

Addendum to the Release Notes for the 1.3 Software Release for Accelar 1000 Series Products

Software Release 1.3.2

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 204767-D
January 1999



Bay Networks



* 2 0 4 7 6 7 - D *

Copyright © 1999 Bay Networks, Inc.

All rights reserved. Printed in the USA. January 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 Bay Networks, Inc.

Trademarks

Bay Networks is a registered trademark of Bay Networks, Inc.

Accelar and the Bay Networks logo are trademarks of Bay Networks, Inc.

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, Bay Networks, Inc. reserves the right to make changes to the products described in this document without notice.

Bay Networks, 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 1.3.2 describes the enhancements and bug fixes to the Bay Networks® Accelar™ software that have been implemented since release 1.3.1. This document is an addendum to the *Release Notes for the Accelar 1000 Series Products Software Release 1.3* (Bay Networks part number 896-00181-D). For information about the changes between release 1.3 and release 1.3.1, refer to the *Addendum to the Release Notes for the 1.3 Software Release for the Accelar 1000 Series Products, Software Release 1.3.1* (Bay Networks part number 204767-C).

Software release 1.3.2 includes updates to the run-time and boot monitor software. The latest software components are:

- Run-Time Software version 1.3.2 (acc1.3.2)
- Boot Monitor Software version 1.3.2 (accboot1.3.2) supplied as a Boot Monitor Updater



Caution: Boot Monitor Software version 1.3.2 is not compatible with the previous version of the run-time image. Always use boot monitor version 1.3.2 in combination with run-time image version 1.3.2 or above.

Software release 1.3.2 can be managed by Device Manager and VLAN Manager version 1.3.3 (dm_133.exe for Windows and dm_1.3.3.tar.z for UNIX).



Caution: Before upgrading your software, back up or save your current configuration file. The version 1.3.2 configuration file contains new configuration options, which are not compatible with 1.3.0 and 1.3.1 run-time images. It is important to back up or save the current configuration file before upgrading in case you must revert to a previous version of the run-time image.

Refer to the version 1.3 release notes for instructions to download the software. Those release notes were provided in hard copy with the 1.3 software. You can also find the release notes on the 1.3.1 Software CD and on the Bay Networks Customer Service Documentation Web page.



Note: The IP behavior has changed in this release. Starting with this version and continuing in software release 2.0, you can decide via configuration whether, when the last active port in the VLAN goes down, the VLAN or isolated routing port IP address should go down (as in version 1.1.x) or remain up (as in version 1.3.1).

As this feature has major implications on IP, management, and routing protocols, it is important to configure the routing switch with the preferred behavior. Please refer to [“New Features and Enhancements”](#) for more information about how to configure this feature. Note that the “IP down” behavior (1.1.x style) is the default and recommended behavior.

When configuring “IP up” behavior (1.3.1 style), route convergence will be slower because it will rely on the timeouts in the routing protocols (RIP/OSPF). In OSPF environments, therefore, there will not be an immediate link state change upon link down; the transition will occur only after the RouterDeadInterval expires. If this delay is a concern, configure for “IP down.”

The addendum includes the following sections:

- [New Features and Enhancements](#) on this page
- [Port Mirroring Clarification](#) on [page 3](#)
- [Bugs Fixed in Release 1.3.2](#) on [page 4](#)
- [Known Issues in Release 1.3.2](#) on [page 6](#)

New Features and Enhancements

The following new features and enhancements were added in release 1.3.2:

- Support for 32 megabytes (MB) of RAM on the new XLR1298 Silicon Switch Fabric (SSF) module. This feature *requires* the new boot monitor software (accboot1.3.2) as well as the new run-time software (acc1.3.2).
- Support for new Quid5/ARU3-based I/O modules has been added (for ARU2 mode only). In order to use the new ARU3 functionality, such as IPX routing, software release 2.0 is required. Quid5/ARU3-based I/O modules *are not supported* in this release.
- The `show sys perf` command now includes frame buffer usage information. (89175)

- The IP behavior is now configurable. (89816)

Using the CLI, you now have the capability to decide via configuration whether or not layer 3 should be notified of VLAN state changes, provided the VLAN is configured as a routable interface. This option is also available for single routable ports, in which case the criteria for state is the operational status of the port.

The VLAN is considered to be up if one of the following is true:

- At least one member of the port-based VLAN has linkup.
- At least one port member of the policy-based VLAN is active.
- At least one static member of the policy-based VLAN has linkup.

Otherwise the VLAN is considered to be down.

Use the following CLI commands to define whether or not state changes will be sent to layer 3:

```
config ethernet <port list> ip advertise-when-down enable|disable
config vlan <vlanid> ip advertise-when-down enable|disable
```

The default for this configuration option is “disable” and relates to IP down behavior as explained above; “enable” relates to IP up behavior.

Port Mirroring Clarification

Some confusion exists regarding the Accelar implementation of port mirroring. The following paragraphs replace the explanation for port mirroring found on page 15-3 of *Reference for the Accelar Management Software* (Bay Networks part number 893-01052-C).

Accelar routing switches support port mirroring for troubleshooting and network traffic analysis.

Using port mirroring, you specify a destination port on which you want to see mirrored traffic and specify the source ports from which traffic is mirrored. Any packets ingressing or egressing the specified ports are forwarded normally, and a copy of the packet is sent out the mirror port. An Accelar 1000 Series routing switch can support mirroring for only two ports. When this feature is active, all packets received or transmitted on the port(s) specified by `MirroredPortOne` and/or `MirroredPortTwo` are copied to `MirrorPort`. The mirroring operation is nonintrusive.

In addition, the port mirroring feature can be used to monitor traffic for MAC addresses where traffic with a given source or destination MAC address is copied to the mirror port. This feature is enabled by setting monitor true for a MAC address in the VLAN forwarding table for the VLAN (from the Device Manager menu, select VLAN > Bridging > Forwarding > Monitor true).



Note: In ARU1 and ARU2 hardware, routed packets are not mirrored in the egress direction.

Bugs Fixed in Release 1.3.2

The following bugs were fixed in release 1.3.2:

- RMON etherHistoryUtilization no longer returns the wrong OID type, counter instead of integer, causing problems in Optivity RMON tools or other similar applications. (85529/85976)
- RMON etherHistoryUtilization also no longer returns invalid information for non-existing modules. (92070)
- On newly installed modules, tagged packets no longer get untagged for access ports. (86535)
- You are now able to ping the Accelar switch when messages are sent to the console at a high rate. This problem was observed when ICMP redirect messages were sent for traffic going through the Accelar switch. (88888)
- The RouterDeadInterval could not be set to less than four times the HelloInterval. You can now configure the RouterDeadInterval to any integer multiple of the HelloInterval. (89394)
- Local routes are no longer cached as “false” in the routing table when the VLSM boundary ARP entry ages out of the ARP cache. (89408)
- Gigabit Ethernet autonegotiation now works correctly with the new Sun Microsystems network interface cards (NICs). (90212)
- Reception of IP datagrams with the IP option length set for 0, or unknown IP options with option fields set for 0, no longer cause continued 100 percent CPU load. (90295)

- The Accelar routing switch did not release frame buffers containing DHCP packets when no default gateway was configured, causing the switch to run low on frame buffers and ultimately to drop traffic directed to the switch. This problem has been corrected. (90443)
- Configuration settings for proxy ARP now persist across a reboot of the Accelar routing switch. (91304)
- Static ARP records are no longer deleted when the ports in a VLAN go down.
- The IP message queue is now the same size as the frame buffer queue, which improves behavior when there is 100 percent CPU load with remaining free frame buffers.
- Putting a no-link port into testing mode no longer causes the IP VLAN to go down. (91819)
- IP filters (source/destination or global) configured with drop action now work correctly when applied to a Gigabit Ethernet interface. (90399)
- MAC forwarding database entries are now cleared for ports removed from a protocol-based VLAN. (89757)
- With version 1.3.1 software on ARU1-based I/O modules, packets were sent to the CPU for routing out of local interfaces. This problem has been corrected in version 1.3.2. (92614)
- The Unknown IP Error code no longer causes a fatal error. (93251)
- The routing switch no longer crashes when attempting to get a frame buffer to read in a packet when no frame buffers are available. (92925)
- When an Accelar 1000 Series switch has one OSPF interface in the backbone area and a manual virtual link is configured on this interface, deleting the interface no longer causes a crash (assertion failure). This problem occurred if the backbone interface was an isolated routing port and link was removed from this port. (90554)

Known Issues in Release 1.3.2

The following issues are known to exist in release 1.3.2:

- Current VRRP implementation does not support authentication, and no provisions/accommodations are made for an authentication field. As a consequence, Accelar routing switches will not interoperate with BN routers or other routers using this field. (91960)
- Multi-Link Trunking 802.1Q trunks are currently supported only for a single spanning tree group. Connectivity failures may result if an 802.1Q tagged MLT is configured for multiple spanning tree groups in this release. (90775)
- The routing switch may fail to send out RIP updates while under an extreme broadcast load. (92426)
- Toggling DHCP or other IP options on a port will disable the port's ability to join a VLAN. (89557)
To recover from this state, configure the port as an isolated router port and enable DHCP. Then disable both conditions.
- Fragmented OSPF packets do not get reassembled. (90895)
- When routing, the Accelar switch does not discard datagrams with a bad destination IP address. Instead, the switch sends an ICMP destination unreachable message. (85280)
- When used as a router, the Accelar switch responds to datagrams that have a bad source IP address. (85281)
- After a failed save to NVRAM (such as with a configuration that is too large), the NVRAMUsed value indicates 0 (zero) K used. (85632)
To recover from this state, perform a successful save to NVRAM or reset the switch.
- 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)
- 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)

- Area range entries are not summarized into the backbone area. (91449)
- OSPF Announce policies “leak” external route information into the link state database (LSDB). (91907)
- ARP entry mismatches occur in VRRP environments. (91573, 91124)
- DHCP server replies with the broadcast flag set are not broadcast by the relay agent in the routing switch. (90926)
The workaround is to configure the Accelar switch to handle all DHCP responses as broadcasts.
- IP global filters can cause the routing switch to crash after a reset. (92168)
- Configuring RMON history to use more than the available RMON memory will cause the Accelar switch to crash with an assertion error after a save and reset. (93738)
- RMON etherHistoryUtilization will not work (always returns 0) after an RMON save and reboot of the routing switch. (93701)

