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

Software Release 2.0.3

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 206494-C
July 1999

NORTEL NETWORKS™

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Introduction

This release note addendum for Accelar™ software release 2.0.3 describes the enhancements and bug fixes to the Accelar software that have been implemented in release 2.0.3. 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 2.0.3 includes updates to the run-time software only. The latest software components are:

- Run-Time Software Version 2.0.3 (acc2.0.3)
- Boot Monitor Software Version 2.0.1 (accboot2.0.1) supplied as a Boot Monitor Updater
- Device Manager Version 2.0.1 (for Microsoft® Windows® 95 or 98 and Windows NT®: dm_201.exe; for UNIX: dm_2.0.1.tar.Z)
- VLAN Manager Version 2.0.1 (for Windows 95/98 and Windows NT: dm_201.exe; for UNIX: dm_2.0.1.tar.Z)

For instructions to download the software, refer to *Upgrading to Accelar 2.0 Software* (Bay Networks part number 206077-A) found on the documentation CD and on the Nortel Networks Customer Service Documentation Web page. For descriptions of Accelar Release 2.0 software features and limitations, refer to the 2.0 release notes (Bay Networks part number 896-00181-E).

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. For details, refer to Accelar 2.0 documentation.

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 2.0.3

Note the following recommendations and miscellaneous information about Accelar software release 2.0.3:

- The new XLR1298SF SSF module has 32 megabytes (MB) of dynamic random access memory (DRAM). Although release 2.0.3 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 module 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 module to increase DRAM to 32 MB.
- When loaded on an XLR1297SF module with 16 MB of DRAM, IPX maximum RIP routes and maximum SAP entries are set to minimum values (RIP 128; SAP 64) to conserve memory. If you are using IPX and require more IPX RIP routes or SAP entries, the values for IPX maximum RIP routes and maximum SAP entries can be reset by using the following CLI commands:

  — **config ipx set max-route <value>**

  — **config ipx set max-sap <value>**

  After resetting the parameters, save the configuration and reboot the switch.
- 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*.

Hardware Multicast Limitations in Release 2.0.3

The ARU3 ASICs (-B version modules and chassis) introduced the capability 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.

This 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.

Bugs Fixed in Release 2.0.3

The following sections list bugs that were fixed in Accelar software release 2.0.3.

General

The following general bugs were fixed in this release:

- The priority of the receive task for incoming traffic toward the CPU can now change dynamically to ensure that other tasks continue to run while the CPU is receiving heavy traffic. (101152)
- The checksum is verified when saving configurations to the standby SSF. (101559)
- The XLR1216FX I/O modules now are initialized properly after switchover to the redundant SSF. (102783)
- Gigabit links will no longer show a status of “half duplex/10Mbps” when the interface is down. (102792)
- 100BASE-FX ports respond properly when reenabled after being administratively brought down. (103954)
- Spanning tree group (STG) settings for tagged ports are now correctly saved in the configuration file. (104262)
- The port status of the LinkSafe™ module now is properly reported. (104495)
- On Accelar 1050 switches and Accelar 11xx switches, it is not possible to set the eoc mode to a higher ARU version than the ARU version on the motherboard, thus avoiding misconfiguration. (104669)

CLI

The following CLI bugs were fixed in this release:

- The **show config** command properly handles brouter ports.
- The “high priority” option can now be properly set on source MAC and IP subnet-based VLANs. (103717)
- The “any” parameter was deleted from the route-source fields in the **config ip policy ospf announce** and the **config ip policy rip announce** commands. (105053)
- The output from the **config ip policy ospf announce info** and the **config ip policy rip announce info** commands now displays the correct information. (105056)
- The new default values for the **config ip policy ospf accept** and the **config ip policy rip accept** commands are “accept,” rather than “ignore.” The new default value for the **config ip policy ospf announce ext-metric** command is “any” rather than “0.” (105066)

IP

The following IP bugs were fixed in this release:

- DHCP-relay can no longer be configured on IP subnet-based VLANs, as it is not supported. (99248)
- The maximum cost of a static route was raised from 15 to 65,535 in the CLI. (103826)
- The ARU now is updated correctly after receiving new RIP information with a different metric. (103986)
- Routed traffic no longer gets forwarded to the CPU after receiving data from the same source IP address from different redundant paths. (104002)
- Local IP traffic for nonexistent or aged-out hosts will no longer be forwarded to the default router. (105488, 105656)

OSPF

The following OSPF bug was fixed in this release:

- Different summary Link State Advertisements (LSAs) containing the same network ID, but different masks, are now handled properly.

IP Multicast

The following IP Multicast bugs were fixed in this release:

- DVMRP probes now are properly handled on non-DVMRP switches when IGMP snooping is enabled.
- The selection of the designated forwarder in DVMRP now functions properly. This correction will ensure that multiple copies of the same multicast packets are no longer forwarded in redundant configurations. (102918)

IPX

The following IPX bugs were fixed in this release:

- When IPX forwarding is manually disabled on a VLAN, it will remain disabled after all ports in that VLAN have been down. (102415)
- When IPX forwarding is disabled on a VLAN and subsequently that IPX network address is learned over RIP, the route now is correctly identified as an externally learned network address. (102419)
- When IPX forwarding is disabled on a VLAN and subsequently that IPX network address is learned over RIP, the IPX broadcasts are no longer routed and switched over tagged ports. (102435)
- When an untagged IPX packet arrives on a tagged port (which also belongs to a VLAN with IPX forwarding disabled), the IPX packet is no longer processed in the context of the VLAN with IPX forwarding disabled. The untagged IPX packet is now processed according to the default VLAN for that port. (102436)
- When a VLAN with IPX forwarding enabled is used to learn IPX routes, and afterwards IPX forwarding is disabled on that VLAN, the learned routes no longer use the disabled IPX VLAN as the next hop. (103146)
- When routing IPX traffic with an IPX source network of “0,” the Accelar routing switch no longer incorrectly modifies the checksum. (103173)

- IPX hop count now is incremented when routing broadcast traffic. (103212)
- IPX routes now are properly updated when learned via another interface. (103815)
- Modifying the SAP update-delay parameter with the CLI no longer incorrectly modifies the RIP update-delay parameter. (104006)
- Static SAP entries can now be removed from the configuration file. (104668)

MLT

The following MLT bugs were fixed in this release:

- The ports within an MLT group now respond correctly when administratively brought up and down.
- The pathcost of an MLT link is calculated dynamically, based on the number of active links.
- The CLI now removes the correct VLAN from an MLT group. (104697)

Known Issues in Release 2.0.3

General

The following general issues exist in this release:

- The trap receiver version cannot be set from the CLI. (105580, 94343)
- User-defined PIDs for protocol-based VLANs are not properly saved in NVRAM. (105202)

VRRP and MLT

The following VRRP and MLT issues exist in this release:

- In an unlikely configuration where you have more than two Accelar units connected with MLTs in a serial configuration running VRRP, it is possible that if an MLT link goes down, then back up, the VRRP advertisement messages will not be seen by all Accelar units. (99150)
- In specific MLT configurations, with DVMRP and VRRP enabled, ARP entries can become corrupted (i.e., have invalid port-entries). (104872)