

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

Software Release 1.3.1

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Part No. 204767-A
November 1998



Bay Networks



* 2 0 4 7 6 7 - A *

Copyright © 1998 Bay Networks, Inc.

All rights reserved. Printed in the USA. November 1998.

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

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, 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 1.3.1 describes the enhancements and bug fixes to the Bay Networks® Accelar™ software that have been implemented since release 1.3. 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).

Software release 1.3.1 includes updates to the run-time, boot monitor, Device Manager, and VLAN Manager software. The latest software components are:

- Run-Time Software Version 1.3.1 (acc1.3.1)
- Boot Monitor Software Version 1.3.1 (accboot1.3.1) supplied as a Boot Monitor Updater
- Device Manager and VLAN Manager Version 1.3.1 or 1.3.2

Refer to the 1.3 Release Notes (part number 896-00181-D) for instructions to download the software and for descriptions of software features and limitations. 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.

The addendum includes the following sections:

- [New Features and Enhancements](#) (including new CLI commands)
- [rlogin Information](#) on [page 7](#)
- [RMON Polling Characteristics](#) on [page 9](#)
- [Bugs Fixed in Release 1.3.1](#) on [page 9](#)
- [Known Issues in Release 1.3.1](#) on [page 12](#)

New Features and Enhancements

The following new features and enhancements were added in release 1.3.1:

- Multi-Link Trunking is completely enabled in the CLI and in Device Manager. This feature is described in the documentation on the 1.3 Documentation CD (*Reference for the Accelar Management Software*, part number 893-01052-C, and *Reference for the Accelar 1000 Series Command Line Interface*, part number 202086-A).

- New run-time CLI commands have been added to increase CLI functionality. These commands are described in [“CLI Changes”](#) on [page 3](#).
- IP filtering has been enhanced to allow multiple IP filters with the same source and destination network addresses. Device Manager supports this new capability and is backward-compatible for the previous MIB supported in 1.3.0. No new CLI commands are required.
- The number of RMON alarms that can be saved to flash or to PCMCIA has been increased.
- Two new trap log messages are sent out when CPU switchover occurs. These traps are the “warm start” trap and an error indication trap to indicate that the CPU has switched to standby mode.
- A trap message is now generated when 90 percent of the CPU’s buffers are filled.
- Messages are now logged when a switch fabric test is started or stopped.
- If someone attempts to log in to the CLI or Web management page with an incorrect password, a “Blocked unauthorized CLI (or Web) access” message will appear on the terminal to alert the system administrator.
- The maximum number of allowable static routes has been increased to 500.
- When a CLI show command is issued for a blank table, the message “no entries in xxxx table” is now displayed.
- The algorithm for aging out the dynamic membership of ports in policy-based VLANs has been improved. In past versions, when the dynamic membership aging timer expired in a policy-based VLAN, all active, dynamic members were deactivated from the VLAN. The port was then moved to standby mode and removed from the VLAN’s multicast groups. This action could result in intermittent loss of connectivity between a client and a server.

In release 1.3.1, dynamic member aging is done in two stages. On the first expiration of the aging timer, the active port is moved into standby mode but is left in the VLAN multicast group. If the port has not rejoined the VLAN by the next expiration of the aging timer, it is removed from the VLAN multicast group.

CLI Changes

The format of the following commands has changed from the 1.3 release:

Command Format in Release 1.3

```
config ip proxy disable [<ipaddr>]
config ip proxy enable [<ipaddr>]
config vlan <vid> create <sid> [<ports>]
[name <value>]
config vlan <vid> add <ports>

config vlan <vid> remove <ports>

config vlan <vid> fdb aging-time <seconds>
config vlan <vid> fdb flush
show ip proxy [<ipaddr>]
```

Command Format in Release 1.3.1

```
config ip proxy disable <ipaddr>
config ip proxy enable <ipaddr>
config vlan <vid> create byport <sid> [name
<value>]
config vlan <vid> ports add <ports> [member
<value>]

config vlan <vid> ports remove <ports> [member
<value>]

config vlan <vid> fdb-entry aging-time <seconds>
config vlan <vid> fdb-entry flush
show ip proxy <ipaddr>
```

The parameters are the same as in the previous release except for the following:

```
config vlan <vid> ports add <ports> [member <value>]
config vlan <vid> ports remove <ports> [member <value>]
```

These commands are used to add and remove ports to/from a VLAN, where `member <value>` is the port member type (portmember/static/not allowed) used only for protocol-based and IP subnet-based VLANs to set a port as always a member, a potential member, or never a member of the VLAN, respectively. For port-based VLANs, this option is not valid because ports are always members of the VLAN. The member subcommand is not allowed for port-based VLANs.

The following sections list the new commands added in the 1.3.1 release.

Port Commands

The following port commands have been added:

- `config ethernet <ports> ip arp-response disable`—disables ARP responses on the port.
- `config ethernet <ports> ip arp-response enable`—enables ARP responses on the port.

- `config ethernet <ports> ip traffic-filter default-action forward`—sets the filter default action on the port to forward.
- `config ethernet <ports> ip traffic-filter default-action drop`—sets the filter default action on the port to drop.
- `config ethernet <ports> stg <sid> faststart <enable|disable>`—enables or disables FastStart for the port and spanning tree group.
- `config ethernet <ports> stg <sid> pathcost <integer>`—sets the contribution of this port to the spanning tree group path cost (1 to 65535).
- `config ethernet <ports> stg <sid> priority <integer>`—sets the priority of this port and spanning tree group (1 to 255).
- `config ethernet <ports> stg <sid> stp <enable|disable>`—enables or disables Spanning Tree Protocol for the port and spanning tree group.
- `show ports info arp [<ports>]`—displays ARP information for the specified ports or for all ports, including the interface index, and indicates if ARP proxy and response are enabled or disabled.
- `show ports info stg [<ports>]`—displays spanning tree group information for the specified ports or for all ports, including the interface index, the priority, the state (enabled or disabled), whether or not Spanning Tree Protocol is enabled, whether or not FastStart is enabled, and the path cost.
- `show ip resp <ipaddr>`—displays whether or not ARP response is enabled for the IP address.
- `show ip traffic-filter active`—displays all active IP filters.

Flow Commands

The following flow commands have been added:

- `config ip ipflow create src-ip <value> src-port <value> dst-ip <value> dst-port <value> protocol <value>`—creates an IP flow with the specified source IP address, source IP port ID (0 to 65535), destination IP address, destination IP port (0 to 65535), and protocol choice (IP, TCP, or UDP).

- `config ip ipflow delete src-ip <value> src-port <value> dst-ip <value> dst-port <value> protocol <value>`—deletes the IP flow with the specified source IP address, source IP port ID (0 to 65535), destination IP address, destination IP port (0 to 65535), and protocol choice (IP, TCP, or UDP).
- `show ip flow`—displays IP flows configured on the switch, including source IP address, source IP port ID, destination IP address, destination IP port, and protocol choice (IP, TCP, or UDP).

RIP Commands

The following RIP commands have been added:

- `config ip rip domain <ipaddr> <value>`—changes the RIP interface configuration domain for the specified IP address to the new domain value (0 to 39321).
- `config ip rip updatetime <seconds>`— sets the RIP update time in seconds (0 to 360)
- `config ip rip receive <ipaddr> mode <value>`—changes the RIP interface configuration receive mode for the specified IP address to a new receive mode value (RIP1, RIP2, or RIP1 or RIP2).
- `config ip rip send <ipaddr> mode <value>`—changes the RIP interface configuration send mode for the specified IP address to a new send mode value (not send, RIP 1, RIP 1 comp, or RIP2).

MLT Commands

The Multi-Link Trunking (MLT) commands were not available in the 1.3 release, but they are documented in the *Reference for the Accelar 1000 Series Command Line Interface*.

- `config mlt <mid> add ports <ports>`—adds ports to a Multi-Link Trunk.
- `config mlt <mid> add vlan <vid>`—adds a VLAN to a Multi-Link Trunk.
- `config mlt <mid> create`—creates a Multi-Link Trunk.
- `config mlt <mid> delete`—deletes a Multi-Link Trunk.
- `config mlt <mid> name <string>`—names a Multi-Link Trunk.
- `config mlt <mid> remove ports <ports>`—removes ports from a Multi-Link Trunk.

- `config mlt <mid> remove vlan <vid>`—removes a VLAN from a Multi-Link Trunk.
- `config mlt <mid> type <access|trunk>`—sets the Multi-Link Trunk as an access or trunk port.
- `show mlt info [<mid>]`—displays Multi-Link Trunking information for the specified MLT or for all MLTs, including the ports that belong to the MLT, the VLANs that belong to the MLT, and if the MLT is configured as an access or trunk port.

VLAN Commands

The following VLAN commands have been added:

- `config vlan <vid> action <action choice>`—sets the action for the specified VLAN to: none, flush MAC forwarding database, flush ARP, flush IP, flush dynamic membership, flush all, flush IGMP snoop membership, trigger RIP updates, or flush IGMP snoop master router.
- `config vlan <vid> agetime <10..100000>`—sets the timeout period in seconds for aging out the dynamic member ports of a policy-based VLAN.
- `config vlan <vid> create byport <sid> [name <value>]`—creates a port-based VLAN.
- `config vlan <vid> create byprotocol <sid> <ip|ipx802dot3|ipx802dot2|ipxSnap|ipxEthernet2|appleTalk|decLat|decOther|sna802dot2|snaEthernet2|netBios|xns|vines|ipV6|usrDefined> [name <value>]`—creates a protocol-based VLAN using the selected protocol or a user-defined protocol with the defined name and value.
- `config vlan <vid> create byipsubnet <sid> <ipaddr/mask> [name <value>]`—creates a source IP subnet-based VLAN.
- `config vlan <vid> fdb-entry monitor <mac> status <value> <true|false>`—sets the VLAN forwarding database monitor MAC address status to: other, invalid, learned, self, or management and enables or disables the monitor.
- `config vlan <vid> fdb-entry priority <mac> status <value> <high|low>`—sets the VLAN forwarding database MAC address status to: other, invalid, learned, self, or management and the priority to high or low.
- `config vlan <vid> fdb-filter add <mac> port <value>`—adds a filter member to the selected VLAN bridge and specified MAC address and port.

- `config vlan <vid> fdb-filter remove <mac>`—removes a filter member from the selected VLAN bridge and specified MAC address.
- `config vlan <vid> fdb-filter notallowfrom add <mac> port <value>`—adds a not-allowed filter member to the selected VLAN bridge and specified MAC address and port.
- `config vlan <vid> fdb-filter notallowfrom remove <mac> port <value>`—removes a nonallowed filter member from the selected VLAN bridge and specified MAC address.
- `config vlan <vid> fdb-static add <mac> port <value>`—adds a static member to the selected VLAN and specified MAC address and port.
- `config vlan <vid> fdb-static remove <mac>`—removes a static member from the specified VLAN and MAC address.
- `config vlan <vid> highpriority <true|false>`—turns the high-priority setting on (true) or off (false) for a VLAN.
- `config vlan <vid> ip resp disable`—disables IP responses for the selected VLAN.
- `config vlan <vid> ip resp enable`—enables IP responses for the selected VLAN.
- `show vlan info arp [<vid>]`—displays ARP information for the specified VLAN or for all VLANs including interface index and whether or not ARP proxy and ARP response is true or false for the VLAN.

rlogin Information

Accelar rlogin access is supported and controlled through the Accelar management access policies. The Accelar switch supports four types of management services:

- Telnet
- SNMP
- HTTP
- rlogin

In the default state, Telnet, SNMP, and HTTP are enabled, allowing switch management through the CLI, Device Manager, and the Web. Also by default, access management filters are globally disabled.

To use rlogin, you must create a new access filter for rlogin management or add rlogin to the default filter and then globally enable filtering. The procedure is as follows.

To enable rlogin through Device Manager:

1. **To view the current default policy, go to the Edit>Security>Access Policy window.**
2. **Create an access policy for rlogin:**
 - a. **Go to Edit>Security>Access Policy>Insert.**
 - b. **Enter a policy ID (1 to 65535) and name.**
 - c. **Click on Active: true, Mode: allow, Service: rlogin.**
 - d. **Set a precedence for the policy (1 to 128 with the lower number being the higher precedence).**
 - e. **Enter a TrustedHostAddr: the IP address of the host from which rlogin is to be allowed.**
 - f. **Enter a TrustedHostUserName: the name of the user to be allowed rlogin access.**
 - g. **Click on Insert.**
3. **Globally enable management access by selecting Edit>Chassis>System and clicking on EnableAccessPolicy true.**

To enable rlogin through the CLI:

1. **To view the current default policy, use the `show ip access info` command.**
2. **Create an access policy for rlogin using the following commands:**
 - a. `config ip access-policy policy <pid> create`
where `pid` is the policy ID (1 to 65535).
 - b. `config ip access-policy policy <pid> rlogin enable`
where `pid` is the policy ID that was just created.
 - c. `config ip access-policy policy <pid> name <name>`
where `pid` is the policy ID and `name` is the name of the policy.
 - d. `config ip access-policy policy <pid> precedence <precedence>`
where `pid` is the policy ID and `precedence` is from 1 to 128 with the lower number being the higher precedence.

- e. `config ip access-policy policy <pid> host <ipaddr>`
where `pid` is the policy ID and `host` is the trusted host address, the IP address of the host from which rlogin is to be allowed.
 - f. `config ip access-policy policy <pid> username <string>`
where `pid` is the policy ID and `username` is the trusted host user name, the name of the user to be allowed rlogin access.
 - g. `config ip access-policy policy <pid> mode allow`
where `pid` is the policy ID.
3. **Globally enable management access using the command:**
`config ip access-policy enable true.`

RMON Polling Characteristics

The following are characteristics of RMON polling in the Accelar software:

- Hardware counters are polled at a frequency of once every 500 milliseconds.
- The smallest resolution supported for Ether stats or alarms defined on Ether stats is 1 second, which allows for double sampling.
- The first value retrieved is not deterministic in the sense that it may include statistics from packets received before Ether stats was enabled. Also, it always includes the SET SNMP PDU that created the Ether stats row (that is, enabled Ether stats on a port). All subsequent values retrieved are deterministic, as long as you wait the full 500 milliseconds for data.

Bugs Fixed in Release 1.3.1

The following bugs were fixed in release 1.3.1:

- A static route can now overwrite an existing RIP route. (77562)
- Static routes are now correctly marked inactive in the system routing table when the interface that is connected to the local network goes down. (77622)
- The switch no longer continuously reboots on a warm start when LinkSafe™ is set up to another board in the same chassis. (78775)
- Under sustained line rate testing at 10 Mb/s, the Accelar 1050 Routing Switch now forwards all frames. (79101)
- Connectivity issues for AppleTalk, VINES, and IPV6 were corrected. (79992)

- You can now reuse a monitor (mirroring) port as an access port after mirroring tagged packets. (83820)
- A static route no longer takes precedence over a local route after a reboot. (84796)
- After a save and reboot, a static route now shows up as valid. (84954)
- If a static route that overrode an OSPF dynamic route was deleted, the OSPF dynamic route was not relearned. Now, after the static route is successfully deleted, the OSPF is notified to run the SPF again to get back the OSPF route that was manually deleted to add the static route. (84980)
- In a nonremovable OSPF backbone area, the switch will determine if the ABR has an active interface in the backbone prior to generating summary LSAs into the backbone area. If the ABR is connected to the backbone area, the switch computes the inter-area routes based on the summary LSAs in the backbone area. Otherwise, it computes the routes based on the summary LSAs present in the nonbackbone area. (85004)
- Broadcast storms and Spanning Tree Protocol loops no longer occur when more than 10 VLANs are configured for a trunk port. The algorithm for handling Spanning Tree Protocol state transitions has been enhanced to provide better scaling when a port belongs to multiple VLANs. (85018)
- The `config ip arp add <ports>` command no longer causes an assertion error when an invalid port is specified. (85094)
- Broadcast storms with redundant trunk links no longer occur when one of the switches is reset. (85436)
- The real-time clock is now supported for the Accelar 1050/1051 Routing Switch (85122), for the Accelar 1100 Routing Switch with hardware version 5.0 and above, and for the XLR1298 SSF module. (88291, 86258)
- The link state database now contains the proper entries for static routes. (85524)
- The L3/L2 community strings now inherit the Read/Write/All password when upgrading from release 1.1.x. (85561)
- The consistency check routine has been modified to prevent using an IP address that belongs to a configured IP subnet-based VLAN unless the user is assigning the address to the IP subnet-based VLAN. (85576)

- When a VRRP failover occurred, the new master was not sending out a gratuitous ARP request as specified in RFC 2338 for VRRP. Thus failover did not invoke a change in the ARP entries of any clients that had the previous master's physical MAC address. Clients are now able to communicate with the new master until an ARP is generated, resulting in the ARP entries being refreshed. (85758)
- The Accelar switch no longer runs out of frame buffers when a broadcast/ARP storm occurs and routing is enabled. (86660)
- Gigabit modules are no longer mistakenly transmitting pause flow control frames during congestion. (86851)
- Potential ports no longer flood traffic when they are inactive. (87267, 87425)
- Multicast group IDs associated with dynamic multicast groups (for IGMP snooping) are now released correctly when a VLAN is deleted. (87270)
- The unnecessary flooding of ports within VLANs has been eliminated. (87653)
- You can now change the OSPF router ID dynamically from the CLI without disabling OSPF. (87772)
- A land.c (denial of service) attack can no longer cause CPU utilization to be 100 percent without recovering. (87867)
- The Web management page now displays the correct sysUpTime after more than 24 hours of operation. (87917)
- The CLI `show tech all` command no longer generates an error on VLAN information. (87956)
- If IP filters were enabled on ports and the configuration was reset, the filters might no longer have been effective. The order of the filter and route records in the ARU was changed so that the filter records are always last, regardless of the order of record creation. (89039)
- In a routed VLAN, MAC forwarding database (FDB) entries are now reliably relearned when their corresponding ARP entries are relearned or refreshed. (90158)

Known Issues in Release 1.3.1

The following issues are known to exist in release 1.3.1:

- You may not always be able to ping the Accelar switch during conditions that cause messages to be sent to the console at a high rate. This problem may occur in a configuration where numerous “not allowed to join vlan xxx” messages are produced on the console, such as when ICMP redirect messages are sent to a host attempting to route through the Accelar switch to an external router attached to the same IP policy-based VLAN. (88888)
- When an ARP entry for a VLSM boundary address ages out in a VLAN, the local route is cached as false in the routing table (ARU1 mode only). (89408)
- 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.
- 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.
- Device Manager RMON history sometimes displays erroneous (usually quite large) values. (86000)
To recover from or work around this problem, in the Device Manager Graph Port window, click on Refresh.
- VLAN Manager 1.3.0 for Microsoft® Windows® 95 or Windows NT® does not print. (87593)
- For ARU1 mode only, VLSM route records may not be cleaned up when a route is learned with a mask length that is 11 or more bits longer than an existing VLAN route. (89006)
- Device Manager 1.1.3 and 1.3.0 will crash on Windows NT 4.0 with Service Pack 4 loaded. (89042)

- The Accelar switch supports only full-duplex mode of operation for Gigabit Ethernet interfaces; half-duplex mode is not supported. If the switch receives an autonegotiation advertisement with the half-duplex bit set, negotiations will fail and the link will not be established, even when there is a common capability advertised, such as when the full-duplex capable bit is also set. (78044, 90212)

In these instances, the device that is advertising half-duplex mode should suppress the half-duplex advertisement, if possible. If this is not possible, you should disable negotiation and set the duplex mode to full-duplex operation on both the originating device and the Accelar port to which it is attached.

For example, with a Sun PCI adapter that supports half-duplex and full-duplex modes, you can temporarily change the advertisement setting to suppress half-duplex advertisement. In a Solaris terminal session, use the following commands:

```
ndd -set /dev/ge adv_1000hdx_cap 0
ndd -set /dev/ge adv_1000autoneg_cap 1
```

Note that these commands are temporary and will be lost on a power reset. Bay Networks recommends that you add these `ndd` commands to the `/etc/rc2.d` file so that the parameters will be reloaded on a reset of the Sun adapter. For more details on configuring Sun Gigabit adapters, refer to the Sun Gigabit Ethernet Adapter user guides.

- 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)
- When a Device Manager session to an Accelar 1050 or Accelar 1051 switch is opened from a Solaris workstation, the bitmap image may not be sized correctly. (85802)
- When the Accelar has one OSPF interface in the backbone area and a manual virtual link is configured on this interface, deleting this interface causes a crash (assertion failure). This problem occurs if the backbone interface is an isolated routing port and link is removed from this port. (90554)
To work around this problem, use the auto virtual link feature instead of a manual virtual link. Or, alternatively, configure the backbone interface as a routable VLAN, because the VLAN's routable interface is not deleted when link goes down on any or all of its ports.
- IP filters do not drop packets when applied to a Gigabit Ethernet interface. (90399)

- Under specific conditions of unique traffic patterns at high frame rates, a temporary condition of high CPU utilization may occur, which may 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)
- 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)
- When using Device Manager 1.3.1 for Windows, the interface statistics counters are always 0 and never increment. (90907)