Part No. 206494-L
March 2000

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 2.0.7.0

NORTEL NETWORKS™

Copyright © 2000 Nortel Networks

Trademarks

Statement of Conditions

Introduction

This release note addendum for Accelar™ software release 2.0.7.0 describes the enhancements and bug fixes to the Accelar software that have been implemented in release 2.0.7.0. This document is an addendum to the *Release Notes for the Accelar 1000 Series Products Software Release 2.0* (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 2.0.7.0 includes updates to the run-time software only. The latest software components are:

- Run-Time Software Version 2.0.7.0 (ac1a2070.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: Before upgrading your software from earlier versions, **back up** your current configuration file. Version 2.0.7.0 configuration files contain configuration options that are not compatible with run-time options in earlier software versions. 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.

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 2.0.7.0. Configurations with boot monitor software versions prior to 2.0.1 must upgrade to Boot Monitor Software Version 2.0.5.

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:

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 2.0.7.0 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 2.0.7.0

Note the following recommendations and miscellaneous information about Accelar software release 2.0.7.0:

- Accelar software release 2.0.7.0 does not support global filters. Configuration information relating to global filters is ignored on boot-up when you use software release 2.0.7.0. Upon booting up with software version 2.0.7.0, 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 2.0.7.0, 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 2.0.7.0 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 Enforce 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 2.0.7.0

DVMRP in the 2.0.7.0 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 that, 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.

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
```

Note: The switch must be rebooted before this change takes effect

Disabling IPX NetBIOS propagation

With the release of Accelar software version 2.0.4 and higher, you can disable IPX NetBIOS (type 20) propagation. You can enable or disable IPX NetBIOS (type 20) propagation globally, that is, on all IPX interfaces in the entire chassis.

Configuring

Configure this feature using the CLI. The CLI command to enable or disable IPX NetBIOS (type 20) propagation is `config ipx set netbios <on/off>`.

To view the current state of IPX NetBIOS propagation, use `config ipx set info`.

Note: The option to enable or disable IPX NetBIOS propagation is associated with IPX routing, so it is relevant only to switches with the ARU3 module and with IPX enabled.

Flash commands

The verbiage in the flash commands `format`, `squeeze`, and `recover` is changed to accurately indicate the behavior when leaving the command—the operation is not canceled when selecting to continue; rather the operation continues in the background. Any attempt to access or manage the flash during processing will fail. (115397-1, 116199-1)

The following is an example of the revised wording:

```
Accelar-1200#

Accelar-1200# format fl

formatting ... Press any key to push operation to
background.
```

When you press any key, the following text appears on the screen:

```
Note: If you push operation to background you will not be
advised as to the result of the operation.

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

formatting ... success

Accelar-1200#



Accelar-1200#

Accelar-1200# format fl

formatting ... Press any key to push operation to
background.
```

When you press any key, the following text appears on the screen:

```
Note: If you push operation to background you will not be
advised as to the result of the operation.

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

formatting ... operation pushed to background

Accelar-1200#
```

Bugs fixed in release 2.0.7.0

The following sections list bugs that were fixed in Accelar software release 2.0.7.0.

General

The following general bugs were fixed in Accelar software release 2.0.7.0:

- The erroneous message upon CPU failover—`agentLoad: problem creating access policy entry`—displayed in the CLI and syslog was removed. (92779-1,100848-1)
- The session no longer becomes unresponsive (hangs) when telnetting into the standby SSF. (92819-1)
- Attempting to create a brouter port using the CLI when the port is already configured as IRP with assigned IP address no longer results in an empty VLAN. (98856-1)
- The default VLAN ID can no longer be changed for access ports. (99201-1)
- Forwarding source MAC VLAN traffic no longer stops when you toggle ports from access to tagged. (103853-1)
- The Enforce Operational Configuration (EOC) mode for ARU2 now also supports ARU2/QUID2 hardware combinations. (107667-1)
- You can no longer add the same MAC address to multiple source MAC VLANs. (108135-1, 115303-1)

- An SNMP trap message was added that indicates when a power supply comes up. (109993-1)
- The internal loopback test no longer fails after boot on a port with no cable attached. (92010-1)
- You can now remove VLANs from MLT groups configured as access and remove the last VLAN from MLT groups configured as tagged. (110636-1)
- Default VLAN information is now maintained when a port is deleted from an MLT group. (114085-1)
- Repeatedly starting and stopping the internal loopback test in a fast sequence no longer causes the switch to restart. (118613-1)
- The MIB-2 variable `ipForwDatagrams` no longer decrements in layer-2-only configurations, that is, in switches with IP forwarding disabled. (111336-1)
- Bridge filters configured for MAC Multicast addresses now show up in Device Manager and CLI. (113440-1)
- Inconsistencies in the unknown-Mac-Discard setting between ports of an MLT group are now recognized. (119705-1)
- Default VLAN ID can now be changed on MLT tagged ports. (120021-1)
- The routing switch no longer reboots when a user with read/write access enters a specific string using the Web interface. (120206-1)

CLI

The following CLI bugs were fixed in Accelar software release 2.0.7.0:

- SourceMac VLANs now report the correct monitor status. (101532-1)
- The `show port info stats vrrp` command now displays results properly. (105335-1)
- The autonegotiation field in the output of the `show port info all` command is now populated with N/A for 100BASE-FX modules. (106974-1)
- You can delete SNMP trap receivers using the CLI. (107290-1)
- The `show port info stg main` command now shows all ports. (107610-1)
- The formatting problem in the output of the `show mlt stats` command was resolved, and the results now display properly. (107678-1)
- The `show config verbose` command now includes SNMP trap receiver information. (109042-1)
- The spelling errors in the `show ports` command are corrected. (109699-1, 110560-1, 113503-1)
- The spelling errors in the `config ethernet unknown-mac-discard info` command are corrected. (110494-1)
- The `show mlt stats` command no longer returns negative results. (112193-1)
- Removing an I/O module in a distributed MLT group no longer causes the MLT counters to show incorrect values and the MLT group to disappear upon reboot. (114817-1)
- The `monitor mlt stats interface` command now returns the correct prompt. (116200-1)
- Output of the `config ethernet <port> info` command no longer includes preferred-phy for 10/100 ports. (116885-1)
- To avoid confusion, additional information has been added to the `config vlan <id> agetime` command and to the `config vlan <id> fdb-entry aging-time` command; and error messages are clearer. (117070-1)

IP

The following IP bugs were fixed in Accelar software release 2.0.7.0:

- ARP entries are no longer generated from IP Multicast traffic. (94717-1, 108340-1)
- Device Manager no longer accepts a HopOrMetric of zero for a static route (108205-1)
- IP filter records are now properly identified as “IP FILTER” in the CLI, rather than as “UNKNOWN.” (113189-1)
- Applying a second destination filter to a specific host address now functions properly. (113944-1)
- ARP requests from MAC multicast addresses are no longer processed. (115513-1)
- The routing switch now properly deals with situations where IP traffic from the same IP source address is received on two different ports and one of these ports goes down. (117515-1)
- The erroneous CLI message—`failed adding a route`—when configuring source/destination filters is now suppressed.
- Forwarding traffic from the same source IP address received on different physical ports between IRPs and VLANs is now handled in the I/O module and no longer requires CPU involvement. (119159-1) Previous releases of software already dealt with VLAN-to-VLAN traffic.
- Inactive static routes are not advertised in RIP or OSPF. (112594-1)
- IP source destination filters with multiple conditions now function properly. (112607-1, 114942-1)
- RIP-2 route convergence now functions properly when in an MLT group—all links in the group are brought down simultaneously. (118211-1)
- Issues in the ARP handling process are resolved. (118718-1) These issues occasionally appeared in situations with very heavy ARP processing. The symptoms included a sustained 100% CPU utilization and no access to the system console and affected switch functionality such as STP. Usually, a switch reset was required to recover.

This issue did only show up occasionally in particular environments—generally those with non-IP layer 2 traffic as well as routed layer 3 IP traffic, which implies overlapping VLANs. It requires very intense ARP use for this issue to occur; adjusting aging timers to minimum values and having high numbers of stations in the network contribute to this situation.

Note: CPU utilization should not be used as the sole symptom to determine if this issue occurs. CPU utilization can reach high levels and spike to 100% for different reasons, even under normal switch operations. The switch OS has been designed to cope with those situations and has built-in facilities to ensure that vital network tasks have the appropriate priorities. Spikes in CPU utilization are not reason for concern, unless CPU utilization is sustained at 100% and does not recover.

- Inconsistency problems with Source/Destination filter IDs were resolved. (120857-1)
- Traffic filters configured with Device Manager now always show up in the `show config` command. (120709-1)

OSPF

The following OSPF bugs were fixed in Accelar software release 2.0.7.0:

- The routing switch now properly handles fragmented OSPF packets, both in transmission and reception.
- ABR routers configured to summarize/aggregate networks and connecting between the same areas no longer learn the same summary-link LSAs of each other. (110040-1)
- The routing table no longer uses less specific AsSummary route when a more specific AsExternal route is learned. (113786-1)
- The `show ip ospf lsdb` command now also displays AsExternal LSAs.

VRRP

The following VRRP bugs were fixed in Accelar software release 2.0.7.0:

- IP packets with the VRRP MAC address as destination, received on two different ports, are now handled properly and do not go to the CPU unnecessarily. (117777-1)
- Configurations with multiple VRRP instances on the same interface now function correctly and no longer result in multiple masters. (117948-1, 118073-1, 118555-1)
- VRRP IP addresses now continue to respond to ICMP after transiting from backup to master. (120331-1)
- The VRRP update task is optimized, reducing CPU load. (120612-1, 121390-1)

IP Multicast

The following IP Multicast bug was fixed in Accelar software release 2.0.7.0:

- The software limitation of 400 entries in the IGMP group table was removed. The software now allows as many IGMP snoop group table entries as the hardware supports. (116285-1)

IPX

The following IPX bugs were fixed in Accelar software release 2.0.7.0:

- Changes made using Device Manager in IPX routing configuration (RIP off or SAP off) are now be reflected in the CLI. (101796-1)
- Static SAP entries no longer take precedence over lower-hop-count dynamic SAP entries. (107364-1)
- Static SAP service names with spaces are now properly saved and restored in ASCII configuration files. (116771-1)
- IPX protocol-based VLANs are now properly activated after resetting the routing switch, and IPX routing functions correctly. (117575-1)

- The routing switch no longer responds to specific requests for another network (N2) on a particular network (N1) if the switch learned the requested network (N2) on that network (N1). The other router that announced N2 will handle the request. This fix avoids routing via CPU on multiencapsulated VLANs.
- The routing switch now learns SAP information from local servers with source/destination network of 0. (119969-1)

Known issues

The following sections list known issues in Accelar software release 2.0.7.0.

General

The following known general issues are in Accelar software release 2.0.7.0:

- 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) When this failure occurs, SNMP does not recover. You must reboot the switch.
- 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, reenable OSPF on the tagged port.
- 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.

- The port configuration in the `show config` command always shows the highest Spanning Tree Group number configured in the switch, even if the port is actually part of another Spanning Tree Group. (120934-1)
- SNA-802.2-Protocol-based VLANs do not support DSAP/SSAP values other than 0x04. (118821-1)
- The internal loopback test will fail on a port that is part of an MLT group. (120931-1)

IP

The following known IP issues are in Accelar software release 2.0.7.0:

- IP policies may become corrupted after they are initially created using an ASCII config file and later edited using Device Manager and saved to NVRAM. (120551-1)
- The routing switch does not use a dynamically learned route (RIP/OSPF) when a static route for that network becomes inactive. (115167-1, 121564-1)

VRRP

The following known VRRP issue is in Accelar software release 2.0.7.0:

- ICMP support for the VRRP virtual IP address is limited. Future releases of software will allow you to disable this functionality to avoid problems with fragmentation (108271-1), traceroute (109230-1), and access to own virtual address (122482-1).

Multicast

The following known multicast issues are in Accelar software release 2.0.7.0:

- 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)
- 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)

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.