Sun Fire™ B1600 Chassis and B100s, B100x, and B200x Blade Product Notes

Adobe PostScript

Contents

Documentation and Miscellaneous Information

These Product Notes contain important information about this release of the Sun Fire B1600 blade system chassis. This is the first release of the product to support Enterprise Linux version 3.0 and SuSE Linux Enterprise Server 8, running on B100x and B200x server blades.

This chapter contains the following sections:

1.1 Upgrading the BIOS on a B200x Server Blade

To run Red Hat Enterprise Linux version 3.0 or SuSE Linux Enterprise Server 8 on a B200x server blade, you must first upgrade the BIOS to version `1.1.32`. This version of the BIOS is available from the following website:

`http://wwws.sun.com/software/download/network.html`

For information on how to upgrade the BIOS on B200x blades with Linux installed, see Section 2.2, "Upgrading the BIOS on B200x Server Blades" on page 2-3 in these Product Notes.

1.2 Before Installing Solaris x86 Onto a Blade

Before you start to install Solaris x86 by following the instructions in the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*, please follow the steps in "Preparing to Install Solaris x86 Onto a Blade" on page 3-1 in these Product Notes.

1.3 Miscellaneous Information

1.3.1 Using N1 Provisioning Software

If you are installing N1 Provisioning software, you do not need to set up a Network Install Server. Before you do the System Chassis software setup, read the *N1 Provisioning Server 3.0 Blades Edition Implementation Guide*. The Implementation Guide explains what you need to do to accommodate the N1 Provisioning software installation.

1.3.2 Using the Sun Fire B10n Content Load Balancing Blade

The Sun Fire B10n Content Load Balancing Blade is now available to provide load balancing across server blades in the Sun Fire B1600 Blade System Chassis and other horizontally scaled Sun platforms.

To use the B10n Content Load Balancing Blade, you need to upgrade the firmware on the System Controller to version 1.1 or later. To perform the upgrade of the System Controller firmware, refer to the *Sun Fire B1600 Blade System Chassis Administration Guide* (Chapter 10).

To configure and use the B10n Content Load Balancing Blade, refer to the *Sun Fire B10n Content Load Balancing Blade Administration Guide*. For the latest information on version 2.1 of the Sun Fire B10n Content Load Balancing Blade, see *Sun Fire B10n Content Load Balancing Blade Version 1.2 Update Product Notes* (817-6211-10)

1.3.3 Downloading New Firmware for Chassis Components

For the latest publicly available firmware, check the following websites:

```
http://wwws.sun.com/software/download/network.html
http://sunsolve.sun.com
http://www.sun.com/servers/entry/b100x/
http://www.sun.com/servers/entry/b200x/
```

1.4 Viewing the Latest Documentation for the Chassis and Its Components

For the most up-to-date documentation, including the most up-to-date Product Notes, visit the following Sun documentation website:

`http://www.sun.com/products-n-solutions/hardware/docs/Servers/`

1.4.1 Documentation in this Release

Documentation for the Sun Fire B1600 blade system chassis and its components is provided on the CD supplied with a chassis or blade. The documentation is in Adobe Acrobat PDF format, therefore you need to use Acrobat Reader to view the files. To download Acrobat Reader (at no cost), go to the following website:

`http://www.adobe.com/products/acrobat`

This release of the Sun Fire B1600 blade system chassis includes the following documents:

- *Sun Fire B1600 Blade System Chassis Software Quick Start Poster*
- *Sun Fire B1600 Blade System Chassis Hardware Quick Start Poster*
- *Sun Fire B1600 Blade System Chassis Hardware Installation Guide*
- *Sun Fire B1600 Blade System Chassis Software Setup Guide*
- *Sun Fire B1600 Blade System Chassis Administration Guide*
- *Sun Fire B1600 Blade System Chassis Switch Administration Guide*
- *Sun Fire B1600 Blade System Chassis Compliance and Safety Manual*
- *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*
- *Sun Fire B10n Content Load Balancing Blade Administration Guide*

To access this PDF documentation, launch Adobe Acrobat Reader and open the file called `HOME.PDF` located in the `DOCS` directory.

1.4.2   Documentation Errata

- In the printed version of the *Sun Fire B1600 Blade System Chassis Compliance and Safety Manua*l (817-2571-10), the contact details given for Zuheir Totari are out of date. The correct contact details are as follows:

Sun Microsystems Ltd Sparc House
Guillemont Park Blackwater
Camberley GU17 9QC
United Kingdom Sun Microsystems
Tel: +44 (0)1252 420113 Fax: +44 (0)1252 421659

Linux

This chapter contains the following sections:

2.1 Installing Linux Onto a B100x or B200x Blade

To install the Linux operating system onto a blade, you must first build a PXE boot installation environment.

The software required to build a PXE boot installation environment is available on the CD supplied with the blade.

To install Linux on B100x or B200x server blades you will need the following:

- The *Sun Fire B1600 Platform Documentation, Drivers, and Installation* CD. This CD includes the drivers required for installing Linux on a server blade, and all the documentation for the B1600 system chassis and its components.
- Installation CDs for the version of Linux you are installing. The following operating systems are supported in this release:
  - Red Hat Enterprise Linux, version 3.0
  - Red Hat Enterprise Linux, Advanced Server 2.1 update 2
  - SuSE Linux Enterprise Server 8, service pack 3
- A PXE boot server machine for installing Linux onto the server blade. This machine must be running one of the following operating systems:
  - Red Hat Enterprise Linux, version 3.0
  - Red Hat Enterprise Linux, Advanced Server 2.1 update 2
  - SuSE Linux Enterprise Server 8, service pack 3
  - Solaris, version 9 or later.

Refer to Chapter 4 of the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* for information on how to perform a PXE boot installation.

2.2 Upgrading the BIOS on B200x Server Blades

To run Red Hat Enterprise Linux version 3.0 or SuSE Linux Enterprise Server 8, service pack 3 on a B200x server blade, the blade must have BIOS version `1.1.32` installed. This version of the BIOS is available from the following website:

`http://wwws.sun.com/software/download/network.html`

You can upgrade the BIOS on the blade using the `biosupdate` utility. This utility loads a device driver called `mtdbios`, performs the BIOS update using the `slflash` utility, and then unloads the `mtdbios` driver.

The `biosupdate` utility is installed on the server blade as part of the PXE boot installation process. For information on the PXE boot installation, see Chapter 4 of the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* (available on the CD supplied with the blade).

Caution – When upgrading the BIOS, do not interrupt the process by resetting or powering down the blade. Interrupting the upgrade will permanently damage the blade.

Note – If the BIOS upgrade fails, a failure message is displayed on the screen and in `/var/log/messages`. If this problem occurs do not reset or power off the blade. Contact your Sun Beta support manager for advise.

2.2.1 To Upgrade the BIOS

1. **Log into the blade for which you want to update the BIOS.**

   At the SC prompt, type:

   ```
   sc> console sn
   ```

   where *n* is the number of the slot containing the blade.

2. **Check the version of the BIOS currently running on the blade, to establish whether the upgrade is necessary:**

```
modprobe mtdbios
cat /proc/BIOS
BIOS Vendor: AMI
BIOS Version: P1.1.32
BIOS Date: 01/19/2004
Manufacturer: Sun Microsystems
Product: Sun Fire B200x
rmmod mtdbios
```

3. **Copy the BIOS image from the beta website to a known location on the blade.**

4. **Run the** biosupdate **command:**

```
biosupdate bios2p.rom-032.bin
```

The blade prompt returns when the update is complete.

Caution – Do not restart the blade while the update is in progress.

---

5. **When the update is complete, reboot the blade:**

```
shutdown -r now
```

You can check the BIOS version when you restart the blade.

2.3 B100x and B200x (Linux) Server Blade Issues

The following problems have been observed on both B100x and B200x server blades:

- 4868095: Red Hat Advanced Server 2.1 does not support layer 2 VLANs

  The Advanced Server kernel (version 2.4.9-e.3) does not support layer 2 VLANs. This means that the `sunconfig` utility is not installed on blades running Red Hat Advanced Server 2.1. If you are using blades with Red Hat Advanced Server 2.1 installed, you must configure the switch to use only untagged VLANs.

  If you require VLAN support, you must install Red Hat Enterprise Linux 3.0 or SuSE Linux Enterprise Server 8 service pack 3.

- 4853227: Spurious interrupt messages appear in `/var/log/messages`

  Spurious interrupt IRQ7 messages may appear in `/var/log/messages`. These messages can be ignored.

2.3.1 Issues Affecting the B200x Server Blade Only

The following problems have been observed on B200x server blades only:

- 44932162: The BSC driver displays names of LEDs not present on the blade

  When the BSC driver is loaded, five extraneous LED files appear in the `/proc/bsc` directory. The LED status for these files is reported as "not present" and they can be ignored.

- 4987508: Kernel panic occurs during booting of B200x server blades (el-3.0u1)

  If you attempt to boot Enterprise Linux 3.0 update 1 with APM enabled, the system will panic. This is due to expectations made by the APM subsystem, which is considered unsafe to run on a multi-processor machine. Without enabling APM (or ACPI), the `/sbin/poweroff` command will not power off the blade.

  This problem has been worked around using the bsc driver. The bsc driver notifies the hardware of the intent to power off, and the hardware powers off the blade 10 seconds later. If the bsc driver is not loaded at power-off time, the system will fail to power off.

- 4991972: B200x blade locks into a repeated "boot net" loop following a BIOS update

  The CMOS footprint may change between revisions of the BIOS. If this is the case, and the CMOS is not reset to its default values during the blade reboot following a BIOS update, the CMOS configuration may become corrupted and result in repeated network boots.

  To avoid this problem, after updating the BIOS on a B200x blade, reset all CMOS settings to their default values when you next reboot the blade. You can do this by typing the `bootmode reset_nvram s`*n* command at the `SC` prompt (where *n* is the number of the slot containing the blade), or by entering the BIOS setup menu and loading the BIOS defaults.

- 5017529: New 1 Gbyte DIMMs used in manufacturing appear as 2x512Mbyte DIMMs

  When new 1 Gbyte DIMMs are inserted in DIMM slots on the B200x blade, the ECC driver incorrectly identifies the memory as 2x512Mbyte DIMMs.

  This is a cosmetic issue, and can be ignored.

- 5015866: Blade does not recognize an SSC which is inserted after a reboot

  When a blade is booted without an SSC present, the network interfaces normally attached to the missing switch will not be operational after the SSC is re-inserted. This problem is fixed when the network interface is brought down and up again.

2.3.2 Issues Affecting the B100x Server Blade Only

The following problem has been observed on B100x server blades only:

- 4915711: The Real Time Clock (RTC) is updated when resetting from the SC

  When you use the `reset` command from the `SC` prompt, the blade updates the Real Time Cock (RTC) from the system controller through the BSC. The RTC should be set only when the blade is powered on.

- 4979474: The bsc driver cannot allocate IRQ9 or IRQ5 (SuSE with kernel 2.4.21)

  The Operating System hangs when the bsc driver attempts to allocate IRQs 9 and 5. This is due to an incorrect ACPI table entry for the bsc hardware.

  You can avoid this problem by booting the kernel using the `pci=noacpi` argument.

- 4906666: The BIOS boot menu displays network names incorrectly

  This is a cosmetic issue and may be ignored.

Solaris x86

This chapter contains the following sections:

Note – If you intend to use the Solaris x86 CD media (instead of the DVD media), you need to use a Solaris x86 system to read the CDs. For more information, refer to Chapter 12 of the *Solaris 9 Installation Guide*. The section of that chapter that you require is called "To Create an x86 Install Server on a SPARC System With x86 CD Media". To use the DVD media, refer to Chapter 11 of the *Solaris 9 Installation Guide*. The section you require is called "Preparing to Install From the Network with DVD Media (Tasks)".

3.1 Preparing to Install Solaris x86 Onto a Blade

Before you start to install Solaris x86 onto a B100x or B200x Blade, please perform the steps in this section.

3.1.1 Solaris x86 Drivers and Documentation

For the first full release of the Solaris x86 software to support B100x or B200x blades, the documentation and some mandatory patches required for Solaris 9 (12/03) are available on the web.

Note – The *Sun Fire B1600 Blade Platform Documentation, Drivers, and Installation* CD that ships with the blade and chassis does not (at the time of writing) contain documentation or patch software for running Solaris 9 x86 on a B100x blade.

To download the documentation and patches you need, do the following:

1. **Go to the website** `http://www.sun.com/servers/entry/b100x/` **or** `http://www.sun.com/servers/entry/b200x/`

2. **In the lefthand column click on the link called "Downloads".**

3. **In the Downloads section, click "Solaris x86 SW Drivers".**

   (If you have not used the Download Service before, you will be invited to register before proceeding.)

4. **Log into the download service.**

5. **Click "Download B100x Solaris x86 Driver Software" and save the packages to the directory** `/var/tmp/blades`**. (The download for the B100x blade is also the download for the B200x blade. They both require the same software.)**

   The file you will download is called `mis.259-4174-11.zip` (The last two digits in this filename indicate the version number. The number is correct at the time of writing, but may have subsequently been incremented).

6. **Return to the web page** `http://www.sun.com/servers/entry/b100x` **or** `http://www.sun.com/servers/entry/b200x/`

7. **Click on the link called "Documentation".**

8. **Download and print out the chapters of the following documents concerning Solaris x86 installation:**
   - *Sun Fire B1600 Chassis, and B100s, B100x, and B200x Blade Product Notes* (this document; you only need to print out the current chapter).
   - *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* (you only need to print the chapter entitled "Installing Solaris x86").

9. **Proceed to Section 3.1.2, "Overview of the Solaris x86 Installation Process" on page 3-3.**

3.1.2   Overview of the Solaris x86 Installation Process

1. **Set up a network install image for Solaris x86 using the Solaris 9 (12/03) Media Kit.**

   For instructions, refer to the *Solaris 9 Installation Guide* supplied with the Media Kit.

   Note – If you intend to use the Solaris x86 CD media (instead of the DVD media), you need to use a Solaris x86 system to read the CDs. For more information, refer to Chapter 12 of the *Solaris 9 Installation Guide*. The section of that chapter that you require is called "To Create an x86 Install Server on a SPARC System With x86 CD Media". To use the DVD media, refer to Chapter 11 of the *Solaris 9 Installation Guide*; the section you require is called "Preparing to Install From the Network with DVD Media (Tasks)".

   ---

2. **Apply the mandatory patches to the Solaris x86 network install image you have created.**

   For instructions, see Section 3.2, "Applying Mandatory Software Patches to the Solaris x86 Install Image" on page 3-3 of these *Product Notes*.

3. **Configure the Network Install Server and the DHCP Server to perform the PXE boot installation for your blade or blades.**

   For instructions, refer to the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*).

3.2   Applying Mandatory Software Patches to the Solaris x86 Install Image

---

Note – The earliest version of the Solaris x86 Operating System supported on the B100x or B200x server blade is Solaris 9 (12/03).

Before following the instructions in the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* for installing Solaris 9 x86 onto a blade, you must follow the instructions in this section of the Product Notes.

This is because the instructions in the blade installation and setup guide assume that you have a later version of Solaris 9 than is currently available. Until the later version becomes available, you need to apply mandatory patches to the Solaris 9 (12/03) install image for the B100x and B200x server blade. This section of the Product Notes tells you how to apply these patches.

Further information about installing the Solaris 9 x86 operating system is available in the *Solaris 9 Installation Guide* supplied with the Solaris 9 media kit. The document can also be downloaded from `http://docs.sun.com`.

3.2.1 Downloading the B100x/B200x Mandatory Software for the Network Install Server

1. **On the Solaris system you are going to use as the Network Install Server, make a directory called** /var/tmp/blades **by typing:**

```
# mkdir -m 755 /var/tmp/blades
```

2. **Download the software if you have not done so already (by following the instructions in Section 3.1.1, "Solaris x86 Drivers and Documentation" on page 3-2).**

3. **Save the download file to the directory** /var/tmp/blades**.**

   The download file is called `mis.259-4174-11.zip`. The filename included here is the correct version number at the time of writing. Because this file is likely to be updated, the final two digits in the name of the file you download may be higher than -11. If so, this indicates that you are downloading a more recent version of the software updates for the B100x and B200x server blades.)

   The download file contains the following B100x- and B200x-specific software:

Patch Number	Patch Title
112234-11 or later	Solaris 9 x86 kernel patch
116485-03 or later	Broadcom Gigabit Ethernet (BGE) driver patch
116483-01 or later	(BSC) Blade Support Chip driver patch
115881-02 or later	`bootconf.exe`, `nbp acpi_intp`, patch

   Note – This download contains all the mandatory patches required for running Solaris 9 (12/03) on the B100x and B200x server blades.

4. **At the Solaris prompt on the system you are going to use as the Network Install Server, unzip the files you have downloaded.**

   To do this, type:

   ```
   # cd /var/tmp/blades
   # unzip mis.259-4174-11.zip
   ```

5. **Proceed to "Creating a Network Install Server" on page 3-6.**

3.2.2　Creating a Network Install Server

To install the Solaris x86 software over the network onto a blade, you must create an install server. This section describes how to set up an install server on the same subnet as the server blade you are about to install, by copying the Solaris x86 CD or DVD images to the hard disk drive on the system that is to perform the role of Network Install Server.

The following procedure refers to Chapter 12 ("Preparing to Install Solaris Software Over the Network") on page 209 of the *Solaris 9 Installation Guide.* This document is supplied with the Solaris 9 media kit. The beginning of the chapter provides background information.

Note – If you are using the Solaris x86 CD media (instead of the DVD media), you need to use a Solaris x86 system available to read the CDs. For more information, refer to Chapter 12 of the *Solaris 9 Installation Guide*. The section of that chapter that you require is called "To Create an x86 Install Server on a SPARC System With x86 CD Media".

---

1. **On the system that is going to be the Solaris x86 Network Install Server, log in and become superuser.**

   This system must include a CD-ROM or DVD drive and be part of the site's network and name service.

2. **Follow the instructions in the** *Solaris 9 Installation Guide* **to set up a Network Install image for Solaris 9 (12/03).**

   These will include the instruction to copy the Solaris 9 image (from its location on a CD or DVD, or from a location on the network) to the install server's hard disk by using the setup_install_server command.

   The setup_install_server command indicates whether or not there is enough disk space available for the Solaris 9 software media images. To determine available disk space, use the df -kl command.

   To run the setup_install_server command, type:

   ```
   # ./setup_install_server install_dir_path
   ```

   where *install_dir_path* specifies the directory that the CD image is to be copied to. The directory must be empty*.

3. **Change to the directory in which you placed** mis.259-4174-11.zip **by typing:**

   ```
   # cd /var/tmp/blades
   ```

4. **Add the patches and packages automatically to the network install server image by typing:**

```
# ./modify_install_server -d install_dir_path
```

where *install_dir_path* is the path to the install image on your install server.

5. **You can now proceed to Chapter 10 of the** *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* **to perform the operating system setup steps for your blades.**

3.3 Issues Affecting B100x and B200x Server Blades That are Running Solaris x86

The known problems listed in this section have been observed to affect both B100x and B200x server blades. They are presented in three groups:

- Section 3.3.1, "Issues for Which You Must Apply a Workaround or Perform an Operational Procedure" on page 3-8
- Section 3.3.2, "Error Messages That Can be Safely Ignored" on page 3-13
- Section 3.3.3, "Other Issues" on page 3-16

3.3.1 Issues for Which You Must Apply a Workaround or Perform an Operational Procedure

4962226: Warning: Jumpstart on Solaris x86 Is Not Single Shot

Caution – In some circumstances a system administrator might choose to boot a blade from the network to recover from possible errors on its hard disk. If you have configured the blade to perform a Jumpstart installation, *any* subsequent network boot of the blade will by default result in a Jumpstart installation being performed. This will erase the contents of the hard disk. Therefore, to prevent the blade from executing a Jumpstart installation (after the first operating system installation), we recommend you remove the SjumpsCF and SsysidCF option names from the blade's client-specific macro after the initial Jumpstart installation has completed. (This network booting behavior is different from that of blades running SPARC Solaris.)

4873161: Need Support for Soft Poweroff in Solaris x86

Solaris x86 does not currently support power button events generated through ACPI. This means that a poweroff command issued on the system controller (sc> `poweroff` s*n*) will not cause an orderly shutdown of the blade's operating system before powering off the blade.

● **Workaround**

To avoid causing possible corruption to the root disk partition by powering off the disk before the operating system has been shutdown in an orderly fashion, first issue a Solaris command to perform an orderly shutdown (for information about different ways to achieve this, refer to the `man` pages for the `shutdown`, `halt`, and `init` commands). For example:

```
# shutdown –i5 –g0
```

The blade can then be safely powered off from the system controller by means of the sc> `poweroff` command. For example:

```
sc> poweroff s2
```

where the '2' indicates the blade in slot 2 of the chassis.

4856947: drv_usecwait is Not Accurate When CPU Frequency Changes

B100x and B200x blades contain CPU processors that go into a power throttling state when only one power supply unit (PSU) is present in the B1600 chassis. During the early stages of the Solaris boot process a number of software timing loops are calibrated. These are affected when the CPU power throttling state changes: they are not currently re-calibrated upon a change of the power throttling state. This means that, if the power throttling state were to change while the blade was running Solaris x86, the timing loops would no longer execute correctly, and the operation of all device drivers making use of critical timing functions would be affected.

In normal use the power throttling state will only change during removal or insertion of a second PSU.

● **Workaround**

If you have removed a second PSU from the B1600 chassis, or if you have inserted a second PSU into the chassis, you can avoid these two issues by rebooting the blades after the PSU insertion or removal.

4856440: Require hostid to Be Set From BSC on Solaris x86 B100x and B200x

The value of the hostid for blades that are running Solaris x86 is different from the hostid value programmed into the B1600 chassis for the blade's physical location. When Solaris x86 is installed for the first time onto a blade, the hostid value is generated by the install process. It is generated as a random unique value, once for the life of the blade. Solaris x86 does not currently support changes to this value under software control. The value is maintained for all subsequent installations of Solaris on the same blade by being stored in an inaccessible location on the hard disk.

If you are replacing a blade, or if you are moving a blade from one slot in the chassis to another, note that the blade does not inherit its hostid from its new physical location in the chassis. (Solaris x86 blades differ in this respect from blades running Linux or SPARC Solaris.)

4945519: Solaris x86 Sometimes Fails to Reboot After Jumpstart Install

During a Jumpstart installation the disk partitioning can get into a state where it causes a validation check to give a false positive. This is indicated by the following WARNING message:

```
Installing 32-bit Solaris Packages
        - Selecting locale (en_US.ISO8859-1)
        - Selecting all disks
        - Configuring boot device
        - Creating "maxfree" Solaris fdisk partition (c0d0)
        - Using existing Solaris fdisk partition (c0d0)
        - Automatically configuring disks for Solaris operating environment

Verifying disk configuration
        - WARNING: Change the system's BIOS default boot device for hands-off
rebooting
```

When this happens the system will not reboot at the end of the Jumpstart installation. However, the error message can be safely ignored and the system simply rebooted.

● **Workaround**

You can workaround this problem by including the reboot command at the end of the Jumpstart x86-finish script. For more information, refer to the chapter entitled "Installing Solaris x86" in the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*.

4975579 Bootpath Line Missing After PXE Install With Single Solaris Partition

During an interactive Solaris installation the install program will prompt you to select the partition layout of the disk. If you choose to layout the disk with a single disk partition that combines x86boot and Solaris partitions, then when the blade is rebooted the 'Device Configuration Assistant' screen will be displayed and you will be prompted to select the boot device.

● **Workaround**

To avoid this, manually select the option to define seperate x86boot and Solaris disk partitions. For instructions about how to do this, refer to the information on disk partitioning in the chapter of the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide* entitle entitled "Installing Solaris x86".

4852503: `break` Command Does Not Work Although Message Confirms a `break` Was Sent

Solaris X86 supports the `break` command when the kernel has been booted under control of the kernel debugger (`kadb`).

Under normal operation the blade will not be booted under `kadb` control and it may appear that the `break` command has had no effect. On a SPARC system the `break` command would cause the operating system to drop to the Open Boot Prom (OBP) `ok` prompt. However, this facility is not available on Solaris x86 systems because they use BIOS instead of OBP.

Nevertheless, if you issue the break command from the System Controller's `sc>` prompt, for example:

```
sc>break s1
Are you sure you want to send break to FRU s1 (y/n)? y
s1: Break sent.
```

then the System Controller will still send the `break` command to the blade even if it is running Solaris x86. Therefore application software that is running on the blade can receive and interpret the command.

4922593: Link Messages Seem to Contradict IPMP State During a Switch Reset

If you have IPMP configured on a blade and you reset the chassis's integrated switch, you will see an error message that appears to contradict the IPMP configuration of the blade. The console output during the reset will possibly display several link up messages while the switch is physically in the process of resetting (see below). There will be no corresponding link down messages. This behavior is a result of the way the bge driver negotiates with the switches. It has no negative impact on network connectivity or on the operation of IPMP. When the switch reset completes, the link status will be correct and IPMP will be functioning correctly.

```
sc>reset -y SSC0/SWT
sc>console s1
Sep 15 15:35:04 bladeS1 bge: NOTICE: bge0: link up 1000Mbps Full-Duplex
Sep 15 15:35:12 bladeS1 in.mpathd[110]: NIC failure detected on bge0 of group
test
Sep 15 15:35:12 bladeS1 in.mpathd[110]: Successfully failed over from NIC bge0
to NIC bge1
Sep 15 15:35:13 bladeS1 bge: NOTICE: bge0: link up 1000Mbps Full-Duplex
Sep 15 15:35:42 bladeS1 bge: NOTICE: bge0: link up 1000Mbps Full-Duplex
Sep 15 15:36:03 bladeS1 in.mpathd[110]: Successfully failed back to NIC bge0
Sep 15 15:36:03 bladeS1 in.mpathd[110]: NIC repair detected on bge0 of group test
```

3.3.2 Error Messages That Can be Safely Ignored

This section lists error messages that will be be observed during a network installation or reboot. In all cases these messages can be safely ignored. They have no impact on B100x and B200x blades.

4903388 `consconfig` *complains on servers with no frame buffer*

When the B100x and B200x blades boot you will see the following message:

```
SunOS Release 5.9 Version Generic_112234-11 32-bit
Copyright 1983-2003 Sun Microsystems, Inc.  All rights reserved.
Use is subject to license terms.
WARNING: consconfig:  could not find driver for screen device /isa/display@1,3b0
WARNING: Could not attach frame buffer to wscons, error 6
configuring IPv4 interfaces: bge0.
starting DHCP on primary interface bge0
```

This message is generated because the blades do not have a frame buffer. It can be safely ignored.

4921001: `/etc/bootrc` *needs to skip 30 second delay when* `boot-args` *is set to* `jumpstart`

When you configure a SPARC Solaris custom Jumpstart installation, the installation program does not prompt you to choose whether to perform an interactive or a Jumpstart installation. Instead it reports that it is proceeding with a Jumpstart installation.

There are two issues, however, that affect Jumpstart installations of Solaris x86. The installation program does prompt you to choose either an interactive or a Jumpstart installation even if you have previously configured a Jumpstart. It pauses for up to 30 seconds and when this time has elapsed the installation program reports that it is 'starting interactive installation' even if you have configured the blade to perform a Jumpstart installation.

Despite this message, the blade will perform a Jumpstart installation if you have configured it to do so.

```
Select the type of installation you want to perform:

        1 Solaris Interactive
        2 Custom JumpStart

Enter the number of your choice followed by the <ENTER> key.
Alternatively, enter custom boot arguments directly.

If you wait for 30 seconds without typing anything,
an interactive installation will be started.

Select type of installation:

<<< starting interactive installation >>>

Booting kernel/unix...
SunOS Release 5.9 Version Generic_112234-11 32-bit
Copyright 1983-2003 Sun Microsystems, Inc.  All rights reserved.
Use is subject to license terms.
```

You can easily confirm that a Jumpstart is in fact taking place because, early in the booting and installation process, the output will contain references to the Jumpstart configuration files (rules.ok, x86-begin, x86-class, x86-finish).

For example:

```
Booting kernel/unix...
:
Starting Solaris installation program...
Searching for JumpStart directory...
Using rules.ok from 123.123.123.163:/export/jumpstart.
Checking rules.ok file...
Using begin script: x86-begin
Using profile: x86-class
Using finish script: x86-finish
Executing JumpStart preinstall phase...
Executing begin script "x86-begin"...
Begin script x86-begin execution completed.
Searching for SolStart directory...
Checking rules.ok file...
```

where the ':' character indicates information that has been removed from the sample output.

4871718: bootstrap complains about devices that have been disabled in the BIOS

On the B100x you will see the following message about /dev/diskette0 during a reboot or PXE boot:

```
CLIENT MAC ADDR: 00 03 BA 29 F0 DE  GUID: 00000000 0000 0000 0000 000000000000
SunOS Secondary Boot version 3.00 255.255.255.0  DHCP IP: 129.156.205.163
GATEWAY IP: 129.156.205.8
/dev/diskette0: device not installed, unknown device type 0


        Solaris Intel Platform Edition Booting System
```

On the B200x you will see the following message about /dev/diskette0 during a reboot or PXE boot:

```
SunOS Secondary Boot version 3.00


/dev/diskette0: can't open - bios configuration error


        Solaris Intel Platform Edition Booting System
```

On the B200x you might also see the following message during a network installation. This message refers to an attempt by Solaris to interact with a floppy disk controller that is disabled by the BIOS.

```
Configuring /dev and /devices
 WARNING: fdgetchng: write protect check failed

 Using DHCP for network configuration information.
 WARNING: fdgetchng: write protect check failed

 WARNING: fdgetchng: write protect check failed

 Searching for configuration file(s)...
 WARNING: fdgetchng: write protect check failed

 WARNING: fdgetchng: write protect check failed

 Search complete.
```

4321917: ACPI Resource Conflicts Unnecessarily Reported

During the booting of a B200x blade you might see the following message briefly
displayed:

```
Warning: Resource Conflict - both devices are added
NON-ACPI device: PNP0C01
   NON-ACPI device: PNP0C01
  Memory: 9FC00-9FFFF, 3FFF0000-3FFFEFFF, 3FFFF000-3FFFFFFF,
       Memory: 9FC00-9FFFF, 3FFF0000-3FFFEFFF, 3FFFF000-3FFFFFFF,
   FEC00000-FECFFFFF, FEE00000-FEE00FFF, FFF00000-FFFFFFFF
ACPI      FEC00000-FECFFFFF, FEE00000-FEE00FFF, FFF00000-FFFFFFFF
device: PNP0C01
ACPI device: PNP0C01
     Memory: 9F000-9FFFF, E0000-EFFFF, 40FFC00-4D8BF373,
     5      Memory: 9F000-9FFFF, E0000-EFFFF, 40FFC00-4D8BF373,
6C5737C-FFFFFFFF
     56C5737C-FFFFFFFF
```

If you see this message in respect of a B200x blade, you can safely ignore it.

3.3.3   Other Issues

*4877872: Difference Between the Network Interface Assignments
Displayed by the BIOS, DCA, and Solaris x86 Software on a B200x*

The B200x blade has two dual-port BCM5704s Gigabit Ethernet chips. Each
individual port is connected to one of the Ethernet switches in the B1600 chassis, and
the BIOS takes responsibility for assigning the MAC addresses to the Ethernet ports
(see FIGURE 3-1).
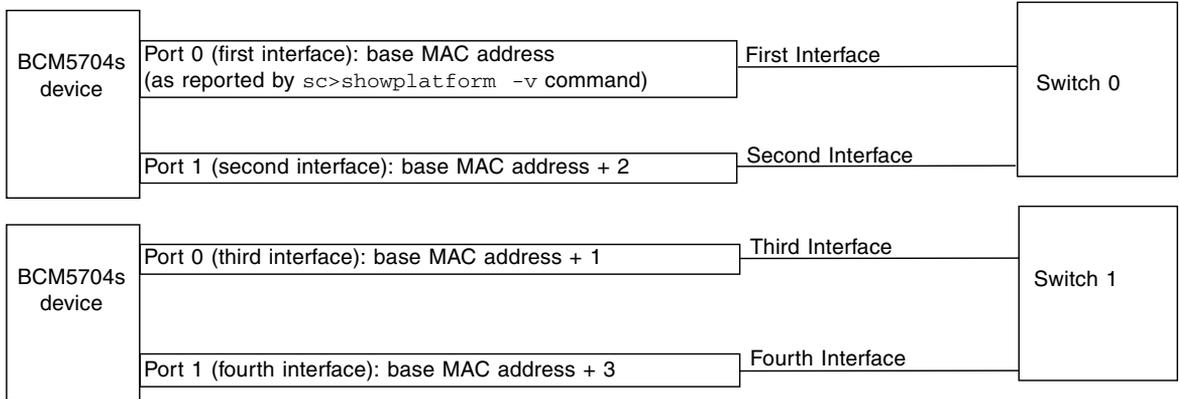
FIGURE 3-1 The Network Interfaces on a B200x Blade

The B200x BIOS displays the network interfaces as shown in FIGURE 3-2.

```
1st Boot Device                    [PM-TOS MK3019GAXB ]
2nd Boot Device                    [SNET0 MBA v6.2.11 ]
3rd Boot Device                    [SNET2 MBA v6.2.11 ]
4th Boot Device                    [SNET1 MBA v6.2.11 ]
5th Boot Device                    [SNET3 MBA v6.2.11 ]
```

FIGURE 3-2 The Network Interfaces Listed by the B200x BIOS

The names that the BIOS uses to identify each network interface correspond with the names used by the System Controller's `bootmode` command (see Section 10.10, "Installing Solaris x86 Onto a Blade by Using the Second, Third, or Fourth Network Interface" in the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*).

If the blade is booted without the bootpath property having been set in `/boot/solaris/bootenv.rc`, the Solaris Device Configuration Assistant (DCA) pauses to request the boot device and display the network devices as shown in FIGURE 3-3.

The network interfaces are displayed by the DCA in the order in which it discovers them when it probes the hardware.

```
[ ]  DISK: Target 0:TOS MK30 19GAXB SUN30G
                on Bus Mastering IDE controller on Board PCI bus 0, at Dev 31,
   [ ]  NET : Broadcom Gigabit Ethernet              <base MAC + 1, snet1>
          on Board PCI bus 3, at Dev 3, Func 0
   [ ]  NET : Broadcom Gigabit Ethernet              <base MAC + 3, snet3>
          on Board PCI bus 3, at Dev 3, Func 1
   [ ]  NET : Broadcom Gigabit Ethernet              <base MAC + 0, snet0>
          on Board PCI bus 4, at Dev 3, Func 0
   [ ]  NET : Broadcom Gigabit Ethernet              <base MAC + 2, snet2>
          on Board PCI bus 4, at Dev 3, Func 1
```

FIGURE 3-3   The Boot Devices as Displayed by the Device Configuration Assistant

When the network interfaces are plumbed under Solaris they are displayed in the order shown in FIGURE 3-4.

```
# ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
        inet 127.0.0.1 netmask ff000000
bge0: flags=1004843<UP,BROADCAST,RUNNING,MULTICAST,DHCP,IPv4> mtu 1500 index 2
        inet 123.123.123.202 netmask ffffff00 broadcast 123.123.123.255
        ether 0:3:ba:2d:d4:a0
bge1: flags=1000842<BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
        inet 123.123.123.203 netmask ffffff00 broadcast 123.123.123.255
        ether 0:3:ba:2d:d4:a2
bge2: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 4
        inet 123.123.123.204 netmask ffffff00 broadcast 123.123.123.255
        ether 0:3:ba:2d:d4:a1
bge3: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 5
        inet 123.123.123.205 netmask ffffff00 broadcast 123.123.123.255
        ether 0:3:ba:2d:d4:a3
```

FIGURE 3-4   The B200x Blade's Network Interfaces as Displayed by Solaris x86

From FIGURE 3-1 and FIGURE 3-4 you can see that the bge0 and bge1 interfaces have even MAC addresses and are connected to Switch 0 (SSC0), and that the bge2 and bge3 interfaces have odd MAC addresses and are connected to Switch 1 (SSC1).

The B100x and B200x blades are factory configured so that the base MAC address is an even MAC address.

If you are configuring IPMP network redundancy, note that the achievement of network resilience (enabling a blade to recover from different hardware and network failures) depends upon each IPMP group containing one connection to each switch. A configuration in which both interfaces in a group of two were connected to the same switch would not continue to transport network traffic if that switch failed.

For information about configuring the blade to use IPMP, refer to the chapter entitled "Configuring IPMP for Network Resiliency on Solaris x86 Blades" in the *Sun Fire B100x and B200x Server Blade Installation and Setup Guide*.

CHAPTER **4**

SPARC Solaris

This chapter contains the following sections:

- Section 4.1, "Installing SPARC Solaris Onto a B100s Server Blade" on page 4-2
- Section 4.2, "B100s (SPARC Solaris) Server Blade Issues" on page 4-2

4.1 Installing SPARC Solaris Onto a B100s Server Blade

To install the Solaris 8 HW 12/02 (Build 5) Operating Environment onto a B100s blade, follow the instructions in the *Sun Fire B1600 Blade System Chassis Software Setup Guide,* the *Solaris 8 Installation Guide* (806-0955), and the *Solaris 8 Advanced Installation Guide* (806-0957).

4.2 B100s (SPARC Solaris) Server Blade Issues

The following problems are known to affect the Sun Fire B100s Server Blade:

- 4877079: Blades running SunVTS sometimes power down when the System Controller is reset.

  This problem is known to affect the Blade Support Chip firmware in version 5.1.3 on the server blades. When a System Controller resets it attempts to read FRUID information from each FRU in the chassis; if it fails to read this information for a particular FRU, it powers that FRU down. It is known that when SunVTS is performing its BSC test on a blade this can cause the FRUID EEPROM to become temporarily unreadable by the System Controller. If the System Controller is reset while this test is executing on a blade, then it might not be able to read the FRU information for that blade. And if it cannot, it will power the blade down. To avoid the possibility of experiencing this problem, you would need to disable the BSC test in Sun VTS.

- <u>4726915: SunVTS will not start because of missing XML libraries</u>
  To run SunVTS you need to load the `SUNWlxml` and `SUNWlxmlx` packages. These are not loaded by default in Solaris 8, therefore you must add them manually from the supplemental CD supplied with Solaris 8 HW 12/02. They are located in the following directory on the CD: `XML_Libraries_2.4.12/Product`.

  To make the contents of this CD available via NFS, on the NFS server, type:

  ```
  # share -F nfs /cdrom/solaris8_hw1202_suppcd
  ```

  Then, to load the packages onto the server blade, access the blade's console, and type:

  ```
  # mount -F nfs server:/cdrom/solaris8_hw1202_suppcd /mn
  #pkgadd -d /mnt/XML_Libraries_2.4.12/Product SUNWlxml SUNWlxmlx
  ```

- <u>4779970: When booting a blade with the kernel debugger (`kadb`) enabled, a warning message appears</u>
  If you boot a blade with the kernel debugger enabled, you will see the message "`WARNING: todblade: kernel debugger detected: hardware watchdog disabled`". Because you have enabled the kernel debugger, the kernel watchdog functionality cannot be used. The message simply warns you of this fact. There is no other effect upon the operating system functionality.

- <u>4803500: If you insert a blade less than 10 seconds after removing it, the service LED is lit on the blade and chassis, and service required messages are logged</u>
  After using the `removefru` command to make a blade safe for removal, remove the blade, and then wait at least 10 seconds before re-inserting it or before inserting another blade into the same slot.

- <u>Shared library patches for C++ are not included in Solaris 8 HW 12/02.</u>
  The 32-Bit and 64-bit shared library patches for C++ (`libC.so.5`, `libCrun.so.1, and libdemangle.so.1`) are not included in the Solaris 8 HW 12/02 release. They are available on `http://sunsolve.sun.com`. The numbers for these patches are: 108434-10 and 108435-10.

- 4811241: When you install Solaris with "Entire distribution plus OEM support" and error message appears
  If you install Solaris with "Entire distribution plus OEM support" onto a server blade, the following error message will appear on the console and in the log file maintained by the Solaris install process (there is no effect on the operation of the server blade and no corrective action is required to be taken by the user):

```
(/var/sadm/system/logs/install_log):
/a/var/sadm/pkg/SUNWidecr/install/postinstall: test: argument
expected pkgadd: ERROR: postinstall script did not complete
successfully
Installation of <SUNWidecr> failed.
```

System Controller

This chapter contains the following sections:

5.1 Upgrading the System Controller (SC) Firmware

To support Sun Fire B200x blades in a Sun Fire B1600 blade system chassis, you must be running System Controller firmware version 1.2. To perform the upgrade of the System Controller firmware, follow the instructions in the *Sun Fire B1600 Blade System Chassis Administration Guide* (Chapter 10).

5.2 What To Do If You Lose Your Password for the System Controller

There is a method described in Chapter 9 of the *Sun Fire B1600 Blade System Chassis Administration Guide* (on the Documentation CD) for regaining access to the System Controller if you have forgotten your password.

Note – It is possible that the method described on the CD you have received is incorrect.

If you lose your password to the System Controller, the correct procedure for regaining access to the device is as follows:

1. **Remove and then re-insert one of the power supplies.**

2. **Within five minutes of re-inserting the power supply, set up a serial connection to the SSC containing the active System Controller and log in as user** admin**.**

   For information about setting up a serial connection to the SSC, see the *Sun Fire B1600 Blade System Chassis Hardware Installation Guide*.

   To log in, type the default user name admin and when prompted for a password, press [ENTER].:

   ```
   username: admin
   password: [ENTER]
   ```

3. **At the** `sc>` **prompt, set a new password for the default user (**`admin`**).**

   To set a new password for user `admin`, type:

   ```
   sc>password
   Enter new password:*****
   Enter new password again: *****
   New password set for user admin successfully
   sc>
   ```

   The password you specify can be up to eight characters long. It must begin with an alphabetic character, and it must contain at least one lower-case alphabetic character. You can use any of the following characters in the password:

   - Alphabetic
   - Numeric
   - Period (.)
   - Underscore (_)
   - Hyphen (-)

4. **Set up a new user name and password for yourself.**

   To do this, follow the instructions in Chapter 3 of the *Sun Fire B1600 Blade System Chassis Administration Guide*.

5.3 System Controller Software Issues

5.3.1 System Controller Firmware 1.2

The following known problems apply to the current release of the System Controller firmware for this product:

■ 4879114: System Controller sometimes resets or fails over if its netmask is invalid: If you have an invalid netmask configured for the System Controller, you may see the following or a similar error message in response to ARP requests generated by the System Controller:

```
arptnew failed on xxxxxxxx
```

where *xxxxxxxx* is the hexadecimal representation of an IP address. If you see an error message like this, it will be followed by a reset of the System Controller (if you are operating the chassis with a single System Controller) or a failover to the standby System Controller (if you are operating a chassis containing a redundant System Controller).

To avoid this problem make sure you configure the System Controller with a netmask that is valid for your network. If you are using DHCP to allocate IP addresses, make sure your DHCP server is configured to provide a netmask that is valid for your network.

■ 4885940: Poweron of Standby SSC fails due to spurious environmental errors: If you run the `removefru` command followed by the `poweron` command on the standby System Controller without physically removing the SSC module containing the standby System Controller, you may see spurious environmental errors being notified for the standby SC. These errors do not indicate any real problem with the SC hardware and the problem can be cleared by issuing a `poweron` command for the standby SC. Alternatively, physically removing and re-inserting that SC will also clear the problem.

5.3.2 System Controller Firmware 1.1

The following known problems applied to release 1.1 of the System Controller firmware (they are fixed in release 1.2):

- 4866668: When the message 'Dynamic FRUID packet discovery failure recovered' shows in the logs for a blade, it is possible that some fault events will be missing or only partially described.

  This problem can occur if the System Controller is reset immediately after an event has been reported on the serial console. In other words, it happens if the System Controller has not completed writing the information to its event logs when the reset is performed.

  To avoid the problem, if you see an event reported on your serial console, wait at least 30 seconds before executing the `resetsc` or `setfailover` command.

- 4868129: Typing `reset -xy s`n (where n is the number of a blade slot) resets all FRUs in a shelf.

  The syntax of the above command is correct, but the software does not process it correctly. The consequence, if you type the command in this way, is that all FRUs in the chassis will be reset.

  To avoid this problem, when you perform an externally initiated reset on a blade (by using the `-x` option) and you specifically do not want to receive the confirmation prompt (the `-y` option), make sure you separate the arguments as follows when you type the command:

  ```
  sc>reset -x -y sn
  ```

  where n is the number of the slot containing the blade you want to reset.

5.3.3　System Controller Firmware 1.0

The following known problems applied to release 1.0 of the System Controller firmware (they are fixed in release 1.1):

- 4810785: Recovery from output rail faults is not correctly reported under some circumstances. A PSU output rail fault is correctly reported when a fault occurs, but recovery from the fault may not be correctly reported under certain circumstances. For instance, when a PSU output rail fault is detected (the most likely reason for this is a blade fault causing a PSU rail to stop providing power), a fault will be logged by the System Controller.

  However, when the underlying fault is removed (in other words, when the blade causing the problem is removed from the chassis), the System Controller will not report the recovery of this rail and will continue to light the service required LED on the affected PSU and on the chassis (even though service is not in fact required).

  When the component causing the underlying problem has been replaced, the `showplatform` and `showenvironment` commands will correctly report that there are no faults on the system. However, the service LED on the PSU will still remain lit. To turn it off after replacing the component that caused the underlying problem, type the following commands from the System Controller's `sc>` prompt:

  ```
  sc>removefru -y psn
  sc>poweron psn
  ```

  where *n* is either 0 or 1 depending on the Power Supply Unit involved. (You do not need to remove the Power Supply Unit.)

  Note – Executing the `removefru` command prepares the PSU for removal but does not stop the PSU from providing power to the chassis. However, it does prevent environmental monitoring of the PSU by the System Controller. This is restored when you run the `poweron psn` command. Therefore, when you do this, the service LED will be reset.

- 4826948: System Controller hangs if you telnet into it, type removefru (without -y) for a specified FRU, and then close the telnet window.
  If you telnet into the System Controller and use the `removefru` command without the `-y` option, you are asked to confirm that you want to remove the FRU. If you do not answer but instead close the telnet window, the System Controller hangs. To restart the System Controller, you must eject it from the chassis and then push it back in.

The System Chassis's Integrated Switch

This chapter contains the following sections:

6.1 Switch Firmware Issues

The following known problems apply to the current release of the switch firmware for this product:

■ 4899178: Blade network traffic is only allowed through the IP filter on the VLAN configured as the management VLAN.
The management VLAN is the VLAN that has been assigned an IP address to allow network access to the switch's management interfaces (by default this is VLAN2). Other VLANs can be assigned to the NETMGT port to allow blades to talk to particular hosts on the management network. The usual way to do this would be to put particular blades and particular hosts that are on the management network onto a tagged VLAN that is separate from the management VLAN. However, the switch's packet filter will not forward any traffic from the blades to the management network unless that traffic is for the management VLAN. This is a problem that will be fixed in the next release of the switch firmware. It means that traffic from the blades will not be seen by hosts on the management network that are external to the chassis (in other words, only other blades in the chassis will see the traffic) unless the blade, the switch, and the external hosts involved are all on the management VLAN. Note that the configurations for multiple tenants described in Chapter 7 of the Sun Fire B1600 Blade System Chassis Software Setup Guide will not be possible until this problem has been fixed.

■ 4854587: It is possible that the System Controller will reset the switch when the switch is executing commands that require unusually intensive processing: The SC continuously polls the switch for status as part of the system healthcheck. It is theoretically possible that, while executing commands requiring unusually intensive processing, the switch will be unable to respond to the SC's status request within the timeout period because it must first complete execution of a process-intensive command.

This should not happen during normal operation. The problem was observed when a user sent a sequence of lengthy commands (for example, commands adding many VLANs to a port) to the switch without waiting for the prompt between each command. This filled the switch's input buffer and blocked the status poll messages.

To avoid the problem, always wait for one command to complete before issuing another command on the CLI. If you are using scripts this is especially important.

- <u>4871779: Blades are unable to receive multicast packets when IGMP querying is enabled.</u>
  Multicasting on the switch does not work correctly if the IP address for the switch is not configured. This is only likely to be the case if the IP address is configured by DHCP and the DHCP process has failed for any reason. If necessary, you can work around the problem by specifying an IP address for the switch manually:

  ```
  Console#configure
  Console(config)#ip address ip address  netmask
  ```

- <u>4885056: The switch does not have ingress filtering enabled by default.</u>
  By default ingress filtering is set to disabled. This allows packets from VLANs other than the VLANs explicitly enabled on each port to pass through the switch. This is a security risk. To enable ingress filtering on the switch, you must enable it on each port individually. The following example demonstrates how to enable it for NETP0. Repeat the commands for each port (from NETP1 through NETP7 and SNP0 through SNP15) and aggregate link (port-channel):

  ```
  Console#configure
  Console(config)#interface ether NETP0
  Console(config-if)#switchport ingress-filtering
  ```

- <u>4894936: Auto-negotiation and speed/duplex mode cannot be configured on the NETMGT port</u>
  The NETMGT port's speed and duplex mode are fixed at 100Mbps and full duplex. This is not manually configurable. However, it is possible to connect to the NETMGT port using a 10BaseT full- or half-duplex connection, or a 100BaseT full- or half-duplex connection. If you do this, the switch's internal hub negotiates the speed and duplex mode automatically with the interface at the other end of the connection. If you attempt to set the speed and duplex mode for NETMGT manually, you will receive an error message to the effect that the `interface ethernet NETMGT speed-duplex` command failed. In future releases of the switch firmware the error message will explain that NETMGT cannot be manually configured using this command.

- <u>4780304: Adding a port to the forbidden VLAN list sometimes fails to remove the VLAN from that port's VLAN list</u>
  This error is only seen when the port you are trying to add to the forbidden list is the last VLAN on that port's VLAN list (apart from the native VLAN, which can never be removed). Firmware that fixes this problem is now available from `http://sunsolve.sun.com` (patch number: 114783-xx).

- 4876495: The port status is unstable

  There is a known problem with Spanning Tree (STP and RSTP) when it is used with aggregated links. When you have an aggregated link (a single link comprised of multiple ports) between two switches and you enable spanning tree on that aggregated link, the spanning tree control packets on the trunk link are not reliably received by the switch in the chassis. The effect of this is that spanning tree does not converge, the port state of the aggregate link keeps changing, and network connectivity is likely to be disrupted. Firmware that fixes this problem is now available from `http://sunsolve.sun.com` (patch number: 114783-xx). Until you have the fixed firmware installed, do not enable spanning tree (either STP or RSTP) on any aggregated link.

- 4804804: no switchport allowed vlan command fails

  There is a known problem with the `no switchport allowed vlan` command that enables you to remove all VLANs except the native VLAN from a particular port. Issuing this command on a port that has learned a VLAN by using GVRP causes the learned VLAN to be assigned to the switch's VLAN database as a static VLAN instead of being removed from the database. If you need to remove static VLANs from a port that has GVRP enabled, we recommend you use the `no switchport allowed vlan remove` *vlanid* command (where *vlanid* is the number identifying the static). If you do use the `no switchport allowed vlan` command, you must delete manually from the VLAN database any unrequired VLANs.

  To do this, type the following:

  ```
  Console#configure
  Console(config)#vlan database
  Console(config)#no vlan vlanid
  ```

  where *vlanid* is the number identifying a VLAN that you want to remove from the switch's database. Firmware that fixes this problem is now available from `http://sunsolve.sun.com` (patch number: 114783-xx).

- Sometimes error messages from the switch refer to ports by number instead of name. The correct mapping of port names to port numbers is:

Port Names	Port Numbers
SNP0-SNP15	1/1-1/16
NETP0-NETP7	1/17-1/24
NETMGT	1/25

- The integrated switches on the Sun Fire B1600 blade system chassis are each composed of two switch chips linked together. It is only possible to mirror the traffic on one port by using another port that is on the same switch chip. And it is only possible to enable flow control between two ports on the same switch chip.

The ports NETP0, NETP1, NETP4, NETP5, and SNP8 through SNP15 are on one switch chip. The ports NETP2, NETP3, NETP6, NETP7 and SNP0 through SNP7 are on the other.

- Multiple Spanning Tree Protocol: the `spanning tree mst` commands are not available in the current release of the switch firmware.

- The switch's DHCP client identifier is set by the System Controller. This means that, if you set it using the switch's command-line interface, web GUI, or by using the SNMP interface to the switch, the setting you specify will be overwritten by the System Controller next time the switch boots. The command to change the DHCP client identifier will be removed from the next release of the switch firmware.

- 4831855: Date set incorrectly on the switch.
  If you set the date on the System Controller (SC) to anything other than the first of the month, then the datestamp on events logged by the switch afterwards will one day behind the current date on the System Controller. For example, if the date according to the System Controller is Wednesday March 26, 2003, then the datestamp on switch event logs will have the datestamp become 3/25/3. The only workaround for this is to wait until the first of the next month, and reset the date on the System Controller.

- 4804197: AN983 internal loopback test reports false failure
  There is a remote possibility that an inaccurate failure report will be generated by the AN983 internal loopback test performed during a switch reset (including following execution of the reload command). If the NETMGT port is accessible from the network, then you can ignore the failure report. The failure will persist over several resets until the SSC undergoes a hard reset (in other words, until you power cycle the SSC. If you close down all web, SNMP, and telnet connections to the switch before you perform the reset, you will not see this problem.

- 4799549: Broadcast ping from a blade on the management network will not receive a response from any external hosts
  If you issue a broadcast ping onto the management network from a server blade, you will not receive any responses from host devices external to the switch (in other words, you will only receive a response from the switch's NETMGT port and from other blades inside the chassis that are also on the management network). This is a known problem and it will be fixed in the next release. However, you can ping known hosts individually on the management network. And if you log into a known host on the management network and issue the broadcast ping from there, you will receive a response from all the host devices on the management network (including all the host devices inside the chassis that are on the management network). Firmware that fixes this problem is now available from `http://sunsolve.sun.com` (patch number: 114783-xx).

- The switch's DHCP client identifier is set by the System Controller. This means that, if you set it using the switch's command-line interface, web GUI, or by using the SNMP interface to the switch, the setting you specify will be overwritten by

the System Controller next time the switch boots. The command to change the DHCP client identifier will be removed from the next release of the switch firmware.

- <u>4795640: Resetting with the factory default configuration causes provisioning errors</u>
  Saving a copy of the switch's factory default configuration file (or saving a modified copy of this file) generates errors if the switch is then rebooted with the saved copy specified as the startup configuration file.You will only see these errors if you press p when asked if you want to view details of the startup provisioning. The errors can be ignored.

- <u>4773404: No traffic statistics available for the NETMGT port</u>
  There is a known problem with the command for viewing traffic statistics. The output for the NETMGT port when you run the `show interfaces counters ethernet NETMGT` command (from the console# prompt) contains zeroes instead of valid data. There is currently no workaround for this problem.

- <u>4773404: No MAC address table available for the NETMGT port</u>
  (This issue has the same Sun number as the previous issue.) There is a problem with the command for displaying the MAC address table for the NETMGT port. The `show mac-address-table interface ethernet NETMGT` command (from the `console#` prompt) always displays an empty table. There is currently no workaround for this problem.

- <u>4789838: LACP sometimes fails if GVRP is enabled</u>
  There is a known problem with the operation of the link aggregation control protocol and the dynamic VLAN configuration protocol GVRP. It is not possible for the LACP protocol to operate reliably if GVRP is enabled. Therefore, if you are using GVRP do not enable LACP. Firmware that fixes this problem is now available from `http://sunsolve.sun.com` (patch number: 114783-xx).

- <u>4773408: Spanning tree mode cannot be set when spanning tree support is disabled</u>
  Setting the spanning tree mode for the switch can only be done when spanning tree is enabled. If you wish to set the initial spanning tree mode for the switch to a particular setting (for example in a configuration file) you must ensure that spanning tree is enabled before issuing the spanning-tree mode command. Type:

```
Console#configure
Console(config)#spanning-tree
Console(config)#spanning-tree mode rstp
```

  To change the initial spanning tree mode from the default (RSTP) to STP with spanning tree disabled, type:

```
Console#configure
Console(config)#spanning-tree
Console(config)#spanning-tree mode rstp
Console(config)#no spanning-tree
```

- <u>4790634: The SSC-to-switch communication protocol might not supply the switch's DHCP client identifier in time for the switch to make a DHCP request</u>
  On a reset of the switch after the SSC unit has been moved into a different slot or into a new chassis (or after the DHCP client identifier has been changed manually from the switch's command-line interface), it is theoretically possible that the switch will make two DHCP requests. This has not been observed in testing. There is currently no workaround for this issue; however, there are no serious consequences of it either.

6.2 Issues Affecting the Web Graphical User Interface to the Switch

A graphical user interface (GUI) is available for configuring the switch. To access it, point a web browser at the host name or IP address you have used for the switch.

The following problems have been observed during testing of the web GUI. Sun bug numbers are included where these are available.

- 4743657, 4744678, 4772618: The Software Download and Upload page gives no indication of progress during the download and upload
  When the Software Download and Upload window refreshes itself, the transfer operation is not fully complete. A further few minutes are required for the new firmware to be programmed into flash memory (from switch RAM). Do not attempt to perform another download until you can see the first file appear in file list when you click the Reload button on your browser.

- 4876509: Minor display problems when Internet Explorer is used to access the web GUI
  There are some minor display problems associated with the use of Internet Explorer to access the web Graphical User Interface. For example, highlighting an item in a list box might cause the disappearance of the next item in the list. If you experience a problem with the display, correct it by refreshing the page.

- 4879052: The switch's web server hangs if you try to configure port mirroring incorrectly
  On the page Sun Fire B1600 => Monitoring => Port Mirroring do not click the "Remove" button without first selecting a port. If you do click the "Remove" button and you have "None" selected, your GUI session will hang, and no other users will be able to log into the GUI. To regain access (for yourself and other users) to the web GUI, you must reset the switch.

- 4743657, 4744678, 4772618: The Software Download and Upload page gives no indication of progress during the download and upload
  When the Software Download and Upload window refreshes itself, the transfer operation is not fully complete. A further few minutes are required for the new firmware to be programmed into flash memory (from switch RAM). Do not attempt to perform another download until you can see the first file appear in file list when you click the Reload button on your browser.

- 4829016: Address tables displayed incorrectly
  This issue concerns the Switch Config=>Address Tables window. If a port has learned some MAC addresses, then querying the address table for a port or VLAN should display the type of address as '`permanent/dynamic/delete on reset`'. If the port is made into a secure port from the Up Links=>Static Addresses window, then any dynamically learned addresses are displayed as '`EMPTY`' when they should be displayed as of type '`Learned-PSEC`'.

- 4828965: Disabling global GVRP state prevents dynamic VLANs from leaving ports

  If you disable GVRP globally on the switch by issuing the following command:

  ```
  Console(config)#no bridge-ext gvrp
  ```

  then VLANs that have been learned dynamically are not dropped even after the GVRP leave-all timer has expired (normally 10 seconds). These VLANS remain active on the ports that learned them, and you must remove them manually. The following sample command removes the dynamically learned VLAN called vlan 3 from NETP4:

  ```
  Console#show vlan
  VLAN Type    Name             Status    Ports/Channel groups
  ---- ------- ---------------- --------- -----------------------
  1   Static   DefaultVlan      Active     SNP0    SNP1    SNP2
                                           SNP3    SNP4    SNP5
                                           SNP6    SNP7    SNP8
                                           SNP9    SNP10   SNP11
                                           SNP12   SNP13   SNP14
                                           NETP0   NETP1   NETP2
                                           NETP3   NETP4   NETP5
                                           NETP6   NETP7
  2   Static   MgtVlan          Active     NETMGT
  3   Dynamic                   Active     NETP4
  Console#configure
  Console(config)#interface ether NETP4
  Console(config-if)#switchport allowed vlan remove 3
  Console(config-if)#exit
  Console(config)#vlan database
  Console(config-vlan)#no vlan 3
  ```

- Error messages are incomplete.

- The web GUI (Monitoring=>Port Statistics=>NETMGT) cannot provide traffic statistics for the NETMGT port. The data for the NETMGT port appears as zeroes instead of valid data. There is currently no workaround for this problem.

- Adding packet filtering rules from the Management Port=>Packet Filtering page must be performed with care, because blank fields on the page will default to a value of zero. Make sure that you have entered a value into every field you require, and check that the rule displayed when you click Add is the rule that you require. In particular, the Protocol Number box next to the Protocol Name box will accept a protocol name (instead of a numerical value) without displaying an error, and if you type a name instead of a numerical value the field defaults to zero.

- The pages of the web GUI include options for configuring an HTTPS server. This functionality is not enabled in the current release of the switch firmware.

- The web GUI (Switch Config=>Address Tables=>NETMGT Port ID) cannot display the MAC address table for the NETMGT port. This table is always empty. There is currently no workaround for this problem.

- The web GUI (Switch Config=>Spanning Tree=>View=>MST instance configuration
=>MST Instance Setup) the multiple `spanning tree options` are not configurable in the current release of the switch firmware.

6.3 The Term "Trunk" Meaning Either an Aggregated Link Or a Tagged VLAN Connection

There is confusion in the networking industry over the term "trunking" because it is used to refer both to link aggregation and to tagged VLAN connections between two switches. In the first of these senses it means a redundant high-bandwidth path between two switches. In the second it means a network connection on a LAN segment that is populated only with VLAN-aware devices.

6.3.1 Aggregated Links

You may have encountered the term "trunking" in the sense of link aggregation if you have used the Sun Trunking 1.2.1 product.

Ports can be statically grouped into an aggregate link to increase the bandwidth of a network connection or to ensure fault recovery. Alternatively, you can use the Link Aggregation Control Protocol (LACP) which automatically negotiates an aggregated link between the switch and another network device. For static aggregated links, the switches must be of the same type. For dynamic aggregated links, the switches simply have to comply with LACP. The switch in the blade system chassis supports up to six aggregated links. An aggregated link consisting of two 1000 Mbps ports can support an aggregate bandwidth of 4 Gbps when operating at full duplex.

To implement a configuration combining the ports NETP0 and NETP1 into an aggregated link called channel group 1, you would type the following commands:

```
Console#configure
Console(config)#interface port-channel 1
Console(config-if)#exit
Console(config)#interface ethernet NETP0
Console(config-if)#channel-group 1
Console(config-if)#exit
Console(config)#interface ethernet NETP1
Console(config-if)#channel-group 1
Console(config-if)#exit
Console(config)#exit
Console#
```

6.3.2 Switch-to-switch Tagged VLAN Trunk Connections

The *Sun Fire B1600 Blade System Chassis Switch Administration Guide* also uses the term "trunking" in the sense of a point-to-point tagged VLAN connection between two switches. Section 4.3.12 tells you how to configure the chassis's end of a connection like this to an external switch, and section 4.3.12.4 tells you how to use the "switchport mode" command to specify that the connection is a "trunk" (as opposed to a "hybrid") connection. If you specify "trunk" the port transmits and receives tagged frames only - in other words, it sends and receives only frames that identify their source VLAN. (However, note that it sends frames belonging to its default VLAN untagged.) If you specify "hybrid" the port will transmit and receive tagged and untagged frames.

To set the configuration mode for port SNP3, and then to set the switchport mode to trunk for VLANs 12 and 22, you would type the following commands:

```
Console#configure
Console(config)#interface ethernet SNP3
Console(config-if)#switchport allowed vlan add 12 tagged
Console(config-if)#switchport allowed vlan add 22
Console(config-if)#switchport native vlan 22
Console(config-if)#switchport allowed vlan remove 1
Console(config-if)#switchport ingress-filtering
Console(config-if)#switchport mode trunk
Console(config-if)#switchport acceptable-frame-types tagged
Console(config-if)#end
```

6.4   Setting up a Tagged VLAN Trunk With Cisco Switches

There is a known problem with setting a switch port to trunk mode if that port is connected to a port on a Cisco switch also in trunk mode (note that we use the word "trunk" in the sense of a point-to-point link, not in the sense of an aggregated link). This is because of a standardization issue (Cisco comply with the Cisco standard whereas the switch in the blade system chassis complies with the IEEE 802.1Q standard). It means that it will drop frames from the Cisco switch port's native VLAN.

To work around this problem, you need to configure the system chassis's switch port to hybrid (not trunk) mode, make sure that it has the same native VLAN Id as the Cisco switch, and also make sure that all the VLANs requiring connection to the Cisco switch have been added to the port. You must also stop packets for VLANs that the port is not a member of from entering the port.

Commands for a sample workaround are printed below. These assume a system chassis port (NETP0) with VLAN 1 as its native VLAN and hybrid as its link mode (this is the factory default configuration for the system chassis's network ports).

The commands for the sample workaround also assume a Cisco switch port with trunk as itslink mode, VLAN 10 as its native VLAN, and additional membership of VLANs 11 and 12.

The commands for the workaround in this scenario are:

```
Console#configure
Console(config)#interface ethernet NETP0
Console(config-if)#switchport allowed vlan add 10
Console(config-if)#switchport native vlan 10
Console(config-if)#switchport allowed vlan remove 1
Console(config-if)#switchport allowed vlan add 11 tagged
Console(config-if)#switchport allowed vlan add 12 tagged
Console(config-if)#switchport ingress-filtering
Console(config-if)#end
Console(config)#
```