically adds the chassis's components, such as the server blades, SSCs, and their interconnections. The discovery process takes a few minutes to execute for each shelf.

   Example:

```
/opt/terraspring/sbin/discoversfb1600
SFB1600 Discovery
Please enter a list of SFB1600 System Controller IP Addresses
End the list with an empty line
IP Address of Blade System Controller #1: 10.5.132.65
IP Address of Blade System Controller #2:

Checking Blade System Chassis...
1: 10.5.132.65 : OK
```

```
Checking Blade System Chassis entries in database....
1: 10.5.132.65 : OK

Using the following IP addresses:
1: 10.5.132.65
Do you want to configure this shelf [yes/no]? yes

Enter common user name: admin
Enter common password:
Are the switch user name and password the same as the management user name
and password? [yY/nN]y

Configuring devices......
Enter default gateway: 10.5.132.1
Enter netmask: 255.255.255.0
Please enter VLAN ranges.
Enter VLAN range lower bound (Press ENTER if done) : 10
Enter VLAN range upper bound (Press ENTER if done) : 255
Enter VLAN range lower bound (Press ENTER if done) :
How many switches has the shelf with controller 10.5.132.65 got? 2
Enter switch IP for switch 1: 10.5.132.68
Enter the SSC number corresponding to this switch [ssc0/ssc1]: ssc0
Configuring switch with IP 10.5.132.68...
done!

Enter switch IP for switch 2: 10.5.132.69
Enter the SSC number corresponding to this switch [ssc0/ssc1]: ssc1

Configuring switch with IP 10.5.132.69...
done!

Configuring blades with controller IP 10.5.132.65...
Warning: assuming SC and Switch have same login username.
Warning: assuming SC and Switch have same login password.
Discovering shelf with IP: 10.5.132.65. Please Wait...
Done!
Connected to jdbc:oracle:thin:@idb:1521:tsprdb as XCF
Processing WML file /var/tmp/devices.20031125.201509.wml ...
Segment name:
Adding device type: sunfireb100x-238-blade
Adding port: p0
Adding port: s0
Adding port: ide
Adding adapter: eth type: ge
Adding port: eth0
Adding port: eth1
Adding local disk: local1 size: 30000000000
Adding disk address, port: ide target: 0 lun: 0
**WARNING** Skipping device type 'sun-b1600-shelf'. It is already defined.
Adding device sun-b1600-shelf:50299
    with attributes: {}
**WARNING** Updating device '50001'. It is already defined.
Adding device sun-b1600-sw:50250
    with attributes: {rack-id=50299, netmask=255.255.255.0,
ipaddress=10.5.132.68, gateway=10.5.132.1}
```

```
Adding device sun-b1600-sw:50251
    with attributes: {rack-id=50299, netmask=255.255.255.0,
ipaddress=10.5.132.69, gateway=10.5.132.1}
Adding device sun-b1600-pwr:50260
    with attributes: {rack-id=50299, netmask=255.255.255.0,
ipaddress=10.5.132.65, gateway=10.5.132.1}Adding device sunfireb100s-95-blade:50200
    with attributes: {rack-id=50299, slot-num=0, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50201
    with attributes: {rack-id=50299, slot-num=1, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50211
    with attributes: {rack-id=50299, slot-num=11, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50202
    with attributes: {rack-id=50299, slot-num=2, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50203
    with attributes: {rack-id=50299, slot-num=3, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50204
    with attributes: {rack-id=50299, slot-num=4, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50205
    with attributes: {rack-id=50299, slot-num=5, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50206
    with attributes: {rack-id=50299, slot-num=6, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50207
    with attributes: {rack-id=50299, slot-num=7, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100s-95-blade:50208
    with attributes: {rack-id=50299, slot-num=8, role=PES}
Adding device port: eth0
Adding device port: eth1
Adding device sunfireb100x-238-blade:50209
    with attributes: {rack-id=50299, slot-num=9, role=PES}
Adding device port: eth0
Adding device port: eth1Adding connection from 50250/snp0 to 50200/eth0
Adding connection from 50251/snp0 to 50200/eth1
Adding connection from 50260/0 to 50200/p0
Adding connection from 50260/s0 to 50200/s0
Adding connection from 50250/snp1 to 50201/eth0
Adding connection from 50251/snp1 to 50201/eth1
Adding connection from 50260/1 to 50201/p0
```

```
Adding connection from 50260/s1 to 50201/s0
Adding connection from 50250/snp11 to 50211/eth0
Adding connection from 50251/snp11 to 50211/eth1
Adding connection from 50260/11 to 50211/p0
Adding connection from 50260/s11 to 50211/s0
Adding connection from 50250/snp2 to 50202/eth0
Adding connection from 50251/snp2 to 50202/eth1
Adding connection from 50260/2 to 50202/p0
Adding connection from 50260/s2 to 50202/s0
Adding connection from 50250/snp3 to 50203/eth0
Adding connection from 50251/snp3 to 50203/eth1
Adding connection from 50260/3 to 50203/p0
Adding connection from 50260/s3 to 50203/s0
Adding connection from 50250/snp4 to 50204/eth0
Adding connection from 50251/snp4 to 50204/eth1
Adding connection from 50260/4 to 50204/p0
Adding connection from 50260/s4 to 50204/s0
Adding connection from 50250/snp5 to 50205/eth0
Adding connection from 50251/snp5 to 50205/eth1
Adding connection from 50260/5 to 50205/p0
Adding connection from 50260/s5 to 50205/s0
Adding connection from 50250/snp6 to 50206/eth0
Adding connection from 50251/snp6 to 50206/eth1
Adding connection from 50260/6 to 50206/p0
Adding connection from 50260/s6 to 50206/s0
Adding connection from 50250/snp7 to 50207/eth0
Adding connection from 50251/snp7 to 50207/eth1
Adding connection from 50260/7 to 50207/p0
Adding connection from 50260/s7 to 50207/s0
Adding connection from 50250/snp8 to 50208/eth0
Adding connection from 50251/snp8 to 50208/eth1
Adding connection from 50260/8 to 50208/p0
Adding connection from 50260/s8 to 50208/s0
Adding connection from 50250/snp9 to 50209/eth0
Adding connection from 50251/snp9 to 50209/eth1
Adding connection from 50260/9 to 50209/p0
Adding connection from 50260/s9 to 50209/s0Marking device '50200' as free.
Marking device '50201' as free.
Marking device '50202' as free.
Marking device '50203' as free.
Marking device '50204' as free.
Marking device '50205' as free.
Marking device '50206' as free.
Marking device '50207' as free.
Marking device '50208' as free.
Marking device '50209' as free.
Marking device '50211' as free.
Marking device '50250' as used.
Marking device '50251' as used.
Marking device '50260' as used.
Marking device '50299' as used.
File '/var/tmp/devices.20031125.201509.wml' processed successfully

Clz6pK2b6qw= 50260
Password set : 50250
```

```
Password set : 50251
Password set : 50260

Validating the blades. This may take upto 30 minutes.
You can choose to skip this step.
Do you want to continue [yes/no]? no
```

8. **Type the `device -l` command to verify that the chassis has been added. See the example in"Understanding I-Fabric Component IDs" on page 26 for details on how to use this command.**

   See the example in"Understanding I-Fabric Component IDs" on page 26 for details on how to use this command.

▼ To Delete Chassis

This procedure describes how to dynamically delete chassis from an I-Fabric.

Before You Begin

All server blade components of the chassis to be deleted must be in `FREE` or `NEW` state. You can change device states by using the `device` command. See the `deletesfb1600` man page for details on how to use that command.

Steps

1. **Type the command:**

   `/opt/terraspring/sbin/deletesfb1600 -d` *chassis-ID*

2. **If the chassis is in a valid state for deletion, you are prompted to confirm deletion of the chassis and all its subcomponents. Type `y` to delete the chassis an its subcomponents.**

3. **Type the `device -l` command to verify that the chassis has been deleted.**

   See the example in "Understanding I-Fabric Component IDs" on page 26 for details on how to use this command.

Example 2–1 Sample Actions for Deleting a Chassis

```
/opt/terraspring/sbin/deletesfb1600 -d 50900
Chassis id 50900 components are:
        DEVICE_ID       STATUS  FARM_ID TYPE
        50902           FREE    -       cpu:sun-svr-blade
        50951           USED    -       sw:stiletto-sw
Chassis  id 50900is in a valid state for DELETION.
Delete chassis  id 50900 and its subcomponents from  the database (Yy/Nn)?y

Deleted chassis id 50900 and all its subcomponents from the database.
```

Adding, Moving, and Deleting Chassis Components

Depending on your computing needs, you might want to add, move, or delete chassis components, such as server blades. This section describes how to add, move and delete chassis components dynamically within an I-Fabric.

Synchronizing the Physical and Database Environment

After you have made changes to the physical state of the chassis, for example, removing or adding a blade, use the `/opt/terraspring/sbin/shelfsync` command to synchronize the CPDB with the physical state of the chassis. Synchronization is preceded by the task of discovering the device to be synchronized. This task is common to adding, removing, moving, and exchanging server blades. Therefore the discovery process and the related example is only shown once in the "Adding, Moving, and Deleting Chassis Components" on page 36 section. Subsequent synchronization sections show only the recommended action and synchronization processes and related examples.

▼ To Add Server Blades and Switch and System Controllers

Steps

1. **Begin synchronization of the CPDB with the discovery of the chassis's physical state by entering the server blade's IP address or device ID when prompted to do so.**

   a. **At the prompt, type your system controller username.**

   b. **At the prompt, type your system controller password.**

      You are asked whether the system controller user name and password are the same as the switch user name and password.

   c. **Indicate whether the user names and passwords are the same.**

      - **If the user names and passwords are different, type `no` and provide the switch login information.**
      - **If the user names and passwords are the same, type `yes`.**

      When the discovery process has completed successfully, the actual synchronization process begins.

2. **When prompted to either type `q` for exiting or the line number for the recommended action, enter the line number.**

3. **When asked whether to add the device, type `y`**

4. **When all recommended actions are completed or you do not want to complete the remaining recommended actions, type `q` to exit.**

5. **When prompted to either resynchronize the shelf with its saved state, to synchronize another chassis, or to exit, make the appropriate selection for synchronizing all the chassis you want to have synchronized.**

6. **Type 3 at the prompt to exit.**

Example 2–2 Sample Actions for Adding Server Blades and SSCs

```
/opt/terraspring/sbin/shelfsync
Begining synchronization of the I-Fabric's database with an SFB1600's
physical state.

Enter SFB1600 management IP address or the device ID: 10.5.132.65
Database query in progress. Please wait.
Found: ipaddr="10.5.132.65" shelfid="50200" powerid="50260"
Checking network connectivity to 10.5.132.65....Done. Network connectivity:
OK


Attempting discovery of the specified SFB1600 device.
Management interface user name: admin
Password:
Are the switch user name and password the same as the management user name
and password? [yY/nN]y

Is there any new switch added to the shelf that you want to bench
configure? [yY/nN]n
Device discovery for 10.5.132.65 in progress. Please wait.

Writing WML output to /var/tmp/deviceWML.29789
Device discovery succeeded.
Fetching database state.....Please Wait.....Done!!

Recommended actions:
                Recommended
#  Cause        Action     Identification               Conflicting IDs
-- -----------  ------- ---------------------- ---------------
 0 Discovered    ADD    50211 00:03:BA:26:91:AD
 1 Discovered    ADD    50251 sw_50251_SSC1
* = action has been run, entering -y will suppress acknowledgement prompts
Enter line number(s) from above or q for quit: 1
ADD device 50251 (Identification: sw_50251_SSC1)? [y/n] y
Adding device 50251 by running /var/tmp/add.withtype.50251:

&lt;segment&gt;
<device-type name="sun-b1600-shelf" parent="bls" desc="Sun B1600 Shelf">
   </device-type>
```

```
...
Processing WML file /var/tmp/add.withtype.50251 ...
**WARNING** Skipping device type 'sun-b1600-shelf'. It is already defined.
**WARNING** Updating device '50001'. It is already defined.
File '/var/tmp/add.withtype.50251' processed successfully
Password set : 50251

Recommended actions:
                Recommended
#  Cause        Action     Identification              Conflicting IDs
-- -----------  ------- ---------------------- ---------------
 0 Discovered    ADD    50250 sw_50250_SSC0
 1 Discovered   *ADD    50251 sw_50251_SSC1
* = action has been run, entering -y will suppress acknowledgement prompts
Enter line number(s) from above or q for quit: 0
ADD device 50211 (Identification: 00:03:BA:26:91:AD? [y/n] y
Adding device 50211 by running /var/tmp/add.withtype.50211:
```

<segment>

```
<device-type name="sun-b1600-shelf" parent="bls" desc="Sun B1600 Shelf">
</device-type> ...
Processing WML file /var/tmp/add.withtype.50211 ...
**WARNING** Skipping device type 'sun-b1600-shelf'. It is already defined.
File '/var/tmp/add.withtype.50211' processed successfully

Recommended actions:
                Recommended
#  Cause        Action     Identification              Conflicting IDs
-- -----------  ------- ---------------------- ---------------
 0 Discovered   *ADD    50211 00:03:BA:26:91:AD
 1 Discovered   *ADD    50251 sw_50251_SSC1
* = action has been run, entering -y will suppress acknowledgement prompts
Enter line number(s) from above or q for quit: q
What do you want to do next:
#1. Re-sync the current SFB1600 device with its saved state.
#2. Synchronize another SFB1600 device with its saved state.
#3. Exit.
Make selection (1-3): 3
bash-2.03#
```

▼ To Delete Server Blades and Switch and System Controllers

The following steps describe how to delete server blades and SSCs in an I-Fabric that are not being used in a farm.

Before You Begin Before physically removing server blades and SSCs, run the `removefru` *blade* command from the system controller. See the Sun Fire™ B1600 Blade System Chassis documentation for details on how to use this command.

All blades to be deleted must be in `FREE` or `NEW` state. You can change device states by executing the `device` command. See the `device` man page for details on how to use that command.

Caution – Do not remove the `SSC0` component, otherwise N1 Provisioning Server software will fail.

Steps

1. **When prompted to either type `q` for exiting or type the line number for the recommended action, type the line number to delete the device.**
2. **When prompted whether to delete the device, type `y`.**
   You are prompted to confirm the deletion.
3. **Type `y` again.**
4. **When all recommended actions are completed or you do not want to complete the remaining recommended actions, type `q` to exit.**
   You are asked whether you want to resynchronize the chassis with its saved state, synchronize another chassis, or exit.
5. **Make the appropriate selection for synchronizing all the chassis you want to have synchronized.**
6. **Type `3` at the prompt to exit.**

Example 2–3 Sample Actions for Deleting server Blades and SSCs

```
Recommended actions:
                Recommended
#  Cause        Action     Identification       Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: Removed       DELETE 50110 00:03:BA:26:91:AD
1: Removed       DELETE 50151 sw_50151_SSC1
* = action has been run
Enter line number from above or q for quit: 0
DELETE device 50110 (Identification 00:03:BA:26:91:AD)? [y/n] y
/opt/terraspring/sbin/device -d 50110
Are you sure to delete (y/n)? y
Deleted device 50110

Recommended actions:
                Recommended
#  Cause        Action     Identification       Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: Removed      *DELETE 50110 00:03:BA:26:91:AD
1: Removed       DELETE 50151 sw_50151_SSC1
* = action has been run
Enter line number from above or q for quit: 1
```

```
DELETE device 50151 (Identification sw_50151_SSC1)? [y/n] y
/opt/terraspring/sbin/device -d 50151
Are you sure to delete (y/n)? y
Deleted device 50151

Recommended actions:
                Recommended
#  Cause        Action     Identification               Conflicting IDs
-- -----------  ------- ---------------------- ---------------
 0 Removed      *DELETE 50110 00:03:BA:2A:14:4A
 0 Removed      *DELETE 50151 sw_50151_SSC1
* = action has been run, entering -y will suppress acknowledgement prompts
Enter line number(s) from above or q for quit: q
What do you want to do next:
#1. Re-sync the current SFB1600 device with its saved state.
#2. Synchronize another SFB1600 device with its saved state.
#3. Exit.
Make selection (1-3): 3
[root@idb1:2]#
```

▼ To Move a Server Blade From One Slot to Another

This scenario describes moving a server blade from one slot to another one within a chassis.

Caution – In this procedure, the order in which tasks are performed is important, otherwise a conflict might occur if the server blade removed from the chassis is still in the CPDB. Ensure that you delete the server blade first from its current slot, then add it to the new one. See the example following the steps for details.

Before You Begin

Before physically removing server blades, execute the `removefru` *blade* command from the system controller. See the Sun Fire B1600 Blade System Chassis documentation for details on how to use this command.

All server blades to be deleted must be in `FREE` or `NEW` state. You can change device states by executing the `device` command. See the `device` man page for details on how to use that command.

Steps

1. **When prompted to either type `q` for exiting or the line number for the recommended action, type the line number.**
2. **When prompted whether to delete or add a device, type `y`.**

   You are prompted to confirm the deletion.
3. **Type `y` again.**

Note – Do `delete` operations first to avoid getting error messages.

Example 2–4 Sample Actions for Moving a Blade Between Slots

In the following example, a server blade was physically moved from slot S10 to S14 of a shelf. The `shelfsync` discovery process recommended two actions: an addition to slot S14 (device ID 50114) and a deletion from slot S10 (device ID 50110).

```
Recommended actions:
                Recommended
#  Cause        Action       Identification      Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT      ADD    50114 00:03:BA:26:91:F5 50110
1: CONFLICT      DELETE 50110 00:03:BA:26:91:F5 50114
* = action has been run
Enter line number from above or q for quit: 1
DELETE device 50110 (Identification 00:03:BA:26:91:F5)? [y/n] y
/opt/terraspring/sbin/device -d 50110
Are you sure to delete (y/n)? y
Deleted device 50110

Recommended actions:
                Recommended
#  Cause        Action     Identification       Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT      ADD    50114 00:03:BA:26:91:F5 50110
1: CONFLICT     *DELETE 50110 00:03:BA:26:91:F5 50114
* = action has been run
Enter line number from above or q for quit: 0
ADD device 50114 (Identification=oo:03:BA:26:91:F5) [y/n] y
Adding device 50114 by running /var/tmp/add.withtype.50114:
&lt;segment&gt;
,device-type name="sun-b1600-shelf" parent="bls" desc="Sun B1600 Shelf">
</device-type>

...
Processing WML file /var/tmp/add.withtype.50114
**WARNING** Skipping device type 'sun-b1600-shelf'. It is already defined.
File '/var/tmp/add.withtype.50114' processed successfully

Recommended actions:
                Recommended
#  Cause        Action     Identification        Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT     *ADD    50114 00:03:BA:26:91:F5 50110
1: CONFLICT     *DELETE 50110 00:03:BA:26:91:F5 50114
* = action has been run
Enter line number from above or q for quit: q

What do you want to do next:
#1. Re-sync the current SFB1600 device with its saved state.
#2. Synchronize another SFB1600 device with its saved state.
```

```
#3. Exit.
Make selection (1-3): 3
[root@idb1:2]#
```

▼ To Exchange Server Blades Within a Chassis

The following scenario describes the exchange of two server blades within a chassis.

Caution – In this procedure, the order in which tasks are performed is important, otherwise a conflict may occur if the server blade removed from the chassis is still in the CPDB. Ensure that you delete the server blade first from its current slot, then add it to the new one. See the example following the steps for details.

Before You Begin

Before physically removing server blades, execute the `removefru` *blade* command from the power controller. See the Sun Fire B1600 Blade System Chassis documentation for details on how to use this command.

All server blades to be deleted must be in `FREE` or `NEW` state. You can change device states by executing the `device` command. See the `delete` man page for details on how to use that command.

Steps

1. **When prompted to either type `q` for exiting or the line number for the recommended action, type the line number.**
2. **When prompted whether to delete the device, type `y`.**
   You are prompted to confirm the deletion.
3. **Type `y` again.**

Example 2–5 Sample Action for Exchanging Server Blades in a Chassis

```
Recommended actions:
                Recommended
#  Cause        Action     Identification       Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT      DELETE 50114 00:03:BA:26:91:F5 50110
1: CONFLICT      ADD    50110 00:03:BA:26:91:F5 50114
2: CONFLICT      ADD    50114 00:03:BA:26:91:AD 50110
3: CONFLICT      DELETE 50110 00:03:BA:26:91:AD 50114
* = action has been run
Enter line number from above or q for quit: 0
DELETE device 50114 (MAC address 00:03:BA:26:91:F5)? [y/n] y
/opt/terraspring/sbin/device -d 50114
Are you sure to delete (y/n)? y
Deleted device 50114
```

```
Recommended actions:
                Recommended
#  Cause        Action     Identification               Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT     *DELETE 50114 00:03:BA:26:91:F5 50110
1: CONFLICT      ADD    50110 00:03:BA:26:91:F5 50114
2: CONFLICT      ADD    50114 00:03:BA:26:91:AD 50110
3: CONFLICT      DELETE 50110 00:03:BA:26:91:AD 50114
* = action has been run
Enter line number from above or q for quit: 3
DELETE device 50110 (Identification 00:03:BA:26:91:AD)? [y/n] y
/opt/terraspring/sbin/device -d 50114
Are you sure to delete (y/n)? y
Deleted device 50110

Recommended actions:
                Recommended
#  Cause        Action     Identification               Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT     *DELETE 50114 00:03:BA:26:91:F5 50110
1: CONFLICT      ADD    50110 00:03:BA:26:91:F5 50114
2: CONFLICT      ADD    50114 00:03:BA:26:91:AD 50110
3: CONFLICT     *DELETE 50110 00:03:BA:26:91:AD 50114
* = action has been run
Enter line number from above or q for quit: 1
ADD device 50110 (Identification=00:03:BA:26:91:F5)? [y/n] y
Adding device 50110 by running /var/tmp/add.withtype.50110:
&lt;segment&gt;
,device-type name="sun-b1600-shelf" parent="bls" desc="Sun B1600 Shelf">
</device-type>

...
Processing WML file /var/tmp/add.withtype.50110...
File '/var/tmp/add.withtype.50110' processed successfully

Recommended actions:
                Recommended
#  Cause        Action     Identification       Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT     *DELETE 50114 00:03:BA:26:91:F5 50110
1: CONFLICT     *ADD    50110 00:03:BA:26:91:F5 50114
2: CONFLICT      ADD    50114 00:03:BA:26:91:AD 50110
3: CONFLICT     *DELETE 50110 00:03:BA:26:91:AD 50114
* = action has been run
Enter line number from above or q for quit: 2
ADD device 50114 (Identification=00:03:BA:26:91:AD)? [y/n] y
Adding device 50114 by running /var/tmp/add.withtype.50114:
&lt;segment&gt;
,device-type name="sun-b1600-shelf" parent="bls" desc="Sun B1600 Shelf">
</device-type>

...
Processing WML file /var/tmp/add.withtype.50114...
File '/var/tmp/add.withtype.50114' processed successfully
```

```
Recommended actions:
                Recommended
#  Cause        Action     Identification               Conflicting IDs
-- -----------  ------- ---------------------- ---------------
0: CONFLICT     *DELETE 50114 00:03:BA:26:91:F5 50110
1: CONFLICT     *ADD    50110 00:03:BA:26:91:F5 50114
2: CONFLICT     *ADD    50114 00:03:BA:26:91:AD 50110
3: CONFLICT     *DELETE 50110 00:03:BA:26:91:AD 50114
* = action has been run
Enter line number from above or q for quit: q
What do you want to do next:
#1. Re-sync the current SFB1600 device with its saved state.
#2. Synchronize another SFB1600 device with its saved state.
#3. Exit.
Make selection (1-3): 3
[root@idb1:2]#
```

Adding a Replacement Server Blade to a Farm

If you exchanged a server blade that was assigned to a farm and you want the newly placed server blade to be assigned to that farm, you need to add the new blade to the farm.

▼ To Add a Replacement Server Blade to a Farm

Steps

1. **Delete the replaceFailedDevice request in the database queue.**
   a. **On the control plane server, type `request -l` to get the request ID.**
   b. **Type `request -d` *request-ID* to delete the request.**
2. **Replace the failed server blade with a new one.**
3. **Get the media access control (MAC) address for the new blade and update the CPDB by typing:**
   `device -s` *device-ID* `-o` *port-name mac-address*
4. **Type `replacedevice -i`**
5. **Logon to the Control Center and submit a request for an update using the Control Center Editor.**
   See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

Replacing Switch and System Controllers

You can have dual SSCs in a chassis and configure them for automatic failover. For details on how to configure SSCs, see the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide*.

The naming convention for dual SSCs is `SSC0` and `SSC1`.

Note – An `SSC0` must always be present for N1 Provisioning Server software to function properly. Therefore, if `SSC0` fails, replace it immediately.

▼ To Replace an SSC

Before You Begin

Before replacing, removing, or adding an SSC, be sure to deactivate all farms associated with that SSC. You can do so either from the Control Center or by running the `farm -d` *farm-ID* command from the N1 Provisioning Server command line.

Steps

1. **Before physically removing the SSC1, type `removefru SSC1` from the power controller.**

   See the Sun Fire B1600 Blade System Chassis documentation for details on how to use this command.

2. **Replace the physical component.**

3. **Type the `shelfsync` command on the control plane server to synchronize the physical state of the chassis with the database. See "Synchronizing the Physical and Database Environment" on page 36 for details.**

4. **Type the `switchsync` command to synchronize the configuration of the switch with the configuration stored in the database for that switch.**

Example 2–6 switchsync Command

```
bash-2.03# switchsync -d 50151
Bench configuring the switch...It may take some time...
done.
Loading switch configuration...
done.
bash-2.03#
```

Adding and Removing SSCs

If you started with one SSC and want to add another SSC, use the same procedure as described in "Synchronizing the Physical and Database Environment" on page 36.

If you started with dual SSCs and want to remove one, remove the `SSC1` component.

Caution – Do not remove the `SSC0` component, otherwise the N1 Provisioning Server software will fail.

▼ To Remove SSC1

Steps

1. **Before physically removing the SSC1, type `removefru SSC1` from the power controller.**

   See the Sun Fire B1600 Blade System Chassis documentation for details on how to use this command.

2. **Remove the physical component.**
3. **At the `SC` prompt, type the command `show SC` to see which SSC is active.**
4. **If SSC1 is active, run the `setfailover` command to make SSC0 active.**
5. **Type `shelfsync` to synchronize the physical state of the chassis with the database.**

   See "Synchronizing the Physical and Database Environment" on page 36for details.

6. **Type `switchsync` to synchronize the configuration of the switch with the configuration stored in the database for that switch.**

Replacing Failed Resources in the Resource Pool

This section describes the procedures for replacing failed resource pool servers.

▼ To Replace Failed Resource Pool Servers

This procedure describes how to replace failed resource pool servers.

Steps

1. **Verify that the device is in a `FAILED` state in the CPDB by typing :**

   `/opt/terraspring/sbin/device -l` *device-ID*

2. **If the device is not in a `FAILED` state, manually fail it by typing:**

   `/opt/terraspring/sbin/device -sB` *device-ID*

3. **Configure an identical replacement device.**

   a. **Collect the interface and port address information, such as Media Access Control (MAC) address for Ethernet ports.**

   b. **Configure the device.**

   c. **For details on how to configure the device see the** *N1 Provisioning Server 3.1, Blades Edition, Installation Guide***.**

      Verify that the firmware is the supported version.

4. **Remove the failed device from the I-Fabric.**

5. **Install the replacement or repaired device into the I-Fabric using existing cables and connectors.**

6. **Label the device to match the failed device.**

7. **Update the CPDB with the information collected in step 3, replacing the old port address information with the new one.**

8. **Change the state of the replacement device in the CPDB to `FREE` by typing:**

   `/opt/terraspring/sbin/device -sF` *device-ID*

▼ To Scrub Disks

Scrubbing disks before using them in an I-Fabric is optional. However, doing so helps to speed up snapshot time and reduces the size of the compressed image. Scrubbing disks also provides additional security. Perform the following steps to scrub all disks before using them in an I-Fabric:

Steps

1. **Gather the disk IDs of all newly added disks by typing:**

   `disk -l|grep local|grep NEW`

2. **Scrub each disk by typing:**

   `diskscrub -d` *disk-ID*

3. **Set the state of the disk to `FREE` by typing:**

   `disk -s FREE` *disk ID*

Adding and Removing Unmanaged Ethernet Devices

You might have Ethernet-capable devices that you want to include in your I-Fabric that are not part of your original I-Fabric configuration. Such devices can include backhaul devices, routers, network attached storage (NAS) devices, printers, and so on. Such devices are not fully managed by N1 Provisioning Server software.

Once you've added the devices to your I-Fabric, the devices are listed when you use the N1 Provisioning Server tools.

Note – After the database is updated, the devices may not be physically wired to the I-Fabric yet. The database is merely a placeholder.

Before you can add unmanaged Ethernet devices, you need to create device types for unmanaged Ethernet devices in the control plane database. You create each device type only once. You can then add multiple unmanaged Ethernet devices to each device type you've created.

▼ To Create Unmanaged Ethernet Device Types

Steps

1. **Create a device type that extends the backhaul type on the control plane using the following command:**

   `devtype -c -n` *devtype-name devtype-description* **`-b`** *backhaul-devtype-ID* **`-a unsupported`**

   Example:

   ```
   devtype -c -n CorporateIntranet "Corporate Intranet" -b 21 -a unsupported
   Aded type: CorporateIntranet, id: 25027
   ```

   To get the ID of an unmanaged device type, type `devtype -l`.

2. **Add Ethernet adapter and Ethernet port attributes to the newly created device type.**

   `devtype -s` *devtype-ID* **`-T ethada+eth`**
   `devtype -s` *devtype-ID* **`-P ethernet-port+eth0+ethada`**

   Example:

   ```
   devtype devtype -s 25027 -T ethada+eth
   Added adapter: 304
   devtype -s 25027 -P ethernet-port+eth0+ethada
   ```

```
Added ethernet port: 305
Reset configuration on 0 device(s)
```

3. **Create an instance of the type on the control plane.**

   `device` *device-ID* **`-c -t`** *devtype-name*

4. **Create its connections to the N1 Provisioning Server fabric layer.**

   `device -sC` *chassis-switch-ID chassis-switch-port device-ID* **`eth0`**

5. **Mark the device as `FREE`.**

   `device sF` *device-ID*

   Example:

```
device 1002 -c -t Corporateintranet
Device 1002 is created.
device -sC 50150 netp6 1002 eth0
Added connection 50150:netp6 --> 1002:eth0
device -sF 1002
```

▼ To Add an Unmanaged Ethernet Device

Steps

1. **Type `devtype -l` command to determine the backhaul type ID.**

2. **Add the Ethernet device.**

   a. **Add an Ethernet adapter to the backhaul type by typing:**

      `devtype -s` *backhaul-type-ID* **`-T ethada+eth -a`** *architecture*

      If this device is a load balancer or a backhaul device, the value for the `-a` option is `unsupported`.

      Example: `devtype -s 25091 -T ethada+eth -a unsupported`

   b. **Add the Ethernet port to the backhaul type by typing:**

      `devtype -s` *backhaul-type-ID* **`-P ethernet-port+eth0+ethada`**

3. **You need to perform the preceding steps for every I-Fabric.**

4. **Add the device to the I-Fabric by typing:**

   `device` *device-ID* **`-c -t backhaul`**

   The device ID is the unique identifier for the unmanaged Ethernet device that you are adding to the I-Fabric.

5. **Add the Ethernet connection for the backhaul device by typing:**

   `device -sC` *chasiss-switch-ID chassis-switch-port backhaul-device-ID* **`eth0`**

   The chassis switch ID is the device ID of the chassis switch where the newly added device is connected.

6. **Mark the device as `FREE` by executing the following command:**
   `device -sF` *device-ID*

▼ To Remove an Unmanaged Ethernet Device

Steps

1. **If the device is allocated to a farm, deactivate the farm.**
2. **Set the device to the `FREE` status by typing the command:**
   `device -s -F` *device-ID*
3. **Remove the device from the CPDB by typing the following command:**
   `device -d` *device ID*
4. **Physically remove the device from the I-Fabric.**

`tspr.properties` Configuration File

The `tspr.properties` ASCII text file resides in the `/etc/opt/terraspring` directory on the control plane server. The file modifies the behaviors of all N1 Provisioning Server software. The file is read line by line. Each line specifies one property and its corresponding value. The property values are configurable. In the following property listing default values are provided only where applicable.

The following sections describe the properties and values of the `tspr.properties` file. Default values are provided where applicable.

- Core Properties
- Dynamic Host Configuration Protocol (DHCP) Properties
- Domain Name Service (DNS) Properties
- Drivers Properties
- Monitoring Properties
- CPDB Properties

Core Properties — `TSPRcs`

The following properties are located in the `TSPRcs` package of the `tspr.properties` file.

`com.terraspring.cs.fm.FM.scrubbing`
Description: Specifies whether to scrub disks during deactivation.

Default: true

`com.terraspring.core.log.CPLog.exthandler`

Description: The name of the class to provide the concrete implementation of the `com.terraspring.core.log.CPLog.exthandler` interface.

Note – Although the core package owns this property, its values are defined by the monitoring package.

Default: `com.terraspring.mlg.MonLogHandler`

Core Properties — `TSPRcore`

The following properties are located in the `TSPRcore` package of the `tspr.properties` file.

`com.terraspring.core.log.CPlog.exthandler`

Description: The class name of the object to provide the extended log handler implementation.

Default: null

`com.terraspring.core.sysGridOS.DBConection`

Description: The name of the class to provide the concrete implementation of the `com.terraspring.core.sysGridOS.DBConection` interface.

Default: null

Note – Although the core package owns this property, its values are defined by the database package (`TSPRdb`).

Driver Properties — `TSPRdrvrs`

The following properties are located in the `TSPRdrvrs` package of the `tspr.properties` file.

`com.terraspring.drivers.sun.HighAvailabilitySFB10LB`

Description: Specifies high availability load balancer pairs.

`com.terraspring.drivers.utils.expect.Expect.print`

Description: Specifies whether to enable logging of the actual telnet or terminal server session between the farm manager and a resource pool server.

Default: no

`com.terraspring.drivers.utils.expect.Expect.output`

Description: The file name to which to print the actual telnet or terminal server session between the farm manager and a resource pool server.

Default: null

`com.terraspring.drivers.utils.snmp.SNMP.retries`

Description: Specifies the number of times to attempt an SNMP connection.

Default: 20

`com.terraspring.drivers.utils.snmp.SNMP.timeout`

Description: The maximum amount of time in milliseconds to attempt an SNMP connection before exiting.

Default: 6000

Monitoring Messages — Message Routing

`com.terraspring.core.log.CPLog.exthandler`

Description: The name of the class to provide the concrete implementation of the `com.terraspring.core.log.CPLog.exthandler` interface.

Note – Although the core package owns this property, its values are defined by the monitoring package.

Default: `com.terraspring.mlg.MonLogHandler`

`com.terraspring.mlg. MonLogPolicy.userMsgMode`

Description: Specifies the location to which to route user messages. You can route user messages to multiple locations by specifying a comma-delimited list of values. Possible values are:

- DB (database)
- NONE (no routing)
- NMS (network management system)
- DB, NMS
- NMS, DB

Default: DB

`com.terraspring.mlg.MonLogPolicy.infoMsgMode`

Description: Specifies the location to which to route informational messages. You can route informational messages to multiple locations by specifying a comma-delimited list of values. Possible values are:

- DB (database)
- NONE (no routing)

- NMS (network management system)
- DB, NMS
- NMS, DB

Default: DB, NMS

Monitoring Properties — SNMP Properties

`com.terraspring.mlg.SnmpConf.SnmpTrapReceivePort`

Description: The port on which the SNMP trap receiver receives traps.

Default: 162

`com.terraspring.mlg.SnmpConf.SnmpTrapSendPort`

Description: The port on which the SNMP trap receiver sends traps.

Default: 162

`com.terraspring.mlg.forwardingIP`

Description: The IP address of the third-party NMS on the N1 Provisioning Server.

Control Plane Database Properties — `TSPRdb`

The following properties are located in the `TSPRdb` package of the `tspr.properties` file.

The following properties apply to the Oracle and the Postgres databases.

`com.terraspring.core.sys.GridOS.DBConnection`

Description: The name of the class to provide the concrete implementation of the `com.terraspring.core.sys.DBConnection` interface.

Note – Although the core package owns this property, its values are defined by the CPDB package (`TSPRdb`).

Default: `com.terraspring.db.oracle.OracleCpdbConnection`

The following properties apply to the Postgres database.

`com.terraspring.core.sys.GridOS.CpdbConnection`

Description: The name of the class to provide the concrete implementation of the `com.terraspring.core.sys.DBConnection` interface.

Note – Although the core package owns this property, its values are defined by the CPDB package (TSPRdb).

Default: `com.terraspring.db.postgres.PostgresCpddConnection`

`com.terraspring.db.CpdbConnection.port`

Description: The port number to which the CPDB is connected.

Default: 5432

CHAPTER 3

Managing Software Images

This chapter describes how the N1 Provisioning Server software enables the management of software images in an Infrastructure Fabric (I-Fabric). Managing software images includes the following tasks:

- Creating images
- Upgrading images
- Importing images
- Validating images
- Deploying images
- Updating image attributes
- Listing images
- Deleting images

This chapter covers the following topics:

Software Images Overview

A software image is an archive of operating environment and N1 Provisioning Server agent software, and it might include application packages. Image creation and deployment enables you to configure the software once and deploy it multiple times. Doing so reduces the chance for error and minimizes system configuration time.

N1 Provisioning Server software supports creation and management of two types of images: global images and account images.

- *Global images* consist of the operating environment (including patches and service packs), integrated N1 Provisioning Server agents, and certain customizations.

  Global images may also contain applications. Global images are created using the image wizard tool or manually from the command line.
- *Account images* are for a particular account and might be one of the following account-specific customizations:

  - Global images with customizations such as application and data packages
  - Blank disks (partitioned and formatted)

  Account images are created in the Control Center or on the command line with the `snapshot` command.

An I-Fabric supports images based on the following operating systems, architectures, and formats:

Operating Environment	Architecture	Image Format
Solaris 9, update 5	SPARC	disk, flash, and JumpStart ™software
Solaris 9, update 5	x86	disk, flash, and JumpStart
Solaris 8	SPARC	disk
Linux RedHat 2.7.1	x86	disk

N1 Provisioning Server software comes with the following preconfigured images that you can import or from which you can create new ones:

- Solaris 9 SPARC flash image
- Solaris 9 x86 flash image

Storing Images and Image Attributes

Images are stored on a local disk of an image server. The image server is a network file system (NFS) file server. The image repository in the control plane database (CPDB) contains attributes of each global image stored in the image server.

Creating and Managing Images

You can create and manage images using one of the following methods:

- Using the image wizard.

- Using image preparation and image release commands on the command line.
- Using the Control Center graphical user interface (GUI). See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

The following diagram provides an overview of the sequence of steps and the commands involved in performing the following image management tasks:

- Creating images
- Upgrading images
- Importing images
- Validating images
- Updating image attributes

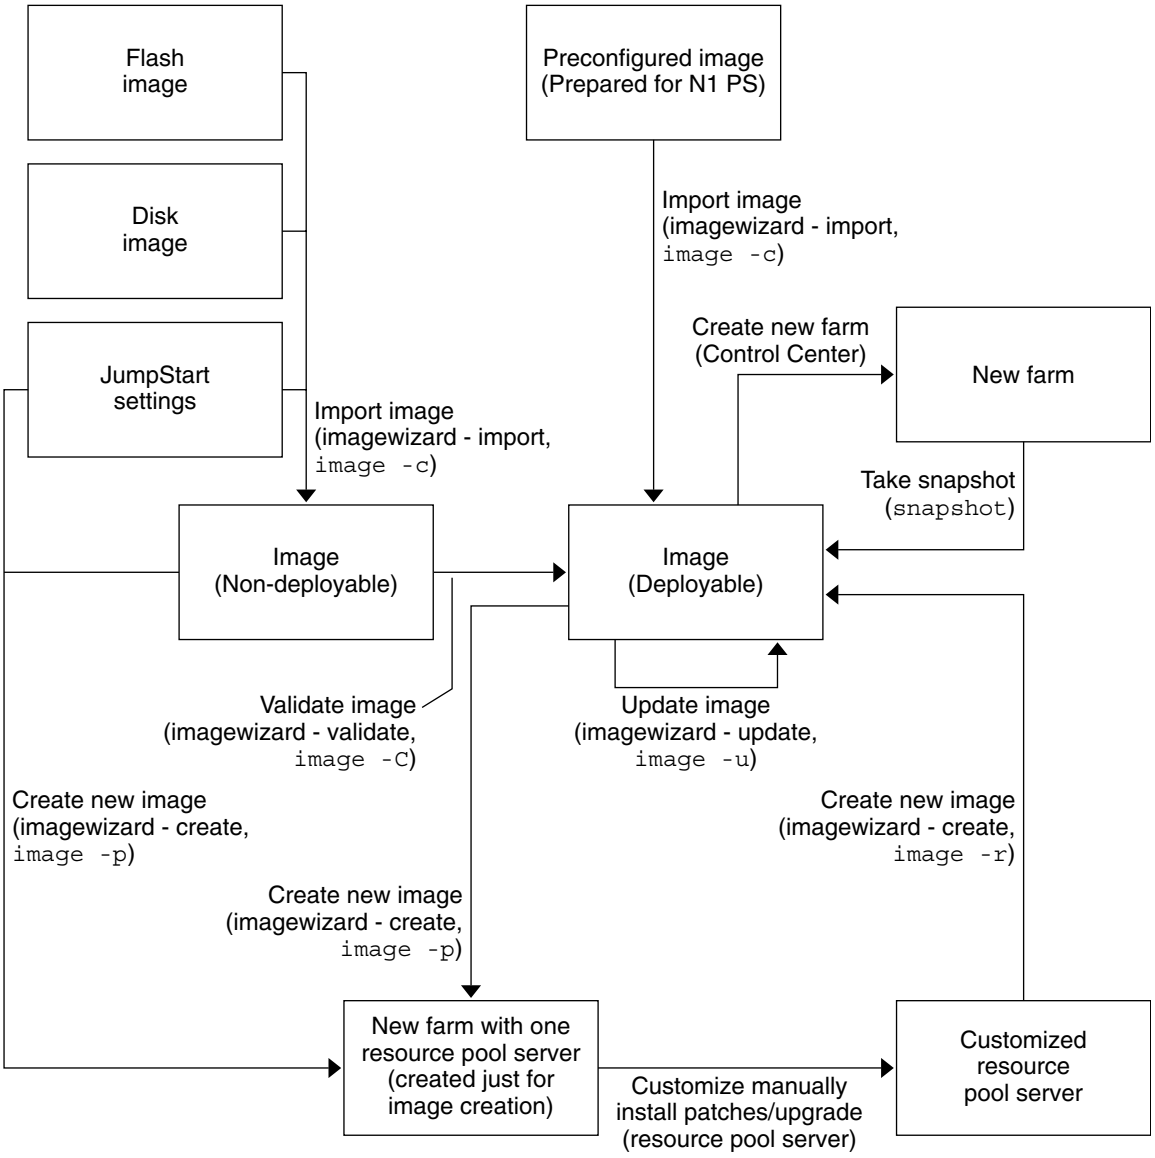

FIGURE 3–1 Image Management Flow Chart

As the diagram shows, you can import the following newly created images from preconfigured images into the image repository using the image wizard user interface or the `image -c` command from the command line.

- disk images
- flash images

- JumpStart images

Note – JumpStart and flash images apply to the Solaris operating environment only.

When an image is imported into the image repository, it is marked as either *deployable* or *non-deployable*. A deployable image contains all relevant N1 Provisioning Server agent software, customizations, and is ready for deployment to a farm using the Control Center. A non-deployable image is not ready for deployment to a farm using the Control Center. All preconfigured images being shipped with the N1 Provisioning Server contain all the relevant software packages and are therefore marked as deployable when imported into the image repository. Images such as disk, flash, and JumpStart images that might not yet contain all relevant software packages are marked as non-deployable when imported into the image repository.

The image validation process validates a non-deployable image and marks it as deployable if the validation is successful. Validating an image ensures that the image can be deployed using the Control Center during farm creation. You can use either the image wizard or the `image -C` command from the command line to validate images.

You can create images and add them to the image repository with either the image wizard user interface or the command line using the `image -p` and `image -r` commands. With either method, image creation follows a two-step process:

1. Create a farm with a single resource pool server and install an operating environment on the resource pool server using one of the following methods:
   - Creating a new image for the resource pool server from an existing image in the repository
   - Installing an operating environment using an existing JumpStart environment
   - Installing an operating environment manually
2. After you have installed the initial operating environment on the resource pool server, you can customize with other applications. You can then release a disk or flash image using the `image -r` command.

You also can take a snapshot of an existing image on a resource pool server in a farm using the `snapshot` command. A disk or flash image created in this way is marked as deployable because the resource pool server in use for a farm has the N1 Provisioning agent installed.

Each image is associated with image attributes, such as image description and size. You can update image attributes using either the image wizard or the `image -u` command on the command line.

For an image upgrade, such as software upgrade or patch installation, you need to create a new image. You can add the software upgrade or patches during the customization step of the image creation process. See “Upgrading Images” on page 96.

Creating and Managing Images Using the Image Wizard

The image wizard runs on the control plane server and provides a menu-driven, command-line interface for creating, importing, and managing images. Using the image wizard you can perform the following tasks:

Task	Image Format	Source	Architecture
Create	Disk or flash	From an existing disk image in the image repository	Solaris SPARC or x86
Create	Disk or flash	From an existing flash image in the image repository	Solaris SPARC or x86
Create	Disk or flash	From a JumpStart installation	Solaris SPARC or x86
Create	Disk or flash	Manual	Solaris SPARC or x86
Create	Disk	From an existing disk image in the image repository	Linux x86
Create	Disk	Manual	Linux x86
Upgrade	All	All	All
Import	Disk or flash	Preconfigured	Solaris SPARC or x86
Import	Disk	Preconfigured	Linux x86
Import	Disk, flash, or JumpStart	All	Solaris SPARC or x86
Import	Disk	All	Linux x86
Import	All	All	All
Validate	All	All	All
List	All	All	All
Update image attributes	All	All	All
Delete	All	All	All

Accessing the Image Wizard

Use the `/opt/terraspring/sbin/imagewizard` command to access the image wizard from the control plane server. See the `imagewizard` man page for details on how to use the command.

Note – You must have root access to the control plane server before you can access the image wizard.

The following is the image wizard's main dialog:

```
Please select,

1) Create - Create a new image and store the image in the repository
2) Import - Import an image into the repository
3) List   - List the images in the repository
4) Update - Update image attribute information in the repository
5) Delete - Delete an image in the repository
6) Validate - Validate imported images
7) Quit
Enter your selection [1-7] >
```

▼ To Create an Image

Steps

1. **Type `/opt/terrasprin/sbin/imagewizard` to access the image wizard's main dialog.**

2. **In the image wizard's main dialog, type `1` to create an image.**

   Make your selections from the following dialog depending on whether you want to create an image from one of the following:

Image Format	Source	Architecture
Disk	Existing disk image	Solaris SPARC
Flash	Existing disk image	Solaris SPARC
Disk	Existing disk image	Solaris x86
Flash	Existing disk image	Solaris x86
Disk	Existing flash image	Solaris SPARC
Flash	Existing flash image	Solaris SPARC
Disk	Existing flash image	Solaris x86

Image Format	Source	Architecture
Flash	Existing flash image	Solaris x86
Disk	JumpStart installation	Solaris SPARC
Flash	JumpStart installation	Solaris SPARC
Disk	JumpStart installation	Solaris x86
Flash	JumpStart installation	Solaris x86
Disk	Manual installation	Solaris SPARC
Flash	Manual installation	Solaris SPARC
Disk	Manual installation	Solaris x86
Flash	Manual installation	Solaris x86
Disk	Existing disk image	Linux x86
Disk	Manual installation	Linux x86

The following dialogs show example selections for creating images from an existing image.

3. **Select the operating environment.**

```
Following are the operating systems supported:

1) Solaris
2) Linux
3) Return to previous option


Please enter your selection [1-3] > 1
```

4. **Select the architecture.**

```
Following architectures are supported:

1) sun4ublade
2) i86pc
3) Return to previous option

Please enter your selection [1-3] > 1
```

5. **Indicate the source for the image.**

```
You can start creating a new image from:

1) Existing image in the repository
2) JumpStart installation server
3) Manual OS installation
4) Return to previous option
```

```
Enter your selection [1-4] > 1
```

6. **Select the image type.**

```
Please select among the following images

--------------------------------------------------------------------------
Sel#    Image Name                   Image Id    Image Description
--------------------------------------------------------------------------

1) solaris9u5-sun4ublade-flash        : 7 : solaris9 update 5 flash image
2) Return to previous option

Please enter your selection [1-2] > 1
```

7. **Type `1` to select the device attributes menu for the image.**

```
Please provide the following information for the new image

1) Device Selection Attributes for the New Image
2) Image Attributes Of the New Image
3) Continue to create the New Image
4) Return to previous option

Enter your selection [1-4] > 1
```

8. **Type `1` to select the server type.**

```
Please enter the following device selection information

1) Server Type:
2) Disk Type:
3) Disk Controller Type:
4) Disk Size: Not specified
5) Return to previous option

Please enter your selection [1-5] > 1
```

9. **Indicate the server type.**

```
Please select the server type
1) sunfireb100x-96-blade
2) sunfireb100s-95-blade
3) Return to previous option

Please enter your selection [1-3] > 2
```

10. **Type `2` to select the disk type.**

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type:
3) Disk Controller Type:
4) Disk Size: Not specified
```

```
5) Return to previous option

Please enter your selection [1-5] > 2
```

11. Indicate the disk type.

```
Please select the disk type
1) local
2) Return to previous option

Please enter your selection [1-2] > 1
```

12. Type `3` to select the disk controller.

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type:
4) Disk Size: Not specified
5) Return to previous option

Please enter your selection [1-5] > 3
```

13. Indicate the disk controller.

```
Please select the disk controller type

1) ide
2) Return to previous option

Please enter your selection [1-2] > 1
```

14. Type `4` to select the disk size.

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type: ide
4) Disk Size: Not specified
5) Return to previous option

Please enter your selection [1-5] > 4
```

15. Indicate the disk size.

```
Please select the disk size

1) 30000000000
2) Return to previous option

Please enter your selection [1-2] > 1
```

16. Type `5` to return to the main image information menu shown in step 5.

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type: ide
4) Disk Size: 30000000000 Bytes
5) Return to previous option

Please enter your selection [1-5] > 5
```

17. Type 2 to select the image attributes menu.

```
Please provide the following information for the new image

1) Device Selection Attributes for the New Image
2) Image Attributes Of the New Image
3) Continue to create the New Image
4) Return to previous option

Enter your selection [1-4] > 2
```

18. Indicate the image name.

```
Please enter the following image attribute information

1) Image name:
2) Image description:
3) Image archive type:
4) Image type: Global Image
5) Image size: Not specified
6) Return to previous option

Please enter your selection [1-6] > 1

Please enter the image name > s9u5-customized-test
```

19. Type 2 to select and specify the image description.

```
Please enter the following image attribute information

1) Image name: s9u5-customized-test
2) Image description:
3) Image archive type:
4) Image type: Global Image
5) Image size: Not specified
6) Return to previous option


Please enter your selection [1-6] > 2
Please enter the image description > created from s9u5 preconfigured image with customiza
```

20. Type 3 to select the archive type.

```
Please enter the following image attribute information
```

```
1) Image name: s9u5-customized-test
2) Image description: created from s9u5 preconfigured image with customizations
3) Image archive type:
4) Image type: Global Image
5) Image size: Not specified
6) Return to previous option

Please enter your selection [1-6] > 3
```

21. Indicate the archive type.

```
Following image archive types are supported:

1) disk_image
2) flash
3) Return to previous option


Please enter your selection [1-3] > 2
```

22. Type `6` to return to the main image information menu shown in step 5.

Note – Selection `4` from the following menu is not available for modification because the image wizard enables creation of global images only.

```
Please enter the following image attribute information

1) Image name: s9u5-customized-test
2) Image description: created from s9u5 preconfigured image with customizations
3) Image archive type: flash
4) Image type: Global Image
5) Image size: Not specified
6) Return to previous option

Please enter your selection [1-6] > 6
```

23. Type `3` to continue the image creation.

```
Please provide the following information for the new image

1) Device Selection Attributes for the New Image
2) Image Attributes Of the New Image
3) Continue to create the New Image
4) Return to previous option

Enter your selection [1-4] > 3
Following is the information you have entered

Create from existing image: solaris9u5-sun4ublade-flash
Operating system: solaris
Architecture: sun4ublade
Image name: s9u5-customized-test
```

```
Image description: created from s9u5 preconfigured image with customizations
Image archive type: flash
Image type: Global Image
Image size: Not specified
Server Type: sunfireb100s-95-blade
Disk Type: local
Disk Controller Type: ide
Disk Size: 30000000000 Bytes
```

24. Type `y` to continue with the image creation procedure.

```
Enter y to continue and n to return to previous option > y
```

25. Verify that you have enough space on the image server for the new image, then type `y`.

```
Please ensure that you have enough space on the image server
to hold the new image that will be created.
The path on image server where the image will be created is:
        /images/master-images

Enter y to continue and n to return to previous option > y
```

26. If you want to use NFS shares from the N1 Provisioning Server, type `y`. If you do not want to use NFS shares from the N1 Provisioning Server, type `n`.

```
Specify whether you want to use NFS
shares from N1 Provisioning Server [default n] > n

Preparing farm for image creation/validation.
This will take some time, please wait till the operation completes.
Do not interrupt the operation.
Please run farm -Lt <farm id> for more information.
```

Example 3–1 Image Creation Output

The following output example also contains instructions on how to prepare the image for release to the resource pool server. See "Creating A Solaris Image from a Preconfigured Image" on page 87, "Creating a Solaris Image Manually from JumpStart" on page 91, "Creating the JumpStart Setup" on page 106, and "Customizing the JumpStart Boot and Configuration Server" on page 108.

If, after starting the `image -r` command, a `replaceFailedDevice` request is queued. You must delete this request for the image creation process to complete successfully. Run the `request -l` command to get the request ID, then use the `request -d` *request_ID* command to do delete the request.

```
Creating farm 'ImageTool1076361866520-751' ...
Request (id: 335) submitted.
Waiting for request 335 to complete...
.
Farm ImageTool1076361866520-751 is created. Farm id = 108
Activating farm 108 ...
```

```
Request (id: 336) submitted.
Waiting for farm request 336 to complete or
farm to reach state 50 ...
.....................................................
Farm 108 is ready for image creation/update.
Please proceed to PES 50101, disk 2 and install/update the software.

After you are finished with the image, please execute the following to release the image:
    image -r [-S image_size] [-T image_archive_type] -f 108 -i <new-image-name>


Completed preparing farm for image creation/validation
Farm ID is: 108

The following resource pool server has been
selected for creating an image.
Resource Pool Server Device ID: 50101

Following is the information on how you can gain console
access to this machine.

In an another window, telnet to the specified system controller (SC)
based on the IP address provided below.

Login into the SC using the username and password of the SC and type

console <id>

The <id> is the id printed below,
(for e.g., sc> console s5)

 Console Information
 ====================
 IP address of Terminal-Server(Service Controller): 10.5.136.20
 Port(Blade) ID: s1



Issue the following commands on the sytem controller to boot the
resource pool server from the disk

sc> bootmode bootscript="boot disk" s<x>.
(where "s<x>" is provided in resource pool server information)

then you can issue the following command to reset

sc> reset -y s<x>

Please enter y when the resource pool server has booted from disk > y


You can customize your installation at this point in time.

Please logon to the resource pool server and customize.
```

```
Please enter y when your customization is completed > y


Please shutdown the resource pool server by issuing the command

For Solaris:
# /usr/sbin/shutdown -y -g0 -i0
For Linux:
# /sbin/shutdown -h now

Once the resource pool server shuts down, for sparc blades set the
OBP of resource pool server to boot using dhcp by issuing the following
command in the OBP prompt (no changes are required for i86pc blades)

ok> setenv boot-device net:dhcp

Please enter y when the resource pool server has completed the shutdown
and when you have configured the resource pool server to boot using dhcp > y

Creating image.

This will take some time, please wait till the operation completes.
Do not interrupt the operation.
Please run farm -Lt <farm-id> for more information.

Disabling monitoring on PES 50101
Powering off PES 50101 for a move to it's original VLAN
Setting PES State to OFF
Moving PES 50101 to its original VLAN
Powering on PES 50101 in it's original VLAN
Setting PES State to ON
Activating farm 108 ...
Request (id: 397) submitted.
Waiting for request 397 to complete...
..
Taking snapshot of target 0 for host server for farm 108 to create final image ...
Request (id: 400) submitted.
Waiting for request 400 to complete...
.....................................................................................
Snapshot was successful.
Deactivating the farm 108
Request (id: 456) submitted.
Waiting for request 456 to complete...
.....................................................................
Farm 108 is deactivated.
Image s9u5-customized-test is ready for use.

Image creation completed.


Updating image attributes.


Updated image: 14
```

```
Image update completed.

WARNING: The farm created for image creation has to be deleted manually.
        Use the farm -D <farm-id> command to delete the farm.
        Please wait for the farm to be deactivated before you issue
        this command.
```

▼ To Upgrade an Image

Before You Begin

For an image upgrade, such as software upgrade or patch installation, you need to create a new image. You can add the software or patches during the customization step of the image creation process.

Steps

1. **Select the Create Image option.**
2. **Select the Create from Existing Image option.**
3. **Provide the information for the new image.**
4. **Answer the following questions with "y".**

   ```
   Specify whether you want to use NFS shares from the N1 Provisioning Server [default n] y.
   ```

5. **The image wizard will prompt you to customize the image. To do so connect to the resource pool server using telnet and install any new software packages and patches.**
6. **Follow the image wizard instructions to create the new image.**

▼ To Import an Image

Steps

1. **From the image wizard's main dialog, type 2 to import an image.**
2. **Specify whether you want to import a preconfigured image or another image.**
   - **Type 1 to import a preconfigured image.**
   - **Type 2 to import another image.**

     ```
     You can start importing images from:

     1) Canned images supplied with this product
     2) Other images
     3) Return to previous option

     Enter your selection [1-3] > 2
     ```

```
CAUTION:  You are about to import an image that is not supplied with
this product. Please refer to the System Administration Guide for importing
images not supplied with the product.
```

3. **Type `y` to continue importing an image.**

```
Please enter y if you want to continue [default y] > y
```

4. **Type `1` to provide operating environment information for the image to be imported.**

```
Please provide information for importing image

1) Operating system:
2) Architecture:
3) Image name:
4) Image description:
5) Image archive type:
6) Image type: Account Image: Account Name:
7) Image size: 0
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 1
```

5. **Select an operating environment from the list.**

```
Following are the operating systems supported:

1) Solaris
2) Linux
3) Return to previous option

Please enter your selection [1-3] > 1
```

6. **Type `2` to provide architecture information for the image to be imported.**

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture:
3) Image name:
4) Image description:
5) Image archive type:
6) Image type: Account Image: Account Name:
7) Image size: 0
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 2
```

7. **Select an architecture from the list.**

```
Following architectures are supported:

1) sun4ublade
2) i86pc
3) Return to previous option

Please enter your selection [1-3] > 1
```

8. **Type `3` to provide a name for the image to be imported.**

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name:
4) Image description:
5) Image archive type:
6) Image type: Account Image: Account Name:
7) Image size: 0
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 3
```

9. **Type the image name.**

```
Please enter the image name > s9u5-jumpstart-test
```

10. **Type `4` to provide a description for the image to be imported.**

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description:
5) Image archive type:
6) Image type: Account Image: Account Name:
7) Image size: 0
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 4
```

11. **Type the image description.**

```
Please enter the image description > s9u5 jumpstart settings
```

12. **Type `5` to provide the archive type for the image to be imported.**

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type:
6) Image type: Account Image: Account Name:
7) Image size: 0
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 5
```

13. Select the image archive type from the list.

```
Following image archive types are supported:

1) disk_image
2) flash
3) jumpstart
4) Return to previous option

Please enter your selection [1-4] > 3
```

14. Type `6` to provide the image type for the image to be imported.

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type: jumpstart
6) Image type: Account Image: Account Name:
7) Image size: Not specified
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 6
```

15. Select the image type from the list.

```
Please enter the image type

1) Global Image
2) Account Image
3) Return to previous option

Enter your selection > 1
```

16. Type `7` to provide the size for the image to be imported.

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type: jumpstart
6) Image type: Global Image
7) Image size: Not specified
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 7
```

17. Specify the image size.

```
Please enter the image size [Not specified] > 30000000000
```

18. Type `8` to provide the location for the image to be imported.

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type: jumpstart
6) Image type: Global Image
7) Image size: 30000000000
8) Image location:
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 8
```

19. Specify the location.

```
Provide the image location URL (<protocol>://<device-id>/<path>)
Make sure the directory specified by <path> exists before going further.
(e.g., nfs://3001/images/master-images/test-image-1
 e.g., ftp://3001/images/master-images/test-image-2
 e.g., nfs://3001/images/jumpstart) > nfs://3001/images/jumpstart
```

20. Type `9` to provide the configuration file for the image to be imported.

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type: jumpstart
6) Image type: Global Image
```

```
7) Image size: 30000000000
8) Image location: nfs://3001/images/jumpstart
9) Image config file:
10) Continue to import image
11) Return to previous option

Please enter your selection [1-11] > 9
```

21. Type `1` to provide the JumpStart installation server for the image to be imported.

```
Please provide the following information for the Jumpstart configuration

1) Solaris install server:Solaris product directory:
2) Solaris boot server:Solaris boot directory:
3) Configuration server:Configuration directory:
4) JumpStart server:JumpStart directory:
5) Return to previous option

Please enter your selection [1-5] > 1
```

22. Specify the JumpStart installation server and path.

```
Specify server and the path to directory on the server where the
solaris distribution can be found. Refer setup_install_server -s option.
[10.42.42.1:/images/s9u5] > 10.42.42.1:/images/s9u5
```

23. Type `2` to provide the JumpStart boot server for the image to be imported.

```
Please provide the following information for the Jumpstart configuration

1) Solaris install server:Solaris product directory: 10.42.42.1:/images/s9u5
2) Solaris boot server:Solaris boot directory:
3) Configuration server:Configuration directory:
4) JumpStart server:JumpStart directory:
5) Return to previous option

Please enter your selection [1-5] > 2
```

24. Specify the JumpStart boot server and path.

```
Specify server and the path to directory on the server where the
solaris miniroot is located. Refer setup_install_server -b -t option.
[10.42.42.1:/tftpboot/terraspring/boot_loader/sun4ublade/Boot] >
10.42.42.1:/tftpboot/terraspring/boot_loader/sun4ublade/Boot
```

25. Type `3` to provide the JumpStart configuration server for the image to be imported.

```
Please provide the following information for the Jumpstart configuration

1) Solaris install server:Solaris product directory: 10.42.42.1:/images/s9u5
2) Solaris boot server:Solaris boot directory:
10.42.42.1:/tftpboot/terraspring/boot_loader/sun4ublade/Boot
3) Configuration server:Configuration directory:
4) JumpStart server:JumpStart directory:
5) Return to previous option
```

```
Please enter your selection [1-5] > 3
```

26. Specify the JumpStart configuration server and path.

```
Specify server and the path to directory on the server where the
user-defined sysidcfg  file  for  pre-configuring system or network
information is located. Refer add_install_client -p option.
[10.42.42.1:/images/jumpstart] > 10.42.42.1:/images/jumpstart
```

27. Type `4` to provide the JumpStart server directory for the image to be imported.

```
Please provide the following information for the Jumpstart configuration

1) Solaris install server:Solaris product directory: 10.42.42.1:/images/s9u5
2) Solaris boot server:Solaris boot directory:
10.42.42.1:/tftpboot/terraspring/boot_loader/sun4ublade/Boot
3) Configuration server:Configuration directory: 10.42.42.1:/images/jumpstart
4) JumpStart server:JumpStart directory:
5) Return to previous option

Please enter your selection [1-5] > 4
```

28. Specify the JumpStart directory location.

```
Specify server and the path to the directory on the server where the
JumpStart files are located. Refer add_install_client -c option.
[10.42.42.1:/images/jumpstart] > 10.42.42.1:/images/jumpstart
```

29. Type `5` to return to the previous menu.

```
Please provide the following information for the Jumpstart configuration

1) Solaris install server:Solaris product directory: 10.42.42.1:/images/s9u5
2) Solaris boot server:Solaris boot directory:
10.42.42.1:/tftpboot/terraspring/boot_loader/sun4ublade/Boot
3) Configuration server:Configuration directory: 10.42.42.1:/images/jumpstart
4) JumpStart server:JumpStart directory: 10.42.42.1:/images/jumpstart
5) Return to previous option

Please enter your selection [1-5] > 5
```

30. Type `10` to continue to import the image.

```
Please provide information for importing image

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: s9u5-jumpstart-test
4) Image description: s9u5 jumpstart settings
5) Image archive type: jumpstart
6) Image type: Global Image
7) Image size: 30000000000
8) Image location: nfs://3001/images/jumpstart
9) Image config file: Has config file
10) Continue to import image
```

```
11) Return to previous option

Please enter your selection [1-11] > 10
```

Example 3–2 Image Importing Output

```
Importing image: s9u5-jumpstart-test
Importing image into image repository.

-------------------- output of internal commands --------------------
Created server image: 16
```

▼ To Validate an Image

Steps

1. **In the image wizard's main dialog, type `6` to validate images.**
2. **Select the image to be validated.**

```
You can validate these images:

Please select among the following images

--------------------------------------------------------------------------
Sel#    Image Name                    Image Id    Image Description
--------------------------------------------------------------------------

1) s9u5-jumpstart-test                : 16 : s9u5 jumpstart settings
2) Return to previous option

Please enter your selection [1-2] > 1
```

3. **Type `1` to provide the server type information.**

```
Please enter the following device selection information

1) Server Type:
2) Disk Type:
3) Disk Controller Type:
4) Disk Size: Not specified
5) Continue to validate image
6) Return to previous option

Please enter your selection [1-6] > 1
```

4. **Select the server type.**

```
Please select the server type
1) sunfireb100x-96-blade
2) sunfireb100s-95-blade
3) Return to previous option
```

```
Please enter your selection [1-3] > 2
```

5. **Type `2` to provide the disk type information.**

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type:
3) Disk Controller Type:
4) Disk Size: Not specified
5) Continue to validate image
6) Return to previous option

Please enter your selection [1-6] > 2
```

6. **Select the disk type.**

```
Please select the disk type
1) local
2) Return to previous option

Please enter your selection [1-2] > 1
```

7. **Type `3` to provide the disk controller type information.**

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type:
4) Disk Size: Not specified
5) Continue to validate image
6) Return to previous option

Please enter your selection [1-6] > 3
```

8. **Select the disk controller type.**

```
Please select the disk controller type

1) ide
2) Return to previous option

Please enter your selection [1-2] > 1
```

9. **Type `4` to provide the disk size information.**

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type: ide
4) Disk Size: Not specified
5) Continue to validate image
6) Return to previous option
```

```
Please enter your selection [1-6] > 4
```

10. Select the disk size.

```
Please select the disk size

1) 30000000000
2) Return to previous option

Please enter your selection [1-2] > 1
```

11. Type `5` to continue to validate the image.

```
Please enter the following device selection information

1) Server Type: sunfireb100s-95-blade
2) Disk Type: local
3) Disk Controller Type: ide
4) Disk Size: 30000000000 Bytes
5) Continue to validate image
6) Return to previous option

Please enter your selection [1-6] > 5
```

12. Confirm whether you want to continue image validation.

- **Type `y` to continue validation.**
- **Type `n` to discontinue validation.**

```
You are about to validate the selected image.

Enter y to continue and n to return to previous option > y
```

Example 3–3 Image Validation Output

```
Preparing farm for image creation/validation.

This will take some time, please wait till the operation completes.
Do not interrupt the operation.
Please run farm -Lt <farm-id> for more information.


-------------------- output of internal commands --------------------
Creating farm 'ImageTool1076366989009-409' ...
Request (id: 468) submitted.
Waiting for request 468 to complete...
.
Farm ImageTool1076366989009-409 is created. Farm id = 110
Activating farm 110 ...
Request (id: 471) submitted.
Waiting for farm request 471 to complete or
farm to reach state 40 ...
.....................................................................................
Farm 110 has reached the ACTIVE/ACTIVE state.
```

```
Deactivating the farm 110
Request (id: 559) submitted.
Waiting for request 559 to complete...
.....................................................................................
Farm 110 is deactivated.
Image s9u5-jumpstart-test is validated and  ready for use.
--------------------------- end of command --------------------------

Completed preparing farm for image creation/validation
Farm ID is: 110

WARNING: The farm created for image validation has to be deleted manually.
        Use the farm -D <farm-id> command to delete the farm.
        Please wait for the farm to be deactivated before you issue
        this command.
```

▼ To List Images

Steps

1. **In the image wizard's main dialog, type `3` to list images.**
2. **From the display options menu make a selection for displaying the images.**

```
Please enter one of the following options for displaying the images

1)  List all images
2)  List images available for use (READY state)
3)  List images marked as deleted
4)  List global images
5)  List information of global images (brief)
6)  List information of global images (verbose)
7)  List images of a specific account
8)  List information of an image of a specific account (brief)
9)  List information of an image of a specific account (verbose)
10) List information of an image specified by id
11) List information of an image specified by id (brief)
12) List information of an image specified by id (verbose)
13) Return to previous option

Enter your selection [1-13] > 1
```

The following output shows the display of images if you had selected `1` from the display options menu.

```
-------------------- output of internal commands --------------------
IMAGE_ID IMAGE_NAME                      CUSTOMER   SIZE        OS       TYPE       STATE
2        solaris9u5-sun4ublade-flash     __grid__   3004789248  solaris  flash      READY
LOCATION  nfs://3001//images/master-images/solaris9u5-sun4ublade-flash

3        solaris9u5-sun4ublade-disk-image __grid__   30000000000 solaris  disk_image READY
LOCATION nfs://3001//images/master-images/solaris9u5-sun4ublade-disk-image

4        solaris9u5-i86pc-disk-image     __grid__   30000000000 solaris  disk_image READY
LOCATION nfs://3001//images/master-images/solaris9u5-i86pc-disk-image
```

```
5        sw@50250                          __grid__   19813                disk_image READY
6        sw@50150                          __grid__   19807                disk_image READY
21       s9u5-disk_image                   __grid__   5000000     solaris  disk_image READY
LOCATION nfs://3001//images/master-images/s9u5-disk_image

22       s9u5-flash                        __grid__   5000000     solaris  flash      READY
LOCATION nfs://3001//images/master-images/s9u5-flash

23       test-other-images-wizard testid             500000      solaris  disk_image NOT READY
LOCATION nfs://3001/images/master-images/test-image
-------------------------- end of command -------------------------
```

3. **Press `y` to continue displaying the list of images available.**

```
Enter y to continue [default y] > y
```

▼ To Update Image Attributes

Steps

1. **In the image wizard's main dialog, type `4` to update image attributes.**
2. **Type the number of the image you want to update.**

```
You are about to update attributes of an existing image in the
image repository.

CAUTION:  Please refer to the System Administration Guide for updating
          images.  Updating wrongly the operating system, architecture,
          image archive type, image size, image location, or image config file
          might lead to failures in deploying the image.


Please select among the following images

-------------------------------------------------------------------------
Sel#    Image Name                    Image Id    Image Description
-------------------------------------------------------------------------

1) solaris9u5-sun4ublade-flash        : 7 : solaris9 update 5 flash image
2) solaris9u5-i86pc-flash             : 8 : solaris9 update 5 flash image
3) sw@50150                           : 10 : sw@50150
4) Return to previous option

Please enter your selection [1-4] > 1
```

3. **Type the number of the image information you want to update.**

```
You can update the following image information

1) Operating system: solaris
2) Architecture: sun4ublade
3) Image name: solaris9u5-sun4ublade-flash
```

```
4) Image description: solaris9 update 5 flash image for sun4ublade with Cassini NIC
5) Image archive type: flash
6) Image size: 30000000000
7) Image location: nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
8) Image config file: Has config file
9) Continue to update image
10) Return to previous option
```

Note – The image location, selection 7 on the menu, is a default value that you cannot modify.

---

4. **You are prompted to continue to make update selections depending on what type of image information you chose to update in the previous step.**

Caution – If you make a wrong choice in the updating process, no alert message will display. The image will not deploy.

Example 3–4 Updating the Image Size

```
Please enter your selection [1-10] > 6

Please enter the image size [30000000000] > 28000000000

Please enter your selection [1-10] > 9

Updating image attributes.


-------------------- output of internal commands --------------------
Updated image: 7
```

▼ To Delete an Image

Steps

1. **In the image wizard's main dialog, type `5` to delete images.**
2. **Make your selections from the following dialogs depending on whether you want to:**
   - Delete a global image.
   - Delete an account-specific image.

   The following example shows how to delete a global image:

   ```
   You can delete these images:

   1) Delete a global image from the repository
   ```

```
2) Delete an account image from the repository
3) Return to the previous option

Enter your selection [1-3] > 1

You are about to delete an existing image in the image repository.

Please select among the following images

------------------------------------------------------------------------
Sel#    Image Name                   Image Id    Image Description
------------------------------------------------------------------------

1) solaris9u5-sun4ublade-flash        : 7 : solaris9u5-sun4ublade-flash
2) solaris9u5-i86pc-flash             : 8 : solaris9 update 5 flash image
3) sw@50150                           : 10 : sw@50150
4) s9u5-customized-test               : 14 : s9u5-customized-test
5) Return to previous option


Please enter your selection [1-5] > 4

Deleting image.

-------------------- output of internal commands --------------------
Delete Image 14 (y/n)? y
Queueing request to delete image ...
Request (id: 612) submitted.
Waiting for request 612 to complete...
.
Deleting image content at: nfs://3001/images/master-images/s9u5-customized-test
size: 759406246   ip: 10.42.42.1   State: done
-------------------------- end of command --------------------------

Image deletion completed.
```

Creating and Managing Images from the Command Line

Using the command line, you can perform the following image management tasks:

Task	Image Format	Source	Architecture
Create	Disk or flash	From an existing disk image in the image repository	Solaris SPARC or x86

Task	Image Format	Source	Architecture
Create	Disk or flash	From an existing flash image in the image repository	Solaris SPARC or x86
Create	Disk or flash	From a JumpStart installation	Solaris SPARC or x86
Create	Disk or flash	Manual	Solaris SPARC or x86
Create	Disk	From an existing disk image in the image repository	Linux x86
Create	Disk	Manual	Linux x86
Upgrade	All	All	All
Import	Disk or flash	Preconfigured	Solaris SPARC or x86
Import	Disk	Preconfigured	Linux x86
Import	Disk, flash, or JumpStart	All	Solaris SPARC or x86
Import	Disk	All	Linux x86
Import	All	All	All
Validate	All	All	All
List	All	All	All
Update image attributes	All	All	All
Delete	All	All	All

The following graphic illustrates the sequence of creating and upgrading a global image:

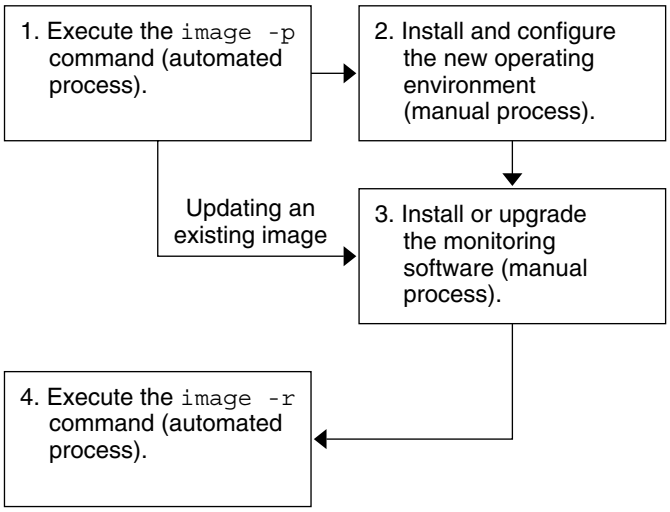

FIGURE 3–2 Creating and Updating Global Images from the Command Line

Images are available to farms through the Control Center once the Control Center is synchronized. When configuring servers for a farm, you can select global images from a drop-down list as described in the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide*.

Note – Before you begin creating global images, ensure that enough storage is available on the N1 Provisioning Server or, if you are using a separate server for storing images, on the image server.

Image Creation Methods and the Steps Involved

The following table summarizes the command line image creation methods supported by the N1 Provisioning Server software.

TABLE 3–1 Image Creation Methods

Steps	Creating from JumpStart Setup	Creating from existing disk, flash, or JumpStart Image
Prerequisite for this operation.	Set up a Solaris JumpStart environment.	Image must exist in the CPDB.

TABLE 3–1 Image Creation Methods *(Continued)*

Steps	Creating from JumpStart Setup	Creating from existing disk, flash, or JumpStart Image
Run the `image -p` command on the control plane server. Wait for the command to complete.	Run the `image -p -j` command on the control plane server. Wait for the command to complete.	Run the `image -p -i` command on the control plane server. Wait for the command to complete.
Open a terminal window (terminal 1) and logon to the system controller.	Open a terminal window and logon to the system controller using the system controller ID (admin) and password (admin).	Open a terminal window and logon to the system controller using the system controller ID (admin) and password (admin).
	At the `sc` prompt type **`console -f`** *blade-ID* to get console access.	At the `sc` prompt type **`console -f`** *blade-ID* to get console access.
Open another terminal window (terminal 2) and logon to the system controller to get console access to the resource pool server.	Open another terminal window and logon to the system controller to get console access to the resource pool server.	Open another terminal window and logon to the system controller to get console access to the resource pool server.
From terminal 1 set up the resource pool server to start the installation from the resource layer network. Wait for this operation to complete.	From terminal 1 set up the resource pool server to start the installation from the resource layer network.	
	At the `sc` prompt type **`bootmode bootscript="boot net:dhcp"`** *blade -D* for x86 architecture or **`bootmode bootscript="boot net:dhcp - install"`** *blade-ID* for SPARC architecture.	
	Type **`reset -y`** *blade-ID*. Wait for this operation to complete.	
From terminal 2 logon to the resource pool server.	From terminal 2 logon to the resource pool server. The default password for root user is `root`.	
From terminal 2 shut down the resource pool server.	From terminal 2 shut down the resource pool server by issuing the command **`/usr/sbin/shutdown -y -g0-i0`** for Solaris.	

TABLE 3–1 Image Creation Methods *(Continued)*

Steps	Creating from JumpStart Setup	Creating from existing disk, flash, or JumpStart Image
From terminal 1 reboot the resource pool server to boot from disk. Wait for this operation to complete.	From terminal 1 configure the resource pool server to boot from disk at the system controller prompt by typing **bootmode bootscript="boot disk"** and **reset -y** *blade-ID*. Wait for this operation to complete.	
From terminal 2 logon to the resource pool server as root user with password `root` and customize the installation.		From terminal 2 logon to the resource pool server as root user with password `root` and customize the installation.
From terminal 2 customize the resource pool server.		From terminal 2 customize the resource pool server.
From terminal 2 shut down the resource pool server.	From terminal 2 shut down the resource pool server by issuing the command **/usr/sbin/shutdown -y -g0 -i0**.	
From terminal 1 and 2 set up the resource pool server to start installation from the resource layer network.	If the resource pool server is a SPARC blade, run the following command in OBP from terminal 1: **setenv boot-device net:dhcp**.	If the resource pool server is a SPARC blade, run the following command in OBP from terminal 1: **setenv boot-device net:dhcp**.
	From terminal 2 set up the resource pool server to boot from the resource layer network.	From terminal 2 set up the resource pool server to boot from the resource layer network.
	At the `sc` prompt type **bootmode bootscript="boot net:dhcp"** *blade-ID*	At the `sc` prompt type **bootmode bootscript="boot net:dhcp"** *blade-ID*
Run the `image -r` command.	Run the `image -r` command.	Run the `image -r` command.

Creating A Solaris Image from a Preconfigured Image

Before creating an image, use the `image -ls` command or the Control Center Administration dialog to check whether the image already exists. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details. The

N1 Provisioning Server software package comes with a preconfigured Solaris image. Use this image as a basis to create your Solaris image. This section describes how to create a Solaris image from the preconfigured Solaris image provided with the N1 Provisioning Server software product.

Creating a Solaris image involves 11 steps. These steps are performed either on the control plane server or on the resource pool server. The following table describes the sequence of the steps to be performed and on which device you need to perform each step for creating an image:

TABLE 3–2 Creating a Solaris Image from a Preconfigured Image

Step No.	Description	Performed on Server
1	Determine the disk type, disk size, server type, and disk controller.	Control plane server
2	Obtain a list of the images.	Control plane server
3	Create a farm with one disk on a resource pool server that contains an appropriate-sized disk using the `image –p -i` command. Wait for this command to complete.	Control plane server
4	Logon to the resource pool server selected by the `image –p -i` command.	Resource pool server
5	Customize the image with any applications and services according to your requirements.	Resource pool server
6	Verify that the N1 Provisioning Server agent is running.	Resource pool server
7	Stop the server.	Resource pool server
8	Configure the server blade to boot from DHCP.	Control plane server
9	Release the farm by running the `image –r` command.	Control plane server
10	Optionally, you can update the image attributes by using the `image -u` command.	Control plane server
11	Delete the farm.	Control plane server

▼ To Create a Solaris Image From a Preconfigured Image

Steps

1. **Determine the server type, disk type, disk size, and disk controller by using the following commands:**

   `/opt/terraspring/sbin/disk -l`

   `/opt/terraspring/sbin/disk -lv`

   `/opt/terraspring/sbin/device -l`

   The output of these commands is used as a parameter to the `image` command in subsequent steps.

2. **Obtain a list of images using the `image -ls` command.**

3. **Run the `image -p -i` command from the N1 Provisioning Server.**

   The command creates a farm with one disk on a resource pool server that contains an appropriate-sized disk.

   The `image -p -i` command for such an image creation procedure takes the form:

   `image -p -i` *image_name server_type disk_type controller_type disk_size*

   For example:

```
image -p -i solaris9u5sun4ublade-flash  sunfire100s-95-blade local ide 30000000000
```

Note – For more detailed information, see the `image` man page.

Note – In addition, if access to the N1 Provisioning Server DVD or another shared medium is required for software installation, use the `-z` option to specify this. You can also specify this option to have access to these shared media during new image creation.

Note – The `image -p` command creates a new farm to be used in creating a new image. The `image -p` tool exits with information regarding the allocated resource pool server and the farm ID of the farm created. The farm ID is required for the eventual release of the allocated resources using the `image -r` command (see below). The `image -r` command releases the image, reactivates the farm, and completes the image creation process.

---

4. **Logon to the resource pool server selected by the `image -p -i` command.**

5. **Customize the image with applications and services according to your requirements.**

Note – To find out which resource pool server is mapped to the farm that was created by the `image -p` command, type **`device -l`**.

---

6. **Verify that the N1 Provisioning Server agent is running by typing from the resource pool server:**

   `ps -aef` | **`grep tspragt`** *IP-address-of-server-used-to-create-image*
   `"tsprnop` *IP-address-of-server-used-to-create-image*`"`

   The output from the `ps` command should look similar to the following:

   ```
   root 361  1 1 00:37:41 ? 0:01 java-Dsun.net.inetaddr.ttl=0
   com.terraspring.mon.client.tspragt start 220.240
   ```

7. **Stop the resource pool server by typing:**

   `/usr/sbin/shutdown -y -g0 -i0` for Solaris

   `/sbin/shutdown –h now` for Linux

8. **Configure the server blade to boot from DHCP.**

   - **For SPARC, at the OBP prompt type:**

     `setenv boot-device net:dhcp`

     Note – x86 servers boot from DHCP by default.

     ---

9. **Run the `image -r` command from the N1 Provisioning Server.**

   This command releases the farm you created with the `image –p` command. Releasing the farm includes the following:

   - Taking a snapshot of the image. The snapshot serves as the new image you have just created.
   - Releasing the resource pool server to the resource pool.
   - Updating the N1 Provisioning Server database with a reference to the image.

   `image -r -f` *farm_ID* **`-S`** *image_size* **`-i`** *image_name* **`-T`** *archive-type*

   For example:

   ```
   image -r -f 171 -S 3000000000 -i new_solaris_image -T flash-image
   ```

Note – If, after starting the `image -r` command, a `replaceFailedDevice` request is queued. You must delete this request for the image creation process to complete successfully. Run the `request -l` command to get the request ID, then use the `request -d` *request_ID* command to do delete the request.

10. **Optionally, you can update the image attributes by using the `image -u` command.**

11. **Delete the farm by typing `farm -Df` *farm-ID*.**

    Example:

    ```
    farm -Df 171
    ```

Make the Image Available to Accounts

To make the image available to accounts, use the Administration dialog of the Control Center to synchronize the Control Center with the CPDB. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

Creating a Solaris Image Manually from JumpStart

Before creating an image, use the `image -ls` command or the Control Center Administration dialog to check whether the image already exists. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

Creating a Solaris image requires setting up a JumpStart server. This section describes how to create a Solaris image using a JumpStart server.

Creating a Solaris image using JumpStart involves 14 steps. These steps are performed either on the control plane server or on the resource pool server. The following table describes the sequence of the steps to be performed and on which device you need to perform each step for creating an image:

TABLE 3–3 Creating a Solaris Image from JumpStart

Step No.	Description	Performed on Server
1	Determine the disk type, disk size, server type, and disk controller.	Control plane server
2	Set up a JumpStart installation server using DHCP.	Control plane server

TABLE 3–3 Creating a Solaris Image from JumpStart *(Continued)*

Step No.	Description	Performed on Server
3	Customize the boot and configuration JumpStart servers for either SPARC or x86 architectures.	
4	Create a farm with one disk on a resource pool server that contains an appropriate-sized disk using the `image -p -j` command.	Control plane server
5	Logon to the chassis that contains the resource pool server.	Resource pool server
6	Logon to the SSC.	Resource pool server
7	Log onto the resource pool server and verify that the N1 Provisioning Server agent is running.	Resource pool server
8	Customize the image.	Resource pool server
9	Stop the resource pool server.	Resource pool server
10	Configure the server blade to boot from DHCP.	Resource pool server
11	Logon to the SSC to set the boot mode for the resource pool server.	Resource pool server
12	Release the farm by running the `image –r` command.	Control plane server
13	Optionally, you can update the image attributes by using the `image -u` command.	Control plane server
14	Delete the farm.	Control plane server

▼ To Create a Solaris Image Manually From JumpStart

Steps **1. Determine the server type, disk type, disk size, and disk controller by using the following commands:**

```
/opt/terraspring/sbin/disk -l
/opt/terraspring/sbin/disk -lv
/opt/terraspring/sbin/device -l
```

The output of these commands is used as a parameter to the `image` command in subsequent steps.

2. **Set up a JumpStart installation Server.**

   See "Creating the JumpStart Setup" on page 106.

3. **Customize the boot and configuration JumpStart servers appropriately for SPARC and x86 architectures.**

   See "Customizing the JumpStart Boot and Configuration Server" on page 108.

4. **Run the `image -p -j` command from the N1 Provisioning Server.**

   The command creates a farm with one disk on a resource pool server that contains an appropriate-sized disk.

   The `image -p -j` command for such an image creation procedure takes the form:

   `image -p -j` *JumpStart_options_file server_type disk_type controller_type disk_size*

   For example:

```
image -p -j /image/jumpstart/js-options sunfire100s-95-blade local ide 30000000000
```

Note – For more detailed information, see the `image` man page.

Note – In addition, if access to the N1 Provisioning Server DVD or another shared medium is required for software installation, use the `-z` option to specify this. You can also specify this option to have access to these shared media during new image creation.

Use the `-j` option to specify the path to a file containing the DHCP name/value pairs pertinent to the Solaris SPARC JumpStart server. This file might look like the following:

```
SinstNM image-server
SinstIP4 10.42.42.1
SinstPTH /images/s9u5s
SrootNM image-server
SrootIP4 10.42.42.1
SrootPTH /images/s9u5s/Solaris_9/Tools/Boot
SsysidCF 10.42.42.1:/images/jumpstart/sun4ublade
SjumpsCF 10.42.42.1:/images/jumpstart
```

The file contains one option pair per line. The Solaris DHCP server options are required and they are typically furnished by the `add_install_client` tool (when used with the `-d` option) for configuring a JumpStart client. See also"Creating the JumpStart Setup" on page 106 for an example of setting up a JumpStart server and running the `add_install_client` script.

Note – The `image -p` command creates a new farm to be used in creating a new image. The farm attempts to activate but will fail because no image has yet been installed on the server. The `image -p` tool exits with information regarding the allocated resource pool server and the farm ID of the farm created. The farm ID is required for the eventual release of the allocated resources using the `image -r` command (see below). The `image -r` command releases the image, reactivates the farm, and completes the image creation process.

The `image` command sets up a resource pool server of the type requested. When the resource pool server is ready, you are notified to connect to the server through the console to initiate the JumpStart installation.

5. **Logon to the system controller of the chassis that contains the resource pool server through telnet.**

Note – To find out which resource pool server is mapped to the farm that was created by the `image -p` command, type **`device -l`**.

Note – If you are using the `postinstall.sh` script and customizing the JumpStart boot and configuration servers, the DHCP client configuration and the N1 Provisioning Server agent installation onto the resource pool server are completed. If you are not using the `postinstall.sh` script, see "Customizing the Resource Pool Server Manually" on page 115.

---

6. **Logon to the blade's SSC to set up the resource pool server to start installation type:**

   For SPARC architecture:

   `bootmode bootscript="boot net:dhcp - install"` *blade-ID*

   `reset -y` *blade-ID*

   For x86 architecture:

   `bootmode bootscript="boot net:dhcp"` *blade-ID*

   `reset -y` *blade-ID*

   This command initiates the JumpStart installation, if the DHCP parameters you provided in previous steps are correct.

Note – If the resource pool server does not boot from disk, from the system controller type:

bootmode bootscript="boot disk" blade *blade-ID*

reset -y

---

7. **Optionally, you can customize the image with applications and services according to your requirements. To customize the image logon to the resource pool server as `root`.**

8. **Verify that the N1 Provisioning Server agent is running by typing:**

   ps -aef | **grep tspragt** *IP-address-of-server-used-to-create-image*
   "tsprnop *IP-address-of-server-used-to-create-image*"

   The output from the `ps` command should look similar to the following:

   ```
   root 361  1 1 00:37:41 ? 0:01 java-Dsun.net.inetaddr.ttl=0
   com.terraspring.mon.client.tspragt start 220.240
   ```

9. **Stop the server by typing:**

   /usr/sbin/shutdown -y -g0 -i0 for Solaris

   /sbin/shutdown -h now for Linux

10. **Configure the server blade to boot from DHCP.**

    - **For SPARC, at the OBP prompt type:**

      setenv boot-device net:dhcp

      Note – x86 servers boot from DHCP by default.

      ---

11. **Logon to the blade's SSC to set the boot mode for the blade to boot from DHCP:**

    bootmode bootscript="boot net:dhcp" *blade-ID*

    reset -y *blade ID*

12. **Run the `image -r` command from the N1 Provisioning Server.**

    This command releases the farm you created with the `image –p` command. Releasing the farm includes the following:

    - Taking a snapshot of the image. The snapshot serves as the new image you have just created.
    - Releasing the resource pool server to the resource pool.
    - Updating the N1 Provisioning Server database with a reference to the image.

`image -r -f` *farm_ID* **`-S`** *image_size* **`-i`** *image_name* **`-T`** *archive-type*

For example:

```
image -r -f 171 -S 3000000000 -i new_solaris_image -T flash
```

Note – If, after starting the `image -r` command, a `replaceFailedDevice` request is queued. You must delete this request for the image creation process to complete successfully. Run the `request -l` command to get the request ID, then use the `request -d` *request_ID* command to do delete the request.

---

13. **Optionally, you can update the image attributes by using the `image -u` command.**
14. **Delete the farm by typing `farm -Df` *farm-ID*.**

Make the Image Available to Accounts

To make the image available to accounts, use the Administration dialog of the Control Center to synchronize the Control Center with the CPDB. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

Creating a Linux Image Manually

For details on how to manually create a Linux image, see the *N1 Provisioning Server 3.1, Blades Edition, Release Notes*.

Upgrading Images

You might want to upgrade an image with patches or software packages. This section describes how to upgrade images. Upgrading images involves the same steps and commands as are used for creating images, except that you use the `-i` *image_name* option with the `image -p` command.

Note – When you upgrade an image, a copy of the original image is placed onto a new disk to allow you to upgrade the newly copied image. Before releasing the new image using the `image -r` command, you must give the new image a unique name.

▼ To Upgrade Images

Steps

1. **Obtain a list of the images using the command:**

```
image -ls
IMAGE_ID  IMAGE_NAME        CUSTOMER    DISK_VOL_ID  SIZE    OS
2             solaris               __grid__    22001/5       8631   Solaris
1087      solaris_dhcpfix    __grid__    22001/22      8631   Solaris
```

2. **Type the following command:**

   image –p [-z] [-i *image_name***]** *server_type disk_type controller_typ disk_size*

   For example:

   ```
   image -p –z –i master_image sun_svr_420R local ide 30004789248
   ```

Note – If access to the N1 Provisioning Server DVD or other shared media is required for software installation, use the `-z` option to specify this.

---

3. **Logon to the resource pool server.**

4. **Modify the image according to your requirements, such as updating the N1 Provisioning Server monitoring package. See "Upgrading the N1 Provisioning Server Agent of an Image" on page 97 for more information.**

5. **Run the `image –r` command from the control plane server.**

   This command releases the farm you created with the `image –p` command. Releasing the farm includes the following:

   - Taking a snapshot of the image. The snapshot serves as the new image you have just created.
   - Releasing the resource pool server to the resource pool.
   - Updating the N1 Provisioning Server database with a reference to the image.

   image -r -f *farm_ID* **[-S** *image_size]* **-i** *image_name*

   For example:

   ```
   image -r -f 171 -S 9050849280 -i image_name
   ```

   The value of the *farm-ID* option is the ID of the farm you created using the `image –p` command in step 3.

6. **Delete the farm by typing:**

   farm -Df *farm-ID*

Upgrading the N1 Provisioning Server Agent of an Image

You can upgrade the monitoring software of an image at any time.

▼ To Upgrade the N1 Provisioning Server Agent Package of an Image

Steps

1. **Use TFTP to get the agent package from the control plane server.**

Note – If you are running the JumpStart post installation script `postinstall.sh`, skip this step.

```
cd var/tmp
tftp>cp
tftp>binary
tftp>get TSPRagsol.pkg
Received 456704 bytes in 0.4 seconds
tftp>quit
```

2. **Type the command `pkgrm` *old-package-name*.**
3. **Type the command `pkgadd -d` *new-package-name*.**

Importing Images

The N1 Provisioning Server supports importing preconfigured and manually created image from the command line.

▼ To Import Preconfigured Images

Use the `image -c` command to import preconfigured images. The `-c` option creates a server image entry in the database.

Before You Begin

Ensure that the image contents are accessible on an NFS-shared directory on the image server.

Ensure that enough storage is available on the N1 Provisioning Server or, if you are using a separate server for storing images, on the image server.

Note – Preconfigured images are packaged with the N1 Provisioning Server agent and other configurations. Do not use the `-n` option with the `image -c` command because it will set the image to a non-deployable state.

Step

- **Type `image -c`**

Example 3–5 Importing an Image

```
image -c -s -L nfs://3001/image/s9u5-img -S 4096M -A sun4ublade
-o solaris s9u5-js 'Solaris 9 update 5 jumpstart for sparc blades'
```

In the example, an image is created and imported with the name `s9u5-img`, the size `4096MB`, located at device `3001` in the directory path `/images/s9u5-js`. The image is for the `sun4ublade` architecture.

▼ To Import JumpStart Images

Steps

1. **Set up the JumpStart server.**
   See "Creating the JumpStart Setup" on page 106.
2. **Customize the boot and configuration JumpStart servers.**
   See "Customizing the JumpStart Boot and Configuration Server" on page 108.
3. **Use the following command to import the JumpStart environment into the N1 Provisioning Server environment:**
   `image -c -s -n -L nfs://3001/images/jumpstart -S` *image size* **`-A sun4ublade -O solaris -T jumpstart -P /images/jumpstart/dhcp-options.txt sun4ublade-js-image "*description*"`**
   Example:

```
image -c -s -n -L nfs://3001//images/jumpstart-3 -S 1 -A sun4ublade -O solaris
-T jumpstart -P dhcp-options.txt solaris-jumpstart-validate-img "Jumpstart validation desc"
Created server image: 92
```

Note – The image locator URL must be unique. The `-n` option is required.

This operation will create a new image in a non-deployable state. You need to validate the image before you can deploy the image to a farm.

Validating Images

▼ To Validate a Solaris JumpStart Image

Validating an image puts the image in a deployable state.

Steps

1. **Import the JumpStart image as described in "Importing Images" on page 98.**
2. **Customize the boot and configuration JumpStart servers as described in "Customizing the JumpStart Boot and Configuration Server" on page 108.**

3. **Obtain a list of images by using the `image -ls` command.**

4. **Use the following command to validate the image:**

   `image -C -i` *image_name server_type disk_type controller_type disk_size*

   Example:

   ```
   image -C -i solaris-jumpstart-validate-img sunfireb100s-95-blade local ide 3004789248
   ```

   The following shows the output of the `image -l` command before image validation:

```
image -l
IMAGE_ID IMAGE_NAME                      CUSTOMER  SIZE    OS              TYPE          STATE
3        sw@50150                        __grid__  19809                   disk_image    READY
92       solaris-jumpstart-validate-img  __grid__  1       solarisjumpstart              NOT READY
LOCATION nfs://3001//images/jumpstart-3
```

   The following shows the output of the `image -l` command after image validation:

```
image -l
IMAGE_ID IMAGE_NAME                     CUSTOMER  SIZE    OS   TYPE          STATE
92       solaris-jumpstart-validate-img __grid__  1       solarisjumpstart   READY
LOCATION nfs://3001//images/jumpstart-3
```

Listing Images from the Command Line

The `image` command provides several options for listing images from the command line:

- Listing all images in the image repository — `image -l -a`
- Listing images available for use— `image -l -y`
- Listing images marked as deleted— `image -l -R`
- Listing global images only — `image -l`
- Listing images of a specific account — `image -lv` *customer-name*
- Listing detailed information of images of a specific account — `image -lV` *customer-name*
- Listing a single image by specifying the image ID — `image -lv` *image-ID*
- Listing detailed information about single image by specifying its ID — `image -lV` *image-ID*

▼ To List all Images in the Image Repository

Step ● **Type `image -l` at the command line.**

Example output:

```
IMAGE_ID  IMAGE_NAME                   CUSTOMER  SIZE        OS       TYPE       STATE
7         solaris9u5-sun4ublade-flash  __grid__  28000000000 solaris  flash      READY
LOCATION: nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
```

```
8         solaris9u5-i86pc-flash       __grid__  30000000000 solaris  flash      READY
LOCATION: nfs://3001//images/master-images/solaris9u5-i86pc-flash
9        js-x86-test-img               __grid__ 30000000000 solaris jumpstart   READY
LOCATION: nfs://3001/images/jumpstartx-GA
10       sw@50150                      __grid__ 19807              disk_image  READY
11       js-sparc-test-img             __grid__ 30000000000 solaris jumpstart   READY
LOCATION: nfs://3001/images/jumpstarts-GA
13       s9u5-x86-disk-jstest-img      __grid__ 30000000000 unknown disk_image  READY
LOCATION: nfs://3001/images/master-images/s9u5-x86-disk-jstest-img
16       s9u5-jumpstart-test           __grid__ 30000000000 solaris jumpstart   READY
LOCATION: nfs://3001/images/jumpstart
```

▼ To List a Single Image by Specifying its ID

Step ● **Type `image -lv` *image-ID* at the command line.**

Example output:

```
IMAGE_ID IMAGE_NAME                CUSTOMER  SIZE         OS       TYPE         STATE
13       s9u5-x86-disk-jstest-img  __grid__  30000000000  unknown  disk_image   READY
LOCATION: nfs://3001/images/master-images/s9u5-x86-disk-jstest-img

Description:   s9u5-x86-disk-jstest-img
Architecture:  i86pc
Last Updated:  2004-02-09 13:48:19.0

Image Locations:
     ID    STATE     SIZE             LOCATION
     58    done      699981071
```

▼ To List Detailed Information about a Single Image by Specifying its ID

Step ● **Type `image -lv` *image-ID* at the command line.**

Example output for a flash archive image:

```
IMAGE_ID IMAGE_NAME                   CUSTOMER  SIZE         OS       TYPE    STATE
7        solaris9u5-sun4ublade-flash  __grid__  28000000000  solaris  flash   READY
LOCATION: nfs://3001//images/master-images/solaris9u5-sun4ublade-flash

Description:   solaris9u5-sun4ublade-flash
Architecture:  sun4ublade
Last Updated:  2004-02-09 13:52:26.0

Image Locations:
     ID    STATE     SIZE             LOCATION
     7     done      759291924
nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
```

```
Image Contents:
================================================
partitioning explicit
filesys any free /
```

Example output for a JumpStart image:

```
IMAGE_ID IMAGE_NAME          CUSTOMER    SIZE          OS        TYPE         STATE
9        js-x86-test-img     __grid__    30000000000   solaris   jumpstart    READY
LOCATION: nfs://3001/images/jumpstartx-GA

Description:   Testing validation of image

Architecture:  i86pc
Last Updated:  2004-02-08 10:25:28.0

Image Locations:
     ID    STATE     SIZE             LOCATION
     9     done      512              nfs://3001/images/jumpstartx-GA
Image Contents:
================================================
SinstNM 10.42.42.1
SinstIP4 10.42.42.1
SinstPTH /images/s9u5x
SrootNM 10.42.42.1
SrootIP4 10.42.42.1
SrootPTH /images/s9u5x/Solaris_9/Tools/Boot
SsysidCF 10.42.42.1:/images/jumpstartx-GA
SjumpsCF 10.42.42.1:/images/jumpstartx-GA
BootFile nbp.SUNW.i86pc
SbootURI tftp://10.42.42.1/SUNW.i86pc
```

Updating Image Attributes from the Command Line

If you upgrade an image, the image attributes should reflect any upgrades you have made.

With the `image -u` command, you can update the following attributes:

- Image name
- Image size
- Architecture
- Operating environment
- Image description
- Archive type
- Image configuration file

Command usage: `image -u -N` *image name* `-S` *image size* `-A` *architecture*`- O` *operating sytem* `-U` *image description* `-T` *image archive type* `-P` *configuration file image ID*

Note – For any image, except for a disk image, if you specify the `-T` option, you must also specify the `-P` option.

EXAMPLE 3–6 `image -u` Command

```
image -u -S 3000000000 101
```

You can specify a new disk layout for a global flash image using the `image -u` command. You can view the default disk layout in the output of the `image -lV` command:

```
image -lV 16
IMAGE_ID IMAGE_NAME                    CUSTOMER  SIZE       OS       TYPE   STATE
16       solaris9u5-sun4ublade-flash   __grid__  3200000000 solaris  flash  READY
LOCATON:  nfs://3001//images/master-images/solaris9u5-sun4ublade-flash

Description:   solaris9 update 5 flash image for sun4ublade with Cassini NIC
Architecture:  sun4ublade
Last Updated:  2004-02-13 23:35:53.0

Image Locations:
    ID    STATE     SIZE             LOCATION
    16    done      759291924
nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
Image Contents:
================================================
partitioning explicit
filesys any free /

================================================
```

Create a new layout file with the following contents:

```
/var/tmp/newdisklayout.sun4ublade
partitioning explicit
filesys any 8096 /
filesys any free /export
```

Then run the command:

```
image -u -P /var/tmp/newdisklayout.sun4ublade 16
Updated image: 16
image -lV 16
IMAGE_ID IMAGE_NAME                    CUSTOMER  SIZE       OS       TYPE   STATE
16       solaris9u5-sun4ublade-flash   __grid__  3200000000 solaris  flash   READY
LOCATION: nfs://3001//images/master-images/solaris9u5-sun4ublade-flash

Description:   solaris9 update 5 flash image for sun4ublade with Cassini NIC
Architecture:  sun4ublade
Last Updated:  2004-02-13 23:35:53.0

Image Locations:
```

EXAMPLE 3–6 `image -u` Command *(Continued)*

```
ID    STATE     SIZE      LOCATION
16    done      759291924  nfs://3001//images/master-images/solaris9u5-sun4ublade-flash
Image Contents:
================================================
partitioning explicit
filesys any 8096 /
filesys any free /export
================================================
```

▼ To Update Image Attributes from the Command Line

Steps

1. **Obtain a list of the images using the command:**

```
image -ls
IMAGE_ID  IMAGE_NAME       CUSTOMER   DISK_VOL_ID  SIZE   OS
2             solaris          __grid__   22001/5      8631   Solaris
1087      solaris_dhcpfix  __grid__   22001/22     8631     Solaris
```

2. **Run the `image -u` command with the options of the attributes you want to update.**

```
image -l 22
IMAGE_ID IMAGE_NAME  CUSTOMER   SIZE        OS       TYPE        STATE
22       js-sample   __grid__   30000000000 unknown  disk_image  NOT READY
LOCATION: nfs://3001/images/master-images/js-sample
image -u -S 40000000000 22
Updated image: 22
image -l 22
IMAGE_ID IMAGE_NAME  CUSTOMER   SIZE        OS       TYPE        STATE
22       js-sample    __grid__  40000000000 unknown  disk_image  NOT READY
LOCATION: nfs://3001/images/master-images/js-sample
```

Deleting Images

When you delete snapshots and images from the Control Center, they are only marked as deleted; they are not yet deleted from the image server and the I-Fabric. Until you delete them from the image server and the I-Fabric, you will not be able to create snapshots and images with the same names as the ones marked as deleted.

To purge snapshots and images from the image server and the I-Fabric, type the `image -lR` command from the control plane server to view a list of the images marked as deleted, then execute the `image -d` command to delete them from the I-Fabric.

▼ To Delete Images

Steps

1. **Obtain a list of the images using the `image` command:**

   `image-ls`

   Example output when listing images:

```
IMAGE_ID  IMAGE_NAME           CUSTOMER  SIZE        OS        STATE
1             solaris8-blade-3001  __grid__    3004789248  solaris8  READY
LOCATION: nfs://3001//images/master-images/solaris8-blade
2         solaris8-blade-3002    __grid__    3004789248  solaris8  READY
LOCATION: nfs://3002//images/master-images/solaris8-blade
3             solaris8-blade-3003    __grid__    3004789248  solaris8  READY
nfs://3003//images/master-images/solaris8-blade
```

2. **Type one of the following commands to delete the image:**

   `image -d {-s` | *customer_name*`} [-L URL]` *name*

   `image -d [-L URL]` *image ID*

   For example:

```
image -d -s js-sample-new
Delete system image js-sample-new (y/n)? y
Queueing request to delete image ...
Request (id: 1633) submitted.
Waiting for request 1633 to complete...
.
Deleting image content at: nfs://3001/images/master-images/js-sample size: 0 ip: 10.2.2.1
State: created
```

   The `image -d` command deletes an image from the image server and removes the entry for that image from the CPDB. It also scrubs and frees the disk. You then use the Administration dialog of the Control Center to synchronize the Control Center with the CPDB. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details.

Creating Account Images Using the Control Center

You can create custom application and data software images using the tools available in the Control Center. Typical operations include creating, deploying, updating, and replacing application+data images, and operating environment+application+data images. You can use the Control Center for managing account images. See the*N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details on how to create and deploy account images.

Alternatively, you can use the `snapshot` command-line tool available on the N1 Provisioning Server to take a snapshot of an existing disk for account image creation. See the `snapshot` man pages for details.

Creating the JumpStart Setup

This section describes how to set up JumpStart for SPARC and x 86 architectures. This section also provides details on how to customize the JumpStart configuration and boot server for SPARC and x86 architectures. Procedures for customizing a resource pool server manually are also included.

▼ To Create the JumpStart Setup for SPARC Architecture

The following procedure only provides the information on JumpStart setup required in the context of the N1 Provisioning Server environment. For details on JumpStart setup see the documentation for Solaris 9. You can set up a JumpStart installation server in several ways. The following procedure is only one example.

Steps

1. **Logon to the image server as superuser.**
2. **Insert the Solaris 9 DVD on the image server.**
3. **Create a directory to contain the DVD image.**
   Example:

   ```
   mkdir -p /images/s9u5s
   ```

   Note – This document assumes the Solaris distribution is available at `/images/s9u5s` (for SPARC) and `/images/s9u5x` (for x86).

4. **Change to the `Tools` directory on the mounted disc.**
   Example:

   ```
   cd /cdrom/cdrom0/s0/Solaris_9/Tools
   ```

5. **Copy the DVD image in the drive to the installation server's hard disk.**
   Example:

   ```
   ./setup_install_server /images/s9u5s
   ```

6. **Verify that the path to the installation server's image is shared appropriately.**

   ```
   share | grep /images/s9u5s
   ```

   Note – If the path is displayed and `anon=0` is displayed in the options, proceed to the next step. If the path is not displayed or you do not have `anon=0` in the options, continue with this step.

a. **Make the installation server available to the boot server by adding this entry to the `/etc/dfs/dfstab` file.**

```
share -F nfs -o rw,anon=0 -d "install server directory"
/images/s9u5s
```

b. **Make sure the installation server's directory path is correctly shared.**

```
ps -ef | grep nfsd
```

If the `nfsd` daemon is not running, start it.

```
/etc/init.d/nfs.server start
```

If the `nfsd` daemon is running, share the installation server.

```
shareall
```

7. **Change directories to root by typing `cd /`.**
8. **Eject the Solaris 9 DVD.**

▼ To Create the JumpStart Setup for the x86 Architecture

Steps

1. **Use the Solaris 9 x86 DVD the `/images/s9u5x` directory to hold the Solaris 9 distribution.**
2. **Follow the instructions given above for SPARC architecture to create a Solaris x86 installation server , but instead of `/images/s9u5s` use the`/images/s9u5x` path.**
3. **Patch the `miniroot` shipped with Solaris 9 12/03 x86 for the drivers for Sun Fire B100x, Sun Fire B200x. See the** *Sun Fire B1600 Chassis, and B100s, B100x, and B200x Blade Product Notes* **and the section** *Applying Mandatory Software Patches to the Solaris x86 Install Image* **for patching details.**
4. **On the image server create the following directory:**

```
mkdir -m 755 /var/tmp/blades
```

5. **Download the software by following the instructions in section** *Solaris x86 Drivers and Documentation* **of the** *Sun Fire B1600 Chassis, and B100s, B100x, and B200x Blade Product Notes***.**
6. **Save the downloaded file to the directory `/var/tmp/blades`.**
7. **At the Solaris prompt on the system you are going to use as the network installation server, unzip the files you have downloaded:**

```
cd /var/tmp/blades
unzip mis.259-4174-11.zip
```

8. **Change to the directory in which you placed `mis.259-4174-11.zip` by typing:**

```
cd /var/tmp/blades
```

9. **Add the patches and packages automatically to the network installation server image by typing:**

```
./modify_install_server -d /images/s9u5x
```

Customizing the JumpStart Boot and Configuration Server

This section describes how to set up and customize the JumpStart boot and configuration server on SPARC and x86 architectures. The JumpStart environment ensures that all the required software is installed on the resource pool server. This section describes customization of the JumpStart boot and configuration server on SPARC and x86 architectures.

If you have specific JumpStart customizations you must merge the scripts provided for N1 Provisioning Server with your customizations. The following example procedures for customizing the JumpStart boot and configuration server for SPARC and x 86 architectures set up a JumpStart environment that uses DHCP.

▼ To Customize the Boot and Configuration Server for the SPARC Architecture

Steps

1. **Create the JumpStart directory on image server.**

   Example:

   ```
   mkdir -p /images/jumpstart-sparc
   ```

2. **Verify that the path to the JumpStart directory is shared appropriately by typing:**

   ```
   share | grep /images/jumpstart-sparc
   ```

   If the path is displayed and `anon=0` is displayed in the options, proceed to the next step. If the path is not displayed or you do not have `anon=0` in the options, continue with this step.

   a. **Make the JumpStart directory NFS shared by adding the following entry to the `/etc/dfs/dfstab` file.**

      ```
      share -F nfs -o anon=0,rw=@10.42.42.0/24 -d "jumpstart server directory sparc"
      /images/jumpstart-sparc
      ```

   b. **Make sure the JumpStart server's directory path is correctly shared.**

      ```
      ps -ef | grep nfsd
      ```

      If the `nfsd` daemon is not running, start it by typing **`/etc/init.d/nfs.server start`**.

      If the `nfsd` daemon is running, share the JumpStart server by typing **`shareall`**.

3. **Use the following two scripts to customize the JumpStart begin and postinstallation for N1 Provisioning Server and copy them to the JumpStart directory:**

```
cp /tftpboot/terraspring/jumpstart/profiles/begin_js.sh
/images/jumpstart-sparc
cp /tftpboot/terraspring/jumpstart/profiles/postinstall.sh /images/jumpstart-sparc
```

See "Customizing the Resource Pool Server Manually" on page 115 if you intend not to use the above two scripts to customize your JumpStart environment.

Note – If you do not have your own specific begin or postinstallation scripts skip to step 4.

Note – If you have your own specific begin or postinstallation scripts incorporate the `begin_js.sh` functionality in your begin scripts and `postinstall.sh` functionality into your postinstallation scripts.

---

4. **Create the following files with the following content for each file:**

```
cat /images/jumpstart-sparc/profile.sparc
```

```
install_type initial_install
     cluster SUNWCXall
     partitioning explicit
     filesys any free /
```

```
cat /images/jumpstart-sparc/rules
```

```
arch sparc begin_js.sh profile.sparc postinstall.sh
```

```
cat /images/jumpstart-sparc/sysidcfg
```

```
system_locale=en US
     timezone=US/Pacific
     timerserver=localhost
     terminal=xterm
     name_service=NONE
     security_policy=NONE
     root_password=WPUDTTMUiG2JI
     network_interface=primary {
             default_route=10.42.42.1
             netmask=255.255.255.0 protocol_ipv6==no
     }
```

5. **Make any other specific JumpStart customizations that you might want to have. See the *Solaris 9 12/03 Installation Guide* for details on how a JumpStart server can be configured.**

6. **Validate the files using the check script.**

```
cd /images/jumpstart-sparc
/images/s9u5s/Solaris_9/Misc/jumpstart_sample/check
```

7. **Ensure that you profile access to the resource pool server through `rsh` by issuing the following commands:**

```
cd /image/9su5s/Solaris_9/Tools/Boot
echo "+" > .rhosts
```

8. **Ensure that a proper boot environment is setup for the platform name and group by issuing the following command:**

```
add_install_client -d -s
Solaris-install-server-ip-address:Solaris-product-directory -c
JumpStart-server-ip-address:JumpStart-directory -p
Jumpstart-Configuration-server-ip-address:Jumpstart-Configuration-directory
platform_name platform_group
```

Note – See the `add_install_client -d -s` for details.

---

   a. **When you install the N1 Provisioning Server software, you specify the image subnet. So when using the N1 Provisioning Server as your JumpStart server, use the IP address on the image subnet. To determine what image IP address the N1 Provisioning Server is using, run the command `ifconfig -a` to get a list of interfaces and their IP addresses. Look for the IP address that is on the image subnet.**

      Example:

```
/images/s9u5s/Solaris_9/Tools/add_install_client -d -s
10.42.42.1:/images/s9u5s -c 10.42.42.1:/images/jumpstart-sparc -p
10.42.42.1:/images/jumpstart-sparc SUNW.Serverblade1 sun4u
```

Note – Each time you run `add_install_client` you are changing the boot file used for booting hardware of the specified `platform_name` and `platform_group`. The resource pool server will boot from the boot file that was setup by the last `add_install_client` command.

---

9. **The `image` command supports the following set of DHCP options:**

   - SrootOpt
   - SrootIP4
   - SrootNM
   - SrootPTH
   - SswapIP4
   - SswapPTH
   - SbootFIL
   - Stz

- SbootRS
- SinstIP4
- SinstNM
- SinstPTH
- SsysidCF
- SjumpsCF
- Sterm

See the *Preconfiguring System Configuration Information* section of the *Solaris 9 12/03 Installation Guide* for more information on the DHCP options.

Following are the contents of the DHCP options file for the example JumpStart configuration created in this procedure:

```
cat dhcp-options.txt
    SinstNM 10.42.42.1
    SinstPTH /image/s9u5s
    SrootNM 10.42.42.1
    SinstIP4M 10.42.42.1
    SinstPTH /images/s9u5/Solaris_9/Tools/Boot
    SsysidCF 10.42.42.1:/images/jumpstart
    SjumpsCFF 10.42.42.1:/images/jumpstart
```

Note – The image server IP address is used for the root server, installation server, and JumpStart configuration.

▼ To Customize the Boot and Configuration Server for the x86 Architecture

Steps

1. **Create the JumpStart directory on image server.**

   Example:

   ```
   mkdir -p /images/jumpstart-x86
   ```

2. **Verify that the path to the JumpStart directory is shared appropriately by typing:**

   ```
   share | grep /images/jumpstart-x86
   ```

   If the path is displayed and `anon=0` is displayed in the options, proceed to the next step. If the path is not displayed or you do not have `anon=0` in the options, continue with this step.

   a. **Make the JumpStart directory NFS shared by adding the following entry to the `/etc/dfs/dfstab` file.**

      ```
      share -F nfs -o anon=0,rw=@10.42.42.0/24 -d "jumpstart server directory x86"
      /images/jumpstart-x86
      ```

   b. **Make sure the JumpStart server's directory path is correctly shared.**

```
ps -ef | grep nfsd
```

If the `nfsd` daemon is not running, start it by typing **`/etc/init.d/nfs.server start`**.

If the `nfsd` daemon is running, share the JumpStart server by typing **`shareall`**.

3. **Use the following two scripts to customize the JumpStart begin and postinstallation for N1 Provisioning Server and copy them to the JumpStart directory:**

```
cp /tftpboot/terraspring/jumpstart/profiles/begin_js.sh
/images/jumpstart-x86
cp /tftpboot/terraspring/jumpstart/profiles/postinstall.sh /images/jumpstart-x86
```

See "Customizing the Resource Pool Server Manually" on page 115 if you intend not to use the above two scripts to customize your JumpStart environment.

Note – If you do not have your own specific begin or postinstallation scripts skip to step 4.

Note – If you have your own specific begin or postinstallation scripts incorporate the `begin_js.sh` functionality in your begin scripts and `postinstall.sh` functionality into your postinstallation scripts.

---

4. **Create the following files with the following content for each file:**

```
cat /images/jumpstart-x86/profile.x86
```

```
install_type initial_install
     cluster SUNWCXall
     partitioning explicit
     filesys any free /
```

```
cat /images/jumpstart-x86/rules
```

```
arch i386 begin_js.sh profile.x86 postinstall.sh
```

```
cat /images/jumpstart-x86/sysidcfg
```

```
system_locale=en US
     timezone=US/Pacific
     timerserver=localhost
     terminal=dtterm
     keyboard=UNKNOWN
     display=UNKOWN
     pointer=UNKNOWN
     monitor=UNKNOWN {
               DisplayChksum=0x0
     }
```

```
root_password=WPUDTTMUiG2JI
name_service=NONE
security_policy=NONE
network_interface=PRIMARY {
          default_route=10.42.42.1
          netmask=255.255.255.0 protocol_ipv6==no
}
```

5. **Make any other specific JumpStart customizations that you might want to have. See the *Solaris 9 12/03 Installation Guide* for details on how a JumpStart server can be configured.**

6. **Validate the files using the check script.**

Note – Run the `check` command from the same Solaris distribution.

```
cd /images/jumpstart-x86
s9u5x/images/s9u5x/Solaris_9/Misc/jumpstart_sample/check
```

7. **Ensure that you profile access to the resource pool server through `rsh` by issuing the following commands:**

```
cd /image/s9u5s/Solaris_9/Tools/Boot
echo "+" > .rhosts
```

8. **Ensure that a proper boot environment is setup for the platform name and group by issuing the following command:**

```
add_install_client -d -s
Solaris-install-server-ip-address:Solaris-product-directory -c
JumpStart-server-ip-address:JumpStart-directory -p
Jumpstart-Configuration-server-ip-address:Jumpstart-Configuration-directory
platform_name platform_group
```

Note – See the `add_install_client -d -s` for details.

---

   a. **When you install the N1 Provisioning Server software, you specify the image subnet. So when using the N1 Provisioning Server as your JumpStart server, use the IP address on the image subnet. To determine what image IP address the N1 Provisioning Server is using, run the command `ifconfig -a` to get a list of interfaces and their IP addresses. Look for the IP address that is on the image subnet.**

      Example:

```
/images/s9u5x/Solaris_9/Tools/add_install_client -d -s
10.42.42.1:/s9u5x/images/s9u5x -c 10.42.42.1:/images/jumpstart-x86 -p
10.42.42.1:/images/jumpstart-x86 SUNW.i86pc i86pc
```

Note – Each time you run `add_install_client` you are changing the boot file used for booting hardware of the specified `platform_name` and `platform_group`. The resource pool server will boot from the boot file that was setup by the last `add_install_client` command.

9. The `image` command supports the following set of DHCP options:

- SrootOpt
- SrootIP4
- SrootNM
- SrootPTH
- SswapIP4
- SswapPTH
- SbootFIL
- Stz
- SbootRS
- SinstIP4
- SinstNM
- SinstPTH
- SsysidCF
- SjumpsCF
- Sterm
- BootFile
- SbootURI

See the *Preconfiguring System Configuration Information* section of the *Solaris 9 12/03 Installation Guide* for more information on the DHCP options.

The following is a sample DHCP options file for the example JumpStart configuration created in this procedure.

```
cat dhcp-options.txt
    SinstNM 10.42.42.1
    SinstPTH /image/s9u5x
    SrootNM 10.42.42.1
    SinstIP4M 10.42.42.1
    SinstPTH /images/s9u5x/Solaris_9/Tools/Boot
    SsysidCF 10.42.42.1:/images/jumpstart-x86
    SjumpsCFF 10.42.42.1:/images/jumpstart-x86
    BootFile nbp.SUNW.i86pc
    SSbootURI tftp://10.42.42.1/SUNW.i86pc
```

Note – The image server IP address is used for the root server, installation server, and JumpStart configuration.

Customizing the Resource Pool Server Manually

This section describes how to manually customize the resource pool server.

▼ To Customize the Resource Pool Server Manually

Steps

1. **Manually complete the JumpStart installation and configure the Solaris image according to the instructions in the Solaris documentation and the following steps.**

   a. **Enable the DHCP on its primary interface (the default).**

      i. **Create the `/etc/hostname.`*interface-identifier* file for the primary interface:**

         For example:

         ```
         /etc/hostname.ce0
         ```

         This file should be empty.

      ii. **Create the `/etc/dhcp.`*interface-identifier* file for the primary interface.**

         This file should contain a single line that reads `wait forever primary`. Do not disable the default `icmp ping` operation on the primary interface from the N1 Provisioning Server software. The monitoring software performs an`icmp ping` operation on the primary interface to check whether a device is running.

      iii. **Create the following file for each additional interface, excluding the loopback interface:**

         `/etc/dhcp.`<*interface_identifier*>

         Each of these files should contain a single line that reads **`wait 60`**.

      iv. **If the image is for a Solaris device with more than two interfaces, also create the following file for each nonprimary and nonloopback interface:**

         `/etc/hostname.`*interface_identifier*

         These files should be empty.

   b. **Change the last line of `/etc/default/dhcpagent` from:**

      ```
      PARAM_REQUEST_LIST=1,3,12,43
      ```

      to

      ```
      PARAM_REQUEST_LIST=1,3,6,12,15,43
      ```

      If the image is for a Solaris device with more than two interfaces (excluding the loopback interface), also make the following modification to the `/etc/default/dhcpagent` file:

      Uncomment the parameter value pair `"RELEASE_ON_SIGTERM=yes"`.

c. Delete all `*.dhc` files in the `/etc/dhcp/` directory.

d. Enable FTP.

Although by default FTP is enabled, the following is how to enable FTP:

i. Uncomment the line in the `/etc/inetd.d.` file that reads:

```
ftpstream tcp6 nowait root /usr/sbin/n.ftpd n.ftpd
```

ii. To determine the ID of the `net` daemon, type the command:

```
ps -ef | grep inetd
```

iii. Type the command:

```
kill -HUP net_pid
```

iv. Ensure that `root` is not in the `/etc/ftpusers` file.

e. Verify that the Java 2 Platform, Standard Edition, on the resource pool server is version 1.4.1_02.

Note – For installation of Solaris 8, include a workaround for Solaris 8 bug No. 4457119. The recommended fix is to create the file **/etc/rcS.d/S99dhcpfix** with the following contents:

```
#
# Name:         /etc/rcS.d/S99dhcpfix
#
# Author:       Chris Morton
#               Ericsson, Inc.
#               (919) 472-6494
#
# This is here to fix a bug in Solaris DHCP  It will remove any lines
# in /etc/inet/hosts that DHCP added, so the
# /sbin/netstrategy program will return the correct values.

HOSTSFILE=/etc/inet/hosts
TEMPHOSTFILE=/tmp/hosts.$$

case "$1" in
"start")
        echo "N1"s Standard DHCP Environment"
        ;; # Fall through -- rest of script is the initialization code
"stop")
        exit 0
        ;;
*)
        echo "Usage: $0 { start | stop }"
        exit 1
        ;;
esac


# What's my hostname?:

hostname=`/sbin/dhcpinfo Hostnamè
if [ -z "$hostname" ]; then
        hostname="unknown"
fi

# What's my IpAddress

ipaddress=`ifconfig ce0 | grep inet | awk '{print $2}'`


# If you can find my hostname anywhere in /etc/inet/hosts, delete
that line.

/usr/bin/rm -f ${TEMPHOSTFILE}
/usr/bin/egrep -v "[    ]${hostname}([  ]|$)" ${HOSTSFILE} >
${TEMPHOSTFILE} 2> /dev/null

# Add a line based on the actuals
```

```
echo "${ipaddress}      ${hostname}  ${TEMPHOSTFILE}

/usr/bin/cp -p ${TEMPHOSTFILE} ${HOSTSFILE}
/usr/bin/rm -f ${TEMPHOSTFILE}
```

f. **Verify that the symbolic link `/usr/java` points to the directory where the Java 2 Platform is installed.**

g. **Reboot the server.**

2. **Copy the N1 Provisioning Server agent package from the N1 Provisioning Server software to the `tftboot` directory on the N1 Provisioning Server for installation on the image.**

Note – If you are running the JumpStart post installation script `postinstall.sh`, skip this step.

---

3. **Use TFTP to get the agent package from the control plane server.**

Note – If you are running the JumpStart post installation script `postinstall.sh`, skip this step.

```
cd var/tmp
tftp>cp
tftp>binary
tftp>get TSPRagsol.pkg
Received 456704 bytes in 0.4 seconds
tftp>quit
```

4. **Manually install the monitoring software package by typing the following command:**

```
pkgadd -d TSPRagsol.pkg
```

Note – If you are running the JumpStart post installation script `postinstall.sh`, skip this step.

Note – Run the command **pkginfo** | **grep** *package_name* to ensure that you have installed the following operating environment packages:

SUNWbzip SUNWbzipx SUNWzip

SUNWgzip SUNWtcsh SUNWscpux

5. **Run the following command to start the monitoring software:**

```
/etc/init.d/N1PSagt start
```

CHAPTER 4

Monitoring and Messaging

This chapter provides an overview of the monitoring and messaging system. The chapter describes the events being monitored within an Infrastructure Fabric (I-Fabric). Topics covered in this chapter include:

Monitoring System Overview

N1 Provisioning Server monitoring software performs the following tasks:

- Monitoring of availability, health, and performance of the following I-Fabric components:
  - The control plane server on which N1 Provisioning Server software runs
  - Resource pool servers
- Providing a message routing framework for events, including the ability to forward messages to a third-party application.

The primary software package for monitoring consists of the following components:

- Monitoring manager processes on the control plane database (CPDB) that collect all messages
- The message repository in the CPDB to which the monitoring manager forwards all messages
- N1 Provisioning Server agents on resource pool servers to monitor their state and health

The N1 Provisioning Server manages all monitoring processes within an I-Fabric. The N1 Provisioning Server actively monitors the state and health of devices within an I-Fabric.

Failing Over or Restarting Failed Processes

Monitoring provides the information needed by the farm manager to make decisions about device failover and recovery or the restarting of failed processes in the control plane, the resource pool, and the fabric layer.

For details on how to define monitoring from the Control Center, see the*N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide*.

Monitoring Within the Control Plane

Monitoring the control plane has three purposes:

- Monitoring I-Fabric devices for network accessibility
- Acting as a single point of contact for all critical events that can occur within an I-Fabric
- Monitoring the control plane server for availability

The monitoring manager on the CPDB monitors the health and performance of the control plane server. Messages are logged to appropriate log files as specified in the `postinstall` scripts for the `TSPRmon` and `TSPRmlg` packages. The log file locations are set using the command `slconfig`, which is not supported for use outside of the packages `postinstall` and `preremove` scripts.

See the `slconfig` man page for details on configuring these log files. The default locations for monitoring log files are `/var/opt/terraspring/log/snmp.log` and `/var/opt/terraspring/log/mon.log`.

▼ To Start and Stop the Monitoring Manager

Steps **1. To start the monitoring manager, type:**

```
/opt/terraspring/sbin/mmd start
```

2. **To stop the monitoring manager, type:**

```
/opt/terraspring/sbin/mmd stop
```

Monitoring Within the Resource Pool

The N1 Provisioning Server is responsible for monitoring its assigned farms as directed by the configuration stored in the CPDB. The N1 Provisioning Server is also the control gateway for some system configuration and control commands, such as configuring additional network interfaces. The monitoring software running on the N1 Provisioning Server forwards monitoring messages to the CPDB.

Monitoring Resource Pool Servers

N1 Provisioning Server agents are deployed on resource pool servers during installation and configuration. The agent software collects various monitoring data from the resource pool server and sends it to the N1 Provisioning Server.

The N1 Provisioning Server processes monitoring data and stores it in the message repository in the CPDB. The CPDB has a local database for messages processed by the servers. You also can route messages to an external application, such as a Network Monitoring System (NMS).

The monitoring system includes the `TSPRagsol` agent package for the Solaris 8 operating environment and the `TSPRaglinx` agent for the RedHat AS2.1 operating environment.

For resource pool servers assigned to farms, the system can monitor the following events:

- Availability of the primary Ethernet interface (`ICMP ECHO` or `ping`)
- CPU utilization
- Disk utilization
- RAM utilization
- Swap memory allocation

Refer to Chapter 4, Building, Updating, and Monitoring Server Farms, of the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for additional monitoring configuration information.

Stopping the Resource Pool Server Monitoring Process

If you need to reboot a resource pool server that is active in a farm, stop the monitoring process. Rebooting is often required after adding applications or other software packages to a server. To stop the monitoring process use the `/opt/terraspring/sbin/tsprmonitor` command to stop the monitoring

process. Doing so prevents `replaceFailedDevice` requests from occurring. The time during which the resource pool server should not be available is configurable. See the `tsprmonitor` man page for details.

Monitoring Within the Fabric Layer

The monitoring manager on the CPDB monitors the health and performance of devices in the fabric layer, such as Ethernet switches. The monitoring manager logs the information according to the logging configuration.

Registering I-Fabric Devices for Monitoring

The equipment that makes up the I-Fabric, such as the control plane server, routers, and switches, must be registered for monitoring with the CPDB. When the system is installed, this registration is done as part of the original deployment. However, you can modify registration if you add, remove, or replace devices.

The detailed process of registering devices with the CPDB is described in the installation procedure in the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide*. After registering the devices, you can view messages about these devices in the CPDB message repository or on the Control Center monitoring screen.

Removing I-Fabric Devices from Monitoring

I-Fabric devices registered for monitoring must be removed when they are no longer needed, or when they are permanently removed (because of hardware failure, for example).

▼ To Remove a Control Plane Device from the CPDB

Step

- **Type the following command:**
  `/opt/terraspring/sbin/cereg -deleteCPDevice -ipaddr` *IP Address*
  This command removes the device with IP address `10.10.10.21`.

Example 4–1 Removing a Control Plane Device from the CPDB

```
/opt/terraspring/sbin/cereg -deleteCPDevice ipaddr 10.10.10.21
```

Automating the Monitoring Process

Monitoring system health requires little manual intervention. The primary activities relevant to monitoring system health that you need to perform include:

- Using the Control Center Monitoring screen to view or configure specific monitoring events for an individual server of a farm.
- Adding and removing resource pool servers from monitoring. This activity is normally handled during initial installation. The addition of new equipment requires some manual intervention.
- Registering or unregistering I-Fabric devices for monitoring if they have been added or removed manually.

Day-to-Day Monitoring Activities

The CPDB provides a single location from where you can observe all critical events and messages.

When a device fails, you must determine the corrective action. Many of the required actions are described in Chapter 6. The section "Troubleshooting Monitoring Problems" on page 196 in Chapter 7 also describes troubleshooting scenarios.

Monitoring and Messaging Log Files

The monitoring log file holds messages from the registration and deregistration of nodes by the Farm Manager. This file also holds all messages that the monitoring daemon generates. Check this log file as part of your day-to-day activities and take the appropriate action to correct any errors as necessary. See Chapter 7for details.

Log file rotation is configured automatically by the `slconfig` command used in the `postinstall` package at installation time. The default log file size in the `/etc/opt/terraspring/logfile_rotation` is 5 MBytes.

The monitoring debug level governs the verbosity of the monitoring system's messages logged in the `/var/opt/terraspring/log/mon.log` file.

By default the debug level is set to `9`. To change the debug level edit the `/etc/opt/terraspring/tspr.properties` file on the N1 Provisioning Server:

`com.terraspring.mon.MonLog.debugLevel=`*debug level value*

Logs of all messages forwarded to an NMS are located in the `/var/opt/terraspring/log/snmp.log` file.

The default Simple Network Management Protocol (SNMP) verbosity level is set to `9`. The SNMP messages are logged into the `/var/opt/terraspring/log/snmp.log` file.

To change the SNMP verbosity level, edit the `/etc/opt/terraspring/tspr.properties` file on the N1 Provisioning Server:

com.terraspring.mon.snmpLog.debugLevel=*debug level value*

Messaging

The N1 Provisioning Server software provides a message routing framework directing messages related to farms and monitoring activities to a central repository on the CPDB or to an external NMS.

N1 Provisioning Server software generates three types of messages:

- Informational messages
- Farm messages
- Billing messages

All messages are forwarded to the N1 Provisioning Server. The N1 Provisioning Server then sends the messages to the message repository in the CPDB. You can view monitoring data for farm servers through the Control Center Monitor screen.

Optionally, you can configure the CPDB to forward messages to an external NMS. See "Management Information Base Definitions" on page 131 for details.

Message Flow to the CPDB

Message flow from the I-Fabric components to the CPDB is automatic. Message flow from the CPDB to an optional third-party NMS is configurable. The following graphic illustrates the flow of messages from I-Fabric elements to the CPDB message repository of an I-Fabric.

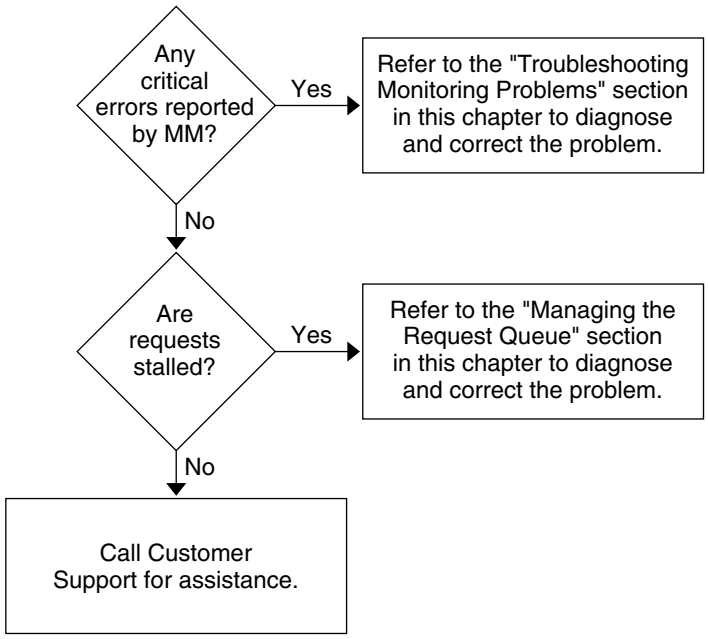

FIGURE 4–1 Message Flow from I-Fabric Elements to the CPDB

Forwarding Messages to an NMS

If you are forwarding messages to an NMS, you need to set up an SNMP connection because all N1 Provisioning Server monitoring messages sent to an NMS are in the form of an SNMP trap. SNMP requires the implementation of a management information base (MIB) for implementing trap versions. The N1 Provisioning Server monitoring mechanism sends only version 2 traps. However the CPDB can convert the trap versions. You can configure the monitoring mechanism to forward messages to an NMS in the `tspr.properties` file on the `/etc/opt/terraspring` directory as follows:

`com.terraspring.mlg.SnmpConf.forwardingIP=`*IP address of NMS*

After configuring this property, restart the `SnmpTrap` daemon as follows:

```
/opt/terraspring/sbin/snmpd stop
/opt/terraspring/sbin/snmpd start
```

See Chapter 2 for more details on monitoring properties.

Configuring Message Routing

You can configure which type of message (farm, billing, or informational) to send to either CPDB or NMS in the `tspr.properties` file in the `/etc/opt/terraspring` directory as follows:

TABLE 4–1 Configuring Message Routing Properties

Property	Description
`com.terraspring.mlg.MonLogPolicy.userMsgMode=DB`	Specifies where to send farm messages. 1. Possible values: `DB, NMS, NONE` 2. Default: `DB`
`com.terraspring.mlg.MonLogPolicy.infoMsgMode=DB`	Specifies where to send informational messages. 1. Possible values: `DB, NMS, NONE` 2. Default: `NMS`

Monitoring Commands

You can use monitoring commands on a command-line interface (CLI) to perform monitoring activities, such as viewing the status of devices, manually registering devices, or troubleshooting. N1 Provisioning Server monitoring software provides the following commands:

- `cereg`
- `cecmd`
- `mls`

Using the `cereg` Monitoring Command

Use the `/opt/terraspring/sbin/cereg` command to manually register and unregister resource pool servers for monitoring. In normal operation, registration of resource pool servers for monitoring is done automatically at I-Fabric installation time, so you would use this command only in an error or troubleshooting situation. See Chapter 7 for details and examples on how to use this command.

Command Usage	Description
`-addAllCPDevices`	Registers all devices on the control plane.
`-deleteCPDevice -ipaddr` *IP address*	Deregisters the control plane server from monitoring.

Command Usage	Description
`-deletenode -ipaddr` *IP address* `-netmask` *netmask* `-nodetype` *node type*	Deregisters the resource pool server from monitoring.
`-addnode -ipaddr` *IP address* `-netmask` *netmask* `-nodetype` *node type*	Adds a resource pool server for monitoring.

Using the `cecmd` Monitoring Command

Use the `/opt/terraspring/sbin/cecmd` command to do one of the following:

- Configure an interface
- Shut down a resource pool server
- Back up a resource pool server

This command is typically used in an error or troubleshooting situation. See Chapter 7 for details and examples on how to use this command.

Command Usage	Description
`-addipinterface -ipAddress` *IPaddress* `-waitTime` *wait-time* `-mac` *MAC-address* `-vlanId` *VLAN-ID* `-ip` *IP-address* `-netMask` *netmask*	Adds an interface to the resource pool server specifying the following information: ■ IP address of the resource pool server ■ Timeout for DHCP ■ MAC address of the interface ■ VLAN ID of the interface
`-deleteipinterface -ipAddress` *IP-address* `-waitTime` *wait-time* `-mac` *MAC-address* `-vlanId` *VLAN-ID* `-ip` *IP-address* `-netMask` *netmask*	Deletes an interface from the resource pool server specifying the following: ■ IP address of the resource pool server ■ Timeout for DHCP ■ MAC address of the interface ■ VLAN ID of the interface
`-shutdown-ipAddress` *IP- address* `-nodetype` *node-type*	Shuts down a resource pool server specifying the IP address and the type of server.
`-backup -ipAdress` *IP-address* [`-start` | `-stop`]	Starts or stops the backup of a resource pool server.
`-addgw -ipAddress` *IP- address* `-ipaddr` *IP-address*	Adds a gateway to the resource pool server specifying the IP address of the resource pool server and the IP address of the default gateway.
`-cleararp -ipAddress` *IP-address*	Clears the ARP table specifying the IP address of the resource pool server.

Command Usage	Description
`-diskChange -ipAddress` *IP-address* `-operation [-add|-delete] -information [-channel:target:LUN:size:name]`	Adds or deletes disk information specifying the IP address and disk information, such as the channel, target, and name.
`-clbconf -ipAddress` *resource-pool-server-IP- address* `-operation [-add|-remove] -ipaddr` *load-balancer -IP-address* `-vip` *VIP-address* `-mask` *VIP-netmask* `-vlanid` *load-balancer-VLAN-ID*	Adds or delete a load balancer specifying the following: ■ Resource pool server IP address ■ Load balancer IP address ■ Virtual IP address ■ Netmask ■ VLAN ID

Using the `mls` Monitoring Command

You can execute the `/opt/terraspring/sbin/mls` monitoring command to view the state of an I-Fabric component, the status of all N1 Provisioning Server agents, or the status of a particular N1 Provisioning Server agent. The `mls` command is useful for troubleshooting as well as day-to-day monitoring of the health of the system.

TABLE 4–2 Using the `mls` Command

Command Usage	Description
`mls -l`	Lists all nodes on the current control plane server.
`mls -all`	Modifies the query scope to the control plane server.
`mls -c`	Lists customer monitoring values for nodes registered on the current control plane server.
`mls -a`	Lists agent status for all registered nodes on the current control plane server.
`mls -x`	Lists the agent version information for the resources controlled by the control plane server.
`mls -f`	Modifies the query scope to a particular farm.
`mls -i`	Modifies the query scope to a particular resource pool server.
`mls -v`	Verbose output.
`mls -d`	Suppresses header information.
`mls -h`	Displays the command's usage information.

Management Information Base Definitions

MIB definitions are described in the following standard SNMP configuration file. The file describes messages supported by the SNMP N1 Provisioning Server agent.

```
-- N1 Provisioning Server  Monitoring MIB.

N1PS-MIB DEFINITIONS ::= BEGIN

IMPORTS
        MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE,
        enterprises, IpAddress
                FROM SNMPv2-SMI
        OBJECT-GROUP, NOTIFICATION-GROUP
                FROM SNMPv2-CONF;

smiModuleIdentity MODULE-IDENTITY
        LAST-UPDATED "200303172342Z"
        ORGANIZATION
                "Sun Microsystems, Inc.
                 4150 Network Circle,
                 Santa Clara, CA 95054 USA"
        CONTACT-INFO
                "http://www.sun.com/service/contacting/index.html"
        DESCRIPTION
                "N1 Provisioning Server MIB."
::= { enterprises 12816 }


-- The root of N1 Provisioning Server MIB Tree.

SMI          OBJECT IDENTIFIER ::= { enterprises 12815 }

-- The version 1 branch.

version1             OBJECT IDENTIFIER ::= { SMI 1 }

-- A tspr.log message.

tsprlogmessage  NOTIFICATION-TYPE
        OBJECTS { severity,
                type,
                farmid,
                message,
                originatingIPAddress }
        STATUS     current
        DESCRIPTION
                "A message with type debug."
        ::= { ifabricLog 11 }

-- A message sent on change of state of a node monitored by the N1 Provisioning Server.

nodeEvents  NOTIFICATION-TYPE
        OBJECTS { severity,
                message,
```

```
                originatingIPAddress }
        STATUS     current
        DESCRIPTION
                "A message sent on change of state of a node monitored by the N1 Provisioning Server."
        ::= { monLog 11 }

-- A message sent on change of Customer Monitoring Values detected  by the N1 Provisioning Server.

cmEvents  NOTIFICATION-TYPE
        OBJECTS { severity,
                message,
                originatingIPAddress }
        STATUS     current
        DESCRIPTION
                "A message sent on change of Customer Monitoring Values detected  by the N1 Provisioning Server."
        ::= { monLog 21 }

-- A message sent on change of the state of the Agent on a Resource Pool Server detected  by the N1
PS.

agentEvents  NOTIFICATION-TYPE
        OBJECTS { severity,
                message,
                originatingTimeStamp }
        STATUS     current
        DESCRIPTION
                "A message sent on change of the state of the Agent on a Resource Pool Server
detected  by the N1 Provisioning Server."
        ::= { monLog 31 }

-- A message received from a Control Plane Node that is to be forwarded.

controPoolTraps  NOTIFICATION-TYPE
        OBJECTS { message,
                originatingIPAddress,
                originatingTimeStamp }
        STATUS     current
        DESCRIPTION
                "A message received from a Control Plane Node that is to be forwarded."
        ::= { deviceTraps 11 }

-- A message received from a Resource Pool Node that is to be forwarded.

resoursePoolTraps  NOTIFICATION-TYPE
        OBJECTS { message,
                originatingIPAddress,
                originatingTimeStamp }
        STATUS     current
        DESCRIPTION
                "A message received from a Resource Pool Node that is to be forwarded."
        ::= { deviceTraps 21 }

-- The type of message.

type  OBJECT-TYPE
```

```
        SYNTAX     OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                "The type of message."
        ::= { trapVariables 1 }

-- The severity of the message.

severity  OBJECT-TYPE
        SYNTAX     OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                "The severity of the message."
        ::= { trapVariables 2 }

-- The Farmid for which the message is sent.

farmid  OBJECT-TYPE
        SYNTAX     OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                "The Farmid for which the message is sent."
        ::= { trapVariables 3 }

-- The actual body of the message.

message  OBJECT-TYPE
        SYNTAX     OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                "The actual body of the message."
        ::= { trapVariables 4 }

originatingIPAddress  OBJECT-TYPE
        SYNTAX     IpAddress
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                ""
        ::= { trapVariables 5 }

-- The TimeStamp of the original Trap.

originatingTimeStamp  OBJECT-TYPE
        SYNTAX     OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS     current
        DESCRIPTION
                "The TimeStamp of the original Trap."
        ::= { trapVariables 6 }
```

```
-- All notifications send for tsprlog messages.

ifabricLog  NOTIFICATION-GROUP
        NOTIFICATIONS { tsprlogmessage }
        STATUS     current
        DESCRIPTION
                "All notifications send for tsprlog messages."
        ::= { version1 100 }

-- All monitoring log messages.

monLog  NOTIFICATION-GROUP
        NOTIFICATIONS { agentEvents,
                cmEvents,
                nodeEvents }
        STATUS     current
        DESCRIPTION
                "All monitoring log messages."
        ::= { version1 200 }

-- All traps received from nodes that need to be forwarded onwards.

deviceTraps  NOTIFICATION-GROUP
        NOTIFICATIONS { controPoolTraps,
                resoursePoolTraps }
        STATUS     current
        DESCRIPTION
                "All traps received from nodes that need to be forwarded onwards."
        ::= { version1 300 }

trapVariables  OBJECT-GROUP
        OBJECTS { severity,
                type,
                farmid,
                message,
                originatingIPAddress,
                originatingTimeStamp }
        STATUS     current
        DESCRIPTION
                ""
        ::= { version1 400 }
END
```

CHAPTER 5

Backing Up and Restoring Components

This chapter describes backup and restoration of Infrastructure Fabric (I-Fabric) components. Sections in this chapter include:

Backup is an essential administration task and one that you will do frequently. This chapter provides general information for backing up and restoring I-Fabric components.

This chapter does not describe particular procedures for backing up because backup policies and procedures are usually specific to an organization. You should have a qualified database administrator (DBA) and a skilled systems administration staff to advise on or implement a local backup procedure. Ensure that backup and storage media rotation schedules are consistent with your best practices and take into account the rate of change within your particular I-Fabric.

General Backup

Back up the `/etc/opt/terraspring/tspr.properties` file immediately following installation for possible later use in system restoration. Because this file is static, you do not need to back it up again.

▼ General Backup

Steps

1. **Copy the Wiring Markup Language (WML) file to a file to be backed up.**

   ```
   wmld > dc.wml
   ```

   For more details on the `wmld` command see the `wmld` man page.

2. **Back up the list of images at regular intervals.**

   ```
   image -lv > images.out
   ```

Backing Up the Control Plane Server

The control plane might consist of one or two servers depending on whether your configuration includes a stand-alone image server or has the image server residing on the control plane server.

The control plane server contains the N1 Provisioning Server software , which fulfills the following roles:

- N1 Provisioning Server database (CPDB)
- Storage manager client (STMC)
- Farm manager
- Segment manager
- DHCP and DNS services
- Monitoring manager
- Control Center

Caution – Regularly back up the N1 Provisioning Server components, especially the CPDB.

The strategy for backing up a control plane server is to designate a server that is not part of the I-Fabric as a backup server and connect it to the control plane network. A variety of commercial backup software is available for performing backup. The only requirement is that it must be able to back up to Oracle or Postgres databases. When you install your N1 Provisioning Server, an IP address and port are designated for the backup server. See the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide* for details.

The following section details the requirements for backing up and restoring a control plane server.

▼ To Configure a Backup Server

To configure a backup server perform the following steps:

Steps

1. **Select a port on the control plane switch to connect a backup server.**
2. **Assign that port into virtual local area network (VLAN) 9.**
3. **Connect an appropriate tape backup device directly to the backup server.**
4. **The IP address `x.y.z.212` has been reserved for the backup server. N1 Provisioning Server software provides the `x`, `y`, and `z` values. The `x.y.z` value is the subnet your N1 Provisioning Server uses. Configure the backup server to use this address.**
5. **Install and configure backup server software on the backup server according to the supplier's instructions and your local requirements.**

Note – The backup server must not route or bridge any Ethernet traffic between VLAN 9 and any other network to prevent unauthorized access to the N1 Provisioning Server network.

Backing Up a CPDB

Every time a farm changes the CPDB changes. Therefore, you should perform frequent, regular backups of the CPDB.

Caution – Restoring the CPDB restores the associated farms to the state of the last backup. Farm changes made subsequent to the most recent backup are lost.

Run the `backupdb` command from the CPDB to back up the CPDB. For details on how to use this command, see the `backupdb` man page.

Note – Backing up the CPDB does not back up the Control Center. See "Backing Up the Control Center Database and the CPDB" on page 138 for details on backing up both, the CPDB and the Control Center, at the same time.

Backing Up the Control Center Database and the CPDB

The Control Center consists of software executable files and a database. Because the software files do not change after they have been installed, you can reinstall the software executables from the DVD to restore them.

This section describes backup of the Control Center and CP server database components running on either an Oracle or a Postgres database depending on which one you are using. The database stores all the data and state information for the Control Center. This is a general description of the Control Center and CPDB component backup process. For details on how to back up an Oracle or Postgres database, refer to the respective database manufacturer's documentation.

Note – We recommend backing up the Control Center, CPDB, Control Center database, data files, control files, `redo` logs, `init.ora` files, and password files at regular intervals.

Note – Database backup and restoration is a complicated task. Therefore, we strongly recommend that an experienced Oracle DBA perform this task.

Backing Up a Control Center and CP Server

This section describes the backup of the Control Center and CP server components running on either an Oracle or a Postgres database depending on which one you are using. The database stores all the data and state information for the Control Center. This is a general description of the Control Center and CPDB backup process. For step-by-step procedures for the full Oracle or Postgres database backup refer to the respective database manufacturer's documentation.

Back up the Control Center, CPDB, Control Center database, data files, control files, `redo` logs, `init.ora` files, `{postgres_data}`, and password files at regular intervals.

Note – Full database backup and restoration is a complicated task. Therefore, an experienced DBA should perform this task.

▼ To Back up the Control Center and Control Plane Server Components

Steps

1. **To back up the control plane server database run the following commands.**

   For backing up the control plane server component running on an Oracle database:

   `/opt/terraspring/sbin/backupdb -o ${Oracle_Home}` *backup file name*

   For backing up the control plane server component running on a Postgres database:

   `/opt/terraspring/sbin/backupdb -o ${Postgres_Home}` *backup file name*

2. **To back up the Control Center database run the following commands.**

   For backing up the Control Center component running on an Oracle or a Postgres database:

   `/opt/terraspring/gwdb/bin/backupDB.sh`

Control Center and CPDB File Locations

For the Oracle database, set the `ORACLE_HOME` variable to a directory, for example, `/opt/oracle/OraHome1`. By default, the Control Center and CPDB data files and control files for the Oracle database reside in the directory `/oracle_data1/tsprdb/DATA`.

The Postgres database is bundled with the N1 Provisioning Server software package and the default location of all the data and control files is `POSTGRES_HOME/data`.

Note – Because the data file and control file locations are configurable, your precise file path may vary.

Data and Control File Location

The data and control file location is configurable.

▼ To Back Up Data and Control Files

Steps

1. **Back up the following Control Center file:**

   `/opt/terraspring/sunone/domain_gw/server_gw/config/cc.properties`

2. **Back up all files in the following directories:**

```
/var/opt/terraspring/gw
/var/opt/terraspring/gwdb
```

Backing Up an Image Server

The image repository can run either on the control plane server or on separate server. If you run the image repository on a separate server, use a file transfer protocol (FTP) server. Images are stored as files. Use standard file system backup procedures as appropriate for your system environment to back up the image server.

Restoring a CPDB

Before restoring a CPDB, ensure that no changes to farm states have been made other than in the database. Otherwise, the database is no longer in sync with the rest of the system.

Caution – If you have created several backups, ensure that you restore from the most current backup file, otherwise the servers will be out of synchronization with the database. Consequently you will lose your most recent changes.

Run the `restoredb` command from the CPDB to back up the CPDB. For details on how to use this command, see the `restoredb` man page.

Note – Restoring the CPDB does not restore the Control Center. See "Restoring a Control Center" on page 142 for details on restoring both, the CPDB and the Control Center, at the same time.

Restoring the Control Center and the CP Server

Restoring to the point of failure on an Oracle or Postgres database requires the following information:

- A database backup for the files that were lost. This backup can be either an online table space backup as described in the previous section or an offline database backup.
- For an Oracle database, a set of archived `redo` logs for the period of time following the backup. In case of an online backup you also need the `redo` logs created during the time of backup.

For an Oracle database, the database looks for archived `redo` logs in the location specified by the `LOG_ARCHIVE_DEST` parameter of the `init`*db-name*`.ora` file. Alternatively, you can specify a location for the archived `redo` logs manually by using the `from` clause in the `alter database recover...` statement.

▼ To Restore the Control Center and the CP Server Component Running on Either an Oracle or Postgres Database

Steps

1. **Run the command:**

   `/opt/terraspring/sbin/restoredb` *restored file name*

   Note – The parameter *restored file name* refers to the backup file created during backup of the database.

   ---

2. **Run the `configure.sh` script and provide the required parameters when prompted.**

   `/opt/terraspring/gwdb/bin/configure.sh`

   For a Postgres database:

```
DB_TYPE                      Postgres (case-sensitive)
       DB Admin User:                postgres
       DB Admin Password:            postgres
       DB User:                      tcc
       DB Password:                  tcc
       DB HOst Name:                 localhost
       DB Instance Name:             tsprdb
       DB Instance Listening Port:   5432
       DB Backup Directory:          backup direcotry name
       DB Backup Filename:           backup file name
```

   For an Oracle database:

```
DB_TYPE                      Postgres (case-sensitive)
       DB Admin User:                system
       DB Admin Password:            manages
       DB User:                      tcc
```

```
DB Password:                  tcc
DB HOst Name:                 localhost
DB Instance Name:             tsprdb
DB Instance Listening Port:   1521
DB Backup Directory:          backup direcotry name
DB Backup Filename:           backup file name
```

3. **To create the database type:**

   `/opt/terraspring/gwdb/bin/createDB.sh -y`

4. **To restore the database type:**

   `/opt/terraspring/gwdb/bin/restoreDB`

Restoring a Control Center

Restore the Control Center by reinstalling the software from the DVD and then reinstalling the applicable backup files.

Restoring an Image Server

Use standard restoration procedures as appropriate for your system to restore an image server.

CHAPTER 6

Error Messages

This chapter describes the critical farm and system error messages with an error ID. The error messages described in this chapter consist of the following categories:

- Core system errors: error ID 1001 through 1022
- Control plane database (CPDB) errors: error ID 2001 through 2043
- Hardware abstraction layer (HAL) errors: error ID 3000 through 3950
- Monitoring errors: error ID 4015 through 4042
- Domain Name Service (DNS) and Dynamic Host Configuration Protocol (DHCP) errors: error ID 5000 through 5021
- Farm management errors: error ID 6100 through 6823 and 9300 through 9316

Note – The error message listing in this chapter is not all inclusive. It lists only the most important messages that you might encounter.

Error Message Overview

During the course of administering an Infrastructure Fabric (I-Fabric), you will encounter error messages. This chapter provides messages, descriptions, and actions required to solve a problem or to gather additional information before reporting a problem. You can locate messages in the log files.

The severity of error messages ranges from merely a minor debug message to a critical alert. All messages from N1 Provisioning Server software are logged in the `/var/adm/tspr.debug` log file. All messages from third-party components (for example, drivers) are logged in the `/var/adm/messages` log file.

As the message is written to the log file, the message is categorized as `okay`, `debug`, `warning`, or `critical`. The following example shows how it appears in the log:

```
Oct 30 00:16:47 sp4 java[ 506]: [ID 289794 user.info] TSPR [sev=okay] [apps=770034] TCPEventHandler:dispatch.
```

Note – [`sev=okay`] denotes the message's severity.

N1 Provisioning Server software logs messages in two separate log files:

- `okay, warning`, and `critical` messages with an error ID are logged in the `/var/adm/messages` log file.
- Along with the `okay, warning`, and `critical` messages, the `/var/adm/tspr.debug` log file also includes the debug messages, which are used for debugging system failures.

Error Message Descriptions

The following table lists messages by their error number. Each message is described in detail and corrective actions are provided. Words enclosed in curly brackets, for example `{application_ID}`, denote variables that are substituted for meaningful values when an error occurs.

Core System Errors (ID 1001 through 1022)

`1001`
`property` *property-name* `is invalid, using default`

Description: The property entered in the `tspr.properties` file for a device is invalid. The system is using the default value.

Solution: Enter a valid value for this property in the `tspr.properties` file or remove the line entirely.

`1002`
`cannot get database access`

Description: The device cannot access the CPDB.

Solution: Ensure that the database settings in `/etc/opt/terraspring/tspr.properties` are correctly pointing to the database. Check that the database is running. See the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide* for details.

`1003`
method name `with null arguments`

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

`1004`
`not registered`

Description: The process is not yet registered.

Solution: Restart all the processes on the control plane server, most important, segment manager. Also ensure that the database is not corrupted by, for instance, manual manipulation of data.

`1005`
application-ID `already registered`

Description: The process has already been registered

Solution: Stop and restart the segment manager.

`1006`
`invalid entry in application table with` *application-ID*

Description: The application table in the CPDB contains an invalid entry for the application.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

`1007`
`cannot dequeue or update request (`*request ID*`)`

Description: The server is unable to process or update the specified request ID.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

`1008`
Object name `not serializable`

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

`1009`
`invalid destination aid`

Description: The database contains an invalid entry of a destination application ID (AID) associated with a particular process.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

`1011`
`cannot send event to` *destination-aid* `(`*exception-details*`)`

Description: A process is unable to send an event to another process with *destination_aid*.

Solution: Ensure that the destination application is running properly.

1012
invalid request id

Description: An invalid request ID has been provided.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

1013
timeout expires

Description: An execution of a request has timed out.

Solution: A request performs a task. When this task is being executed, a timeout value is usually associated with the request. When the time since execution exceeds this value, this message is printed. Ensure that the destination application is running. Also ensure that no blocked requests for the destination application exist. Try to reissue the command.

1014
application-ID: max number of digit exceeded

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

1015
application-ID: not a digit

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

1016
Invalid IP address: *IP-address*

Description: An invalid IP address has been provided.

Solution: Check that all IP addresses in the CPDB have been correctly entered.

1017
method name:*value* invalid argument

Description: An invalid argument has been entered.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

1022
Could NOT create the LogHandler, *classname*.

Description: An invalid `LogHandler` has been entered for the `com.terraspring. core.log.exthandler.`

Solution: Ensure that the `exthandler` property is set to a valid class name that implements the `LogHandler` interface.

CPDB Errors (ID 2001 through 2043)

2002

Not enough *resource-name* resources available to complete allocation.

Description: The farm requires resources that are currently unavailable.

Solution: Place more devices into the resource pool or wait for more resources to become available.

2005

object-name value is not in correct state for this operation.

Description: The farm is currently not in the correct state for executing the specified operation.

Solution: Ensure that the farm is allowed to make the transition intended (refer to Chapter 7).

2008

Request *request-ID* could not be materialized from the database.

Description: This could be caused by incompatible objects on the client and server side.

Solution: Ensure that the packages are up to date. Delete the request and restart the Segment Manager (refer to Chapter 7). Retry the operation or reissue the command that previously failed.

2009

Invalid database connection parameters. Check the tspr.properties file.

Description: The `/etc/opt/ terraspring/ tspr.properties` file has an invalid entry.

Solution: In the `tspr.properties` file, check the lines that begin with `com.terraspring.db. CpdbConnection` and enter the correct connection parameter.

2010

More than one Segment Manager is running on the host.

Description: You may only have one segment manager running.

Solution: Stop all the segment managers by issuing `/opt/terraspring/sbin/sm -stop`. Restart the segment manager using `/opt/terraspring/sbin/sm -start`.

2011

Invalid application state for this operation.

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

2016

Image size is greater than the requested disk size.

Description: Internal error.

Solution: Ensure that the specified image is of a smaller or equal size than the specified disk. Choose a smaller image or increase the disk size.

2017
Invalid special device type: *type*

Description: The type must be BACKHAUL.

Solution: This message is associated with the farm's Farm Markup Language (FML) and Wiring Markup Language (WML). Contact http://sun.com/service/contacting/index.html for assistance.

2031
IP does not belong to subnet

Description: The specified IP address does not belong to the specified subnet.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

2033
Invalid deQueue request status: *request-ID*

Description: The queue is in QUEUED, QUEUED_BLOCKED, or INPROGRESS state.

Solution: Refer to Chapter 7 for details on dealing with requests.

2035
Duplicate MAC address *mac-address* cannot be added

Description: The media access control (MAC) address is already in use.

Solution: Assign a MAC address that is not yet in use.

2036
Subnet cannot be added because it overlaps with an existing one

Description: This subnet is already defined.

Solution: Assign a subnet that is not yet defined.

2037
CPDB server is probably dead or not allowing connections

Description: The CPDB connection is failing.

Solution: Check whether the server is running. If not, start it.

Check the listener for the network connection. If the connection has failed, reestablish the network connection.

2038
Farm name: *farm-name* does not follow the DNS domain label format

Description: An invalid DNS name has been specified.

Solution: DNS names accept only alphanumeric characters and the minus (-) sign. Specify a proper DNS name.

2039
Customer name: *customer-name* does not follow the DNS domain label format

Description: An invalid DNS name has been specified.

Solution: DNS names accept only alphanumeric characters and the minus (-) sign. Specify a proper DNS name.

2040
CP name cannot be null

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

2041
Monitor Node with *monitoring_node* not found

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

2042
IP Address *IP_address* is not a valid IP

Description: An invalid IP address has been assigned.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

Hardware Abstraction Layer (HAL) Errors (ID 3000 through 3950)

3000
Missing attributes *attribute-name*

Description: An attribute, such as an IP address for the server or a password, is missing in the CPDB.

Solution: Use the `device` command to add the missing attribute to the CPDB. For details on how to use this command see the `device` man page.

3100
Network interface can only be configured by root.

Description: The Network Virtual Interfaces (NVIs) cannot be created because the farm manager is not run as `root`.

Solution: Ensure that the user has root access and retry the operation, or reissue the command that previously failed. Refer to the Solaris 8 documentation for details on root access.

3101
Solaris network interface is already in use.

Description: The NVI intended to be created has already been created. This error can be caused by resources not being cleared up completely after a failed farm operation.

Solution: Refer to Chapter 7 for details on how to clear up NVIs and such resources.

3102
ifconfig failed on Solaris

Description: `ifconfig` is a standard Solaris command. Failure to issue this command implies some misbehavior in the operating environment.

Solution: Ensure that the Solaris software has been set up properly.

3103
unknown exit code from ifconfig

Description: `ifconfig` is a standard Solaris command. Failure to issue this command implies some misbehavior in the operating environment.

Solution: Ensure that the Solaris software has been set up properly.

3202
Login/passwd incorrect

Description: Cannot log into a device, such as a server.

Solution: Ensure that the user names and passwords on the database are current.

3204
Operation not supported on Solaris

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

3300
methodname The MAC address for *key* not found on device

Description: The MAC address for the device is missing.

Solution: This message is associated with the farm's WML. Contact http://sun.com/service/contacting/index.html for assistance.

3500
Warning: Cannot initialize interfaces on *ifname* (*error*)

Description: The system is unable to bring up an interface, such as a card or driver.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

3600
SQLException: *exception*

Description: Database error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

3700
This feature is not supported.

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

3800
Physical Device is either null or not valid

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
3900
Invalid Power Controller Port
```

Description: Inconsistency with the ports and the power controller specified in the database.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
3950
Could not power down the specified port.
```

Description: Power controller could not power down the device on a port.

Solution: Check the physical device and ensure accessibility and its connection to the power controller.

Monitoring Errors (ID 4015 through 4042)

```
4015
Monitoring [start] Error: No database connection.
```

Description: The database connection is unavailable now.

Solution: Check the CPDB connection and redo the operation.

```
4016
Monitoring [start] Error: No farm id.
```

Description: The farm ID is not valid. The CPDB has incorrect values for this farm.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
4017
Monitoring [start] Error: No node id.
```

Description: The node ID is not valid. The CPDB has incorrect values for this node.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
4019
Monitoring [start] Warning: Empty MML.
```

Description: The Monitoring Markup Language (MML) file is empty.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

```
4022
Monitoring [start] Error: Database error occurred (exception details).
```

Description: A database error occurred while accessing MML from the CPDB.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
4024
Monitoring [start] Error: Error occurred (exception-details).
```

Description: A monitoring exception occurred.

Solution: View details of the message, then contact http://sun.com/service/contacting/index.html for assistance.

`4028`
`MML Parsing Error: Parsing error occurred (`*exception details*`).`
Description: The MML is invalid.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

`4029`
`MML Parsing Error: I/O error occurred (`*exception-details*`).`
Description: An I/O error occurred during parsing of the MML.

Solution: View details of the message, then contact http://sun.com/service/contacting/index.html for assistance.

`4030`
`MML Parsing Error: MML content is empty.`
Description: The MML is empty.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

`4033`
`MML Parsing Warning: Invalid threshold value` *threshold-volume*
Description: The MML is wrong.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

`4034`
`MML Parsing Warning: Illegal` *node-name* `child node` *node-ID* `found.`
Description: The MML is wrong.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

`4035`
`MML Parsing Warning: Invalid monitor` *monitor-name* `encountered.`
Description: The MML is wrong.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

`4036`
`MML Parsing Warning: No value for` *integer*`.`
Description: The MML is wrong.

Solution: This message is associated with the farm's MML. Contact http://sun.com/service/contacting/index.html for assistance.

`4037`
`Remote Command for shutdown invoked successfully, but shutdown has a problem, check the console for device` *device-ID*

Description: The system in unable to shut down the resource pool server.

Solution: Ensure that the server can be pinged and the N1 Provisioning Server agent is active.

4038
Shutdown of device *device-ID* failed. Interrupted exception while pinging the device.

Description: The Java Virtual Machine (JVM) is unable to execute sleep mode while waiting for the device to be stopped.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

4039
MML Parser Error: Cannot create parser with requested configuration (*parser-configuration_ exception*)

Description: An error exists in the MML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

4040
MML Parser Warning: The object *warning* has an invalid value (*threshold*)

Description: An error exists in the MML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

4042
Monitoring [*start*] Error: Illegal access error occurred (*illegal_access_exception*)

Description: Unable to open the connection to the CPDB.

Solution: Ensure that the CPDB is running and accessible.

Domain Name Service (DNS) and Dynamic Host Configuration Protocol (DHCP) Errors (ID 5000 through 5021)

5000
unable to look up localhost info

Description: The local server's IP address cannot be determined or obtained.

Solution: Ensure that the local server has an IP address.

5001
IOException *exception details*

Description: Error writing to disk.

Solution: Ensure that both `/etc/named.conf` and `/etc/dhcpd.conf` are not write-protected.

5002
IOException cannot restart DHCP server

Description: Internal error

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

5003
unable to instantiate subnet in dhcp sync

Description: Error when allocating internal objects.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

5004
problem reading the "leasetime" property

Description: The `leasetime` property in `/etc/opt /terraspring/tspr.properties` is not properly formatted or unreadable.

Solution: Ensure that the line `com.terraspring.dns.leasetime` exists in the property file, and that it is followed by an equal sign (=) and then an integer.

5005
Interrupted Exception: *exception details*

Description: The restart of the `dhcpd` daemon has been interrupted.

Solution: Issue `dhcpdmod sync` manually. If this fails, check `/etc/dhcpd.conf` and try running `dhcpd` manually in debug mode to see details of the failure.

5007
encountered a null host object while writing dhcp config

Description: Internal error.

Solution: Restart the segment manager and retry the operation, or reissue the command that encountered this error.

5008
cannot restart DNS server

Description: The restart of the `BIND` daemon has failed.

Solution: Issue the `dnsmod sync` command manually. If this fails, check `/etc/named.conf`. Try running `named` manually in debug mode to see details of the failure. Read `/var/adm/messages`.

5011
SQLException: *exception details*

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

5012
SQLException: encountered a null *attribute* object while writing dns config

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

5013
SQLException: encountered a bad *attribute* object while writing dns config

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

5016
DNS Error: Cannot create zone file directory *directory-name*.

Description: For DNS to function properly, a directory must be created for DNS database files. The default directory is `/var/opt/ terraspring/named`.

Solution: Ensure that the disk is not write-protected.

5017
DNS Error: Zone file directory *directory name* is not a directory.

Description: For DNS to function properly, a directory must be created for DNS database files. The default directory is `/var/opt/ terraspring/named`.

Solution: Ensure that the directory has not been previously created as a file. If so, delete the file and create the directory manually.

5018
DNS Warning: Error occurred (*exception details*).

Description: An unexpected exception occurred internally.

Solution: Restart the Segment Manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

5019
DHCP Warning: Error occurred (*exception details*).

Description: An unexpected exception occurred internally.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

5020
Unable to open file:

Description: The system is unable to open the DHCP configuration file.

Solution: Check the file permissions of `/etc/dhcp.conf`.

5021
Unable to write to file:

Description: The system is unable to write to the DHCP configuration file.

Solution: Check the file permissions of `/etc/dhcp.conf`.

Farm Management Errors (ID 6120 through 6823 and 9300 through 9316)

6120
FixedServerConfig.:{*error-details*}

Description: Internal Error.

Solution: Check the logs for earlier messages that might help explain the root of the problem.

6130
{*method-name*}: No common subnet found between lb {*load-balancer-ID*} and bind elements.

Description: This message is associated with the farm's FML and WML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6131
LBConfig. {*method-name*}: No such subnet {*subnet-ID*} found on lb {*load-balancer-ID*}.

Description: This message is associated with the farm's FML and WML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6132
LBConfig. {*method-name*}: Cannot instantiate class: {*exception-details*}

Description: Internal error.

Solution: Restart the segment manager and retry the operation or reissue the command that previously failed.

6133
LBConfig. {*method-name*}: Cannot access database {*exception-details*}

Description: Database connection error.

Solution: Ensure that the database is running and the connection values are properly specified in `/etc/opt/terraspring/tspr.properties`.

6134
LBConfig. {*method-name*}: No such if {*logical-network-interface-name*} found for device *device name*.

Description: This message is associated with the farm's FML and WML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6135
LBConfig: {*method-name*}: bind list empty for lb {*load-balancer-ID*}

Description: This message is associated with the farm's FML and WML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6136
LBConfig {*method-name*}: no interface found for lb {*load-balancer-ID*}

Description: This message is associated with the farm's FML and WML.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6137
LBConfig: no interface found for bind element: {*bind-element*}

Description: The load balancer cannot find the interfaces to the `bind` device.

Solution: This message is associated with the farm's FML. Contact http://sun.com/service/contacting/index.html for assistance.

6300
Farm {*farm-id*} cannot access the database during allocation (*exception-details*).

Description: Database connection error.

Solution: Ensure that the database is running and the connection values are properly specified in /etc/opt/terraspring/tspr.properties.

6307
SQL Error: Farm {*farm-ID*} cannot re-allocate physical devices (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6308
Config Error: Insufficient resource to re-allocate physical devices in farm {*farm-ID*} (*exception-details*).

Description: No more free devices are available for reallocation to the farm.

Solution: Free up resources that are occupied by other farms.

6309
Config Error: Error occurred while re-allocating physical devices in farm {*farm-ID*} (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6310
SQL Error: Farm *farm-ID* cannot replace failed devices (*device-ID*).

Description: The farm manager was trying to replace a failed device but failed.

Solution: Use the device command to check the device manually.

6311
Config Error: insufficient resource to replace failed device in Farm *farm-ID* (*exception-details*).

Description: No more free devices are available for replacing the failed device.

Solution: Free some resources by deactivating unused farms or add more devices to the resource pool.

6312
Config Error: Error occurred while replacing failed devices in Farm *farm-ID* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6313
Config Error: Database error occurred while allocating Farm *farm-ID* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6314
Config Error: Insufficient resource to allocate Farm *farm-ID* (*exception-details*).

Description: No more free devices are available for allocation to the farm .

Solution: Free some resources by deactivating unused farms or add more devices to the resource pool.

6315
Config Error: Invalid argument error while allocating Farm *farm-ID* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6316
Config Error: Error occurred while allocating Farm *farm-ID* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6317
FML Error: Subnet *subnet-ID* referred in the Interface *interface-ID* is not defined in the FML.

Description: The subnet ID is incorrect or does not exist.

Solution: From the Control Center ensure that the correct subnet ID was defined. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details. If the error persists after redefining the correct subnet ID from the Control Center, contact Customer Support for assistance.

6318
FML Error: Subnet {*subnet-ID* }referred in VIP {*VIP ID*} for load balancer {*load-balancer-ID*} is not defined in the FML.

Description: The subnet ID is incorrect or does not exist.

Solution: From the Control Center ensure that the correct subnet ID was defined. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details. If the error persists after redefining the correct subnet ID from the Control Center, contact Customer Support for assistance.

6319
FML error: Subnet {*subnet-ID*} referred in the intferface {*interface-ID*} of load balancer {*load-balancer-ID*} is not defined in the FML.

Description: The subnet ID is incorrect or does not exist.

Solution: From the Control Center ensure that the correct subnet ID was defined. See the *N1 Provisioning Server 3.1, Blades Edition, Control Center Management Guide* for details. If the error persists after redefining the correct subnet ID from the Control Center, contact Customer Support for assistance.

6320
Alloc Error: Cannot find the FML file for Farm *farm-ID* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6329
FML Parser: Null or Empty content

Description: The FML specifies an invalid empty string.

Solution: Check the FML for syntax accuracy.

6330
FML Parser: Parsing Error (*error*)

Description: A parsing error occurred when the FML was parsed.

Solution: Check the FML for syntax accuracy.

6331
FML Parser: I/O error (*error*)

Description: An I/O error occurred during FML parsing.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6332
FML Parser: DTD error (*error*)

Description: An error occurred while retrieving or using the FML document type definition (DTD) file.

Solution: Check the permissions of the FML DTD file found under `/opt/terraspring/etc/fml.dtd`.

6333
FML Parser: *device-name* is not a valid DNS name

Description: An invalid device name has been entered into the FML.

Solution: Change the device name to a valid DNS name through the Control Center or modify the FML.

6334
FML Parser: Cannot create parser with requested configuration

Description: An FML parser error occurred.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6335
FML Parsor: Unknown Protocol.

Description: Internal error.

Solution: This message is associated with the farm's FML. Contact http://sun.com/service/contacting/index.html for assistance.

6340
Wiring Error: Invalid device id *device-ID* in FAILED state, needs reallocation.

Description: A failed device has been allocated to a farm.

Solution: Check the database error log for related messages that might help explain the root of the problem.

6341
Wiring Error: Invalid internal state for device id *device-ID* with FREE device in wired phase.

Description: Inconsistencies exist in the database.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

6342
Wiring Error: Invalid internal state for device id *device-ID* marked OFF with power.

Description: Inconsistencies exist with the physical device and its state in the database.

Solution: Properly shut down the device and try to reactivate the farm.

6343
Wiring Error: Cannot control power of device id *device-id* (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6344
Wiring Error: Interface *interface-ID* not found in device id *1*.

Description: Inconsistencies exist in the database.

Solution: This error concerns the device tables that hold all the information on all devices. Refer to Chapter 7 for details on how to remedy database problems.

6345
Wiring Error: I/O error while moving cpu *cpu-ID*:*1* to VLAN *VLAN-ID* (*exception-details*).

Description: Error configuring the switch to move the CPU to the indicated virtual local area network (VLAN).

Solution: Check the accessibility of the switch connected to the CPU device.

6346
Wiring Error: I/O error while hiding cpu *cpu-ID*:*1* to VLAN *VLAN-ID* (*exception-details*).

Description: Error configuring the switch to move the CPU to the idle VLAN.

Solution: Check the accessibility of the switch connected to the CPU device.

6347
Wiring Error: Cannot find device id *device-ID*.

Description: Inconsistencies exist in the database.

Solution: This error concerns the device tables that hold all the information on all devices. Refer to Chapter 7 for details on how to remedy database problems.

6350
Wiring Error: Image not found for role *role-name*.

Description: The image specified in the FML does not exist in the database.

Solution: Use the `image` command to list all images. This error concerns the disk and image tables that collectively hold all the information on all images. Refer to Chapter 7 for details on how to remedy database problems.

`6352`
`Wiring Error: I/O error occurred during wiring (`*exception-details*`).`

Description: The CPU or disk might be inaccessible.

Solution: Check the physical CPU and disk as well as their connection to ensure accessibility.

`6353`
`Wiring Error: I/O error while unbinding disk` *error-details*`.`

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

`6354`
`Wiring Warning: Device` *device-id* `is marked ON with no power, resetting the DB state to OFF`

Description: The device was marked `ON` but was not powered on.

Solution: Ignore this error.

`6355`
`Wiring Error: Cannot get vlan` *VLAN-id* `for farm` *farm-id*

Description: The control plane VLAN cannot connect to the resource layer VLAN for a device in the farm.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

`6356`
`Wiring Error:Image (name:` *name*`, id:` *id*`) is not ready (`*error*`)`

Description: The specified image is not in the ready state.

Solution: Use the `image -p` and `image -r` commands to properly prepare and release the image, or select another image that has been properly prepared. See the `image` man page or Chapter 3 for details on how to use this command.

`6357`
`Wiring [`*methodName*`] Warning: Cannot power off device` *deviceid*

Description: An error occurred when the system tried to power off the device.

Solution: Use the `power` tool to check that the device can be powered off by the power controller.

`6360`
`Wiring Error: Cannot remove scsi bindings due to error` *error-details*

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

`6361`
`Wiring Error: Copying All Disks for LogicalStorage Device,` *deviceid*`, Failed due to error (`*error*`)`

Description: An error occurred while copying disks for a storage device.

Solution: Examine the logs for the cause of the failed copy operation. Make sure that the storage manager client (STMC) process is running.

```
6362
Wiring Warning: A Failed Server has No Snapshot Image, using the original Role Image
```

Description: A warning that a failed resource pool server has no snapshot image available to copy to the replacement server.

Solution: Ignore this message.

```
6363
Wiring Warning: A Standby Farm has No Snapshot Image, using the original Role Image
```

Description: A warning that a server that was in standby has no snapshot image available to copy to the new server.

Solution: Check that the local disk was properly snapshot to a standby image. Contact http://sun.com/service/contacting/index.html for assistance if an image is available and done.

```
6364
Wiring Error: Error Deleting the StandbyImage due to error (error).
```

Description: An error occurred while deleting the standby image after activating a farm in standby mode.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6365
Wiring Error: Chosen DiskImage can not be NULL for LogicalDisk [name=diskname, imid=imageId.
```

Description: An error occurred while retrieving the image specified for the logical disk.

Solution: Use the `image -l` command to check for the existence of the image.

```
6366
Wiring Error: Error Copying Global Disks due to error (error).
```

Description: An error occurred while copying an image to the global disks.

Solution: Examine the logs for a possible cause of the failed operation. Check the STMC for indications that might help explain the root of the problem.

```
6367
Wiring Error: Error Configuring Disks due to (error).
```

Description: An error occurred while configuring the disks.

Solution: Check the STMC for indications that might help explain the root of the problem.

```
6368
Wiring Error: Configuring switch fabric for logical device logical-device-ID failed.
```

Description: The switch fabric configuration in the fabric layer for a device failed.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6370
Dispatch Error: Cannot get NIC interface (*exception-details*).

Description: An interface specified as associated with a particular farm does not exist in the database.

Solution: Bring the farm to STANDBY and reactivate.

6371
Dispatch Error: Farm *farm-id* is already dispatched on *cp*

Description: A farm is already dispatched on the N1 Provisioning Server.

Solution: Ignore this message.

6373
Dispatch Error: DNS error (*error*)

Description: A DNS error occurred during dispatch.

Solution: Check the log for earlier error messages that might help explain the root of the problem.

6374
Dispatch Error: DHCP error (*error*)

Description: A DHCP error occurred during dispatch.

Solution: Check the log for earlier error messages that might help explain the root of the problem.

6375
Dispatch Error: Cannot find subnet id *subnet-ID*.

Description: A subnet specified as associated with a particular farm does not exist in the database.

Solution: A database inconsistency exists. Contact http://sun.com/service/contacting/index.html for assistance.

6377
Dispatch Error: Illegal argument error (*exception-details*).

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

6379
Insufficient resource (*exception-details*)

Description: Insufficient network interfaces available to complete deployment of the farm.

Solution: Free some of the network interfaces.

6380
FM Dispatch: Cannot find VLAN id for Subnet *subnet-ID* (VLAN name = *VLAN-name*).

Description: Inconsistencies exist between the database and the FML.

Solution: Compare the subnet ID specified in the FML to that in the database. Contact http://sun.com/service/contacting/index.html for assistance.

6381
FM Dispatch: NIC interface for IP address *ipaddr* does not exist on *CP*

Description: Internal error

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6384
Dispatch Warning: No subnet found.

Description: The farm has no subnet.

Solution: FML must declare at least one subnet. Contact http://sun.com/service/contacting/index.html for assistance if a subnet exists.

6385
Dispatch Error: Farm *farm-id* has no owner.

Description: The farm has no owner. The owner might have been removed manually from the CPDB.

Solution: Reactivate the farm. Doing so enables the segment manager to assign an owner automatically.

6386
Dispatch Error: Farm *farm-id* has a null *Runtime/-Logical* list.

Description: The farm has no associated logical resources during dispatch.

Solution: A database inconsistency exists. Contact http://sun.com/service/contacting/index.html for assistance.

6387
Idle Error: Unable to deregister device *deviceid*

Description: The system is unable to unregister the device from monitoring during idle state.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6388
FMDispatch Error: Cannot stop customer monitoring for farm *farm-id* device *devicename* (*deviceid*) due to error (*error*).

Description: The system cannot stop account monitoring for the device.

Solution: Make sure that the TSPR N1 Provisioning Server agent is running on the device.

6389
Idle Error: Farm *farm-id* has a null *Runtime/Logical* list.

Description: The farm has no associated logical resources during idle state.

Solution: A database inconsistency exists. Contact http://sun.com/service/contacting/index.html for assistance.

6390
Farm *farm-id* failed to dispatch.

Description: The farm failed to dispatch.

Solution: Check the log for earlier error messages that might help explain the root of the problem.

6392
Farm *farm-id* failed to idle.

Description: The farm failed to idle.

Solution: Check the log for earlier error messages that might help explain the root of the problem.

6400
Activation Error: Cannot register for monitoring (*exception-details*)

Description: A device cannot be registered to the monitoring manager for monitoring.

Solution: Refer to Chapter 7.

6401
Activation Error: Cannot activate device *device-name* (*device-id*).

Description: The farm failed to activate.

Solution: Check the log for earlier error messages that might help explain the root of the problem.

6402
Activation Error: Cannot shutdown device *device-ID* device id *device-ID* (*exception-details*).

Description: The device or its power controller might be inaccessible.

Solution: Check the physical device or its power controller to ensure accessibility.

6403
Activation Error: Cannot unconfigure device *device-ID* device id *device-ID* (*exception-details*).

Description: The device might be inaccessible.

Solution: Check the physical device's power controller to ensure accessibility.

6404
method name Warning: Cannot clean device *device-ID* (*exception-details*).

Description: The system cannot scrub the device.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed. Check the physical device.

6406
Activation Error: No physical device found for *logical-id*; cannot deregister with.

Description: No physical devices were found for the associated logical device.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6407
Arp cache of router(s) may need to be cleared.

Description: Informational message.

Solution: Check whether the ARP cache or router need to be cleared.

6408
method name Error: Cannot power on device *device-ID* (*exception-details*).

Description: An error occurred while trying to power on the device.

Solution: Check the physical device and ensure accessibility and its connection to the power controller.

6409
Activation Error: Waiting for UP Device Status has timed out.

Description: Some of the devices failed to come up and send an UP event to the farm manager.

Solution: Check whether the device can be pinged and make sure that the monitoring process is running on the CP.

6410
Activation Warning: Property *name* value *value* is invalid, using default (*default*).

Description: The property value specified is invalid, and a default value will be used.

Solution: Check the tspr.propperties file to see if the property entered is invalid.

6411
Activation Warning: Error occurred while notifying disk addition (*exception*).

Description: An error occurred while notifying monitoring of a new disk addition.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6430
Backup Error: SQL Exception while retrieving the DNS name for the fixed-server name=*server-name* id=*server-ID* of Farm *farm-ID* (*exception-details*).

Description: The database could not retrieve a valid DNS name for the specified server.

Solution: This error concerns the database table Request that holds all the requests. Refer to Chapter 7 for details on how to remedy database problems.

6460
FMShutdown Warning: Cannot power off device *device-ID* due to error (*error*).

Description: The system cannot power off the device.

Solution: Make sure that the device power controller is configured properly.

6461
method name Error: Cannot power off device *device-ID* (*exception-details*)

Description: The system is unable to power off the device.

Solution: Check the physical device.

6490
FMDeactivate Error: Disks cannot be unconfigured (*error*)

Description: The disk cannot be unconfigured.

Solution: Check the STMC for indications that might help explain the root of the problem.

6491
FMDeactivate Error: Disk *diskid* has error (*error*).

Description: An error occurred during farm deactivation.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6492
FMdeactivate Warning: Farm *fmid* cannot delete image *imgname* (*imgid*) of device *devname* (*devid*) due to error (*error*).

Description: Warning that an image cannot be deleted.

Solution: Check CPDB availability and rerun the command.

6493
FMDeactivate Error: Scrubbing all the Disks due to error (*error*)

Description: An error occurred while scrubbing the disks during deactivation or standby.

Solution: Check the STMC for indications that might help explain the root of the problem.

6494
Config Error: Cannot unconfigure the disk *disk-ID* (*exception-details*)

Description: The disk is in an invalid state, it must be in `CONFIGURED` state.

Solution: Use the `disk` command to change the state of the disk. See the `disk` man page.

6495
FM Warning: Scrubbing bypassed for disk (*disk-ID*)

Description: Disk scrubbing during deactivation is ignored because the `com.terraspring.cs.FM.scrubbing` property is set to false.

Solution: If you want disks to be scrubbed during deactivation, set the scrubbing property to `true`.

6496
FM Warning: Scrubbing bypassed for all disks in farm *farm-id*.

Description: Disk scrubbing during deactivation is ignored because the `com.terraspring.cs.FM.scrubbing` property is set to false.

Solution: If you want disks to be scrubbed during deactivation, set the scrubbing property to `true`.

6510
FMStandby Warning: Farm *0* cannot rename image *imgid* of device *devname* (*devid*) from *origName* to *newName* due to error (*error*).

Description: An error occurred while renaming a device's image during standby.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6511
FMStandby Error: Farm farm-id server serverName has failed Scrubbing Disk due to error (error).
```

Description: An error occurred while scrubbing a disk during standby.

Solution: Check the `tspr.debug` log for indications that might help explain the root of the problem.

```
6512
FMStandby Error: Farm farm-id server serverName has failed due to error (error).
```

Description: An error occurred during standby.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6540
Snapshot Error: Cannot find the source disk disk-ID on server server-ID of farm farm-ID.
```

Description: A snapshot has been issued on a nonexistent source disk or target.

Solution: The command was probably issued with the wrong parameters. Reissue the command using the correct parameters.

```
6541
Snapshot Error: Cannot find the destination disk with target disk-ID on host hostname of Farm farm-ID.
```

Description: A snapshot has been issued on a nonexistent destination disk or target.

Solution: The command was probably issued with the wrong parameters. Reissue the command using the correct parameters.

```
6543
Snapshot Error: Image image name already exists.
```

Description: An attempt to overwrite an existing snapshot image was made.

Solution: Delete the existing image with the `image` command or reissue the command using a different name for the image.

```
6544
Snapshot Error: Cannot create image disk.
```

Description: The destination image could not be created for the snapshot.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6545
Snapshot Error: Error occurred (error).
```

Description: An error occurred during a snapshot operation.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6546
Snapshot Error: Cannot quiesce fixed-server.
```

Description: An error occurred while trying to prepare the device for shutdown.

Solution: Check the log for earlier messages that might help explain the root of the problem.

6547
Snapshot Error: Cannot copy disk.

Description: An error occurred while trying to copy the disks during snapshot.

Solution: Check the log for earlier messages that might help explain the root of the problem.

6548
Snapshot Error: Cannot shut down fixed-server (name:*server-name*, id:*server-ID*) of farm *farm-ID* due to error (*error*).

Description: An error occurred while trying to power off the server for the snapshot operation.

Solution: Check the physical device and ensure accessibility and its connection to the power controller.

6549
Snapshot Error: Cannot re-enable monitoring for CPUDev: *cpu-dev-id* due to error (*error*).

Description: A snapshot error occurred while re-enabling monitoring for the device.

Solution: Ensure that the device is running and that it can be pinged. Ensure that the N1 Provisioning Server agent on the device is running.

6550
Snapshot Error: Cannot query Power for CPUDev: *cpu-dev-id* due to error (*error*).

Description: An error occurred while querying the power status for the device.

Solution: Check the physical device and ensure accessibility and its connection to the power controller.

6551
Snapshot Error: Cannot safely power off Server *server* because server is not monitored.

Description: Monitoring is not enabled so the system cannot send a remote procedure call (RPC) message to safely shut down the server.

Solution: Manually shut down the server.

6552
Snapshot Error: Cannot set the Last Snapshot Image: *imgageid* for LogicalDisk *diskid* due to error (*error*).

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

6553
Snapshot Error: Device *device-ID* is not a valid device ID in *farm-ID*.

Description: The snapshot operation cannot proceed because the device is not in the specified farm.

Solution: Check the resources allocated to the farm and correct the inputs to the snapshot.

6554
Snapshot Error: FixedServer *server-ID* of Farm *farm-ID* does not contain CPUs.

Description: Inconsistencies exist in the database.

Solution: This error concerns the device tables that hold all the information on all devices. Refer to Chapter 7 for details on how to remedy database problems.

6555
FM Error: Snapshot failed (farm id: *farm-ID*, hostname: *hostanme*, target: *disk-ID*, image name: *image-ID*, console id: *console-ID*) with the following message: *exception-details*.

Description: A snapshot operation failed.

Solution: Check the log for earlier messages that might help explain the root of the problem.

6580
FMMonitor Error: Error occurred.

Description: Error occurred during an `FMMonitor` operation.

Solution: Check the log for earlier messages that might help explain the root of the problem.

6581
FMMonitor Error: Cannot start monitoring for farm {fmid} on device {deviceid} due to error ({exception}).

Description: Monitoring for the device cannot be started.

Solution: Check that the monitoring service is running. Contact http://sun.com/service/contacting/index.html for assistance.

6582
FMMonitor Error: Cannot stop monitoring for farm {fmid} on device {deviceid} due to error {exception}.

Description: Monitoring for the device cannot be stopped.

Solution: Check that the monitoring service is running. Contact http://sun.com/service/contacting/index.html for assistance.

6583
FM Warning: Cannot quiesce storage device (name: {device name}, id: {deviceid}) due to error ({error}).

Description: The resource pool server cannot be prepared for shutdown.

Solution: Ensure that the N1 Provisioning Server agent is running on the resource pool server and that the device can be pinged.

6584
Monitoring Warning: Server {servername} cannot be notified of {adding/remove} disk {diskinfo} due to error ({error}).

Description: A problem with notifying the monitoring system of a change in a resource pool server's disk occurred.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6585
Monitoring Warning: Cannot find image for disk {diskid} due to error ({error}).
```

Description: Monitoring cannot find the image for the disk. Inconsistencies in the database exist.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6586
Monitoring Warning: Cannot find DNS name for device {deviceid} due to error ({error}).
```

Description: Monitoring cannot get the DNS name for the device.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6620
Farm {fmid}: needs to unblock an update request for failover of device {devid}.
```

Description: A request to replace a failed device has been queued and blocked for a device in the farm.

Solution: Check the device. If the device needs to be replaced, unblock the queued request.

```
6621
Error occurred when retrieving Farm {fmid}'s context onto CP {cp name}.
```

Description: An error occurred while trying to redispatch the farm onto the control plane (CP).

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6622
Farm {fmid}'s old InProgress Requests were Interrupted.
```

Description: An in-progress farm request was abruptly interrupted.

Solution: Reissue the request to the farm.

```
6623
Warning: Device {device ID} Device id {device ID} is not a network device,
   so cannot act as a gateway.
```

Description: The FML specifies that the device act as a gateway.

Solution: Ensure that a network device is specified as a gateway within the FML.

```
6661
FM Warning: Storage-device (name: {dev name}, id: {dev id}) has no power controller;
   cannot power it {errMsg}.
```

Description: The storage device has no power controller associated.

Solution: Use the `device -lv` command to check the device's connections to the power controller.

```
6662
FM Warning: Storage-device (name: {dev name}, id: {dev id} has no configuration class;
   cannot shut it down.
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6663
FM Warning: Cannot power off storage-device (name: {dev name}, id: {dev id})
   due to error ({error Msg}).
```

Description: An error occurred while powering off the storage device.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6664
FM Warning: Cannot power on storage-device (name: {dev name}, id: {dev id})
   due to error ({error Msg}).
```

Description: An error occurred while powering on the storage device.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6701
FarmManagerImpl: SQLException in method {method name}, Farm {farm ID}:
   ({Exception Details}).
```

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

```
6702
FarmManagerImpl: IllegalArgument Exception in method {method name}, Farm {farm ID}:
   ({Exception Details}).
```

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

```
6703
FarmManagerImpl: IllegalState Exception in method {method name}, Farm {farm ID}:
   ({Exception Details}).
```

Description: The farm manager for a farm has been instructed to make an illegal state transition.

Solution: Do not issue illegal state transition instructions.

For activating a farm, the following values are legal states:

- CREATED
- NEW CONFIG
- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- SHUTDOWN
- IDLE
- UPDATED

- STANDBY

For updating a farm, the following values are legal states:

- ACTIVE
- CREATED
- NEW CONFIG

If a previous attempt of updating a farm failed and no changes were made, the following states are also valid:

- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- UPDATED
- SHUTDOWN
- IDLE

For deactivating a farm, the following values are legal states:

- NEW CONFIG
- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- SHUTDOWN
- IDLE
- UPDATED
- UNWIRED
- STANDBY

For deleting a farm, the following values are legal states:

- CREATED
- DEACTIVATED

For putting a farm into standby mode, the following values are legal states:

- ACTIVE
- SHUTDOWN
- IDLE
- UNWIRED

```
6704
FarmManagerImpl: Failed to reach state {State No}.
```

Description: The farm manager has encountered errors.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6705
FarmManagerImpl: Failed to get subnet number from {subnet}.
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6706
FarmManagerImpl: Error occurred in method {method} farm {fmid}: ({error})
```

Description: Internal error.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6707
FM Error: Cannot snapshot image {imageid} on node {dnsname} target {target} of
   farm {fmid} due to error ({error}).
```

Description: An error occurred while taking a snapshot of an image on the node's target disk.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6708
FM Error: Cannot snapshot configuration for device {dev id} of farm {farm id}
   due to error ({error}).
```

Description: An error occurred while taking a snapshot the configuration of a device.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6801
SQLException in method {method name}, Farm {farm ID}, at host {hostname}:
   {exception details}.
```

Description: Internal error.

Solution: Restart the segment manager (refer to Chapter 7) and retry the operation, or reissue the command that previously failed.

```
6802
IllegalArgument Exception in method {method name}, Farm {farm ID}.
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6803
IllegalState Exception in method {method name}, Farm {farm ID}.
```

Description: The farm manager for a farm has been instructed to make an illegal state transition.

Solution: Do not issue illegal state transition instructions.

For activating a farm, the following values are legal states:

- `CREATED`
- `NEW CONFIG`

- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- SHUTDOWN
- IDLE
- UPDATED
- STANDBY

For updating a farm, the following values are legal states:

- ACTIVE
- CREATED
- NEW CONFIG

If a previous attempt of updating a farm failed and no changes were made, the following states are also valid:

- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- UPDATED
- SHUTDOWN
- IDLE

For deactivating a farm, the following values are legal states:

- NEW CONFIG
- ALLOCATED
- WIRED
- DISPATCHED
- ACTIVE
- SHUTDOWN
- IDLE
- UPDATED
- UNWIRED
- STANDBY

For deleting a farm, the following values are legal states:

- CREATED
- DEACTIVATED

For putting a farm into standby mode, the following values are legal states:

- ACTIVE
- SHUTDOWN
- IDLE
- UNWIRED

```
6804
PostService Exec found Farm {farm ID} with error {error NO} in {State NO}.
```

Description: Internal error.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6808
SM: on {CP server ID} can not start, ({Segment Manager ID}) not removed.
```

Description: A segment manager is already running.

Solution: Stop and restart the segment manager.

```
6809
SM: {Aid} cannot create ({Aid Filename}).
```

Description: A file must be created when the segment manager starts and this file cannot be created.

Solution: Check that the disk has read/write permissions.

```
6812
SM: moveFarm ({fmid}) to same CP ({cp}) not allowed.
```

Description: An attempt was made to move the farm to the same CP on which it already resides.

Solution: Use the `farm -m` command to specify another CP.

```
6813
SM: moveFarm ({fmid}) to CP machine ({cp}) does not have enough interfaces.
```

Description: The CP does not have enough interfaces to support the move of the farm.

Solution: Use the `farm -m` command and move the farm to an CP with enough free interfaces.

```
6814
SM: moveFarm ({fmid}) failed to remove the CP Context from SM {sm}: {error}.
```

Description: The source CP had an error removing the farm's CP interface from the CP.

Solution: Check the source CP's log for earlier messages that might help explain the root of the problem.

```
6818
{method}: SM {smname} encountered error ({error}).
```

Description: The segment manager process experienced a monitoring error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6819
{method}: SM {smname} has no timestamp.
```

Description: The segment manager has no timestamp.

Solution: Ensure that the segment manager is running. Stop and restart the segment manager if necessary.

```
6820
{method}: SM {smname} encountered SQL error ({error}).
```

Description: Internal error.

Solution: Check the database. Stop and restart the segment manager if necessary.

```
6821
{method}: SM {smname} encountered illegal access error ({error}).
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
6822
FMException in method {method} from farm {fmid}.
```

Description: Internal error.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
6823
SM {spname} stat request error ({error}).
```

Description: An error occurred while processing the `sminfo` request.

Solution: Retry the `sminfo` command. Stop and restart the segment manager if necessary.

```
9300
FM Error: Internal error, cannot set state {state NO}({Exception Details}).
```

Description: Unable to update an entry in the `Farm` table.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

```
9302
FM Error: Farm id {farm ID} does not exist.
```

Description: The specified farm does not exist in the farm table.

Solution: Refer to Chapter 7 for details on how to remedy database problems.

```
9303
{classname} invalid classname ({error}).
```

Description: The specified class name is invalid.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9304
Unexpected exception: {error}.
```

Description: An unexpected error has occurred.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9305
FM Error: {method} for farm {fmid} encountered error ({error}).
```

Description: Internal error.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
9306
FMDispatch Error: Cannot clear config.
```

Description: An error occurred while removing the network interface card (NIC) interfaces on the CP for a farm.

Solution: Check the log for earlier messages that might help explain the root of the problem.

```
9307
FMDispatch Error: Cannot get DNS/DHCP handle for farm {fmid} due to error ({error}).
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9308
FMDispatch Warning: {DNS/DHCP} sync for farm {fmid} failed due to error ({error});
   continuing
```

Description: An error occurred while synchronizing the configuration for the DNS or DHCP daemons.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9309
FMDispatch Warning: Cannot sync customer monitoring for farm {fmid}
   due to error ({error}).
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9310
FMDispatch Error: Cannot delete node from monitoring for farm fmid device devname
   (id=devid, ip=ipaddr) due to error (error).
```

Description: The system cannot delete a device node from monitoring.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9311
FM Warning: Farm fmid had error state state in db; resetting to state.
```

Description: Informational warning.

Solution: Ignore this message.

```
9312
FM Error: state state (stateid) is not valid.
```

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

```
9313
FMDispatch Error: Cannot check monitoring for registration of the farm fmid, node nodeid
```

due to error (*error*).

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

9314
FM Error: Error occurred (*error*)

Description: An error occurred during a farm action.

Solution: Check the log for earlier messages that might help explain the root of the problem.

9315
Error: Database error occurred (*error*).

Description: An error occurred while accessing the database.

Solution: Check the log for earlier messages that might help explain the root of the problem.

9316
FM Error: Invalid Goal State *goal-state* from state *current-state*.

Description: Internal error.

Solution: Contact http://sun.com/service/contacting/index.html for assistance.

CHAPTER 7

Troubleshooting

As with any software, you might have occasional difficulties with N1 Provisioning Server software or an Infrastructure Fabric (I-Fabric). This chapter helps you to identify and resolve such problems. This chapter describes the potential problems, their likely causes, and practical solutions to these problems. It features useful troubleshooting tables and flowcharts. The following major topics are covered in this chapter:

Failover of resource pool servers within an I-Fabric is automated, and requires you to perform few or no manual tasks after a failover and startup of a new device. Any tasks to be performed after a failover are described within the sections dedicated to specific I-Fabric devices within this chapter.

Debug Levels

You can set the debug level in the `tspr.properties` configuration file according to your preference of debugging detail. The higher the debug levef, the more logging information you receive. A debug level setting of 9 is recommended. You can view

debug information in the `tspr.debug.log`file. The following is an example `tspr.properties`file the debug level set to 9. All entries in the `tspr.properties`file must have the following format:

package name . class name . attribute=value

```
com.terraspring.core.sys.GridOS.debuglevel=9
```

Handling Failed Farm Devices

N1 Provisioning Server software actively and automatically monitors the availability of all devices within a farm to enable automated failover of logical server farm devices. No special configuration is necessary to enable this monitoring. When the software detects that a device is no longer available, the software can take the necessary automated steps to replace the physical device with an identical device from the pool of unused devices. In replacing the failed physical device, the system duplicates configuration and logically reattaches storage to the new device so that the new device can take over the role of the failed device.

Whenever a failure is detected, the system starts the failover process by placing a failover "job" request into the N1 Provisioning Server queuing mechanism. This same queuing mechanism is used to process all farm activation and farm update requests. You can view the current set of tasks and their status by running the `request -l` command.

Two options for specifying failover behavior when the monitoring system detects a device failure:

- Automatic failover consists of the N1 Provisioning Server placing and automatically processing a `replacefaileddevice` request in the queue. You are not required to intervene.
- Manual failover consists of the N1 Provisioning Server placing a `replacefaileddevice` request as a blocked request in the request queue for the replacement of the failed device. You must unblock the request before it can be processed and the device replaced. Run the command `request -u` *request-ID* to do this.

Note – A blocked failover request blocks all subsequent requests for that farm, including farm update requests made through the Control Center. The system does not process any other changes to the farm until the failover request is either unblocked or deleted.

The automatic option enables immediate processing of the failover when the request enters the system's request queue. The manual option causes the failover request to be blocked in the system's request queue. As a result, the blocked failover request and any subsequent request for the failed device's farm are not processed until you either unblock the request to allow the failover or you delete the request to abort the failover.

Automatic failover can be costly if the server has not really failed, but was unavailable. In this case, the device might get replaced unnecessarily.

To configure the behavior of the failover mechanism, alter the property in the `/etc/opt/terraspring/tspr.properties` file.

If the `com.terraspring.cs.services.DeviceStatus.blockReqFailedDevice` property is set to `false`, any detected device failures will not result in blocked failover requests. If this property is set to `true`, any detected device failures will result in blocked failover requests. The property is set to `true` by default, which is the recommended setting.

Troubleshooting Farm Device Failure

Upon the detection of a farm device failure, the monitoring software sends a device `DOWN` event to the segment manager of the N1 Provisioning Server that manages the farm.

Note – The monitoring software sends a device `DOWN` event only for active farms.

When the segment manager receives a `DOWN` event from the monitoring software, the segment manager performs the following procedures:

1. A blocked `replacePhysicalDevices` request is sent to the farm manager of the farm that owns that device.
2. A critical error message is logged into the log file to alert operations. The critical message that the segment manager generates contains the failed device's name and the corresponding farm ID.

▼ To Respond to Farm Device Failure

Perform the following procedure on an control plane server if you receive a critical farm device failure message.

Note – If the automatic failover property is set to `true`, no action needs to be taken.

Steps

1. **List the current requests for the farm:**

   `request -lf` *farm ID*

2. **Review the list and obtain the `requestID` of the blocked `replacePhysicalDevices` request generated by the segment manager for the farm.**

   You can identify the `requestID` by the `replacePhysicalDevice` request where state is listed as `QUEUED_BLOCKED`. The second argument of the `replacePhysicalDevices` request specifies the IDs of the devices that failed.

3. **Verify that the physical device has actually failed and that it is not a spurious error. See "Handling a Failed Control Plane Server" on page 185 for details. Temporary network failures can cause spurious errors.**

   - **If the device has not failed, or you do not want to replace the device, delete the request by typing `request -d` *request ID*.**
   - **If only one device failed, start the device replacement by unblocking the `replacePhysicalDevices` request typing `request -u` *request-ID*.**
   - **If multiple devices failed, you will see many `replacePhysicalDevices` requests.**

     a. **Identify the device ID of each of the failed devices.**

     b. **Run the `replacedevice` command with the device IDs of all the devices that you want to replace:**

        `replacedevice` *farm-ID failed-device-ID failed-device-ID failed-device-ID*

        For example:

        ```
        replacedevice 5 19001 37001 2003
        ```

4. **After replacing failed devices, delete the `replacePhysicalDevices` requests by typing:**

   `request -d` *request-ID*

Handling a Failed Control Plane Server

If the control plane server fails, see the server documentation for details on replacing the server. If the Oracle database on the control plane database (CPDB) has been damaged, contact http://sun.com/service/contacting/index.html for assistance.

See Chapter 5 for a description of the files that you need to restore. See the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide* for details on how to reinstall the software that runs on the control plane, such as the Control Center software and the N1 Provisioning Server software.

Handling Failed Software on the Control Plane

This section describes how to troubleshoot the failure of software processes on the control plane.

Monitoring Manager Failover

If the monitoring manager fails, attempt to restart it by running the `/opt/terraspring/sbin/mmd start` command. If the monitoring manager does not restart, see the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide* for details on how to reinstall the monitoring manager software and restore the most recent backup.

Control Center Failure

If the Control Center fails, attempt to restart it by running the `/opt/terraspring/sunone/bin/appserv start` command. If the Control Center does not restart, see the *N1 Provisioning Server 3.1, Blades Edition, Installation Guide* for details on how to reinstall the Control Center software and restore the most recent backup.

Diagnosing N1 Provisioning Server Problems

Almost all farm operations are carried out asynchronously, that is, a request (message) is queued for the N1 Provisioning Server, which processes them in the order received. Use the `request` command to view all pending and in-progress requests in the N1 Provisioning Server.

Any critical error occurring during the execution of a farm operation is reported by the Monitoring Manager through the monitoring system. Critical messages are also logged to the `/var/adm/messages` file on the N1 Provisioning Server. The critical error causes the current farm operation to exit and moves the farm into an error state. When the farm is in an error state, requests in the queue cannot be processed until the error is cleared manually. After the problem causing the error is resolved, you must reset the farm error state so that the stopped operation can be restarted.

In addition to the `/var/adm/messages` file, various debug messages are logged in the file `/var/adm/tspr.debug`. The amount of information in this file is determined by the debug level set in the `/etc/opt/terraspring/tspr.properties` file. The default debug level is 9. When a critical error is encountered, the additional information in the `/var/adm/tspr.debug` file is useful in determining the cause of the error.

To diagnose a problem with the N1 Provisioning Server, follow the steps illustrated in Figure 7–1.

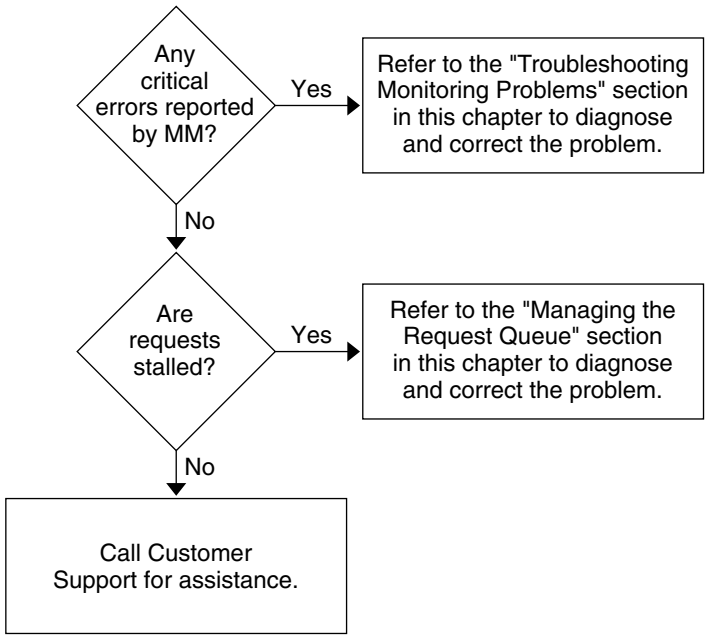

FIGURE 7–1 Diagnosing N1 Provisioning Server Problems

Message Log File Format

All messages logged into the `/var/adm/messages` and the `/var/adm/tspr.debug` files are in a standard format. The following example is a typical message:

```
Nov 12 17:22:57 sp3 java[23033]: [ID 398540 user.info] TSPR
[sev=crit] [fmid=1211] [MSG6718] FM Activate: Ready timeout
expired.
```

TABLE 7–1 Message Log File Format

Message Element	Description
`Nov 12 17:22:57`	The date and timestamp of the message.
`sp3`	The name of the control plane server.
`java[23033]`	The unique ID of the Solaris process.
`[ID 398540`	The unique ID of the system log.

TABLE 7–1 Message Log File Format *(Continued)*

Message Element	Description
`user.info]`	The type of system log entry.
`TSPR`	The type of message. In the example, the message is an N1 Provisioning Server message.
`[sev=crit]`	Indicates the severity of the message. In the example, `crit` means this error message is critical. Other severity types are `warn`, `okay`, and `dbug`.
`[fmid=1211]`	Indicates the application ID that generated the message. In the example, the application is the farm manager for farm `1211`.
	The segment manager application ID is in the format `[smid=27003]`.
	The command-line tools application ID is in the format `[apps=26765]`.
`[MSG6718]`	The message ID assigned by the N1 Provisioning Server.
`FM Activate: Ready timeout expired`	The message itself.

When an error or exception occurs during execution, the name of the exception is logged as part of the message along with a description of the exception.

Managing the Request Queue

The farm operations initiated using the Control Center or the `farm` command are queued for the segment manager. The segment manager processes the request, and dispatches the request to the farm manager, which actually does the work. When investigating problems, ensure that your request has not stalled along the way to the farm manager.

If there are no critical errors or messages relating to requests that are being processed in the `tspr.debug` file of the segment manager, check the request queue to verify that your request has been processed.

The `request` command is used to view the request queue:

- To list all current requests in the N1 Provisioning Server, run the command `request -l`.

- To list all current requests for a farm, run the command `request -lf` *farm-ID*.
- To list all completed requests for a farm, run the command `request -lcf` *farm-ID*.
- To list all requests for a farm, run the command `request -laf` *farm-ID*.

For more information on the `request` command and its usage, see the `request` man page.

If your request is still in the queue in the `QUEUED` state, the request has not been processed yet. If no other requests are ahead of this request for your farm, the request queue might have stalled. The next section describes how to resolve this problem.

Handling Stalled Requests

A request can sometimes fail to be processed by its intended server and might stay in the queue unattended. This section explains how to diagnose and solve this problem.

Check for Blocked Requests

If starting the segment manager does not process all requests on the queue, some requests might be blocked. Blocked requests require manual intervention prior to their processing. When you have reviewed the blocked requests, you can unblock them to process them or delete them. You can run the `request -u` *request-ID* command to unblock or the `request -d` *request-ID* command to delete blocked requests.

After the blocked requests are cleared, the requests in the queue are processed in the order in which they were received.

Verify That the Farm Queue Handler Is Active

Sometimes the request is not blocked but just queued, and the farm manager is still not processing the request. In this case, ping the farm. Issuing the `ping` command to the farm activates the farm manager queue handler to process requests. This command especially applies to requests for farms that are in an error state. Run the following command to ping the farm to activate the request queue handler:

`farm -p` *farm-ID*

If the farm was in an error state (that is, a previous operation ended in error), reset the error to move the queue:

`farm -pf` *farm-ID*

Managing Failed Power and Switch Operation Requests

You can look for failed power or switch operation requests by checking the `tspr.debug` log for `ExpectTimedOut` messages. Power operation requests include `powerUp`, `ispowerUp`, and `powerDown` operations. Switch operation requests include `addPort`, `removeVlan`, and `removeAll` operations.

A failure can occur if the N1 Provisioning Server is unable to access a device or the N1 Provisioning Server receives unexpected output from a device. If a failure occurs, check the following items:

1. Verify that the IP address for the device is set correctly and that communication from the N1 Provisioning Server to the device exists.
2. Run the `device -lv` *device-ID* command from the N1 Provisioning Server to verify that all login names and passwords for this device are correct.
3. Ensure that the firmware is the supported version and that any changes made to the default settings are acceptable.

If you need to debug further to resolve the issue, enable the `expect` log in the following properties:

- `com.terraspring.drivers.util.expect.Expect.print=true`
- `com.terraspring.drivers.util.expect.Expect.output=`*log-name-path*

Run the `tail -f` *log-name-path* command and reissue your request to view the operation.

Troubleshooting Farm Operations

The following table describes N1 Provisioning Server troubleshooting issues related to farm operations. This list is not inclusive. The message log file indicates whether your problem relates to farm operations.

TABLE 7–2 Troubleshooting Farm Operations

Problem	Possible Cause	Action
Configuration exception	Invalid Farm Markup Language (FML) for the farm.	Contact http://sun.com/service/contacting/index.html for assistance.
Dynamic Host Configuration Protocol (DHCP) exception	Cannot create DHCP configuration for the farm.	This problem can occur if the host name of the control plane server cannot be determined. Make sure the server has the proper network configuration.
Domain Name Service (DNS) exception	Cannot create the DNS configuration for the farm.	This problem can occur if the host name of the control plane server cannot be determined. Make sure the server has the proper network configuration. This problem also can occur if the database connectivity is lost. To verify that the control plane database (CPDB) is running and is accepting connections, run any command, such as `farm -l`. If no proper output displays, the CPDB is not running.
`NoMoreResources` exception	Not enough resources are available to allocate the farm.	Check the message to see which resource is exhausted. Use one of the following methods to provision more resources ■ Add resources and add their information to the database ■ Free up resources by deactivating existing farms ■ Flex down a farm to free up the required devices. Restart the farm operation.
I/O exception	An I/O error occurred while performing a task such as manipulating disks, configuring monitoring, or accessing files.	Depending on the specific place where the exception is thrown, take appropriate action that should be part of the message with the exception, such as checking network connections.
SQL exception	A database error occurred.	This problem can occur if the database connectivity is lost. To verify that the CPDB is running and is accepting connections, run any command, such as `farm -l`. If no proper output displays, the CPDB is not running. Report all other database errors to http://sun.com/service/contacting/index.html.

TABLE 7–2 Troubleshooting Farm Operations *(Continued)*

Problem	Possible Cause	Action
Exceptions other than SQL, such as`IllegalState`, `IllegalArgument`, `IllegalAccess`, and so on	Internal error.	Contact http://sun.com/service/contacting/index.html for assistance.
Farm activation fail.	Check the log file for any critical errors that might point to the cause.	Depending on what type of error message was received, take appropriate action to activate the farm.
Farm activation fails during allocation because not enough subnets are available	The subnet size might be too large.	You have the option to choose the size of the external subnet for the farm. If you select a size that is not currently in the database, you need to add a subnet by running the `subnet` command.

Farm deployment fails because `named` could not restart	The JVM™ software caches the configuration of the `nsswitch.conf` file, which describes which database to use for host lookups. If DNS was not part of the `nsswitch.conf` file's host entry at the time the segment manager was started, all host lookups that cannot be resolved using the `/etc/hosts` file will fail. See the `tspr.debug` log for a detailed message describing the error.	1. Ensure that the entry for hosts lookup in the `/etc/nsswitch.conf` file reads as follows:
		`hosts: files dns`
		2. Restart the segment manager by running the command
		`/opt/terraspring/sbin/sm -start`
		3. Reactivate the farm.

TABLE 7–2 Troubleshooting Farm Operations *(Continued)*

Problem	Possible Cause	Action
Server is booting, but is unable to get its IP address through DHCP.	The DHCP daemon might not be running or the media access control (MAC) address of the server is incorrect.	1. Ensure that the details listed in the database are correct for that server. Run the `device -lv` *device-ID* command and check the MAC addresses and switch ports. 2. Verify that the DHCP daemon is running and is answering requests by first running `ps -ef |grep dhcp`, then look in the `tspr.debug` file to see if there are DHCP messages logged. 3. Connect to the switch and ensure that those ports are connected. Depending on the switch, run either the `sh cam dyn` *port* or `sh.mac dyn` command to ensure that the correct MAC address of the server appears. 4. Check that the Ethernet interface on the control plane server appears as connected on the switch and that it is running as an 802.1q trunk port, with a native virtual local area network (VLAN) of `1`.
The control plane server cannot create the DHCP configuration for the farm and receives a message indicating an unknown host.	The network configuration of the control plane server might be incorrect.	1. Check the network configuration of the control plane server. 2. Check database connectivity and the file system that contains the DNS configuration.
The control plane server cannot create the DNS configuration for the farm and receives a message indicating an unknown host.	The network configuration of the control plane server might be incorrect.	1. Check the network configuration of the control plane server. 2. Check database connectivity and the file system that contains the DNS configuration.

TABLE 7–2 Troubleshooting Farm Operations *(Continued)*

Problem	Possible Cause	Action
The Control Center does not display a farm correctly after it has been updated.	Two possible causes: 1. The update request has not yet completed in the CPDB. This may take a few minutes. 2. The update request has failed.	1. Reconfigure the farm in the Editor dialog and resubmit the update request. 2. Reissue the the `farm -af` *farm ID* command from the command line. Then submit another farm update request from the Control Center Editor dialog.
Replace device requests are generated intermittently even though the devices are running and able to `ping` successfully.	The Ethernet port speeds might not be set to the correct value.	Ensure that the Ethernet port speeds and duplex setting are set to the same values on all sides (control plane server, switch, and device).

Farm Error Status Codes

Every farm has an error status code that is associated with the farm to indicate whether the farm is currently in an abnormal state.

- The error status code of `0` represents a health state.
- The error status code of `1000` means that the farm manager is processing a request.
- An error status code other than `0` or `1000` means that the farm has an error.

During the request process, the farm's internal state changes whenever the transition process is completed with a success. If the farm fails to transition from one internal state to another internal state, the farm's internal state is not changed and the farm is set with an error status. The value of the error status code is the failed internal state value.

For example, if the farm failed the transition from state `ALLOCATED (20)` to state `WIRED (30)`, the farm is still left with the internal state `ALLOCATED (20)` and the farm error status code is set to `30` to represent the failed state `WIRED`. Just before the code is set to `30`, it is `1000` to indicate that the request is in progress.

Whenever a farm error occurs, a critical error message is generated in the system log file `/var/adm/messages`.

Caution – The farm manager will not process any further farm requests until the error condition is changed and the error status code is cleared to `0`.

Troubleshooting Image Problems

The following table contains descriptions of troubleshooting scenarios involving images. This list is not inclusive. The message log file indicates whether your problem relates to images.

TABLE 7–3 Troubleshooting Image Problems

Problem	Possible Cause	Solution
While allocating the resources for a farm, the N1 Provisioning Server generates a message indicating that it cannot find the named image.	The image ID in the FML file is incorrect.	Contact http://sun.com/service/contacting/index.html for assistance. Synchronize the Control Center. Also check the images list on the CPDB by running the command `image -ls`, and ensure that this list matches what displays in the Control Center.
A farm update request failed after making a snapshot.	The server was not completely backed up and running yet when the update request was made.	When you issue a farm update request after taking a snapshot, make sure that the server is completely up and running before issuing a farm update request. Otherwise, the farm update will fail. To ensure that the server is up and running, run the command `ping` *server IP-address*.
Farm activation fails after image creation.	Not enough external subnets.	The farm that is automatically created by the image creation process includes an external subnet. Define a new external subnet using the command `subnet -cx -m` *netmask-size network-IP*.
Farm standby request failed.	Not enough space on the image server for storing the image.	Go to the directory where all images are stored and issue the command `df -k` to determine how much space is available for storing image. If the capacity is above 85 percent, add more storage space for images.

Troubleshooting the Image Server

The image server manages images. The image server can be either a stand-alone network file server (NFS) or it can run on a control plane server. See Chapter 3 for details.

Gigabit Ethernet Support

If your image requires Gigabit Ethernet support for the Solaris operating environment, ensure that the appropriate drivers are loaded each time the system boots. You can initialize the list of managed interfaces by running the following commands:

```
devfsadm
ifconfig -a plumb
ifconfig -a | grep flags | cut -d: -f1 | grep -v lo0> /etc/opt/terraspring/managed_interfaces
```

Now edit the contents of `/etc/opt/terraspring/managed_interfaces` by commenting out any interfaces that should not be managed by N1 Provisioning Server software.

Note – The instance assigned to any gigabit Ethernet card must always be `0`. For example, `ce0, skge 0, alt0`.

Troubleshooting Monitoring Problems

This section describes some common monitoring problems. It shows how to diagnose these problems and suggests possible corrective actions. See Chapter 4 for an introduction to monitoring concepts.

The most common problems are as follows:

- An `UP` message is not received from one or more resource pool servers.
- Monitors do not show colors on the Element Monitor window of the Control Center.
- A farm does not activate even though an `UP` message was sent.
- Frequent `UP` and `DOWN` messages are received.

Many of these problems have interconnected root causes. The most common causes are:

- Network, DNS, or DHCP issues.
- Monitoring processes are not running on the control plane server.
- Agent processes are not running on the resource pool servers.

This section describes how to diagnose these symptoms. “Corrective Actions for Monitoring” on page 202 describes the corrective action to take to resolve these problems.

An UP Message Not Received From One or More Servers

You can confirm an UP message problem on the control plane server by checking whether the following conditions exist:

- In the `/var/adm/tspr.debug` file, a message is listed similar to the following:

```
"Still waiting for 1 device(s) in 2879974 ms"
```

- The farm activation shows `ERROR 50`, as shown in the following example:

```
FARM_ID    FARM_NAME   CUSTOMER   STATE   ISTATE       ERROR
123        Farm_Name   Customer   NEW      DISPATCHED   50
```

The following figure shows the steps needed to diagnose and resolve this problem.

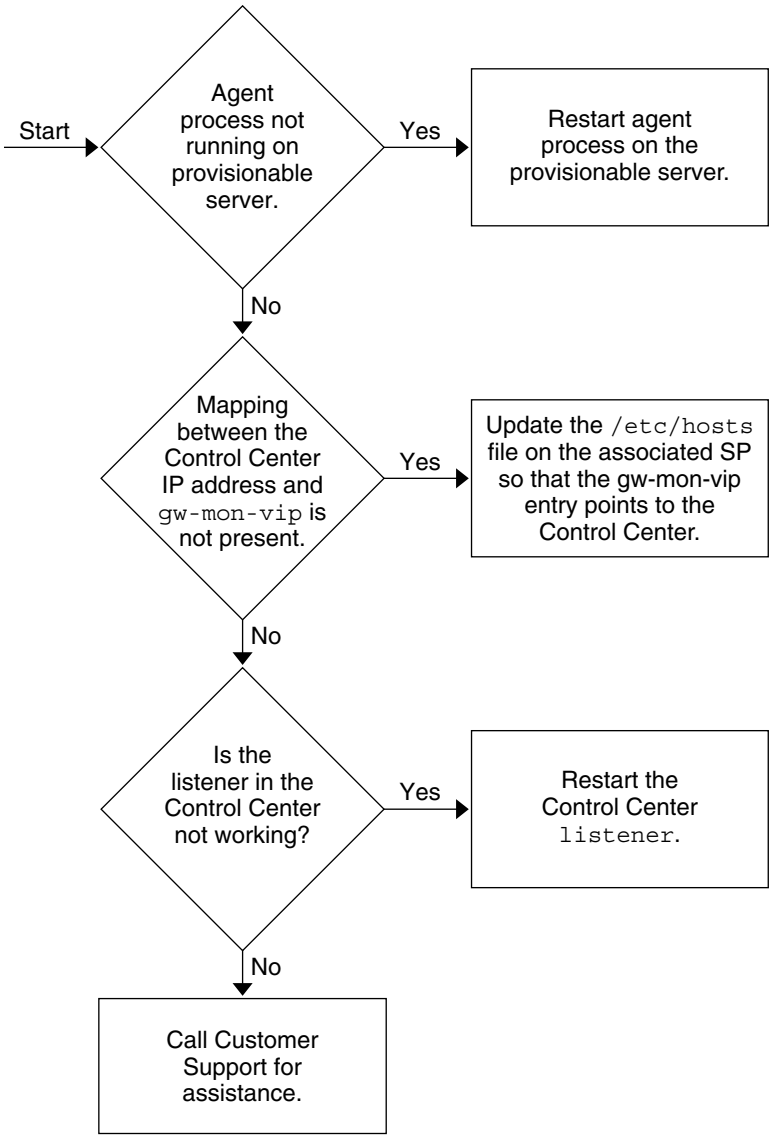

FIGURE 7–2 Resolving Monitoring Problems

The preceding illustration shows the following troubleshooting sequence:

1. Check for a network, DNS, or DHCP problem. See "Network, DNS, or DHCP Problems" on page 200 for details on how to do this.
2. Check that the monitoring processes are running on the control plane server. See "Monitoring Processes Are Not Running on the Control Plane" on page 200 for details. Follow the instructions in this section to restart the processes.

3. Check that the agent processes are running on the resource pool server. See "Agent Processes Are Not Running on a Resource Pool Server" on page 201 for details. If the agent processes are not running, follow the instructions in this section to restart them.

Monitors Do Not Show Colors in the Element Monitor Window of the Control Center

Farm-specific monitors might not appear on the Control Center. This condition could be caused by one of the following problems:

- Agent processes are not running on the servers.
- The mapping between the `gw-mon-vip` and the IP address of the Control Center server software is not set in the `/etc/hosts` file on the control plane server.
- The `listener` on the Control Center is not running. See "Control Plane Server-to-Control Center Messages Not Working" on page 201 for information on how to verify this condition.

Figure 7–2 shows the sequence of steps for you to follow to diagnose and resolve the above error condition. See "Control Plane Server-to-Control Center Messages Not Working" on page 201 for details on how to resolve these problems.

Farm Does Not Activate

Even though the `UP` message was sent by the monitoring system, the segment manager might not be running. In this case, restart the segment manager. See "Check for Blocked Requests" on page 189 for details on this procedure.

Frequent UP and Down Messages Received

A number of `UP` and `DOWN` messages received for a server might be received as a result of incorrect configuration of the interfaces on the N1 Provisioning Server.

Clear the duplicate Ethernet interfaces on the control plane server by running the `clearNicInterface` command. See the man pages for details on using this command.

Diagnosing Common Monitoring Symptoms

A number of symptoms are common to a number of problems. This section describes how to diagnose the following symptoms.

- "Network, DNS, or DHCP Problems" on page 200

Network, DNS, or DHCP Problems

Perform the checks in the following table for network, DNS, or DHCP problems:

TABLE 7–4 Checking for Errors

Error Check	Error Confirmation
Verify that all the resource pool servers can receive `ping`signals by running the following command on the control plane server: `/opt/terraspring/sbin/mls -lf` *farm-ID*. **Note –** This command lists all the servers in the farm that can receive `ping` signals.	Any of the servers are listed as `ADDED`
Verify that all the resource pool servers are reachable by performing a `telnet` to each of the servers.	Any of the servers are not reachable with `telnet`

Note – Sometimes a server can receive `ping` signals but is not reachable with `telnet` when in a single-user mode. To resolve this problem, connect to the console port and boot into multiuser mode.

Monitoring Processes Are Not Running on the Control Plane

After you determine a diagnosis for a monitoring process run the command:

```
/usr/ucb/ps -auxww | grep MM
```

If the monitoring process is running, you will see an output similar to this example:

```
USER     PID %CPU %MEM SZ  RSS TT   S START  TIME  COMMAND
root 14540 0.2    1.14 485 620 608? S Mar 05     18:32 /bin/../java/bin/..
/bin/sparc/native_threads/java -Dsun.net.inetaddr.ttl=0 com.
terraspring.mon.MM
root 9529  0.1    0.1      976 672 pts/2 S 11:49:40 0:00 grep MM
```

If the monitoring process is not running, you will see an output similar to this example:

```
USER PID %CPU %MEM  SZ  RSS TT     S  START TIME     COMMAND
root 9565 0.1  0.1  976 672 pts/2  S  11:50:28  0:00 grep MM
```

See "Restart the Monitoring Processes on the Control Plane Server" on page 202 for details on how to restart the process.

Agent Processes Are Not Running on a Resource Pool Server

Agent processes might not be running on a resource pool server. You can verify this condition by one of two methods:

- On the control plane server run the following command:

  /opt/terraspring/sbin/mls -a *IP address of host*

  To be able to use this command, you must know the IP address of the server.

- On the server on which the agent you want to verify is running, run the following command:

  ```
  /usr/ucb/ps -auxww | grep tspragt
  ```

  If the agent processes are running, you will see output similar to the following example:

```
root 7652  0.1  0.1  976  656 pts/1 S 11:37:30  0:00 grep tspragt

root 321  0.1  0.73167213816 ? S 16:26:37  0:10 /usr/bin/../java/bin/..
/bin/sparc/native_threads/java -Dsun.net.inetaddr.ttl=0
com.terraspring.mon.client.tspragt start 10.42.14.2
```

  If the agent processes are not running, you will see output similar to the following example:

  ```
  root      7709  0.1  0.1  976  656 pts/1    S 11:39:54  0:00 grep tspragt
  ```

See "Restart the Agent Processes on a Resource Pool Server" on page 202 for details on how to restart the process.

Control Plane Server-to-Control Center Messages Not Working

For a number of reasons messages between the control plane server and Control Center might not work. The most common reasons include:

- The mapping between the `gw-mon-vip` to the IP address of the Control Center server software is not set in the `/etc/hosts` file on the control plane server. Verify that a suitable entry is present to check this condition.

  For example:

  ```
  10.5.131.19 gw-mon-vip
  ```

- The `listener` on the Control Center server software is not running. You can verify this condition by running `finger test@gw-mon-vip` on the control plane server. The expected sample output is similar to the following examples:

  ```
  [gw-mon-vip]
  ```

  or

  ```
  [hostname]
  ```

Corrective Actions for Monitoring

This section describes a number of corrective actions that you can take to resolve a monitoring problem.

Restart the Monitoring Processes on the Control Plane Server

To restart the monitoring process run the following commands on the control plane server:

`/opt/terraspring/sbin/mmd stop`

This command ensures that all relevant processes are stopped. Restart the monitoring processes with the following command:

`/opt/terraspring/sbin/mmd start`

Restart the Agent Processes on a Resource Pool Server

If the control agent process terminated on the server, start the process on the server with the following command:

`/etc/init.d/N1PSagt start`

To verify that the processes have restarted, run the following command from the control plane server:

`/opt/terraspring/sbin/mls -a` *server IP address*

If the agent is running you will see output similar to the following:

```
FARM_ID IP_ADDRESS TYPE STATE DB_STATE SINCE 134 10.9.0.35 Server UP UP Feb 05 14:15:32
```

If the agent is down in real-time (STATE) it might still be marked as being UP in the database (DB_STATE because the database state is updated every five minutes. Therefore it will be up in real time, but still down in the database state. You will see output similar to the following:

```
FARM_ID IP_ADDRESS TYPE STATE DB_STATE SINCE 134 10.9.0.35 Server DOWN UP Feb 10 14:20:33
```

Troubleshooting Control Center Problems

The main problem you might encounter when working with the Control Center relates to the CPDB connection: either the connection between the Control Center and the CPDB is down, or parameters are incorrectly configured in the Control Center database.

To determine which of these two problems you might be encountering:

1. Log in as an administrator in the Control Center using the browser.
2. In the configuration tools section, select I-Fabrics to bring up a list of I-Fabrics.
3. On the I-Fabrics List, select the I-Fabric whose connection needs to be checked.
4. Click OK when the following dialog displays: Improper Configuration of these Properties may disrupt System Operation. Proceed with Caution.
5. In the Property Name Property Value dialog, take note of the IP address of the database server (defined at the field primary) and its port.
6. If no IP address is listed, telnet to the hostname of the URL on which the Control Center runs.
7. Try the telnet command (assuming 10.0.0.18 and 1521 are the IP address and port obtained at step 4):

   telnet 10.0.0.18 1521

   ```
   Trying 10.0.0.18...
   ```

   If you see the following response:

   ```
   Connected to cpdb
   Escape character is "^]"
   ```

the connection is okay. Type Ctrl-C to cancel the telnet session.

If you see the following response:

```
telnet: Unable to connect to remote host: Connection refused
```

the IP is okay, but the CPDB database server is not running (or it is running on a different port). Contact your DBA or consult the database manufacturer's documentation for information on how to verify the listener port.

```
ping 10.0.0.18
```

If you see the following result:

```
no answer from 10.0.0.18
```

no communication exists between the Control Center and the database server hosts exists. This error might be caused by a routing problem or by the database server being down.

Connecting to the CPDB

The connection from the Control Center to the CPDB might be too slow at times causing a timeout. This condition might be caused by either the database or the network being slow due to overwhelming workload. To prevent a timeout when the connection is slow, increase the value in the timeout field in the I-Fabric Properties screen in the Control Center. The default is 30 seconds.

▼ To Increase the Timeout Interval

Steps

1. **Login to the Administration screen of the Control Center.**
2. **Select I-Fabrics under Configuration Tools.**
   An I-Fabrics List and I-Fabrics Properties screen appear.
3. **Increase the value in the Timeout field.**
   If this property does not yet exist you can create it by clicking the Add Property button.

APPENDIX **A**

Chargeable Events

This appendix describes chargeable events within an Infrastructure Fabric (I-Fabric) and provides definitions for these events..

Chargeable Event Overview

The farm manager generates chargeable events that you can use to bill customers for usage of certain devices within their farm (for example, resource pool servers).

- You can use commercially available billing systems to correlate events into usage records and charge customers based on these usage records.
- You can compile a usage record using a commercially available reporting system and generate reports.
- You can create your own custom billing tools. If you create your own billing tool, the definitions provided in this appendix will be useful.

Definitions of Events

An N1 Provisioning Server generates chargeable events at the farm level and at the resource pool level. All chargeable event messages include the time expressed in UTC, a sequence number and details of the event. All chargeable event messages have the following common format:

```
event-message  ::= utc-date-time ”,” sequence-info ”,” event-info ”,” event-details

event-info     ::= ”event” ”=” event
```

```
event          ::= “resource” | “control” | “farm”

event-details  ::= rl-event-desc
                   | cl-event-desc
                   | farm-event-desc
```

The sequence number uniquely identifies the event within an N1 Provisioning Server. A sequence has the following definition:

```
sequence-info    ::=    “seq”    “=”    fabric-name    “:”    sequence-id
```

The chargeable event records contain the event timestamp in the time zone of the control plane server that generated the data.

Typically, events are described as a sequence of name/value pairs. Unless explicitly indicated, all values are enclosed in double quotation marks to allow for easy transition to XML formats.

General Definitions

The following definitions apply to the values of all farm-level and resource pool definitions described in the preceding sections:

```
number ::= digit { digit }

digit  ::= "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"
```

where *number* represents a decimal number.

```
name    ::= letter { letter | digit | special }    | quoted-string

special ::= "." | "-" | "_"
```

where *value* `::=` represents a contiguous sequence of nonwhite space characters.

You must enclose a name or value that includes a comma (`","`) in double quotation marks. In such cases, a double quotation mark is escaped as two contiguous double quotations (that is, `""` represents a quoted double quote).

fabric-name, *account-name*, *image-name,* and *dns-name* are defined to be a *name*.

Thus *sequence-id, account-id, device-id, farm-id, device-id, disk-id, vlan-id*, and *disk-size* are all defined to be a *number*. Both *disk-size* and *mask-len* are also defined to be a *number*.

IDs are set and generated by the N1 Provisioning Server. Names are generated by users of the N1 Provisioning Server. For example, *image-name* is specified by a user operating the Control Center.

subnet-mask is a quoted string representing a TCP/IP subnet mask. Typically, the string has the format of an IP address.

```
ipaddr    ::=    number    "."    number    "."    number    "."    number
```

Where *number* is in the range *0 ... 255*. An IP address is defined to have the same format as a subnet mask.

The definition of *utc-date-time* is as follows:

```
utc-date-time ::= utc-date utc-time

utc-date      ::= year    "-"    month    "-"    day

utc-time      ::= hours    ":"    minutes    ":"    seconds    [ "." subseconds ]
```

where *year* is a *number*, *month* is a *number* in the range `1...12`, and *day* is a number in the range `1..31`.

hours is a *number* in the range `0...23`, *minutes* and *seconds* are *numbers* in the range `0...59`. *subseconds* is a *number*.

A *wml-description* is the WML architecture for the specified device. The architecture is defined to be a quoted string.

Farm-Level Events

Farm-level events indicate when a farm is created or deactivated. The event messages indicate the customer with which the farm ID is associated so that you can subsequently associate resource pool events with a farm-level event. Farm-level events are structured like the following example:

```
farm-event-desc ::= farm-operation    ","    farm-info    ","    farm-state-info
","    account-info

farm-operation  ::= "op" "=" farm-op

farm-op         ::=     "add" | "del" | "update"

farm-info       ::= "farm-id" "="     farm-id

farm-state-info ::= "state" "=" farm-state

farm-state      ::= "new" | "active" | "standby" | "deactivated"

account-info    ::= "account-id" "=" account-id [ "," "name" "=" account-name ]
```

```
2003-01-20
14:00:01.0,newyork:1001,event="farm",op="add",farm-id="99",state="ACTIVE",
account-id="jdoe"
```

`1001` is the sequence ID of this event. Farm `99` was added and it is in state `ACTIVE`. This event occurred at 14:00:01 on Jan. 20, 2003.

```
2003-01-20
19:10:01.0,newyork:1022,event="farm",op="update",farm-id="99", state="UPDATE",account-id="jo
```

1022 is the sequence ID of this event. Farm 99 was updated, it is in state UPDATE. This event occurred at 19:10:01 on Jan, 20, 2003.

Resource Pool Events

Resource pool events indicate when a resource is allocated, made available, fails, and is deallocated. The definition of a resource pool event is as follows:

```
rl-event-desc ::= rl-operation "," farm-info "," category-info "," resource-desc [ "," account-info ]

rl-operation  ::= "op" "=" rl-op

rl-op         ::= "add" | "del" | "update" | "avail" | "fail" | "reboot"

farm-info     ::= "farm-id" "=" farm-id

category-info ::= "category" "=" category

category      ::= "device" | "disk" | "vlan" | "subnet" | "ipaddress"

resource-desc ::= device-desc | disk-disc | vlan-desc | subnet-desc | ipaddr-desc
```

The following table describes the types of event and when they occur:

Event	Occurs
"add"	When a device, such as a server, is allocated
"avail"	When a device is made available
"fail"	When a device fails
"del"	When a device is deallocated
"add" and "avail"	When a device is restored

All events referring to devices include the device ID. If the events refer to disks, the events include the disk ID. The device IDs and disk IDs enable you to accurately track all devices and disks.

Devices

The definition of a device event is as follows:

```
device-desc      ::= deviceclass-info "," devicetype-info "," deviceid-info

deviceclass-info ::= "class" "=" device-class

device-class     ::= "lb" | "fw" | "server" | "nas" | "ethport"

devicetype-info  ::= "type" "=" wml-description
```

```
deviceid-info    ::= "device-id" "=" device-id
```

Note – `"ethport"` is the name given to unmanaged Ethernet devices or other such devices with an IP address.

```
2003-02-01

10:00:00.0,newyork:2998,event="resource",op="add",farm-id="99",
category="device",class="server",type="sun-svr-blade",device-id="50101",
service-units="0"
```

On Feb 1, 2003 at 10:00:00, in the `newyork` I-Fabric, sequence id `2998`, a resource was added to farm `99`. This resource was a `device`. The class of `device` was a server. The type of the server was a `sun-svr-blade` server. The ID of the actual device allocated was `50101`.

Disks

The definition of a disk event is as follows:

```
disk-desc      ::= diskclass-info "," disktype-info "," diskid-info size-info "," [ "," image-info]

diskclass-info ::= "location" "=" disk-location

disk-location  ::= "internal" | "external"

disktype-info  ::= "type" "=" wml-description

diskid-info        ::= "disk-id" "=" disk-id

size-info      ::= "size" "=" disk-size

image-info     ::= "image-name     "=" image-name

               2003-02-01

               10:00:10.0,newyork:2999,event="resource",op="add",farm-id="99",
               category="disk",location="internal",type="local",disk-id="62",
               size="1000000000",image-name="small_solaris_blade"
```

On February 1, 2003 at 10:00:10 in the `newyork` I-Fabric, `sequence id 2999`, a resource was added to farm `99`. This resource was a `disk`. The disk's location is `internal` (that is, physically located inside of a server). The disk type is `local`. The physical disk ID is `62`. The size in bytes is `1000000000`. The image placed on the disk is `small_solaris_blade`.

Virtual Local Area Networks (VLANs)

The definition of VLAN events is as follows:

```
vlan-desc ::= "vlan" "=" vlan-id
```

```
2003-02-01

10:00:25.0,newyork:3002,event="resource",op="add",farm-id="99",
category="vlan",vlan="22"
```

On Feb 1, 2003 at 10:00:25, in the `newyork` I-Fabric, a VLAN was added to farm 99. The VLAN ID is 22.

Subnets

The definition of subnet events is as follows:

```
subnet-desc      ::=  subnet-mask-info "," masklen-info "," addr-type-info

subnet-mask-info ::= "subnet-mask" "," subnet-mask

masklen-info     ::= "mask-len" "," mask-len

addr-type-info   ::= "type" "," addr-type

addr-type        ::= "internal" | "external"

2003-02-01

10:00:15.0,newyork:3000,event="resource",op="add",farm-id="99",
category="subnet",subnet="10.10.0.81",mask-len="28",type="external"
```

On Feb 1, 2003 at 10:00:15 in the `newyork` I-Fabric, event `3000` was generated. A resource was added to farm `99`. The resource was a `subnet`. The subnet is located at `10.10.0.81`, the mask length of the subnet is `28`, and the type of the subnet is `external`.

IP Addresses

The definition of IP address events is as follows:

```
ipaddr-desc ::= ipaddr-info "," addr-type-info [ "," dns-name]

ipaddr-info ::= "ipaddress" "="     ipaddress

2003-02-01

10:00:20.0,newyork:3001,event="resource",op="add",farm-id="99",
category="ipaddress",ipaddress="10.10.0.83",type="external",server1
```

On Feb 2. 2003 at 10:00:20 in the `newyork` I-Fabric event 3001 was generated. The event was a `resource` event. A resource was added to farm `99`. The resource was an IP address. The IP address was `10.10.0.83`, it is an external IP address, and the DNS name that maps to that IP address is `server1`.

APPENDIX **B**

Command-Line Tools

N1 Provisioning Server command-line tools comprise two categories:

- Maintenance commands that you use to perform daily maintenance of an N1 Provisioning Server
- Administrator tools that you use during set up and initialization of N1 Provisioning Server components

The following table lists all maintenance command-line tools with a brief description. For details regarding the options and parameters for each command, and what type of value each command takes, see the man page for that command.

TABLE B–1 Daily Maintenance Command-Line Tools

Command	Description
`aps`	Lists details of N1 Provisioning Server applications
`bill`	Manages billing events for a specified account
`console`	Gets console access to a device
`device`	Lists and manages devices in the control plane database (CPDB)
`devtype`	Lists and manages device types
`disk`	Lists and manages disks in the CPDB
`diskscrub`	Scrubs disk to all zeroes
`encrypter`	Generates an encrypted string from a clear text string
`farm`	Farm management tool

TABLE B–1 Daily Maintenance Command-Line Tools *(Continued)*

Command	Description
`image`	Manages images within the N1 Provisioning Server
`imagecopy`	Copies images to disks and vice versa
`lockfarm`	Locks and unlocks a farm to prevent modifications to it
`lr`	Lists the resources used by a farm
`mls`	Lists monitoring information
`power`	Manages power for devices
`replacedevice`	Replaces a device in a farm
`request`	Lists and manages requests in the N1 Provisioning Server
`resetpasswd`	Restores the password on a farm, device, or device type
`rsck`	Checks the availability of resources required to allocate the farm
`snapshot`	Takes image snapshots
`subnet`	Lists and manages subnets and IP addresses
`vlan`	Manages VLANs in the CPDB
`vlanconfig`	Gets VLAN for a device, lists devices in a VLAN, and sets a VLAN for a device

The following table lists all administrator command-line tools with a brief description. For details regarding the options and parameters for each command, and what type of value each command takes, see the man page for that command.

TABLE B–2 Administrator Command-Line Tools

Command	Description
`backupdb`	Backs up the CPDB.
`cecmd`	Executes a remote command using the N1 Provisioning Server agent.
`cereg`	Registers a device for monitoring.

TABLE B–2 Administrator Command-Line Tools *(Continued)*

`check_rel`	Checks the version of the installed N1 Provisioning Server software and verify that the correct packages and versions are installed for the server's assigned role. The following roles are available, and they are listed in the `/opt/terraspring/etc/Roles` file: ■ `CC` – Control Center ■ `CPDB` – control plane database ■ CP – control plane server ■ STMC – Storage manager client (displays global images onto server blades and administers snapshots)
`deletesfb1600`	Deletes server blade configurations from the CPDB.
`dhcpmod`	Modifies the DHCP configuration.
`discoversfb1600`	Discovers the server blade configuration and adds it to the CPDB.
`dnsmod`	Modifies the DNS configuration.
`fmrouter`	Provides information to set up external routers to farms.
`imageserversetup`	Sets up the image server.
`imagesync`	Marks an image as done in the database.
`imagewizard`	Launches the image management tool.
`mmd`	Starts or stops the monitoring manager.
`newtsprdebug`	Rotates the log files specified in the configuration file `/etc/opt/terraspring/logfile_rotation`.
`pestest`	Tests free resource pool servers to see whether they can be used in a farm.
`pesupdate`	Add, updates, or removes agent and application packages on resource pool servers in an active farm.
`restoredb`	Restores the CPDB from a previous backup.
`segment`	Lists and sets a data center name.
`setupimgsubnet`	Sets up or removes the image server subnet.
`shelfsync`	Synchronizes the physical state of resources with the database.

TABLE B–2 Administrator Command-Line Tools *(Continued)*

`sm`	Starts and stops a segment manager.
`snmpd`	Starts or stops the SNMP daemon.
`switchsync`	After a failed switch has been replaced, this command synchronizes the new switch with the configuration information stored in the database for the failed switch.
`uninstall_PS`	Uninstalls the N1 Provisioning Server software.
`wmld`	Prints CPDB information in Wiring Markup Language (WML) format.
`wmlp`	Processes a WML file to create CPDB data.

Index

L

M

N

P

Q

R

S

T

U

V

W