

Part No. 206494-H
December 1999

4401 Great America Parkway
Santa Clara, CA 95054

Addendum to the Release Notes for the 2.0 Software Release for Accelar 1000 Series Products Software Release 2.0.5.5



NORTEL
NETWORKS™

Copyright © 1999 Nortel Networks

All rights reserved. Printed in the USA. December 1999.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks NA Inc.

Trademarks

NORTEL NETWORKS is a trademark of Nortel Networks.

Bay Networks is a registered trademark, and Accelar, LinkSafe, and Nortel Networks are trademarks of Nortel Networks.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks NA Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks NA Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Introduction

This release note addendum for Accelar™ software release 2.0.5.5 describes the enhancements and bug fixes to the Accelar software that have been implemented in release 2.0.5.5. This document is an addendum to the *Release Notes for the Accelar 1000 Series Products Software Release 2.0* (Bay Networks® part number 896-00181-E). The 2.0 release notes and addendums are available on the 2.0 Software CD and on the Nortel Networks Customer Service Documentation Web page (<http://support.baynetworks.com/library/tpubs/nav/rtswitch/accelar.htm>).

Software release 2.0.5.5 includes updates to the run-time software only. The latest software components are:

- Run-Time Software Version 2.0.5.5 (ac1a2055.img)
- Boot Monitor Software Version 2.0.5 (ac10b205.img) supplied as a Boot Monitor Updater
- Device Manager and VLAN Manager Version 2.0.5 (for Microsoft® Windows® 95 or Windows 98 and Windows NT®: dm_205.exe; for UNIX: dm_2.0.5.tar.Z)



Note: Boot Monitor Software Version 2.0.5 is equivalent to Boot Monitor Software Version 2.0.1. Existing configurations with Boot Monitor Software Version 2.0.1 can continue to use this boot monitor with the Run-Time Software Version 2.0.5. Configurations with boot monitor software versions prior to 2.0.1 must upgrade to Boot Monitor Software Version 2.0.5.



Note: Before upgrading your software from either version 2.0.4 or earlier, back up your current configuration file. Version 2.0.5.5 configuration files contain configuration options that are not compatible with version 2.0.4 or earlier run-time options. It is important to back up the current configuration file before upgrading, in case you must revert to a previous version of the run-time image.

For the latest information about software issues, always refer to the Accelar Products site from the Nortel Networks™ Web page (www.nortelnetworks.com) or contact Nortel Networks Customer Support at 1-800-2LANWAN.

This addendum includes the following sections:

- [Recommendations and Information About Release 2.0.5.5 \(page 3\)](#)
- [Multicast Limitations in Release 2.0.5.5 \(page 4\)](#)
- [New Features in Release 2.0.5.5 \(page 5\)](#)
- [STG and BPDU Clarification \(page 7\)](#)
- [High-Priority Switching \(page 8\)](#)
- [Bugs Fixed in Release 2.0.5.5 \(page 8\)](#)
- [Known Issues \(page 16\)](#)
- [Related Publications \(page 18\)](#)



Note: Many of the new features in release 2.0 and above 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.



Warning: Software release 2.0.5.5 requires 32 MB of DRAM. If you do not have 32 MB of DRAM, an error message appears when you boot up the Accelar switch.

The memory upgrade kit (AA0011017) is available for the XLR1297SF module and increases DRAM to 32 MB. If your Accelar 105x or 11x0 Routing Switch has 16 MB of DRAM, contact your Nortel Networks sales representative or authorized reseller to upgrade your switch.

Recommendations and Information About Release 2.0.5.5

Note the following recommendations and miscellaneous information about Accelar software release 2.0.5.5:

- Accelar software release 2.0.5.5 does not support global filters. Configuration information relating to global filters is ignored on boot-up when you use software release 2.0.5.5. Upon booting up with software version 2.0.5.5, the following message appears on the screen:

```
Global filters are not supported in this release.
```

If you attempt to configure global filters using software version 2.0.5.5, the following error message appears on the screen:

```
Operation not allowed.
```

- DVMRP support requires chassis and modules that are -B versions. The -A version chassis and modules do support IGMP snooping.
- When you create a Multi-Link Trunking (MLT) group, the resulting MLT is put into the default VLAN (VLAN 1). The MLT should then be assigned to other VLANs as appropriate.
- The new XLR1298SF SSF module has 32 megabytes (MB) of dynamic random access memory (DRAM). Release 2.0.5.5 requires 32 MB of DRAM, so you must upgrade your XLR1297SF module to increase memory. If you do not have 32 MB of DRAM, an error message appears on boot-up. A memory upgrade kit (AA0011017) is available for the XLR1297SF module to increase DRAM to 32 MB.
- Always set a specific enforced operational configuration (eoc) mode (refer to the Accelar software release 2.0 release notes for more information) instead of allowing the default eoc mode (which is to the lowest-level module in the switch) in order to avoid losing functionality in case a lower-revision module is installed in the switch.
- Terminology has been modified in Device Manager and the command line interface (CLI) so that “trunk” is used only in reference to Multi-Link Trunking (MLT). What were previously referred to as *trunk ports* (in contrast to access ports) are now referred to as *tagged ports*.

- Gigabit LinkSafe™ configurations must have autonegotiation enabled. Setting autonegotiation to False is not supported on Gigabit LinkSafe modules in *redundant* configurations. However, autonegotiation can be set to False if a Gigabit LinkSafe module is connected in a nonredundant setup to a Gigabit module not supporting autonegotiation.
- Nortel Networks recommends against configuring VRRP on IP-subnet-based VLANs as there is no hardware support for this configuration in the I/O modules and all traffic forwarding must be handled by the CPU. This situation can cause high CPU utilization and affect performance. (105851)

Multicast Limitations in Release 2.0.5.5

DVMRP in the 2.0.5.5 release has known issues when running with other features such as OSPF and VRRP. These issues may cause high CPU utilization in meshed networks. The resulting high CPU utilization can cause general operational issues with the routing switch.

The ARU3 ASICs (-B version modules and chassis) introduced the ability to replicate a multicast stream over a tagged port by generating one copy for each VLAN that requires receipt of the multicast stream. This feature also works when deployed over an MLT link.

The above feature is limited to -B version modules and chassis; therefore, using this feature may affect the suitability of -A modules and chassis when deploying a multicast-enabled network.



Note: DVMRP is not supported on ARU2/QUID4 Enforce Operational Configuration (EOC) mode. ARU2/QUID4 mode is considered suitable for IGMP snooping and proxy operation.

An additional consideration is because some IP multicast MAC addresses share the MAC address used by the reserved range of 224.0.0.x, IP multicast sessions with destination MAC 01-00-5E-00-00-xx are not processed and are flooded in the VLAN. The affected address range is 225-239.0.0.x and 224-239.128.0.x (108919, 108920). Whenever possible, configure IP multicast applications to not use these address ranges.

New Features in Release 2.0.5.5

This section discusses the new features in release 2.0.5.5:

- [ARP Learning from DHCP](#) (this page)
- [Discarding Traffic for Unknown Routes](#) (page 6)

ARP Learning from DHCP

The LearnArp Bootstrap Protocol/Dynamic Host Configuration Protocol (BootP/DHCP) Relay parameter for an IP interface controls whether the Accelar routing switch attempts to validate the binding of an IP address to the BootP/DHCP client by initiating an ARP request for the IP address. By default, LearnArp is disabled. LearnArp should be enabled on the IP interface if there are workstations initiating the BootP/DHCP process that do not reply to ARP requests. One situation where this may happen is in a diskless workstation where the boot ROM neither replies to an ARP request nor issues an ARP request.

When LearnArp is disabled (default value), the Accelar routing switch creates an ARP entry only if the end station validates the ARP information by replying to an ARP request. This is done for security reasons, to prevent address spoofing. When forwarding a BootP/DHCP reply from the server back to the client, the Accelar routing switch creates a temporary ARP entry based on the information in the server's reply. The Accelar routing switch then attempts to validate the ARP information by sending an ARP request to the client. If an ARP reply is received, the ARP entry is validated and the appropriate entry is created in the ARP cache. If no ARP reply is received, the temporary ARP entry is deleted; and all further communication to the workstation fails because there is no ARP entry.

When LearnArp is enabled, an ARP entry is created for the workstation in the ARP cache based on the information received in the BootP/DCHP server's reply. In this case, the Accelar routing switch does not require the ARP entry to be validated. Note that this is not a significant security hole because the Accelar routing switch only populates an ARP entry based on a reply received for an outstanding BootP/DHCP request for a specific MAC address.

ARP entries learned using LearnArp age as all other ARP entries.

Diskless Workstation

One case where you must enable LearnArp is when a diskless workstation obtains an IP address and boot image from servers and the boot ROM will not respond to ARP requests. In this situation, the diskless workstation obtains an IP address through BootP/DHCP; but if LearnArp is disabled, the Accelar routing switch does not receive a reply to its ARP request, and the switch deletes the ARP entry for the workstation. When the workstation attempts to retrieve its boot image (for example, using TFTP) the transfer fails because there is no ARP entry for the diskless workstation in the Accelar routing switch. Enabling LearnArp on the IP interface will allow the diskless workstation to boot successfully.

Configuring Using the CLI

To enable or disable LearnArp for DHCP relay on a VLAN using the CLI, enter:

```
config vlan <vlan#> ip dhcp-relay learnarp <enable|disable>
```

To enable or disable LearnArp for DHCP relay on an isolated router port using the CLI, enter:

```
config ethernet <port#> ip dhcp-relay learnarp <enable|disable>
```

Discarding Traffic for Unknown Routes

Discarding traffic for unknown routes allows you to disable the generation of ICMP net unreachable messages by the switch for traffic for an unknown route. When this feature is set to disabled, traffic for an unknown route is discarded at the I/O module, which preserves CPU resources.

The default setting is disabled. You must enable this feature to use the functionality.

Configuring Using the CLI

To use the CLI, enter the following IP configuration command:

```
config ip icmp-net-unreach <disable|enable>
```

STG and BPDU Clarification

The following two controls regulate the behavior of the Spanning Tree Protocol (STP) in a Spanning Tree Group (STG) on an Accelar switch:

- A global parameter to enable or disable STP at the STG level
- Port parameters to enable or disable STP on individual ports

When the STP is globally disabled on the STG, received bridge protocol data units (BPDUs) are handled like a MAC-level multicast and flooded out the other ports of the STG. Note that an STG can contain one or many VLANs. Remember that MAC broadcasts are flooded out all ports on a VLAN; a BPDU is a MAC-level message, but the BPDU is flooded out all ports on the STG, which may encompass many VLANs.

When STP is globally enabled on the STG, BPDU handling depends on the STP setting of the port:

- When STP is enabled on the port, received BPDUs are processed in accordance with STP.
- When STP is disabled on the port, the port will always be in a forwarding state, received BPDUs are dropped and not processed, and no BPDUs are generated.

To configure STP on STGs with the CLI, use the command:

```
config stg <sid> group-stp <enable|disable>
```

To configure STP on a port with the CLI, use the command:

```
config ethernet <ports> stp <sid> <enable|disable>
```

To configure STGs with Device Manager, choose VLAN > Stg (Spanning Tree Groups) > Configuration. To configure STP on a port with Device Manager, choose the port and the Spanning Tree tab.

High-Priority Switching

The Accelar routing switch operates in either of two modes: Best Effort or Priority mode. The factory default setting is Best Effort mode; in this 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.

Nortel Networks recommends that you enable Priority mode on switches in very heavy traffic situations. Enabling Priority avoids delaying vital high-priority network traffic, including BPDUs and routing protocol information. To enable Priority using the CLI, enter:

```
config sys sets flags highpriomode true
```

Bugs Fixed in Release 2.0.5.5

The following sections list bugs that were fixed in Accelar software release 2.0.5.5. Many of these bugs that were fixed addressed high CPU utilization situations, allowing greatly increased CPU utilization.

General

The following general bugs were fixed in Accelar software release 2.0.5.5:

- Spanning tree information relating to root bridge and interface state is now properly updated when a spanning tree group is disabled. (99271)
- Internal and external loopback tests are now possible on tagged ports. (101884)
- The “ifInNUcastPkts” counter now returns the correct value. (102836)
- Brouter ports no longer fail to process broadcasts after reset. (107335)
- The autolearn function of the Unknown MAC Discard feature also functions when station ARP entries are aged out or deleted. (107649)
- The number of allowed MAC addresses for the Unknown MAC Discard feature is increased to 1,000.
- Gigabit interfaces set for autonegotiate=false now initialize properly upon SSF failover. (107670)

- The switch now properly autonegotiates when the remote link of a 10/100 Mb/s connection comes back up. (103788)
- Allowed MAC addresses are now properly saved in ASCII configuration files.
- FDB and ARP cache are now cleared when enabling “unknown MAC discard.”

This fix avoids previously learned addresses remaining active.

- During boot-up, the initialization of an XLR1216TX interface with disabled ports will no longer cause a reboot of the switch. (108077)
- SSF failover now supports script files. (108173)
- Breaking and restoring MLT connections will no longer cause connectivity problems in spanning tree configurations. (108225)
- Interfaces configured for 100 Mb/s full-duplex operation no longer switch to half-duplex mode after CPU failover. (108374)
- Communication to and from the CPU no longer fails under very high loads of traffic to the CPU. (108721)
- ARU record corruption issues have been resolved. (108993)

This problem manifested itself by inconsistency messages in the log file.

- The maximum number of static MAC entries in the FDB has been increased to 1000 for run-time operations and ASCII configuration files; NVRAM configurations are still limited to 100 entries. (109690)
- Static MAC entries in the FDB can now be viewed with the `show config` command. (109691)
- Static MAC entries in the FDB are now cleared properly when a VLAN is deleted. (109712)
- Traps for spanning tree events are now generated from a separate task and no longer influence the timing of BPDU transmissions. (110399)
- The timing of BPDU transmissions is no longer influenced by other spanning tree groups. (110461)
- An MLT connectivity problem between the BayStack 450 switches and the Accelar routing switches in which the MLT ports would remain in listening state has been corrected. (108245)

- The Activity LEDs on the XLR1216-TX module now resume normal operation when the ports are enabled after being administratively shut down. (110628)
- MLT groups without ports can now be saved in the configuration file. (112310)
- Spanning Tree Protocol is no longer reenabled on MLT ports when an ASCII configuration file is loaded. (112179)
- The Ethernet port linktrap status is now properly saved in ASCII configuration files. (112425)

CLI

The following CLI bugs were fixed in Accelar software release 2.0.5.5:

- The `show config` command now properly handles VLAN IDs that are greater than 383 and no longer generates an error message “ERROR Code =0xffffffff Task=tShell: rcIpGetIf: Bad IfIndex <vid>!” (104820)
- The `show config` command no longer fails with a machine check error when dumping large configurations. (105593)
- The `show config` command no longer includes learned VLAN FDB entries. (106751)
- The port mirroring configuration is included in the `show config` command. (108742)
- The `show vlan info fdb-entry` command now includes information on the number of entries in the table. (104204)
- The `show port info vlan` command now displays the correct value for the “Discard Untagged Frames” configuration. (106916)
- The `ext-metric-type info` command for OSPF accept policies now shows the correct type information when the actual type is “any.”
- The `show config` command now includes IGMP snoop information for all VLANs. (108826)
- Blocked ports are now properly added when executing the command `config/vlan/n/igmp-snoop/static-members/x.x.x.x/add <ports> blocked`.
- A successful save to standby SSF no longer generates an error message on the console of the standby SSF. (108376)

- The `show config` command now includes inactive static routes and static routes overridden by dynamically learned routes with a lower cost. (109584, 109585)
- Route sources in announce policies are now correctly ordered in the `show config` command. (109809)
- The rip-metric information for RIP announce policies is now displayed by the `show config` command. (109811)
- The `show ip policy netlist` command now displays the name and ID of the list when an id-number is specified. (109815)
- The `show ip policy rip/ospf accept/announce info` command no longer displays all information when an id-number is specified. (109817)
- Disabling DVMRP on an interface using CLI no longer disables DVMRP globally. (110171)
- VRRP configurations are now properly saved in ASCII configuration files. (110626)
- IP multicast route configuration information is now properly saved in ASCII configuration files. (110790)
- The OSPF default metric can now be changed for multiple network types with a single command. (108663)

An example command follows:

```
config ip ospf default-metric ethernet 122 fast-ethernet 45
gig-ethernet 22
```

- The `show config` command no longer reports 100 Mb/s speed for Gigabit ports with autonegotiation disabled. (109579)
- The `config ethernet <port> info` command now includes only parameters that are configurable for that port type. (109989)
- The `config ip static-route info` command now shows only static routes. Previously it listed all routes.(110435)
- The `active` and `destination` options for the `show ip traffic-filters` command now properly return requested information. Before, these options returned no information. (111151)
- The VLAN aging time can now be set up to 1000000 using the CLI.

- The `config log` screen status (on/off) is now restored after a logout on the console.
- TFTP file transfer of large files to PCMCIA or flash memory no longer time out. (112557)
- The `show test loopback` command can now be run without specifying the optional parameters. (113727)
- The `config sys set config` command now accepts only valid choices (flash, pcmcia, nvram, or file), and file names are now correctly displayed. (113116)

IP

The following IP bugs were fixed in Accelar software release 2.0.5.5:

- DHCP packets with a source port other than 68 are now properly handled. (84890)
- ARP entries are no longer created based on information in multicast packets. (94717, 108340)
- Static routes remain active for CPU-generated traffic when IP forwarding is disabled. (96221)

This change allows you to manage the switch from a different IP-net in layer 2 only configurations.

- Source/Destination global filters now also filter out ICMP echo packets. (106846)
- It is now possible to configure static routes for routes dynamically learned in OSPF. (103495)
- VRRP on brouter ports no longer fails on reboot. (106956)
- IP traffic destined for a route with a next-hop address that is learned on different physical ports no longer forwards to CPU. (108424)
- ARP aging timer no longer counts negative (108471), and ARP entries age out correctly. (106992)
- Static routes no longer disappear when the next-hop address ages out in the ARP table and gets relearned. (108738)
- IP policy address lists are now correctly restored from ASCII configuration files. (108759)

- UDPFWD information is now properly saved in ASCII configuration files. (108831)
- IP filters are now correctly saved in ASCII configuration files and no longer cause problems upon boot. (109139)
- VRRP IP address reachability issues have been resolved. (109228)
- ARP timeout issues for next-hop addresses have been corrected. (109285)
This fix avoids routed traffic being directed to the CPU unnecessarily.
- It is now possible to configure static routes when the next-hop address is not in the ARP cache. (109320)
However, the static route becomes active only when the next-hop address is available.
- Isolated router ports (IRPs) on modules that have been hot-swapped no longer become inactive after CPU failover. (109468)
- Static default route now becomes active after RIP-supplied default route ages out. (104701)
- IP filter records are now properly initialized, which avoids problems when deleting the filters.
- Noncontiguous subnet masks are no longer allowed in RIP Accept Policies. (103228)
- A default route using a port on an ARU3-based I/O module no longer fails when the module is swapped in a chassis also containing ARU2-based I/O modules and running in the EOC-mode default.
- An ARP table indexing bug has been resolved, which avoids problems when processing ARP requests and responses.
- A static ARP entry used as a next hop for a static route no longer fails to load when the static route is loaded first.
- Inactive static routes are no longer advertised. (112594)
- Under heavy CPU load, packets with TTL expired are discarded and no ICMP processing is done.
- IP traffic is now handled properly in situations where an ARP entry exists in the AR table but no MAC entry exists.
- ASCII configuration files now handle UDP forwarding configurations with more than one portfwddlist. (112088)

- Problems related to deleting static ARP entries when used as a next hop for dynamic routes are resolved.
- Static routes created with a next hop that is already dynamically learned are no longer cached as false. (114030)
- Static routes now return properly after the physical link to the next hop device goes down and comes back up. (114031)
- VRRP hello packets are now generated by a high-priority task. This change avoids delay in transmission under heavy CPU load.

OSPF

The following OSPF bugs were fixed in Accelar software release 2.0.5.5:

- The OSPF interface parameters “poll interval,” “priority,” “retransmit delay,” and “transit delay” are properly saved and restored upon reboot of the switch. (103830, 104190)
- Disabling OSPF on a VLAN on a tagged port no longer causes OSPF to be disabled on all VLANs configured on that port. (107204)
- OSPF counters no longer increment for ports with no IP circuit configured. (108257)
- OSPF passive port configuration for tagged MLT ports is now properly saved in binary and ASCII configuration files. (108724, 109577)
- The seq numbers in LSA headers are now properly updated. (103228)
- OSPF now supports LSA aggregation and LSA ACK aggregation.
- OSPF now properly checks for virtual links when backbone area connections are lost. (114380)
- OSPF `my router id` warning messages now include the ingress port, source and destination IP addresses, and OSPF message type. These inclusions ease troubleshooting.

IP Multicast

The following IP multicast bugs were fixed in Accelar software release 2.0.5.5:

- Multicast traffic is properly forwarded over tagged Gigabit links (106116) as well as untagged access Gigabit links. (107809)
- DVMRP can no longer be configured on IPX VLANs. (91878)
- Multicast traffic over a Gigabit link will no longer interfere with unicast traffic on the same link with possible drop of unicast traffic. (108280)
- DVMRP configuration settings can now be saved in ASCII configuration files. (109010)
- DMVRP interoperability issues with third-party routers that do not send DVMRP probes have been resolved.
- Noncontiguous subnet masks are no longer allowed in VLAN multicast snoop access entries. (103229)
- Single sender entries can now be flushed from snoop VLANs using the CLI. (110619)
- Configurations with DVMRP disabled are now properly saved. (111727)
- Mroute configuration information is now saved correctly in ASCII configuration files. (112158)
- AR inconsistency problems when disabling IP Multicast on Accelar routing switches with priority queueing enabled have been resolved.
- IP Multicast packets with TTL of 1 are no longer dropped but are routed.

IPX

The following IPX bugs were fixed in Accelar software release 2.0.5.5:

- IPX traffic is no longer forwarded on spanning tree blocked ports after port state changes. (107714)
- IPX NetBIOS (type 20) traffic no longer causes OSPF traffic to be reflected back out of the receiving port. (108344)
- IPX traffic is no longer routed out of spanning tree blocked ports.
- Queueing priority for IPX packets sent out from the CPU is now always properly set; this change avoids possible out-of-sequence packets.

- The `show ipx route` command now displays the correct port number when an IPX route is relearned on a different port.
- The occasional dropping of IPX packets when routing between gigabit connections due to misinterpretation of NCP sequence numbering as hop count is resolved. (113700, 114510) This fix addresses the performance problems seen with IPX file transfer over gigabit links.

Known Issues

The following sections list known issues in Accelar software release 2.0.5.5.

General

The following known general issues are in Accelar software release 2.0.5.5:

- Some resources are reserved when using software release 2.0.x in QUID5/ARU3 mode. As a consequence, this configuration will support a maximum of 100 VLANs where software release 1.3.x supports up to 124 VLANs.

In both cases (software versions 1.3.x and 2.0.x), the maximum VLAN number is reduced by the number of STG groups (1 per STG group) and MLT links (4 per MLT link). Using software version 1.3.1, the maximum VLAN number is further reduced by the number of IGMP-snoop groups (1 per group).

- SNMP may fail after receiving an invalid SNMP get request. (111019) Once this failure occurs, SNMP does not recover.
- The `ipForwDatagrams` counter returns invalid data (decrementing number) when queried on a switch with IP forwarding disabled. (111336)
- The `rcStatBridgeOutBroadcastFrames` counter is not supported. (113124)
- Disabling OSPF on a VLAN may cause OSPF to be disabled on a tagged port if there are other VLANs with OSPF still enabled.

To recover from this situation, reenables OSPF on the tagged port.

- When you use more than a single destination filter to a host address, destination filters added after the first one fail; only the first applied filter functions correctly.

- When heavily oversubscribed, the 2-port Accelar Gigabit Ethernet module may experience intermittent connectivity loss.

To avoid this issue, distribute traffic over multiple Accelar Gigabit Ethernet modules.

Multicast

The following known multicast issues are in Accelar software release 2.0.5.5:

- IGMP snooping may forward multicast data to the wrong VLAN in a situation when multiple Snoop VLANs exist and a multicast data stream first ingresses a Snoop VLAN that does not have the lowest VLAN ID. The multicast data gets forwarded to the receiver's VLAN with the lowest VLAN ID. (109720)
- When ports are moved between VLANs, the multicast data stream for the moved port may be dropped. (109721)
- If there are multiple snoop-enabled VLANs and the VLAN that a multicast stream first ingressed gets disabled and then reenabled, that VLAN may never learn the sender(s). (109822)
- Using DVMRP, senders are aged out at 5-minute intervals rather than aged out dynamically. This situation may cause a periodic interruption of multicast sessions. (110522)
- Software currently limits the combined number of multicast senders and receivers to a total of 400. (109932, 108438)
- When using IGMP snooping and a querier moves from an active multicast router port to a statically configured port, the old querier port may be left in an active multicast router state after the move.

The workaround is to disable and then reenable IGMP snooping on the VLAN. (109510)

- ARP entries can be removed from the ARP table after the multicast stream is started on a given port. This situation may cause a loss of subsequent unicast traffic. (110042)
- You cannot add a static multicast receiver after inserting a multicast access filter for the same multicast group. (97499)
- Deleting a VLAN does not remove IGMP access filters. (97500)

IPX

The following known IPX issue is in Accelar software release 2.0.5.5:

- When multiple encapsulation types are configured on links between Accelar routing switches with IPX routing enabled, CPU intervention may be required to forward traffic, depending on the encapsulation type through which the routes are learned. (112681) This situation can cause high CPU utilization and affect performance.

The workaround is not to configure multiple encapsulation types on those links.

Related Publications

For additional information about the Accelar 1000 Series products, refer to the documents found at <http://support.baynetworks.com/library/tpubs/nav/rtswitch/> on the World Wide Web.