Release Notes for the Accelar 1000 Series Products

Software Release 2.0

Accelar 1050/1051 Routing Switch
Accelar 1100/1150 Routing Switch
Accelar 1200/1250 Routing Switch
Accelar Boot Monitor Software Version 2.0
Accelar Run-Time Software Version 2.0
Accelar Device Manager Version 2.0
Accelar VLAN Manager Version 2.0

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 896-00181-E
March 1999

NORTEL NETWORKS™

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Contents

Introduction 1
Related Publications 2
Recommendations and Information About Release 2.0 3
New Features in Release 2.0 4
ARU3 Hardware Support 4
Enforced Operational Configuration (eoc) Mode 5
CLI Enhancements 6
IPX Routing 7
IP Multicasting 8
DVMRP 8
Layer 3 IGMP 9
UDP Broadcast Forwarding 10
RARP Support 11
OSPF Enhancements 11
Not So Stubby Area (NSSA) Support 12
MD5 Authentication 12
Differentiated Service 12
VLANs 13
Source MAC Address-Based VLANs 13
Brouter Support 14
Other Enhancements 14
Network Advertisement When Down 14
VRRP Enhancement 15
IP Prefix Flow Filter Enhancements 15
Unknown MAC Discard 15
Mirroring Egress Traffic 15
Default Route to Non-Directly Connected Subnet and Supernet Addresses 16
SSF Enhancements 16
Latest CLI Changes 16

Bugs Fixed in Release 2.0 17
Known Software Limitations or Differences 17
Traffic Prioritization 18
VRRP 18
IP Multicasting 18
Accelar Spanning Tree 19
VLANs 19
Mirroring 20
OSPF 20
Counters 20
Filters 21
RIP2 Next Hop Support 22
Router and ICMP Redirect Messages 22
Static Routes 23
Routing Broadcast Format 23
Network Management Functionality and Limitations 23
Port Statistics Support 23
SNMP Traps 25
Device Manager Limitations 26
RMON Counter Support 26
Solaris and HP/UX Platforms 26
Windows 95, Windows 98, and Windows NT 27
Known Problems in Release 2.0 29
General 29
Device Manager 30
Multi-Link Trunking 31
OSPF 31
IP Routing 31
IPX Routing 32
VLANs 32
IP Filters 33
IP Multicasting 33

Introduction

These release notes provide the latest information about the Bay Networks® Accelar™ Software Release 2.0. For procedures to install the software, refer to *Upgrading to Accelar 2.0 Software* (Bay Networks part number 206077-A) on the Documentation CD. Currently released components include:

- Boot Monitor Software Version 2.0 (accboot2.0.0)
- Run-Time Software Version 2.0 (acc2.0.0)
- Device Manager Version 2.0 (for Windows 95/98/NT: dm_200.exe; for UNIX: dm_2.0.0.tar.Z)
- VLAN Manager Version 2.0 (for Windows 95/98/NT: dm_200.exe; for UNIX: dm_2.0.0.tar.Z)

These release notes contain the following major sections:

- Related publications (page 2)
- Recommendations and general information about this release (page 3)
- Descriptions of the new and enhanced features of Accelar software (page 4)
- Latest CLI changes (page 16)
- Bugs fixed in release 2.0 (page 17)
- Known software limitations or differences at the time of release (page 17)
- Known problems in release 2.0 (page 29)

Note: Many of the new features in release 2.0 require modules and chassis (Accelar 1100/1150 routing switches) to be -B versions or above with ASICs that are ARU3 or above. Hardware with ARU1 or ARU2 ASICs does not support these features.

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

Reference for the Accelar Management Software (Bay Networks part number 893-01052-C) has been discontinued for this release. The information previously in that document, as well as new information pertinent to this release, has been divided into the following three documents:

- *Networking Concepts for the Accelar Series 1000 Routing Switch* (Bay Networks part number 205588-A) provides general information and descriptions of how the Accelar routing switch handles various networking features, such as VLANs, Multi-Link Trunking, OSPF, RIP, and IPX.
- *Reference for the Accelar Management Software Switching Operations* (Bay Networks part number 205586-A) describes how to use Device Manager to configure and manage layer 2 (switching) functions in the Accelar switch.
- *Reference for the Accelar Management Software Routing Operations* (Bay Networks part number 205587-A) describes how to use Device Manager to configure and manage layer 3 (routing) functions in the Accelar switch.

Refer also to the following documents on the Accelar documentation CD:

- *Reference for the Accelar 1000 Series Command Line Interface Release 2.0* (Bay Networks part number 202086-B)
- *Installing the Accelar 1000 Series Chassis* (Bay Networks part number 893-01051-D)
- *Using the Accelar 1050/1051 Routing Switch* (Bay Networks part number 201603-C)
- *Using the Accelar 1100/1150 Routing Switch* (Bay Networks part number 893-01050-C)
- *Using the Accelar 1200/1250 Routing Switch* (Bay Networks part number 893-01049-C)
- *Upgrading to Accelar 2.0 Software* (Bay Networks part number 206077-A)

Recommendations and Information About Release 2.0

Note the following recommendations and miscellaneous information about Accelar software release 2.0:

- The new XLR1298SF SSF module has 32 megabytes (MB) of dynamic random access memory (DRAM). Although release 2.0 does not require 32 MB of DRAM, if you will be using RMON or are in a large OSPF routing environment and your switch SSF module is an XLR1297SF with only 16 MB of DRAM, you should upgrade your SSF module to increase memory size to improve performance. A memory upgrade kit (AA0011017) is available for the XLR1297SF to increase DRAM to 32 MB.
- If you have configured IPX parameters, but are not currently using IPX routing, reducing IPX parameters will conserve switch memory. Set IPX maximum route, sap, static route, and static sap entries to the minimum values using the following CLI commands:
  - `config ipx set max-route 128`
  - `config ipx set max-route 64`
  - `config ipx set max-route 64`
  - `config ipx set max-route 32`

  After resetting the parameters, save the configuration and reboot the switch. Refer to "IPX Routing" on page 7 for more information about the IPX protocol.
- Always set a specific enforced operational configuration (eoc) mode (refer to "Enforced Operational Configuration (eoc) Mode" on page 5), instead of allowing the default eoc mode (which is to the lowest level module in the switch) to avoid losing functionality in case a lower revision module is installed in the switch.
- Terminology in this release has been modified in Device Manager and the CLI so that "trunk" is used only in reference to Multi-Link Trunking. What were previously referred to as *trunk ports* (in contrast to access ports) are now referred to as *tagged ports*.

New Features in Release 2.0

Release 2.0 of the Accelar Management Software includes the following new features:

- ARU3 hardware support (page 4)
- CLI enhancements (page 6)
- IPX Routing routing (page 8)
- IP multicasting (page 8)
- UDP broadcast forwarding (page 10)
- RARP support (page 11)
- OSPF enhancements (page 11)
- Differentiated Service (page 12)
- VLANs (page 13)
- Other enhancements (page 14)

ARU3 Hardware Support

Release 2.0 of the Accelar Management Software supports ARU1, ARU2 (-A hardware), and ARU3 (-B hardware) ASICs. It also supports a mixed environment that includes multiple hardware levels. Certain features require ARU3 hardware. All operating modules in a chassis must be -B modules in order to use these features.

Accelar 2.0 software supports all -B revisions of existing modules, as well as a new module, the XLR1208-FL-B 8-port 10BASE-FL module.

Note: In a chassis with mixed version modules and eoc mode set to the default, the features of the lowest common denominator are in effect.

In Device Manager, fields pertaining to unsupported features are grayed out. In the CLI, attempts to use these features result in an “incompatible hardware” message.

Enforced Operational Configuration (eoc) Mode

By default, the switch operates in the mode of the lowest version ASIC present in any module at boot time. If you replace a module with a lower version and reboot the switch, the entire switch will operate with the functionality of the lower version. However a new feature with release 2.0 is enforced operational configuration (eoc) mode. This mode allows you to lock in a mode of operation. Then, if a lower version module is inserted, the module will not be initialized and error messages will indicate that the module is not operable.

Note: If an ARU2 module is inserted while the switch is in ARU3 mode, the module is disabled, but the warning message indicates a consistency check failure instead of that the board is not operational.

Bay Networks highly recommends that you set eoc mode to the highest level of hardware (ARU1, ARU2 or ARU3) in the chassis, so that you do not inadvertently cause downward revision of your switch by installing a lower revision module and losing any hardware-dependent functionality or configurations.

Warning: Be careful not to set eoc mode to a version greater than the switch can support or problems may occur.

To set eoc mode in Device Manager, select the chassis and choose Edit > Chassis > Chassis. Select EocModeAdminStatus as default, aru1Quid4, aru2Quid4 or aru3Quid5. In the CLI, use the command:

```
config sys set eoc-mode <eocmode> {default|aru1quid4|aru
2quid4|aru3quid5}.
```

After saving this configuration, it is necessary to reboot the switch to change the current ARU mode to a different mode than what is currently running. If you are enforcing the current mode, rebooting is not required. Use the show sys info command to determine the current mode.

For ASICs below Quid4, use the equivalent Quid4 setting (ARU1 or AUR2).

Note: Some features require specific hardware versions: -A (ARU2) or -B (ARU3). If a version -A or lower module is installed in the switch, in order to utilize a feature requiring ARU3, you must remove the module or set eoc status to aru3Quid5 and reboot, which enables ARU3 features but leaves the lower version module inoperable.

Use the command line interface (CLI) or Device Manager to determine ARU status. In Device Manager, choose Edit Chassis > Chassis and view EocModeOperStatus.

In the CLI, use the command `show sys info`. This example shows a partial display:

```
Accelar-1100# show sys info

       General Info  :
       SysName       : Accelar-1100
       SysUpTime     : 8 day(s), 23:59:21
       SysContact    : support@baynetworks.com
       SysLocation   : 4401 Great America Pkwy, Santa Clara CA 95052

Chassis Info :
        Chassis : 1100
        Serial# : 0008
        HwRev   : v3.0
        NumSlots: 3
        AruMode : AruTwo
        EocMode : default
```

CLI Enhancements

The CLI has been further revised and expanded in release 2.0 with more than 200 new commands and tabular output for show commands. You can also copy a running config file or copy a file to a running configuration.

New monitor commands are essentially self-updating show commands. These commands have the following subsets:

- monitor mlt error commands
- monitor mlt stats commands
- monitor port error commands
- monitor port stats commands

CLI functionality is now nearly as complete as Device Manager, with the exception of configuring RMON. Refer to *Reference for the Accelar 1000 Series Command Line Interface Release 2.0* for a complete list of commands and definitions of parameters. The complete command list in Appendix A of that document indicates which commands are new to this release.

IPX Routing

The Internetwork Packet Exchange (IPX) Protocol is the Novell, Inc. adaptation of the Xerox Network System (XNS) Protocol. IPX is the network layer routing protocol used in the Novell NetWare environment. The primary tasks of IPX are addressing, routing, and switching information packets from one location to another on a network. The network interface card (NIC) in a client provides network node addressing.

The Accelar implementation of IPX supports four Ethernet frame formats:

- Ethernet II
- 802.2-LLC
- 802.3-RAW
- 802.3-SNAP

Frame translations from one frame format to another frame format are supported. However, the IPX host format must match the defined protocol frame format of the interface to which the host is connected. In Accelar switches, IPX is supported on routed VLANs. IPX routing is not available on isolated routing ports.

You cannot configure the IPX routing preference (by ticks then hop count) in this release. If two paths exist, the best route is the one with the fewest ticks (1/18th of a second). If the number of ticks is equal, the best route is the one with fewer hops. A hop is counted each time the packet passes through a router.

Note: IPX requires ARU3 hardware (-B modules).

An Accelar router running IPX provides the following network-layer support:

- Dynamic routing of IPX packets
- Up to four IPX network addresses to an interface for port-based VLANs, one per frame type (On IPX VLANs, only one IPX network address can be defined.)
- Routing Information Protocol (RIP) to exchange routing information
- Service Advertisement Protocol (SAP) to advertise services and addresses of service-providing nodes
- Static route support
- Default route support

To configure IPX in Device Manager, choose Routing > IPX > IPX and Routing >IPX > RIP or Routing > IPX > SAP. In the CLI, use the config ipx commands.

Note: When not using IPX, reducing IPX parameters will conserve memory. See "Recommendations and Information About Release 2.0" on page 3.

IP Multicasting

Two IP multicasting protocols are supported in Accelar switches:

- Internet Group Management Protocol (IGMP) is used by hosts to report their multicast group memberships to neighbor multicast routers.
- Distance Vector Multicast Routing Protocol (DVMRP) is used between routers to exchange multicast routing information.

DVMRP

DVMRP is a distance vector type of multicast routing protocol. It advertises shortest-path routes to multicasting *source networks*—that is, any network containing hosts that have the capability to issue multicast datagrams. (In this respect, DVMRP is the opposite of RIP, which advertises routes to destination networks.) Coupled with IGMP, membership for a multicast stream is learned from both the routers and directly attached hosts.

DVMRP in Accelar switches fully supports multiaccess networks. The forwarding entries for the receivers on multiaccess networks are port based rather than network based. IP Multicast routing is supported on ports with port-based or IP subnet-based VLANs enabled.

The DVMRP router listens to all IGMP host membership reports, even if it is not the designated querier, and keeps a local group database of every host membership reporter. When a multicast stream (UDP packets) first enters the switch, if DVMRP is enabled for the interface, then DVMRP will process this packet as necessary and create a hardware cache entry to handle subsequent UDP packets for the same multicast destination. The packets are discarded if there are no members; otherwise, they are forwarded.

The Accelar 2.0 implementation does not support DVMRP tunneling.

Using Device Manager to set up DVMRP, choose Routing >IP > DVMRP. In the CLI, use the config ip dvmrp commands.

Layer 3 IGMP

The layer 3 IGMP has the following characteristics:

- It allows a host to register group memberships with the local querier router to receive any datagrams sent to this router and targeted to a group with a specific IP Multicast address.
- It allows a router to learn the existence of group members on its directly attached networks by periodically sending a general query message to each of its local networks. Any host that is a member of any multicasting group identifies itself by sending a response. The Accelar switch supports both version 1 and version 2 queries.

Note: ARU3 (-B version) mode is required to support version 2 queries.

If multiple IGMP routers exist on the network, one router is designated to send queries, using the following rules:

- Choose a router that generates version 1 queries over a router that generates version 2 queries.
- Choose the router with the lowest IP address when running version 2. In version 1, the router with the highest IP address becomes the IGMP querier.

The Accelar switch supports IGMPv1 and IGMPv2 registration protocols and will generate IGMP queries on all subnets and interfaces for which IP multicasting is enabled. Multicast frames that arrive from an interface are forwarded on all interfaces/subnets on which IGMP reports have been received for the multicast group indicated in the destination IP address.

Multicast packets forwarded within the same VLAN remain unchanged, and packets are not forwarded to networks with no members of the multicast group indicated in the destination IP address.

Multicast routing can be enabled and disabled on an interface basis. If multicast routing is disabled on an interface, IGMP queries are not generated. When used as a switch, Accelar 1000 series products support IGMPv1 and IGMPv2 to prune group membership per port within a VLAN (IGMP snooping). If the switch is in IGMP router behavior mode, IGMP snooping is not configurable.

If the Accelar switch receives multiple reports for the same multicast group, it forwards only the first report. If there is new information that another multicast group has been added or that a query has been received since the last report was transmitted upstream, then the report will be forwarded onto the multicast router ports. This feature is known as IGMP proxy.

In Device Manager, to configure layer 3 IGMP, choose Routing > IP > L3 IGMP. In the CLI, use the config ip l3 igmp commands.

UDP Broadcast Forwarding

Some network applications, such as the NetBIOS name service, rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server. By default, broadcasts are not forwarded by a router. UDP broadcast forwarding is a general mechanism for selectively forwarding UDP broadcasts received on an IP interface out other router IP interfaces as a rebroadcast or to a configured IP address.

- If the destination address is that of a server, the packet will be sent as a unicast packet to this address.
- If the destination address is that of an interface on the router, the frame will be rebroadcast.

The basic steps for setting up UDP broadcast forwarding are:

1. Enter protocols into a table.
2. Create policies (protocol/server pairs).
3. Assemble these policies into lists or profiles.
4. Apply the list to the appropriate interfaces.

To set up UDP forwarding from Device Manager, choose Routing > IP > UDP Fwd. In the CLI, use the config ip udp commands.

RARP Support

Reverse Address Resolution Protocol (RARP) is a protocol used by some devices to obtain an IP address by providing their MAC layer address information to a RARP server. In previous versions of Accelar software, RARP was broadcast along with ARP and IP on all ports associated with an IP protocol-based or port-based VLAN. Therefore, it was not possible for a host to reach a RARP server outside the IP VLAN to get its IP address.

RARP has the format of an Address Resolution Protocol (ARP) frame but its own Ethernet type (8035), which makes it possible for RARP to be removed from the IP protocol-based VLAN definition and treated as a standalone protocol. By doing so, the concept of a RARP protocol-based VLAN is created.

In Device Manager, choose VLAN > VLANs > Insert and click on byProtocolID. Then select rarp. In the CLI, use the command:

```
config vlan <vid> create byprotocol <sid> rarp [name <value>].
```

OSPF Enhancements

This release complies with Open Shortest Path First (OSPF) RFC2178. The Accelar software does not support overlapping areas. In addition, the following enhancements have been made to the OSPF protocol in this release.

Not So Stubby Area (NSSA) Support

In an OSPF network, a stub area does not receive advertisements for external routes, which reduces the size of the link state database. A stub area has only one area border router. Any packets destined outside the area are simply routed to that area border exit point, examined by the area border router, and forwarded to a destination.

A not so stubby area (NSSA) also prevents the flooding of AS-External Link State advertisements into the area by replacing them with a default route. This feature is supported in release 2.0. The added feature of NSSAs is the ability to import small stub (non-OSPF) routing domains into OSPF.

In Device Manager to create a stub area or NSSA, choose Routing > IP > OSPF > Area and under the ImportASExtern field, select the area you want to change to a stub or not so stubby area and use the pull-down menu to select importNoExternal (stub area) or importNssa (not so stubby area), and click Apply. In the CLI, use the command: `config ip ospf area <area> nssa <true|false>.`

MD5 Authentication

Previous versions of Accelar supported only no authentication or simple password authentication. Release 2.0 also supports message digest (MD5) encrypted authentication.

MD5 authentication can be configured only with the CLI, not with Device Manager. To set up MD5 authentication for an OSPF virtual link or interface, use the CLI commands:

- `config ip ospf area <area> virtual-interface <nbr> add-message-digest-key <md5-key-id> md5-key <value>`
- `config ip ospf interface <ipaddr> add-message-digest-key <md5-key-id> md5-key <value>`

Differentiated Service

Differentiated Service as defined in RFCs 2474 and 2475 provides an architecture for scalable service differentiation in the Internet. The Differentiated Services (DiffServ) specification defines a *codepoint*, which is a 6-bit value that has historically been known as the 8-bit Type of Service (TOS) field in an IP protocol header. Within the DiffServ architecture, setting this codepoint provides a means of delivering a differentiated or better class of service for these IP packets.

Accelar 2.0 software provides the capability of setting the DiffServ bits on an IP frame using an IP filter mechanism.The DiffServ AND rule is first applied to the 8-bit field and acts as a mask. This value can be used to protect or mask bits that may already be set. The DiffServ OR rules provide three values that can be used to set the DiffServ bits. The chosen rule will be logically ORed with the intermediate result after the original ANDing. The final result will be set as the new DiffServ codepoint in the IP header of the filtered frame.

In Device Manager, choose Routing > IP > IP DiffServ to set the decimal values that will be used in an IP protocol filter to set the DiffServ bits. Select the OR rule to be applied in the bottom field of the IP Filter window accessed by choosing Routing > IP > IP Filter > Filters > Insert. In the CLI, use the config ip diffserv-rule commands.

Note: Differentiated Service requires ARU3 hardware (-B modules).

VLANs

The following paragraphs describe VLAN enhancements in the Accelar switch. VLAN default behavior has been changed in that a VLAN interface is up if one of the following is true:

- Any port in a port-based VLAN is up.
- A port is active in a policy-based VLAN.

Source MAC Address-Based VLANs

In release 2.0, you can create policy-based VLANs by source MAC address. As with all policy-based VLANs, using source MAC address VLANs allows the Accelar routing switch to associate frames with a VLAN based on the frame content. With source MAC-based VLANs, a frame is associated with a VLAN if the source MAC address is one of the MAC addresses explicitly associated with the VLAN by adding it to a list of MAC addresses that comprise the VLAN. However, because it is necessary to explicitly associate MAC addresses with a source MAC-based VLAN, the administrative overhead can be quite high.

Source MAC-based VLANs are used in situations where users want to enforce a MAC level security scheme to differentiate groups of users.

Note: Routing on a source MAC address-based VLAN is not supported.

In Device Manager, choose VLAN > VLANs > Basic and click Insert. Then click in the SrcMac Type field. In the CLI, use the command:

```
config vlan <vid> create bysrcmac <sid> [name <value>]
```

Brouter Support

A special type of VLAN supported by the Accelar switch is a brouter port, which is actually a one-port VLAN. A brouter port differs from an isolated routing port in that it can route IP packets as well as bridge all other traffic. The only difference between a brouter port and a standard IP protocol-based VLAN configured to do routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still be able to route the IP traffic. A brouter port is actually a one-port VLAN; therefore, each configured brouter port uses a VLAN ID and decreases the number of available VLANs by one. Brouter ports are supported on single ports only; they are not supported on Multi-Link Trunks.

In Device Manager, create a brouter port by creating an IP protocol-based VLAN in spanning tree group (STG) 0 with the specified port as always a member and all other ports never members. In the CLI, use the command:

```
config ethernet <ports> ip create-brouter <ipaddr/mask> <tag-id>
```

where `tag-id` is the VLAN ID (1 to 1024) and must be a new ID number.

Other Enhancements

The following sections describe other general enhancements included in this release.

Network Advertisement When Down

You have the new option to advertise the network on an interface, even if the interface is down. When you configure an interface without any link and enable AdvertiseWhenDown, it will not advertise the route until the interface is up. Then the route will be advertised even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown. The default is disabled.

This feature is configured in Device Manager in the windows accessed by choosing Port > IP > OSPF, Port > IP > RIP, VLAN > IP > OSPF, or VLAN > IP > RIP.

In the CLI, use the following commands:

```
config ethernet <ports> ip ospf advertise-when-down <enable|disable>
config ethernet <ports> ip rip advertise-when-down <enable|disable>
config vlan <vid> ip rip advertise-when-down <enable|disable>
config vlan <vid> ip ospf advertise-when-down <enable|disable>
```

VRRP Enhancement

In Accelar version -A hardware (ARU2 ASICs), four VRRP interfaces (isolated routing ports *and* VLANs) are allowed per switch, and all virtual router IDs (VRIDs) must be unique. In version -B hardware and release 2.0, a maximum of 253 VRIDs can be configured. Note that it is still important not to duplicate VRRP ID numbers in the same broadcast subnet, even if the VRRP IP addresses are different.

IP Prefix Flow Filter Enhancements

This release includes support for IP prefix flow filters that are more or less specific than routing table entries. Using the CLI, you can display and clear a filter statistics table that displays filters and counters by filter ID. Use the show ip traffic-filter stats and the config ip traffic-filter clear commands, respectively. A port default action of drop is also supported in ARU3 hardware.

Unknown MAC Discard

This release includes layer 2 support for MAC-based security where bridged frames from any unknown MAC address are discarded. This feature is enabled in the CLI using the command:

```
config ethernet <ports> unknown-mac-discard <enable|disable>.
```

In Device Manager, choose Edit > Port > Interface window.

Mirroring Egress Traffic

In ARU1 and ARU2 hardware, routed packets are not mirrored in the egress direction. Enabling mirroring on an ARU2 port will cause IP routed traffic ingressing that port to be copied to the mirror port. In ARU3 hardware and release 2.0, IP routed traffic both ingressing and egressing a mirrored port will be copied to the mirroring port.

Default Route to Non-Directly Connected Subnet and Supernet Addresses

This release includes support for a default route on a non-directly connected subnet to simplify administration on complex routed networks. It also supports supernet address configuration on local router interfaces.

SSF Enhancements

With this release, the procedure for SSF failover and reset has been improved to reduce the failover time. In addition, the software recognizes the new XLR1298SF SSF module with 32 megabytes (MB) of dynamic random access memory (DRAM). The earlier version of the SSF module, the XLR1297SF, differs from the XLR1298SF in that it has no real-time clock and only 16 MB of DRAM. A memory upgrade kit (AA0011017) is available for the XLR1297SF to increase DRAM to 32 MB.

Latest CLI Changes

The following commands have changed since the *Reference for the Accelar 1000 Series Command Line Interface Release 2.0* was completed.

Command listed in Reference Guide	Revised Command
`ping <ipaddr> [<datasize>] [<count>] [-s] [-I <value>] [-t <value>] [-d]`	`ping <ipaddr> [datasize <value>] [count <value>] [-s] [-I <value>] [-t <value>] [-d]`
`config vlan <vid> igmp-snoop static-members <GroupAddress> add <ports> <static|blocked>`	`config vlan <vid> igmp-snoop static-members <GroupAddress> add <ports> [<static|blocked>]`
`config vlan <vid> igmp-snoop static-members <GroupAddress> create <ports> <static|blocked>`	`config vlan <vid> igmp-snoop static-members <GroupAddress> create <ports> [<static|blocked>]`
`config vlan <vid> igmp-snoop static-members <GroupAddress> remove <ports> <static|blocked>`	`config vlan <vid> igmp-snoop static-members <GroupAddress> remove <ports> [<static|blocked>]`

Bugs Fixed in Release 2.0

All bugs fixed through release 1.3.3 (unless listed under "Known Problems in Release 2.0" on page 29) are fixed in this release. If you need more information about which bugs were fixed in releases 1.3.1 through 1.3.3, refer to *Addendum to the Release Notes for the 1.3 Software Release for Accelar 1000 Series Products* (Bay Networks part numbers 204767-A through D). These documents are 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 release 1.3.3:

- Connectivity failures no longer result if an 802.1Q tagged Multi-Link Trunk is configured for multiple spanning tree groups. (90775)
- Administratively bringing down a port with an MLT group no longer causes local ARP entries to be cleared. (96038)
- BPDUs are no longer forwarded over all links of an MLT group when spanning tree is disabled. (95421)
- Accelar routing switches and Bay Networks routers are now able to interoperate for VRRP. (91960)

Warning: VRRP implementation in release 2.0 is not compatible with VRRP in software release 1.3.3 or earlier.

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.

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.

IPX Routing

You cannot configure non-ASCII characters into a static SAP entry in this release.

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. Version 2 implementation of VRRP supports no authentication.

Warning: VRRP implementation in release 2.0 is not compatible with VRRP in software release 1.3.3 or earlier.

IP Multicasting

IP multicasting is supported on port-based and IP subnet-based VLANs. It is currently not supported on protocol-based VLANs.

When configuring IP multicasting over a Multi-Link Trunk (MLT), the traffic is not distributed over all MLT links but is transmitted over the link where a report is received. This situation may cause an interruption of traffic when that link is disconnected. The length of the disruption is determined by the timing of a report in response to a query.

Source MAC addresses of all multicast packets are not learned. However, with ARU2 hardware, source MAC addresses for untagged multicast packets and source MAC addresses for all queries are correctly learned and displayed.

In Quid5 mode with a local multicast sender and no local receiver, the multicast data will be forwarded to the querier. In Quid4 mode, you must create a static receiver on the sender port for data to be multicast routed.

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.

VLANs

The unknown MAC discard feature works only for layer 2 (bridged) traffic. The layer 3 (routed) traffic will be addressed in a later release.

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.

Mirroring

In ARU1 and ARU2 hardware, the Accelar switch supports mirroring of IP routed (layer 3) traffic only on ingress. ARU3 hardware supports mirroring of egress traffic. Mirroring of multicast and broadcast traffic and CPU originated traffic is not supported but will be available in a future release. Port mirroring on an MLT is not supported.

OSPF

The following are known limitations with OSPF:

- 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 or 2.0. If the option is set to true, it will be 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.
- Host routes can only be created in area 0.

Counters

Egress counters are inconsistent with ingress counters. The egress unicast counters miss a small percentage (7 in 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 and ARU3 hardware.

Filters

The following limitations apply to IP filtering:

- This release supports IP filtering for ARU2 and ARU3 hardware revisions, although some filtering limitations continue to apply when the switch is running in ARU2 mode (with ARU2, ARU3, or mixed ARU revisions).
- In ARU2 mode, the only port default action that is supported is “forward.” This default action causes the switch to forward (route) IP packets that ingress on a given filtered port if these packets match no filters applied to that port.
- In ARU3 mode, port default actions of forward and drop are supported. The port default action of drop causes the switch to drop (discard) any IP packets that ingress a given filtered port but do not match any IP filters applied to that port.
- Packets that ingress a filtered port and match multiple IP filters applied to the port will be forwarded or dropped, depending on the actions of the filters they match in conjunction with the default action configured for the port. When the port default action is forward (the only default supported in ARU2 mode), one matching IP filter with a drop action causes the packet to be dropped, even if additional filters with an action of forward match the packet. Similarly when the port default action is drop, one matching IP filter with an action of forward will cause the packet to be forwarded, even if additional filters with an action of drop match the packet.
- Because IP filters are only applied at ingress, the IP filter mirror action can be used to mirror only IP traffic ingressed on a particular port or set of ports. In order to mirror both directions of an IP flow, you will typically have to use multiple mirroring IP filters on multiple ports or complement a mirroring IP filter with another mirroring operation, such as port or MAC-address based mirroring.
- The IP filter packet limit feature and the TCP connect features are only applicable when the filter has an effective action mode of forward. Such a filter must either have an explicit action mode of forward or have an action mode of default and be applied to a port with a default action of forward.
- 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.
- Not more than eight Global Filters may be applied to any port or any set of ports sharing an ARU.

- A filter set may contain a maximum of 32 Source and/or Destination Filters. A Global filter set may contain a maximum of eight Global Filters.
- The minimum size mask length for a Source or Destination Filter is 8 bits (255.0.0.0). The maximum size for all filters is 32 bits (255.255.255.255).
- IP Destination Filters are ineffective in filtering traffic destined to a broadcast address.
- The IP Filter log-stats feature is not enabled in this release.
- Source and Destination Filters may specify a TCP or UDP Source Port or Destination Port to match. When this port is specified, the implied operation is always equals.
- Source and Destination Filter actions other than drop and forward (for example, mirror, change priority, or modify DiffServ value) are applied to all matching traffic ingressing any routed port, even if the port is not configured for filtering.
- 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.
- When changing the characteristics of any given applied filter, you must disable and reenable this filter before these changes take effect on any given set of ports.

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.

Routing Broadcast Format

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

Network Management Functionality and Limitations

This section describes the latest functionality and known limitations of managing an Accelar 1000 Series release 2.0 routing switch. Unless otherwise noted, the information is valid across all management platforms.

You cannot ping or telnet from the Accelar switch to another device if the device is reachable through a statically configured route with a nonlocal next hop. However, this restriction does not affect forwarding through the Accelar switch using that static route.

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/XD
InOctets	X (ARU2/3 only)	X (ARU2/3 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
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/XD
etherStatsDropEvents	Not Applicable	Not Applicable	Not Applicable
etherStatsOctets	X (ARU2/3 only)	X (ARU2/3 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 2.0 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.
rcErrorNorification	An error has occurred with error level, code, and text.
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 2.0.

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 25. 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 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, DM 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 DM 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, DM may occasionally fail when RMON alarm traps are received. When this problem occurs, the host displays a message indicating that 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 2.0

The following sections list problems known to exist in release 2.0.

General

The following are known general problems:

- You cannot enable RMON on an Accelar 1100 or Accelar 1200 routing switch with 16 MB of DRAM. (98061, 97467)
- When operating in ARU3 mode and an ARU2 module is inserted, the module is disabled as it should be, but the warning message indicates a consistency check failure rather than that the board is not operational. (96060)
- If a module is offline because of the eoc setting (that is, the module is a lower revision) and an active link is connected to the module, the port LEDs will continue to function. The module On-Line LED will be off. (96098)
- Copying an ASCII-based configuration file from TFTP to the flash file system will cause the last 169 bytes of the file to be corrupted. The effect of this corruption can be negated by padding the original ASCII configuration file with sufficient comments (each line preceded by a pound sign #) to absorb the effects of this file corruption. (99973)
- When rebooting the Accelar switch, you may see the message: WARNING: code=0x0 Task=tTimerTask. This message is a benign log file entry and can be ignored. (94322)
- When using TFTP to transmit files from SSF module to SSF module, files that are more than 10,000 bytes may generate an erroneous error message indicating that the transfer failed. Verify that the file has been transferred. (98117)
- On an Accelar 1100/1150 switch, if the eoc mode is set higher than the ASIC version in the baseboard (slot 3), the switch will fail to boot. In other words, the switch will not boot if ARU3 modules are inserted into slot 1 or 2 of an Accelar 1100/1150 switch with an ARU2 baseboard and the eoc mode is set to aru3/Quid5. (99360)
- VRRP state change counters are not incrementing. (97619)
- Syslog stops sending messages to the host if the local log file gets too full or otherwise cannot write to the flash file system. (85398)

Device Manager

The following are known problems in Device Manager:

- When configuring a static route from Device Manager using supernetted addresses, DM displays an erroneous error message. Refreshing the display removes the message and displays the static route. (98317)
- When using the Netscape 4.06 browser and Microsoft Windows 95/98/NT, the Device Manager Help screen successfully launches for the initial Help screen, but does not launch for subsequent help screens for context Help. You must exit from the current browser before clicking on a Help button to launch a context-relevant Help page. (99987)
- When using the Internet Explorer 3.0 browser and Microsoft Windows 95/98/NT, the Device Manager Help screen successfully launches for the initial Help screen, but may not launch for subsequent help screens for context Help. You can click on a given hyperlinked Help topic in the initial page to access Help information for the selected topic.
- After a failed save to NVRAM (such as with a configuration that is too large), the NVRAMUsed value indicates 0 (zero) K used. To recover from this state, perform a successful save to NVRAM or reset the switch. (85632)
- When exiting from Device Manager, DM will sometimes inform you that your configuration has changed and ask if you want to save it, even if the configuration has not been changed. (96710)
- Attempting to add significantly more than the maximum supported number of filters to a filter set (32 for Source and Destination Filters or eight for Global Filters) may cause Device Manager to crash. (99759)
- In Device Manager, attempting to specify a filter with an invalid IP address (that is, an address having a first octet larger than 255) will result in a corrupted IP address for that filter. (95634)
- The online IP Filter help screen in Device Manager 2.0 incorrectly states that “A ‘global’ filter is a filter that is applied to all packets regardless of the packet's source and destination.” The phrase “is applied” should be changed to “may be applied.” (99336)

Multi-Link Trunking

The following are known problems with MLT:

- When forwarding IP Multicast traffic across an MLT group, the links do not load share the traffic. (97562)
- A port “flapping” in an MLT group causes BPDUs to be sent out irregularly on all ports. (96527)

OSPF

The following are known problems with OSPF:

- If you delete an OSPF announce policy, it can take up to one minute for the link state update to be propagated to other routers. (98154)
- If any OSPF interface link on a switch goes down, OSPF adjacencies may be dropped on all interfaces on the switch. (98229)
- When you change an OSPF not so stubby area (NSSA) to a normal area, external routes may not be propagated. For proper operation, you should disable and then reenable OSPF on the router. (96596)
- Resetting OSPF transit routers can cause routers to crash. (97818)
- If the switch loses an intra-area route because of an interface going down, but still has an inter-area route, it does not generate a new Type 3 summary-LSA into the other attached non-backbone areas reflecting the cost of this inter-area route. It does generate a MaxAge summary-LSA for the intra-area route. (90581)

IP Routing

The following are known problems with IP routing:

- When routing, the Accelar switch does not discard datagrams with a bad destination IP address. Instead, the switch sends an ICMP destination unreachable message. (90775)

- Under specific conditions of unique traffic patterns at high frame rates, a temporary condition of high CPU utilization may occur, which could temporarily affect other switch operations. The condition can occur if a high level of routable IP packets transmitted by hosts residing on an ingress port of a local routable VLAN are destined for a remote network and the packets are layer 2 switched to a remotely attached router and then routed back through the same switch in which the client resides. (89959)

IPX Routing

The following are known problems with IPX routing:

- IPX allocates memory even if IPX routing is not enabled, which reduces the available memory for other processes. You should reduce IPX parameters when not using the IPX protocol. Refer to "Recommendations and Information About Release 2.0" on page 3. (97008)
- The pingipx command does not work correctly on tagged ports. (98499)
- The pingipx command cannot be used to ping local IPX interfaces. (99086)
- In the help commands for the config ipx set commands, the displayed ranges are incorrectly displayed as 1 to 100000. They should be as follows:
  - `config ipx set max-route <max-entries>` 128 to 8000
  - `config ipx set max-sap <max-entries>` 64 to 8000
  - `config ipx set max-static-route <max-entries>` 64 to 500
  - `config ipx set max-static-sap <max-entries>` 32 to 500

VLANs

The following are known problems with VLANs:

- You are allowed to configure DHCP relays on IP subnet-based VLANs, even though DHCP is not supported. DHCP relays operate correctly only on isolated routing ports or port-based VLANs. (99248)
- VLAN MAC addresses for IPX-routed VLANs may not be correctly displayed using the show vlan info command (CLI) or the VLAN >VLANs > Advanced screen (Device Manager). (93941)

- If a Bridge Protocol Data Unit (BPDU) with the topology change notification (TCN) bit set is received on a port with Spanning Tree Protocol enabled, the Accelar switch processes the BPDU. (99060)
- A user-defined VLAN configuration is not saved after a switch reset. (94644)

IP Filters

The following are known problems with IP filters:

- Issuing the CLI show ip traffic-filter stats command causes the internal packet counters for each active filter to be reset to zero. For filters that have the packet limit option configured, issuing this command will have the effect of resetting the counter used in evaluating whether the packet limit has been reached for this filter. (98069)
- A maximum of 1023 IP filter definitions may be defined on a running Accelar switch, but only 357 IP filter definitions may be saved to NVRAM or the flash file system in this release. To save or restore IP filter definitions beyond this limitation, Bay Networks recommends that you use the ASCII-based configuration features available in this release. (85613)
- When using the packet limit option in a Global Filter, you must issue a config ip traffic-filter show stats command after applying the filter before the packet limit functionality will take effect. (98069)

IP Multicasting

The following are known problems with IP multicasting:

- InQuid5 mode, port mirroring does not mirror multicast streams on Quid5 ASICs. (97088)
- In IGMP snoop mode, with -A hardware or mixed mode, duplicate queries may be forwarded. This problem does not affect the operation of the multicast stream. (97204)
- Manually flushing snoop members causes static receivers to be removed from internal tables. (97493)
- An IGMP snoop access list entry may not be saved in the switch binary configuration file. These entries can be restored properly using ASCII configuration. (98865)

- In Quid4 hardware configured in IGMP snoop mode, static receivers must be configured before the first query is processed or only the first group will be learned. To recover from this condition, save the configuration with the static receivers configured and reset the switch. (99984)
- In Quid4 mode, receivers may continue to receive multicast streams even if you have executed a manual flush of the multicast router. (99985)
- With Quid5 hardware and multiple snooped VLANs, when the querier ages while there is a sender without a local receiver, or if you manually flush snoop multicast routers or senders, you cannot reestablish the source/group combination on the switch. (99369, 99370)
- In Quid4 mode, Quid5 operation will not handle IGMP snoop reports properly.
- When ports are moved into an IGMP snoop enabled VLAN that has a querier and senders, the newly added ports will not receive multicast packets. The workaround is to move ports into an IGMP-snoop enabled VLAN only when there is no querier present or before IGMP snoop is enabled. (99371)
- In IGMP snooping mode, if there is a sender or static receiver with a group address that is less than the lowest group address defined on an access filter, the entire access list fails to display in Device Manager. The problem can be avoided by adding a dummy entry. The list displays correctly in the CLI. (99017)
- In IGMP snooping mode, not all receivers are always displayed in the Receiver table. (99018)
- When running in eoc mode of ARU2/Quid4, the software displays that layer 3 IGMP version is incorrectly defaulted to version 2 although the switch is running in version 1. (97216)
- In IGMP mode, if an access filter is defined on a multicast group, you cannot add a static receiver defined on the same group. The workaround is to delete the access filter entry with the same group address, insert the static receiver, and then reinsert the access filter entry. (97499)
- In DVMRP multicast mode, the router may intermittently continue forwarding UDP packets even if IGMP membership expires. (99263)
- In DVMRP mode, with multiple routers on one leaf network, it is possible that both routers may forward multiple streams. (98678)

- If the Accelar switch sends a prune (DVMRP message to stop sending) and the remote switch does not receive the prune, the multicast stream will continue. The switch will correctly discard the stream, but the prune message will not be reissued. (97095)
- In DVMRP routing, with multiple multicast applications, some requested applications may not be forwarded to multicast clients if clients or servers are moved around within the network, or if the network is reconfigured. (99637)