Release Notes for the Accelar 1000 Series Products Software Release 1.3

Accelar 1050/1051 Routing Switch
Accelar 1100/1150 Routing Switch
Accelar 1200/1250 Routing Switch
Accelar Boot Monitor Software Version 1.3
Accelar Run-Time Software Version 1.3
Accelar Device Manager Version 1.3
Accelar VLAN Manager Version 1.3

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 896-00181-D
September 1998

Bay Networks

896-00181-D

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 1.3. Procedures for installing the software are included at the end of these release notes. Currently released components include:

- Boot Monitor Software Version 1.3 (accboot1.3.0)
- Run-Time Software Version 1.3 (acc1.3.0)
- Device Manager Version 1.3 (for Windows 95/98/NT: dm_130.exe; for UNIX: dm_1.3.0.tar.Z)
- VLAN Manager Version 1.3 (for Windows 95/98/NT: dm_130.exe; for UNIX: dm_1.3.0.tar.Z)

This document contains the following major sections:

Note: Many of the new features in Release 1.3 require modules and chassis to be A versions or above with ASICs that are ARU2 or above. Hardware with ARU1 ASICs does not support these features.

Note: Multi-Link Trunking is in the documentation but is not supported in this release. It will be supported in future releases.
IP traffic filters are supported on a limited basis. Refer to "IP Prefix Flow Filters" on page 10.

For the latest information about software issues, refer to the Accelar Products site from the Bay Networks Web page (*www.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-C)
- *Using the Accelar 1050/1051 Routing Switch* (Bay Networks part number 201603-B)
- *Using the Accelar 1100/1150 Routing Switch* (Bay Networks part number 893-01050-B)
- *Using the Accelar 1200/1250 Routing Switch* (Bay Networks part number 893-01049-B)
- *Reference for the Accelar Management Software Release 1.3* (Bay Networks part number 893-01052-C)
- *Reference for the Accelar 1000 Series Command Line Interface Release 1.3* (Bay Networks part number 202086-A)

New Features in Release 1.3

Release 1.3 of the Accelar Management Software includes the following new features:

ARU2 Hardware Support

Release 1.3 of the Accelar Management Software supports ARU1 and ARU2 hardware. It also supports a mixed environment that includes both hardware levels. Certain features require ARU2 hardware. All modules in a chassis must be ARU2 in order to use ARU2 features. In Device Manager, fields pertaining to these features will be grayed out in ARU1 devices.

The following new devices are supported:

- Accelar XLR1216-FX and XLR1216-TF modules
- Accelar XLR1102XD and XLR1202XD modules
- Accelar 1200/1250 XLR1296PS DC power supply

Accelar software also supports a mixed ARU environment. ARU status can be displayed from the command line interface (CLI) or Device Manager. In the CLI, use the command:

`show sys info` (version 1.3 syntax) or `sys show info` (version 1.1 syntax)

This partial display indicates “AR1” (ARU1) as the ASIC version:

```
Card Info:
Slot# Type         Part#  Serial#   HwRev    Oper      Asic Version
                                              Status
3  16x100BaseTXWG  40193      43     v3.0 up                 SQ2   Xy15  SW1
                                                           QUID2   PIC3  AR1
```

In Device Manager, go to Edit Chassis>Chassis. One entry on the screen is AruMode.

Note: In a mixed chassis of ARU1 and ARU2 modules, the features of the lowest common denominator will be in effect. In a chassis with ARU1 and ARU2 modules, the features will be grayed out in ARU1 and in mixed mode (ARU1 + ARU2).

CLI Enhancements

The CLI has been greatly expanded in Release 1.3. The CLI now includes most of the configuration capabilities that are available through Device Manager. In addition, the syntax of the CLI commands has been changed to conform to other Bay Networks devices. In general, the new syntax contains the basic action (config, show, clear), followed by the recipient of the action (ethernet port, vlan, ip access policy), followed by the specific action or parameter (ospf enable, create, mode). Refer to *Reference for the Accelar 1000 Series Command Line Interface Release 1.3* for a complete list of commands and definitions.

The CLI also supports sourcing of configuration files using the `boot` and `copy` commands. The `config` option allows a script file to be read in off of a TFTP server and applied to a default configuration.

In addition to the new commands documented in the manual, the `save nvram standby` and `save standby` commands have been added to allow simultaneously saving configuration changes to both active and standby SSF modules in an Accelar 1200 switch.

The `show ip route info` command now has an optional parameter to show ranges of subnets. For example, `show ip route info -s 8.8.0.0/15` will show all subnets within the given subnet mask.

Virtual Router Redundancy Protocol

This feature requires ARU2 hardware. Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure in static default-routed environments. End stations are often configured with the IP address of a router on the local subnet as the default gateway. If the default gateway fails, VRRP introduces the concept of a virtual IP address (transparent to users) shared between two or more routers connecting the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides a dynamic default gateway redundancy in the event of a failure.

Four VRRPs (isolated routing ports and VLANs) are allowed per Accelar switch, and all VRIDs must be unique.

Using Device Manager, to set up VRRP parameters:

- On a port, select Edit->Port->VRRP.
- On a VLAN, select VLAN->VLANs->Basis->IP->VRRP.
- To globally enable, go to Routing->VRRP.

Refer to Chapter 9, "IP Interfaces and Router Management," in the *Reference for the Accelar Management Software*, for more information.

Using the command line interface, set up VRRP parameters with these commands:

- `config ethernet <ports> ip vrrp....`
- `config vlan <vid> ip vrrp...`

To view VRRP status, use these commands:

- `show ip vrrp info`
- `show ip vrrp stats`

Refer to *Reference for the Accelar Series 1000 Command Line Interface* for more information.

Accept and Announce Route Policies

Route policies provide control over information exchange or redistribution between routing protocols. Policies can also be used for manual route aggregation, route generation (including a default route), RIP metric manipulation, and OSPF external metric manipulation. An IP routing policy consists of matching criteria, set parameters, and action. The matching criteria are used to decide whether or not a policy should be applied to a certain route. Once a policy is selected, if the action is to announce, the parameters are used to construct the route advertisement. Accelar supports the following policies:

- OSPF Announce policies are applied for non-OSPF routes in an Autonomous Boundary Router.
- OSPF Accept policies are applied when the OSPF engine computes the external routes due to a topology change or an external link-state advertisement.
- RIP Announce policies are applied while sending a RIP update.
- RIP Accept policies are applied when the router receives a RIP update.

Refer to Chapter 12, “IP Policies,” in the *Reference for the Accelar Management Software*, for more explanation and examples of how these policies are used. In Device Manager, IP policies are set up by selecting Routing->IP Policy. Chapter 12 explains how this is done.

In the CLI, use the `config ip policy` commands. The commands are listed and parameters defined in *Reference for the Accelar Series 1000 Command Line Interface.*

IP Multicast Optimization at Layer 2

This feature requires ARU2 hardware. IP multicasting provides services such as the delivery of information to multiple destinations with a single transmission and the solicitation of servers by clients. When used as a switch, an Accelar routing switch supports Internet Group Management Protocol (IGMP) to prune group memberships per port within a VLAN. This feature is called IGMP snooping. It allows users to optimize the multicast data flow for a group within a VLAN only to members of the group.

In Device Manager, you can set up IGMP snooping by selecting VLAN->VLANs->Snoop. This feature is described in Chapter 13, “IP Multicasting,” of *Reference for the Accelar Management Software*.

In the CLI, use the `config vlan <vid> igmp-snoop` commands described in *Reference for the Accelar Series 1000 Command Line Interface.*

Access and Security Features

The following features and changes have been made to Accelar access and security.

Access Levels

Two new access levels have been added to the Accelar login through Device Manager or the CLI. (These levels are not used on the Accelar Configuration Page on the World Wide Web.) The new levels are:

- Layer 2 (l2) Read-Write access allows the user to change only layer 2 (switching) parameters on the switch.
- Layer 3 (l3) Read-Write access allows the user to change layer 2 *and* layer 3 (routing) parameters on the switch.
- The default logins and passwords are l2, l2 and l3, l3.

Management Access

The Management Access feature allows you to control system access by setting access lists for services to prevent or allow access to the switch. You can specify which hosts or networks can access the switch through HTTP, rlogin, SNMP, and Telnet. If no access policies exist for a given service, the default is to allow the connection. Access policies apply to connections via the I/O ports and not to the console or modem port. The lowest number indicates the highest precedence.

In Device Manager, access policies are set through the Edit->Security->Access Policy window. In the CLI, use the *config ip access-policy* commands. Refer to the reference guides for more details.

Port Locking

In this release, you can lock a port so that its configuration cannot be changed. In Device Manager, this is done through the Edit->Security->Port Lock window. In the CLI, use the `config sys set portlock <on|off>` command.

RIP Enhancements

The following paragraphs describe the Routing Information Protocol (RIP) features and enhancements included in Release 1.3.

Autoaggregation

This feature has been added to the RIP interface parameters to allow automatic route aggregation if RIP v2 is used. When enabled and RIP v2 is used, the router automatically aggregates routes to their natural mask when they are being advertised out an interface in a different class network.

In Device Manager, autoaggregation is enabled on the Edit->Port->RIP window (for ports) or the VLAN->Basic->IP->RIP window (for VLANs). In the CLI, use the `config ethernet <ports> ip rip auto-aggr <enable|disable>` or `config vlan <vid>ip rip auto-aggr <enable|disable>` command.

Enable/Disable Triggered Updates and Manually Trigger Updates

The previous release supported triggered updates for RIP. This feature allows the user to disable triggered updates. It would be desirable in unstable networks where problems could occur if too many updates are generated. Triggered updates are enabled or disabled on a per interface basis, with the default being disabled.

In Device Manager, select TriggeredUpdateEnable true or false in the Edit ->Port ->RIP window. This feature does not apply to VLANs. The manual trigger is not available in DM.

From the CLI, you can also select to manually trigger updates. The appropriate commands are:

```
config ethernet <ports> ip rip manualtrigger
config ethernet <ports> ip rip trigger <enable|disable>
config vlan <vid> ip rip manualtrigger
config vlan <vid> ip rip trigger <enable|disable>
```

Adjustable Holddown Timer

After RIP has determined that a network is unreachable, it will continue to advertise a route to that network for a default holddown period of 90 seconds with a metric of infinity (16). The range is 0 to 360 seconds.

In Device Manager, go to Routing>RIP. On the Globals window, enter the Holddown Time parameter. In the CLI, use the command:

```
config ip rip holddown <seconds>
```

OSPF Enhancements

The following paragraphs describe OSPF enhancements.

OSFP Holddown Timer

In addition to the RIP holddown timer, you can also set maximum holddown time for OSPF. In Device Manager, go to Routing->OSPF->General. The last item in the window is the Holddown counter. In the CLI, use this command:

```
config ip ospf holddown <seconds>.
```

The range is 3 to 60 with a default of 3 seconds.

OSPF Area ID

The AreaId field has been added to the Port->OSPF window to allow the user to specify the Area ID prior to assigning the IP address. The default value is the 0.0.0.0 backbone area.

In addition, an area option import summary has been added to the OSPF area to provide control for the user to specify whether or not the summary link state advertisements (LSAs) should be imported.

The user can also delete an area without disabling the OSPF admin status if the area to be deleted does not have any interfaces.

Other OSPF Features

The router ID can be dynamically changed.

The Authentication Key is now entered as an ASCII value.

VLAN Enhancements

Release 1.3 supports IP subnet-based policy VLANs, as well as additional protocol-based VLANs.

IP Subnet (Source Address) Policy for VLANs

This feature applies only to ARU2 hardware. IP subnet-based VLANs allow network managers to create VLANs based on IP address prefixes. Each IP subnet or network number can be associated with a dynamic VLAN. All IP traffic with source addresses within the policy address range become part of the dynamic VLAN. Port membership is timed out if traffic meeting the policy is not seen on that port for a specified time interval. Multiple IP subnet policies can be applied to all ports on each switch.

When using IP subnet-based VLANs, the port connecting to existing routers should always be configured as a dedicated routing port or with 802.1Q tagging (if supported by the router). Otherwise, traffic originating from sources that do not fall in the address range of the defined IP subnet policies will be dropped.

In Device Manager, select VLAN->VLANs->Basic->Insert and choose Type “byIpSubnet.” This feature is not configurable through the CLI.

Protocol-Based VLAN Enhancements

In this release, you can select IPv6 as the protocol type for a protocol-based VLAN. All Ethernet traffic of Ether Type 0x86DD must belong to the IPv6 VLAN.

User-defined protocol type policies for VLANs allow you to specify any protocol based on Ether Type, IEEE 802.2 SSAP/DSAP, or PID field in the SNAP header; frames with the matching value are then mapped to the specified VLAN. This feature is helpful to users who have portions of their network running old protocols for which there is no direct support under the protocol definition types on the switch. Any Ether Type can be specified in the user-definable field, one Ether Type per VLAN. For information about limitations on available protocol IDs, refer to the sections on "User-Defined Protocols" in Chapters 1 and 8 in *Reference for the Accelar Management Software*. To configure IPv6 VLANs or a user-defined protocol, go to VLAN->VLANs->Basic->Insert. Select UserDefinedPID and enter the ID, or select ipV6.

These protocol-based VLANs are not configurable through the CLI.

In ARU2 hardware only, this release also supports NetBIOS and SNA protocols.

IP Prefix Flow Filters

Note: Please see "Known Software Limitations or Differences" on page 15, for a listing of limitations.

This feature requires ARU2. IP Prefix Flow Filters extend the application intelligence of the Accelar routing switches by enabling network managers to manage aggregate traffic flows using filters, based on source and destination IP address prefixes, source and destination TCP/UDP ports, and the IP protocol type (also referred to as an aggregate IP Microflow). An IP Microflow is defined as a single instance of an application-to-application flow of packets, which is identified by source address, source port, destination address, destination port, and the protocol ID.

Each filter is defined as a combination of match criteria and the associate actions. The match criteria allow users to specify an IP Microflow or an aggregate IP Microflow. Actions that can be associated with each Microflow include: Drop, Prioritize, Mirror, set VLAN tag priority bits, or allow TCP connections. Users can specify multiple IP Microflows as the match criteria and a different action for each Microflow. For example, filtering (denying) http traffic originating from the engineering subnet/VLAN destined for the HR subnet/VLAN and allowing http traffic between the engineering subnet/VLAN and the marketing subnet/VLAN.

IP Prefix Flow Filters apply only to the traffic that is routed by the Accelar switch.

When no filters are applied, IP packets ingressing a port are forwarded (routed) through the switch. If a filter is applied to a port and an IP packet that matches the filter criteria ingresses on this port, the packet will be dropped or forwarded (routed) depending on the action mode (drop, forward, or default) of the matching IP Filter record. In this release, the action mode of default is equal to the action mode of forward.

In addition to forwarding or dropping the matching IP packet, any of the mirror, priority, or IEEEvlan priority actions specified in the filter record will be taken.

If multiple filters are applied to a port and an IP packet matching more than one of these filters ingresses on this port, the packet will be dropped if any one of these matching filters has an action mode set to drop. When an IP packet matches multiple IP Filters, the cumulative mirror, priority, and IEEEvlan priority actions specified in the matching filters will also be taken.

In Device Manager, these filters are edited through Routing->IP Filters->Filters window. In the CLI, the following command sets are used to set filters:

- `config ethernet <ports> ip traffic-filter`
- `config ip traffic-filter`

Syslog

Syslog is a facility on UNIX machines for message logging. Messages sent to the system daemon either locally or remotely have an associated facility and severity.

The Accelar switch supports sending messages to as many as 10 different syslog hosts (with a default of 5). Accelar software log messages have four severity levels (INFO, WARNING, ERROR, and FATAL). The syslog feature is used primarily for remote logging of routing switch system messages.

In Device Manager, select Edit->Diagnostics->Syslog Info to set parameters. In the CLI, use the `config sys syslog` commands to set.

Bay Topology Discovery Protocol

Bay Topology Discovery Protocol allows the Accelar switches to exchange topology information with other Bay Networks devices. You can now set topology through the Accelar software. In Device Manager, go to Edit->Diagnostics->Topology Info to view or set status. In the CLI, use the command `config sys set topology <on|off>` to turn topology on or off.

Enable/Disable IP Forwarding

In release 1.3, enabling IP forwarding replaces the Port and VLAN fields to turn routing on or off. This feature globally enables or disables IP forwarding. In Device Manager, go to the Routing->IP window. In the CLI, use the following commands:

- `config ip forwarding disable`
- `config ip forwarding enable`

System Clock

This release supports the system clock hardware on new Accelar 1200 modules. The time stamp is used in the log files and is viewable and settable from the CLI. It is available only if a real-time clock is present. (The Accelar 1050 and 1051 switches do not support a real-time clock at this time; it will be supported in a future release.)

- `config setdate` sets the calendar time.
- `date` displays the set time.

Other Enhancements

The following list contains other general enhancements included in this release:

- CIDR support (ARU2 only) is now available.
- VLSM enhancements include allowing IP subnet ranges to exceed 11 bits.

- Device Manager and VLAN Manager are now downward compatible. When Device Manager or VLAN Manager is launched, it automatically determines what version of software is running on a device and opens the appropriate version of DM/VM.
- Device Manager and VLAN Manager can now be run on IBM AIX operating system.
- Redundant CPU enhancements add support for Telnet, rsh, rlogin, and TFTP to the standby CPU.

Usage Tips or New Features

The following sections contain tips or procedures not documented elsewhere.

Copy Script File to a Running Configuration

An extension of the `copy` command allows a script file (an ASCII-based text file containing CLI commands) to be read by the switch and the commands executed as though they were typed at a console session. By default, script execution does not display at the device where the command was issued. However, if the optional `debug` parameter is used, then the execution of the command in the script file and the results are output to the device from which the command was executed.

The script file itself is an ASCII text file. The first line of the file must include a pound sign (#) followed by a carriage return, with the remaining lines containing valid CLI commands, one per line.

The format of the command is:

```
copy <sourcedevice:filename> running-config [debug]
```

where:

`sourcedevice` may be a flash, PCMCIA, or TFTP-based file server. If “tftp” is specified, you will be prompted for the server IP address and the file name.

`filename` is the name of the file to be copied.

`[debug]` is the optional parameter that allows viewing execution of the script.

Note: Exercise care when executing script files from within the CLI. The command execution will reference from your current position in the directory structure.

Discard Tagged Frames

To optimize performance, on access ports with users where you do not expect to see tagged frames, you should set the discardtaggedframes option to true. However, for access ports for interconnecting switches, it is probably better to keep tagged frames (that is, set discardtaggedframes to false). In that way, if you should convert an inter-switch port from an access port to a trunk port, you will not lose connectivity.

Bugs Fixed in Release 1.3

All bugs fixed through Release 1.1.6 are fixed in this release. If you need more information about which bugs were fixed in releases 1.1.2 through 1.1.6, refer to *Addendum to the Release Notes for the Accelar 1000 Series Product Software Release 1.1.6* (Bay Networks part number 202693-B). This document is located on the Web under Software Release Notes at the URL:

http://support.baynetworks.com/library/tpubs/nav/rtswitch/accelar.htm

The following additional bugs have been fixed since the 1.1.6 release:

- Routes “flapping” in redundant networks could cause the Accelar switch to forward packets to an all 0s MAC address. (81722)
- For non-real-time clock systems, trace and log messages now show days in system uptime. (82311)
- In some VLSM configurations, directed broadcasts were routed back to the original network. (83353)
- The CLI will now check for the maximum number (500) of static routes allowed. (83456)
- Static routes could be overwritten by RIP-learned routes. (82849)
- Gigabit links could fail to initialize after booting. (82301)

Known Software Limitations or Differences

This section describes the latest functionality and known limitations for the Accelar 1000 Series software and provides additional information not documented elsewhere.

Routing 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 zeros” are not supported.

Traffic Prioritization

An Accelar 1000 Series switch can operate in either of two modes: Best Effort mode or Priority mode. The factory default setting is Best Effort mode. 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 run-time CLI using the `config sys set flags highpriomode` command. Note that after changing this setting, you MUST save the configuration and reboot the switch before the change takes effect. For information about changing priority in Device Manager, refer to Chapter 16, “Prioritization,” in *Reference for the Accelar Management Software*.

VRRP

If the same VRRP IP address has different virtual router IDs (VRIDs) on different switches, problems can occur in differentiating between master and backup router. VRIDs must always be the same for the same IP address. In addition, the VRRP IP must not be the IP address used to manage the switch for Device Manager. VRRP IP addresses do not respond to ping requests whether they are local or remote interfaces.

VRRP is not fully supported in this release. The feature can be used with isolated routed ports connected to shared media; however, the following configurations are not supported:

- Configurations where an isolated routing port with VRRP enabled is connected to the port of a switch (for example, the BayStack 350 switch)
- Configurations where VRRP is enabled on the virtual router ports of VLANs

IGMP

IGMP snooping is only supported for IGMP version 1 and IGMP version 2 packets without the "IP Router Alert Option." Refer to rfc2236 (Internet Group Management Protocol, Version 2) and rfc2113 (IP Router Alert Option) for details. Support for IGMP version 2 packets with the router alert option will be provided in a future release.

In a VLAN with multicast enabled, the switch flushing database learns the source MAC address of IGMP queries and IGMP data (UDP) and the member (source IP address) associated with the multicast group in the IGMP report, but it does not learn the source MAC address of IGMP reports.

Accelar Spanning Tree

The behavior of an Accelar switch is consistent with the latest revision of the Spanning Tree Protocol (STP) code from the IEEE 802.1D, but the operation is slightly different from other Bay Networks devices. When an Accelar switch receives BPDUs from another switch with a lower bridge priority with a message age of 5 and a maximum age of 6, the Accelar switch assumes itself to be the root and transmits its own timer values in its BPDUs, rather than incrementing the message age and sending BPDUs with the new message age value using the other device as root.

VLAN

When routing is enabled and IP addresses are added to the VLAN, these routes show up in the active routing table and appear as routes in the Address Resolution (AR) table dump. In this release, VLANs remain in the routing table at all times, even if no ports are assigned.

The Accelar switch supports mirroring only of bridged (layer 2) traffic. Support of IP routed (layer 3) traffic will be available in a future release. Please note that attempts to mirror routed traffic will result in the routed traffic being redirected to the mirror port.

OSPF

The following are known limitations with OSPF:

- In an Accelar switch using OSPF protocol, the dead interval set must be four times the hello interval.
- When the next hop for an external route is reachable by an intra-area or inter-area OSPF route, the router does not at this time provide the next-hop address in the Forwarding Address field of an ASE LSA.
- If OSPF redistribution in release 1.1 has the StaticToOspfAdverDefRte option set to false, it is not correctly converted to the OSPF Announce policy in release 1.3. If the option is set to true, it is converted correctly because the OSPF Announce policy is created to announce all static routes including the default route.

  To work around this problem:

  - Create an OSPF Announce policy to announce all static source routes.
  - Create a network entry with 0.0.0.0/0.0.0.0.
  - Create a network list entry to include the network entry.
  - Create a policy to ignore 0.0.0.0/0.0.0.0.
  - Set the ignore policy at a higher precedence.

Counters

Egress counters are inconsistent with ingress counters. The egress unicast counters miss a small percentage (7/648,000) of packets egressed. Unknown unicast packets get counted as unicast packets, and some packets that should be counted are missed. Multicast and broadcast egress counters with ARU1 hardware have greater inaccuracies, but these are accurate with ARU2 hardware.

Filters

The Accelar switch supports mirroring only of bridged (layer 2) traffic. Support for mirroring of IP routed (layer 3) traffic will be available in a future release. Please note that attempts to mirror routed traffic using IP Filters result in the routed traffic being redirected to the mirror port; also, this function is currently possible only with a filter action of drop.

Source and Destination Filters may not be more specific, in terms of mask length, than their corresponding routes. For example, a Destination Filter of 100.100.100.0/24 requires that the route to the 100.100.100.0 network be learned with a mask of 24 or more. If you need a more specific Source or Destination Filter, you must create a matching route entry having a mask length no smaller than that of the filter entry.

When applying a Source or Destination Filter to a port, you may have to flush the learned ARP entries on that port before the filter will take effect.

In this release, packets not matching a specified filter (or set of filters) are always forwarded, and so exclusion filters (forward packets to x.x.x.x. and drop everything else) are not possible.

Each filter is uniquely identified by filter type, source address/mask, and a destination address/mask. For this reason, you may not have multiple IP Filters of the same type with the same address/mask pairs.

Not more than eight Global Filters may be applied to any port or any set of ports sharing an ARU.

The minimum size mask length for a Source or Destination Filter is 8 bits (2555.0.0.0). The maximum size for all filters is 32 bits (255.255.255.255).

Source and Destination Filters may specify a TCP or UDP Source Port or Destination Port to match. When this is specified, the implied operation is always equals.

When specifying a TCP or UDP Source Port or Destination Port for a Global Filter, the matching operation may be equals, not equal, greater than, or less than.

RIP2 Next Hop Support

The next-hop field helps to eliminate packets being routed through extra hops (RFC 1723). The current version of Accelar software does not support this feature.

Router and ICMP Redirect Messages

When the Accelar switch is used as a router, it does not send ICMP redirect messages to a host that has an incorrect default-gateway entry. Routing is performed by the address resolution unit (ARU). As long as there is a valid route in the ARU table to the destination, the ARU will route the frame.

The Accelar switch does not generate ICMP destination unreachable error code 13 messages when a host or network is reachable but is blocked administratively by IP filtering.

Static Routes

When static routes are configured pointing to a specific next hop and the interface to reach the next hop is deleted, the static routes disappear from the routing table and will not show up again unless the switch is rebooted or reconfigured.

When an ARP request is seen, the ARP subnet gateway can determine if it knows a route to the target host by looking in the routing table. The Accelar switch responds to the ARP request if the destination is reachable by a dynamic route, but it does not respond if it is reachable by a static route.

Although the RFC states that the default route should not be used when checking for a route to the target host of an ARP request, the Accelar switch does respond to requests for a target reachable only by the default route.

The behavior of routable interfaces has changed in this release. In the case of an isolated routed port, if the link on the port goes down or is administratively brought down, the IP circuit will be brought down and the route will be deleted from the routing table. All static routes associated with the interface will also be deactivated. For routable (port-based) VLANs, if the link of the last port in a VLAN goes down or is administratively brought down, the IP circuit associated with the VLAN will still be active in the switch. The route will still be in the routing table, and all static routes will also be active. Thus, all traffic that exercises this deactivated route will no longer be forwarded and will be dropped by the switch.

Network Management Functionality and Limitations

This section describes the latest functionality and known limitations of managing an Accelar 1000 Series release 1.3 routing switch. Unless otherwise noted, the information is valid across all management platforms.

Port Statistics Support

The following tables indicate which counters are supported in the Accelar 1000 Series. Table 1 contains support information for the port interface counters; Table 2 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 1. Accelar 1000 Series Interface Counter Support

Counter	10BASE-T	100BASE-TX/FX	1000BASE-SX/LX
InOctets	X (ARU2 only)	X (ARU2 only)	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

Table 1. Accelar 1000 Series Interface Counter Support (continued)

Counter	10BASE-T	100BASE-TX/FX	1000BASE-SX/LX
OutBroadcastPkts	Not Available*	Not Available*	X
OutDiscards	Not Available*	Not Available*	Not Available*
OutErrors	Not Applicable	Not Applicable	Not Applicable

* These counters are forced to zero (0).

Table 2. Accelar 1000 Series RMON Counter Support

Counter	10BASE-T	100BASE-TX/FX	1000BASE-SX/LX
etherStatsDropEvents	Not Applicable	Not Applicable	Not Applicable
etherStatsOctets	X (ARU2 only)	X (ARU2 only)	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 1.3 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 3)

Table 3. 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.
rcLinkOscillation	Excessive link state transitions on a port.

Device Manager Limitations

This section describes the latest functionality and known limitations of Accelar Device Manager (DM) version 1.3.

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 2 on page 21. This limitation applies to all platforms.

Solaris and HP/UX Platforms

The following functionality applies to Solaris, HP/UX, and AIX 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 (Solaris)

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 not automatically size the resized window. The subwindow will automatically size if it is closed and reopened.

Windows 95, Windows 98, and Windows NT

The following information applies only to Microsoft® Windows® 95 and 98 and Windows NT® platforms.

Context-Sensitive Online Help

Device Manager displays the online Help screens using the default Web browser. With 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 takes you to the top of the HTML file.

RMON Alarm Traps

On Windows hosts, Device Manager may occasionally fail when RMON alarm traps are received. When this happens, the host displays a message indicating the "NP_WSX32.exe" driver has crashed. If you see this problem, restart DM. If the problem persists, reconfigure the RMON feature so that alarm traps are not sent to the Windows host.

SNMP Trap Support

Device Manager under Windows supports only 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 run-time error does not cause any corruption of data. To recover from the abnormal termination, terminate the NetPlus/32 task before restarting DM.

Known Problems in Release 1.3

The following sections list problems known to exist in release 1.3.

Management

- Valid packets with more than 1514 bytes on a trunk port are counted as oversized.
- ARU2-mode Gigabit access ports treat all inunicast traffic as inbroadcast traffic. (81616)
- Some file details displayed in DM and CLI are incorrect; they are showing as release 1.1 files instead of release 1.3 files. (83274)
- The “Card Removal INFO” message sometimes fails to display. Card Removal INFO displays correctly during a single card reseat but fails to display when a second card is removed and reseated very quickly after the first card is reseated. However, if you remove and reseat cards several times, Card Removal INFO then sometimes fails to display at all for a given slot, until a card in another gets reseated. (83892)
- The error message “alsWriteClientRecord () failed” sometimes occurs while a configuration is being saved. (84590)
- The LastSaveToNVRam value is incrementing incorrectly if the device is reset with Device Manager open. (83351)
- When a power supply is brought down, the syslog is updated with a power supply down message; but when the power supply is brought back up, the syslog file does not get updated. (84695)
- When the state of a port is changed to “testing” from either the CLI or DM, the local screen displays an error. (84974)
- If a link-state transition (port going up or down) occurs during a squeeze flash operation, the Accelar switch may crash when a port is subsequently made active (link up). (85664)

VLAN

- Some connectivity problems exist with VINES and AppleTalk protocol-based VLANs. (79992)

  The workaround for this problem is to add all the ports into the AppleTalk protocol-based VLAN and then remove those ports you do not want in the VLAN. The remaining ports should be the ports that are members of the AppleTalk protocol-based VLAN.

- When the port priority is set to 0, the rcStgPortDesignatedPort counter returns no value. Also in Device Manager, in the Edit Port->Spanning Tree menu, the Designated Port field shows no value when port priority is set to 0. (83473)
- Incorrect warning error messages sometimes occur about mismatched VLAN PIDs. (84077)

IP

- The Accelar switch might unnecessarily send ARPs for the source address of ingressing broadcast or directed broadcast packets. (82284)

  The workaround for this problem is to reboot the Accelar switch.

- CPU throttling for IP traffic is not working properly. When two switches are connected with two access ports in a routing VLAN, if the link is dropped, the ARP entry goes away, all traffic is flooded to the CPU, and CPU utilization goes to 100 percent. (83680)
- With proxy ARP enabled, an Accelar switch in an IP subnet VLAN with all ports as potential members can send ARP requests for its own IP address. (84945)
- The CLI and DM are inconsistent when adding static ARP entries. When entering the same IP address, Device Manager will indicate that you are using incorrect values, while CLI will allow you to add the IP address. (84979)
- The error message “rcIpAddRoute: null pRent” may occur when an ARP entry is added either statically or dynamically. (82152)
- In the “Insert Filters” dialog box of Device Manager, selecting “Global” followed by one of the ProtocolType buttons (icmp, tcp, or udp) will cause the SrcOption and DstOption fields to be erroneously grayed out. (84643)

  To work around this problem, select the “Global” button again to refresh the dialog box.

- If a packet is dropped because of an IP Filter, an ICMP destination unreachable message is not generated. (84359)

DM/VM

- A DC power supply does *not* display as a DC supply in Device Manager on an Accelar 1250 switch. However, a DC power supply does display correctly as a DC supply on an Accelar 1200 switch.

  The workaround for this issue is to use the CLI. A DC power supply is correctly identified as a DC supply through the CLI on an Accelar 1250 switch.

- Device Manager does not install correctly on NT Server 4.0.
- When all modules in a switch are ARU1 modules, the Device Manager and the CLI displays indicate ARU mode is ARUmixed, instead of ARU1. (83671)
- Isolated routing ports are not reflecting Device Manager changes in policy-based VLANs. If you create a policy-based VLAN with some ports as not allowed and then edit one of these ports and assign it an IP address, the allowed port still shows as red in the policy-based VLAN but gray in the default VLAN. (80633)
- When running multiple DM sessions to different agents running the same software code image, the session that gets connected first comes up with the requested port information, whereas the other and all successive DM sessions do not provide any information on the agent. (81938)
- Device Manager allows setting *less than*, *greater than*, and *not equal to* as port match values for traffic filters. Source and destination traffic filters actually only support *equal to* match values. (82147)
- ARU2 counters increment incorrectly in Device Manager. Using the Device Manager Edit->Diagnostic->ARStats screen, the wrong type of deleted VLAN counters were incrementing. (83442)
- Launching PC-based Device Manager against an unsupported device results in a “timeout” error instead of an “unsupported” error message. (84129)
- In Device Manager, in the Routing->IP->ARP window, the wrong interface value was displayed, or “Err” was displayed for the interface. (84658)

- If DHCP-relay is enabled, the mode parameter is changed, and then the user displays the DHCP-relay interface, an error message “Unable to find Ip Record for Ifindex 16” is displayed. (83888)
- In Device Manager, in the Edit->Port->Interface window, when attempting to select the action FlushMacFdb at a port, the error message “rcPort locked” appears even if port locking is disabled. (85003)

OSPF

- You cannot create a static route that overrides an OSPF dynamic route. To work around this problem, first bring down a link to delete the OSPF route and then create the static route. In addition, if the static route is subsequently deleted, the dynamic route is not relearned. (84509, 84980)
- OSPF frames across Q trunks are not set as high priority. With OSPF frames, the precedence bit is set in the IP header so they can be prioritized. When these frames are transmitted out a trunk port, the priority bits for the Q header are set to 0, which cause the Q tagged OSPF frames to be placed in a low-priority queue instead of a high-priority queue. (84923)
- Configured as a nonbackbone ABR, without any interface in the backbone area, with interfaces in a contiguous nonbackbone area, and in a Discontiguous area, respectively, the Accelar switch will appoint itself as a backbone OSPF ABR. (85004)

RIP

- You cannot not add a static route if a RIP route exists for the same destination. Also, a RIP route for the same destination overwrites an existing static route. The static route should have higher preference in Accelar software. (83012)

IGMP

- A host sending an IGMP version 2 report with the IP Router Alert option may cause host members to leave and rejoin the multicast group. (83719)

VRRP

- The switch does not become VRRP master if the IP address and the virtual router IP address (VRID) are the same for the IP-subnet VLAN. (84946)

- VRRP advertisements are sent by the CPU even if the link is down. (81351)
- VRRP does not work in a switched layer 2 environment because it does not use VRRP virtual MAC addresses for VRRP advertisement packets and because the backup Accelar switch does not do a gratuitous ARP request upon transition. This not an issue if the switched media is replaced with a shared media (hub) environment.(85755 and 85758)
- A VRRP router port does not issue a broadcast ARP request for a host on a network that is directly connected to the port. However, the VRRP router port does learn the ARP information for the host promiscuously. When the VRRP router port sees a broadcast MAC frame issued by the host (for example, an ARP request) or receives a unicast MAC frame from the host, it will learn the host ARP information. (85855)

RMON

- RMON Traps are showing negative integer values for CPU utilization. (82501)
- When graphing RMON history for a port, the pie chart will not graph more than 33 samples. (84966)

Downloading the Software

This section includes instructions for downloading the Accelar run-time and boot monitor image software.

Prerequisites for Upgrading

To upgrade the flash images, you need the following:

- New version 1.3 run-time 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 boot monitor image is stored in a reserved area in flash memory that cannot be accessed directly. The boot monitor image is updated by executing a boot monitor updater that in turn updates the boot monitor image. The run-time image is stored in flash memory in an area that is accessible using the command line interface (CLI).

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

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

Version 1.3 Flash Memory Update Procedure

The procedure assumes that you will be updating both the boot monitor image and the run-time image in flash memory to version 1.3. Be sure to read and understand the procedure *before* attempting to upgrade the software.

The procedures for upgrading from version 1.0.0 or from version 1.1.6 are basically the same, except for the resulting screens. Examples from both upgrades are included. The procedure for reverting to a pre-upgrade state follows this update procedure. The downgrade procedure is supported only for version 1.1.6 upgrades. 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.

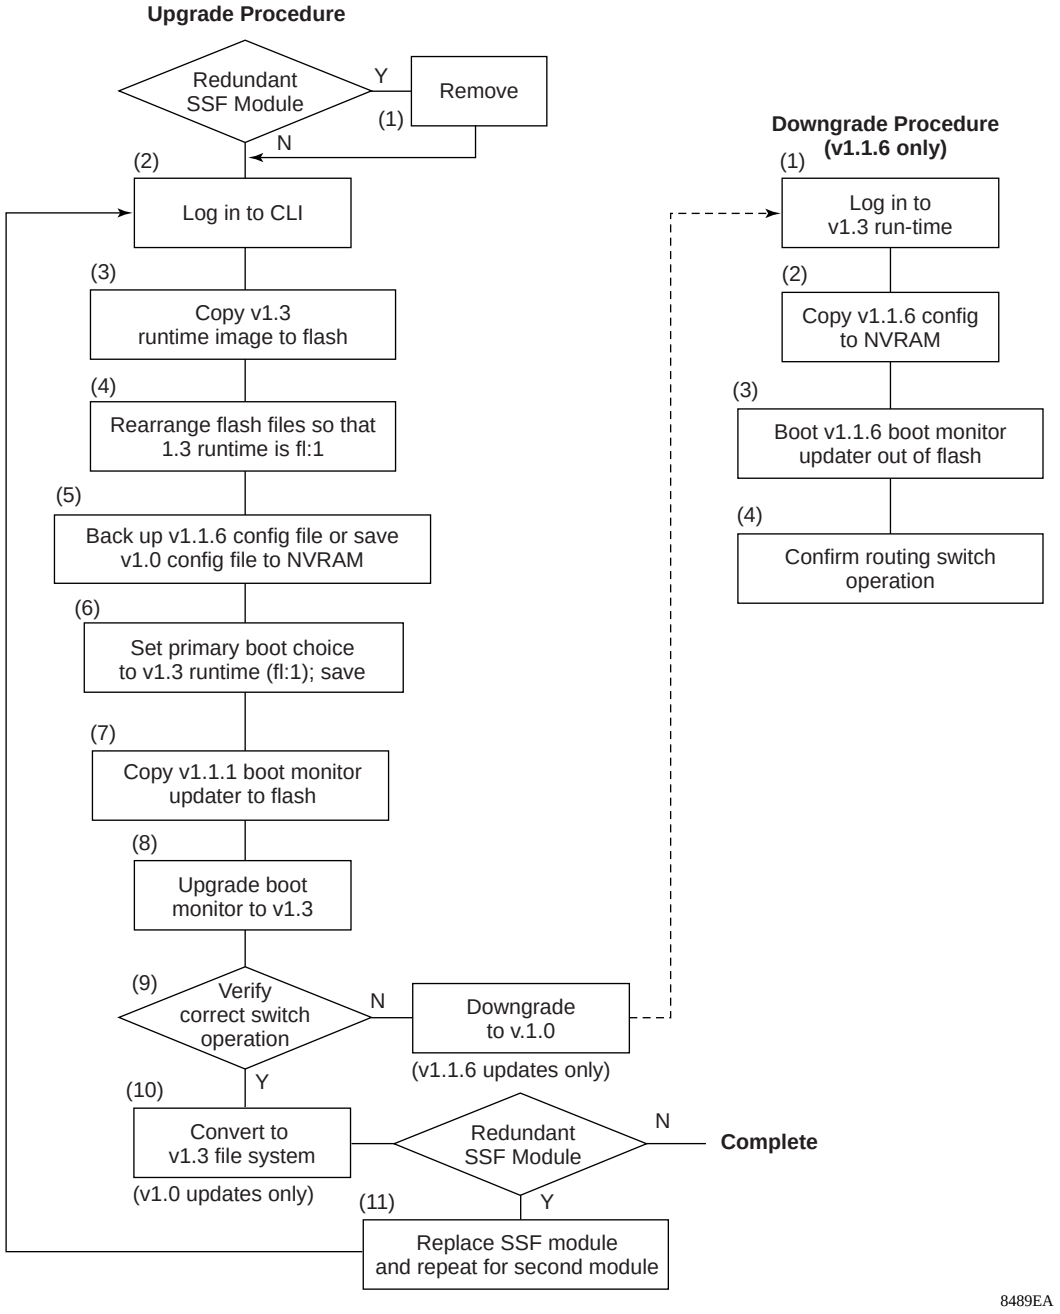

Figure 1. Software Upgrade/Downgrade Flow Chart

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

Caution: It is *very important* that you follow the instructions below in the indicated sequence. When upgrading from version 1.1.6, following the instructions simplifies downgrading the software if the upgrade is not successful. When upgrading from version 1.0.0, downgrading is not possible. Contact the Bay Networks Technical Support Center if the upgrade is not successful.

To upgrade the flash software to version 1.3:

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

   For Accelar 1200 chassis with redundant Silicon Switch Fabric (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 module will cause SSF modules to swap active and standby roles, complicating the upgrade process.

2. **Log in to the run-time 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 1.1.6 if necessary.

   At the CLI `Login:` prompt, enter the password for Read-Write or Read-Write-All access privileges.

```
*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1997     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.x.x      *
*******************************
Login: ***
Password: ***

Bay>
```

3. Copy the version 1.3.x run-time image into flash.

The version 1.3.x run-time 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 run-time CLI `copy` command allows the copying of files to and from flash memory, PCMCIA cards, and TFTP servers. To copy the new run-time 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 run-time image file as follows:

```
Bay> copy tftp flash
Enter source tftp server address [0.0.0.0]: 10.10.20.100
Enter source file []: acc1.3.x
programming flash:3 ... xxxxxx bytes
verifying flash:3 ... done
tftp result: success
Bay>
```

The run-time image is a large file and can take about one minute to download. The run-time 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 run-time 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.

In this step, the current (version 1.1.6 or 1.0.0) run-time image (flash:1) and the log file (flash:2) will be copied to the end of flash; the version 1.3.0 run-time image will remain as file three (flash:3). There are some differences in appearance when upgrading from version 1.1.6 or 1.0.0, as indicated below:

When upgrading from version 1.1.6:

```
Bay> directory flash
Device: flash
FN NameFlagsLength
-- ---------------
1  acc1.1.6XZN1000443
2  syslogLN131072
3  acc1.3.0 XZN1362138
--     ------
3   files bytes used=  2642004 free=1552300

Bay>
```

When upgrading from version 1.0.0:

```
Bay> dir fl
Device: flash
------

File 1: /export/acura3/projects/first/rel1.0/rel1.0.0/main/hw/accelar.st
Version: rel1.0/rel1.0.0/main on Fri Jan  9 11:28:28 PST
Length: 895550 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x6e79

File 2: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 41,0 free: 214,256

File 3: 0.b23.st on Fri Sep  4 16:15:46 PDT 1998
Version:
/export/home2/projects/first/rel1.3/rel1.3.0/main/hw/acc1.3.0st on Fri Sep
4 16:15:46 PDT 1998
Length: 1361962 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0x4065

Bay> copy flash:1 flash
 programming . . . flash:acc as file#4 1000443 bytes

Bay> copy flash:2 flash
 programming . . . flash:acc as file#4 1000443 bytes
```

```
Bay> directory flash
 Device: flash

 FNNameFlagsLength
 -----------------
 1accl.1.6XZN1000443
 2syslogLN131072
 3accl.3.0XZN1362138
 4accl.1.6XZN1000443
 5syslogLN131072
 ---------
 5filesbytes used=3735987free=458217
```

Now delete files 1 and 2, and squeeze the flash file system. This step has the effect of renumbering the files so that the version 1.3.0 run-time image is now the first file in flash (flash:1), the version 1.1.6 run-time 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 . . . success
```

Upgrade from version 1.1.6:

```
 Bay> directory flash
 Device: flash
 FN NameFlags  Length
 -- ---- ------ -------
 1  acc1.3.0XZN 1362138
 2accl.1.6XZN1000443
 3  syslogLN 131072
 --        ------
 3   files     bytes used=  2642004 free=1552300
```

Upgrade from version 1.0.0:

```
 Bay> dir fl
 Device: flash
 ------

 File 1: rel.13.0.st on Fri Sep  4 16:15:46 PDT 1998
 Version:
 /export/home2/projects/first/rel1.3/rel1.3.0/main/hw/acc1.3.0.st on
 Fri Sep  4 16:15:46 PDT 1998
 Length: 1361962 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0x4065
 File 2: /export/acura3/projects/first/rel1.0/rel1.0.0/main/hw/
 accelar.st
 Version: rel1.0/rel1.0.0/main on Fri Jan  9 11:28:28 PST
```

```
Length: 895550 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x6e79

File 3: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 41,0 free: 214,256
```

5. **Back up or save the current configuration file.**

   From version 1.1.6:

   The version 1.1.6 configuration file can be used by the version 1.3.0 run-time image, but the version 1.1.6 run-time image will not read version 1.3.0 configuration files. For this reason, you should copy the current version 1.1.6 configuration file to flash in case you must revert to the version 1.1.6 run-time image.

   Before backing up the current configuration, change the primary boot choice to flash:<acl.1.6>, which is the version 1.1.6 run-time image. This selection will then be saved in the configuration you copy to the flash file system.

   ```
   Bay> sys set primary flash:<accl.1.6>
   Bay> save
   Bay> copy config flash
   programming flash:4 ... 4376 bytes

   Bay>
   ```

   From version 1.0.0:

   The version 1.0.0 configuration file can be used by the version 1.3.0 run-time image. The version 1.3.0. configuration cannot be run by the version 1.0.0 run-time image. For this reason, save the version 1.0.0 configuration in the NVRAM before upgrading the switch to version 1.3.0.

   ```
   Bay>save
   ```

6. **Set the primary boot choice to v1.3.0 runtime.**

   Set the primary boot choice to the first file in the flash file system; that is, the version 1.3.0 run-time image.

   From version 1.1.6:

   ```
   Bay> sys set primary flash:1
   Bay> save
   Bay>
   ```

From version 1.0.0:

```
Bay> sys set primary flash:1
Bay> save
Bay>
```

7. **Copy the version 1.1.1 boot monitor updater to flash memory.**

   The version 1.1.1 boot monitor updater is a required file should a downgrade to version 1.1.6 be necessary. *You cannot downgrade to version 1.0.0.*

```
Bay> copy tftp flash
Enter source tftp server address [10.10.20.100]: 10.10.20.100
Enter source file [acc1.3.0.b23]: accboot1.3.0.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 tftp ip <tftp_server> file <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.

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 1.3.0 run-time 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 run-time 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 tftp ip 10.10.20.100 accboot1.3.0
tftp result: success
```

If you invoked the `boot tftp` command from a Telnet session, your connection will be dropped at this point. When the upgrade process completes in a few seconds, the switch will reboot into the new run-time image, retaining your previous switch configuration. When this reboot is complete, typically in a few minutes, you should be able to re-establish your Telnet session with the Accelar switch.

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

Upgrade from version 1.1.6:

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

File ACCBOOT.ROM found in loaded image

Boot Monitor Version v1.3.0

Image size: text 132988 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
Accelar Monitor v1.3.0

CPU: 60Mhz PPC 403GCX Type 2 Rev 1 in Slot 5
DRAM: 16M
Chassis: 1200
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 1 is occupied
Slot 7 is occupied
Power supply 2 not present

Press any key to stop autoboot
```

After the boot monitor is upgraded, the routing switch will automatically reboot and load the version 1.3.0 run-time image.

```
Accelar Monitor v1.3.0

CPU: 60Mhz PPC 403GCX Type 2 Rev 1 in Slot 5
DRAM: 16M
Chassis: 1200
Reason for Last Reset: PWR/BUTTON/SOFT
Slot 1 is occupied
Slot 7 is occupied
Power supply 2 not present

running boot script at 0x8000f004

monitor> boot flash:1
Booting from [flash:1] on-board flash memory
Configuration from [nvram] ...
Unzipping file acc1.3.0
Details /export/home2/projects/first/rel1.3/rel1.3.0/main/hw/
acc1.3.0.b23.st on Fri Sep  4 16:15:46 PDT 1998
from 0x761000b0 to 0x80010000 1361962 to 7257934 bytes
Attaching network interface nicEvb0... done.
Null 0.0.0.0 inet address for interface nicEvb.
Attaching network interface lo0... done.

Initializing card in slot #1 ... OK
Initializing card in slot #7 ... OK
[000 00:00:16:850] INFO: Code=0x0 Task=rcStart: System is ready

Accelar System Software Release 1.3.0
Copyright (c) 1996-1998 Bay Networks, Inc.

[000 00:00:02:383] INFO: Code=0x0 Task=rcStart: System boot
[000 00:00:02:416] INFO: Code=0x0 Task=rcStart: Accelar System
Software Release 1.3.0
[000 00:00:02:450] INFO: Code=0x0 Task=rcStart: System log file
flash:system log file:0:43

*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1200                *
* Software Release 1.3.0      *
*******************************
Login: rwa
Password: ***
```

Upgrade from version 1.0.0:

```
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/accboot1.3.0 on Tue Mar  3 16:
Version /export/acura3/projects/first/rel1.3/rel1.3.0
/boot/403_EVB/openbios/flash/accboot1.3.0 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 v1.3.0

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 will automatically reboot and load the version 1.3.0 run-time image.

```
Accelar Monitor v1.3.0

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 acc1.3.0
Details /export/home2/projects/first/rel1.3/rel1.3.0/main/hw/
acc1.3.0.b23.st on Fri Sep  4 16:15:46 PDT 1998
from 0x761200b0 to 0x80010000 1361962 to 7257934 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 1.3.0
Copyright (c) 1996-1998 Bay Networks, Inc.

[000 00:00:02:383] INFO: Code=0x0 Task=rcStart: System boot
[000 00:00:02:416] INFO: Code=0x0 Task=rcStart: Accelar System
Software Release 1.3.0
[000 00:00:02:450] INFO: Code=0x0 Task=rcStart: System log file
flash:system log file:0:43

** Using configuration from nvram **

[000 00:00:06:316] INFO: Code=0x0 Task=rcStart: System is ready

*******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1200                *
* Software Release 1.3.0      *
*******************************
Login: ***
Password: ***
```

9. Verify routing switch operation.

After performing the boot monitor upgrade, the Accelar switch should reboot into the version 1.3.0 run-time image (flash:1) and use the configuration stored in NVRAM.

At this point, you should log in to your switch through either the local console or Telnet and verify that your switch is configured and operating properly in the new 1.3.0 environment.

The version 1.3.0 run-time image has five levels of access (Read-Only, Layer 2 Read-Write, Layer 3 Read-Write, 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.

```
******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.3.0      *
******************************
Login: rw
Password: **

Accelar-1200#
```

If you upgraded from version 1.1.6 and the routing switch is not operating properly, use the version 1.1.6 downgrade procedure described below. **You cannot downgrade to version 1.0.0.** In case of problems, call Bay Networks Technical Support Center.

Note: You must be at the routing switch console to downgrade the software to version 1.1.6.

For version 1.0.0 upgrades, perform step 10 to convert the file system.

For version 1.1.6 upgrades, there is no difference in the format of the v1.1.6 & v1.3. x file system. Proceed to step 11.

10. **Convert to version 1.3.0 file system (1.0.0 upgrades only).**

Note: Before completing this step, verify that the routing switch is operating properly with the version 1.3.0 software. After the file system is converted to the version 1.3.0 format, the version 1.0 run-time image will not load. The procedure to get a version 1.0 run-time image on a routing switch with a version 1.3.0 file system requires local console access and that the flash memory be completely erased.

The flash file system has been enhanced in version 1.3.0, so the existing files must be upgraded to the new format. The version 1.3.0 firmware supports both file numbers and file names in the flash file system, whereas version 1.0 only supported file numbers. In the sample directory below, the N flag indicates that the file has been updated to the new file naming format.

```
Bay> dir fl
Device: flash
------

File 1: rel.13.0.st on Fri Sep  4 16:15:46 PDT 1998
Version:
/export/home2/projects/first/rel1.3/rel1.3.0/main/hw/acc1.3.0.b23.st on
Fri Sep  4 16:15:46 PDT 1998
Length: 1361962 EntryPoint: 0x10000 Flags: XZ (0x700) CRC: 0x4065

File 2: /export/acura3/projects/first/rel1.0/rel1.0.0/main/hw/
accelar.st
Version: rel1.0/rel1.0.0/main on Fri Jan  9 11:28:28 PST
Length: 895550 EntryPoint: 0x10000 Flags: XZ (0x300) CRC: 0x6e79

File 3: system log file
Version: 1.0
Length: 130896 EntryPoint: 0x0 Flags: L (0x102) CRC: 0x0
blocks: 256 used: 41,0 free: 214,256
```

To convert all files in flash memory (or PCMCIA) to the version 1.3.0 file system format, use the recover command.

```
Accelar-1200# recover flash
recovering files ... success
Accelar-1200# directory flash
Device: flash
  FN NameFlags Length
  -- ---------------
  1acc1.3.0XZN944692
  2accelar.st XZN 895483
  3syslog LN 130896
  4configCN4376
  --    ------
  4   files  bytes used=  1047447 free=2097152

  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 an 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 1.1.6

The downgrade procedure from version 1.3.0 to version 1.1.6 is described below.

To downgrade to version 1.1.6:

1. **Log in to version 1.3.0 Run-Time CLI.**

   The version 1.3.0 run-time CLI has five levels of access (Read-Only, Layer 2 Read-Write, Layer 3 Read-Write, Read-Write and Read-Write-All). Both a user name and password must be specified. The default user name and password for Read-Write access is rw.

```
******************************
* Bay Networks,Inc.           *
* Copyright (c) 1996-1998     *
* All Rights Reserved         *
* Accelar 1100                *
* Software Release 1.3.0      *
******************************

Login: rw
Password: **
Accelar-1200#
```

2. **Copy the version 1.1.6 configuration to NVRAM.**

   During the upgrade procedure, the version 1.1.6 configuration with primary boot choice of flash:<accl.1.6> was saved to flash memory. You can use the directory command to locate the file:

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

In the directory listing above, the saved configuration is flash:4. Copy the configuration to NVRAM:

```
Accelar-1200# copy flash:4 nvram
nvram configuration updated
Accelar-1200#
```

3. **Boot the version 1.1.6 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 1.1.6 boot monitor updater was copied to flash memory. In the directory listing above, it is flash:5. Boot flash:5 to downgrade to the version 1.1.6 boot monitor.

```
Accelar-1200# boot flash:5

Accelar Monitor v1.1.6

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 flash:5
Booting from [flash:5] on-board flash memory
Configuration from [nvram] ...
Unzipping file /export/acura3/projects/first/rel1.1/rel1.1.6
/boot/403_EVB/openbios/flash/flshbo
Details rel1.1/rel1.1.6/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 v1.1.6

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
```

4. **Confirm routing switch operation.**

   After rebooting the routing switch, you will be running the same configuration and run-time software run prior to performing the version 1.3.0 upgrade. The flash file system will have additional files, and the primary boot choice may be different, but the switch will be operating the same way prior to upgrading the software.

Installing Accelar Management Software

Release 1.3 of the Accelar Management Software includes version 1.3 of Device Manager (DM) and VLAN Manager (VM) software. An enhancement in this release is that when Device Manager or VLAN Manager is launched, it automatically determines which version of software the selected device is running and launches the appropriate version of management software. This feature eliminates the need to maintain past versions of DM/VM along with the current version.

The following sections are included:

System Requirements

Before installing Accelar Device Manager and Accelar VLAN Manager, note the system requirements listed in Table 4.

Table 4. System Requirements

Operating System	Workstation	DRAM Required	Space Required on Hard Drive	No. of Devices Supported Simultaneously VLAN Manager
Solaris 2.4 or higher	Any SPARC workstation	32 MB	12 MB	100
HP-UX 10.20	HP	32 MB	12 MB	100
Windows NT 3.1, 4.x, Windows 95 or 98	75 MHz Pentium (minimum) or 100 MHz 486	16 MB	8 MB	50
AIX 4.1 or higher	IBM	32 MB	12 MB	100

Installing on UNIX Systems

This section provides instructions for installing the Accelar Device Manager and Accelar VLAN Manager software for AIX, HP-UX, and Solaris UNIX platforms.The standard installation procedure involves copying the compressed tar file with the DM/VM software from the CD-ROM onto the hard disk. The file is then uncompressed and untared. Finally, add the “DMPATH” environment variable to the shell file.

If you are running HP OpenView (without Optivity) and want to launch DM/VM software directly from the HP OpenView windows, execute an additional script (refer to “Integrating into HP OpenView” on page 49).

Installing DM/VM Software

To install Accelar Management Software on Solaris, HP-UX, and AIX platforms:

1. **Insert the *Accelar 1000 Series Software* 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 systems administration guide specific to your operating system.**

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

   This location can be anywhere you choose. Typically, third-party software is installed under the /opt directory (for example, /opt/bay 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/bay.

   ```
   % mkdir /opt/bay
   ```

4. **Go to the directory on the CD-ROM where the DM/VM software is located:**

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

5. **Execute the install_dm_vm script by providing the version of the software and the name of the target directory:**

   ```
   % install_dmvm dm_v130 /opt/bay
   ```

6. **Add the DMPATH variable to your shell file. If you use the C shell, add the following to your .cshrc:**

   ```
   % setenv DMPATH= /opt/bay/dm_v
   ```

   where `v` is the version number of the DM/VM software you just installed.

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

   ```
   $ DMPATH /opt/bay/dm_v; export DMPATH
   ```

   where `v` is the version number of the DM/VM software you just installed.

7. **For launching Accelar Device Manager (dm) to manage an operational switch:**

   ```
   % dm a.b.c.d
   ```

   where `a.b.c.d` is the IP address of the Accelar device. For example, `% dm 132.32.1.5` for the Accelar device with IP address 132.32.1.5.

Integrating into HP OpenView

Before integrating the Accelar Management Software into HP OpenView (HPOV), make sure that the following prerequisites are satisfied:

- HPOV must be installed.
- Optivity® network management software should *not* be installed. If Optivity is installed, follow the Optivity installation procedures to upgrade Optivity software with the newer Accelar Management Software.

Note: Using the following instructions with Optivity installed may cause Optivity to have uninstall problems.

- Accelar Management Software must be installed.
- You must be logged on as root.

To integrate Accelar Management Software into HPOV:

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

   ```
   % cd /opt/bay/dm_1.x.x
   ```

   where 1.x.x is the latest version of the software.

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 Help from the menu bar and select the Run HP OpenView Install on Console option.**

   This option will copy the Accelar-specific icons and MIB files to the standard HPOV 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 5 contains a brief description of each command.

Table 5. Command Menu

You Can Run...	To...
`ovw -fields`	Update the fields database.
`ovstop netmon`	Stop the network monitor to allow the database to update.
`ovtopofix`	Update the topology database.
`ovstart netmon`	Restart the network monitor.
`ovw`	Start HP OpenView.

4. **After HPOV is started, load the Accelar-specific MIBs by selecting Options->Load MIBs. Load the rapid_city.mib file located in the /var/opt/OV/share/snmp_mibs/Vendor/Bay Networks/Accelar directory.**

Installing on Windows Platforms

To install Accelar Management Software in a Windows environment:

1. **Start Microsoft Windows.**

Note: If Windows is already running on your system, exit all applications running on Windows before starting the Accelar Management Software setup program.

2. **Insert the *Accelar 1000 Series Software* CD into the CD-ROM drive.**

   For the purposes of these instructions, the CD-ROM drive is drive D.

3. **Click on the Windows Start button and click on Run.**

   The Run dialog box opens.

4. **In the Open field, type `d:\dm-windows\dm_130.exe` and click on OK.**

   A welcome screen opens for the Accelar Management Software setup program.

5. **Follow the instructions appearing on the screen to complete the installation.**
6. **To open the Accelar Management Software, click on the Windows Start button and then select Programs->Accelar Management Software.**
7. **Select Device Manager or VLAN Manager.**