Release Notes for the Accelar 1000 Series Products Software Release 1.1

Accelar 1100/1150 Routing Switch
Accelar 1200/1250 Routing Switch
Accelar Boot Monitor Software Version 1.1
Accelar Runtime Software Version 1.1
Accelar Device Manager Version 1.1
Accelar VLAN Manager Version 1.1

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 896-00181-B
April 1998

Bay Networks

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Contents

Introduction

These release notes provide the latest information for the Bay Networks® Accelar™ Software Release 1.1.0. Currently released components include:

- Boot Monitor Software Version 1.1.0
- Runtime Software Version 1.1.0
- Device Manager Version 1.1.0.b6
- VLAN Manager Version 1.1.0.b6

This document contains the following major sections:

- Flash file system information (beginning on page 2), including compatibility information (page 4), software upgrade/downgrade instructions (page 5), and management software upgrade instructions (page 21)
- Descriptions of the new and enhanced features of Accelar software (page 28)
- Known software limitations at the time of release (page 41)
- Bugs fixed in release 1.1 (page 51)
- Known problems in release 1.1 (page 54)

For the latest information about software issues, refer to the Accelar site on the Bay Networks Web page (baynetworks.com) or contact Bay Networks Customer Support at 1-800-2LANWAN.

Related Publications

For more information about the Accelar products, refer to the following documents on the Accelar documentation CD:

- *Installing the Accelar 1000 Series Chassis* (Bay Networks part number 893-01051-B)
- *Using the Accelar 1100/1150 Routing Switch* (Bay Networks part number 893-01050-A)
- *Using the Accelar 1200/1250 Routing Switch* (Bay Networks part number 893-01049-A)
- *Reference for the Accelar Management Software* (Bay Networks part number 893-01052-B)

Accelar 1000 Series Version 1.1 Flash File System

This section gives a brief description of the flash file system in an Accelar 1000 Series routing switch, along with the procedure for upgrading the flash software image from version 1.0 to version 1.1. For more information about the file system, including file types, file naming, and file system commands, refer to Chapter 2, "Accelar Management Basics," of the *Reference for the Accelar Management Software.*

Flash Organization

There are two onboard flash memory devices on the Accelar routing switch: the Boot Flash and the System Flash. On Accelar 1200 Series switches, optional PCMCIA flash cards can be used. The flash file system on an Accelar 1000 Series routing switch holds executable images and the switch configuration.

Accelar 1000 Series routing switches have two software images that reside in flash memory: the boot monitor and the runtime image. The boot monitor image is low-level code that initializes the devices on the Silicon Switch Fabric (SSF) module and starts the boot process. The runtime image is the code that executes initializing the I/O modules and provides full routing switch functionality.

Boot Flash

The Boot Flash is 512 kilobytes (KB) and is divided into reserved areas for the boot monitor image and the routing switch configuration. The boot monitor image is not directly user accessible. It is updated using a special boot monitor updater that writes to the area reserved for the boot image.

By default, the routing switch configuration is stored in a reserved area in the boot flash, although it is possible to specify alternative locations in the file system for the switch configuration. The routing switch configuration is written whenever a save operation is performed on the configuration of the device.

The area reserved in Boot Flash for the switch configuration is accessed by the file system commands using the "config" or "nvram" file names. Both config and nvram refer to the same file. Note that the switch configuration is read only when the runtime image loads.

System Flash

The System Flash is 4 megabytes (MB) and is primarily used for runtime images, the system log, configuration files, and other general storage. The System Flash is divided into 64K blocks. Files stored in the System Flash are always stored in an integral number of blocks.

Files stored in the System Flash are numbered sequentially starting with the number one (1). Files can be assigned names by the user or referenced by file number in flash memory. The file name for System Flash files is `flash:<filename>` or `flash:<file#>` where `flash` is the device name, `<filename>` is the file name, and `<file#>` is the number of the file.

PCMCIA

Accelar 1200 Series routing switches can use an optional PCMCIA flash memory card. PCMCIA cards can be used for general storage for all file types and are a convenient way of moving files between switches because they are portable.

System Flash File Types

Although System Flash and PCMCIA are primarily used for runtime images, configuration files, and the system log, they are also used to store other types of files. *Executables* are images that are executed by the Accelar 1000 Series CPU. The two most common executables needed by users are runtime images and boot monitor updaters. These executables can be stored in the flash file system in zipped format to conserve space. If necessary, the routing switch will automatically unzip the file upon execution.

The *runtime image* is an executable file that executes after the boot monitor image, initializing the I/O modules and providing full routing switch functionality. A runtime image can be stored and executed from System Flash and PCMCIA. It is stored in flash memory in an area accessible through the command line interface (CLI).

The *boot monitor image* is low-level code that initializes devices on the SSF module and starts the boot process. The boot monitor image is stored in a reserved area in flash memory that cannot be accessed directly. It is updated by executing a boot monitor updater that in turn updates the boot monitor image.

In addition to the area in Boot Flash reserved for the switch configuration, *configuration files* can be stored on PCMCIA and flash memory. A new feature in version 1.1 allows you to source the configuration file from these types of media when booting. Thus you can now source the configuration file from PCMCIA, flash, or nonvolatile RAM (NVRAM).

Prerequisites for Upgrading

To upgrade the flash images, you need the following:

- New version 1.1 runtime and boot monitor image files
- TFTP server
- Network connection to the TFTP server
- Direct console connection or a Telnet connection

Note: Bay Networks recommends a direct console connection for performing the upgrade procedure because it is REQUIRED if it is necessary to reverse the upgrade procedure.

The Accelar 1000 Series routing switch has 4 MB of flash memory for images. Typical sizes for the various files mentioned in this document are as follows: a runtime image is less than 1 MB, a boot monitor updater is less than 100K, and the log file is 128K.

Note: Upgrading flash memory requires rebooting the switch and will disrupt network operation. The upgrade procedure should be performed during network downtime.

Compatibility

Version 1.1 of the Accelar software CAN use a version 1.0 configuration stored in the Boot Flash. Version 1.1 CANNOT use a version 1.0 configuration file stored in System Flash or PCMCIA. Although one of the new features of version 1.1 software is the ability to use configurations stored in areas other than the reserved configuration area in the Boot Flash, this feature works only with version 1.1 configuration files. If you configure version 1.1 to use a version 1.0 configuration file in flash memory or PCMCIA, the configuration will not load and the routing switch will boot with a factory default configuration.

A version 1.0 configuration file is converted to version 1.1 format if it is copied into the reserved “config” or “nvram” Boot Flash area. Once converted to version 1.1 format, the configuration file can then be copied and loaded in flash memory or PCMCIA.

Caution: Version 1.1 configuration files ARE NOT compatible with version 1.0 software. Loading a version 1.1 configuration with a version 1.0 runtime image may result in the failure of the switch to boot.

The following update procedure is structured with these compatibility issues in mind to allow proper upgrading and downgrading, if required.

Before beginning the upgrade, you should also be aware that Accelar Management Software (Device Manager and VLAN Manager) and system software images must be the same level to be compatible. Release 1.1 management software IS NOT compatible with Accelar version 1.0 runtime and boot images; release 1.0 management software IS NOT compatible with version 1.1 runtime and boot images. Therefore, for switch management, you should plan to upgrade management software at the same time you upgrade the runtime and boot image software. You should also plan to keep a copy of the version 1.0 management software on your system until all switches have been updated to the version 1.1 software image.

Version 1.1 Flash Memory Update Procedure

This procedure is a generic upgrade procedure that will work on all Accelar switches. For more product-specific procedures that may involve fewer steps, refer to the Accelar site on the Bay Networks Web page. The procedure assumes that you will be updating both the boot monitor image and the runtime image in flash memory to version 1.1. Be sure to read and understand the procedure BEFORE attempting to upgrade the software.

The procedure for reverting to a pre-upgrade state follows this update procedure. Bay Networks recommends reviewing that procedure also prior to starting the software upgrade.

Figure 1 is a flow chart of the recommended software upgrade sequence. Numbers in parentheses ( ) indicate the number of the detailed step that follows.

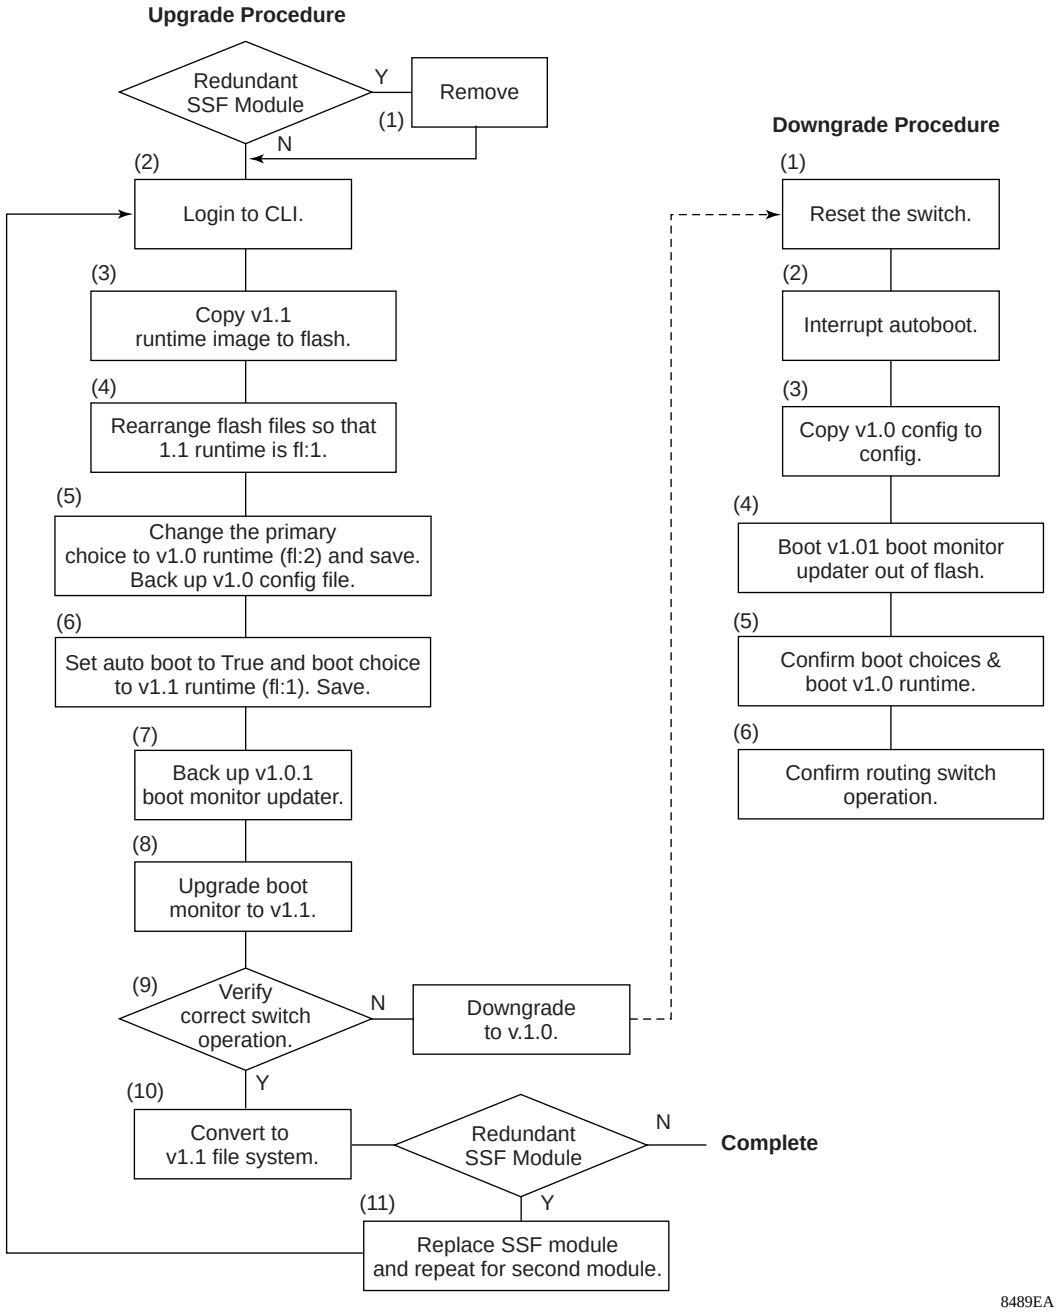

Figure 1. Software Upgrade/Downgrade Flow Chart

Caution: It is VERY IMPORTANT that you follow the upgrade instructions carefully in order to simplify the downgrading of the software to version 1.0 in case that becomes necessary. In particular, DO NOT convert the flash file system to version 1.1 format until AFTER you have verified proper routing switch operation under version 1.1.

Note: Before beginning any upgrade procedure, whenever possible, it is good practice to back up your configuration files, as well as the current runtime and boot image to another location, using TFTP or PCMCIA.

To upgrade the flash software from version 1.0 to version 1.1:

1. **Remove the redundant SSF module, if present.**

   For Accelar 1200 chassis with redundant SSF modules, the upgrade should be focused on one SSF module at a time to avoid confusion. Remove the standby SSF module, if present, from the chassis. You will later need to repeat the flash memory update procedure on the second SSF module to ensure that the flash image versions are synchronized.

   The upgrade procedure requires rebooting the SSF module. If the redundant SSF module is not removed, rebooting the active SSF will cause SSF modules to swap active and standby roles, complicating the upgrade process.

2. **Log in to the runtime CLI.**

   The update procedure can be performed at the local console or through a Telnet connection. The console port on the SSF module is a serial DTE device operating at 9600 bps, 8 data bits, no parity, and 1 stop bit.

Note: The console port is the preferred method because you must be at the console to downgrade to version 1.0.1 if necessary.

   At the CLI `Login:` prompt, enter the password for read-write or read-write-all access privileges. By default for versions 1.0.0 and 1.0.1, the read-write access password is “private.”

```
********************************
* Bay Networks,Inc.            *
* Copyright (c) 1996-1997      *
* All Rights Reserved          *
* Accelar 1100                 *
* Software Release 1.0.0       *
********************************

Login: *******

Bay>
```

3. Copy the version 1.1 runtime image into flash.

The version 1.1 runtime image is loaded into the Accelar switch by a TFTP copy operation initiated at the Accelar CLI. Before attempting a TFTP copy, it is useful to verify that the TFTP server is reachable and responding. You can do so by using the `ping` command, initiated at the CLI prompt. In the following example, the TFTP server has an address of 10.10.20.100:

```
Bay> ping 10.10.20.100
10.10.20.100 is alive
Bay>
```

The runtime CLI `copy` command allows the copying of files to and from flash memory, PCMCIA cards, and TFTP servers. To copy the new runtime image file from the TFTP server to flash memory, use the `copy` command as shown here, specifying the TCP/IP address of the TFTP server and the path to the new runtime image file as follows:

```
Bay> copy tftp flash
Enter source tftp server address [0.0.0.0]: 10.10.20.100
Enter source file []: acc1_1_0
programming flash:3 ... xxxxxx bytes
verifying flash:3 ... done
tftp result: success
Bay>
```

The runtime image is a large file and can take about one minute to download. The runtime CLI copy command works silently and does not indicate the progress of the TFTP copy.

Note: Some Windows-based TFTP servers do not support long file names even though long file names are displayed. If your TFTP server does not support long file names, you may experience TFTP failures. You can rename the image files without affecting the upgrade process.

4. Rearrange files in the flash system.

To simplify the upgrade and downgrade process, you should move the new runtime image to the first file position in the flash file system. The files are rearranged in flash memory using the `copy`, `delete`, and `squeeze` commands.

The flash directory currently looks like:

```
Bay> directory flash
Device: flash
------

File 1:
/export/acura3/projects/first/rel1.0/rel1.0.1/main/hw/accelar.s
t
Version: rel1.0/rel1.0.1/main on Thu Mar 12 17:44:58 PST
Length: 897670 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x1cc8

File 2: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 3,0 free: 252,256

File 3: .0.st on Sat Mar 28 23:33:07 PST 1998
Version:
/export/acura3/projects/first/rel1.1/rel1.1.0/main/hw/acc1.1.0.
st on Sat Mar 28 23:33:07 PST 1998
Length: 948247 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0xd6af

Files: 3 BlocksUsed: 31 BytesUsed: 2031616 BlocksFree: 33
BytesFree: 2162688

Bay>
```

In the copy step, the current (1.0.1) runtime image (flash:1) and the log file (flash:2) will be copied to the end of flash; the 1.1 runtime image will remain as file three (flash:3).

```
Bay> copy flash:1 flash
programming flash:4 ... 895659 bytes
verifying flash: 4 ... done
Bay> copy flash:2 flash
programming flash:5 ... 131072 bytes
verifying flash:5 ... done

Bay> directory flash
Device: flash
------
```

```
File 1:
/export/acura3/projects/first/rel1.0/rel1.0.1/main/hw/accelar.s
t
Version: rel1.0/rel1.0.1/main on Thu Mar 12 17:44:58 PST
Length: 897670 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x1cc8

File 2: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 53,0 free: 202,256

File 3: .0.st on Sat Mar 28 23:33:07 PST 1998
Version:
/export/acura3/projects/first/rel1.1/rel1.1.0/main/hw/acc1.1.0.
st on Sat Mar 28 23:33:07 PST 1998
Length: 948247 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0xd6af

File 4:
/export/acura3/projects/first/rel1.0/rel1.0.1/main/hw/accelar.s
t
Version: rel1.0/rel1.0.1/main on Thu Mar 12 17:44:58 PST
Length: 897670 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x1cc8

File 5: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 53,0 free: 202,256

Files: 5 BlocksUsed: 47 BytesUsed: 3080192 BlocksFree: 17
BytesFree: 1114112

Bay>
```

Now delete files 1 and 2 and squeeze the flash file system. This process has the effect of renumbering the files so that the version 1.1 runtime image is now the first file in flash (flash:1), the version 1.0.1 runtime image is the second file (flash:2), and the system log file is the third file (flash:3). The `delete` command marks a file for deletion; the `squeeze` command actually deletes the files and compresses the file system, renumbering the files in the process.

```
Bay> delete flash:1
Bay> delete flash:2
Bay> squeeze flash
recovering deleted file space
formatting flash ...  done
programming flash:1 ... 948423 bytes
verifying flash:1 ... done
```

```
programming flash:2 ... 897850 bytes
verifying flash:2 ... done
programming flash:3 ... 131072 bytes
verifying flash:3 ... done
2 files squeezed from device flash
Files: 3 BlocksUsed: 31 BytesUsed: 2031616 BlocksFree: 33
BytesFree: 2162688
```

Now the directory looks like:

```
Bay> directory flash
Device: flash
------
File 1: .0.st on Sat Mar 28 23:33:07 PST 1998
Version:
/export/acura3/projects/first/rel1.1/rel1.1.0/main/hw/acc1.1.0.
st on Sat Mar 28 23:33:07 PST 1998
Length: 948247 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0xd6af

File 2:
/export/acura3/projects/first/rel1.0/rel1.0.1/main/hw/accelar.st
Version: rel1.0/rel1.0.1/main on Thu Mar 12 17:44:58 PST
Length: 897670 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x1cc8

File 3: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 3,0 free: 252,256

Files: 3 BlocksUsed: 31 BytesUsed: 2031616 BlocksFree: 33
BytesFree: 2162688

Bay>
```

5. **Set the primary boot choice to version 1.0 runtime and back up the version 1.0 configuration file.**

   The version 1.0 configuration file can be used by the version 1.1 runtime image, but the version 1.0 runtime image will not read the version 1.1 configuration files. Therefore, you should copy the current version 1.0 configuration file to flash memory in case you need to revert to the version 1.0 runtime image.

Caution: Attempting to use version 1.1 configuration files with the version 1.0 runtime image may crash or freeze your system.

However, before backing up the current configuration, change the primary boot choice to flash:2, which is the version 1.0 runtime image. This selection will then be saved in the configuration you copy to the flash file system, allowing you to boot from the version 1.0 runtime image if necessary.

```
Bay> sys set primary flash:2
Bay> save
Bay> copy config flash
programming flash:4 ... 4376 bytes
Bay>
```

6. **Set Autoboot to True and set the primary boot choice to version 1.1 runtime.**

   Enable autoboot and set the primary boot choice to the first file in the flash file system (that is, the version 1.1 runtime image):

```
Bay> sys set flags autoboot true
Autoboot is enabled
Bay> sys set primary flash:1
Bay> save
Bay>
```

7. **Copy the version 1.0.1 boot monitor updater to flash memory.**

Note: Copy the version 1.0.1 boot monitor updater (accboot.gz). The version 1.0.1 boot monitor updater can update a broader range of flash memory chips than can the version 1.0.0 boot monitor updater.

This step is also a precautionary backup step in case of downgrade. The version 1.0.1 boot monitor updater is a required file if a downgrade to version 1.0 is necessary.

```
Bay> copy tftp flash
Enter source tftp server address [10.10.20.100]: 10.10.20.100
Enter source file [acc1_1_0]: accboot.gz
programming flash:5 ... 85543 bytes
verifying flash:5 ... done
tftp result: success
Bay>
```

8. **Upgrade the boot monitor.**

Caution: This step reboots the routing switch and can disrupt network operation.

The boot monitor upgrade process involves running a boot monitor updater image (an executable) from the TFTP server. This updater image will load the new boot monitor into the NVRAM of your Accelar switch.

You can upgrade your boot monitor by invoking the boot monitor updater image located on your TFTP server using the boot net command. The syntax of this command is `boot net <tftp_server> <image_file>`, where `<tftp_server>` is the IP address of the TFTP server and `<image_file>` is the file to download and execute (in that order) on the TFTP server. Note that in version 1.1, the boot command has additional features, so the syntax is different.

If you are performing the upgrade when connected to the switch with a Telnet connection, you will lose connectivity after the boot command executes. Additionally, when the boot monitor is updated and the switch boots with the version 1.1 runtime image, the banners will not be visible through Telnet, but they will appear on the local console (if attached). When the upgrade process completes, the switch will reboot into the new runtime image, retaining your previous switch configuration. When this reboot is complete, typically in a couple of minutes, you should be able to reestablish your Telnet session with the Accelar switch.

```
Bay> boot net 10.10.20.100 accboot1_1_0
tftp result: success
```

Note: If you are running boot monitor version 1.0.1, the system will NOT automatically reset after the TFTP. You must continue on in the process by typing “reset” at the next `Bay>` prompt.

```
Bay> reset
```

If you invoked the `boot net` command from the local console, then you will see upgrade and boot messages similar to the following:

```
Accelar Monitor v1.0.0

CPU: 50Mhz PPC 403GCX Type 3 Rev 2 in Slot 4
DRAM: 16M
Chassis: 1100
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 3 is occupied
Power supply 2 not present
```

```
running boot script at 0x8000f004

monitor> boot dram 0x80c54538
Booting from [dram] System Memory...
Unzipping file enbios/flash/accboot1.1.0 on Tue Mar 3 16:
Version /export/acura3/projects/first/rel1.1/rel1.1.0
/boot/403_EVB/openbios/flash/accboot1.1.0 on Tue Mar 3 16:
from 0x804000b0 to 0x80010038 87041 to 576992 bytes

############### ACCELAR CPU BOOT FLASH Update ###############

File ACCBOOT.ROM found in loaded image

Boot Monitor Version v1.1.0

Image size: text 132924 data 26848 bss 11240
Number of flash sectors to be programmed: 3

WARNING:  You are about to re-program your Boot Monitor FLASH
          image. Do NOT turn off power or press reset
          until this procedure is completed. Otherwise
          the card may be permanently damaged!!!

Do you wish to continue? (y or n)y

Erase of 3 sectors completed
-Verifying new FLASH Image...
196608 matches, 0 mismatches

Update complete!

Press any key to reboot
```

After the boot monitor is upgraded, the routing switch automatically reboots and loads the version 1.1 runtime image.

```
Accelar Monitor v1.1.0

CPU: 50Mhz PPC 403GCX Type 3 Rev 2 in Slot 4
DRAM: 16M
Chassis: 1100
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 3 is occupied
Power supply 2 not present

Press any key to stop autoboot
User Selected Boot Sources
 Primary   = flash:1
 Secondary = flash:
 Tertiary  = net
 Config    = nvram
```

```
Booting from [flash:1] on-board flash memory
Configuration from [nvram] ...
Unzipping file acc1.1.0
Details /export/acura3/projects/first/rel1.1/rel1.1.0
/main/hw/acc1.1.0.st
on Wed Mar 4 15:19:49 PST 1998
from 0x760000b0 to 0x80010000 944692 to 4998914 bytes
Attaching network interface nicEvb0... done.
Null 0.0.0.0 inet address for interface nicEvb.
Attaching network interface lo0... done.

Accelar System Software Release 1.1.0
Copyright (c) 1996-1998 Bay Networks, Inc.

INFO : code=0x0  Task rcStart:  System log file flash:system log
file:0:1

Loading configuration from nvram
Initializing card in slot #3 ... OK

The system is ready.
```

9. Verify routing switch operation.

After performing the boot monitor upgrade, the Accelar switch should reboot into the version 1.1 runtime image (flash:1) and use the configuration stored in NVRAM.

At this point, you should log in to your switch by way of either the local console or Telnet and verify that your switch is configured and operating properly in the new 1.1 environment.

The version 1.1 runtime image has three levels of access (read-only, read-write, and read-write-all). Both a user name and a password must be specified. The default user name and password for read-write access is “rw.”

Note: If you entered a Web password under version 1.0, this password becomes the default for Web, console, and Telnet under version 1.1. Enter this password to access the CLI to verify operation.

```
*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.1.0      *
*******************************
```

```
Login: rw
Password: **

Accelar-1200#
```

If the routing switch is not operating properly, use the version 1.0 downgrade procedure described starting on page 17. You will need to be at the routing switch console to downgrade the software to version 1.0.

10. **Convert to the version 1.1 file system.**

Note: Before completing this step, verify that the routing switch is operating properly with the version 1.1 software. **After the file system is converted to the version 1.1 format, the version 1.0 runtime image will not load.** The procedure to get a version 1.0 runtime image on a routing switch with a version 1.1 file system requires local console access and completely erasing the flash image.

The flash file system has been enhanced in version 1.1, so the existing files need to be upgraded to the new format. The version 1.1 image supports both file numbers and file names in the flash file system whereas version 1.0 supported only file numbers. In the following sample directory, the N flag indicates that the file has been updated to the new file naming format:

```
Accelar-1200# directory flash
Device: flash
FN Name                                Flags      Length
-- ----                                -----      ------
1  acc1.1.0                              XZN      948247
2 /export/acura3/projects/first/
  rel1.0/rel1.0.1/main/hw/accelar.st     XZ      897670
3  system log file                        L      130896
4  config                                 C        4200
5  /export/acura3/projects/first
  rel1.0/rel1.0.1/boot/403_EVB/
  openbios/flash/fls hbo                 XZ       85367
                                         --      ------
5   files           bytes used=    2228224  free=1966080

Accelar-1200#
```

To convert all files in flash (or PCMCIA) to the version 1.1 file system format, use the recover command.

```
Accelar-1200# recover flash
recovering files ... success
Accelar-1200# directory flash
Device: flash
FN Name                                 Flags      Length
-- ----                                 -----      ------
1  acc1.1.0                              XZN      948247
2  accelar.st                            XZN      897670
3  syslog                                 LN      130896
4  config                                 CN        4200
5 accboot1.1.0                           XZN       85367
                                          --      ------
4   files          bytes used=      2228224  free=1966080

Accelar-1200#
```

11. **Repeat the process for the second SSF module.**

    Following the preceding steps has updated the flash images on one SSF module. For Accelar 1200 chassis with redundant SSF modules, repeat the process on the second SSF module to ensure synchronization of the flash image versions.

Downgrade Procedure to Version 1.0

The downgrade procedure from version 1.1 to version 1.0.1 described here assumes that the file system has not been converted to a version 1.1 format. If the file system has been converted to a version 1.1 format, contact Bay Networks Technical Support for information about the downgrade process.

Note: You MUST perform the downgrade procedure at the local routing switch console.

To downgrade the software image from version 1.1 to version 1.0:

1. **Reset the routing switch.**

   You can reset the switch either by logging in to the CLI or by pressing the hardware reset button.

2. **Interrupt autoboot.**

   When the switch begins to boot, press any key to interrupt the autoboot and bring you to the boot monitor prompt:

   ```
   Accelar Monitor v1.1.0
   ```

```
CPU: 50Mhz PPC 403GCX Type 3 Rev 2 in Slot 4
DRAM: 16M
Chassis: 1100
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 3 is occupied
Power supply 2 not present

Press any key to stop autoboot

monitor>
```

3. **Copy the version 1.0 switch configuration to config.**

   During the upgrade procedure, the version 1.0 configuration with primary boot choice of flash:2 was saved to flash memory. You can use the directory command to locate the file:

```
monitor> directory flash
Device: flash
FN Name                                Flags      Length
-- ----                                -----      ------
1  acc1.1.0                              XZN      948247
2 /export/acura3/projects/first/
  rel1.0/rel1.0.1/main/hw/accelar.st  XZ      897670
3  system log file                        L      130896
4  config                                 C        4200
5  /export/acura3/projects/first
  rel1.0/rel1.0.1/boot/403_EVB/
  openbios/flash/fls hbo                XZ       85367
                                        --      ------
5   files          bytes used=   2228224  free=1966080

monitor>
```

   In the preceding directory listing, the saved configuration is flash:4.

Caution: If you load a version 1.1 configuration file using a version 1.0 runtime image, the switch may fail to boot.

   Copy the configuration to config:

```
monitor> copy flash:4 config
erase of config complete
write of config complete
monitor>
```

4. **Boot the version 1.0.1 boot monitor updater out of flash.**

Note: You must perform this step at the routing switch local console because you must confirm downgrading the boot monitor image and press a key to reset the routing switch.

During the upgrade procedure, the version 1.0.1 boot monitor updater was copied to flash memory (flash:5 in step 3). Boot flash:5 to downgrade to the version 1.0.1 boot monitor, responding to the prompts as needed:

```
monitor> boot flash:5
Booting from [flash:5] on-board flash memory
Configuration from [nvram] ...
Unzipping file
/export/acura3/projects/first/rel1.0/rel1.0.1/boot/403_EVB/open
bios/flash/flshbo
Details rel1.0/rel1.0.1/boot/403_EVB/openbios/flash on
from 0x762000b0 to 0x80010038 85367 to 576992 bytes

############### ACCELAR CPU BOOT FLASH Update ################

File ACCBOOT.IMG found in loaded image

Boot Monitor Version v1.0.1

Image size: text 129428 data 26448 bss 10872
Number of flash sectors to be programmed: 3

WARNING:  You are about to re-program your Boot Monitor FLASH
          image. Do NOT turn off power or press reset
          until this procedure is completed. Otherwise
          the card may be permanently damaged!!!

Do you wish to continue? (y or n)y
Erase of 3 sectors completed
-Verifying new FLASH Image...
196608 matches, 0 mismatches

Update complete!

Press any key to reboot
```

Press any key to reboot and again press any key to stop the boot process at the boot monitor prompt.

```
Accelar Monitor v1.0.1

CPU: 50Mhz PPC 403GCX Type 3 Rev 2 in Slot 4
DRAM: 16M
Chassis: 1100
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 3 is occupied
Power supply 2 not present

Press any key to stop autoboot

monitor>
```

5. **Confirm boot choices and boot the version 1.0 runtime image.**

   After the upgrade procedure, the version 1.0 runtime image should be the second file in flash (flash:2). Use the `choices` command to verify that the primary boot choice is the second flash image:

```
monitor> choices
---   CHANGE BOOT CHOICES  ---
Current Boot Choices :
 Primary   = flash:2
 Secondary = flash:1
 Tertiary  = net
  1 - Change primary bootsource
  2 - Change secondary bootsource
  3 - Change tertiary bootsource
  0 - cancel
select boot choice to change [1]: 0
monitor>
```

   Continue the boot process to boot the version 1.0 runtime image:

```
monitor> boot

The system is ready.

*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1997     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.0.1      *
*******************************

Login:
```

6. **Confirm routing switch operation.**

   After rebooting the routing switch, you will be running the same runtime software and configuration as prior to the version 1.1 upgrade. The flash file system will have additional files and the primary boot choice may be different, but the switch will operate as it did prior to the upgrade.

Installing Accelar Management Software

Release 1.1 of the Accelar Management Software includes version 1.1 of Device Manager and VLAN Manager software. These applications MUST be installed to manage switches that have been upgraded to version 1.1 runtime and boot monitor software images. Management software and system software images must be the same level to be compatible. Release 1.1 management software IS NOT compatible with Accelar version 1.0 runtime and boot images; release 1.0 management software IS NOT compatible with version 1.1 software images.

Bay Networks recommends that when you install the version 1.1 management software, you do not delete release 1.0 management software from your system until you are sure that all switches have been upgraded to version 1.1 software and you have verified proper switch operation. Because VLAN Manager operates across switches, you may be managing routing switches with different software versions, which will require different versions of the management software.

If you upgrade all switches to the version 1.1 software image, you can uninstall the version 1.0 management software once you have verified that the new version is working correctly. However, Bay Networks recommends that you uninstall by deleting the version 1.0 directory, and NOT by running the Uninstall program, to avoid deleting common files.

Installation on Windows NT and Windows 95 Platforms

The minimum system requirements for installing Accelar Management Software in a Microsoft® Windows® environment are:

- 75 MHz Pentium or 100 MHz 486
- 16 MB DRAM
- 8 MB space on hard drive

Caution: By default, Windows will install newer software over an older version. Therefore, to ensure that version 1.0 of the management software remains on your system until you can verify that all switches are running version 1.1 runtime and boot images, be sure that you install version 1.1 of the management software in a different directory from the 1.0 software.

To install Device Manager on a Windows platform:

1. **Click the Windows Start button and select Run.**

   The Run dialog box opens.

2. **In the Open field, type** `E:\DM-Windows\setup.exe` **and click on OK.**

3. **Follow the instructions appearing on the screen to complete the installation.**

4. **To run the Accelar Management Software, click on the Windows Start button and click on Programs. Select Accelar Device Manager from the Programs menu.**

Note: When all switches are verified to be correctly running version 1.1 runtime and boot images and management software, you can uninstall the version 1.0 management software by deleting the version 1.0 directory. DO NOT run the Windows Uninstall program.

Installation on UNIX Systems

The minimum system requirements for installing Accelar Management Software in a UNIX environment are:

- SPARC workstation running the Solaris 2.5.x Operating System
- HP workstation running the HP/UX 10.20 Operating System
- 32 MB DRAM
- 12 MB space on hard disk, 4 MB available in a temporary directory and 8 MB free in the directory where you want to install the Accelar Management Software.

The Accelar Management Software is available on the Bay Networks support World Wide Web site at:

http://support.baynetworks.com/software/routingswitches/accelar_1000.html

If you are downloading the software from the Bay Networks Web site, you will need the following two files:

- install_accelar_dmvm
- dm_1.1.0.tar.Z

These files can be found under the "Script for Installing DM Solaris & HP/UX using .Z file (v1.1.0)" and "Device Manager Solaris & HP/UX v1.1.0 (compressed)" links, respectively.

Note: If you currently have version 1.0 of Accelar Management Software on your system, by default the 1.1 version will install to a different subdirectory, leaving version 1.0 intact.

Preparing for Installation

To prepare to install Accelar Management Software on Solaris and HP/UX platforms:

1. **Place the two files `dm_1.1.0.tar.Z` and `install_accelar_dmvm` into a temporary directory on your Solaris or HP/UX workstation.**
2. **Log in as `root` or `su` to the root account.**
3. **Change directory (cd) into the directory where you placed the two files.**
4. **Execute the following command:**

   ```
   # chmod 755 ./install_accelar_dmvm
   ```

5. **Proceed to the "Running the Install Script" instructions.**

Running the Install Script

If installing from the CD-ROM, these instructions assume that you have mounted the CD-ROM drive onto the `/cdrom` directory. For instructions on mounting the CD-ROM drive, refer to the user guide for your operating system.

To run the Install script:

1. **Create a directory on the host's hard disk where you want the Accelar DM/VM software to reside.**

   The location can be anywhere of your choice. Typically, third-party software is installed under the `/opt` directory. Depending on how your system is configured, you may need to be the root user to have write permission to the `/opt` directory. If you do not have root access, you can install the software under `/usr/local/Accelar`.

   The directory should be similar to the following:

   ```
   % mkdir /opt/Accelar
   ```

2. **If installing from the CD-ROM, change directory into the `dm-unix` directory on the CD-ROM.**
   If you obtained the compressed software file from the Web, change directories into the directory where you placed the file. For example:

   ```
   % cd /cdrom/dm-unix
   ```

   or `% cd /var/tmp/accelar`

3. **Execute the `install_accelar_dmvm` script by providing the version of the software you want to install and the target directory where you decide to install the Accelar DM/VM software.**

   ```
   install_accelar_dmvm <filename> <target_directory>
   ```

   where:

   `filename` is the name of Device Manager file, omitting the `.tar.Z` extension.

   `target_directory` is the name of the directory where DM is to be installed (the directory created in step 1).

   For example:

   ```
   % install_accelar_dmvm dm_1.1.0  /opt/Accelar
   ```

   `dm_1.1.0` is the name of the dm tar file, (that is, `dm_1.1.0.tar.Z`)

   `/opt/Accelar` is the target directory where dm_1.1.0 will be installed.

4. **After installation, each Accelar software user's environment must be set correctly. Assuming that you installed the software into the directory** `/opt/Accelar/dm_1.1.0`**, set the DMPATH variable using one of the following instructions:**

   a. **If you use the C-shell, add the following to your .cshrc:**

   ```
   % setenv DMPATH /opt/Accelar/dm_1.1.0
   ```

   b. **If you use the Bourne shell or Korn shell, add the following to your .profile:**

   ```
   $ DMPATH=/opt/Accelar/dm_1.1.0; export DMPATH
   ```

Note: If you installed into a different directory, substitute the directory where you installed the software for `/opt/Accelar/dm_1.1.0`. After modifying the appropriate file, you must source the file for the changes to take effect.

Manual Installation Procedure

If you have difficulty installing the Accelar Management Software using the installation script provided, you can use the following procedure to manually install the software:

1. **Insert the Accelar Series 1000 Software 1.0 (Windows 95/NT, Solaris, and HP-UX) CD into the CD-ROM drive.**
2. **Log in as the root user or su to root and mount the CD onto the /cdrom directory. If you are unfamiliar with this procedure, refer to the user guide for your operating system.**

   a. **For Solaris systems, change directory to** `/cdrom/cdrom0/dm-unix`**.**

   b. **For HP-UX systems, change directory to** `/cdrom/dm-unix`**.**

3. **Copy the software into /tmp with the following command:**

   ```
   # cp dm_1.1.0.tar.Z  /tmp/dm_1.1.0.tar.Z
   ```

4. **Change directory into the directory where you want the Accelar Management Software to be installed. The software will automatically be loaded into a subdirectory called** `dm_1.1.0`**. For example, if you want to install the software into the** `/opt/dm_1.1.0` **directory, change directory to** `/opt`**.**
5. **Uncompress and untar the software with the following command:**

   ```
   # zcat /tmp/dm_1.1.0.tar.Z | tar xvf -
   ```

6. **Once installed, change directory into the dm_1.1.0 directory and execute:**

```
# ./standalone_setup
```

7. **After installation, each Accelar software user's environment must be set correctly. Assuming that you installed the software into the `/opt/dm_1.1.0` directory, set the DMPATH variable using one of the following instructions:**

   a. **If you use the C-shell, add the following to your .cshrc:**

```
% setenv DMPATH /opt/dm_1.1.0
```

   b. **If you use the Bourne shell or Korn shell, add the following to your .profile:**

```
$ DMPATH=/opt/dm_1.1.0; export DMPATH
```

Note: If you installed into a different directory, substitute the directory where you installed the software for /opt/dm_1.1.0. After modifying the appropriate file, you must source the file for the changes to take effect.

Installing DM/VM into HP OpenView (Optional)

Prerequisites for installing DM/VM into HP OpenView are:

- HP OpenView must be installed.
- Accelar Management Software must be installed.
- You must be logged in as root.

To integrate Accelar Management Software into HP OpenView:

1. **Go to the directory where the Accelar Management Software is installed.**

```
# cd /opt/dm_1.1.0
```

2. **As the root user, start Accelar Device Manager and point it to an Accelar device (for example, `dm <a.b.c.d>` where `a.b.c.d` is the IP address of the Accelar device or the name of the device if your network is running a name service).**

```
# dm 132.32.1.5
```

Note: You MUST launch `dm` against an Accelar device to activate the appropriate menus for step 3.

3. **From Accelar Device Manager, click on Help from the menu bar and select the Run HP OpenView Install on Console option.**

   Selecting this option will copy the icons specific to Accelar and MIB files to the standard HP OpenView software directories. When this task is completed, a list of commands will be shown on the console. Manually execute the commands as root to complete the installation process.

   Table 1 gives a brief description of each command and its function.

Table 1. Commands for Installing DM/VM into HP OpenView

Command	Description
ovw -fields	Update the fields database.
ovstop netmon	Stop network monitor to allow the database to update.
ovtopofix	Update topology database.
ovstart netmon	Restart network monitor.
ovw	Start HP OpenView.

4. **After HP OpenView is started, load the MIBs specific to Accelar by selecting the Options->Load MIBs option.**

   Load the `rapid_city.mib` file located in the following directory:

   ```
   /var/opt/OV/share/snmp_mibs/Vendor
   ```

Note: When all switches are verified to be correctly running version 1.1 runtime and boot images and management software, you can uninstall the version 1.0 management software by deleting the 1.0 directory. DO NOT run the UNIX Uninstall program.

New Features in Release 1.1

Release 1.1 of the Accelar Management Software includes the following new features:

- Console and Telnet login enhancements (starting on page 28)
- CLI enhancements (page 30)
- Telnet client and server enhancements (page 31)
- BootP/DHCP relay (page 32)
- Proxy ARP (page 38)
- Default VLAN for nontagged frames on trunk ports and tagged frames on access ports (page 40)
- Policy VLAN for VINES IP (page 41)

Console and Telnet Login Enhancements

The default values for login name and password for the console and Telnet sessions have been changed to the values in Table 2.

Table 2. Login and Password Default Values

Access Level	Default Login	Default Password
read/write/all	rwa	rwa
read/write	rw	rw
read only	ro	ro

Note: If you entered a Web password under version 1.0, this password becomes the default password for Web, console, and Telnet under version 1.1.

The command line interface (CLI) has three login/password pairs, one for each access level:

- rwa: read/write/all
- rw: read/write
- ro: read only

The CLI passwords are now independent of the SNMP community strings and can be set or changed from the CLI prompt. Login and password settings can be changed only if you log in with read/write/all privileges (that is, the rwa access level). You can also change the login names for different modes.

The console, Telnet, and Web access all use the same login names and passwords. When you change a password, it is changed for all three interfaces. Web access is enabled by default. When the CLI prompts for login and password, the access level is set corresponding to the login and password pair entered.

The “login” command allows you to log in again with a different login access by entering the user name and password. The prompt remains at the same level that you were before logging in again, but privileges are determined by the current access level.

The “logout” command allows the user to log out and reenter at the top level prompt. If you connect to the routing switch through Telnet, logout terminates the Telnet session.

CLI Commands to Change Console/Telnet Password

The following commands can be used to change the login name and the password for each different login access level:

```
cli password ro <username> [<password>]
cli password rw <username> [<password>]
cli password rwa <username> [<password>]
```

To display information about the access levels for login names and passwords, use this command:

```
cli password show
```

Example

```
*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.1.0      *
*******************************
```

```
Login: rwa
Password: ***

Accelar-1100# cli password
Accelar-1100/cli/password# ?
cli password ro <username> [<password>]
cli password rw <username> [<password>]
cli password rwa <username> [<password>]
cli password show
Accelar-1100/cli/password# rwa test bay
Accelar-1100/cli/password# show
Access    Login            Password
rwa       test             bay
rw        rw               rw
ro        ro               ro
Accelar-1100/cli/password#
```

CLI Enhancements

The CLI is organized into a tree data structure. Help can be accessed from any level of the tree by typing a question mark (?) or the word `help`. Once you are within a given branch of the tree, you only need to type the subcommand for that level. For example, to get information about a port from the top level, type: `port show info`. When you are already in the "port" branch, you need only type: `show info`.

To avoid having to type complete commands, you can type part of a command and then press the Tab key to complete the command. If the letters you typed are unique to a command, the command will be automatically completed. If not, a bell will sound to indicate that more information is necessary.

Throughout the CLI, the following keystrokes are available:

- Control-P: to view and scroll through the previous history commands
- Control-N: to view and scroll through the next history commands
- Control-U: to delete a line character
- Control-C: to abort a line character
- Control-D: to log off the system
- Control-I (or the Tab key): for command completion
- Control-H: to backspace

In addition, certain commands are used for navigation through the CLI:

- `back` or **..** to take you back up one level
- `box` or `top` to take you to the box or top level
- `pwc` to display the current working level

Parameter values in the CLI are indicated by carets < >. Parameters can be optional or required. Required parameters must be in the specified order, followed by optional parameters. Optional parameters are displayed in brackets [ ].

When entering multiple CLI commands, you can terminate a command within a single line of input by using the semicolon (;) as the separator. A semicolon is treated like a carriage return by the CLI.

For more information about CLI commands, refer to Chapter 13, "Command Line Interface," in the *Reference for the Accelar Management Software*.

Telnet Client and Server Enhancements

Accelar switches support eight incoming Telnet sessions, in addition to the console serial port. The modem port on the Accelar 1200/1250 also has CLI access, allowing a total of 10 incoming sessions on the Accelar 1200/1250. One outbound session is also supported.

By default, eight incoming Telnet sessions are allowed. To change the number of allowed incoming Telnet sessions, use the command:

```
cli telnet-sessions <nsessions>
```

For example:

```
Accelar-1100# cli telnet-sessions 1
Accelar-1100# cli show

cli configuration

more             : true
screen-lines     : 23
telnet-sessions  : 1
timeout          : 900 seconds
```

BootP/DHCP Relay

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), is used to dynamically provide host configuration information to the workstations. One or a few centrally located DHCP servers are preferred by network managers because of the lower administrative overhead. Using few DHCP servers requires the routers connecting to the subnets or VLANs/bridge domains to support BootP/DHCP relay so that hosts can get the configuration information from servers several router hops away.

Differences Between DHCP and BootP as Specified in RFC 2131

The following differences are specified in RFC 2131, which BootP did not address:

- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters needed to operate.

DHCP uses the BootP message format defined in RFC 951. A packet is classified as DHCP if the first four octets in the options field are 99, 130, 83, 99 and the fifth octet is 53. The first four octets are referred to as the “Magic Cookie,” and the fifth is the DHCP message type code. The remainder of the options field consists of a list of tagged parameters that are called “options” (RFC 2131).

Summary of DHCP Relay Operation

BootP/DHCP clients (workstations) generally use UDP/IP broadcasts to determine their IP addresses and configuration information. If such a host is on a network or a subnet segment (or VLAN) that does not include a DHCP server, the UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. The Accelar routing switches can be configured to overcome this issue by forwarding the broadcasts to the server through isolated or virtual IP router interfaces. The IP router interfaces can be configured to forward DHCP broadcasts to other locally connected network segments or directly to the server’s IP address.

In Figure 2, an end station is connected to subnet 1, corresponding to VLAN 1. Assume that the Accelar routing switch connects two subnets via the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255), the Accelar routing switch forwards DHCP requests to subnet 2 or to the host address of the DHCP server.

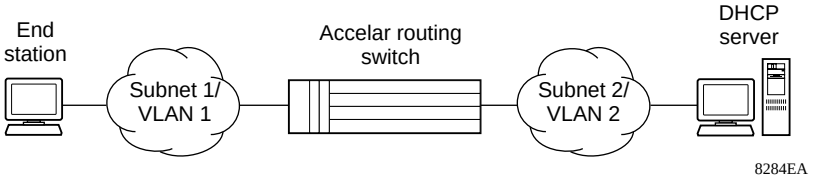

Figure 2. Example of DHCP Operation

Related RFCs

The following RFCs refer to DHCP operation:

- RFC 1542
- RFC 951
- RFC 2131

Command Line Interface

This section describes the CLI commands used for the BootP/DHCP features. For an explanation of how to use BootP/DHCP features with Device Manager, refer to "BootP/DHCP Relay" in Chapter 9 of *Reference for the Accelar Management Software*.

The relaying function is a two- or three-step process:

1. Enabling the BootP/DHCP relay on an interface
2. Changing the default parameters (optional)
3. Setting the forwarding path

You must specify the interfaces on which the received BootP/DHCP packets should be processed for relaying, and then specify the forwarding path.

Enable BootP/DHCP Relay on an Interface

The following commands are used to enable or disable the processing of BootP/DHCP broadcast messages on an IP interface (where an IP interface can be either an isolated or a VLAN router port):

```
dhcp enable <ipaddr>

   Required parameters:

   <ipaddr>         = ip address {a.b.c.d}

   default max hop value is 4
   default min sec value is 0

   default mode is both DHCP and BootP, i.e., both BootP and DHCP
   requests are accepted on this interface.
```

This command enables the processing of BootP/DHCP broadcast packets received on the IP interface specified as `<ipaddr>` for the relaying function.

Note: The `<ipaddr>` parameter specifies the IP address of the input interface on which the relaying of received BootP/DHCP packets must be enabled.

```
dhcp disable <ipaddr>
```

This command disables the relaying of BootP/DHCP broadcast packets received on the IP interface specified as `<ipaddr>`.

Changing Default Parameters

The BootP/DHCP parameters have the following default values:

- max-hop value
  Before the routing switch relays or forwards the BootP/DHCP packets to the specified interface or IP address, it increments the “hops” field in the BootP/DHCP packets and discards the packets if the hops field is greater than the “max hop” value. The default max hop value is 4, and the maximum value is 16. The max-hop value can be changed by using the `dhcp max-hop` command.

- min-sec value
  According to RFC 1542, “The ‘secs’ value of a BootPREQUEST message SHOULD represent the elapsed time, in seconds, since the client sent its first BootREQUEST message. Note that this implies that the ‘secs’ field of the first BootREQUEST message SHOULD be set to zero.” The Accelar routing switch relays or forwards a BootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped. The default min-sec value is 0; it can be changed by using the `dhcp min-sec` command.
- dhcp mode
  The `dhcp mode` command is used to specify if only DHCP, only BootP, or both types of packets received on the specified interface should be processed for the relay function. This parameter can be changed only if the BootP/DHCP function has been enabled on an interface. The default is to process both BootP and DHCP packets for relaying.

The following three commands are used to change the default processing for packets received on BootP/DHCP enabled interfaces:

```
dhcp max-hop <ipaddr> <max-hop>
   Required parameters:
   <ipaddr>          = ip address {a.b.c.d}
      <max-hop>        = max hop count {1.16}
      default max-hop value is 4

dhcp min-sec <ipaddr> <min-sec>
   Required parameters:
   <ipaddr>          = ip address {a.b.c.d}
   <min-sec>         = min second count {0.65535}
   default min-sec value is 0

dhcp mode <ipaddr> <mode>
   Required parameters:
   <ipaddr>          = ip address {a.b.c.d}
   <mode>            = mode value {bootp|dhcp|bootp_dhcp}
   default mode is both DHCP and BootP, i.e., both BootP and DHCP
   requests are accepted on this interface.
```

Setting the Forwarding Path

To set up a forwarding path for BootP/DHCP packets received on an interface enabled for relaying, use the following command:

```
dhcp from <ipaddr> to <ipaddr> <op>
    (default mode for the path is both DHCP and BootP)
    Required parameters:
    <ipaddr>          = ip address {a.b.c.d}
    to                = keyword
    <ipaddr>          = ip address {a.b.c.d}
    <op>              = operation {disable|enable|delete}
```

The `from <ipaddr>` parameter is one of the addresses of an IP interface of the switch. DHCP must be enabled on that interface through the `dhcp enable` command before adding the forwarding path. The to <*ipaddr*> parameter is either the IP address of the BootP/DHCP server or the address of another local interface of the switch. If the `to <ipaddr>` parameter is the address of the BootP/DHCP server, then the request is unicast to the server's address. If the `to <ipaddr>` parameter is one of the IP addresses of an interface on the switch, the BootP/DHCP requests will be broadcast out of that local interface.

To modify the default mode for the forwarding path, use the following command (this command would most likely be used when the BootP and DHCP packets are to be forwarded to independent locations or servers):

```
dhcp relay <ipaddr> to <ipaddr> <mode>
   Required parameters:
   <ipaddr>          = ip address {a.b.c.d}
   to                = keyword
   <ipaddr>          = ip address {a.b.c.d}
   <mode>            = mode value {bootp|dhcp|bootp_dhcp}
```

Show Commands

The following show commands are used with DHCP:

`dhcp show interface`

This command displays the ingress interfaces on which the processing of BootP/DHCP broadcast has been enabled.

`dhcp show route`

This command displays the configured forwarding information for BootP and DHCP.

```
dhcp show counters
```

This command displays the number of BootP/DHCP requests and replies that have been received by all interfaces enabled by BootP/DHCP.

The following example displays the show commands:

```
Accelar-1100# dhcp show interface
>       Interface   Enable   MaxHop   MinSec          Mode
>   134.177.16.23     TRUE        4        0  DHCP & BOOTP

Accelar-1200# dhcp show route
>       Interface          Server   Enable           Mode
>     10.10.20.59  192.168.20.161     TRUE    DHCP & BOOTP
>   192.168.20.43  192.168.20.161     TRUE    DHCP & BOOTP

> Accelar-1200# dhcp show counters
>       Interface  Requests  Replies
>   192.168.20.43         0        0
>     10.10.20.59         0        0
```

Example

In the example in Figure 3, to configure the Accelar routing switch to forward DHCP packets from the end station to the server, issue these CLI commands:

```
dhcp enable 10.10.1.2
dhcp from 10.10.1.2 to 10.10.2.1 enable
```

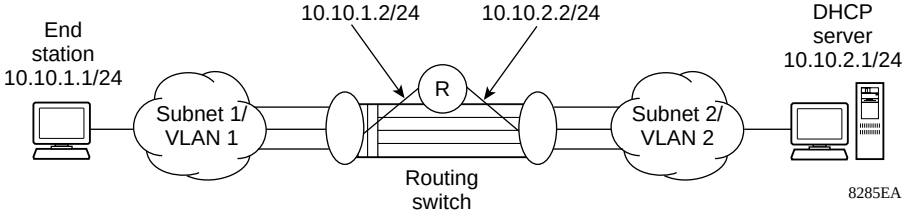

Figure 3. Forwarding DHCP Packets

All BootP broadcast packets, including DHCP packets, which appear on the VLAN 1 router interface (10.10.1.2), will be forwarded to the DHCP server. In this case, the DHCP packets will be forwarded as unicast to the DHCP server's IP address.

To forward BootP/DHCP packets as broadcast packets to VLAN 2, specify the IP address of the switch VLAN2 router interface (10.10.2.2.) as the “to ip address” as follows:

```
dhcp enable 10.10.1.2
dhcp from 10.10.1.1 to 10.10.2.2 enable
```

Proxy ARP

Proxy Address Resolution Protocol (Proxy ARP) allows the Accelar routing switches to respond to an ARP request from a locally attached host or end station for a remote destination. It does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the Accelar routing switch has an active route to the destination network.

Figure 4 is an example of Proxy ARP operation. In the figure, Host B could send an ARP request for Host C. The Accelar routing switch would respond to the ARP request with Host C’s IP address but with its own MAC address.

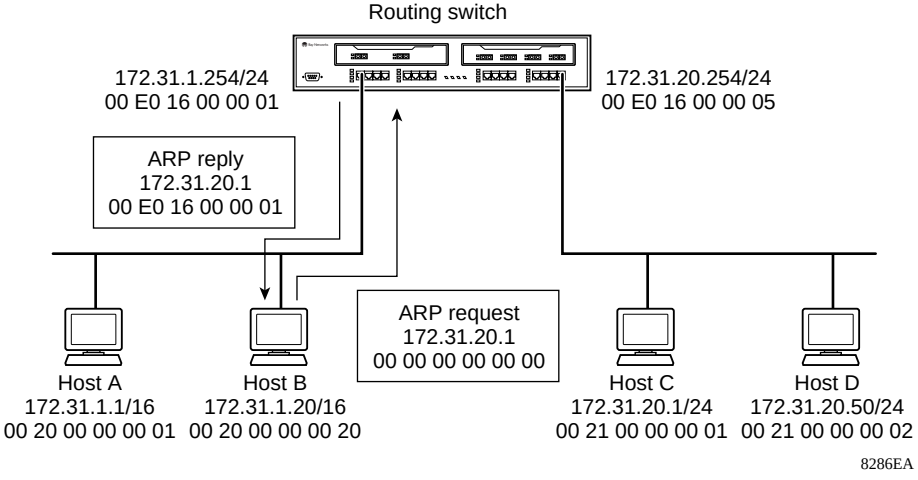

Figure 4. Proxy ARP Operation

Related RFCs

The following RFCs refer to Proxy ARP operation:

- RFC 826
- RFC 925
- RFC 950
- RFC 1027

Command Line Interface

This section describes the CLI commands used for Proxy ARP operation. For an explanation of how to use Proxy ARP with Device Manager, refer to "Using Proxy ARP" in Chapter 9 of *Reference for the Accelar Management Software*.

Enable or Disable Proxy ARP on Selected Interfaces

Use the following commands to enable or disable Proxy ARP on an interface:

```
arp proxy enable    [<ipaddr>]
arp proxy disable   [<ipaddr>]
```

The `<ipaddr>` parameter specifies the IP address of the interface on which Proxy ARP must be enabled or disabled. This parameter is optional.

Enable or Disable Proxy ARP on all IP Interfaces

The `<ipaddr>` parameter is optional. If no IP address is specified, Proxy ARP is enabled or disabled on all configured IP interfaces on the switch. An IP interface refers to the VLAN router interface, as well as the isolated router port.

Show Commands

Use the following show command with Proxy ARP:

```
arp proxy show [<ipaddr>]
```

This command displays the status of the Proxy ARP for all configured IP interfaces on the routing switch, for example:

```
Accelar-1100#    arp proxy show
       Interface           Enable
       10.10.1.1           TRUE
       159.23.207.100      FALSE
```

To display the status of Proxy ARP on a specific IP interface, specify the IP address of the interface as a parameter for the show command:

```
Accelar-1100#    arp proxy show 10.10.1.1
       Interface           Enable
       10.10.1.1           TRUE
```

Default VLAN for Untagged Frames on Trunk Ports and Tagged Frames on Access Ports

In release 1.1, the user has the option to configure the switch to discard or to accept tagged frames received on an access port and untagged frames (except for BPDUs) received on a trunk port. The default behavior is to forward these frames to the default VLAN. You can also configure a port-based VLAN to which the frames will be forwarded.

In release 1.0, tagged frames received on an access port and untagged frames (except for BPDUs) received on a trunk port were discarded by the switch. In some situations, an access port may be temporarily connected to a trunk port. In these scenarios, the routing switch should not discard the tagged frames received on the access port or the untagged frames received on the trunk port because doing so could cause a temporary loss of connectivity from the management station.

Exceptions to the described behavior where the tagged frames received on the access ports will get dropped include:

- Tagged frames with unknown VLAN ID
- VLAN ID known but not assigned to the port on which the tagged frame is received

Command Line Interface

This section explains the CLI commands used to configure the switch to accept untagged frames on trunk ports and tagged frames on access ports. For an explanation of configuring with Device Manager, refer to “Accepting Tagged Frames on Access Ports and Untagged Frames on Trunk Ports” in Chapter 8 of *Reference for the Accelar Management Software*. You can also set the VLAN to which the frames will be forwarded (the DefaultVlanId) using the same window. The default value of this variable is 1.

Note: The VLAN specified by DefaultVlanId must be a port-based VLAN in the Default Spanning Tree Group (STG).

To direct the switch to discard tagged frames on an access port (default is to NOT discard), use the command:

```
port set disctagged <ports> <true|false>
```

To direct the switch to discard untagged frames on a trunk port (default is to NOT discard), use the command:

```
port set discuntagged <ports> <true|false>
```

To direct the switch to send these frames to a particular VLAN, use the command:

```
port set defvid <ports> <vid>
```

where `<vid>` must be a port-based VLAN in the Default STG.

Policy VLAN for VINES IP

The definition of protocol-sensitive VLANs has been enhanced to include Banyan VINES Internetwork Protocol (VINES IP). The VLAN policy for VINES IP is applied in the hardware, similar to that for other protocol type policies in the Accelar routing switch hardware. The software checks for an Ethertype value of hexadecimal 0BAD.

Known Software Limitations

This section describes the latest functionality and known limitations for the Accelar 1000 Series hardware and system software.

Gigabit Ethernet Ports

For information about the limitations of the Gigabit Ethernet ports, refer to the *Release Notes for Accelar 1000 Series Gigabit Ethernet Modules* (Bay Networks part number 201539-B), shipped with the Gigabit Ethernet modules.

Policy-based VLANs

The following policy-based virtual LAN (VLAN) configurations are not supported in this release:

- IP subnet-based VLANs
- Protocol-based VLANs for SNA and NetBIOS

These configurations will be supported in a future release.

10/100BASE-TX Autonegotiation

The 10/100BASE-TX ports may not autonegotiate correctly with older 10/100BASE-TX equipment. In some cases, the older devices can be upgraded with new firmware or driver revisions. If an upgrade does not allow autonegotiation to correctly identify the link speed and duplex settings, these settings can be manually configured for the particular link in question using Device Manager or CLI. See Chapter 6, "Port Configuration and Graphing," in the *Reference for the Accelar Management Software* manual. Check the Bay Networks Web site (baynetworks.com) for the latest compatibility information.

Routing

The following paragraphs describe routing limitations.

Broadcast Format

Accelar 1000 Series routing switches support IP broadcast addresses in the "all ones" format. This format includes both the directed and local IP broadcast addresses. IP broadcast addresses with "all zeroes" are not supported.

RIP Route Redistribution

New Routing Information Protocol (RIP) global parameters allow the user to configure which routes are advertised by the routing switch through RIP.

Note: RIP route redistribution parameters are available only in Device Manager under Routing->RIP and are not documented in *Reference for the Accelar Management Software.* The runtime CLI does not allow the configuration of RIP route redistribution.

RIP route redistribution parameters are global RIP parameters. The RIP configuration for each router interface will behave in accordance with the global RIP parameters.

There are three types of RIP route redistribution:

- Redistribute OSPF to RIP—Routes learned through OSPF can be advertised through RIP.
- Redistribute Direct to RIP—Any locally connected routes are advertised through RIP. This is the normal operation mode for RIP route distribution.
- Redistribute Static to RIP—Any static routes defined will be advertised through RIP. Advertising static routes can easily cause routing problems in a network. For this reason, you must exercise great care when redistributing static routes.

Note: By default when RIP is globally enabled, no route redistribution is enabled even for direct routes. In this mode, the router will learn about new RIP routes and will advertise RIP learned routes, but it will not advertise any direct or OSPF routes. To advertise direct routes, “Redistribute Direct to RIP” must be enabled and the RIP interface configuration needs to have RIP enabled and “talk” enabled.

Management Ports

The following information applies to the Accelar management ports.

PCMCIA Card Compatibility

For Accelar 1200 and 1250 routing switches, the XLR1299PC PCMCIA memory card is the only card approved for use in the XLR1297SF slot. Other PCMCIA flash memory cards may not be compatible with the XLR1297SF module and should not be used.

A PCMCIA flash memory card compatible with the XLR1299PC memory card requirements is the Intel Value Series 100, 4 MB, 5-volt memory card (Intel part number iMC004FLSC-10).

Remote Console Management

For remote management, you can connect a modem to the console port to gain remote access to the boot monitor or runtime CLIs. This section describes issues relating to configuring the modem for this mode of operation.

The console serial port is a data terminal equipment (DTE) device operating at 9600 bps, 8 data bits, no parity, and one stop bit. Because the console port expects to receive Data Set Ready (DSR) and Clear To Send (CTS) signals before transmitting, these control lines are required in the cabling. The console port does not support any inbound flow control; that is, the port does not toggle control lines to indicate an input buffer full condition.

The following modem settings should be saved to modem configuration as the power-on default:

- Auto Answer enabled
- No echo
- Quiet Mode (no return codes)

The following other relevant settings are generally modem specific:

- Communication parameters fixed at 9600 bps, 8 data bits, no parity, and one stop bit
- Ignore Data Terminal Ready (DTR)
  If the modem is set to ignore DTR, it will not hang up if DTR is lost, and it will auto-answer regardless of the DTR state. Ignore DTR is the desired setting when there is a single silicon switch fabric (SSF) module and a single modem.
  If you have redundant SSF modules and modems cabled in parallel for remote management (as described in *Using the Accelar 1200/1250 Routing Switch*), the modems should NOT be configured to Ignore DTR. This precaution will ensure that the modem on the active SSF module answers when dialed.

If the modem does not properly detect the console interface settings or if flow control problems occur, you should reconfigure the modem interface characteristics. A sample modem setup sequence with comments is as follows:

1. AT&F1<cr> (restore factory defaults)
2. ATQ1E0S0=1<cr> (quiet mode, echo off, auto answer on one ring)
3. AT&W0<cr> (write to profile 0, the normal power-on default profile)

Modem Port

The modem port on the Accelar 1200 and 1250 SSF module is not supported. However, this port does have CLI access. In a future release, this port will be used for remote Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) management.

Traffic Prioritization

An Accelar 1000 switch can operate in either of two modes: Best Effort mode or Priority mode. The factory default setting is Best Effort mode.

Note: Best Effort and Priority modes and the associated command line interface (CLI) commands are not documented in the manuals. Configuring high-priority traffic is described in Chapter 14, "Prioritization," of the *Reference for the Accelar Management Software.*

The following differences exist between the Best Effort and Priority modes:

- In Best Effort mode, all traffic is treated with the same priority.
- In Priority mode, high-priority traffic flows through the switch fabric using a high-priority data path; output buffers are reserved for high-priority traffic.

You can change the operating mode of the switch from the runtime CLI using the `sys set flags` command. Note that after changing this setting, you MUST save the configuration and reboot the switch before the change takes effect.

To enable Priority mode:

1. **Enter the following commands at the CLI:**

```
Accelar-1100#> sys set flags highpriomode true
Accelar-1100#> save
```

2. **Reboot the routing switch to enable the configuration change to take effect.**

High-priority traffic can be enabled on a per-port, per-MAC address, per-VLAN, or per-flow basis.

- When a port is set to High Priority mode, all traffic received on this port is assigned a high priority.
- When a MAC address is set to High Priority mode, all traffic from the MAC address is assigned a high priority.
- When a VLAN is set to High Priority mode, frames received on any of the active ports of the VLAN are assigned a high priority.
- Finally, an IP Flow record can be used to assign high switching priority to an IP packet based on its source and destination IP addresses, protocol type, source port number, and destination port number.

When a high-priority frame is sent out a trunk port, the 3-bit User Priority field in the IEEE VLAN Tag is set to 7. A normal priority frame has a User Priority of 0.

Any received tagged frames with a User Priority greater than 0 are treated as high priority.

Network Management Functionality and Limitations

This section describes the latest functionality and known limitations of managing an Accelar 1000 Series Release 1.1 routing switch. Unless otherwise noted, the information is valid across all management platforms.

Multisegment Autotopology

The Bay Networks Multisegment Autotopology™ protocol is not supported; therefore, an Optivity® management console will not be able to determine the physical connectivity for an Accelar routing switch. Multisegment Autotopology will be supported in a future release.

Port Statistics Support

The following tables indicate which counters are supported in the Accelar 1000 Series. Table 3 contains support information for the port interface counters; Table 4 contains support information for the port RMON counters.

An “X” in a column indicates that the counter is supported and works correctly. “Not Available” indicates that the counter is not available, and “Not Applicable” means that the counter is not applicable to the specific interface.

Note: Even when a counter is supported by the Accelar routing switch, the counters available on a given network management platform will depend on the capabilities of that platform.

Table 3. Accelar 1000 Series Interface Counter Support

Counter	10BASE-T	100BASE-TX/FX	1000BASE-SX/LX
InOctets	Not Available	Not Available	X
InPackets	X	X	X
InUnicastPkts	X	X	X
InNUnicastPkts	X	X	X
InMulticast	X	X	X
InBroadcastPkts	X	X	X
InDiscards	X	X	X
InErrors	X	X	X
OutOctets	X	X	X
OutPackets	X	X	X
OutUnicastPkts	X	X	X
OutNUnicastPkts	X	X	X
OutMulticast	Not Available	Not Available	X
OutBroadcastPkts	Not Available	Not Available	X
OutDiscards	Not Available	Not Available	Not Available
OutErrors	Not Applicable	Not Applicable	Not Applicable

Table 4. Accelar 1000 Series RMON Counter Support

Counter	10BASE-T	100BASE-TX/FX	1000BASE-SX/LX
etherStatsDropEvents	Not Applicable	Not Applicable	Not Applicable
etherStatsOctets	Not Available	Not Available	X
etherStatsPkts	X	X	X
etherStatsBroadcastPkts	X	X	X
etherStatsMulticastPkts	X	X	X
etherStatsCRCAlignErrors	X	X	X
etherStatsUndersizePkts	X	X	X
etherStatsOversizePkts	X	X	X
etherStatsFragments	X	X	X
etherStatsJabbers	Not Available	Not Available	Not Available
etherStatsCollisions	X	X	Not Applicable
etherStatsPkts64Octets	Not Available	Not Available	X
etherStatsPkts65to127Octets	Not Available	Not Available	X
etherStatsPkts128to255Octets	Not Available	Not Available	X
etherStatsPkts256to511Octets	Not Available	Not Available	X
etherStatsPkts512to1023Octets	Not Available	Not Available	X
etherStatsPkts1024to1518Octets	Not Available	Not Available	X

Counters that are not available or not applicable will always show a 0 (zero) value when viewed with network management software.

SNMP Traps

In the version 1.1 release, the Accelar 1000 Series routing switches support the following SNMP traps:

- MIB2 traps (RFC1213)
- OSPF traps (RFC1850)
- RMON alarm traps (RFC1271)
- Enterprise traps (summarized in Table 5)

Table 5. Accelar 1000 Enterprise Traps

Enterprise Trap	Description
rcCardDown	Card is down.
rcCardUp	Card is up.
rcErrorTrap	An error has occurred with error code.
rcStpNewRoot	New spanning tree root bridge exists.
rcStpTopologyChange	Spanning Tree Protocol topology is changed.
rcChasPowerSupplyDown	Power supply is down.
rcChasFanDown	Fan is down.

Device Manager Limitations

This section describes the latest functionality and known limitations of Accelar Device Manager (DM) version 1.1.

RMON Counter Support

Device Manager does not support the packet size distribution RMON counters. For a list of the RMON counters supported by the Accelar chassis in the MIB, refer to Table 4 on page 48. This limitation applies to all platforms.

Solaris and HP/UX Platforms

The following functionality applies only to Solaris and HP/UX platforms.

Context-Sensitive Online Help

Under UNIX, Device Manager (DM) displays the online Help screens using the Netscape Web browser. DM assumes that Netscape is in the current directory or in the path. If DM cannot find Netscape, it will return a message indicating that it could not find or execute Netscape when online Help is accessed.

Receiving Traps

To receive SNMP traps when running Device Manager, you must execute with root user privileges. If you do not run with root privileges, Device Manager will report a “Can't open trap port, Permission denied” error on startup, which indicates that you do not have sufficient privileges to receive traps.

Use with HP OpenView

When using Device Manager with HP OpenView (HPOV), note the following:

- HPOV 4.x can only relay SNMPV1 traps. You must ensure that trap v1 format is configured in Edit Chassis->TrapReceiver for any HPOV v4.x trap receivers.
- When launched from the command line, Device Manager will default to the community strings in dm.ini (public, private). If you launch DM within HPOV, it uses the community strings HPOV has configured for that device.

Manually Resizing Windows

If Device Manager subwindows are manually resized, DM will no longer automatically size the resized window. The subwindow will automatically size if it is closed and reopened.

Windows 95 and NT

The following information applies only to Windows 95 and Windows NT® platforms.

Context-Sensitive Online Help

Device Manager displays the online Help screens using the default Web browser. Using Netscape Navigator, online Help is context sensitive in that it brings up the correct part of the Help HTML file. If Microsoft Internet Explorer® version 4.0 or earlier is the default browser, online Help only takes you to the top of the HTML file.

SNMP Trap Support

Device Manager under Windows only supports SNMP v2c traps, which is the default trap type. SNMP v1 traps sent to DM are not displayed in the Trap Log.

Low Memory Errors

When Device Manager runs low on memory, you will get a “WINSNMP error #99” (Internal error) message. To work around this problem, either reduce the number of running processes or increase the Windows swap space.

Abnormal Termination Recovery

When started, Device Manager automatically launches the NetPlus/32.dll. If DM terminates abnormally, the NetPlus/32 task may still be running. You should terminate the NetPlus/32 task before restarting DM.

Intermittent "bitmap 'gray50' not defined" Error

An intermittent "bitmap 'gray50' not defined" error can occur when opening a new device in one Device Manager session. Closing and restarting the DM session will correct this behavior.

Runtime Error Changing VLAN Colors

Intermittent "Runtime Error!" messages in wish42.exe can occur when you attempt to change the VLAN color after graphing data in Device Manager. The runtime error does not cause any corruption of data. To recover from the abnormal termination, terminate the NetPlus/32 task before restarting DM.

Bugs Fixed in Release 1.1

This section lists bugs from previous versions fixed for Release 1.1. Bug categories include OSPF, RIP, IP, bridging, Device Manager, and miscellaneous bugs.

OSPF

The following OSPF bugs were fixed:

- If an Accelar routing switch was an Area Border Router (ABR) and IP interfaces were removed, the box could prematurely change status to non-ABR.
- Designated router election (DR-Elect) always occurred after 40 seconds, regardless of the dead time interval setting.
- In OSPF environments with many VLANs, "arQuidCheckRec" errors occurred.
- Virtual links did not reach full adjacency.

- During the Hello interval, “WARNING: code=0x0 Task tRCIPTask...” error messages occurred. OSPF Hello packets were being received on spanning tree blocked ports, resulting in the following log message: “285: WARNING : code=0x0 Task tRcIpTask: The routerId=X.X.X.X in the received ospf pkt & myrtid=X.X.X.X are identical - Port=(258).”
- OSPF Hellos were sent when the interface was disabled.
- A new router link was not generated when the only OSPF interface went down.
- The Hello packet did not contain IP Precedence.
- The configured Hello interval did not take effect.
- Link up/Link down in certain OSPF environments could produce 100 percent CPU utilization.

IP

The following IP bugs were fixed:

- The routing switch was vulnerable to the land.c TCP/IP hacker attack.
- Inconsistencies existed between MAC entries and the corresponding ARP and route entries.
- Static ARP table entries were not correctly restored after a reboot.
- The source IP address was incorrect for Telnet and ping replies from routes more than one hop away.
- Functionality and performance have been improved in certain VLSM environments.

RIP

The following RIP bugs were fixed:

- RIPv1 compatible mode did not include the correct subnet mask in updates.
- RIPv1 aggregations that contained local interfaces were not rejected, causing instability.

Management

The following management bugs were fixed:

- SNMP and ping performance was poor under certain CPU load situations.
- You could not ping a VLAN, but you could route through a VLAN.
- Routing protocol traffic addressed to multicast or broadcast addresses could result in 100 percent CPU utilization.
- 10 Mb/s port counters returned inaccurate information, causing improper switch fabric utilization statistics (less than 100 percent).
- No traps were sent when the primary CPU went down.
- When a defective module was replaced, the module configuration was not restored.

Bridging

The following bridging bugs were fixed:

- Spanning tree convergence was improved in certain cases that previously took longer than one minute.
- Spanning tree blocked ports routed IP traffic on routed VLANs.
- When a VLAN was deleted, an “arDeleteMacEntry: Multicast Group ID not found” error was generated.
- Configuration of multiple protocol-based VLANs could corrupt the configuration of a previously configured protocol-based VLAN.
- Deleting a VLAN from the CLI resulted in a machine check.
- Restoring IP protocol VLAN with static members did not work correctly.

Device Manager

The following bugs were fixed in Device Manager:

- When protocol-based VLANs were added to a trunk port, an “undoFailed” message resulted.
- Traps were not sent properly if there was a trap receiver with an invalid IP address.

Miscellaneous

The following miscellaneous bugs were fixed in this release:

- The occurrence of “arQuidChecBin” and “arCheckConsistency” errors.
- The 10 Mb ports violated the Inter Packet Gap.
- A lengthy CLI output was handled incorrectly.

Known Problems in Release 1.1

The following are known problems in the 1.1 release:

- The MAC address of a Proxy ARP router is learned on the incorrect port.
- Multiple policy-based VLANs cannot be assigned to a trunk port.
- Some large OSPF environments result in 100 percent CPU utilization.
- Deleting the OSPF area results in loss of a routable interface.
- If a routing switch running a version 1.0 configuration and high-priority VLANs is upgraded to a version 1.1 configuration, “arQuidCheckRec” and “arDelMacByPortTask” errors are seen.
- An Address Resolution (AR) record is incorrectly set when a static ARP record for the next hop of a static route is inserted after that static route is inserted.