System Configuration and Administration Guide

AX Series Advanced Traffic Manager

Document No.: D-030-01-00-0024

Ver. 2.6.6-GR1 5/8/2013

Headquarters

A10 Networks, Inc.
3 West Plumeria Dr.
San Jose, CA 95134

Tel: 408 325 8668 (main)
Tel: 408 325 8676 (support)
Fax: 408 325 8666

www.a10networks.com

Information in this document is subject to change without notice.

Trademarks

A10 Networks, A10 Thunder, vThunder, the A10 logo, aACI, aCloud, ACOS, aDCS, aDNS, aELB, aFleX, aFlow, aGalaxy, aPlatform, aUSG, aVCS, aWAF, aXAPI, IDAccess, IDSENTRIE, IP to ID, SmartFlow, SoftAX, Thunder, Unified Service Gateway, Virtual Chassis, VirtualADC, and VirtualN are trademarks or registered trademarks of A10 Networks, Inc. All other trademarks are property of their respective owners.

Patents Protection

A10 Networks products including all AX Series products are protected by one or more of the following US patents and patents pending: 20120216266, 20120204236, 20120179770, 20120144015, 20120084419, 20110239289, 20110093522, 20100235880, 20100217819, 20090049537, 20080229418, 20080148357, 20080109887, 20080040789, 20070283429, 20070282855, 20070271598, 20070195792, 20070180101, 8387128, 8332925, 8312507, 8291487, 8266235, 8151322, 8079077, 7979585, 7716378, 7675854, 7647635, 7552126

Confidentiality

This document contains confidential materials proprietary to A10 Networks, Inc. This document and information and ideas herein may not be disclosed, copied, reproduced or distributed to anyone outside A10 Networks, Inc. without prior written consent of A10 Networks, Inc.

A10 Networks Inc. Software License and End User Agreement

Software for all A10 Networks products contains trade secrets of A10 Networks and its subsidiaries and Customer agrees to treat Software as confidential information.

Anyone who uses the Software does so only in compliance with the terms of the End User License Agreement (EULA), provided later in this document or available separately. Customer shall not:

1) reverse engineer, reverse compile, reverse de-assemble or otherwise translate the Software by any means

2) sublicense, rent or lease the Software.

Disclaimer

This document does not create any express or implied warranty about A10 Networks or about its products or services, including but not limited to fitness for a particular use and non-infringement. A10 Networks has made reasonable efforts to verify that the information contained herein is accurate, but A10 Networks assumes no responsibility for its use. All information is provided "as-is." The product specifications and features described in this publication are based on the latest information available; however, specifications are subject to change without notice, and certain features may not be available upon initial product release. Contact A10 Networks for current information regarding its products or services. A10 Networks’ products and services are subject to A10 Networks’ standard terms and conditions.

Environmental Considerations

Some electronic components may possibly contain dangerous substances. For information on specific component types, please contact the manufacturer of that component. Always consult local authorities for regulations regarding proper disposal of electronic components in your area.

Further Information

For additional information about A10 products, terms and conditions of delivery, and pricing, contact your nearest A10 Networks location, which can be found by visiting www.a10networks.com.

End User License Agreement

IMPORTANT: PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY. DOWNLOADING, INSTALLING OR USING A10 NETWORKS OR A10 NETWORKS PRODUCTS, OR SUPPLIED SOFTWARE CONSTITUTES ACCEPTANCE OF THIS AGREEMENT.

A10 NETWORKS IS WILLING TO LICENSE THE PRODUCT TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS LICENSE AGREEMENT. BY DOWNLOADING OR INSTALLING THE SOFTWARE, OR USING THE EQUIPMENT THAT CONTAINS THIS SOFTWARE, YOU ARE BINDING YOURSELF AND THE BUSINESS ENTITY THAT YOU REPRESENT (COLLECTIVELY, "CUSTOMER") TO THIS AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS AGREEMENT, THEN A10 NETWORKS IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND DO NOT DOWNLOAD, INSTALL OR USE THE PRODUCT.

The following terms of this End User License Agreement ("Agreement") govern Customer's access and use of the Software, except to the extent there is a separate signed agreement between Customer and A10 Networks governing Customer's use of the Software

License. Conditioned upon compliance with the terms and conditions of this Agreement, A10 Networks Inc. or its subsidiary licensing the Software instead of A10 Networks Inc. ("A10 Networks"), grants to Customer a nonexclusive and nontransferable license to use for Customer's business purposes the Software and the Documentation for which Customer has paid all required fees. "Documentation" means written information (whether contained in user or technical manuals, training materials, specifications or otherwise) specifically pertaining to the product or products and made available by A10 Networks in any manner (including on CD-Rom, or on-line).

Unless otherwise expressly provided in the Documentation, Customer shall use the Software solely as embedded in or for execution on A10 Networks equipment owned or leased by Customer and used for Customer's business purposes.

General Limitations. This is a license, not a transfer of title, to the Software and Documentation, and A10 Networks retains ownership of all copies of the Software and Documentation. Customer acknowledges that the Software and Documentation contain trade secrets of A10 Networks, its suppliers or licensors, including but not limited to the specific internal design and structure of individual programs and associated interface information. Accordingly, except as otherwise expressly provided under this Agreement, Customer shall have no right, and Customer specifically agrees not to:

a. transfer, assign or sublicense its license rights to any other person or entity, or use the Software on unauthorized or secondhand A10 Networks equipment

b. make error corrections to or otherwise modify or adapt the Software or create derivative works based upon the Software, or permit third parties to do the same

c. reverse engineer or decompile, decrypt, disassemble or otherwise reduce the Software to human readable form, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction

d. disclose, provide, or otherwise make available trade secrets contained within the Software and Documentation in any form to any third party without the prior written consent of A10 Networks. Customer shall implement reasonable security measures to protect such trade secrets.

Software, Upgrades and Additional Products or Copies. For purposes of this Agreement, "Software" and "Products" shall include (and the terms and conditions of this Agreement shall apply to) computer programs, including firmware and hardware, as provided to Customer by A10 Networks or an authorized A10 Networks reseller, and any upgrades, updates, bug fixes or modified versions thereto (collectively, "Upgrades") or backup copies of the Software licensed or provided to Customer by A10 Networks or an authorized A10 Networks reseller.

OTHER PROVISIONS OF THIS AGREEMENT:

a. CUSTOMER HAS NO LICENSE OR RIGHT TO USE ANY ADDITIONAL COPIES OR UPGRADES UNLESS CUSTOMER, AT THE TIME OF ACQUIRING SUCH COPY OR UPGRADE, ALREADY HOLDS A VALID LICENSE TO THE ORIGINAL SOFTWARE AND HAS PAID THE APPLICABLE FEE FOR THE UPGRADE OR ADDITIONAL COPIES

b. USE OF UPGRADES IS LIMITED TO A10 NETWORKS EQUIPMENT FOR WHICH CUSTOMER IS THE ORIGINAL END USER PURCHASER OR LEASEE OR WHO OTHERWISE HOLDS A VALID LICENSE TO USE THE SOFTWARE WHICH IS BEING UPGRADED

c. THE MAKING AND USE OF ADDITIONAL COPIES IS LIMITED TO NECESSARY BACKUP PURPOSES ONLY.

Term and Termination. This Agreement and the license granted herein shall remain effective until terminated. All confidentiality obligations of Customer and all limitations of liability and disclaimers and restrictions of warranty shall survive termination of this Agreement.

Export. Software and Documentation, including technical data, may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licenses to export, re-export, or import Software and Documentation.

Trademarks

A10 Networks, A10 Thunder, vThunder, the A10 logo, aACI, aCloud, ACOS, aDCS, aDNS, aELB, aFleX, aFlow, aGalaxy, aPlatform, aUSG, aVCS, aWAF, aXAPI, IDAccess, IDSENTRIE, IP to ID, SmartFlow, SoftAX, Unified Service Gateway, Virtual Chassis, VirtualADC, and VirtualN are trademarks or registered trademarks of A10 Networks, Inc. All other trademarks are property of their respective owners.

Patents Protection

A10 Networks products are protected by one or more of the following US patents and patents pending: 20120216266, 20120204236, 20120179770, 20120144015, 20120084419, 20110239289, 20110093522, 20100235880, 20100217819, 20090049537, 20080229418, 20080148357, 20080109887, 20080040789, 20070283429, 20070282855, 20070271598, 20070195792, 20070180101, 8387128, 8332925, 8312507, 8291487, 8266235, 8151322, 8079077, 7979585, 7716378, 7675854, 7647635, 7552126

Limited Warranty

Disclaimer of Liabilities. REGARDLESS OF ANY REMEDY SET FORTH FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL A10 NETWORKS OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE PRODUCT OR OTHERWISE AND EVEN IF A10 NETWORKS OR ITS SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

In no event shall A10 Networks’ or its suppliers' or licensors' liability to Customer, whether in contract, (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim or if the Software is part of another Product, the price paid for such other Product.

Customer agrees that the limitations of liability and disclaimers set forth herein will apply regardless of whetherCustomer has accepted the Software or any other product or service delivered by A10 Networks. Customer acknowledges and agrees that A10 Networks has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the parties.

The Warranty and the End User License shall be governed by and construed in accordance with the laws of the State of California, without reference to or application of choice of law rules or principles. If any portion hereof is found to be void or unenforceable, the remaining provisions of the Agreement shall remain in full force and effect. This Agreement constitutes the entire and sole agreement between the parties with respect to the license of the use of A10 Networks Products unless otherwise supersedes by a written signed agreement.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid A10 Networks Regular and Technical Support service contracts, the A10 Networks Technical Assistance Center provides support services online and over the phone.

Corporate Headquarters

A10 Networks, Inc.
3 West Plumeria Dr
San Jose, CA 95134 USA

Tel: +1-408-325-8668 (main)
Tel: +1-888-822-7210 (support – toll-free in USA)
Tel: +1-408-325-8676 (support – direct dial)
Fax: +1-408-325-8666

www.a10networks.com

Collecting System Information

The AX device provides a simple method to collect configuration and status information for Technical Support to use when diagnosing system issues.

To collect system information, use either of the following methods.

Using the GUI (Recommended)

1. Log into the GUI.
2. On the main page (Monitor Mode > Overview > Summary), click 🛟. This option downloads a text log file.
3. Email the file as an attachment to support@A10Networks.com.

Using the CLI

1. Log into the CLI.
2. Enable logging in your terminal emulation application, to capture output generated by the CLI.
3. Enter the **enable** command to access the Privileged EXEC mode of the CLI. Enter your enable password at the Password prompt.

4. Enter the **show techsupport** command.
5. After the command output finishes, save the output in a text file.
6. Email the file as an attachment to support@A10Networks.com.

Note: As an alternative to saving the output in a log file captured by your terminal emulation application, you can export the output from the CLI using the following command:

```
show techsupport export [use-mgmt-port] url
```

(For syntax information, see the *AX Series CLI Reference*.)

About This Document

This document describes features of the A10 Networks AX Series.

FIGURE 1 AX 5630 (front panel view)

Information is available for AX Series products in the following documents. These documents are included on the documentation CD shipped with your AX Series product, and also are available on the A10 Networks support site:

- *AX Series Installation Guides*
- *AX Series LOM Reference*
- *AX Series System Configuration and Administration Guide*
- *AX Series IPv4-to-IPv6 Transition Solutions Guide*
- *AX Series Traffic Logging Guide for IPv6 Migration*
- *AX Series GUI Reference*
- *AX Series CLI Reference*
- *AX Series MIB Reference*

Make sure to use the basic deployment instructions in the *AX Series Installation Guide* for your AX model, and in the *AX Series System Configuration and Administration Guide*. Also make sure to set up your device's Lights Out Management (LOM) interface, if applicable.

Note: Some guides include GUI configuration examples. In these examples, some GUI pages may have new options that are not shown in the example screen images. In these cases, the new options are not applicable to the examples. For information about any option in the GUI, see the *AX Series GUI Reference* or the GUI online help.

Audience

This document is intended for use by network architects for determining applicability and planning implementation, and for system administrators for provision and maintenance of A10 Networks AX Series products.

Documentation Updates

Updates to these documents are published periodically to the A10 Networks support site, on an updated documentation CD (posted as a zip archive). To access the latest version, please log onto your A10 support account and navigate to the following page: Support > AX Series > Technical Library.

http://www.a10networks.com

A10 Virtual Application Delivery Community

You can use your A10 support login to access the A10 Virtual Application Delivery Community (VirtualADC). The VirtualADC is an interactive forum where you can find detailed information from product specialists. You also can ask questions and leave comments. To access the VirtualADC, navigate here:

http://www.a10networks.com/adc/

System Overview

This chapter provides a brief overview of the AX Series system and features. For more information, see the other chapters in this guide.

AX Series Features

Key features of the AX Series include the following.

Note: Support for some features depends on AX model or software version. The specific models and software versions required are listed in the detailed sections for the features.

- Comprehensive IPv4/IPv6 Support
  - Easy deployment into existing infrastructures

Note: For IPv6 migration features, the AX device must be deployed in route mode (Layer 3). IPv6 migration features are not supported in transparent mode (Layer 2) deployments.

  - Static trunking, and dynamic trunking with Link Aggregation Control Protocol (LACP)
  - Standard Network Address Translation (NAT) – IPv4-IPv4, IPv6-IPv6
  - RIPv2 for IPv4 and RIPng IPv6
  - BGP4 for IPv4 and IPv6
  - OSPFv2 for IPv4, OSPFv3 for IPv6
  - IS-IS for IPv4 and IPv6
  - IPv4/IPv6 static routes
  - Bidirectional Forwarding Detection (BFD)
  - Duplicate Address Detection (DAD)
  - DHCP relay
- Advanced NAT for IPv4-IPv6 migration:
  - Large Scale NAT (LSN)
  - Dual-Stack Lite (DS-Lite)

  - DNS64 / NAT64
  - IPv6 Rapid Deployment (6rd)
  - Stateless NAT46
- High Availability (HA)
  - Active-Active and Active-Standby deployments with sub-second failover
  - Layer 4 session synchronization
  - Configuration synchronization
- Acceleration and Security
  - Traffic security
  - Management access security – Local admin database and support for optional remote RADIUS or TACACS+ AAA
  - DoS attack detection and prevention
  - Access Control Lists (ACLs)
- System Management
  - Secured console access
  - Dedicated IP and IPv6 management interfaces
  - Lights Out Management (on some models)
  - Multiple access methods – SSH, Telnet, HTTPS
  - Web-based Graphical User Interface (GUI) with language localization
  - Industry-standard Command Line Interface (CLI) support
  - On-demand backup of configuration files, logs, and system files
  - SNMP, syslog, alerting
  - Network Time Protocol (NTP) server (automatically enabled)
- Troubleshooting tools
  - Port mirroring
  - AXdebug subsystem for packet capture

ACOS Architecture

AX Series devices use embedded Advanced Core Operating System (ACOS) architecture. ACOS is built on top of a set of Symmetric Multi-Processing CPUs and uses shared memory architecture to maximize application data delivery.

AX Software Processes

The AX software performs its many tasks using the following processes:

- a10mon – Parent process of the AX device. This process is executed when the system comes up. The a10mon process does the following:
    - Brings AX user-space processes up and down.
    - Monitors all its child processes and restarts a process and all dependent processes if any of them die.
- syslogd – System logger daemon that logs kernel and system events.
- a10logd – Fetches all the logs from the AX Log database.
- a10timer – Schedules and executes scheduled tasks.
- a10stat – Monitors the status of all the main processes of the AX device, such as a10switch (on models AX2200 and higher) and a10lb.

  The a10stat process probes every thread within these processes to ensure that they are responsive. If a thread is deemed unhealthy, a10stat kills the process, after which a10mon restarts the process and other processes associated with it.
- a10switch – Contains libraries and APIs to program the Switching ASIC to perform Layer 2 and Layer 3 switching at wire speed.
- a10hm – Performs health-checking for real servers and services. This process sends pre-configured requests to external servers at pre-defined intervals. If a server or individual service does not respond, it is marked down. Once the server or service starts responding again, it is marked up.
- a10rt – Routing daemon, which maintains the routing table with routes injected from OSPF, as well as static routes.
- a10rip – Implements RIPv1 and v2 routing protocols.
- a10ospf – Implements the OSPFv2 routing protocol.
- a10snmpd – SNMPv2c and v3 agent, which services MIB requests.

- a10wa – Embedded Web Server residing on the AX device. This process serves the Web-based management Graphical User Interface (GUI).
- a10snpm_trapd – Handles SNMP traps initiated by a10lb.
- a10lb – The heart of the AX device. This process contains all the intelligence to perform network resource management.
- rimacli – This process is automatically invoked when an admin logs into the AX device through an interface address. The admin is presented a Command Line Interface (CLI) that can issue and save commands to configure the system.

Hardware Interfaces

See the *AX Series Installation Guide* for your AX model.

Software Interfaces

The AX device supports the following management interfaces:

- Graphical User Interface (GUI)
- Command Line Interface (CLI) accessible using console, Telnet, or Secure Shell (v1 and v2)
- Simple Network Management Protocol (SNMP) v1, v2c, and v3

The configuration examples in this manual show how to configure the AX Series using the CLI and GUI. For more information about the AX management interfaces, see the following documents:

- *AX Series GUI Reference*
- *AX Series CLI Reference*
- *AX Series MIB Reference*

IPv4 Preservation / IPv6 Migration Features

The AX Series provides advanced Network Address Translation features to help service providers migrate legacy IPv4 networks to IPv6.

In the current release, NAT features for IPv4-IPv6 migration are supported only on the following 64-bit ACOS AX models: AX 5630, AX 5200-11, AX 5200, AX 5100, AX 3400, AX 3200-12, AX 3530, AX 3030, AX 3000-11, AX 3000, AX 2600, and AX 2500.

Note: The AX 5630 can be deployed as a high-performance, carrier-class, efficient 64-bit AX model delivering service provider solutions for IPv4 preservation, IPv4/IPv6 translation, and full IPv6 migration.

Large Scale NAT

Large Scale NAT (LSN) enables service providers to use scarce IPv4 resources to provide NAT service for large numbers of IPv4 enterprise and residential clients.

LSN provides the following key benefits:

- Multiple clients can be allocated to a single IPv4 address.
- Standardized behavior for IPv4 NAT devices provides consistent behavior and expectations for applications.
- Efficiency through use of:
  - Hairpinning
  - Full-cone NAT
  - Fairness through admin-allocated port quotas

DNS64 / NAT64

NAT64 and DNS64 are complementary features that enable IPv6 clients to access IPv4 servers.

- DNS64 – Performs IPv4 and IPv6 DNS queries on behalf of IPv6 clients, and synthesizes IPv6 replies based on the IPv4 replies as required.
- NAT64 – Translates IPv6 packets from clients into IPv4 packets for communication with IPv4 servers. Likewise, NAT64 translates the IPv4 packets in server replies into IPv6 packets to send to the client.

Dual-Stack Lite

Dual-stack Lite (DS-Lite) enables the AX device to act as an end-point for IPv4 traffic tunneled through an IPv6 link.

Client IPv4 traffic is tunneled through the client’s gateway over the service provider’s IPv6 network. Within the service provider’s IPv6 network, the

AX device acts as an endpoint for the IPv6 tunnel, and decapsulates the client's IPv4 request from the IPv6 tunnel. The AX device then sends the request to the IPv4 server. On the return trip, the AX device encapsulates the IPv4 server reply in the IPv6 tunnel and sends the reply to the client.

DS-Lite provides the following benefits:

- Allows incremental IPv6 deployment
- Allows use of a single IPv6 network to serve IPv4 and IPv6 clients
- Simplifies carrier deployment and management

Stateless NAT46

Stateless NAT46 enables IPv4 clients to reach IPv6 servers, without the need to maintain per-connection information on the AX device.

Stateless NAT46 uses statically configured IPv4-IPv6 mappings. When an IPv4 client sends a request to a server, the destination address of the request is an IPv4 address. If the destination IPv4 address is statically mapped to the server's IPv6 address, stateless NAT46 NATs the request and forwards it to the server.

IPv6 Rapid Deployment

IPv6 Rapid Deployment (6rd) enables IPv6 clients and IPv6 servers separated by IPv4 networks to communicate, without the need to make changes to the IPv4 network.

Where Do I Start?

- To configure basic system settings, see "Basic Setup" on page 25.
- To configure network settings, see "Network Setup" on page 61.
- To configure security features, see the following chapters:
  - "Management Security Features" on page 239
  - "Traffic Security Features" on page 301
- To configure IPv4 preservation or IPv6 migration features, see the *AX Series IPv4-to-IPv6 Transition Solutions Guide*.

Basic Setup

This chapter describes how to log onto the AX device and how to configure the following basic system parameters:

- Hostname and other Domain Name Server (DNS) settings
- CLI banner messages
- Date/time settings
- System log (Syslog) settings
- Simple Network Management Protocol (SNMP) settings

After you are through with this chapter, go to “Network Setup” on page 61.

Note: The only basic parameters that you are required to configure are date/time settings. Configuring the other parameters is optional.

Note: This chapter does not describe how to access the serial console interface. For that information, see the installation guide for your AX model.

Caution: **When you make configuration changes, be sure to remember to save the changes. Unsaved configuration changes will be lost following a reboot. To save changes, click Save on the top row of the GUI window or enter the write memory command in the CLI.**

Logging On

AX Series devices provide the following management interfaces:

- Command-Line Interface (CLI) – Text-based interface in which you type commands on a command line. You can access the CLI directly through the serial console or over the network using either of the following protocols:
  - Secure protocol – Secure Shell (SSH) version 2
  - Unsecure protocol – Telnet (if enabled)
- Graphical User Interface (GUI) – Web-based interface in which you click to access configuration or management pages and type or select values to configure or manage the device. You can access the GUI using either of the following protocols:

- Secure protocol – Hypertext Transfer Protocol over Secure Socket Layer (HTTPS)
- Unsecure protocol – Hypertext Transfer Protocol (HTTP)

Note: By default, Telnet access is disabled on all interfaces, including the management interface. SSH, HTTP, HTTPS, and SNMP access are enabled by default on the management interface only, and disabled by default on all data interfaces.

Note: The AX device supports a maximum of 128 simultaneous management sessions. This includes any combination of CLI and GUI sessions.

Federal Information Processing Standards (FIPS) Compliance

To comply with FIPS security standards, beginning in AX Release 2.4.2, management access to the AX device has the following requirements:

- Web access to GUI – The browser used to access the AX GUI must support encryption keys of 128 bits or longer. Shorter encryption keys (for example, 40 bits) are not supported. The browser also must support TLS 1.0. SSL access is not supported.
- SSH access to CLI – The SSH client used to access the CLI must support SSHv2. SSHv1 is not supported. The SSHv2 client must support one of the following encryption ciphers:
  - 3des-cbc
  - aes128-cbc
  - aes192-cbc
  - aes256-cbc

  Other ciphers are not supported.

Logging Onto the CLI

Note: The AX Series provides advanced features for securing management access to the device. This section assumes that only the basic security settings are in place. (For more information about securing management access, see “Management Security Features” on page 239.)

To log onto the CLI using SSH:

1. On a PC connected to a network that can access the AX device’s management interface, open an SSH connection to the IP address of the management interface.
2. Generally, if this the first time the SSH client has accessed the AX device, the SSH client displays a security warning. Read the warning

carefully, then acknowledge the warning to complete the connection. (Press Enter.)

3. At the `login as:` prompt, enter the admin username.

4. At the `Password:` prompt, enter the admin password.

   If the admin username and password are valid, the command prompt for the User EXEC level of the CLI appears: AX>

   The User EXEC level allows you to enter a few basic commands, including some **show** commands as well as **ping** and **traceroute**.

Note: The “AX” in the CLI prompt is the hostname configured on the device, which is “AX” by default. If the hostname has already been changed, the new hostname appears in the prompt instead of “AX”.

5. To access the Privileged EXEC level of the CLI and allow access to all configuration levels, enter the **enable** command.

   At the `Password:` prompt, enter the enable password. (This is not the same as the admin password, although it is possible to configure the same value for both passwords.)

   If the enable password is correct, the command prompt for the Privileged EXEC level of the CLI appears: AX#

6. To access the global configuration level, enter the **config** command. The following command prompt appears: `AX(config)#`

Logging Onto the GUI

Web access to the AX device is supported on the Web browsers listed in Table 1.

TABLE 1 GUI Browser Support

	Platform		
Browser	**Windows**	**Linux**	**MAC**
IE 6.0-8.0	**Supported**	N/A	N/A
Firefox 2.x-3.x	**Supported**	**Supported**	N/A
Safari 3.0	Not Supported	N/A	**Supported**
Safari 2.0	Not Supported	N/A	Not Supported
Chrome 5.0	**Supported**	**Supported**	**Supported**

A screen resolution of at least 1024x768 is recommended.

1. Open one of the Web browsers listed in Table 1.

2. In the URL field, enter the IP address of the AX device's management interface.

3. If the browser displays a certificate warning, select the option to continue to the server (the AX device).

Note: To prevent the certificate warning from appearing in the future, you can install a certificate signed by a Certificate Authority. See "Replacing the Web Certificate" on page 54.

A login dialog is displayed. The name and appearance of the dialog depends on the browser you are using.

FIGURE 1 GUI Login Dialog (Internet Explorer)

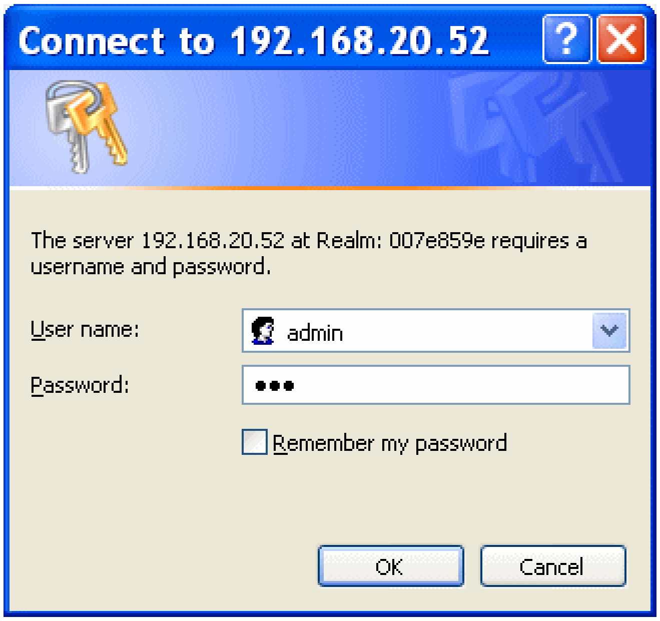

Connect to 192.168.20.52

The server 192.168.20.52 at Realm: 007e859e requires a username and password.

User name: admin

Password: •••

□ Remember my password

4. Enter your admin username and password and click OK.

Note: The default admin username and password are "admin", "a10".

The Summary page appears, showing at-a-glance information for your AX device.

You can access this page again at any time while using the GUI, by selecting Monitor > Overview > Summary.

FIGURE 2 Monitor > Overview > Summary

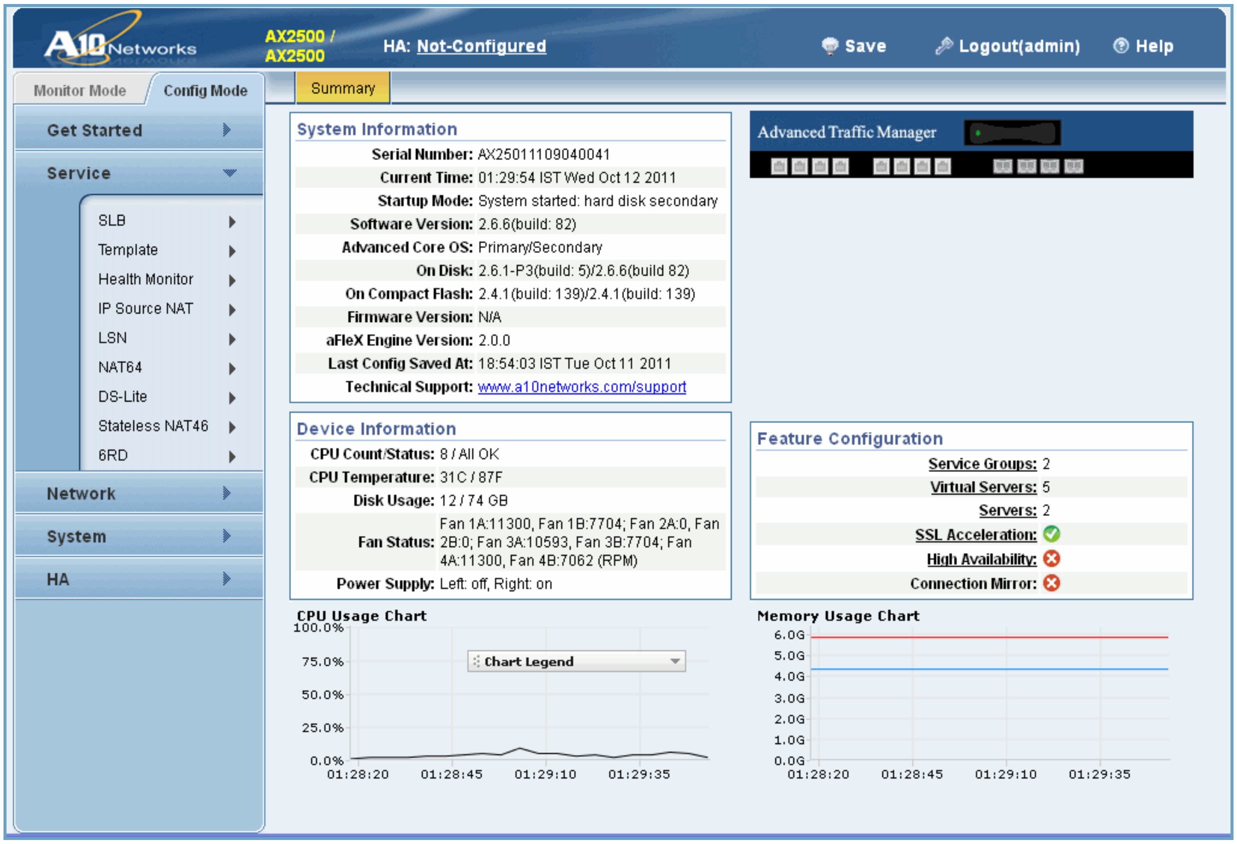

System Information

Serial Number:	AX25011109040041
Current Time:	01:29:54 IST Wed Oct 12 2011
Startup Mode:	System started: hard disk secondary
Software Version:	2.6.6(build: 82)
Advanced Core OS:	Primary/Secondary
On Disk:	2.6.1-P3(build: 5)/2.6.6(build 82)
On Compact Flash:	2.4.1(build: 139)/2.4.1(build: 139)
Firmware Version:	N/A
aFleX Engine Version:	2.0.0
Last Config Saved At:	18:54:03 IST Tue Oct 11 2011
Technical Support:	www.a10networks.com/support

Device Information

CPU Count/Status:	8 / All OK
CPU Temperature:	31C / 87F
Disk Usage:	12 / 74 GB
Fan Status:	Fan 1A:11300, Fan 1B:7704; Fan 2A:0, Fan 2B:0; Fan 3A:10593, Fan 3B:7704; Fan 4A:11300, Fan 4B:7062 (RPM)
Power Supply:	Left: off, Right: on

Feature Configuration

Service Groups:	2
Virtual Servers:	5
Servers:	2
SSL Acceleration:	✓
High Availability:	✗
Connection Mirror:	✗

Note: GUI management sessions are not automatically terminated when you close the browser window. The session remains in effect until it times out. To immediately terminate a GUI session, click Logout.

Note: For more information about the GUI, see the *AX Series GUI Reference* or the GUI online help.

Configuring Basic System Parameters

This section describes the basic system parameters and provides CLI and GUI steps for configuring them.

Note: If you plan to use the GUI, the Basic System page under Config Mode also provides configuration access to most of the system parameters described in this chapter. For information, navigate to Config Mode > Basic System, then click Help.

Setting the Hostname and Other DNS Parameters

The default hostname is “AX”. To change the hostname, use either of the following methods.

Note: Do not use a period (.) in the AX hostname. If you do use a period, the AX device will use the text after the period as the DNS suffix instead of the DNS suffix you configure.

Using the GUI

1. Select Config > Network > DNS. The DNS page appears.
2. In the Hostname field, edit the name to one that will uniquely identify this particular AX device (for example, “AX-SLB1”).
3. In the DNS Suffix field, enter the domain name to which the host (AX Series) belongs.
4. In the Primary DNS field, enter the IP address of the external DNS server the AX Series should use for resolving DNS queries.
5. In the Secondary DNS field, enter the IP address of an external backup DNS server the AX Series should use if the primary DNS server is unavailable.
6. Click OK.

Using the CLI

1. Access the global configuration level of the CLI:

```
AX>enable
Password:enable-password
AX#config
AX(config)#
```

2. Use the following command to change the hostname:

```
hostname string
```

   After you enter this command, the command prompt should change to the same value as the new hostname.

Note: The “ > ” or “ # ” character and characters in parentheses before “ # ” indicate the CLI level you are on and are not part of the hostname.

3. To set the default domain name (DNS suffix) for hostnames on the AX device, use the following command:

   ip dns suffix *string*

4. To specify the DNS servers the AX should use for resolving DNS requests, use the following command:

   ip dns {**primary** | **secondary**} *ipaddr*

   The **primary** option specifies the DNS server the AX device should always try to use first. The **secondary** option specifies the DNS server that the AX device should use if the primary DNS server is unavailable.

Setting the CLI Banners

The CLI displays banner messages when you log onto the CLI. By default, the messages shown in bold type in the following example are displayed:

```
login as: admin
Welcome to AX
Using keyboard-interactive authentication.
Password:
Last login: Thu Feb  7 13:44:32 2008 from
192.168.1.144

[type ? for help]
```

You can format banner text as a single line or multiple lines.

If you configure a banner message that occupies multiple lines, you must specify the end marker that indicates the end of the last line. The end marker is a simple string up to 2-characters long, each of the which must be an ASCII character from the following range: 0x21-0x7e.

The multi-line banner text starts from the first line and ends at the marker. If the end marker is on a new line by itself, the last line of the banner text will be empty. If you do not want the last line to be empty, put the end marker at the end of the last non-empty line.

Using the GUI

1. Select Config > System > Settings.

2. On the menu bar, select Terminal > Banner.

3. To configure a banner:

   a. Select the banner type, single-line or multi-line.

   b. If you selected multi-line, enter the delimiter value in the End Marker field.

   c. Enter the message in the Login Banner or Exec Banner field.

      If the message is a multi-line message, press Enter / Return at the end of every line. Do not type the end marker at the end of the message. The GUI automatically places the end marker at the end of the message text in the configuration.

4. If you are configuring both messages, repeat step 3 for the other message.

5. Click OK.

USING THE CLI

To change one or both banners, use the following command:

```
[no] banner {exec | login} [multi-line end-marker]
line
```

The **login** option changes the first banner, which is displayed after you enter the admin username. The **exec** option changes the second banner, which is displayed after you enter the admin password.

To use blank spaces within the banner, enclose the entire banner string with double quotation marks.

Setting Time/Date Parameters

To configure time/date parameters:

- Set the timezone.
- Set the system time and date manually or configure the AX device to use a Network Time Protocol (NTP) server. This release enabled other devices in the network to use the AX device as their NTP server.

The default timezone is Europe/Dublin, which is equivalent to Greenwich Mean Time (GMT). The time and date are not set at the factory, so must manually set them or configure NTP.

Note: You do not need to configure Daylight Savings Time. The AX device automatically adjusts the time for Daylight Savings Time based on the timezone you select.

Note: When you change the AX timezone or system time, the statistical database is cleared. This database contains general system statistics (performance, and CPU, memory, and disk utilization) and statistics. For example, in the GUI, the graphs displayed on the Monitor > Overview page are cleared.

USING THE GUI

1. Select Config > System > Time. The Date/Time page appears.
   - To set the time and date by synchronizing them with the time and date on the PC from which you are running the GUI, click Sync Local Time.
   - To configure the time and date manually:

   a. Enter the date in the Date field or select the date using the calendar.

   b. Enter the time in the Time field.
   - To set the time and date using NTP:

   a. Select the Automatically Synchronize with Internet Time Server checkbox.

   b. In the NTP Server field, enter the NTP server's IP address.

   c. In the Update System Clock Every field, enter the number of minutes you want the AX device to wait between synchronizations with the NTP server.

2. To select the timezone:

   a. Click Time Zone.

   b. From the Time Zone Name pull-down list, select the time zone.

   c. Click OK.

   d. Click Date/Time to re-display the section, if not already displayed.

3. Click OK.

USING THE CLI

To set the timezone

Enter the following command at the global configuration level of the CLI:

```
clock timezone timezone [nodst]
```

The **nodst** option disables Daylight Savings Time (DST) for the zone. DST is enabled by default, if applicable to the timezone.

To view the available timezones, enter the following command:
clock timezone ?

To configure the AX device to use NTP

1. To specify the NTP server to use, enter the following command at the global configuration level of the CLI:

   ```
   ntp server {hostname | ipaddr}
   ```

   You can configure a maximum of 4 NTP servers.

2. To enable NTP and synchronize the AX clock with the NTP server, enter the following command:

   ```
   ntp enable
   ```

To set the time and date manually

1. Return to the Privileged EXEC level of the CLI by entering the **exit** command.

2. Enter the following command at the Privileged EXEC level of the CLI:

   ```
   clock set time day month year
   ```

   Enter the time and date in the following format:

   time – *hh*:*mm*:*ss*

   day – 1-31

   month – January, February, March, ...

   year – 2008, 2009 ...

Note: The clock is based on 24 hours. For example, for 1 p.m., enter the hour as "13".

3. To display clock settings, use the following command:

   ```
   show clock [detail]
   ```

Configuring Syslog Settings

The AX device logs system events with Syslog messages. The AX device can send Syslog messages to the following places:

- Local buffer
- Console CLI session
- Console SSH and Telnet sessions
- External Syslog server
- Email address(es)
- SNMP servers (for events that are logged by SNMP traps)

Logging to the local buffer and to CLI sessions is enabled by default. Logging to other places requires additional configuration. The standard Syslog message severity levels are supported:

- Emergency – 0
- Alert – 1
- Critical – 2
- Error – 3
- Warning – 4
- Notification – 5
- Information – 6
- Debugging – 7

Table 2 lists the configurable Syslog parameters.

TABLE 2 Configurable System Log Settings

Parameter	Description	Supported Values
Disposition (message target)	Output options for each message level. For each message level, you can select which of the following output options to enable: • Console – Messages are displayed in Console sessions. • Buffered – Messages are stored in the system log buffer. • Email – Messages are sent to the email addresses in the Email To list. (See below.) • SNMP – SNMP traps are generated and sent to the SNMP receivers. • Syslog – Messages are sent to the external log servers specified in the Log Server fields. (See below.) • Monitor – Messages are displayed in Telnet and SSH sessions. **Note:** For information about emailing log messages, see "Emailing Log Messages" on page 56.	The following message levels can be individually selected for each output option: • Emergency (0) • Alert (1) • Critical (2) • Error (3) • Warning (4) • Notification (5) • Information (6) • Debug (7) Only Emergency, Alert, and Critical can be selected for SNMP. Only Emergency, Alert, Critical, and Notification can be selected for Email.
Logging Email Filter	Settings for sending log messages by email.	See "Emailing Log Messages" on page 56.
Logging Email Buffer Number		
Logging Email Buffer Time		
Facility	Standard Syslog facility to use.	Standard Syslog facilities listed in RFC 3164.
Log Buffer Entries	Maximum number of log entries the log buffer can store.	10000 to 50000 entries Default: 30000
Log Server	IP addresses or fully-qualified domain names of external log servers. Only the message levels for which Syslog is selected in the Disposition list are sent to log servers. **Note:** By default, the AX device can reach remote log servers only if they are reachable through the AX device's data ports, not the management port. To enable the AX device to reach remote log servers through the management port, see "Using the Management Interface as the Source for Management Traffic" on page 325.	Any valid IP address or fully-qualified domain name. Default: None configured
Log Server Port	Protocol port to which log messages sent to external log servers are addressed.	Any valid protocol port number Default: 514

TABLE 2 Configurable System Log Settings (Continued)

Parameter	Description	Supported Values
Email To	Email addresses to which to send log messages. Only the message levels for which Email is selected in the Disposition list are sent to log servers.	Valid email address. Click the down arrow next to the input field to add another address (up to 10). Each email address can be a maximum of 31 characters long.
SMTP Server	IP address or fully-qualified domain name of an email server using Simple Message Transfer Protocol. **Note:** By default, the AX device can reach SMTP servers only if they are reachable through the AX device's data ports, not the management port. To enable the AX device to reach SMTP servers through the management port, see "Using the Management Interface as the Source for Management Traffic" on page 325.	Any valid IP address or fully-qualified domain name. Default: None configured
SMTP Server Port	Protocol port to which email messages sent to the SMTP server are addressed.	Any valid protocol port number Default: 25
Mail From	Specifies the email From address.	Valid email address Default: Not set
Need Authentication	Specifies whether access to the SMTP server requires authentication.	Selected (enabled) or unselected (disabled) Default: disabled
Username	Username required for access to the SMTP server.	Valid username Default: Not set
Password	Password required for access to the SMTP server.	Valid password Default: Not set

Log Rate Limiting

The AX device uses a log rate limiting mechanism to ensure against overflow of external log servers and the internal logging buffer.

The rate limit for external logging is 15,000 messages per second from the AX device.

The rate limit for internal logging is 32 messages per second from the AX device.

- If the number of new messages within a one-second interval exceeds 32, then during the next one-second interval, the AX sends log messages only to the external log servers.

- If the number of new messages generated within the new one-second interval is 32 or less, then during the following one-second interval, the AX will again send messages to the local logging buffer as well as the external log server. In any case, all messages (up to 15,000 per second) get sent to the external log servers.

USING THE GUI

1. Select Config > System > Settings.
2. Select Log on the menu bar.
3. Change settings as needed. (For descriptions of the settings, see Table 2.)
4. Click OK.

USING THE CLI

1. To change the severity level of messages that are logged in the local buffer, use the following command:

   ```
   logging buffered severity-level
   ```

2. To change the severity level of messages that are logged in other places, use the following command:

   ```
   logging target severity-level
   ```

   The *target* can be one of the following:

   - **console** – Serial console
   - **email** – Email
   - **monitor** – Telnet and SSH sessions
   - **syslog** – external Syslog host
   - **trap** – external SNMP trap host

Note: Only severity levels **emergency**, **alert**, **critical**, and **notification** can be sent by email. Sending log messages by email requires additional configuration. See “Emailing Log Messages” on page 56.

3. To configure the AX device to send log messages to an external Syslog server, use the following command to specify the server:

   ```
   logging host ipaddr [ipaddr...]
   [port protocol-port]
   ```

   You can enter more than one server IP address on the same command line. The default protocol port is 514. You can specify only one protocol port with the command. All servers must use the same protocol port to listen for syslog messages.

If you use the command to add some log servers, then need to add a new log server later, you must enter all server IP addresses in the new command. Each time you enter the **logging host** command, it replaces any set of servers and syslog port configured by the previous **logging host** command.

4. To configure the AX device to send log messages by email, use the following commands to specify the email server and the email addresses:

   ```
   smtp {hostname | ipaddr} [port protocol-port]
   ```

   The **port** option specifies the protocol port to which to send email. The default is 25.

   ```
   logging email-address address [...]
   ```

   To enter more than one address, use a space between each address.

5. To send event messages to an external SNMP server, see "Enabling SNMP" on page 39.

Enabling SNMP

AX devices support the following SNMP versions: v1, v2c, v3. SNMP is disabled by default.

You can configure the AX device to send SNMP traps to the Syslog and to external trap receivers. You also can configure read (GET) access to SNMP Management Information Base (MIB) objects on the AX device by external SNMP managers.

Note: SNMP access to the AX device is read-only. SET operations (write access) are not supported.

The AX device supports the following SNMP-related RFCs:

- RFC 1157, A Simple Network Management Protocol (SNMP)
- RFC 1213, Management Information Base for Network Management of TCP/IP-based internets: MIB-II

  The sysService object returns a value that indicates the set of services the AX device offers. For the AX device, the sysService object always returns the following value: 76

  This value indicates that the AX device offers the following services:

  - datalink/subnetwork – 0x2
  - internet – 0x4
  - end-to-end – 0x8
  - applications – 0x40

For information about how the value is calculated, see the RFC.

- RFC 1850, OSPF Version 2 Management Information Base
- draft-ietf-ospf-ospfv3-mib-08, OSPF Version 3 Management Information Base
- RFC 1901, Introduction to Community-based SNMPv2
- RFC 2233, The Interfaces Group MIB using SMIv2
- RFC 2576, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework
- 2790, Host Resources MIB

  The following subtrees are supported:
  - hrSystem: .1.3.6.1.2.1.25.1
  - hrStorage: .1.3.6.1.2.1.25.2
  - hrDeviceTable: .1.3.6.1.2.1.25.3.2
  - hrProcessorTable: .1.3.6.1.2.1.25.3.3
- RFC 3410, Introduction and Applicability Statements for Internet Standard Management Framework
- RFC 3411, An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
- RFC 3412, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
- RFC 3413, Simple Network Management Protocol (SNMP) Applications
- RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
- RFC 3415, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
- RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types

SNMP Traps

Table 3 lists the SNMP traps supported by the AX device. All traps are disabled by default.

TABLE 3 AX SNMP Traps

Trap Category	Trap	Description
SNMP	Link Up	Indicates that an Ethernet interface has come up.
	Link Down	Indicates that an Ethernet interface has gone down.
System	Start	Indicates that the AX device has started.
	Shutdown	Indicates that the AX device has shut down.
	Restart	Indicates that the AX device is going to reboot or reload.
	Control CPU utilization	Indicates that the control CPU utilization is higher than 90%.*
	Data CPU utilization	Indicates that data CPU utilization is higher than 90%.*
	High Temperature	Indicates that the temperature inside the AX chassis is higher than 68 C.* If you see this trap, check for fan failure traps. Also check the installation location to ensure that the chassis room temperature is not too high (40 C or higher) and that the chassis is receiving adequate air flow.
	Fan Failure	Indicates that a system fan has failed. Contact A10 Networks.
	Power Supply Failure	Indicates that a power supply has failed. Contact A10 Networks.
	Primary Hard Disk	Indicates that the primary Hard Disk has failed or the RAID system has failed. Contact A10 Networks. In dual-disk models, the primary Hard Disk is the one on the left, as you are facing the front of the AX chassis.
	Secondary Hard Disk	Indicates that the secondary Hard Disk has failed or the RAID system has failed. Contact A10 Networks. The secondary Hard Disk is the one on the right, as you are facing the front of the AX chassis. **Note:** This trap does not apply to newer models.
	High Disk Usage	Indicates that hard disk usage on the AX device is higher than 85%.*
	High Memory Usage	Indicates that the memory usage on the AX device is higher than 95%.*
	Packet Buffer drop	Indicates that the AX device is dropping too many packets.*
Network	Trunk Ports Threshold	Indicates that the trunk ports threshold feature has disabled trunk members because the number of up ports in the trunk has fallen below the configured threshold.

TABLE 3 AX SNMP Traps (Continued)

Trap Category	Trap	Description
Routing	IS-IS	Indicates routing events for Intermediate System To Intermediate System (IS-IS) routing. (For the complete list, enter the see the GUI or enter the following: **snmp-server enable traps routing isis ?**
	OSPF	Indicates routing events for Open Shortest Path First (OSPF) routing. (For the complete list, enter the see the GUI or enter the following: **snmp-server enable traps routing ospf ?**
High Availability (HA)	Active	Indicates that the AX device is going from HA Standby mode to Active mode.
	Standby	Indicates that the AX device is going from HA Active mode to Standby mode.
	Active-Active	Indicates that an Active-Active HA configuration has been enabled.

TABLE 3 AX SNMP Traps (Continued)

Trap Category	Trap	Description
Server Load Balancing (SLB)	Server Up	Indicates that an SLB server has come up.
	Server Down	Indicates that an SLB server has gone down.
	Service Up	Indicates that an SLB service has come up.
	Service Down	Indicates that an SLB service has gone down.
	Server Connection Limit	Indicates that an SLB server has reached its configured connection limit.
	Server Connection Resume	Indicates that an SLB server has reached its configured connection-resume value.
	Service Connection Limit	Indicates that an SLB service has reached its configured connection limit.
	Service Connection Resume	Indicates that an SLB service has reached its configured connection-resume value.
	Virtual Server Connection Limit	Indicates that the connection limit configured on a virtual server has been exceeded.
	Virtual Port Connection Limit	Indicates that the connection limit configured on a virtual port has been exceeded.
	Virtual Server Connection-Rate Limit	Indicates that the connection rate limit configured on a virtual server has been exceeded.
	Virtual Port Connection-Rate Limit	Indicates that the connection rate limit configured on a virtual port has been exceeded.
	Virtual Port Up	Indicates that an SLB virtual service port has come up. An SLB virtual server's service port is up when at least one member (real server and real port) in the service group bound to the virtual port is up.
	Virtual Port Down	Indicates that an SLB virtual service port has gone down.
	Application Buffer Threshold	Indicates that the configured SLB application buffer threshold has been exceeded.*
Large Scale NAT (LSN)	Per IP Port Usage Threshold	Indicates that an LSN global IP address has reached its configured port usage threshold.
	Total Port Usage Threshold	Indicates that the AX device has reached its configured system-wide port usage threshold for Large Scale NAT (LSN) global IP addresses.
	Traffic Exceeded	Indicates that an LSN IP address pool has reached its threshold of available addresses.

* This threshold is configurable. To use the GUI, navigate to Config > System > Settings > General > Threshold. In the CLI, use the **monitor** command at the global configuration level.

SNMP Communities and Views

You can allow external SNMP managers to access the values of MIB objects from the AX device. To allow remote read-only access to AX MIB objects, configure one or both of the following types of access.

SNMP Community Strings

An SNMP *community string* is a string that an SNMP manager can present to the AX device when requesting MIB values.

Community strings are similar to passwords. You can minimize security risk by applying the same principles to selecting a community name as you would to selecting a password. Use a hard-to-guess string and avoid use of commonly used community names such as “public” or “private”.

You also can restrict access to specific Object IDs (OIDs) within the MIB, on an individual community basis. OIDs indicate the position of a set of MIB objects in the global MIB tree. The OID for A10 Networks AX Series objects is 1.3.6.1.4.1.22610.

SNMP Views

An SNMP *view* is like a filter that permits or denies access to a specific OID or portions of an OID. You can configure SNMP user groups and individual SNMP users, and allow or disallow them to read specific portions of the AX MIBs using different views.

When you configure an SNMP user group or user, you specify the SNMP version. SNMP v1 and v2c do not support authentication or encryption of SNMP packets. SNMPv3 does. You can enable authentication, encryption, or both, on an individual SNMP user-group basis when you configure the groups. You can specify the authentication method and the password for individual SNMP users when you configure the users.

SNMP Configuration Steps

To configure SNMP:

1. Optionally, configure location and contact information.
2. Optionally, configure external SNMP trap receivers.
3. Optionally, configure one or more read-only communities.
4. Optionally, configure views, groups, and users.
5. Enable the SNMP agent and SNMP traps.
6. Save the configuration changes.

You are not required to perform these configuration tasks in precisely this order. The workflow in the GUI is slightly different from the workflow shown here.

Note: By default, the AX device can reach remote logging and trap servers only if they are reachable through the AX device's data ports, not the management port. To enable the AX device to reach remote logging and trap servers through the management port, see "Using the Management Interface as the Source for Management Traffic" on page 325.

USING THE GUI

Note: To configure support for SNMPv3 or to configure views, groups, and users, use the CLI. The current release does not support configuration of SNMPv3 using the GUI.

1. Select Config > System > SNMP.
2. In the General section, configure general settings:
   a. To enable SNMP, select Enabled next to System SNMP Service.
   b. In the System Location field, enter a description of the AX device's location.
   c. In the System Contact field, enter the name or email address of the AX administrator to contact for system issues.
3. In the Community section, configure community strings:
   a. Click Community.
   b. In the SNMP Community field, enter a community name.
   c. To restrict SNMP access to a specific host or subnet, enter a hostname or an IP address and network mask in the Hostname (IP/Mask) field.

      By default, any host can access the SNMP agent on the AX device.
   d. In the Object Identifier field, enter the OID at which SNMP management applications can reach the AX device.
   e. Click Add.
   f. Repeat step b through step e for each combination of community string, management host, and OID.
4. In the Trap section, specify external trap receivers:
   a. Click Trap.
   b. In the Community field, enter the name of the community sending the traps.
   c. In the IP Address (host) field, enter the IP address or fully-qualified hostname of the SNMP trap receiver.

d. If the trap receiver does not use the standard protocol port to listen for traps, change the port number in the Port field.

e. Select SNMP the version from the Version drop-down list:

- V1
- V2c

f. Click Add to add the receiver.

g. Repeat step b through step f for each trap receiver.

5. In the Trap List section, enable traps:

   a. Click Trap List.

   b. To enable all traps, select All Traps. Otherwise, select the individual traps you want to enable.

6. Click OK.

7. To save the configuration changes, click the Save button.

Note: When there are unsaved configuration changes on the AX device, the Save button flashes.

USING THE CLI

All SNMP configuration commands are available at the global configuration level of the CLI.

1. To configure location and contact information, use the following commands:

```
snmp-server location location

snmp-server contact contact-name
```

2. To configure external SNMP trap receivers, use the following command:

```
snmp-server host trap-receiver
[version {v1 | v2c}]
community-string
[udp-port port-num]
```

3. To configure one or more read-only communities, use the following command:

```
snmp-server community read ro-community-string
[oid oid-value]
[remote {hostname | ipaddr mask-length |
ipv6-addr/prefix-length}]
```

4. To configure views, groups, and users, use the following commands:

```
snmp-server view view-name oid [oid-mask]
{included | excluded}

snmp-server group group-name
{v1 | v2c | v3 {auth | noauth | priv}}
read view-name

snmp-server user username group groupname
{v1 | v2 | v3 [auth {md5 | sha} password
[encrypted]]}
```

5. To enable the SNMP agent and SNMP traps, use the following command:

```
[no] snmp-server enable
     [
      traps [
             routing {isis | ospf} [trap-name] |
             snmp [trap-name] |
             ha [trap-name] |
             network [trap-name] |
             slb [trap-name] |
             lsn [trap-name] |
             system [trap-name]
             ]
      ]
```

6. To save the configuration changes, use the following command at the Privileged EXEC level or any configuration level of the CLI:

```
write memory
```

Configuration Examples

The following examples show how to configure the system settings described in this chapter.

GUI Example

The following examples show the GUI screens used for configuration of the basic system settings described in this chapter.

Note: The GUI does not support configuration of SNMPv3 settings.

FIGURE 3 Config > Network > DNS > DNS

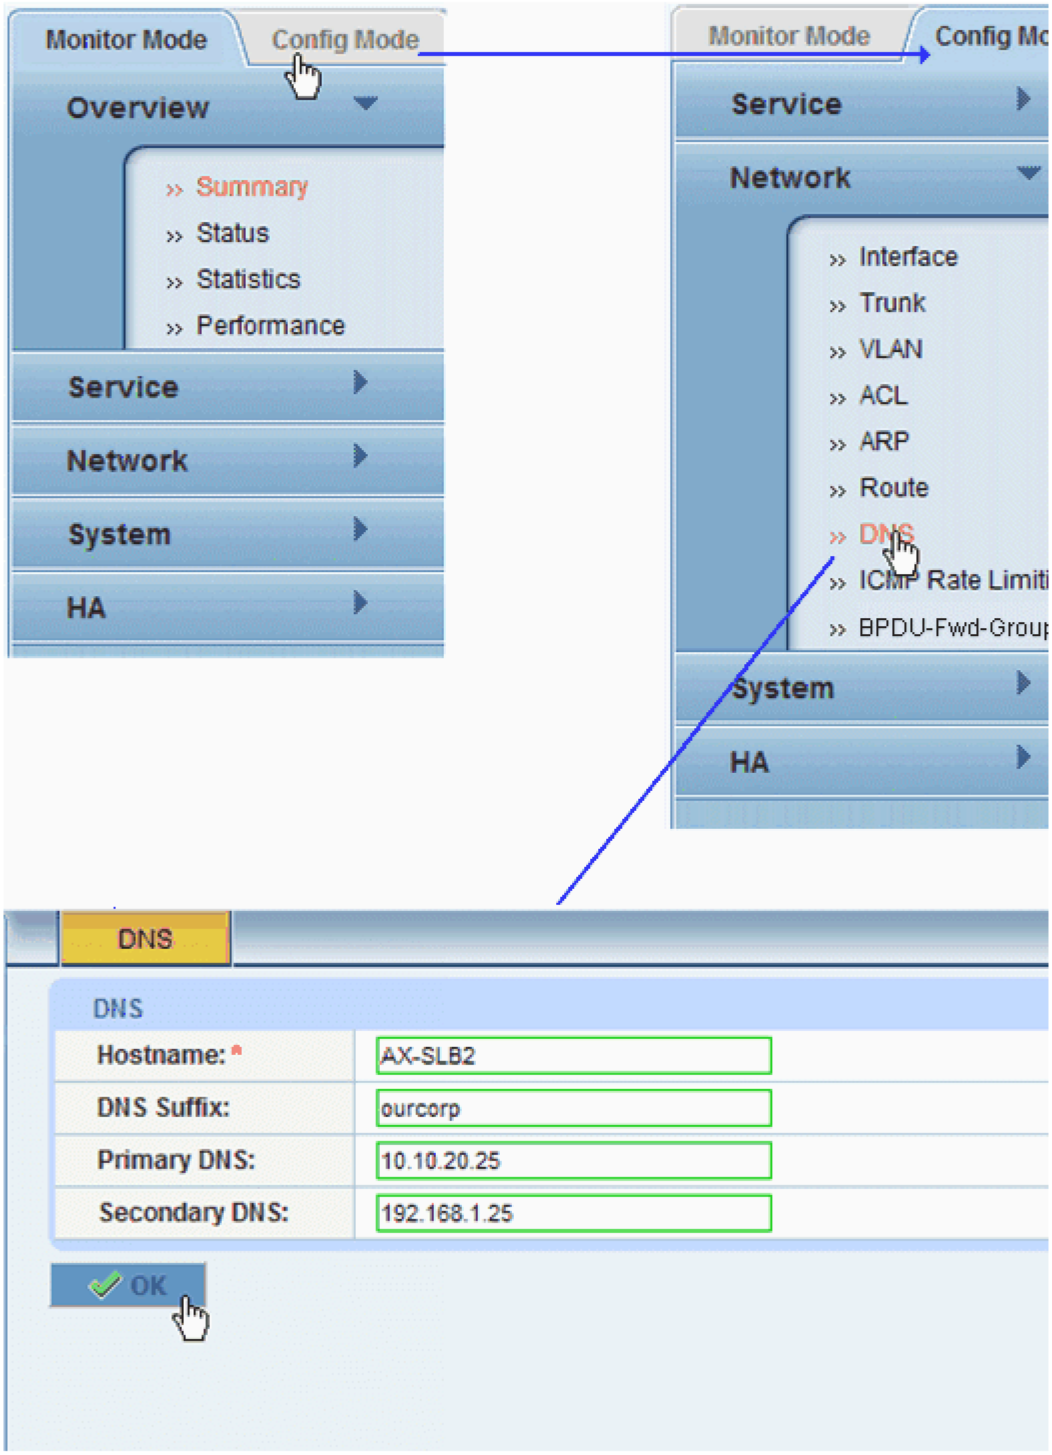

Hostname: *	AX-SLB2
DNS Suffix:	ourcorp
Primary DNS:	10.10.20.25
Secondary DNS:	192.168.1.25

FIGURE 4 Config > System > Time

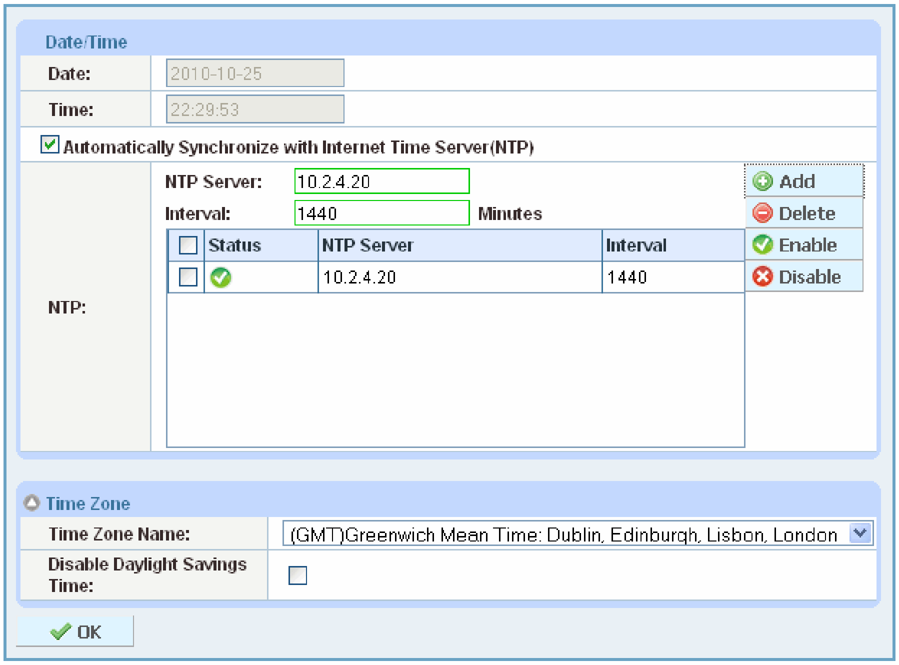

Date/Time

Date:	2010-10-25
Time:	22:29:53

☑ Automatically Synchronize with Internet Time Server(NTP)

NTP:

NTP Server: 10.2.4.20

Interval: 1440 Minutes

□	Status	NTP Server	Interval
□	✓	10.2.4.20	1440

Time Zone

Time Zone Name:	(GMT)Greenwich Mean Time: Dublin, Edinburgh, Lisbon, London
Disable Daylight Savings Time:	□

FIGURE 5 Config > System > Settings > Log

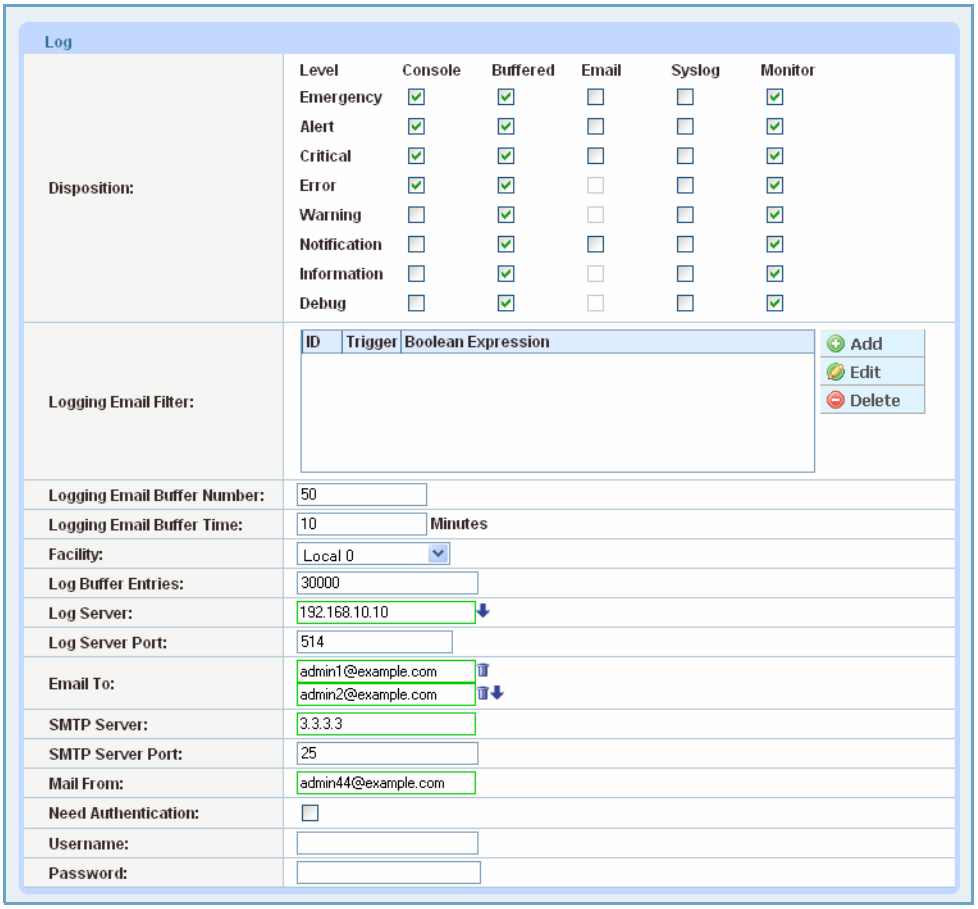

Log

Disposition:

Level	Console	Buffered	Email	Syslog	Monitor
Emergency	☑	☑	□	□	☑
Alert	☑	☑	□	□	☑
Critical	☑	☑	□	□	☑
Error	☑	☑	□	□	☑
Warning	□	☑	□	□	☑
Notification	□	☑	□	□	☑
Information	□	☑	□	□	☑
Debug	□	☑	□	□	☑

Logging Email Filter:

ID	Trigger	Boolean Expression

Logging Email Buffer Number:	50
Logging Email Buffer Time:	10 Minutes
Facility:	Local 0
Log Buffer Entries:	30000
Log Server:	192.168.10.10
Log Server Port:	514
Email To:	admin1@example.com admin2@example.com
SMTP Server:	3.3.3.3
SMTP Server Port:	25
Mail From:	admin44@example.com
Need Authentication:	□
Username:	
Password:	

FIGURE 6 Config > System > SNMP

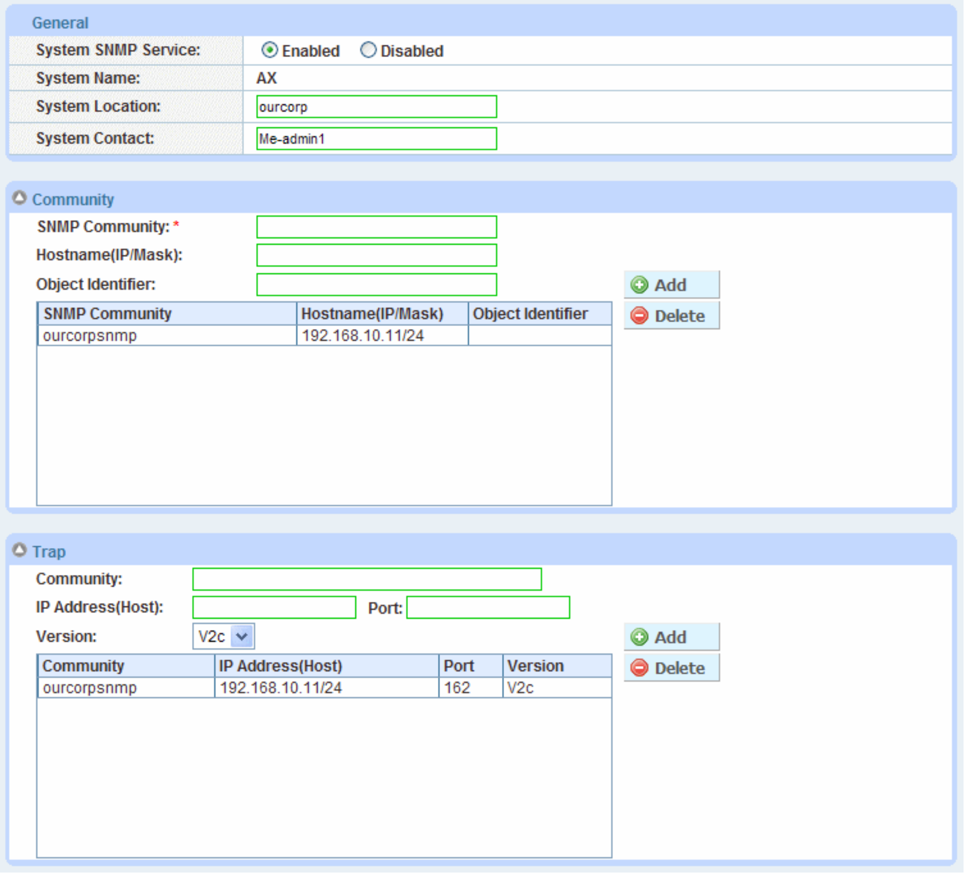

System SNMP Service:	● Enabled ○ Disabled
System Name:	AX
System Location:	ourcorp
System Contact:	Me-admin1

SNMP Community	Hostname(IP/Mask)	Object Identifier
ourcorpsnmp	192.168.10.11/24	

Version: V2c

Community	IP Address(Host)	Port	Version
ourcorpsnmp	192.168.10.11/24	162	V2c

FIGURE 7 Config > System > SNMP > Trap List

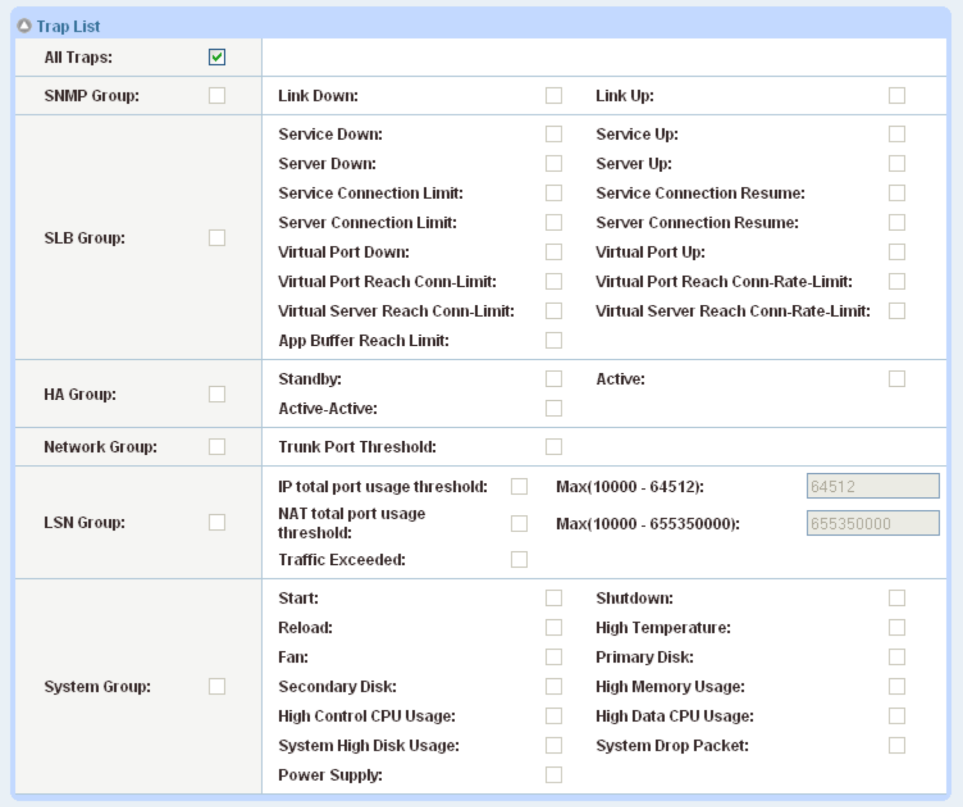

Trap List

All Traps:	☑				
SNMP Group:	□	Link Down:	□	Link Up:	□
SLB Group:	□	Service Down:	□	Service Up:	□
		Server Down:	□	Server Up:	□
		Service Connection Limit:	□	Service Connection Resume:	□
		Server Connection Limit:	□	Server Connection Resume:	□
		Virtual Port Down:	□	Virtual Port Up:	□
		Virtual Port Reach Conn-Limit:	□	Virtual Port Reach Conn-Rate-Limit:	□
		Virtual Server Reach Conn-Limit:	□	Virtual Server Reach Conn-Rate-Limit:	□
		App Buffer Reach Limit:	□		
HA Group:	□	Standby:	□	Active:	□
		Active-Active:	□		
Network Group:	□	Trunk Port Threshold:	□		
LSN Group:	□	IP total port usage threshold:	□	Max(10000 - 64512):	64512
		NAT total port usage threshold:	□	Max(10000 - 655350000):	655350000
		Traffic Exceeded:	□		
System Group:	□	Start:	□	Shutdown:	□
		Reload:	□	High Temperature:	□
		Fan:	□	Primary Disk:	□
		Secondary Disk:	□	High Memory Usage:	□
		High Control CPU Usage:	□	High Data CPU Usage:	□
		System High Disk Usage:	□	System Drop Packet:	□
		Power Supply:	□		

FIGURE 8 Save Button

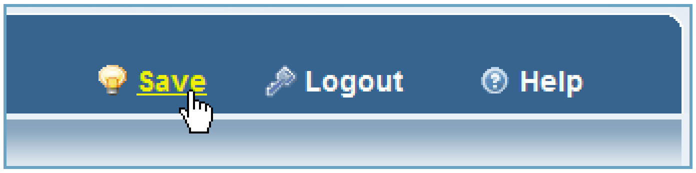

CLI Example

The following commands log onto the CLI, access the global configuration level, and set the hostname and configure the other DNS settings:

```
login as: admin
Welcome to AX
Using keyboard-interactive authentication.
Password:********
Last login: Tue Jan 13 19:51:56 2009 from 192.168.1.144

[type ? for help]

AX>enable
Password:********
AX#config
AX(config)#hostname AX-SLB2
AX-SLB2(config)#ip dns suffix ourcorp
AX-SLB2(config)#ip dns primary 10.10.20.25
AX-SLB2(config)#ip dns secondary 192.168.1.25
```

The following examples set the login banner to “welcome to login mode” and set the EXEC banner to “welcome to exec mode”:

```
AX-SLB2(config)#banner login “welcome to login mode”
AX-SLB2(config)#banner exec “welcome to exec mode”
```

The following commands set the timezone and NTP parameters:

```
AX-SLB2(config)#clock timezone ?
  Pacific/Midway              (GMT-11:00)Midway Island, Samoa
  Pacific/Honolulu            (GMT-10:00)Hawaii
  America/Anchorage           (GMT-09:00)Alaska
  America/Tijuana             (GMT-08:00)Pacific Time(US & Canada); Tijuana
  America/Los_Angeles         (GMT-08:00)Pacific Time
...

AX-SLB2(config)#clock timezone America/Los_Angeles
AX-SLB2(config)#ntp server 10.1.4.20
AX-SLB2(config)#ntp server enable
```

The following commands configure the AX device to send system log messages to an external syslog server and to email Emergency messages to the system admins. In this example, the message levels sent to the external server are left at the default, Error (3) and above. By default, the same message levels are sent to the management terminal in CLI sessions. The message level emailed to admins is set to Emergency (0) messages only.

```
AX-SLB2(config)#logging host 192.168.10.10
AX-SLB2(config)#smtp ourmailsrvr
AX-SLB2(config)#logging email-address admin1@example.com admin2@exam-
ple.com
AX-SLB2(config)#logging email 0
```

The following commands enable SNMP and all traps, configure the AX device to send traps to an external trap receiver, and configure a community string for use by external SNMP managers to read MIB data from the AX device.

```
AX-SLB2(config)#snmp-server location ourcorp-HQ
AX-SLB2(config)#snmp-server contact Me_admin1
AX-SLB2(config)#snmp-server enable trap
AX-SLB2(config)#snmp-server community read ourcorpsnmp
AX-SLB2(config)#snmp-server host 192.168.10.11 ourcorpsnmp
```

The following command saves the configuration changes to the startup-config. This is the file from which the AX device loads the configuration following a reboot.

```
AX-SLB2(config)#write memory
```

Replacing the Web Certificate

You can replace the web certificate shipped with the AX device. Replacing the certificate with a CA-signed certificate prevents the certificate warning from being displayed by your browser when you log onto the GUI.

USING THE GUI

1. Select Config > System > Settings > Web Certificate.
2. Click Add.
3. Select the file type from the Type drop-down list:
   - Child Certificate
   - Child Key
   - Chain Certificate

4. Select the location of the file to be imported:
   - Local – The file is on the PC you are using to run the GUI, or is on another PC or server in the local network. Go to step 5.
   - Remote – The file is on a remote server. Go to step 7.

5. Click Browse and navigate to the location of the class list.

6. Click Open. The path and filename appear in the Source field. Go to step 12.

7. To use the management interface as the source interface for the connection to the remote device, select Use Management Port. Otherwise, the AX device will attempt to reach the remote server through a data interface.

8. Select the file transfer protocol: FTP, TFTP, RCP, or SCP.

9. In the Host field, enter the directory path and filename.

10. If needed, change the protocol port number n the port field. By default, the default port number for the selected file transfer protocol is used.

11. In the User and Password fields, enter the username and password required for access to the remote server.

12. Click OK.

13. Click Reload Web Server.

USING THE CLI

The current release does not support this option using the CLI.

Emailing Log Messages

You can configure the AX device to email log messages, using email log filters. By default, emailing of log messages is disabled.

Log email filters consist of the following parameters:

- Filter ID – Filter number, 1-8.
- Conditions – One or more of the following:
  - Severity – Severity levels of messages to send in email. If you do not specify a message level, messages of any severity level match the filter and can be emailed.
  - Software Module – Software modules for which to email messages. Messages are emailed only if they come from one of the specified software modules. If you do not specify a software module, messages from all modules match the filter and can be emailed.
  - Regular Expression (Patterns and Operators) – Message text to match on. Standard regular expression syntax is supported. Only messages that meet the criteria of the regular expression can be emailed. The regular expression can be a simple text string or a more complex expression using standard regular expression logic. If you do not specify a regular expression, messages with any text match the filter and can be emailed.

    The operators (AND, OR, NOT) specify how the conditions should be compared. (See "“Boolean Operators” on page 56".)
- Trigger option – Specifies whether to buffer matching messages or send them immediately.

Boolean Operators

A logging email filter consists of a set of conditions joined by Boolean expressions (AND / OR / NOT).

The CLI Boolean expression syntax is based on Reverse Polish Notation (also called Postfix Notation), a notation method that places an operator (AND, OR, NOT) after all of its operands (in this case, the conditions list).

After listing all the conditions, specify the Boolean operator(s). The following operators are supported:

- AND – All conditions must match in order for a log message to be emailed.
- OR – Any one or more of the conditions must match in order for a log message to be emailed.
- NOT – A log message is emailed only if it does not match the conditions

(For more information about Reverse Polish Notation, see the following link: http://en.wikipedia.org/wiki/Reverse_Polish_notation.)

Using the GUI

1. Select Config > System > Settings > Log.

2. In the Logging Email Filter section, click Add. A configuration page for the filter appears.

3. In the ID field, enter the filter ID, 1-8.

4. To immediately send matching messages in an email instead of buffering them, select Trigger. Otherwise, matching messages are buffered until the message buffer becomes full or the send timer for emailed log messages expires.

5. Construct the rest of the filter by selecting the conditions.

Note: The conditions must be selected in the order described here. Otherwise, the filter will be invalid. If you accidentally configure an invalid filter, you can click Clear to remove the filter conditions and start again.

   a. Select the message severity level from the Level drop-down list, and click Add. To add more severity levels, repeat this step for each severity level.

   b. Optionally, select a software module from the Module drop-down list, and click Add. To add more modules, repeat this step for each module.

   c. Optionally, enter a regular expression in the Pattern field to specify message text to match on, and click Add.

   d. Select the operator from the Operator drop-down list, and click Add.

6. Click OK. The new filter appears in the Logging Email Filter section on the Log page.

7. Optionally, to change the maximum number of log messages to buffer before sending them in email, edit the number in the Logging Email Buffer Number field. You can specify 16-256 messages. The default is 50.

8. Optionally, to change the number of minutes the AX device waits before sending all buffered messages, edit the number in the Logging Email Buffer Time field. This option takes affect if the buffer does not reach

the maximum number of messages allowed. You can specify 10-1440 minutes. The default is 10.

9. When finished configuring log settings, click OK.

FIGURE 9 Config > System > Settings > Log - Add (Logging Email Filter section)

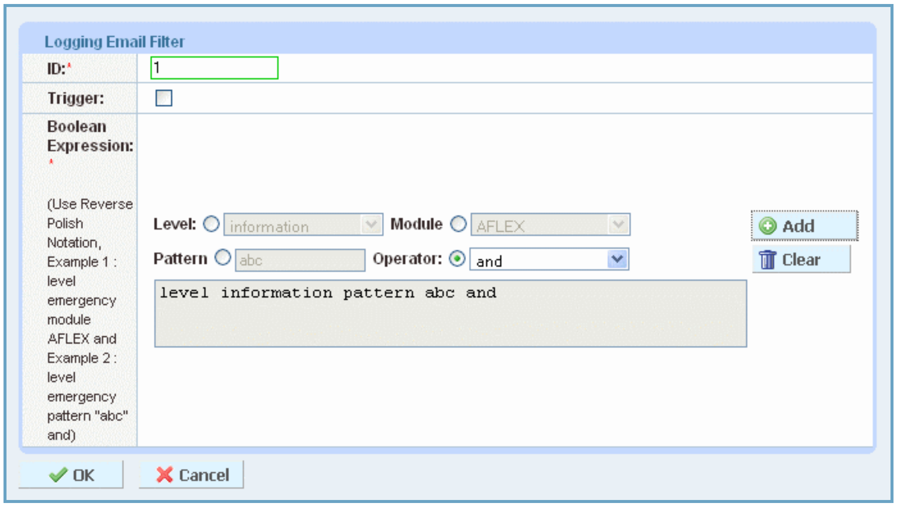

Logging Email Filter

ID:*	1
Trigger:	□
Boolean Expression: * (Use Reverse Polish Notation, Example 1 : level emergency module AFLEX and Example 2 : level emergency pattern "abc" and)	Level: ○ information Module ○ AFLEX Pattern ○ abc Operator: ● and level information pattern abc and

FIGURE 10 Config > System > Settings > Log (Logging Email Filter added)

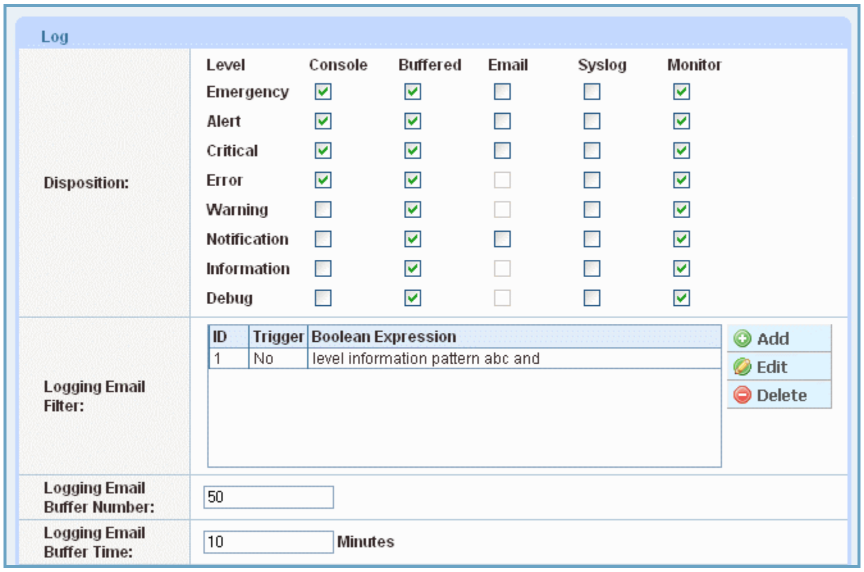

Log

Disposition:

Level	Console	Buffered	Email	Syslog	Monitor
Emergency	☑	☑	□	□	☑
Alert	☑	☑	□	□	☑
Critical	☑	☑	□	□	☑
Error	☑	☑	□	□	☑
Warning	□	☑	□	□	☑
Notification	□	☑	□	□	☑
Information	□	☑	□	□	☑
Debug	□	☑	□	□	☑

Logging Email Filter:

ID	Trigger	Boolean Expression
1	No	level information pattern abc and

Logging Email Buffer Number: 50

Logging Email Buffer Time: 10 Minutes

Using the CLI

To configure log email settings, use the following commands at the global configuration level of the CLI:

```
[no] logging email buffer [number num]
[time minutes]
```

This command configures message buffering. The **number** option specifies the maximum number of messages to buffer. You can specify 16-256. The default is 50. The **time** option specifies how long to wait before sending all buffered messages, if the buffer contains fewer than the maximum allowed number of messages. You can specify 10-1440 minutes. The default is 10.

Whenever an email is triggered, the email will contain all buffered log messages.

```
[no] logging email filter filter-num conditions
operators [trigger]
```

The *filter-num* option specifies the filter number, and can be 1-8.

The *conditions* list can contain one or more of the following:

- **level** *severity-levels* – Specifies the severity levels of messages to send in email. You can specify the severity levels by number (0-7) or by name: **emergency**, **alert**, **critical**, **error**, **warning**, **notification**, **information**, or **debugging**.
- **mod** *software-module-name* – Specifies the software modules for which to email messages. Messages are emailed only if they come from one of the specified software modules. For a list of module names, enter **?** instead of a module name, and press Enter.
- **pattern** *regex* – Specifies the string requirements. Standard regular expression syntax is supported. Only messages that meet the criteria of the regular expression will be emailed. The regular expression can be a simple text string or a more complex expression using standard regular expression logic.

The *operators* are a set of Boolean operators (AND, OR, NOT) that specify how the conditions should be compared.

The **trigger** option immediately sends the matching messages in an email instead of buffering them. If you omit this option, the messages are buffered based on the **logging email buffer** settings.

Considerations

- You can configure up to 8 filters. The filters are used in numerical order, starting with filter 1. When a message matches a filter, the message will be emailed based on the buffer settings. No additional filters are used to examine the message.
- A maximum of 8 conditions are supported in a filter.
- The total number of conditions plus the number of Boolean operators supported in a filter is 16.
- For backward compatibility, the following syntax from previous releases is still supported:

  logging email *severity-level*

  The *severity-level* can be one or more of the following: **0**, **1**, **2**, **5**, **emergency**, **alert**, **critical**, **notification**.

  The command is treated as a special filter. This filter is placed into effect only if the command syntax shown above is in the configuration. The filter has an implicit trigger option for emergency, alert, and critical messages, to emulate the behavior in previous releases.

CLI Example

The following command configures the AX device to buffer log messages to be emailed. Messages will be emailed only when the buffer reaches 32 messages, or 30 minutes passes since the previous log message email, whichever happens first.

```
AX(config)#logging email buffer number 32 time 30
```

The following command resets the buffer settings to their default values.

```
AX(config)#no logging email buffer number time
```

The following command configures a filter that matches on log messages if they are information-level messages ***and*** contain the string “abc”. The **trigger** option is not used, so the messages will be buffered rather than emailed immediately.

```
AX(config)#logging email filter 1 level information pattern "abc" and
```

The following command reconfigures the filter to immediately email matching messages.

```
AX(config)#logging email filter 1 level information pattern "abc" and trigger
```

Network Setup

This chapter describes how to add the AX device to your network.

Note: The AX device must be deployed in route mode (Layer 3). IPv6 migration features are not supported in transparent mode (Layer 2) deployments.

Overview

AX Series devices can be inserted into your network with minimal or no changes to your existing network. You can insert the AX device into your network as a Layer 3 router (route mode).

The AX device has a dedicated Ethernet management interface, separate from the Ethernet data interfaces. You can assign an IPv4 address and an IPv6 address to the management interface.

This chapter provides the following configuration examples:

- “Management Interface Configuration” on page 64
- “Route Mode Deployment” on page 68

Trunk Support

The AX device supports aggregation of multiple Ethernet data ports into logical links, called “trunks”. Trunks can enhance performance by providing higher throughput and greater link reliability.

You can configure the following types of trunks:

- Static trunks – You can configure up to 16 static trunks. Each trunk can contain 2-8 Ethernet data ports.
- Dynamic trunks – You can enable Link Aggregation Control Protocol (LACP) on Ethernet data interfaces, to make those interfaces candidate members of dynamically configured trunks.

Although this chapter does not show trunking examples, the following chapter does: “Link Trunking” on page 79

Virtual LAN Support

A VLAN is a Layer 2 broadcast domain. MAC-layer broadcast traffic can be flooded within the VLAN but does not cross to other VLANs. For traffic to go from one VLAN to another, it must be routed.

You can segment the AX device into multiple VLANs. Each Ethernet data port can be a member of one or more VLANs, depending on whether the port is tagged or untagged:

- Tagged – Tagged ports can be members of multiple VLANs. The port can recognize the VLAN to which a packet belongs based on the VLAN tag included in the packet.
- Untagged – Untagged ports can belong to only a single VLAN. By default, all AX Ethernet data ports are untagged members of VLAN 1.

Note: A port actually can be an untagged member of a single VLAN, and also a tagged member of one or more other VLANs. For example, to add a port to a new VLAN, it is not required either to remove the port from VLAN 1, or to change its status in VLAN 1 from untagged to tagged.

Default VLAN (VLAN 1)

By default, all the AX device’s Ethernet data ports are members of a single virtual LAN (VLAN), VLAN 1.

On a new or unconfigured AX device, as soon as you configure an IP interface on any individual Ethernet data port or trunk interface, Layer 2 forwarding on VLAN 1 is disabled.

When Layer 2 forwarding on VLAN 1 is disabled, broadcast, multicast, and unknown unicast packets are dropped instead of being forwarded. Learning is also disabled on the VLAN. However, packets for the AX device itself (ex: LACP, HA, OSPF) are not dropped.

To re-enable Layer 2 forwarding on VLAN 1, use the following command at the global configuration level of the CLI:
enable-def-vlan-l2-forwarding

Note: Configuring an IP interface on an individual Ethernet interface indicates you are deploying in route mode (also called “gateway mode”). If you deploy in transparent mode instead, in which the AX device has a single IP address for all data interfaces, Layer 2 forwarding is left enabled by default on VLAN 1.

Virtual Ethernet Interfaces

You can configure IP interfaces on VLANs. To configure an IP interface on a VLAN, add a Virtual Ethernet (VE) interface to the VLAN, then assign the IP interface to the VE.

Each VLAN can have one VE. The VE ID must be the same as the VLAN ID. For example, VLAN 2 can have VE 2, VLAN 3 can have VE 3, and so on.

IP Subnet Support

The AX device has a management interface and data interfaces. The management interface is a physical Ethernet port. A data interface is a physical Ethernet port, a trunk group, or a Virtual Ethernet (VE) interface.

The management interface can have a single IPv4 address and a single IPv6 address.

An AX device can have separate IP addresses on each data interface. No two interfaces can have IP addresses that are in the same subnet. This applies to the management interface and all data interfaces.

Routing Support

The AX device supports the following dynamic routing protocols:

- Routing Information Protocol (RIP) version 2 for IPv4 and RIP Next Generation (RIPng) for IPv6
- Open Shortest Path First (OSPF) version 2 for IPv4 and OSPFv3 for IPv6
- Intermediate System to Intermediate System (IS-IS) for IPv4 and IPv6
- Border Gateway Protocol version 4 with multiprotocol extensions (BGP4+), for IPv6 and IPv4

For OSPF information, see the following:

- "Open Shortest Path First (OSPF)" on page 107
- *AX Series CLI Reference*

For RIP, IS-IS, and BGP information, see the *AX Series CLI Reference*.

Gateway Health Monitoring

You can configure the AX device to monitor the health (Layer 3 reachability) of its nexthop gateways. Gateway health monitoring provides more efficient routing by sending traffic only to available gateways.

For more information, see “Gateway Health Monitoring” on page 103.

High Availability

You can deploy the AX device as a single unit or as a High Availability (HA) pair. Deploying a pair of AX devices in an HA configuration provides an extra level of redundancy to help ensure your site remains available to clients. For simplicity, the examples in this chapter show deployment of a single AX device. For information about HA, see “High Availability” on page 169.

Management Interface Configuration

The management interface (MGMT) is an Ethernet interface to which you can assign a single IPv4 address and a single IPv6 address. The management interface is separate from the Ethernet data interfaces.

Figure 11 shows an example of the management interface on an AX Series device.

FIGURE 11 AX Deployment Example – Management Interface

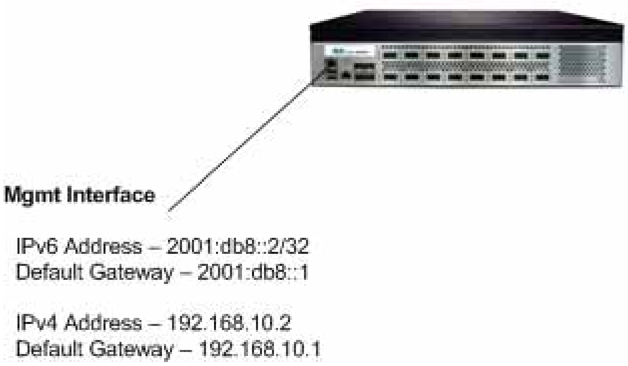
Mgmt Interface
IPv6 Address – 2001:db8::2/32
Default Gateway – 2001:db8::1
IPv4 Address – 192.168.10.2
Default Gateway – 192.168.10.1

Note: By default, the AX device attempts to use a route from the main route table for management connections originated on the AX device. You can enable the AX device to use the management route table to initiate management connections instead. (For information, see “Using the Management Interface as the Source for Management Traffic” on page 325.)

Configuration Example

This section shows the GUI screens and CLI commands needed to configure the management interface as shown in Figure 11.

Using the GUI

1. Log into the GUI.

Note: Unless you have already configured an IP interface, navigate to the default IP address: http://172.31.31.31.

2. Select Config > Network > Interface > Management. (See Figure 12.)

   Figure 12 shows the GUI screen used to configure the management IP addresses shown in Figure 11. Here and elsewhere in this guide, the command paths used to access a GUI screen are listed in the figure caption.

3. In the IPv4 section, enter the IPv4 address, network mask, and default gateway address.

4. In the IPv6 section, enter the IPv6 address, prefix length, and default gateway address.

5. Click OK.

FIGURE 12 Config > Network > Interface > Management

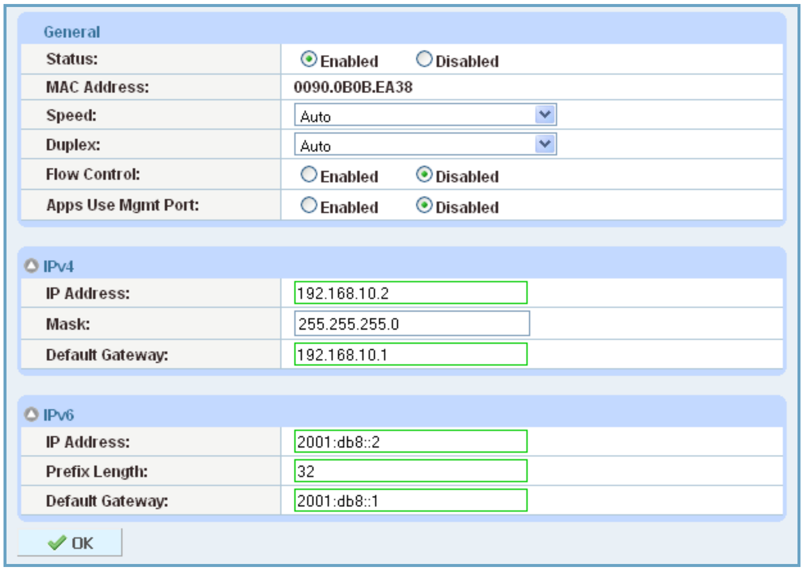

General

Status:	● Enabled ○ Disabled
MAC Address:	0090.0B0B.EA38
Speed:	Auto
Duplex:	Auto
Flow Control:	○ Enabled ● Disabled
Apps Use Mgmt Port:	○ Enabled ● Disabled

IPv4

IP Address:	192.168.10.2
Mask:	255.255.255.0
Default Gateway:	192.168.10.1

IPv6

IP Address:	2001:db8::2
Prefix Length:	32
Default Gateway:	2001:db8::1

Using the CLI

The following commands access the global configuration level of the CLI:

```
login as: admin
Welcome to AX
Using keyboard-interactive authentication.
Password:***
[type ? for help]
AX>enable
Password:(just press Enter on a new system)
AX#
AX#config
AX(config)#
```

The following commands configure IPv6 access on the management interface:

```
AX(config)#interface management
AX(config-if:mgmt)#ipv6 address 2001:db8::2/32
AX(config-if:mgmt)#ipv6 default-gateway 2001:db8::1
```

The following commands configure IPv4 access on the management interface:

```
AX(config)#interface management
AX(config-if:mgmt)#ip address 192.168.2.228 /24
AX(config-if:mgmt)#ip default-gateway 192.168.2.1
```

The following commands verify the configuration:

```
AX(config-if:mgmt)#show interface management
GigabitEthernet 0 is up, line protocol is up.
  Hardware is GigabitEthernet, Address is 0090.0b0b.ea38
  Internet address is 192.168.10.2, Subnet mask is 255.255.255.0
  Internet V6 address is 2001:db8::2/32
  Configured Speed auto, Actual 1000, Configured Duplex auto, Actual fdx
  Flow Control is disabled, IP MTU is 1500 bytes
  781 packets input,  58808 bytes
  Received 33 broadcasts,  Received 66 multicasts,  Received 662 unicasts
  0 input errors,  0 CRC  0 frame
  0 runts  0 giants
  924 packets output 3549 bytes
  Transmitted 157 broadcasts  7 multicasts  770 unicasts
  0 output errors  0 collisions
```

Route Mode Deployment

Figure 13 shows an example of an AX device deployed in route mode.

Note: Route mode is also called “gateway” mode.

FIGURE 13 AX Deployment Example – Route Mode

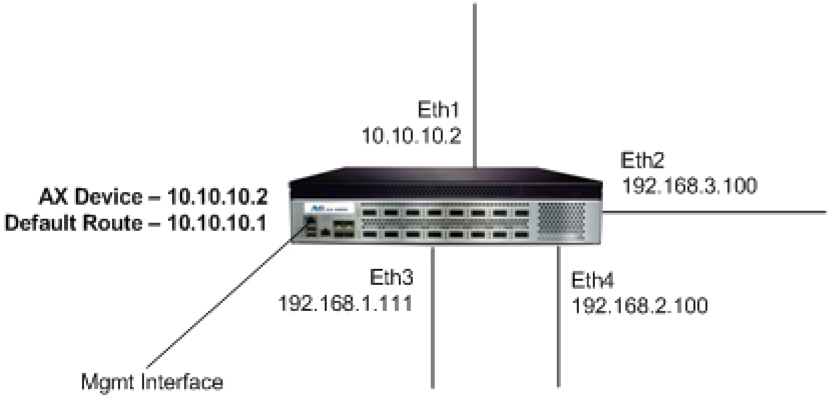

In this example, the AX device has separate IP interfaces in different subnets on each of the interfaces connected to the network. The AX device can be configured with static IP routes and can be enabled to run OSPF and IS-IS. In this example, a static route is configured to use as the default route through 10.10.10.1.

Although this example shows single physical links, you could use trunks as physical links. You also could use multiple VLANs. In this case, the IP addresses would be configured on Virtual Ethernet (VE) interfaces, one per VLAN, instead of being configured on individual Ethernet ports.

Since the AX device is a router in this deployment, downstream devices can use the AX device as their default gateway. For example, devices connected to Ethernet port 2 would use 192.168.3.100 as their default gateway, devices connected to port 3 would use 192.168.1.111 as their default gateway, and so on.

If a pair of AX devices in a High Availability (HA) configuration is used, the downstream devices would use a floating IP address shared by the two AX devices as their default gateway. (For more on HA, see “High Availability” on page 169.)

Configuration Example

This section shows the GUI screens and CLI commands needed to implement the configuration shown in Figure 13.

USING THE GUI

1. Select Config > Network > Interface > LAN.
2. Click on the interface name (for example, “e1”). The configuration page for the port appears. (See Figure 14.)
3. To assign an IPv4 address, click “IPv4” to display the configuration fields for that section, and enter the address information..
4. To assign an IPv6 address, click “IPv6” to display the configuration fields for that section, and enter the address information..
5. Click OK.

Configuring the Default Route

1. Select Config > Network > Route > IPv4 Static (or IPv6 Static) > ⊕ .
2. Enter the route information. (For an IPv4 example, see Figure 15.)
3. Click OK.

FIGURE 14 Config > Network > Interface > LAN - Ethernet data port 1 selected

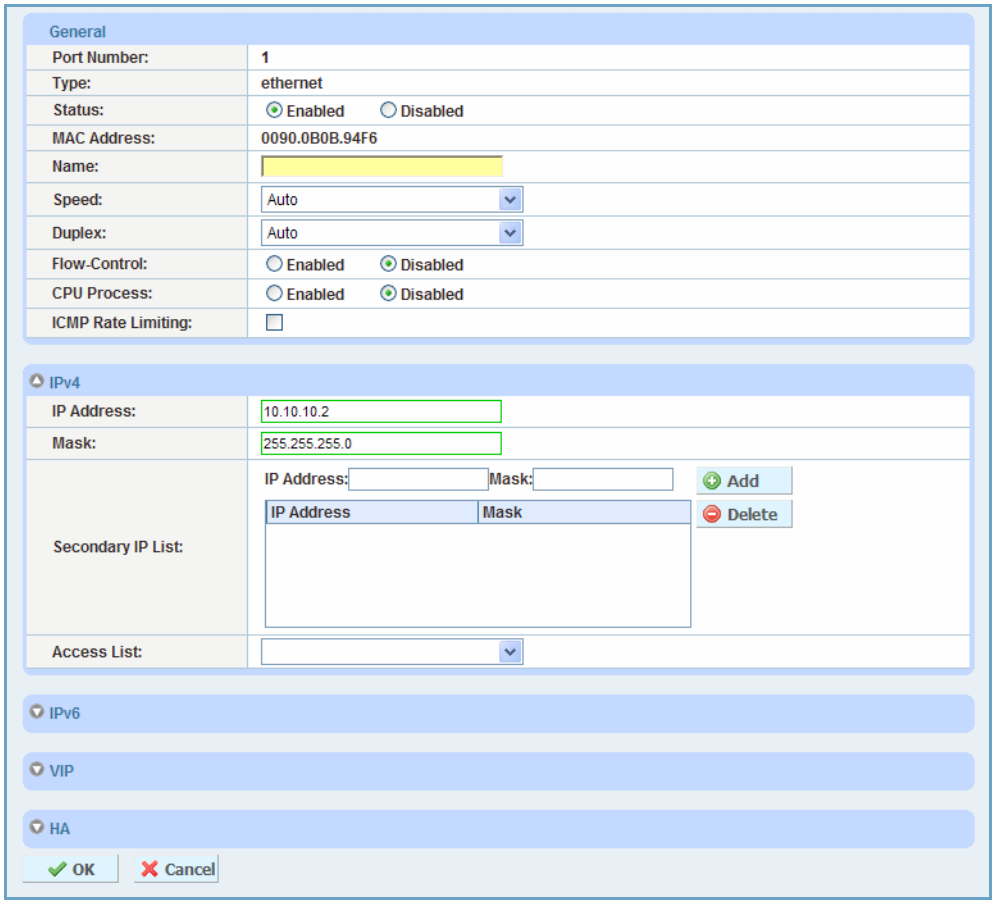

General

Port Number:	1
Type:	ethernet
Status:	● Enabled ○ Disabled
MAC Address:	0090.0B0B.94F6
Name:	
Speed:	Auto
Duplex:	Auto
Flow-Control:	○ Enabled ● Disabled
CPU Process:	○ Enabled ● Disabled
ICMP Rate Limiting:	□

IPv4

IP Address:	10.10.10.2
Mask:	255.255.255.0
Secondary IP List:	IP Address: Mask: Add IP Address / Mask Delete
Access List:	

FIGURE 15 Config > Network > Route > IPv4 Static

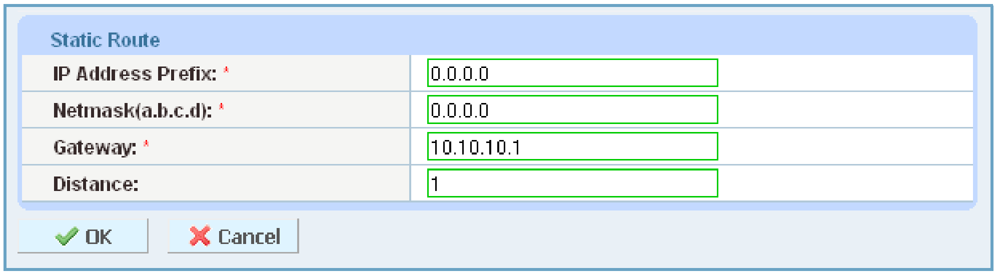

Static Route

IP Address Prefix: *	0.0.0.0
Netmask(a.b.c.d): *	0.0.0.0
Gateway: *	10.10.10.1
Distance:	1

USING THE CLI

The following commands enable the Ethernet interfaces used in the example and configure IP addresses on them:

```
AX(config)#interface ethernet 1
AX(config-if:ethernet1)#enable
AX(config-if:ethernet1)#ip address 10.10.10.2 /24
AX(config-if:ethernet1)#interface ethernet 2
AX(config-if:ethernet2)#enable
AX(config-if:ethernet2)#ip address 192.168.3.100 /24
AX(config-if:ethernet2)#interface ethernet 3
AX(config-if:ethernet3)#enable
AX(config-if:ethernet3)#ip address 192.168.1.111 /24
AX(config-if:ethernet3)#exit
AX(config-if:ethernet3)#interface ethernet 4
AX(config-if:ethernet4)#enable
AX(config-if:ethernet4)#ip address 192.168.2.100 /24
AX(config-if:ethernet4)#exit
```

The following command configures the default route through 10.10.10.1:

```
AX(config)#ip route 0.0.0.0 /0 10.10.10.1
```

Jumbo Frames

The current release adds support for jumbo frames. In this release, a jumbo frame is an Ethernet frame that is more than 1522 bytes long.

By default, the maximum transmission unit (MTU) on all AX physical Ethernet interfaces is 1500 bytes. The default Ethernet frame size is 1522 bytes, which includes 1500 bytes for the payload, 14 bytes for the Ethernet header, 4 bytes for the CRC, and 4 bytes for a VLAN tag. Jumbo support is disabled by default.

You can enable jumbo support on a global basis. In this case, the MTU is not automatically changed on any interfaces, but you can increase the MTU on individual interfaces.

- On FPGA models, you can increase the MTU on individual Ethernet interfaces, Virtual Ethernet (VE) interfaces, and trunks up to 12000 bytes. This applies to the following AX models: AX 5630, AX 5200-11, AX 3400 and AX 3200-12.
- On non-FPGA models, you can increase the MTU on individual Ethernet interfaces, VE interfaces, and trunks up to 9216 bytes. This applies to the following AX models: AX 3530, AX 3030, AX 3000-11-GCF, and AX 3000.

Notes:

- Enabling jumbo support does not automatically change the MTU on any interfaces. You must explicitly increase the MTU on those interfaces you plan to use for jumbo packets.
- Jumbo support is not recommended on 10/100 Mbps ports.
- On FPGA AX models, for any incoming jumbo frame, if the outgoing MTU is less than the incoming frame size, the AX device fragments the frame into equal-size fragments that are less than 1500 bytes. For example, if the MTU on the AX device is 3000 bytes and the frame size of the incoming package is 4000 bytes, the AX device divides the frame into almost three equal fragments of 1330 bytes, 1330 bytes, and the remaining 1340 bytes.

  On non-FPGA models,. if the outgoing MTU is less than the incoming frame size, the AX device fragments the frame into equal-size fragments that are greater than 1500 bytes. For example, if the MTU on the AX device is 3000 bytes and the frame size of the incoming package is 4000 bytes, the AX device divides the frame into two 2000-byte fragments.

- Setting the MTU on an interface indirectly sets the frame size of incoming packets to the same value. (This is called the Maximum Receive Unit [MRU]).
- The default MTU is 1500 and can not be set to a higher value.

Caution: **On non-FPGA models, after you enable (or disable) jumbo frame support, you must save the configuration and reboot to place the change into effect.**

If jumbo support is enabled on a non-FPGA model and you erase the startup-config, the device is rebooted automatically after the configuration is erased.

Using the GUI

Enabling Jumbo Support (FPGA models)

1. Select Config Mode > Network > Interface > Global.
2. Select Enabled next to Jumbo Frame.
3. Click OK.

Enabling Jumbo Support (non-FPGA models)

If you are using a non-FPGA model):

1. Select Config Mode > Network > Interface > Global.
2. Select Enabled next to Jumbo Frame.
3. Click OK.
4. Click Save at the top of the GUI window to save the configuration change.
5. Select Config Mode > System > Settings > Action > Reboot.

Changing the MTU on an Interface

1. Select Config Mode > Network > Interface > LAN.
2. Click on the interface number, in the Interface column. The configuration page for the interface appears.
3. Edit the value in the MTU field.
4. Click OK.

Disabling Jumbo Support (FPGA models)

On each interface with an MTU higher than 1500 bytes:

1. Select Config Mode > Network > Interface > LAN.
2. Click on the interface number, in the Interface column. The configuration page for the interface appears.
3. Edit the value in the MTU field to be 1500 (or less).
4. Click OK.
5. Repeat for each interface on which the MTU is greater than 1500 bytes.
6. Select Config Mode > Network > Interface > Global.
7. Select Disabled next to Jumbo Frame.
8. Click OK.

Disabling Jumbo Support (non-FPGA models)

If you are using a non-FPGA model:

1. Select Config Mode > Network > Interface > LAN.
2. Click on an interface number, in the Interface column. The configuration page for the interface appears.
3. Edit the value in the MTU field to be 1500 (or less).
4. Click OK.
5. Repeat for each interface on which the MTU is greater than 1500 bytes.
6. Select Config Mode > Network > Interface > Global.
7. Select Disabled next to Jumbo Frame.
8. Click OK.
9. Click Save at the top of the GUI window to save the configuration change.
10. Select Config Mode > System > Settings > Action > Reboot.

Caution: **On non-FPGA models, you must save the configuration and reboot after entering the “no enable-jumbo” command. If you reload or reboot without first saving the configuration, the feature can not be**

re-enabled until you first repeat the procedure above to disable it. Then, you can re-enable the feature.

USING THE CLI

Enabling Jumbo Support (FPGA models)

To globally enable jumbo support, use the following command at the global configuration level of the CLI:

```
enable-jumbo
```

Enabling Jumbo Support (non-FPGA models)

If you are using a non-FPGA model (see above):

1. Use the following command at the global configuration level of the CLI: **enable-jumbo**

2. Enter the following command to save the configuration: **write memory**

3. Use the following command at the Privileged EXEC level to reboot: **reboot**

Changing the MTU on an Interface

To change the MTU on an interface, use the following command at the configuration level for the interface:

```
mtu bytes
```

Disabling Jumbo Support (FPGA models)

1. On each interface with an MTU higher than 1500 bytes, enter the following command: **no mtu**

2. Use the following command at the global configuration level of the CLI: **no enable-jumbo**

Disabling Jumbo Support (non-FPGA models)

If you are using a non-FPGA model:

1. On each interface with an MTU higher than 1500 bytes, enter the following command: **no mtu**

2. Use the following command at the global configuration level of the CLI: **no enable-jumbo**

3. Enter the following command to save the configuration: **write memory**

4. Use the following command at the Privileged EXEC level to reboot:
   reboot

Caution: **On non-FPGA models, you must save the configuration and reboot after entering the “no enable-jumbo” command. If you reload or reboot without first saving the configuration, the feature can not be re-enabled until you first repeat the procedure above to disable it. Then, you can re-enable the feature.**

Link Trunking

This chapter describes how to configure trunk links on the AX device.

Overview

The AX device supports aggregation of multiple Ethernet data ports into logical links, called “trunks”. Trunks can enhance performance by providing higher throughput and greater link reliability.

Higher throughput is provided by the aggregate throughput of the individual links in the trunk. Greater link reliability is provided by the multiple links in the trunk. If an individual port in the trunk goes down, the trunk link continues to operate using the remaining up ports in the trunk.

You can configure the following types of trunks:

- Static trunks – You can configure up to 16 static trunks. Each trunk can contain 2-8 Ethernet data ports.
- Dynamic trunks – You can enable Link Aggregation Control Protocol (LACP) on Ethernet data interfaces, to make those interfaces candidate members of dynamically configured trunks. You can configure up to 16 dynamic trunks with a maximum of 8 Ethernet data member ports per trunk.

Interface parameters for a trunk apply collectively to the entire trunk, as a single interface. For example, IP addresses and other IP parameters apply to the entire trunk as a single interface.

Static Trunk Configuration

You can configure up to 16 static trunks. Each trunk can contain 2-8 Ethernet data ports.

USING THE GUI

To configure a static trunk:

1. Go to Config Mode > Network > Trunk.
   The following window appears:

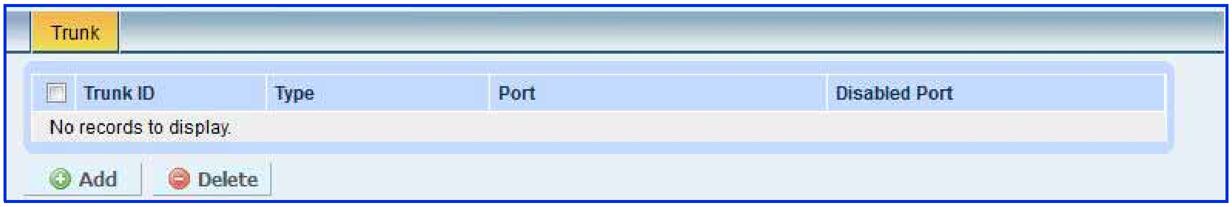

☐	Trunk ID	Type	Port	Disabled Port
No records to display.				

2. Click on Add to be able to create a static trunk.
   The following window appears:

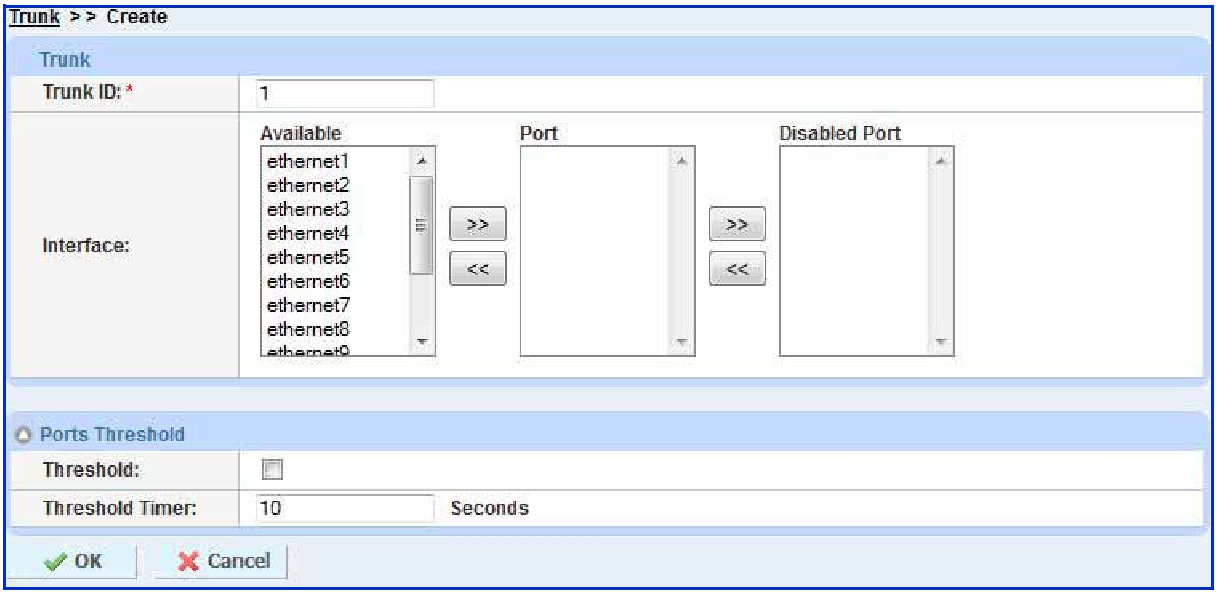

Trunk ID: * 1

Interface:
Available: ethernet1, ethernet2, ethernet3, ethernet4, ethernet5, ethernet6, ethernet7, ethernet8, ethernet9

Threshold: ☐

Threshold Timer: 10 Seconds

3. Indicate a Trunk ID from 1-16.

4. In the Interface section, click on the interface you want to associate with the trunk and move the interface from the Available column using the greater than redirect arrows (>>). If you want to move them back to the Available column, select the interface, and move it back using the less than redirect arrows. Similarly, you can move the interface from the Port to the Disabled Port column.

5. In the Ports Threshold section, do the following:

   a. Click the checkbox in the Threshold field and from the drop down menu, choose a value of 2-8. This field specifies the minimum number of ports that must be up in order for the trunk to remain up

   b. In the Port Threshold Timer field, indicate a timer value from 1-300 seconds. This threshold timer specifies how many seconds to wait after a port goes down before marking the trunk down, if the threshold is not met. You can set the ports-threshold timer to 1-300 seconds. The default is 10 seconds.

6. Click on OK.
   Your static trunk will be added:

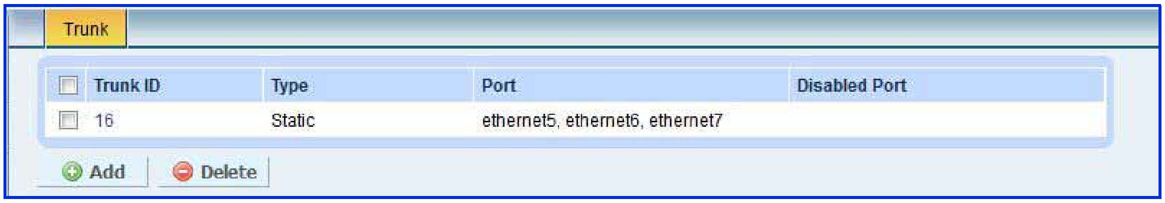

□	Trunk ID	Type	Port	Disabled Port
□	16	Static	ethernet5, ethernet6, ethernet7	

Using the CLI

To configure a static trunk:

1. Add the trunk and configure trunk parameters:
   - Port membership. You can add 2-8 Ethernet data ports to the trunk.
   - Optionally, configure port-threshold parameters. The port threshold specifies the minimum number of member ports in the trunk that must be up, for the trunk itself to remain up.

2. Configure interface-level parameters on the trunk, if applicable:
   - Name – You can assign a name to the trunk, in addition to the numeric ID you specify when you create the trunk.
   - IPv4 and IPv6 parameters – You can assign one or more IPv4 and IPv6 addresses, and configure other IP-related parameters such as IP helper or IPv6 neighbor discovery.
   - Dynamic routing – You can configure interface-level OSPF and IS-IS parameters.
   - Access list (ACL) – You can filter incoming traffic based on source and destination IPv4 or IPv6 address and protocol port, as well as additional parameters such as ICMP type and code or VLAN ID.
   - ICMP rate limiting – You can enable protection against denial-of-service (DoS) attacks such as Smurf attacks, which consist of floods of spoofed broadcast ping messages.
   - Layer 3 forwarding – Layer 3 forwarding is enabled by default. You can disable it.

     If you want to allow Layer 3 forwarding *except* between VLANs, a separate option allows you to disable Layer 3 forwarding between VLANs.

Trunk Parameters and Trunk Interface Parameters

You can configure trunk-related parameters at the following configuration levels:

- Trunk
- Interface

At the trunk configuration level, you can enable or disable the trunk, add or remove trunk members, and set port-threshold parameters. Using the **disable** command at the trunk configuration level completely disables all ports in the trunk.

At the interface configuration level for a trunk, you can configure interface-related parameters such as IP, IPv6, and routing settings. You also can disable or re-enable Layer 3 forwarding on the trunk.

Using the **disable** command at the interface configuration level for a trunk disables Layer 3 forwarding on the trunk but does not completely disable the trunk.

Configuring Trunk Parameters

To configure trunk parameters, use the following commands.

```
[no] trunk trunk-id-num
```

Enter this command at the global configuration level of the CLI. This command changes the CLI to the configuration level for the specified trunk. The *trunk-id-num* can range from 1-16.

Adding Ports to a Trunk

To add Ethernet data ports to the trunk, use the following command:

```
[no] ethernet portnum
[to portnum] [ethernet portnum] ...
```

Configuring Port-threshold Parameters

By default, a trunk's status remains UP so long as at least one of its member ports is up. You can change the ports threshold of a trunk to 2-8 ports.

If the number of up ports falls below the configured threshold, the AX automatically disables the trunk's member ports. The ports are disabled in the running-config. The AX device also generates a log message and an SNMP trap, if these services are enabled.

Note: After the feature has disabled the members of the trunk group, the ports are not automatically re-enabled. The ports must be re-enabled manually after the issue that caused the ports to go down has been resolved.

In some situations, a timer is used to delay the ports-threshold action. The configured port threshold is not enforced until the timer expires. The ports-threshold timer for a trunk is used in the following situations:

- When a member of the trunk links up.
- A port is added to or removed from the trunk.
- The port threshold for the trunk is configured during runtime. (If the threshold is set in the startup-config, the timer is not used.)

To configure port-threshold parameters, use the following commands at the configuration level for the trunk:

```
[no] ports-threshold num
```

This command specifies the minimum number of ports that must be up in order for the trunk to remain up. You can specify 2-8.

If the number of up ports falls below the configured threshold, the AX automatically disables the trunk's member ports. The ports are disabled in the running-config. The AX device also generates a log message and an SNMP trap, if these services are enabled.

```
[no] ports-threshold-timer seconds
```

This command specifies how many seconds to wait after a port goes down before marking the trunk down, if the threshold is not met. You can set the ports-threshold timer to 1-300 seconds. The default is 10 seconds.

Configuring Interface-level Parameters on a Static Trunk

To configure interface-level parameters for the trunk, use the following command to access the *interface* configuration level for the trunk:

```
interface trunk num
```

Enter this command at the global configuration level of the CLI. This command changes the CLI to the interface configuration level for the specified trunk, where the following commands are available:

```
[no] access-list acl-num in
[no] bfd options
clear
disable
do
enable
end
exit
[no] icmp-rate-limit options
[no] interface options
[no] ip options
[no] ipv6 options
[no] isis options
[no] l3-vlan-fwd-disable
[no] mtu options
[no] name string
[no] ospf options
show
write
```

Note: The **disable** and **enable** commands at the interface configuration level for the trunk control Layer 3 forwarding on the trunk but do not completely disable the trunk. To control all forwarding on the trunk, use the **disable** or **enable** command at the trunk configuration level instead.

For more information about these commands, see the “Config Commands: Interface” chapter of the *AX Series CLI Reference*.

Static Trunk Configuration Example

The following commands configure trunk 7 with ports 5-8, and verify the configuration:

```
AX(config)#trunk 7
AX(config-trunk:7)#ethernet 5 to 8
AX(config-trunk:7)#show trunk
Trunk ID         : 7       Member Count: 4
Trunk Status      : Up
Members           : 5   6   7   8
Cfg Status        : Enb Enb Enb Enb
Oper Status       : Up  Up  Up  Up
Ports-Threshold   : None
Working Lead      : 5
AX(config-trunk:7)#exit
```

The following commands access the interface configuration level for the trunk and assign an IPv6 address to the trunk interface:

```
AX(config)#interface trunk 7
AX(config-if:trunk7)#ipv6 address 2001:db8::7/32
```

Dynamic Trunk Configuration

Link Aggregation Control Protocol (LACP) dynamically creates trunk links.

The AX implementation of LACP is based on the 802.3ad IEEE specification.

You can configure a maximum of 16 LACP trunks on an AX device. An interface can belong to a single LACP trunk.

LACP Parameters

The following LACP parameters are configurable.

Global LACP Parameter

- LACP system priority – Specifies the LACP priority of the AX device. In cases where LACP settings on the local device (the AX device) and the remote device at the other end of the link differ, the settings on the device with the higher priority are used.

  You can specify 1-65535. A low priority number indicates a high priority value. The highest priority is 1 and the lowest priority is 65535. The default is 32768.

Global Parameters for Individual LACP Trunks

- Trunk state – Enabled or disabled. LACP trunks are enabled by default.
- Ports threshold – Minimum number of ports that must be up in order for the trunk to remain up. If the number of up ports falls below the configured threshold, the AX automatically disables the trunk’s member ports. The ports are disabled in the running-config. You can specify 2-8. By default, no port threshold is set. A trunk remains up if even 1 of its ports is up.
- Ports threshold timer – Number of seconds to wait after a port goes down before marking the trunk down, if the configured threshold is not met. You can set the ports-threshold timer to 1-300 seconds. The default is 10 seconds.

Interface-Level LACP Parameters

- LACP trunk ID – ID of a dynamic trunk. Adding an interface to an LACP trunk makes that interface a candidate for membership in the trunk. During negotiation with the other side of the link, LACP selects the interfaces to actively participate in the link.

  You add up to 8 interfaces to an LACP trunk.

  When you add an interface, you must specify whether LACP will run in active or passive mode on the interface. Active mode initiates link formation with the other end of the link. Passive mode waits for the other end of the link to initiate link formation.

  The admin key must match on all interfaces in the trunk. The value can be 10000-65535.
- LACP port priority – Priority of the interface for selection as an active member of a link. If the LACP trunk has more candidate members than are allowed by the device at the other end of the link, LACP selects the

interfaces with the highest port priority values as the active interfaces. The other interfaces are standbys, and are used only if an active interface goes down.

You can specify 1-65535. A low priority number indicates a high priority value. The highest priority is 1 and the lowest priority is 65535. The default is 32768.

- LACP timeout – Aging timeout for LACP data units from the other end of the LACP link.

  You can specify short (3 seconds) or long (90 seconds). The default is long.

- Mode – Indicate whether you want LACP to operate in Active or Passive Mode. The Active mode initiates link formation with the other end of the link. In this case, the AX device will send the LACP frame to its link partner. Passive mode waits for the other end of the link to initiate link formation. In this case, the AX device will only send an LACP frame if it receives an LACP frame from the link partner.
- Admin Key – The admin key must match on all interfaces in the trunk. The value can be 10000-65535.
- Unidirectional Link Detection (UDLD) – UDLD checks the links in LACP trunks to ensure that both the send and receive sides of each link are operational. UDLD is disabled by default on AX LACP trunk links. You can enable UDLD on individual LACP trunk interfaces. (For more information, see “UDLD” on page 87.)

UDLD

When UDLD is enabled, the AX Series UDLD uses LACP protocol packets as heartbeat messages. If an LACP link on the AX device does not receive an LACP protocol packet within a specified timeout, LACP blocks traffic on the port. This corrects the problem by forcing the devices connected by the non-operational link to use other, fully operational links.

A link that is blocked by LACP can still receive LACP protocol packets but blocks all other traffic.

UDLD is disabled by default on AX LACP trunk links. You can enable UDLD on individual LACP trunk interfaces.

Heartbeat Timeout

The local port waits for a configurable timeout to receive an LACP protocol packet from the remote port. If an LACP protocol packet does not arrive before the timeout expires, LACP disables the local port. You can set the timeout to 1-60 seconds (slow timeout) or 100-1000 milliseconds (fast timeout). The default is 1 second.

If the remote port begins sending LACP protocol packets again, LACP on the local port re-enables the port.

Requirements

To operate properly, UDLD must be supported and enabled on both devices that are using LACP trunk links.

It is recommended to use auto-negotiation on each end of the link to establish the mode (half duplex or full duplex). Auto-negotiation helps ensure link bidirectionality at Layer 1, while UDLD helps at Layer 2.

Configuration

Using the GUI

Configuring LACP via the GUI

Configure the global LACP parameter:

1. Go to Config Mode > Network > Interface > LAN.

2. Choose an Ethernet number.
   The following window appears:

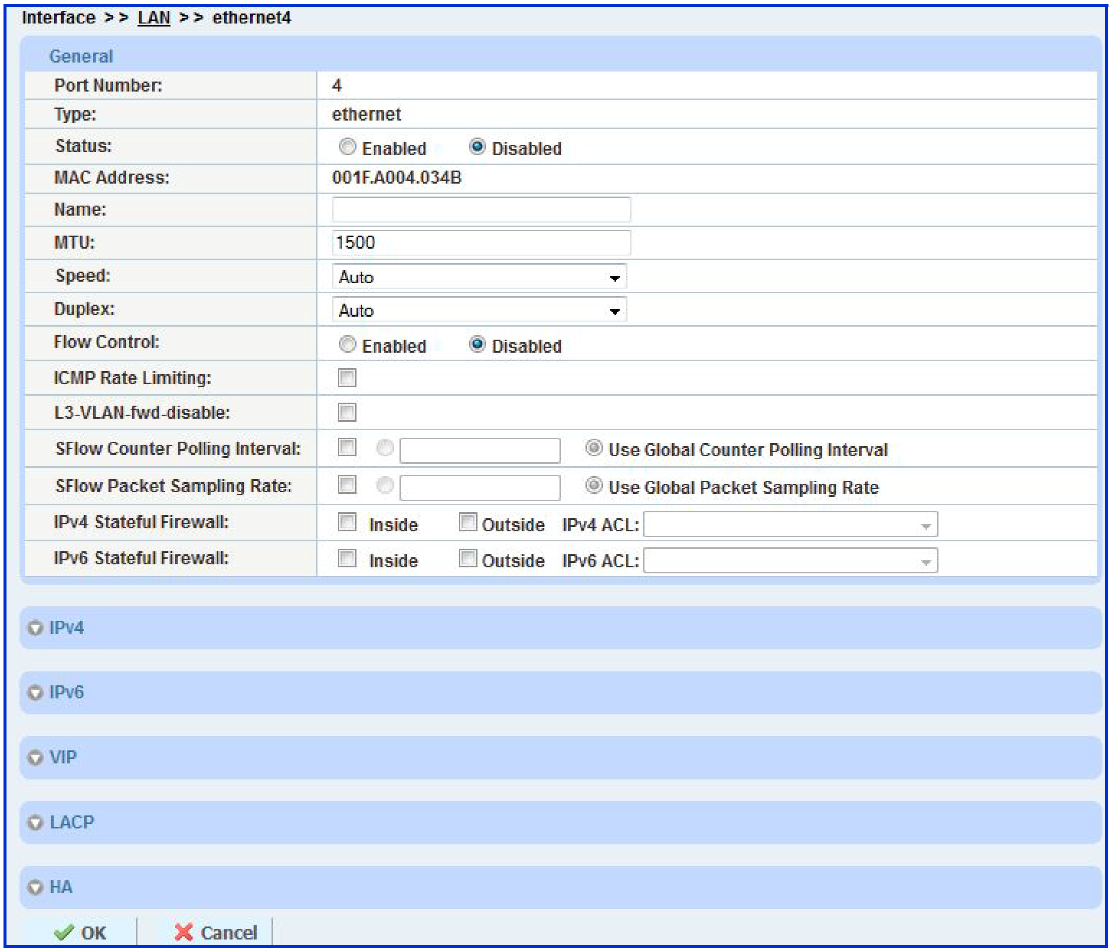

Interface > > LAN > > ethernet4

General

Port Number:	4
Type:	ethernet
Status:	○ Enabled ● Disabled
MAC Address:	001F.A004.034B
Name:	
MTU:	1500
Speed:	Auto
Duplex:	Auto
Flow Control:	○ Enabled ● Disabled
ICMP Rate Limiting:	□
L3-VLAN-fwd-disable:	□
SFlow Counter Polling Interval:	□ ○ ● Use Global Counter Polling Interval
SFlow Packet Sampling Rate:	□ ○ ● Use Global Packet Sampling Rate
IPv4 Stateful Firewall:	□ Inside □ Outside IPv4 ACL:
IPv6 Stateful Firewall:	□ Inside □ Outside IPv6 ACL:

3. Scroll down and click on the LACP tab.
   The following window appears:

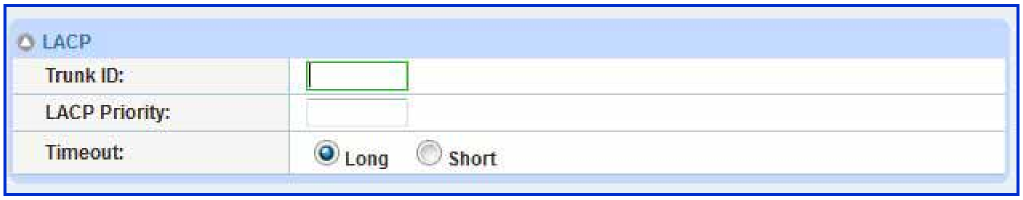

LACP

Trunk ID:	
LACP Priority:	
Timeout:	● Long ○ Short

4. Enter the Trunk ID.
   The LACP window refreshes and displays additional configuration options:

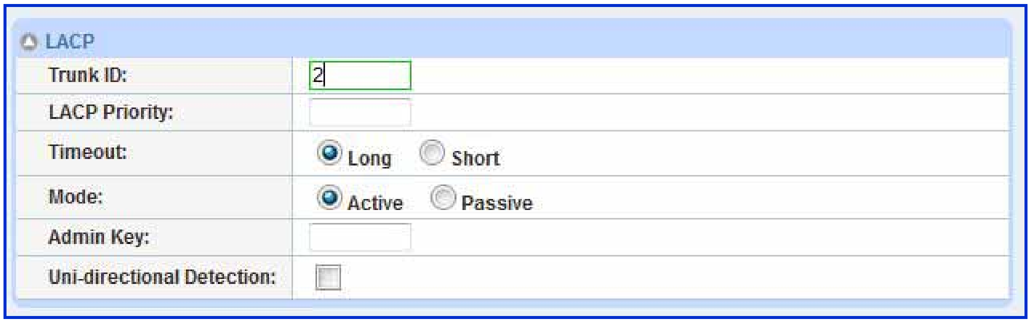

LACP

Trunk ID:	2
LACP Priority:	
Timeout:	● Long ○ Short
Mode:	● Active ○ Passive
Admin Key:	
Uni-directional Detection:	□

5. Enter the LACP priority.

6. Choose a Timeout value of Long or Short.

7. Click the radio button to indicate if LACP should operate in Active or in Passive mode.

8. Specify an Admin Key.

9. Click the checkbox for Uni-directional Detection.

10. Click OK.

To configure the LACP system priority, follow these steps:

1. Go to Config Mode > Network > LACP.
   The following window will appear:

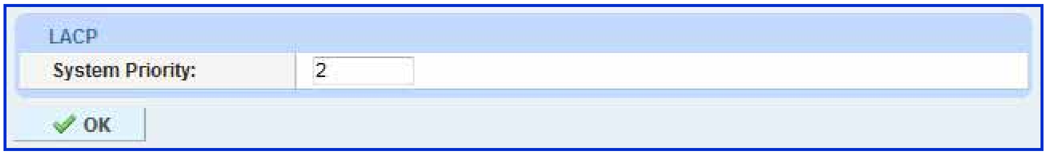

LACP

System Priority:	2

2. Accept the default LACP system priority or change it to the desired value.

3. Click OK.

Minimum Port Threshold for LACP

This release enhances LACP with support for minimum port thresholds. Previous releases support minimum port thresholds for static trunks but not for dynamic (LACP) trunks.

By default, a trunk's status remains Up so long as at least one of its member ports is up. You can change the ports threshold of a trunk to 2-8 ports.

Since a trunk comprises of several member links, if the number of operational members of a trunk goes below the configured threshold value, the remaining member links are automatically marked as "blocked" and the trunk is considered non--operational. When the down link is functional again, the remaining links that were marked blocked are also operational again, making the trunk available for use.

Note: If you administratively disable the LACP feature from members of the trunk group, the links are not automatically re-enabled. The links must be re-enabled manually after the issue that caused the links to go down has been resolved.

The LACP feature can also be administratively disabled on links in a trunk.

In some situations, a timer is used to delay the ports-threshold action. The configured port threshold is not enforced until the timer expires. The ports-threshold timer for a trunk is used in the following situations:

- When a member of the trunk link is up.
- A port is added to or removed from the trunk.
- The port threshold for the trunk is configured during runtime. (If the threshold is set in the "startup-config" file, the timer is not used.)

Using the GUI

To configure the port threshold parameters for LACP trunks, do the following.

Note: These steps assume that you have already created an LACP dynamic trunk using Config Mode > Interface > LAN (from the LACP tab). For details, refer to the *System and Administration Guide*.

1. Go to Config Mode > Network > Trunk.
   The following window will appear:

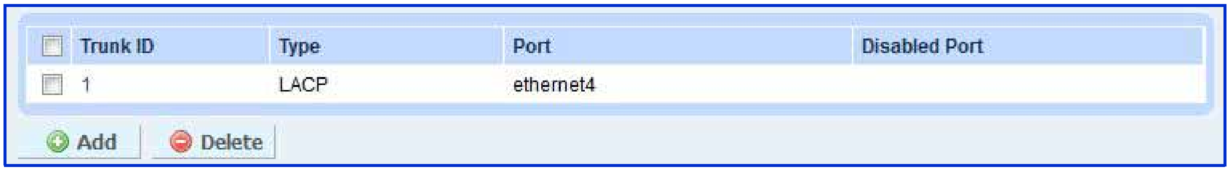

□ Trunk ID	Type	Port	Disabled Port
□ 1	LACP	ethernet4	

2. Click on the existing LACP Trunk 1.
   The Trunk >>Create window will appear. You will not be able to edit any interface information in this window. It is displayed for informa-

tional purposes only. If you want to add any interfaces to this LACP trunk, you must add these interfaces from the System > LAN > Interfaces menu where you originally created the LACP trunk:

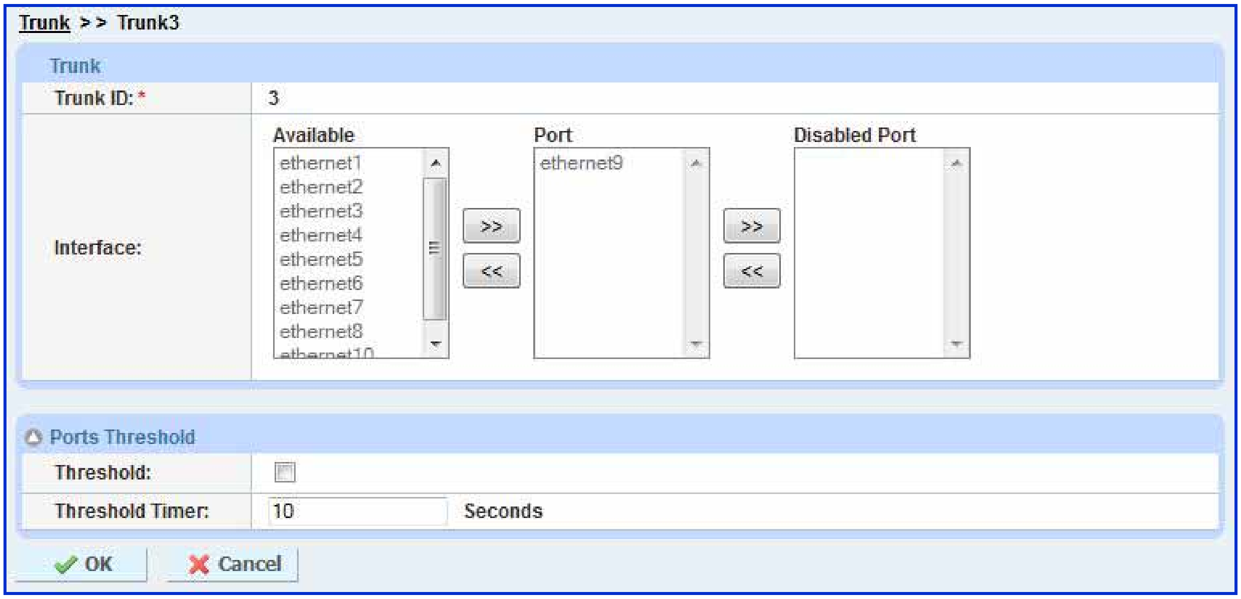

Trunk >> Trunk3

Trunk

Trunk ID: * 3

Interface:

Available: ethernet1, ethernet2, ethernet3, ethernet4, ethernet5, ethernet6, ethernet7, ethernet8, ethernet10

Port: ethernet9

Disabled Port:

Ports Threshold

Threshold: □

Threshold Timer: 10 Seconds

3. In the Ports Threshold section, do the following:

   a. Click the checkbox in the Threshold field and from the drop down menu, choose a value of 2-8. This field specifies the minimum number of ports that must be up in order for the trunk to remain up

   b. In the Port Threshold Timer field, indicate a timer value from 1-300 seconds. This threshold timer specifies how many seconds to wait after a port goes down before marking the trunk down, if the threshold is not met. You can set the ports-threshold timer to 1-300 seconds. The default is 10 seconds.

4. Click on OK.

Verifying Port Threshold Configuration

To verify your LACP configuration of the Port Threshold and the Port Threshold Timer, do the following:

Go to Monitor Mode > Network > Trunk.
The following window will display the configured Port Threshold Timer of 10 for ethernet3 that belongs to Trunk ID 1:

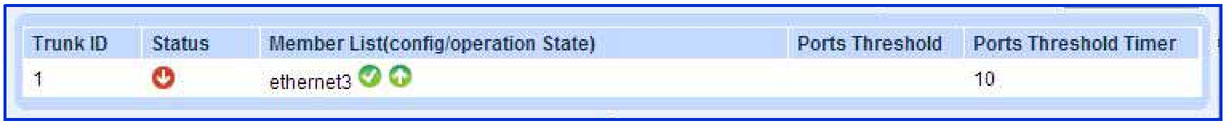

Trunk ID	Status	Member List(config/operation State)	Ports Threshold	Ports Threshold Timer
1	↓	ethernet3 ✓ ↑		10

Monitoring LACP via the GUI

You can monitor LACP from Monitoring Mode. View the following LACP information:

- Display LACP system IDs,
- Display LACP counters.
- Display LACP trunk administration key list details.
- Display summarized and detailed information on LACP ports and 1-16 trunk lists.

To monitor LACP, do the following:

1. Go to Monitor Mode > Network > LACP > System ID.
   The following window will appear displaying information on the system ID:

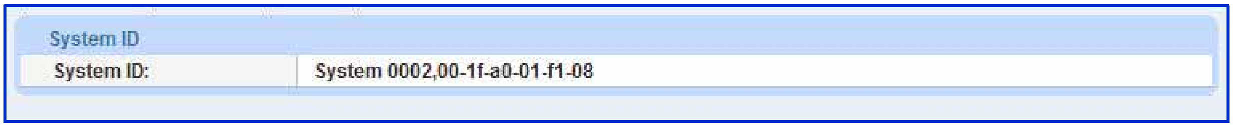

System ID

System ID:	System 0002,00-1f-a0-01-f1-08

2. Go to Monitor Mode > Network > LACP > Trunk.
   This menu will display information on Admin Key List, Summary, and Detail:

   a. View information in the trunk Admin Key List tab:

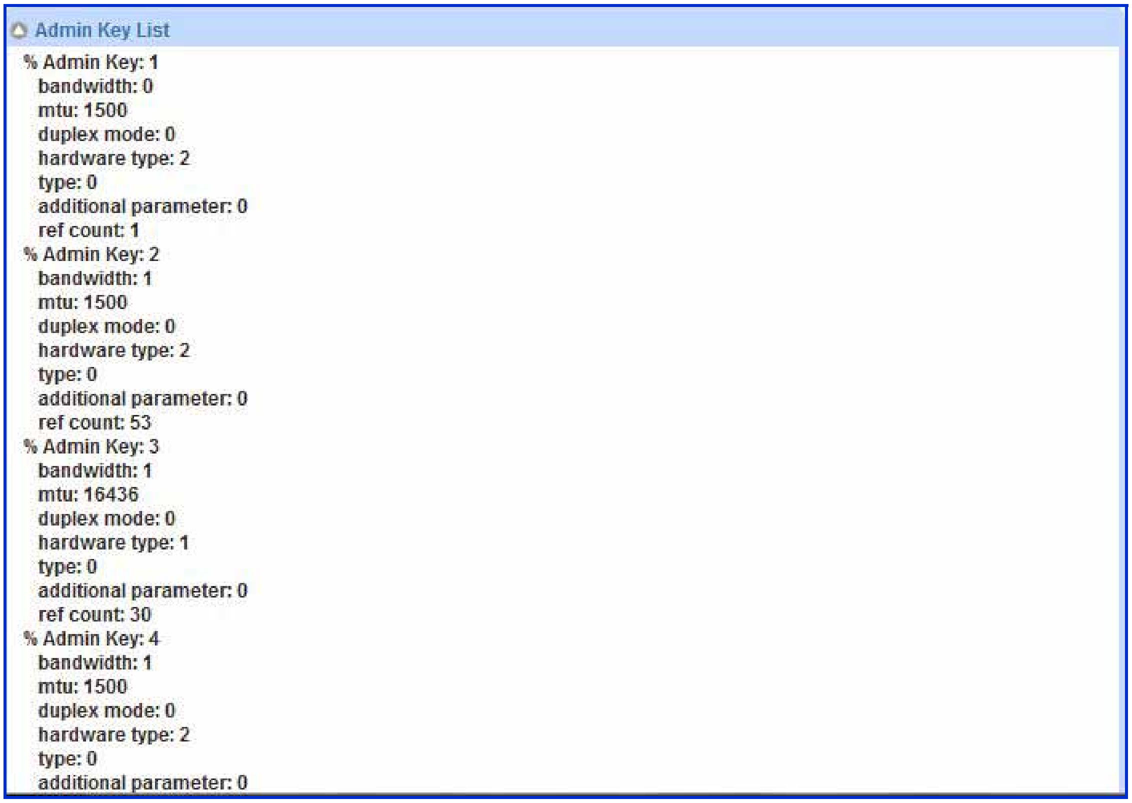

Admin Key List

```
% Admin Key: 1
   bandwidth: 0
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 0
   ref count: 1
% Admin Key: 2
   bandwidth: 1
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 0
   ref count: 53
% Admin Key: 3
   bandwidth: 1
   mtu: 16436
   duplex mode: 0
   hardware type: 1
   type: 0
   additional parameter: 0
   ref count: 30
% Admin Key: 4
   bandwidth: 1
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 0
```

   b. View information in the trunk Summary tab:

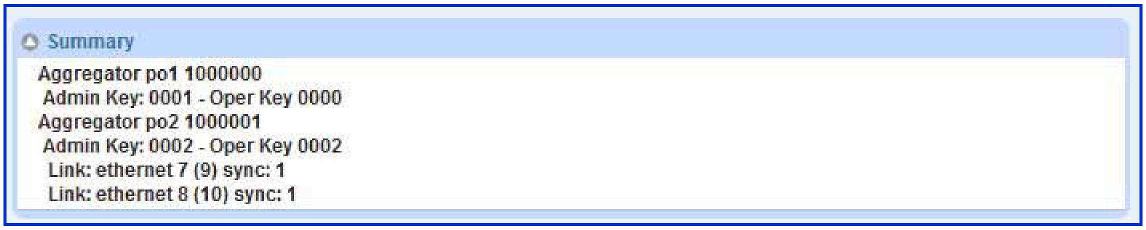

Summary

```
Aggregator po1 1000000
 Admin Key: 0001 - Oper Key 0000
Aggregator po2 1000001
 Admin Key: 0002 - Oper Key 0002
  Link: ethernet 7 (9) sync: 1
  Link: ethernet 8 (10) sync: 1
```

   c. View information in the trunk Detail tab:

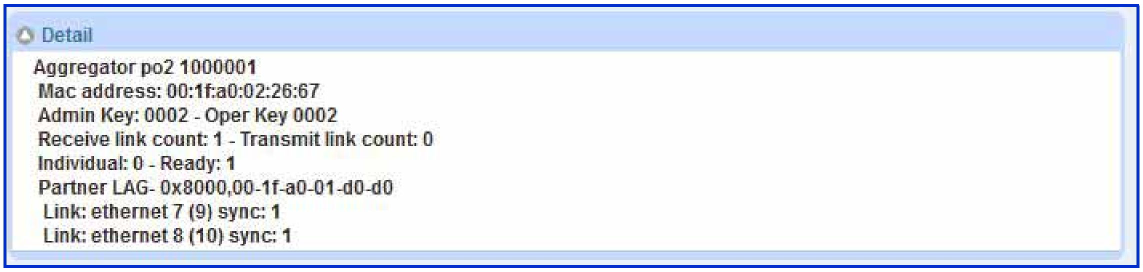

Detail

```
Aggregator po2 1000001
 Mac address: 00:1f:a0:02:26:67
 Admin Key: 0002 - Oper Key 0002
 Receive link count: 1 - Transmit link count: 0
 Individual: 0 - Ready: 1
 Partner LAG- 0x8000,00-1f-a0-01-d0-d0
  Link: ethernet 7 (9) sync: 1
  Link: ethernet 8 (10) sync: 1
```

USING THE CLI

Configuring each Interface

1. Change the CLI to the configuration level for the interface.

2. Assign the interface to the LACP trunk, using the following command:

   ```
   [no] lacp trunk lacp-trunk-id [admin-key num]
   mode {active | passive}
   [unidirectional-detection]
   ```

   The interface becomes a candidate member of the trunk.

3. (Optional) Specify the LACP priority of the interface, using the following command:

   ```
   [no] lacp port-priority num
   ```

   You can specify 1-65535. The default is 32768.

4. (Optional) Specify the aging timeout for LACP data units from the other end of the LACP link, using the following command:

   ```
   [no] lacp timeout {short | long}
   ```

   You can specify **short** (3 seconds) or **long** (90 seconds). The default is **long**.

5. (Optional) Specify the UDLD aging timeout, using the following command:

   ```
   [no] lacp udld-timeout {fast | slow} num
   ```

   You can specify **fast** (100-1000 milliseconds) or **slow** (1-60 seconds). The default is **slow 1**.

Configuring LACP

1. (Optional) Set the LACP system priority, using the following command at the global configuration level of the CLI:

   ```
   [no] lacp system-priority num
   ```

   You can specify 1-65535. The default is 32768.

2. (Optional) Configure ports-threshold settings:

   a. Specify the minimum number of ports that must remain up, using the following command at the LACP trunk configuration level of the CLI:

      ```
      [no] ports-threshold num [do-manual-recovery]
      ```

      You can specify 2-8 ports.

The **do-manual-recovery** option disables automatic recovery of the trunk when the required number of ports come back up. If you use this option, the trunk remains disabled until you re-enable it.

b. Specify how many seconds to wait after a port goes down before marking the trunk down, if the configured threshold is not met.

```
[no] ports-threshold-timer seconds
```

You can specify 2-8 ports. You can set the ports-threshold timer to 1-300 seconds. The default is 10 seconds.

Configuring Interface-level Parameters on an LACP Trunk

To configure interface-level parameters for the trunk, use the following command to access the *interface* configuration level for the trunk:

```
interface trunk num
```

Enter this command at the global configuration level of the CLI. This command changes the CLI to the interface configuration level for the specified trunk, where the following commands are available:

```
[no] access-list acl-num in
[no] bfd options
clear
disable
do
enable
end
exit
[no] icmp-rate-limit options
[no] interface options
[no] ip options
[no] ipv6 options
[no] isis options
[no] l3-vlan-fwd-disable
[no] mtu options
[no] name string
```

```
[no] ospf options
show
write
```

Note: The **disable** and **enable** commands at the interface configuration level for the trunk control Layer 3 forwarding on the trunk but do not completely disable the trunk. To control all forwarding on the trunk, use the **disable** or **enable** command at the trunk configuration level instead.

For more information about these commands, see the “Config Commands: Interface” chapter of the *AX Series CLI Reference*.

Displaying LACP Information

To display LACP information, use the following commands:

```
show lacp sys-id
```

This command displays the AX device’s LACP system ID. The LACP system ID consists of the LACP system priority and the system MAC address.

On AX devices, the system MAC address is the lowest numbered MAC address on the device.

```
show lacp counter [lacp-trunk-id]
```

This command displays statistics.

```
show lacp trunk
[
admin-key-list-details |
detail |
summary |
lacp-trunk-id
]
```

This command shows configuration and status information for LACP.

Clearing LACP Statistics

To clear LACP statistics counters, use the following command at the Privileged EXEC level of the CLI:

```
clear lacp counters [lacp-trunk-id]
```

Configuration Example

This example enables LACP on physical Ethernet data ports 1-3 and 6.

The following commands configure LACP parameters on the ports:

```
AX-1(config)#interface ethernet 1
AX-1(config-if:ethernet1)#lacp trunk 5 admin-key 10001 mode active
AX-1(config-if:ethernet1)#lacp timeout long
AX-1(config-if:ethernet1)#interface ethernet 2
AX-1(config-if:ethernet2)#lacp trunk 5 admin-key 10001 mode active
AX-1(config-if:ethernet2)#lacp timeout long
AX-1(config-if:ethernet2)#interface ethernet 3
AX-1(config-if:ethernet3)#lacp trunk 10 mode active
AX-1(config-if:ethernet3)#lacp timeout short
AX-1(config-if:ethernet3)#interface ethernet 4
AX-1(config-if:ethernet4)#lacp timeout long
AX-1(config-if:ethernet4)#interface ethernet 6
AX-1(config-if:ethernet6)#lacp trunk 10 mode active
AX-1(config-if:ethernet6)#lacp timeout short
AX-1(config-if:ethernet6)#end
```

Show Commands

The following command shows the LACP system ID:

```
AX-1#show lacp sys-id
 System 0064,00-1f-a0-01-d4-f0
```

The following command shows details about the LACP admin keys:

```
AX-1#show lacp trunk admin-key-list-details
% Admin Key: 1
   bandwidth: 0
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 10001
   ref count: 2
% Admin Key: 2
   bandwidth: 1
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 0
   ref count: 451
% Admin Key: 3
   bandwidth: 1
   mtu: 16436
   duplex mode: 0
   hardware type: 1
   type: 0
   additional parameter: 0
   ref count: 14
% Admin Key: 4
   bandwidth: 1
   mtu: 1500
   duplex mode: 0
   hardware type: 2
   type: 0
   additional parameter: 0
   ref count: 6
```

The following command shows summary trunk information:

```
AX-1#show lacp trunk summary
 Aggregator po5 1000000
  Admin Key: 0005 - Oper Key 0005
   Link: ethernet 1 (3) sync: 1
   Link: ethernet 2 (4) sync: 1
 Aggregator po10 1000001
  Admin Key: 0010 - Oper Key 0010
   Link: ethernet 6 (8) sync: 1
```

The following command shows information for trunk 5:

```
AX-1#show lacp trunk 5
 Aggregator po5 1000000 Admin Key: 0005 - Oper Key 0005 Partner LAG: 0x0064,00-
1f-a0-01-dc-60 Partner Oper Key 0005
   Link: ethernet 1 (3) sync: 1
   Link: ethernet 2 (4) sync: 1
```

The following command shows detailed information for all LACP trunks:

```
AX-1#show lacp trunk detail
 Aggregator po5 1000000
  Mac address: 00:1f:a0:02:1e:48
  Admin Key: 0005 - Oper Key 0005
  Receive link count: 1 - Transmit link count: 0
  Individual: 0 - Ready: 1
  Partner LAG- 0x0064,00-1f-a0-01-dc-60
   Link: ethernet 1 (3) sync: 1
   Link: ethernet 2 (4) sync: 1
 Aggregator po10 1000001
  Mac address: 00:1f:a0:02:1e:4d
  Admin Key: 0010 - Oper Key 0010
  Receive link count: 1 - Transmit link count: 0
  Individual: 0 - Ready: 1
  Partner LAG- 0x8000,00-1f-a0-10-19-66
   Link: ethernet 6 (8) sync: 1
```

The following command shows LACP information for Ethernet data port 1:

```
AX-1#show lacp trunk port ethernet 1
 LACP link info: ethernet 1 - 3
 LAG ID: 0x8000,00-1f-a0-02-1e-48
 Partner oper LAG ID: 0x8000,00-1f-a0-01-dc-60
 Actor priority: 0x8000 (32768)
 Admin key: 0x0005 (5) Oper key: 0x0005 (5)
 Physical admin key:(1)
 Receive machine state : Current
 Periodic Transmission machine state : Slow periodic
 Mux machine state : Collecting/Distributing
 Oper state: ACT:1 TIM:0 AGG:1 SYN:1 COL:1 DIS:1 DEF:0 EXP:0
 Partner oper state: ACT:1 TIM:0 AGG:1 SYN:1 COL:1 DIS:1 DEF:0 EXP:0
 Partner link info: admin port 0
 Partner oper port: 3
 Partner admin LAG ID: 0x0000-00:00:00:00:0000
 Admin state: ACT:1 TIM:0 AGG:1 SYN:0 COL:0 DIS:0 DEF:1 EXP:0
 Partner admin state: ACT:0 TIM:0 AGG:1 SYN:0 COL:0 DIS:0 DEF:1 EXP:0
 Partner system priority - admin:0x8000 - oper:0x0064
 Aggregator ID: 1000000
```

IPv6 Duplicate Address Detection and Logging

When an IPv6 address is configured on the AX device, Duplicate Address Detection (DAD) checks to ensure that the same address is not already in use by another device. If the address is already in use, a warning message such as the following appears in the log.

```
Sep 30 2011 05:33:29 Warning [IPV6]:(10 times) Duplicate IPv6 Address Detected: Neighbor
Advertisement received for e101::1112 from Port: 5 Mac: 021f.a000.0021 Vlan: 2
```

If a neighbor tries to configure or use an IPv6 address that is already in use by the AX device, a warning such as the following appears in the log.

```
Sep 30 2011 08:10:32 Warning [IPV6]:(8 times) Duplicate IPv6 Address Detected: Neighbor
Solicitation received for e101::1112 from Port: 5 Mac: 021f.a000.0021 Vlan: 2
```

Notes

- In the current release, detection of duplicate IPv6 addresses is based only on incoming neighbor solicitations and advertisements. In future releases, when an IPv6 address is configured on the AX device that is already in use by another device, the address configuration will fail and an error message will be displayed.
- DAD log messages are rate limited. In the event of duplicate address detections, the warning messages are added to the log once every 10 seconds.

Displaying DAD Statistics

To display DAD statistics, use the following command:

show slb switch

Here is an example:

```
AX#show slb switch
                         Total
------------------------------------------------------------------
L2 Forward               2793
L3 IP Forward            0
...
IPV6 DAD on Solicits     8
IPV6 DAD on Adverts      10
IPV6 DAD MAC Conflicts   2
IPV6 DAD out-of-memory   0
```

Gateway Health Monitoring

This chapter describes gateway health monitoring and how to configure it.

Notes

- Gateway health monitoring is useful in cases where there is more than one route to a destination. In this case, the AX device can discard the routes that use unresponsive gateways. If there is only one gateway, this feature is not useful.
- Gateway health monitoring and SLB server health monitoring are independent features. If a gateway fails its health check, a server reached through the gateway is not immediately marked down. The status of the server still depends on the result of the SLB server health check.
- If you plan to use gateway health as a failover trigger for High Availability (HA), a different configuration option is required. See “Gateway-based Failover” on page 180.

Overview

Gateway health monitoring uses ARP to test the availability of nexthop gateways. When the AX device needs to send a packet through a gateway, the AX device begins sending ARP requests to the gateway.

- If the gateway replies to any ARP request within a configurable timeout, the AX device forwards the packet to the gateway.
- The ARP requests are sent at a configurable interval. The AX device waits for a configurable timeout for a reply to any request. If the gateway does not respond to any request before the timeout expires, the AX device selects another gateway and begins the health monitoring process again.

The following parameters are used for gateway health monitoring:

- Interval – The interval specifies the amount of time between health check attempts (ARP requests), and can be 1-180 seconds. The default is 5 seconds.
- Retries – The retries specifies the number of seconds the AX device waits for a response before re-attempting the health check (sending another ARP request). The maximum value is 3 seconds and is not configurable.

- Timeout – The timeout specifies how long the AX device waits for a reply to any of the ARP requests, and can be 1-60 seconds. The default is 15 seconds.

Using the default gateway health monitoring settings, a gateway must respond to a gateway health check within 15 seconds. Figure 16 shows how a gateway health check times out using the default settings.

Note: It is recommended not to use a timeout value smaller than 3 times the interval value. This is especially true for short interval values.

FIGURE 16 Gateway Health Check Using Default Settings – Timeout

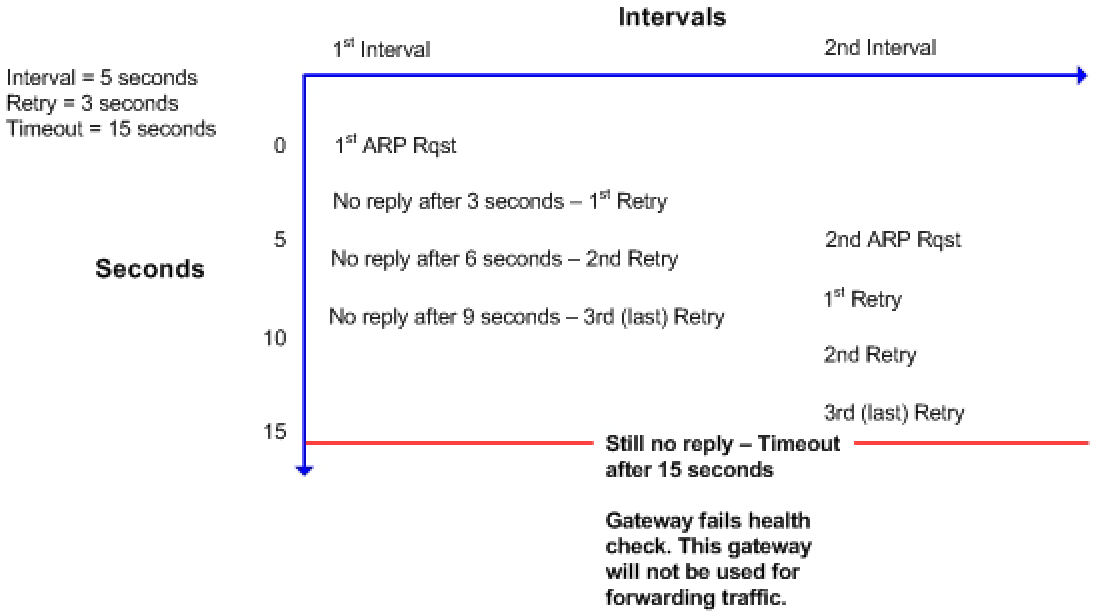
Intervals
1st Interval
2nd Interval
Interval = 5 seconds
Retry = 3 seconds
Timeout = 15 seconds
Seconds
0
5
10
15
1st ARP Rqst
No reply after 3 seconds – 1st Retry
No reply after 6 seconds – 2nd Retry
No reply after 9 seconds – 3rd (last) Retry
2nd ARP Rqst
1st Retry
2nd Retry
3rd (last) Retry
Still no reply – Timeout after 15 seconds
Gateway fails health check. This gateway will not be used for forwarding traffic.

Figure 17 shows an example in which a gateway responds before the timeout.

FIGURE 17 Gateway Health Check Using Default Settings – Gateway Responds

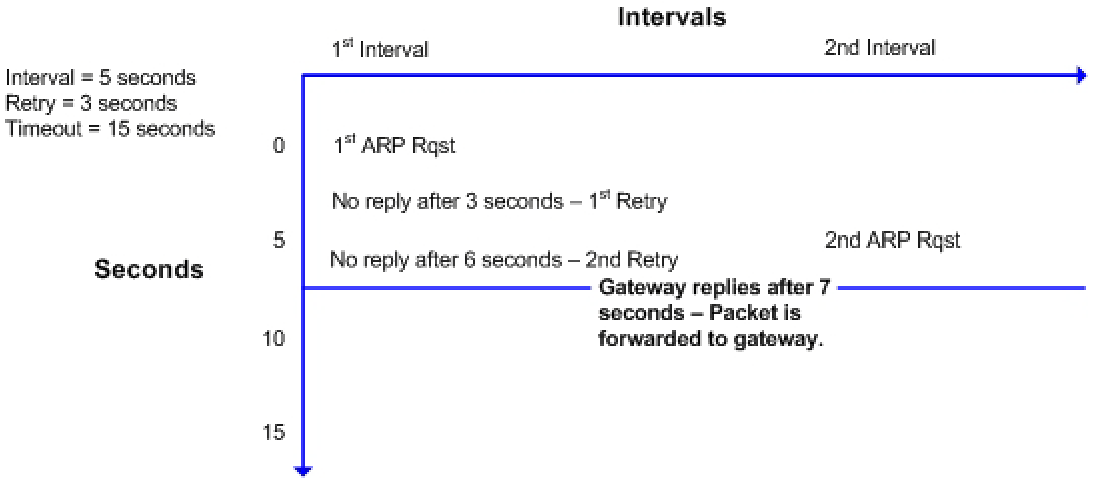
Intervals
1st Interval
2nd Interval
Interval = 5 seconds
Retry = 3 seconds
Timeout = 15 seconds
0
1st ARP Rqst
No reply after 3 seconds – 1st Retry
5
2nd ARP Rqst
Seconds
No reply after 6 seconds – 2nd Retry
Gateway replies after 7 seconds – Packet is forwarded to gateway.
10
15

Configuration

Gateway health monitoring is disabled by default. To enable it, use either of the following methods.

Using the GUI

The current release does not support configuration of this feature using the GUI.

Using the CLI

To enable gateway health monitoring, use the following command at the global configuration level of the CLI:

```
slb gateway-health-check
[interval seconds [timeout seconds]]
```

The **interval** and **timeout** options are described above.

To display gateway health monitoring statistics, use the following command:

show health gateway

CLI Example

The following command enables gateway health monitoring with the default settings:

```
AX(config)#slb gateway-health-check
```

The following command displays gateway health monitoring statistics:

```
AX(config)#show health gateway
Gateway health-check is enabled
Interval=5, Timeout=15
Total health-check sent        :  10
Total health-check retry sent  :  2
Total health-check timeout     :  1
```

Open Shortest Path First (OSPF)

The AX device supports the following OSPF versions:

- OSPFv2 for IPv4
- OSPFv3 for IPv6

This chapter provides configuration examples. For detailed CLI syntax information, see the *AX Series CLI Reference*.

Support for Multiple OSPFv2 and OSPFv3 Processes

The AX device supports up to 65535 OSPFv2 processes on a single AX device. Only a single OSPFv2 process can run on a given interface.

Each IPv6 link can run up to 65535 OSPFv3 processes, on the same link.

Each OSPF process is completely independent of the other OSPF processes on the device. They do not share any information directly. However, you can configure redistribution of routes between them.

Support for OSPFv2 and OSPFv3 on the Same Interface or Link

You can configure OSPFv2 and OSPFv3 on the same interface or link. OSPFv2 configuration commands affect only the IPv4 routing domain, while OSPFv3 configuration commands affect only the IPv6 routing domain.

OSPF MIB Support

The following OSPF MIBs are supported:

- RFC 1850 – OSPFv2 Management Information Base
- draft-ietf-ospf-ospfv3-mib-08 – OSPFv3 Management Information Base

OSPF Configuration Example

The configuration excerpts in this example configure OSPFv2 and OSPFv3 on an AX device.

Interface Configuration

The following commands configure two physical Ethernet data interfaces. Each interface is configured with an IPv4 address and an IPv6 address. Each interface also is added to OSPF area 0 (the backbone area).

The link-state metric (OSPF cost) of Ethernet 2 is set to 30, which is higher than the default, 10. Based on the cost difference, OSPF routes through Ethernet 1 will be favored over OSPF route through Ethernet 2, because the OSPF cost of Ethernet 1 is lower.

```
interface ethernet 1
 ip address 2.2.10.1 255.255.255.0
 ipv6 address 5f00:1:2:10::1/64
 ipv6 router ospf area 0 tag 1
!
interface ethernet 2
 ip address 3.3.3.1 255.255.255.0
 ipv6 address 5f00:1:2:20::1/64
 ip ospf cost 25
 ipv6 router ospf area 0 tag 1
```

The following commands configure two Virtual Ethernet (VE) interfaces. On VE 3, an IPv4 address is configured. On VE 4, an IPv4 address and an IPv6 address are configured.

OSPFv2 authentication is configured on VE 3, and the OSPF cost is set to 20.

On VE 4, the OSPF cost is set to 15.

```
interface ve 3
 ip address 1.1.1.2 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 abc
 ip ospf cost 20
!
```

```
interface ve 4
 ip address 1.1.60.2 255.255.255.0
 ipv6 address 5f00:1:1:60::2/64
 ip ospf cost 15
```

Global OSPF Parameters

The following commands configure global settings for OSPFv2 process 2. The router ID is set to 2.2.2.2. Subnets 1.1.x.x, 2.2.10.x, and 3.3.3.x are added to the backbone area. Redistribution is enabled for static routes, routes to VIPs, IP source NAT addresses, and floating IP addresses (used by HA). In addition, an extra HA cost is configured, and the SPF timer is changed.

```
router ospf 2
 ospf router-id 2.2.2.2
 ha-standby-extra-cost 25
 timers spf exp 500 50000
 redistribute static metric 5 metric-type 1
 redistribute vip metric 500 metric-type 1
 redistribute ip-nat
 redistribute floating-ip metric-type 1
 network 1.1.0.0 0.0.255.255 area 0
 network 2.2.10.0 0.0.0.255 area 0
 network 3.3.3.0 0.0.0.255 area 0
```

The following commands configure global settings for OSPFv3 process 1. The router ID is set to 3.3.3.3. A stub area is added, redistribution is enabled, and the SPF timer is changed.

```
router ipv6 ospf 1
 router-id 3.3.3.3
 redistribute static metric 5 metric-type 1
 redistribute ip-nat
 redistribute floating-ip
 area 1 stub
 timers spf exp 500 50000
```

OSPF Logging

Router logging is disabled by default. You can enable router logging to one or more of the following destinations:

- CLI terminal (stdout)
- Local logging buffer
- Local file
- External log servers

Note: Log file settings are retained across reboots but debug settings are not.

Note: Enabling debug settings that produce lots of output, or enabling all debug settings, is not recommend for normal operation.

Configuring Router Logging for OSPF

To configure router logging for OSPF:

1. Enable output options.
2. Set severity level and facility.
3. Enable debug options to generate output.

For additional syntax information, including **show** and **clear** commands for router logging, see the *AX Series CLI Reference*.

1. Enable output options:

To enable output to the terminal, use the following command at the global configuration level of the CLI:

```
router log stdout
```

To enable output to the local logging buffer, use the following command at the global configuration level of the CLI:

```
router log syslog
```

To enable output to a local file, use the following command at the global configuration level of the CLI:

```
[no] router log file
{
name string |
per-protocol |
rotate num |
size Mbytes
}
```

To enable output to a remote log server, use the following command at the global configuration level of the CLI:

```
logging host ipaddr [ipaddr...]
[port protocol-port]
```

Up to 10 remote logging servers are supported.

2. Set severity level and facility:

The default severity level for router logging is 7 (debugging). The default facility is local0.

To change set the severity level for messages output to the terminal, use the following command at the global configuration level of the CLI:

```
logging monitor severity-level
```

The severity-level can be one of the following:

- **0** or **emergency**
- **1** or **alert**
- **2** or **critical**
- **3** or **error**
- **4** or **warning**
- **5** or **notification**
- **6** or **information**
- **7** or **debugging**

To change the severity level for messages output to the local logging buffer, use the following command at the global configuration level of the CLI:

```
logging buffered severity-level
```

To change the severity level for messages output to external log servers, use the following command at the global configuration level of the CLI:

```
logging syslog severity-level
```

To change the severity level for messages output to a file, use the following command at the global configuration level of the CLI:

```
router log trap severity-level
```

To change the facility, use the following command at the global configuration level of the CLI:

```
logging facility facility-name
```

The *facility-name* can be one of the following:

- `local0`
- `local1`
- `local2`
- `local3`
- `local4`
- `local5`
- `local6`
- `local7`

3. Enable debug options to generate output:

To enable debugging for OSPF, use the following commands at the global configuration level or Privileged EXEC level of the CLI:

```
debug a10 [ipv6] ospf
debug [ipv6] ospf type
```

The **ipv6** option enables debugging for OSPFv3. Without the **ipv6** option, debugging is enabled for OSPFv2.

The *type* specifies the types of OSPF information to log, and can be one or more of the following:

- **all** – Enables debugging for all information types listed below.
- **events** – Enables debugging for OSPF events.
- **ifsm** – Enables debugging for the OSPF Interface State Machine (IFSM).
- **lsa** – Enables debugging for OSPF Link State Advertisements (LSAs).
- **nfsm** – Enables debugging for the OSPF Neighbor State Machine (NFSM).

- **nsm** – Enables debugging for the Network Services Module (NSM). The NSM deals with use of ACLs, route maps, interfaces, and other network parameters.
- **packet** – Enables debugging for OSPF packets.
- **route** – Enables debugging for OSPF routes.

For each level, both **debug** commands are required.

CLI Example

The following commands configure OSPFv2 logging to a local file.

```
AX(config)#router log file name ospf-log
AX(config)#router log file per-protocol
AX(config)#router log file size 100
AX(config)#debug a10 ospf all
AX(config)#debug ospf packet
```

These commands create a router log file named “ospf-log”. The **per-protocol** option will log messages for each routing protocol separately. The log file will hold a maximum 100 MB of data, after which the messages will be saved in a backup and the log file will be cleared.

The following command displays the contents of the local router log file:

```
AX(config)#show router log file ospfd
2010/04/21 09:57:20 OSPF: IFSM[ve 3:1.1.1.2]: Hello timer expire
2010/04/21 09:57:20 OSPF: SEND[Hello]: To 224.0.0.5 via ve
3:1.1.1.2,
length
64
2010/04/21 09:57:20 OSPF:
-----------------------------------------------------
2010/04/21 09:57:20 OSPF: Header
2010/04/21 09:57:20 OSPF:   Version 2
2010/04/21 09:57:20 OSPF:   Type 1 (Hello)
2010/04/21 09:57:20 OSPF:   Packet Len 48
2010/04/21 09:57:20 OSPF:   Router ID 2.2.2.2
2010/04/21 09:57:20 OSPF:   Area ID 0.0.0.0
2010/04/21 09:57:20 OSPF:   Checksum 0x0
2010/04/21 09:57:20 OSPF:   Instance ID 0
2010/04/21 09:57:20 OSPF:   AuType 2
2010/04/21 09:57:20 OSPF:   Cryptographic Authentication
2010/04/21 09:57:20 OSPF:   Key ID 1
```

```
2010/04/21 09:57:20 OSPF:   Auth Data Len 16
2010/04/21 09:57:20 OSPF:   Sequence number 1271830931
2010/04/21 09:57:20 OSPF: Hello
2010/04/21 09:57:20 OSPF:   NetworkMask 255.255.255.0
2010/04/21 09:57:20 OSPF:   HelloInterval 10
2010/04/21 09:57:20 OSPF:   Options 0x2 (-|-|-|-|-|-|E|-)
2010/04/21 09:57:20 OSPF:   RtrPriority 1
2010/04/21 09:57:20 OSPF:   RtrDeadInterval 40
2010/04/21 09:57:20 OSPF:   DRouter 1.1.1.200
2010/04/21 09:57:20 OSPF:   BDRouter 1.1.1.2
2010/04/21 09:57:20 OSPF:   # Neighbors 1
2010/04/21 09:57:20 OSPF:     Neighbor 31.31.31.31
2010/04/21 09:57:20 OSPF:
-----------------------------------------------------
2010/04/21 09:57:21 OSPF: IFSM[ethernet 2:3.3.3.1]: Hello timer
expire
2010/04/21 09:57:21 OSPF: SEND[Hello]: To 224.0.0.5 via ethernet
2:3.3.3.1,
length 48
2010/04/21 09:57:21 OSPF:
-----------------------------------------------------
2010/04/21 09:57:21 OSPF: Header
2010/04/21 09:57:21 OSPF:   Version 2
2010/04/21 09:57:21 OSPF:   Type 1 (Hello)
2010/04/21 09:57:21 OSPF:   Packet Len 48
2010/04/21 09:57:21 OSPF:   Router ID 2.2.2.2
2010/04/21 09:57:21 OSPF:   Area ID 0.0.0.0
2010/04/21 09:57:21 OSPF:   Checksum 0x49eb
2010/04/21 09:57:21 OSPF:   Instance ID 0
2010/04/21 09:57:21 OSPF:   AuType 0
2010/04/21 09:57:21 OSPF: Hello
2010/04/21 09:57:21 OSPF:   NetworkMask 255.255.255.0
2010/04/21 09:57:21 OSPF:   HelloInterval 10
2010/04/21 09:57:21 OSPF:   Options 0x2 (-|-|-|-|-|-|E|-)
2010/04/21 09:57:21 OSPF:   RtrPriority 1
2010/04/21 09:57:21 OSPF:   RtrDeadInterval 40
2010/04/21 09:57:21 OSPF:   DRouter 3.3.3.2
2010/04/21 09:57:21 OSPF:   BDRouter 3.3.3.1
2010/04/21 09:57:21 OSPF:   # Neighbors 1
2010/04/21 09:57:21 OSPF:     Neighbor 81.81.81.81
...
```

Bidirectional Forwarding Detection

The current release provides additional support for Bidirectional Forwarding Detection (BFD).

BFD provides very fast failure detection for routing protocols. When BFD is enabled, the AX device periodically sends BFD control packets to the neighboring devices that are also running BFD. If a neighbor stops sending BFD control packets, the AX device quickly brings down the BFD session(s) with the neighbor, and recalculates paths for routes affected by the down neighbor.

BFD provides a faster failure detection mechanism than the timeout values used by routing protocols. Routing protocol timers are multiple seconds long, whereas BFD provides sub-second failover.

The A10 implementation of BFD is based on the following RFCs:

- RFC 5880, Bidirectional Forwarding Detection (BFD)
- RFC 5881, Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)
- RFC 5882, Generic Application of Bidirectional Forwarding Detection (BFD)
- RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths

Support in this Release

The current release has the following BFD support:

- Basic BFD protocol (packet processing, state machine, and so on)
- BGP client support
- Multihop
- BFD Asynchronous mode
- OSPFv2/v3 client support
- Static route support
- IS-IS client support
- BFD Demand mode
- Full Echo function support

- Authentication

BFD Parameters

BFD is disabled by default. You can enable it on a global basis.

BFD Echo

BFD echo enables a device to test data path to the neighbor and back. When a device generates a BFD echo packet, the packet uses the routing link to the neighbor device to reach the device. The neighbor device is expected to send the packet back over the same link.

BFD Timers

You can configure BFD timers at the following configuration levels:

- Global
- Interface

If you configure the timers on an individual interface, the interface’s settings are used instead of the global settings. Likewise, if the BFD timers are not set on an interface, that interface uses the global settings. For BGP loopback neighbors, BFD always uses the global timer.

The DesiredMinTXInterval, RequiredMinRxInterval and DetectMult timer fields can be configured at the interface and the global configuration level. However, the actual timer will vary depending on the Finite State Machine (FSM) state, through negotiation, and whether or not echo has been enabled.

BGP Support

If you run BGP on the AX device, you can enable BFD-based fallover for individual BGP neighbors.

Static Route Support

A static route flap can occur when you enable BFD in global mode or when you configure a static BFD session.

In the following example, you will see that the static routes experience a flap when BFD is enabled. The fields to note are flagged in bold:

```
AX(config)#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       i - IS-IS, B - BGP
Timers: Uptime

C   3ffe:100::/64 via ::, ve 10, 00:01:28
C   3ffe:1111::/64 via ::, loopback 1, 00:01:30
S   3ffe:2222::/64 [1/0] via 3ffe:100::20, ve 10, 00:00:25  <<===value before flap
timer
C   3ffe:3333::/64 via ::, loopback 2, 00:01:30
AX(config)#bfd enable   <<<===enable BFD
AX(config)#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       i - IS-IS, B - BGP
Timers: Uptime

C   3ffe:100::/64 via ::, ve 10, 00:01:32
C   3ffe:1111::/64 via ::, loopback 1, 00:01:34
S   3ffe:2222::/64 [1/0] via 3ffe:100::20, ve 10, 00:00:01  <<<===value after flap
C   3ffe:3333::/64 via ::, loopback 2, 00:01:34
AX(config)#
```

Using the GUI

The current release does not support BFD configuration using the GUI.

Using the CLI

To enable BFD, use the following command at the global configuration level of the CLI:

[no] **bfd enable**

To enable BFD echo, use the following command at the global configuration level of the CLI:

[no] **bfd echo**

To configure BFD timers, use the following commands. These commands are available at the global configuration level and at the configuration level for individual interfaces.

[no] **bfd interval** *ms* **min-rx** *ms* **multiplier** *num*

The **interval** value can be 48-1000 ms, and is 800 ms by default. The **min-rx** value can be 48-1000 ms, and is 800 ms by default. The **multiplier** value can be 3-50 and is 4 by default.

Configuring BFD Parameters for BGP

To enable BFD-based fallover for a BGP neighbor, use the following command at the BGP configuration level:

```
[no] neighbor ipaddr fall-over bfd [multihop]
```

To display BFD information for BGP neighbors, use the following command:

```
show ip bgp neighbor
```

Displaying BFD Information

To display summarized BFD neighbor information, use the following command:

```
show bfd neighbors
```

To display detailed BFD neighbor information, use the following command:

```
show bfd neighbors detail
```

To display BFD statistics, use the following command:

```
show bfd statistics
```

To display BFD statistics, use the following command:

```
show bfd statistics
```

To clear BFD statistics, use the following command:

```
clear bfd statistics
```

Disable BFD

To disable BFD, enter the following command in global configuration mode:

```
  AX(config)#no bfd enable
```

Enter the command to stop processing all BFD packets.

Configure BFD with OSPF (for IPv4)

To enable BFD with OSPF on an interface, enter one of the following sets of commands:

To enable BFD on an individual interface:

```
AX(config)# interface ethernet 1
AX(config-if)# ipv6 router ospf area 0 tag 1
AX(config-if)# ip address 20.0.0.1 255.255.255.0
AX(config-if)# ip ospf bfd
```

- To enable BFD on a virtual interface:

```
AX(config)# interface ve 100
AX(config-if)# ip ospf bfd
```

- To enable BFD on a trunk:

```
AX(config)#interface trunk 1
AX(config-if)#ip ospf bfd
```

To enable BFD for all OSPF-enabled interfaces, enter the following commands:

```
AX(config)#router ospf 1
AX(config-router)# bfd all-interfaces
```

To selectively disable BFD per interface, enter the following command:

```
AX(config)#interface ethernet 1
AX(config)#ip ospf bfd disable
```

To configure a multihop neighbor over a virtual-link, enter the following command:

```
AX(config-router)# area 1 virtual-link 40.0.0.1
fall-over bfd
```

Sample Configuration

Your running configuration will display your current BFD with OSPF configuration:

```
!
interface ethernet 1
 ipv6 router ospf area 0 tag 1
 ip address 20.0.0.1 255.255.255.0
 ip ospf bfd
!
interface ethernet 2
```

```
 ipv6 router ospf area 0 tag 1
 ip address 30.0.0.1 255.255.255.0
!
!
router ospf 1
 bfd all-interfaces
 network 20.0.0.0/24 area 0
 network 30.0.0.0/24 area 0
 area 1 virtual-link 40.0.0.1 fall-over bfd
!
!
bfd enable
!
```

Configure BFD with OSPF (for IPv6)

To enable BFD with OSPF for IPv6 support on an interface, enter one of the following sets of commands:

To enable BFD on an individual interface:

```
AX(config)# interface ethernet 1
AX(config-if)#ipv6 address 2001::1/64
AX(config-if)#ipv6 router ospf area 0 tag 1
AX(config-if)#ipv6 ospf bfd
```

- To enable BFD on a virtual interface:

```
AX(config)# interface ve 100
AX(config-if)# ipv6 ospf bfd
```

- To enable BFD on a trunk:

```
AX(config)#interface trunk 1
AX(config-if)#ipv6 ospf bfd
```

To enable BFD for all OSPFv3-enabled interfaces, enter the following commands:

```
AX(config)#router ipv6 ospf 1
AX(config-router)#bfd all-interfaces
```

To selectively disable BFD per interface, enter the following command:

```
AX(config)#interface ethernet 1
AX(config-if)#ipv6 ospf bfd disable
```

To configure a multihop neighbor over a virtual-link, enter the following command:

```
AX(config-router)# area 1 virtual-link 2.2.2.2
fall-over bfd
```

Sample Configuration

Your running configuration will display your current BFD with OSPF for IPv6 configuration:

```
!
interface ethernet 1
 ipv6 address 2001::1/64
 ipv6 router ospf area 0 tag 1
 ipv6 ospf bfd
!
interface ethernet 2
 ipv6 router ospf area 0 tag 1
 ipv6 address 3001::1/64
!
!
router ipv6 ospf 1
 router-id 1.1.1.1
 bfd all-interfaces
 area 1 virtual-link 2.2.2.2 fall-over bfd
!
!
bfd enable
!
```

Configure BFD with IS-IS (for IPv4)

To enable BFD with ISIS on an interface, enter one of the following sets of commands:

- To enable BFD on an individual interface:
```
AX(config)#interface ethernet 1
AX(config-if)#ip address 20.0.0.1 255.255.255.0
AX(config-if)#ip router isis
AX(config-if)#isis bfd
```
- To enable BFD on a virtual interface:
```
AX(config)#interface ve 100
AX(config-if)#isis bfd
```
- To enable BFD on a trunk:
```
AX(config)#interface trunk 1
AX(config-if)#isis bfd
```

To enable BFD for all IS-IS-enabled interfaces, enter the following commands:

```
AX(config)#router isis
AX(config-router)#bfd all-interfaces
AX(config-router)#net 49.0001.0000.0000.0001.00
```

To selectively disable BFD per interface, enter the following command:

```
AX(config)#interface ethernet 1
AX(config-if)#isis bfd disable
```

Sample Configuration

Your running configuration will display your current BFD with ISIS configuration:

```
!
interface ethernet 1
 ip address 20.0.0.1 255.255.255.0
 ip router isis
 isis bfd
!
interface ethernet 2
 ip address 30.0.0.1 255.255.255.0
 ip router isis
 isis bfd
!
```

```
!
router isis
 bfd all-interfaces
 net 49.0001.0000.0000.0001.00
!
!
bfd enable
!
```

Configure BFD with IS-IS (for IPv6)

To enable BFD with ISIS for IPv6 support on an interface, enter one of the following sets of commands:

- To enable BFD on an individual interface:
  ```
  AX(config)#interface ethernet 1
  AX(config-if)# isis bfd
  ```
- To enable BFD on a virtual interface:
  ```
  AX(config)#interface ve 100
  AX(config-if)#ipv6 address 2ffe:123::1/64
  AX(config-if)#ipv6 router isis
  AX(config-if)#isis bfd
  ```
- To enable BFD on a trunk:
  ```
  AX(config)#interface trunk 1
  AX(config-if)#isis bfd
  ```

To enable BFD for all IS-IS-enabled interfaces, enter the following commands:

```
AX(config)#router isis
AX(config-router)#bfd all-interfaces
AX(config-router)#net 49.0001.0000.0000.0002.00
```

To selectively disable BFD per interface, enter the following command:

```
AX(config)#interface ethernet 1
AX(config-if)#isis bfd disable
```

Sample Configuration

Your running configuration will display your current BFD with ISIS (for IPv6 support) configuration:

```
!
interface ve 100
ipv6 address 2ffe:123::1/64
ipv6 router isis
isis bfd
!
router isis
 bfd all-interfaces
 net 49.0001.0000.0000.0002.00
!
!
bfd enable
```

Configure BFD with BGP

When BFD is configured with BGP, it is configured on a per neighbor basis. This is different from the OSPF or ISIS configuration with BFD. Use the following commands to configure BFD with BGP:

```
AX(config)# router bgp 1
AX(config-router)# neighbor 1.2.3.4 fall-over bfd
```

To configure a multihop BFD neighbor, use the following command:

```
AX(config-router)# neighbor 1.2.3.4 fall-over bfd
multihop
```

Sample Configuration

Your running configuration will display your current BFD with BGP configuration:

```
!
router bgp 1
 neighbor 1.2.3.4 remote-as 2
 neighbor 1.2.3.4 fall-over bfd multihop
!
!
bfd enable
!
```

Configuring Static BFD

IPv4 Static BFD (Global)

From the global configuration mode, use the following command to add a static BFD entry for the specified IPv4 nexthop:

```
AX(config)# ip route static bfd 20.0.0.1 20.0.0.2
```

In the above command, the first parameter is the IPv4 address of the local interface. You can only use the IP addresses for interfaces to setup the BFD session. The second parameter is the IPv4 address of the remote interface that serves as the gateway for the static route.

IPv6 Static BFD (Global)

From the global configuration mode, use the following command to add a static BFD entry for the specified IPv6 nexthop:

```
AX(config)# ipv6 route static bfd 2001::1 2001::2
```

In the above command, the first parameter is the IPv6 address of the local interface. You can only use the IP addresses for interfaces to setup the BFD session. The second parameter is the IPv6 address of the remote interface that serves as the gateway for the static route.

IPv6 Static BFD (Link-Local)

From the global configuration mode, use the following command to add a static BFD entry for the specified link-local IPv6 nexthop:

```
AX(config)# ipv6 route static bfd ve 100 fe80::1
```

In the above command, the first parameter is the local interface name (Ethernet, VE, or a specified trunk), and the second parameter is the remote link-local IPv6 address that serves as the gateway.

Configuration of BFD Intervals

Global Interval Configuration

From the global configuration mode, use the following command to modify the global interval timer values:

```
AX(config)# bfd interval 500 min-rx 500 multiplier
4
```

This command will help configure the interval for any one of the following three parameters and will be applied to all BFD sessions:

- DesiredMinTxInterval
- RequiredMinRxInterval
- Multiplier

Interface Interval Configuration

From the interface configuration mode, use the following command to modify the interface interval timer values:

```
AX(config)# interface ve 10
AX(config-if)# bfd interval 500 min-rx 500 multi-
plier 4
```

Note: For a BFD session for BGP using a loopback address, for an OSPFv2 virtual link, and for an OSPFv3 virtual link, the AX device will always use the global timer configuration, immaterial of the timer that is configured at the interface level.

Authentication

Authentication Per interface

To configure authentication per interface, from the interface configuration mode, apply one of the following authentication schemes to OSPF, OSPFv3, IS-IS, or static BFD neighbors.

```
AX(config-if)# bfd authentication 1 md5
password-string
```

You may choose an authentication method from the following available choices:

- Simple password
- Keyed MD5
- Meticulous Keyed MD5
- Keyed SHA1
- Meticulous Keyed SHA1

Authentication Per Neighbor (for BGP only)

The following command is configured under the BGP (`router bgp`) configuration mode:

```
AX(config-router)# neighbor 1.2.3.4 fall-over bfd authentication 1 md5
password-string
```

Enable Echo and Demand function

Enable the Echo Function

From the global configuration mode, enable the BFD echo:

```
AX(config)# bfd echo
```

Enable the Echo Function Per Interface

After you configure the global BFD echo, from the interface configuration mode, you can enable BFD echo on a per interface basis using the following command:

```
AX(config-if)# bfd echo
```

Enable Demand Mode

From the interface configuration mode, you can enable the demand mode to work in conjunction with the echo function using the following command:

```
AX(config-if)# bfd echo demand
```

When demand mode is enabled, after a BFD session is established, a system will be able to verify connectivity with another system at will instead of routinely. Instead of constantly receiving BFD control packets, the system will request that the other system stop sending BFD Control packets. To verify connectivity again, the system will explicitly send a short sequence of BFD Control packets to the other system and receive a response. Demand mode can be configured to work either independently in each direction, or bidirectionally at the same time.

Asynchronous Mode

The Asynchronous mode is the default mode of operation for BFD. In this mode, systems establish connectivity and know of each other’s existence by periodically exchanging BFD Control packets. A session between two connected systems is only declared down after several packets in a row are not received by the other system. BFD will operate in this mode if you do not configure or enable echo or demand.

BFD Summary Output

To check the BFD neighbor status, from the EXEC mode, execute the **show BDF neighbors** command:

```
AX# show bfd neighbors
Our Address            Neighbor Address          State   Holddown txint mult diag
219.0.0.1              219.0.0.2                  Up           150    50    3 3/0
219.0.1.1              219.0.1.2                  Up           150    50    3 3/0
219.0.2.1              219.0.2.2                  Up           150    50    3 0/0
219.0.3.1              219.0.3.2                  Up           150    50    3 0/0
219.0.4.1              219.0.4.2                  Up           150    50    3 3/0
219.0.5.1              219.0.5.2                  Up           150    50    3 3/0
219.0.6.1              219.0.6.2                  Up           150    50    3 0/0
219.0.7.1              219.0.7.2                  Up           150    50    3 3/0
```

Table 4 describes the fields in the command output.

TABLE 4 show bfd neighbors fields

Field	Description
Our Address	AX interface associated with the BFD session.
Neighbor Address	Neighbor interface associated with the BFD session.
State	Shows the local state of the session.
Holdtime	Maximum amount of time the AX device waits for a BFD control packet from the neighbor.
txint	Configured interval at which the AX device sends BFD control packets to the neighbor.
mult	Maximum number of consecutive times the AX device will wait for a BFD control packet from the neighbor.
diag	Diagnostic codes for the local and remote ends of the BFD session. For information, contact A10 Networks.

BFD Detailed Output

To check the BFD neighbor detailed status, from the EXEC mode, execute the **show BDF neighbors detail** command:

```
AX# show bfd neighbors detail
Our Address      219.0.0.1
Neighbor Address 219.0.0.2
 Clients OSPFv2, IS-IS
 Singlehop, Echo disabled, Demand disabled, UDP source port 53214
 Asynchronous mode, Authentication None
 CPU ID 2, Interface index 93
 Local State Up, Remote State Up, 2h:29m:45s up
 Local discriminator 0x00000fdf, Remote discriminator 0x0000006f
 Config DesiredMinTxInterval 50 milliseconds, RequiredMinRxInterval 50 milli-
seconds
 Local DesiredMinTxInterval 50 milliseconds, RequiredMinRxInterval 50 millisec-
onds
 Remote DesiredMinTxInterval 50 milliseconds, RequiredMinRxInterval 50 milli-
seconds
 Local Multiplier 3, Remote Multiplier 3
 Hold Down Time 150 milliseconds, Transmit Interval 50 milliseconds
 Local Diagnostic: Neighbor Signalled Session Down(3)
 Remote Diagnostic: No Diagnostic(0)
 Last sent echo sequence number 0x00000000
 Control Packet sent 215226, received 215195
 Echo Packet sent 0, received 0
```

Table 5 describes the fields in the command output.

TABLE 5 show bfd neighbors detail fields

Field	Description
Our Address	AX interface associated with the BFD session.
Neighbor Address	Neighbor interface associated with the BFD session.
Clients	Protocol that initiates this BFD session. It can be one or more of the following: Static, OSPFv2, OSPFv3, IS-IS, or BGP.
Singlehop (or Multihop)	BFD session can be either singlehop or multihop.
Echo	Indicates whether Echo functionality has been enabled or disabled.
Demand	Indicates whether Demand mode functionality has been enabled or disabled.
UDP source port	UDP source port used for this BFD session.

TABLE 5 show bfd neighbors detail fields (Continued)

Field	Description
Asynchronous mode (or Demand) mode	If configured and running, indicates whether BFD is operating in Asynchronous mode or Demand mode.
Authentication	Authentication method. This can be either “None” (if it is not configured) or one of the following supported authentication schemes: • Simple password • Keyed MD5 • Meticulous Keyed MD5 • Keyed SHA1 • Meticulous Keyed SHA1
CPU ID	Since BFD traffic is distributed across multiple data CPUs, this CPU ID refers to the one associated with the current BFD session.
Interface index	Interface index associated with the current BFD session. This index is used mostly for debugging purposes
Local State	Shows the local state the session. The state can be one of the following: • Init • Up • AdminDown • Down
Remote State	Shows the remote state the session. The state can be one of the following: • Init • Up • AdminDown • Down
Local discriminator	The local discriminator value that the AX device assigns for the current BFD session.
Remote discriminator	The remote discriminator value that the neighboring router claims.
Config	The configured timer values.
Local	The configured timer values sent in the last BFD control packet. This value is determined based on BFD package exchange and negotiation.
Remote	The timer values received in the last BFD control packet from the BFD neighbor.
Local Multiplier	The local multiplier sent in the last BFD packet.
Remote Multiplier	The remote multiplier received in the last BFD packet from the neighbor.

TABLE 5 show bfd neighbors detail fields (Continued)

Field	Description
Hold Down Time	The expiration time after which the BFD session will be brought down. This value is determined with the negotiated interval value and the remote multiplier value.
Transmit Interval	The periodic interval to send BFD control packets.
Local Diagnostic:	The diagnostic value sent in the last BFD control packet.
Remote Diagnostic:	The diagnostic value received in the last BFD control packet from the neighbor.
Last sent echo sequence number	A10 Network’s proprietary sequence number sent in the last echo packet.
Control Packet sent....received	Statistics of control packets for this BFD session.
Echo Packet sent...received	Statistics of echo packets received for this BFD session.

Viewing BFD Statistics

The following command shows BFD statistics:

```
AX(config)# show bfd statistics
IP Checksum error                        0
UDP Checksum error                       0
No session found with your_discriminator 39958
Multihop config mismatch                 0
BFD Version mismatch                     0
BFD Packet length field is too small     0
BFD Packet data is short                 0
BFD Packet DetectMult is invalid         0
BFD Packet Multipoint is invalid         0
BFD Packet my_discriminator is invalid   0
BFD Packet TTL/Hop Limit is invalid      0
BFD Packet auth length is invalid        0
BFD Packet auth mismatch                 0
BFD Packet auth type mismatch            0
BFD Packet auth key ID mismatch          0
BFD Packet auth key mismatch             103
BFD Packet auth seq# invalid             0
BFD Packet auth failed                   0
BFD local state is AdminDown             2
BFD Destination unreachable              1
BFD Other error                          0
```

Table 6 describes the fields in the command output.

TABLE 6 show bfd statistics fields

Field	Description
IP Checksum error	Number of BFD packets that had an invalid IP checksum.
UDP Checksum error	Number of BFD packets that had an invalid UDP checksum.
No session found with your_ discriminator	Number of BFD packets whose Your Discriminator value did not match a My Discriminator value on the AX device.
Multihop config mismatch	A multihop configuration mismatch occurs when an AX device receives a BFD packet with a source or destination that matches an existing BFD session. It can also be caused in two other scenarios: • Local is configured as singlehop, but the packet is received on the UDP port for multihop. • Local is configured as multihop, but packet is received on the UDP port for singlehop.
BFD Version mismatch	Number of BFD packets with a different BFD version than the one in use by the AX device.
BFD Packet length field is too small	Number of BFD packets whose Length field value was shorter than the minimum BFD packet length (24 bytes without authentication or 26 bytes with authentication).
BFD Packet data is short	The packet payload size is smaller than the BFD length value.
BFD Packet DetectMult is invalid	The value of the received DetectMult is “0”.
BFD Packet Multipoint is invalid	The value of the received multipoint flag is set to “1”.
BFD Packet my_ discriminator is invalid	Number of BFD packets whose My Discriminator value was invalid.
BFD Packet TTL/Hop Limit is invalid	In a singlehop BFD session, the IP time-to-live or IPv6 hop limit value must be 255. If a value other than 255 is detected, this field is incremented.
BFD Packet auth length is invalid	The BFD length without the BFD packet header does not match the expected authentication length byte value. The number of BFD control packets have wrong authentication lengths in bytes
BFD Packet auth type mismatch	Number of BFD packets carrying an authentication type that does not match the BFD authentication type configured on the AX device.
BFD Packet auth key ID mismatch	This field is incremented when the key ID in the authentication header does not match the one configured on the AX device.

TABLE 6 show bfd statistics fields (Continued)

Field	Description
BFD Packet auth key mismatch	This field is incremented when the received authentication key does not match the one configured on the AX device.
BFD Packet auth seq# invalid	This field is incremented when the received authentication sequence number is not equal to or greater than the sequence number received previously.
BFD Packet auth failed	Number of BFD packets with an incorrect authentication value.
BFD local state is AdminDown	Number of BFD packets received while the BFD session was administratively down.
BFD Destination unreachable	Number of times the destination IP address for a BFD neighbor was unreachable while the AX device was attempting to transmit a BFD packet to the neighbor.
BFD Other error	Number of BFD errors not counted in any of the fields above.

DHCP Relay

Dynamic Host Configuration Protocol (DHCP) is a mechanism commonly used by clients to auto-discover their addressing and other configuration information when connected to a network.

You can configure the AX device to relay DHCP traffic between DHCP clients and DHCP servers located in different VLANs or subnets.

DHCP relay is supported only for the standard DHCP protocol ports:

- Boot protocol server (BOOTPS) – UDP port 67
- Boot protocol client (BOOTPC) – UDP port 68

DHCP relay service is supported for IPv4 and IPv6.

DHCP is a Client-Server protocol and relies on broadcast communication between the client and server for packet exchanges. Accordingly, the clients and the servers must be in the same broadcast domain (Layer 2 VLAN) for this to work, since Layer 3 routers typically do not forward broadcasts. However, in most deployments it is not practical to have a DHCP server in each Layer 2 VLAN. Instead, it is typical to use a common DHCP server for all VLANs and subnets in the network.

To enable DHCP communication between different VLANs or subnets, you can use a DHCP relay. A DHCP relay acts as a mediator between the DHCP client and the DHCP server when they are not in the same broadcast domain.

To configure the AX device as a DHCP relay, configure the DHCP server IP address as a helper address on the AX IP interface connected to DHCP clients. The AX device intercepts broadcast DHCP packets sent by clients on the interface configured with the helper address.

The AX device then places the receiving interface’s IP address (not the helper address) in the relay gateway address field, and forwards the DHCP packet to the server. When the DHCP server replies, the AX device forwards the response to the client.

Notes

- In the current release, the helper-address feature provides service for DHCP packets only.
- The AX interface on which the helper address is configured must have an IP address.
- The helper address can not be the same as the IP address on any AX interface or an IP address used for SLB.

USING THE GUI

The current release does not support this feature in the GUI.

USING THE CLI

To configure a helper address, use the following command at the configuration level for the AX IP interface connected to the DHCP clients:

```
[no] ip helper-address ipaddr
```

To display DHCP relay information, use the following command:

```
show ip helper-address [detail]
```

To clear statistics counters for DHCP relay, use the following command:

```
clear ip helper-address statistics
```

CLI Example

The following commands configure two helper addresses. The helper address for DHCP server 100.100.100.1 is configured on AX Ethernet interface 1 and on Virtual Ethernet (VE) interfaces 5 and 7. The helper address for DHCP server 20.20.20.102 is configured on VE 9.

```
AX(config)#interface ethernet 1
AX(config-if:ethernet1)#ip helper-address 100.100.100.1
AX(config-if:ethernet1)#interface ve 5
AX(config-if:ve5)#ip helper-address 100.100.100.1
AX(config-if:ve5)#interface ve 7
AX(config-if:ve7)#ip helper-address 100.100.100.1
AX(config-if:ve7)#interface ve 9
AX(config-if:ve9)#ip helper-address 20.20.20.102
```

The following command shows summary DHCP relay information:

```
AX3200(config)#show ip helper-address
Interface  Helper-Address            RX            TX      No-Relay         Drops
---------  --------------  ------------  ------------  ------------  ------------
eth1       100.100.100.1              0             0             0             0
ve5        100.100.100.1           1669          1668             0             1
ve7                                1668          1668             0             0
ve8        100.100.100.1              0             0             0             0
ve9        20.20.20.102               0             0             0             0
```

Table 7 describes the fields in the command output.

TABLE 7 show ip helper-address fields

Field	Description
Interface	AX interface. Interfaces appear in the output in either of the following cases: • A helper address is configured on the interface. • DHCP packets are sent or received on the interface.
Helper-Address	Helper address configured on the interface.
RX	Number of DHCP packets received on the interface.
TX	Number of DHCP packets sent on the interface.
No-Relay	Number of packets that were examined for DHCP relay but were not relayed, and instead received regular Layer 2/3 processing. Generally, this counter increments in the following cases: • DHCP packets are received on an interface that does not have a helper address and the packets are not destined to the relay. • DHCP packets are received on an interface that does have a helper address, but the packets are unicast directly from the client to the server and do not need relay intervention.
Drops	Number of packets that were ineligible for relay and were dropped.

The following command shows detailed DHCP relay information:

```
AX#show ip helper-address detail
IP Interface: eth1
------------
  Helper-Address: 100.100.100.1
  Packets:
            RX: 0
               BootRequest Packets : 0
               BootReply Packets   : 0
            TX: 0
               BootRequest Packets : 0
               BootReply Packets   : 0
  No-Relay: 0
  Drops:
            Invalid BOOTP Port  : 0
            Invalid IP/UDP Len  : 0
            Invalid DHCP Oper   : 0
            Exceeded DHCP Hops  : 0
            Invalid Dest IP     : 0
            Exceeded TTL        : 0
            No Route to Dest    : 0
            Dest Processing Err : 0

IP Interface: ve5
------------
  Helper-Address: 100.100.100.1
  Packets:
            RX: 16
               BootRequest Packets : 16
               BootReply Packets   : 0
            TX: 14
               BootRequest Packets : 0
               BootReply Packets   : 14
  No-Relay: 0
  Drops:
            Invalid BOOTP Port  : 0
            Invalid IP/UDP Len  : 0
            Invalid DHCP Oper   : 0
            Exceeded DHCP Hops  : 0
            Invalid Dest IP     : 0
            Exceeded TTL        : 0
            No Route to Dest    : 2
            Dest Processing Err : 0

IP Interface: ve7
------------
  Helper-Address: None
  Packets:
            RX: 14
               BootRequest Packets : 0
               BootReply Packets   : 14
            TX: 14
               BootRequest Packets : 14
               BootReply Packets   : 0
  No-Relay: 0
  Drops:
            Invalid BOOTP Port  : 0
            Invalid IP/UDP Len  : 0
            Invalid DHCP Oper   : 0
            Exceeded DHCP Hops  : 0
            Invalid Dest IP     : 0
            Exceeded TTL        : 0
            No Route to Dest    : 0
            Dest Processing Err : 0
```

Table 8 describes the fields in the command output.

TABLE 8 show ip helper-address detail fields

Field	Description
IP Interface	AX interface.
Helper-Address	IP address configured on the AX interface as the DHCP helper address.

TABLE 8 show ip helper-address detail fields (Continued)

Field	Description
Packets	DHCP packet statistics: • RX – Total number of DHCP packets received on the interface. • BootRequest Packets – Number of DHCP boot request packets (Op = BOOTREQUEST) received on the interface. • BootReply Packets – Number of DHCP boot reply packets (Op = BOOTREPLY) received on the interface. • TX – Total number of DHCP packets sent on the interface. • BootRequest Packets – Number of DHCP boot request packets (Op = BOOTREQUEST) sent on the interface. • BootReply Packets – Number of DHCP boot reply packets (Op = BOOTREPLY) sent on the interface.
No-Relay	Number of packets that were examined for DHCP relay but were not relayed, and instead received regular Layer 2/3 processing. Generally, this counter increments in the following cases: • DHCP packets are received on an interface that does not have a helper address and the packets are not destined to the relay. • DHCP packets are received on an interface that does have a helper address, but the packets are unicast directly from the client to the server and do not need relay intervention.

TABLE 8 show ip helper-address detail fields (Continued)

Field	Description
Drops	Lists the following counters for packets dropped on the interface: • Invalid BOOTP Port – Number of packets dropped because they had UDP destination port 68 (BOOTPC). • Invalid IP/UDP Len – Number of packets dropped because the IP or UDP length of the packet was shorter than the minimum required length for DHCP headers. • Invalid DHCP Oper – Number of packets dropped because the Op field in the packet header did not contain BOOTREQUEST or BOOTREPLY. • Exceeded DHCP Hops – Number of packets dropped because the number in the Hops field was higher than 16. • Invalid Dest IP – Number of packets dropped because the destination was invalid for relay. • Exceeded TTL – Number of packets dropped because the TTL value was too low (less than or equal to 1). • No Route to Dest – Number of packets dropped because the relay agent (AX device) did not have a valid forwarding entry towards the destination. • Dest Processing Err – Number of packets dropped because the relay agent experienced an error in sending the packet towards the destination.

The following command clears the DHCP relay counters:

```
AX#clear ip helper-address statistics
```

NetFlow v9 and v10 (IPFIX)

The AX device can act as a NetFlow exporter. The NetFlow exporter (AX device) monitors traffic and sends the data to one or more NetFlow collectors, where the information can be stored and analyzed by a network administrator.

FIGURE 18 NetFlow Architecture with AX device as Exporter

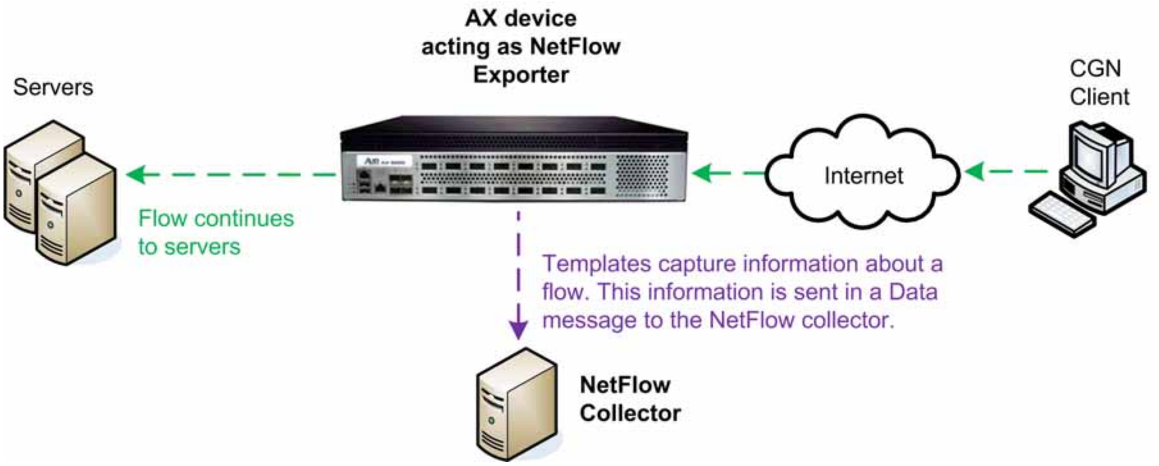
AX device
acting as NetFlow
Exporter
Servers
CGN
Client
Internet
Flow continues
to servers
Templates capture information about a
flow. This information is sent in a Data
message to the NetFlow collector.
NetFlow
Collector

Caution: **NetFlow is a heavy user of system resources and requires more memory for each session than is typically the case. When NetFlow is enabled, the session table capacity is reduced to just one-third (1/3) of its original amount.**

NetFlow Parameters

On the AX device, you can configure up to 64 NetFlow monitors.
A NetFlow monitor consists of the following protocol parameters, which can be used to configure the AX device to export data in the format of NetFlow v9 or NetFlow v10 (IPFIX). The default protocol is netflow v9.

- Export destination – External devices to export the collected data. You can specify the IP address of a single NetFlow collector, or configure a service group that comprises multiple collectors.
  - To achieve load balancing of NetFlow traffic among two or more collectors, they must be placed within the same service group.
  - You can give one collector a higher priority within the service group than the other collector. This will direct NetFlow traffic to the primary collector, and the lower priority collector will only be used as a backup.
  - If two or more NetFlow collectors are configured using only IP addresses and are not included in a service group, and if they are configured with the same NetFlow properties (record types), then NetFlow traffic will be duplicated to both places and the NetFlow traffic will not be load balanced.

Note: NetFlow information is sent from the AX device through a data port that is dynamically selected and is based upon information in the routing table.

- Record type – Types of data to export. NetFlow exporters use the following types of messages to send collected data to a collector server:
  - Templates – A NetFlow template defines the set of data to be collected, and the order in which that information will appear in the data messages.
  - Data – NetFlow data messages contain the collected data, such as flow information. Packets for data messages can contain data for more than one flow.

  Each NetFlow monitor can use one or more NetFlow templates. This release includes some predefined NetFlow templates. (See “Predefined NetFlow Templates” on page 146.)

- Monitoring filters – Specific type and subset of resources to monitor. You can specify monitoring of one ***or*** the other of the following resources:
  - Ethernet data ports – Specify the list of ports to monitor. Flow information for the monitored interfaces is sent to the NetFlow collector(s).
  - NAT pools – Specify the pool of NAT addresses to monitor. Flow information for the monitored IP addresses is sent to the NetFlow

collector(s). Currently, only CGN pools can be entered; standard (non-CGN pools) are not supported.

By default, no filters are in effect. Traffic on all interfaces and for all NAT pools is monitored.

- Flow timeout – This is the interval for sending flow records for long-lived sessions. (For short-lived sessions, any flow records are sent upon termination of the session.) For long-lived sessions, the flow timeout default value is 10 minutes. After this amount of time has elapsed, the AX will send any flow records to the NetFlow collector, even if the flow is still active. The flow timeout can be set to 0-1440 minutes. If this is set to 0, this essentially disables the flow timeout feature. Regardless of how long-lived a flow might be, the AX device waits until the flow has ended and the session is deleted before it sends any flow records for it.

Note: This parameter applies only to templates for flows. These are the templates listed in the “Templates for A10 Flow Records with NAT Addresses” section of Table 9 on page 146.

- Template transmission options – The AX device periodically resends the NetFlow templates to the collector(s). The following counters control when the templates are resent:
  - Number of data records sent – this is a running counter of the total number of data messages that have been sent to the NetFlow collector. After the specified number of data records are sent, the AX device resends the template that describes the data (as a way to refresh the template). The default is 1000 records. You can set the set template interval to 0-1000000 records.
  - Number of seconds since the last time the template was sent – After the specified number of seconds has passed, the AX device resends the template to perform a refresh of the template on the collector. The default is 1800 seconds. You can set it to 0-86400 seconds. After the template is resent, this counter is set back to 0 seconds.

  You can disable either counter by setting it to 0.

- Management interface – Uses the IP of the AX management interface, instead of the IP of the data interfaces when sending traffic to the NetFlow collectors. By default, the AX device sends NetFlow traffic out the data interface, but when the Management Interface option is enabled, the NetFlow information will still be sent out a data interface that is dynamically (and automatically) selected based upon the routing table, but the source IP of the packets will be the IP of the management port.

- Monitor state – Enabled or disabled. By default, a NetFlow monitor is enabled.

NetFlow Versions Supported

The current release supports NetFlow version 9 and NetFlow version 10 (IPFIX). This version is described in RFC 3954, Cisco Systems NetFlow Services Export Version 9.

Predefined NetFlow Templates

The AX device includes some pre-defined NetFlow templates. Table 9 lists the pre-configured NetFlow templates available in the current release.

TABLE 9 AX NetFlow Template Types

Template Name	Data Fields
Templates for A10 Flow Records with NAT Addresses	
These template types are bidirectional. One session results in one flow record.	
nat44	**Key fields** • IP Protocol • IPv4 Source Address • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) **Non-key fields** • IPv4 NAT source address • IPv4 NAT dest address • NAT source port • NAT dest port • Interface Input • Interface Output • Fwd Bytes • Fwd Packets • Rev Bytes • Rev Packets • Start time (msec) • Duration (msec)

TABLE 9 AX NetFlow Template Types (Continued)

Template Name	Data Fields
nat64	**Key fields** • IP Protocol • IPv6 Source Address • IPv4 Destination Address • IPv6 Destination Address (hairpin) • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) **Non-key fields** • IPv4 NAT source address • IPv6 NAT source address • IPv6 NAT dest address • IPv4 NAT dest address • NAT source port • NAT dest port • Interface Input • Interface Output • Fwd Bytes • Fwd Packets • Rev Bytes • Rev Packets • Start time (msec) • Duration (msec)

TABLE 9 AX NetFlow Template Types (Continued)

Template Name	Data Fields
dslite	**Key fields** • IP Protocol • IPv6 Source Address • IPv4 Source Address • IPv6 Destination Address • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) **Non-key fields** • IPv6 NAT source address • IPv4 NAT source address • IPv6 NAT dest address • IPv4 NAT dest address • NAT source port • NAT dest port • Interface Input • Interface Output • Fwd Bytes • Fwd Packets • Rev Bytes • Rev Packets • Start time (msec) • Duration (msec)
Templates for NAT Event Records	
sesn-event-nat44	• IP Protocol • IPv4 Source Address (inside client address) • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) • IPv4 NAT source address • IPv4 NAT dest address • NAT source port • NAT dest port • timestamp (msec) • sesnEvent (Create, Delete)

TABLE 9 AX NetFlow Template Types (Continued)

Template Name	Data Fields
sesn-event-nat64	• IP Protocol • IPv6 Source Address • IPv4 Source Address (inside client address) • IPv6 Destination Address • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) • IPv6 NAT source address • IPv4 NAT source address • IPv6 NAT dest address • IPv4 NAT dest address • NAT source port • NAT dest port • sesnEvent (Create, Delete)
sesn-event-dslite	• timeStamp • IP Protocol • IPv6 Source Address • IPv4 Source Address (inside client address) • IPv6 Destination Address • IPv4 Destination Address • Source Port • Destination Port • Flow Direction (inbound, outbound, or hairpin) • IPv6 NAT source address • IPv4 NAT source address • IPv6 NAT dest address • IPv4 NAT dest address • NAT source port • NAT dest port • timestamp (msec) • sesnEvent (Create, Delete)
port-mapping-nat44	• IP Protocol • IPv4 Source Address • Source Port • IPv4 NAT source address • NAT source port • timestamp (msec) • natEvent (Create, Delete)

TABLE 9 AX NetFlow Template Types (Continued)

Template Name	Data Fields
port-mapping-nat64	• IP Protocol • IPv6 Source Address • IPv4 Source Address • Source Port • IPv4 NAT source address • NAT source port • timestamp (msec) • natEvent (Create, Delete)
port-mapping-dslite	• IP Protocol • IPv6 Source Address • IPv4 Source Address • Source Port • IPv4 NAT source address • NAT source port • timestamp (msec) • natEvent (Create, Delete)

Configuring NetFlow

1. If using multiple NetFlow collectors, create a server configuration for each collector, and add the server configurations to a service group.

   Make sure to disable the Layer 4 health check on the UDP port.

2. Configure a NetFlow monitor. Within the monitor, specify the following:

   - The destination, which can be one of the following:
     - Host address, if using a single NetFlow collector
     - Service-group name, if using multiple NetFlow collectors
   - The record types to export. (Specify them by NetFlow template type.)
   - (Optional) The Ethernet interfaces from which to collect NetFlow information. By default, information is collected for all interfaces.
   - (Optional) Adjust the flow timeout.
   - (Optional) Adjust the template resend counters.

Note: If you plan to use only a single NetFlow collector, you do not need to perform step 1. You can specify the NetFlow collector’s IP address when configuring the NetFlow monitor (in step 2).

USING THE GUI

1. Select Config Mode > Service > NetFlow Monitor.

2. Click the Add button. The NetFlow Monitor Create page appears.

FIGURE 19 Config Mode > Service > NetFlow Monitor

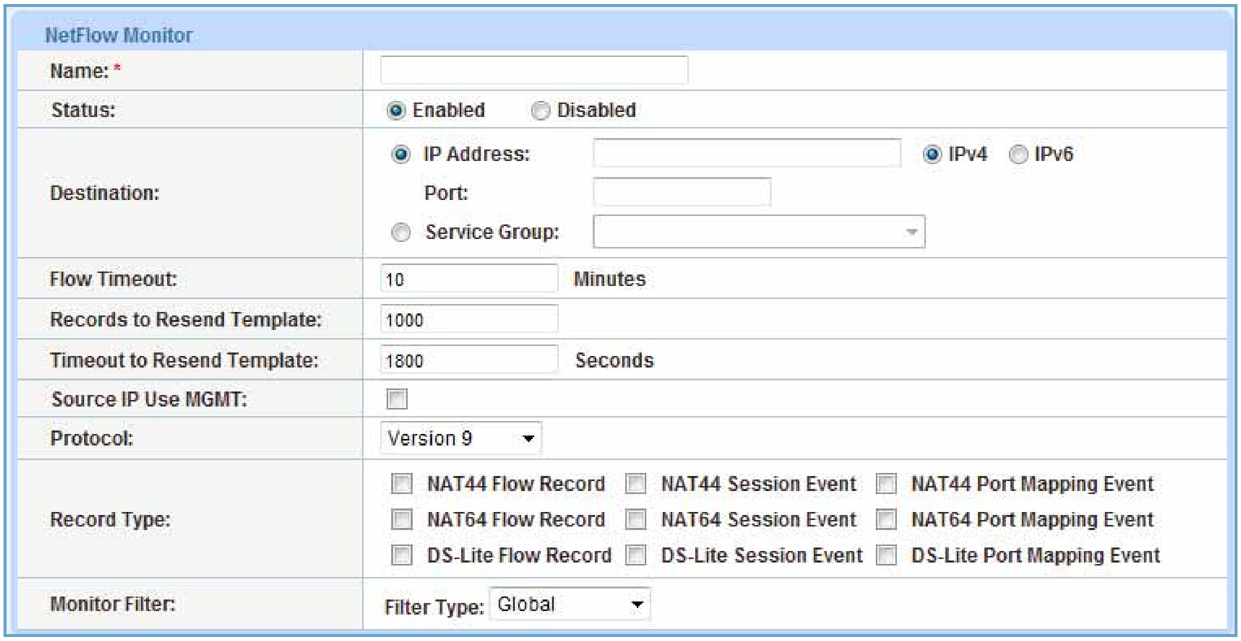

NetFlow Monitor

Name: *	
Status:	● Enabled ○ Disabled
Destination:	● IP Address: ● IPv4 ○ IPv6 Port: ○ Service Group:
Flow Timeout:	10 Minutes
Records to Resend Template:	1000
Timeout to Resend Template:	1800 Seconds
Source IP Use MGMT:	□
Protocol:	Version 9
Record Type:	□ NAT44 Flow Record □ NAT44 Session Event □ NAT44 Port Mapping Event □ NAT64 Flow Record □ NAT64 Session Event □ NAT64 Port Mapping Event □ DS-Lite Flow Record □ DS-Lite Session Event □ DS-Lite Port Mapping Event
Monitor Filter:	Filter Type: Global

3. Enter a name for the NetFlow monitor in the Name field.

4. Configure the following fields and options:
   - The Destination, which can be one of the following:
     - Host address, if using a single NetFlow collector
     - Service-group name, if using multiple NetFlow collectors
   - (Optional) Adjust the flow timeout. Default is 20 minutes.
   - (Optional) Adjust the Records to Resend Template. The default is 1000 records, after which the template will be re-sent to the collector and the counter will be reset to 0.
   - (Optional) Adjust the Timeout to Resend Template. The default is 1800 seconds, after which the template is re-sent to the collector.
   - (Optional) Select the Source IP Use MGMT checkbox to use the IP address of the AX management port as the source IP of the NetFlow packets, even when packets are sent out the data interface.

- Select the Protocol version. The default is NetFlow Version 9, but you can also select NetFlow Version 10.
- The record types to export. (Specify them by NetFlow template type.)

5. When finished, click OK to save your changes.

Using the CLI

Configure the NetFlow Servers and Service Group (if using multiple NetFlow Collectors)

Note: If you plan to use only a single collector, skip this section and go to "Configure the NetFlow Monitors" on page 153.

1. Use the following commands to create the real server configurations.

   ```
   [no] slb server server-name ipaddr
   ```

   This command creates the server and accesses the configuration level for it.

   Use the following command to add the protocol port on which the collector listens for NetFlow traffic.

   ```
   [no] port portnum {tcp | udp}
   ```

   This command adds the port to the server configuration, and accesses the configuration level for the port. At this level, use the following command to disable the Layer 4 health check for the port:

   ```
   no health check
   ```

Note: Layer 4 health checks are not supported by NetFlow, so if a Layer 4 health check were sent to the collector, it would be rejected and the AX device, (failing to receive the expected response from the collector), would mark the collector as "down".

2. Use the following commands to create a service group and add the NetFlow servers to it.

   ```
   [no] slb service-group group-name {tcp | udp}
   ```

   For **tcp** | **udp**, specify the same protocol specified for the ports on the servers.

This command creates the service group and accesses the configuration level for it. At this level, use the following command to add the servers and NetFlow ports to the service group:

```
[no] member server-name:portnum
```

Configure the NetFlow Monitors

The following command configures a NetFlow monitor and changes the CLI to the configuration level for the NetFlow monitor.

```
[no] netflow monitor name
```

At the NetFlow configuration level, the following command specifies the type of NetFlow records to export.

```
[no] record netflow-template-type
[both | creation | deletion]
```

The *netflow-template-type* refers to the NetFlow template that defines the NetFlow records to export, and it includes the following template types:

- nat44
- nat64
- dslite
- sesn-event-nat44
- sesn-event-nat64
- sesn-event-dslite
- port-mapping-nat44
- port-mapping-nat64
- port-mapping-dslite

See Table 9 on page 146 for more details about netflow template types.

The option to specify **both**, **creation**, and **deletion** allows you to determine which types of events will be exported:

- **both** – Export both creation and deletion events (default)
- **creation** – Export only creation events
- **deletion** – Export only deletion events

Note: The option to specify **both**, **creation**, and **deletion** is only available for session event and port mapping event templates. It is not available for NAT44, NAT64, or DS-Lite.

The command below configures filters to specify the type of traffic to be monitored:

```
[no] monitor
{
ethernet portnum [...] |
global |
nat-pool pool-name
}
```

The traffic monitors include the following:

- The **ethernet** option monitors specific Ethernet data ports.
- The **global** option monitors all interfaces and NAT pools. This is the default configuration.
- The **nat-pool** option monitors the NAT addresses in the specified pool.

The following command specifies the NetFlow collectors upon which to export the data.

```
[no] destination
{ipaddr portnum | service-group group-name}
```

- If you plan to use only one NetFlow collector, you can specify its IP address and NetFlow port number with this command.
- If you plan to use multiple NetFlow collectors, use the **service-group** *group-name* option to specify the name of the group that contains the NetFlow collectors.

```
[no] source-ip-use-mgmt
```

This command does not change the AX port from which NetFlow traffic is exported, but configures the AX device to use the management port's IP as the source IP for the NetFlow packets.

```
[no] flow-timeout minutes
```

This command specifies the interval for sending data records for long-lived flows. (See "NetFlow Parameters" on page 144.)

[no] resend-template records *num*

[no] resend-template timeout *seconds*

These commands specify the counters by which the AX device resends the templates to the collectors. (See "NetFlow Parameters" on page 144.)

enable | disable

Enables or disables the NetFlow monitor. By default, NetFlow monitors are enabled.

Displaying NetFlow Information

To display NetFlow information, use the following command:

show netflow monitor [*monitor-name*]

CLI Examples

Example 1

The following example shows one single monitor that is used to collect all nat44/nat64/dslite records and export them to the host at IP 10.1.1.10.

```
AX(config)#netflow monitor 1
AX(config-netflow-monitor)#record nat44
AX(config-netflow-monitor)#record nat64
AX(config-netflow-monitor)#record dslite
AX(config-netflow-monitor)#destination 10.1.1.10
```

Example 2

The following example shows two monitors configured to collect records from different interfaces and export them to different hosts. In this example, since the record types are the same, identical NetFlow traffic will be sent to both places.

```
AX(config)#netflow monitor 1
AX(config-netflow-monitor)#monitor ethernet 1
AX(config-netflow-monitor)#record nat44
AX(config-netflow-monitor)#record nat64
AX(config-netflow-monitor)#record dslite
AX(config-netflow-monitor)#destination 10.1.1.10

AX(config)#netflow monitor 2
AX(config-netflow-monitor)#monitor ethernet 3
AX(config-netflow-monitor)#record nat44
AX(config-netflow-monitor)#record nat64
AX(config-netflow-monitor)#record dslite
AX(config-netflow-monitor)#destination 10.1.1.11
```

Example 3

The following example shows load balancing configured between multiple netflow collectors. Layer 4 health monitors are disabled in order to prevent the collector from being marked down. The two NetFlow collectors are included in the same service group, and no distinction is made to give one member a higher priority than the other. Therefore, traffic will be load balanced with no priority given to either collector.

```
AX(config)#slb server server1 10.1.1.101
AX(config-real server)#port 9996 udp
AX(config-real server-node port)#no health check
AX(config)#slb server server2 10.1.1.102
AX(config-real server)#port 9996 udp
AX(config-real server-node port)#no health check

AX(config)#slb service-group netflow-server-farm udp
AX(config-slb svc group)#member server1:9996
AX(config-slb svc group)#member server2:9996

AX(config)#netflow monitor 1
AX(config-netflow-monitor)#record nat44
AX(config-netflow-monitor)#destination service-group netflow-server-farm
```

sFlow

The AX device can act as an sFlow agent by sampling random packets and sending statistics in an sFlow datagram to an external sFlow collector for analysis.

sFlow Sampling Types

sFlow supports two types of sampling. One type of sampling uses a time-based approach to retrieve statistics for a specific interface, while the other approach samples information from the packet header of every *Nth* packet.

Counter Polling Interval

This is a counter sampling method that is based on time. Statistics for an interface are gathered periodically and sent to the sFlow collector. You can specify the time interval (frequency) with which the counter interfaces statistics are gathered and sent. This global configuration will apply to all interfaces where sFlow is enabled unless a more granular value is configured at the interface level. You can enter a value ranging from 1–200 seconds. By default, this interval is set to 20 seconds.

Once the AX device has sampled statistics from a target interface, the information is collected and sent in an sFlow datagram to one or more sFlow collectors. The sFlow datagrams are listed in the Received and Transmitted counter fields in **show interface** CLI output, or on the Monitor Mode > Network > Interface page of the GUI.

Packet Sampling Rate

This is a sampling method that is based on the number of incoming packets. This sampling rate value essentially means that one packet is sampled out of every *N* packets. When expressed as a ratio, the packet sampling rate looks like 1/*N*. You can enter a value for *N* (the denominator) ranging from 10–1000000 packets. By default, *N* is equal to 1000, meaning that one packet is sampled out of every 1000 packets arriving at that interface. This global configuration will apply to all interfaces where sFlow data is collected, unless a more granular value has been configured at the interface level.

Unlike the other time-based sampling method, which gathers counter statistics for an interface, this packet-volume sampling approach gathers data about specific packets arriving at an interface. Information is extracted from the first 128 bytes in the header of the sampled packet, beginning with the MAC header. Once the AX device has sampled packets from a specified target interface, the information is collected and sent in an sFlow datagram to one or more sFlow collectors.

Details:

- You can enable one or both sampling types on a single Ethernet data port – the sampling types are not mutually exclusive.
- The sFlow datagram includes information about the incoming interface but not the outgoing interface where sampling occurred.
- sFlow data can be exported to up to 4 sFlow collectors. This offers the benefit of redundancy, as well as the ability to send sFlow datagrams to different destinations.
- By default, the sFlow datagrams use the management IP of the AX device as the source address, but you can modify the exported sFlow datagrams to the source address of your choice.

Limitations

- sFlow data collection is supported only for individual Ethernet data ports and cannot be performed on VEs, trunk interfaces, loopback interfaces, or on the management interface of the AX device.
- Sampling of CGN packets is performed only on pre-translation packets. The samples do not carry post-translation header information.
- None of the following are supported:
  - Host resource sampling
  - Application behavior sampling
  - Configuration of sFlow agent behavior using SNMP

Information Included in sFlow Datagrams

- **Discarded packets** – Information about the discarded packets is included in the sFlow datagrams. For a list of Destination Unreachable codes associated with discarded packets, see section “Input/ Output Port Information” in the following RFC: http://sflow.org/sflow_version_5.txt
- **Extended NAT Layer 4 port data** – NAT port translation (mapping) information is included in the exported sFlow datagrams. For example, if client (A) sends a packet to the AX device (B) and then the packet is sent to the server (C), and if the AX device is performing source NAT to replace the client’s IP address and port with its own, then this port translation and IP address mapping information will be included in the sFlow datagram that is sent to the sFlow collector. The source NAT mappings will appear in the “Extended NAT data” section of the sFlow datagram.

- **Export CPU and Memory information** – CPU and memory information are included in the "Processor information" section of the exported sFlow datagram, as shown in Figure 20 below.

FIGURE 20   *Processor information section of sFlow datagram*

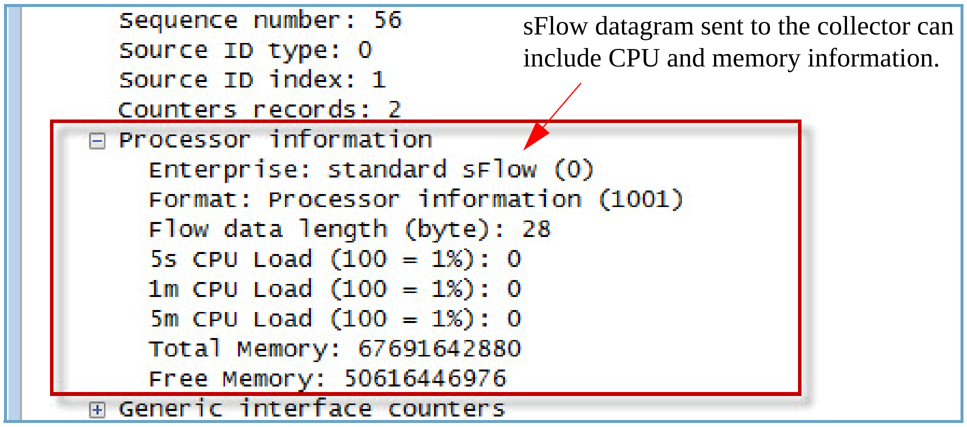

```
Sequence number: 56
Source ID type: 0
Source ID index: 1
Counters records: 2
⊟ Processor information
    Enterprise: standard sFlow (0)
    Format: Processor information (1001)
    Flow data length (byte): 28
    5s CPU Load (100 = 1%): 0
    1m CPU Load (100 = 1%): 0
    5m CPU Load (100 = 1%): 0
    Total Memory: 67691642880
    Free Memory: 50616446976
⊞ Generic interface counters
```

sFlow datagram sent to the collector can include CPU and memory information.

- **Export LSN NAT pool usage information** – The "application information" section of the exported sFlow datagram includes LSN NAT pool usage information, as shown in Figure 21.

FIGURE 21   *Application information section in sFlow datagram*

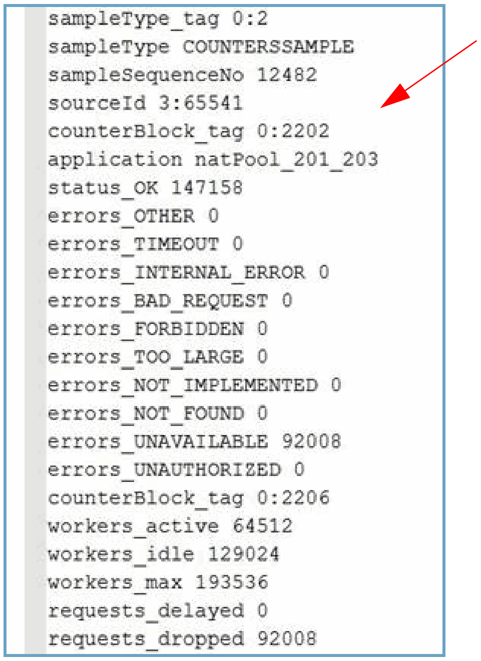

```
sampleType_tag 0:2
sampleType COUNTERSSAMPLE
sampleSequenceNo 12482
sourceId 3:65541
counterBlock_tag 0:2202
application natPool_201_203
status_OK 147158
errors_OTHER 0
errors_TIMEOUT 0
errors_INTERNAL_ERROR 0
errors_BAD_REQUEST 0
errors_FORBIDDEN 0
errors_TOO_LARGE 0
errors_NOT_IMPLEMENTED 0
errors_NOT_FOUND 0
errors_UNAVAILABLE 92008
errors_UNAUTHORIZED 0
counterBlock_tag 0:2206
workers_active 64512
workers_idle 129024
workers_max 193536
requests_delayed 0
requests_dropped 92008
```

Application information section of sFlow datagram includes LSN NAT pool usage information.

Configuration

Below are the high-level steps involved in configuring the sFlow data collection feature on an AX device:

1. Specify the sFlow collector where data will be exported.
2. (Optional) Enable use of the management interface’s IP as the source address for outbound sFlow packets. This may be beneficial for filtering at the collector or to maintain consistency in the source address of the sFlow packets.
3. Specify the individual Ethernet data interfaces that will be sampled.
4. (Optional) Change the default data sampling rate or polling interval.

USING THE GUI

1. Select Config Mode > Service > sFlow. The sFlow create page appears.

FIGURE 22 Config Mode > Service > sFlow create page

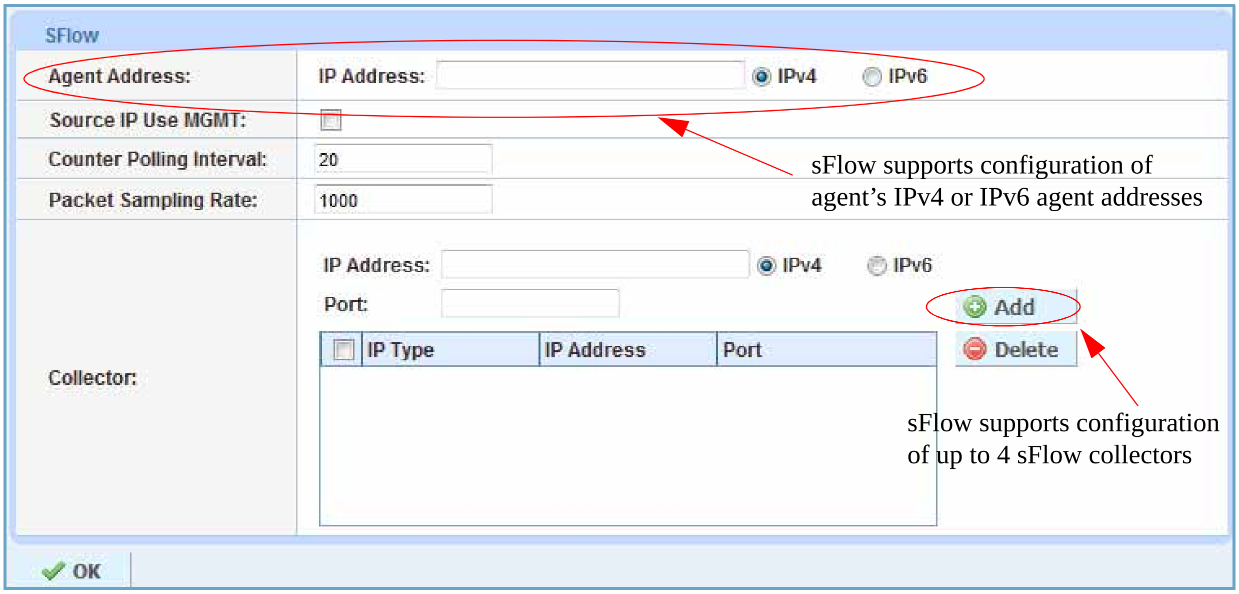

SFlow

Agent Address:	IP Address: ● IPv4 ○ IPv6
Source IP Use MGMT:	□
Counter Polling Interval:	20
Packet Sampling Rate:	1000
Collector:	IP Address: ● IPv4 ○ IPv6 Port:

□	IP Type	IP Address	Port

sFlow supports configuration of agent’s IPv4 or IPv6 agent addresses

sFlow supports configuration of up to 4 sFlow collectors

2. (Optional) Enter an IP address for the sFlow agent. By default, the management IP of the AX device is used, but you may enter a different address if desired.

Note: This information will appear in the Layer 4 information section of the sFlow datagram. Although the information is “textual” and is not used for routing decisions, it may be helpful in identifying which sFlow agent a particular packet came from, particularly in complex networks that have more than one sFlow agent.

3. (Optional) Select the Source IP Use MGMT checkbox if you wish to use the AX device’s management IP as the source address for exported sFlow datagrams. This changes the source address on the sFlow datagrams but has no effect on which interface the AX device selects for exporting sFlow datagrams.

4. (Optional) In the Counter Polling Interval field, specify the time interval at which the counter of interface statistics will be sampled. (See “Counter Polling Interval” on page 157 for more information.)

5. (Optional) In the Packet Sampling Rate field, alter the default value if desired. Smaller numbers increase the sampling frequency, and larger numbers decrease the sampling frequency. (See “Packet Sampling Rate” on page 157 for more information.)

6. In the Collector area of the page, enter an IPv4 or IPv6 address in the IP address field, and then select the IPv4 or IPv6 radio button.

7. Enter a value in the Port field. This is the port on the collector where sFlow traffic will be sent. By default, traffic is sent to UDP port 6343.

8. Click Add to add the sFlow collector’s information, and then click OK to save your changes.

9. Next, select Config Mode > Network > Interface > LAN, and then click on the interface upon which you would like to enable sampling.
   A page similar to the one shown below appears.

FIGURE 23 Config Mode > Network > Interface > LAN

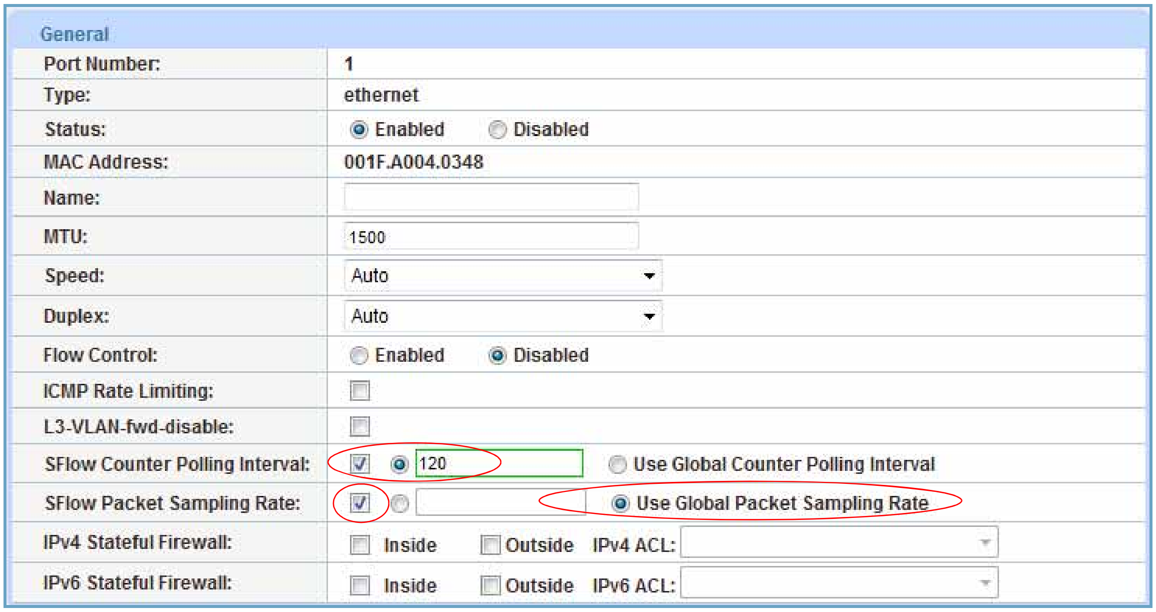

General	
Port Number:	1
Type:	ethernet
Status:	● Enabled ○ Disabled
MAC Address:	001F.A004.0348
Name:	
MTU:	1500
Speed:	Auto
Duplex:	Auto
Flow Control:	○ Enabled ● Disabled
ICMP Rate Limiting:	□
L3-VLAN-fwd-disable:	□
SFlow Counter Polling Interval:	☑ ● 120 ○ Use Global Counter Polling Interval
SFlow Packet Sampling Rate:	☑ ○ ● Use Global Packet Sampling Rate
IPv4 Stateful Firewall:	□ Inside □ Outside IPv4 ACL:
IPv6 Stateful Firewall:	□ Inside □ Outside IPv6 ACL:

10. Select the **sFlow Counter Polling Interval** checkbox to use a time-based approach sampling. By default, the globally configured polling interval (configured in step 4) is used. To configure a value that is specific to this interface, enter a new value ranging from 1–200 seconds in the field and select the associated radio button.

11. Select the **sFlow Packet Sampling Rate** checkbox to use a traffic volume-based approach to sampling. By default, the globally configured packet sampling rate (configured in step 5) is used. You can enter a new sampling rate value of one packet per *N* packets (where *N* ranges from 10–1000000).

12. When finished, scroll down and click OK to save your changes.

Use the following commands to configure sFlow support on the AX device.

Configuring an sFlow Agent

To configure an sFlow agent address, use the following command at the global configuration level of the CLI:

```
[no] sflow agent address ipaddr
```

The *ipaddr* value can be any valid IPv4 or IPv6 address. By default, sFlow datagrams use the management IP of the AX device as the source address, but you can specify a different IP address, if desired. The information will appear in the Layer 4 information section of the sFlow datagram, and it is not used to make routing decisions.

Configuring an sFlow Collector

To configure an sFlow collector, use the following command at the global configuration level of the CLI:

```
[no] sflow collector ipaddr [portnum]
```

The *ipaddr* value can be any valid IPv4 or IPv6 address, and you can configure up to a total of 4 collectors. For the *portnum* parameter, you can enter any valid port number between 1-65535. The default protocol port number for sFlow is 6343.

Configuring the Global Counter Polling Interval

By default, the **counter-polling-interval** is set to 20 seconds. If desired, you can modify this value by using the following command at the global configuration level of the CLI:

```
[no] sflow counter-polling-interval seconds
```

The *seconds* option specifies the interval at which an incoming packet will be sampled from the interface. The interval value can range from 1-200 seconds.

Configuring the Global Packet Sampling Rate

By default, the **packet-sampling-rate** samples one packet out of every 1000 incoming packets. If desired, you can modify this value by using the following command at the global configuration level of the CLI:

```
[no] sflow packet-sampling-rate num
```

The *num* option specifies the value of *N*, where *N* is the value of the denominator in the ratio 1/*N*. This ratio says that one packet will be sampled out of *N* packets (denominator), and the value of *N* can range from 10-1000000. The default is 1000, meaning one packet out of every 1000 packets will be sampled.

Enabling sFlow On a Specific Interface

You can enable sFlow on an interface (while relying on the globally configured sFlow sampling values) using one of the following commands.

To enable time-based sFlow sampling, use the following command at the global configuration level of the CLI:

```
[no] sflow polling ethernet port-num
```

To enable packet volume-based sFlow sampling, use the following command at the global configuration level of the CLI:

```
[no] sflow sampling ethernet port-num
```

Note: Despite the differences in syntax (“polling” versus “sampling”), both of the commands above enable sampling on an interfaces. The time-based sampling approach (with “polling” syntax), samples a counter of interface statistics at a regular time interval, while the other command (with “sampling” syntax), samples the headers from packets as they arrive at an interface.

With either command, the `port-num` option specifies the target Ethernet interface.

Changing the Default Counter Polling Interval for an Interface

If sFlow is enabled on an interface, but the counter polling interval has not been specified, then the globally-configured counter polling interval will be applied. You can configure a different counter polling interval for this specific interface by using the following command at the global configuration level of the CLI:

```
[no] sflow polling ethernet port-num
[to port-num]
interval seconds
```

The *port-num* option specifies the target ethernet interface. The *seconds* option specifies the frequency with which the counter data for an interface will be sampled, and it can range from 1-200 seconds.

Changing the Default Packet Sampling Rate for an Interface

If sFlow is enabled on an interface but the packet sampling rate has not been specified, then the globally-configured packet sampling rate is applied. You can configure a different packet sampling rate value for a specific interface by using the following command at the global configuration level of the CLI:

```
[no] sflow sampling ethernet port-num
[to port-num]
rate sampling-num
```

The *port-num* option specifies the target ethernet interface. The *sample-num* option specifies the denominator in the ratio for the packet sampling, 10-1000000. The globally configured packet sampling rate is one packet out of every 1000.

Configuring Management IP as Source Address

To enable use of the management interface's IP as the source IP for the sFlow datagrams sent to collectors, use the following command:

```
[no] sflow source-ip-use-mgmt
```

Including LSN NAT Pool and CPU Usage in Exported Datagrams

By default, the LSN NAT pool usage and CPU usage information are included in exported sFlow datagrams. However, you can use the following commands at the global config level to enable or disable the appearance of this information:

```
[no] sflow polling lsn-pool-usage
```

```
[no] sflow polling cpu-usage
```

Displaying sFlow Information

To display the sFlow configuration, use the following command:

```
show sflow configuration [ethernet portnum]
```

To display sFlow statistics, use the following command:

```
show sflow statistics [ethernet portnum]
```

To clear sFlow statistics, use the following command:

```
clear sflow statistics
```

CLI Examples

The following commands specify the sFlow collector, and enable use of the management interface's IP as the source IP for the data samples sent to the sFlow collector:

```
AX(config)#sflow collector 192.168.100.3
AX(config)#sflow source-ip-use-mgmt
```

The following command enables counter polling for several Ethernet data interfaces, and uses the globally configured sampling rate by default:

```
AX(config)#sflow polling ethernet 1 2 5 6 8 10 11
```

The following command enables counter polling for several Ethernet data interfaces, but uses a custom sampling rate of 24 seconds:

```
AX(config)#sflow polling ethernet 3 4 7 9 12 interval 24
```

The following command enables packet flow sampling for two Ethernet interfaces, with a sampling rate of 500, meaning that one packet out of every 500 will be sampled:

```
AX(config)#sflow sampling ethernet 1 2 rate 500
```

The following command enables packet sampling for a range of Ethernet interfaces, with a packet sampling rate of 12, meaning that one packet out of every 12 will be sampled:

```
AX(config)#sflow sampling ethernet 3 to 5 rate 12
```

The following command displays the sFlow configuration:

```
AX(config)#show sflow configuration
sFlow collector(s)
   address                        0: 192.168.100.3 /6343
   source-ip-use-mgmt:            true

sFlow agent address:              not set, use management ip address

sFlow default parameter
   counter polling interval:      20
   packet sampling rate:          1000

sflow polling ethernet 1 2 5 6 8 10 11
sflow polling ethernet 3 4 7 9 12 interval 24

sflow sampling ethernet 1 2 rate 500
sflow sampling ethernet 3 to 5 rate 12
```

The following command displays sFlow data collection statistics:

```
AX(config)#show sflow statistics
Interface      Packet Sample Records      Counter Sample Records
-------------------------------------------------------------------
1              3461                       81
2              20801                      81
3              0                          81
4              0                          81
5              0                          81
6              0                          81
7              0                          81
8              0                          81
9              0                          81
10             0                          81
11             0                          81
12             0                          81
-------------------------------------------------------------------
sflow total statistics
   Packet sample records:         24262
   Counter sample records:        972
   Sflow packets sent:            16257
```

The following command displays sFlow data collection statistics specifically for Ethernet data port 4:

```
AX(config)#show sflow statistics ethernet 4
   sflow statistics on ethernet 4
       packet sample records:        0
       counter sample records:       82
```

High Availability

This chapter describes High Availability (HA) and how to configure it.

Overview

High Availability (HA) is an AX feature that provides AX-level redundancy to ensure continuity of service to clients. In HA configurations, AX devices are deployed in pairs. If one AX device in the HA pair becomes unavailable, the other AX device takes over.

Note: Models AX 5200 and AX 5200-11 can be used together as an HA pair. For all other models, AX devices must be the same model and must be running the same software version.

Note: For complete syntax information on all configurable options, see the *AX Series CLI Reference*.

Figure 24 shows an example of an HA deployment for Dual-Stack Lite (DS-Lite).

FIGURE 24 HA deployment for DS-Lite

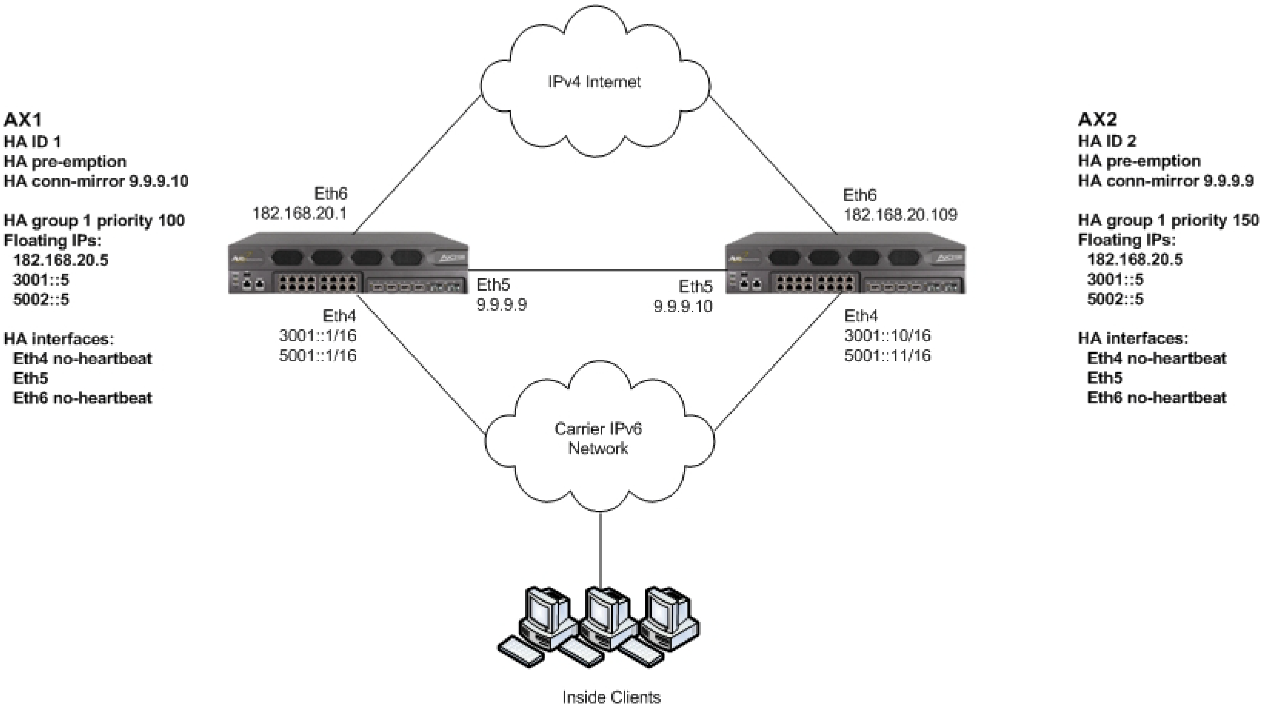

Each AX device in this HA pair has the following HA configuration elements:

- HA ID – The HA ID of AX1 is 1 and the HA ID of AX2 is 2. Each AX device in an HA deployment must have a unique HA ID. The ID must be different on each AX device. The ID can be used as a tie breaker to select the Active AX device. (See "How the Active AX Device Is Selected" on page 183.)
- Pre-emption – Configuration option that allows failover based on configuration changes.
- Connection mirroring – Synchronizes session information from the Active device to the Standby device, so that the Standby device can continue service for the sessions.

- HA group – Set of IP resources that are backed up by HA. You can assign the following types of resources to an HA group:
    - Network Address Translation (NAT) resources, such as IP address pools
    - Floating IP addresses (described below)
    - IPv6 router advertisements

  Each HA group must be configured with a priority. The priority can be used as a tie breaker to select the Active AX device for a VIP.

  Each HA group has a shared MAC address, 021f.a0000.00*xx*. The *xx* portion of the address is unique to the HA group. The shared MAC address is used for all IP addresses for which HA is provided (source NAT addresses, floating IP addresses, and so on).

- Floating IP addresses – IP addresses that reside on the Active AX device, and move to the Standby device if failover occurs.

- HA interfaces – AX interfaces upon which an AX device's HA status depends, or that are used by the pair of AX devices for session and configuration synchronization.

The commands used to configure this example are shown in "HA with DS-Lite" on page 188.

Note: A floating IP address can not be the same as an address that already belongs to a device. For example, the IP address of an AX interface can not be a floating IP address.

HA Messages

The AX devices in an HA pair communicate their HA status with the following types of messages:

- HA heartbeat messages

- Gratuitous ARP requests and replies

HA Heartbeat Messages

Each of the AX devices regularly sends HA heartbeat messages out its HA primary interfaces. The Standby AX device listens for the heartbeat messages. If the Standby AX device stops receiving heartbeat messages from the Active AX device, one of the following things occurs:

- If one or more HA interfaces are designated as redundant (backup) interfaces, the Standby AX device starts sending heartbeat messages to the Active AX device on the redundant interfaces.

  - If the Active AX device responds, the device remains Active.
  - If the Active AX device does not respond, failover occurs instead. (See below.)
- If no HA interfaces are designated as redundant interfaces, the Standby AX device transitions to Active.

By default, heartbeat messages are sent every 200 milliseconds. If the Standby AX device does not receive a heartbeat message for 1 second (5 times the heartbeat interval), the Standby AX device transitions to Active. The heartbeat interval and retry count are configurable.

Gratuitous ARPs (IPv4) and ICMPv6 Neighbor Advertisement (IPv6)

For IPv4, when an AX device transitions from Standby to Active, the newly Active AX device sends gratuitous ARP requests and replies (ARPs) for the IPv4 address under HA control.

Similarly, for IPv6, when an AX device transitions from Standby to Active, the newly Active AX device sends ICMPv6 neighbor advertisement for the IPv6 address under HA control.

Gratuitous ARPs or ICMPv6 neighbor advertisements are sent for the following types of addresses:

- Virtual server IP addresses, for the VIPs that are assigned to an HA group.
- Floating IP address, if configured
- NAT pool IP addresses, for NAT pools assigned to an HA group

Devices that receive the ARPs or ICMPv6 neighbor advertisements learn that the MAC address for the AX HA pair has moved, and update their forwarding tables accordingly.

The Active AX device sends the gratuitous ARPs or ICMPv6 neighbor advertisements immediately upon becoming the Active AX device. To make sure ARPs or ICMPv6 neighbor advertisements are being received by the target addresses, the AX device re-sends them 4 additional times, at 500-millisecond intervals.

After this, the AX device sends gratuitous ARPs or ICMPv6 neighbor advertisements every 30 seconds to keep its IP information current.

The retry count is configurable.

HA Interfaces

When configuring HA, you specify each of the AX interfaces to use as HA interfaces. An HA interface is an interface that is connected to an upstream router, a real server, or the other AX device in the HA pair.

Notes

- HA heartbeat messages can be sent only on HA interfaces.
- When you configure an HA interface that is a tagged member of one or more VLANs, you must specify the VLAN on which to send the heartbeat messages.
- The maximum number of HA interfaces you can configure is the same as the number of Ethernet data ports on the AX device.
- If the heartbeat messages from one AX device to the other will pass though a Layer 2 switch, the switch must be able to pass UDP IP multicast packets.

Note: If a tracked interface is a member of a trunk, only the lead port in the trunk is shown in the tracking configuration and in statistics. For example, if a trunk contains ports 1-3 and you configure tracking of port 3, the configuration will show that tracking is enabled on port 1. Likewise, tracking statistics will show port 1, not port 3. Similarly, if port 1 goes down but port 3 is still up, statistics still will show that port 1 is up since it is the lead port for the trunk.

HA Status for AX Interfaces

Changes to the state of an HA interface can trigger a failover. By default, the HA state of an interface can be Up or Down. Optionally, you can specify the HA interface type as one of the following:

- Server interface – A real server can be reached through the interface.
- Router interface – An upstream router (and ultimately, clients) can be reached through the interface.
- Both – Both a server and upstream router can be reached through the interface.

If you specify the HA interface type, the HA status of the AX device is based on the status of the AX link with the real server and/or upstream router. The HA status can be one of the following:

- Up – All configured HA router and server interfaces are up.

- Partially Up – Some HA router or server interfaces are down but at least one server link and one router link are up.
- Down – All router interfaces, or all server interfaces, or both are down. The status also is Down if both router interfaces and server interfaces are not configured and an HA interface goes down.

If both types of interfaces (router interfaces and server interfaces) are configured, the HA interfaces for which a type has not been configured are not included in the HA interface status determination.

During selection of the active AX, the AX with the highest state becomes the active AX and all HA interfaces on that AX become active. For example, if one AX is UP and the other AX is only Partially Up, the AX that is UP becomes the active AX.

If each AX has the same state, the active AX is selected as follows:

- If HA pre-emption is disabled (the default), the first AX to come up is the active AX.
- If HA pre-emption is enabled, the AX with the higher HA group priority becomes active for that group. If the group priorities on the two AX devices are also the same, the AX that has the lowest HA ID (1 or 2) becomes active.

Note: You can configure up to 31 HA groups on an AX, and assign a separate HA priority to each. For Active-Standby configurations, use only one group ID. For Active-Active configurations, use multiple groups IDs and assign VIPs to different groups.

Redundant HA Interfaces

Each HA interface can be a *primary* HA interface or a *redundant* HA interface:

- Primary HA interface – Primary HA interfaces normally send or receive HA heartbeat messages.
- Redundant HA interface – Redundant HA interfaces send or receive HA heartbeat messages only if a primary HA interface stops sending them.

The AX device can be an initiator or a receiver of HA heartbeat messages.

- Initiator – The AX device becomes an initiator of heartbeat messages on redundant HA interfaces, if the AX device does not receive the specified number of consecutive heartbeat messages, on any of the primary HA interfaces.

- Receiver – The AX device becomes a receiver of heartbeat messages on redundant HA interfaces, if the AX device receives a heartbeat message on any redundant HA interface.

Once transmission of heartbeat packets on redundant HA interfaces is triggered, the AX device continues sending heartbeat messages to the redundant HA interfaces until any of the following occurs:

- If an initiator, the AX device receives at least the minimum specified number of heartbeat messages.
- If a receiver, the AX device stops receiving heartbeat messages from the other AX device on the redundant HA interfaces.

To stop sending heartbeat messages on redundant HA interfaces, the AX device must not be an initiator or a receiver of heartbeat messages on any redundant HA interfaces.

There are no redundant HA interfaces by default. When you configure an HA interface, you can specify whether it is redundant.

Interface-specific IPv6 Link-local Floating IP Addresses

You can associate an IPv6 link-local floating IP address with specific AX data interfaces. This option is useful in cases where you do not want the floating IP address to be associated with all AX interfaces. For example, you can restrict an IPv6 link-local floating IP address to server interfaces only, and prevent the floating IP address from being associated with client interfaces.

Note: In the current release, this option does not apply to IPv4 interfaces or to any types of IPv6 interfaces other than link-local interfaces.

Note Regarding Upgrade to AX Release 2.6.6-P2 or Later

If an IPv6 link-local interface is already configured as a floating IP address for an HA group, the floating IP address must be reconfigured following upgrade. You will need to specify the data interface when you reconfigure the link-local floating IP address. (See the syntax information below.)

The following configuration still works following after the upgrade:

```
AX(config)#floating-ip 2001::24 ha-group 1
```

The following configuration does not work after the upgrade. The configuration needs to be re-added with the correct interface number:

```
AX(config)#floating-ip fe80::def ha-group 1 ethernet 1
```

Using the GUI

1. Select Config > HA > Setting > HA Global.
2. Select Floating IP Address to display the configuration fields for that section.
3. Select the HA group ID from the Group Name drop-down list.
4. Select IPv6, and enter the IPv6 address.
5. Click Add.
6. Click OK.

Using the CLI

To specify the AX data interfaces to associate with an IPv6 link-local floating IP address, use one of the following options at the end of the command line when configuring the floating IPv6 address:

- **ethernet** *port-num*
- **ve** *ve-num*
- **trunk** *num*

If you do not specify any interfaces, the floating IP address is associated with all data interfaces. This is the same behavior as in previous releases.

CLI Example

The following commands configure 2 floating IPv6 addresses for HA group 1. Each floating IPv6 address is assigned to a specific IPv6 link-local data interface.

```
AX(config)#floating-ip fe80::def ha-group 1 ethernet 1
AX(config)#floating-ip fe80::de2 ha-group 1 ve 200
```

Session Synchronization

HA session synchronization sends information about active client sessions to the Standby AX device. If a failover occurs, the client sessions are maintained without interruption. Session synchronization is optional. Without it, a failover causes client sessions to be terminated. Session synchronization can be enabled on individual virtual ports.

Session synchronization applies primarily to Layer 4 sessions. Session synchronization does not apply to DNS sessions. Since these sessions are typically very short lived, there is no benefit to synchronizing them. Likewise, session synchronization does not apply to NATted ICMP sessions or to any static NAT sessions. Synchronization of these sessions is not needed since the newly Active AX device will create a new flow for the session following failover.

To enable session synchronization, see “Enabling Session Synchronization” on page 195.

Session synchronization is required for config sync. Config sync uses the session synchronization link. (For more information, “Manually Synchronizing Configuration Information” on page 198.)

Note: Session synchronization is also called “connection mirroring”.

Session Synchronization Between Different AX Models

This release allows the following AX device models to synchronize its sessions with other supported AX device models:

- AX 5200 devices and AX 5200-11 devices
- AX 5200 devices and AX 5630 devices
- AX 5200-11 devices and AX 5630 devices

This support simplifies migration of HA pairs for the supported models and eliminates the need to schedule downtime while migrating to a newer AX model.

For Nat64, DS-Lite, and CGN/LSN implementations, HA session synchronization is possible across the supported heterogeneous and homogeneous AX devices, provided the devices are running Release 2.6.6-P4 software.

Note: HA session synchronization is not supported when LSN port-batching is enabled.

Migration of AX Devices in an HA Pair

This release supports synchronization of the HA sessions across the listed platforms to facilitate the migration process of old AX device models in an HA pair to newer AX device models. Ideally, this ensures that the Active and the Standby AX devices stay in sync while they are both being migrated from an old model to a newer model. This feature is available only after both of the Active and the Standby AX devices are upgraded to Release 2.6.6-P4.

Migration to Release 2.6.6-P4 will allow for two disparate AX devices in an HA configuration to continue to synchronize HA sessions. For example, if you have two 5200 AX devices in an HA pair with synchronized sessions, you can replace one of the devices with a 5200-11 AX device and still synchronize sessions between the 5200 AX device and the 5200-11 AX device. After you replace one 5200 AX device with a 5200-11 and test to see if the HA sessions are synchronized, you can migrate the other 5200 in the HA pair to a 5200-11 AX device as well. This capability is applicable for all the listed pairs in the bulleted list above. This feature facilitates an easy migration path from older devices to new devices without a disruption of HA session synchronization capabilities.

To migrate, issue the following commands. It is best to begin with upgrading your Standby AX device first:

Note: **For ease of explaining the migration process, the steps assume that you are migrating the AX devices from 5200 devices to 5200-11 devices. The same procedures will apply for the other supported platforms as well.**

1. On your Standby AX device, issue the **write memory** command to save the current working configuration on the local AX device and transfer this file to a remote machine:

   ```
   AX-Standby# write memory filename
   ```

   When you have access to Release 2.6.6-P4 software image file on a server, upgrade your Standby AX device to download the new image to and replace the old image, using the following command:

   ```
   upgrade location use scp://username@ip-address/
   directory/Ax_2_6_6-P4.tgz
   ```

   where *location* is the image area (primary or secondary) from where you will access the AX 2.6.6-P4 image for the upgrade process. Each image area has its own separate startup configuration.

   At the completion of the upgrade process, the Standby AX device will prompt you to reboot from the upgraded Release 2.6.6-P4 image.

Note: At this point the Active AX device will still be running Release 2.4.3 software while the Standby AX device will be running the Release 2.6.6-P4 software.

2. Wait for all HA sessions to be synced from the Active to the Standby device and ensure that both devices are operational as usual.

3. Failover from the Active AX device to the Standby AX device. The migration process requires an administratively triggered failover using one of the following ways:

   - Force the active AX 5200 into self-standby. (Use the **ha force-self-standby** command or GUI page Config Mode > HA > Setting HA Global.)
   - Make sure that pre-emption is enabled on the Active AX device and the Standby AX device using the **ha preemption-enable** command, and then change the priority of each HA group the AX 5200 is active for to a lower value that on the standby AX 5200. (For example, set the priority to 1).
     - In the GUI – Use the page Config Mode > HA > Setting HA Global.
     - In the CLI – Use the **ha group** *group-id* **priority** *num* command.

   In either case, the failover must apply to all HA groups for which the AX device is active.

4. On the previous 5200 Active AX device that is now in Standby mode, repeat the procedures in step 1 to ensure that you have a way to revert to the previous Release 2.4.3 configuration.

5. Proceed with step 1 to upgrade to Release 2.6.6-P4. Save your configuration.

6. When you are ready to replace your 5200 AX device with a 5200-11 AX device, copy the configuration file from your 5200 AX device to your 5200-11 AX device. Make sure that your replacement 5200-11 AX device has the identical configuration as your older 5200 AX device.

7. Remove the previous Active 5200 AX device from service and replace it with a 5200-11 AX device with the 2.6.6-P4 image.

8. Ensure that all HA sessions synchronize and the Active and the Standby devices are operational.

9. (Optionally) Fail over to the AX 5200-11 so that the Standby device resumes its role as the Active device.

10. Repeat step 4 to step 7 on the new Standby device to upgrade that to a 5200-11 as well.

Once upgraded, both the Active and the Standby AX devices will be 5200-11 AX devices and will be running 2.6.6-P4. All the HA sessions will be synchronized on both devices. Since all existing connections are mirrored, you will not need to retransmit or re-establish their connections.

Optional Failover Triggers

In addition to HA interface-based failover, you can configure failover based one any of the following:

- Inactive VLAN (VLAN-based failover)
- Unresponsive gateway router (gateway-based failover)
- Unresponsive real servers (VIP-based failover)

VLAN-based Failover

You can enable HA checking for individual VLANs. When HA checking is enabled for a VLAN, the active AX device in the HA pair monitors traffic activity on the VLAN. If there is no traffic on the VLAN for half the duration of a configurable timeout, the AX device attempts to generate traffic by issuing ping requests to servers if configured, or broadcast ARP requests through the VLAN.

If the AX device does not receive any traffic on the VLAN before the timeout expires, a failover occurs. The timeout can be 2-600 seconds. You must specify the timeout. Although there is no default, A10 recommends trying 30 seconds.

This HA checking method provides a passive means to detect network health, whereas heartbeat messages are an active mechanism. You can use either or both methods to check VLAN health. If you use both methods on a VLAN, A10 recommends that you specify an HA checking interval (timeout) that is much longer than the heartbeat interval.

For a configuration example, see “VLAN-Based Failover Example” on page 190.

Gateway-based Failover

Gateway-based failover uses ICMP health monitors to check the availability of the gateways. If any of the active AX device’s gateways fails a health check, the AX device changes its HA status to Down. If the HA status of the other AX device is higher than Down, a failover occurs.

Likewise, if the gateway becomes available again and all gateways pass their health checks, the AX device recalculates its HA status according to the HA interface counts. If the new HA status of the AX device is higher than the other AX device's HA status, a failover occurs.

Configuration of gateway-based failover requires the following steps:

1. Configure a health monitor that uses the ICMP method.
2. Create a server configuration for the gateway and apply the ICMP health monitor to it.
3. Enable HA checking for the gateway.

For a configuration example, see "Gateway-Based Failover Examples" on page 191.

Route-based Failover

Route-based failover reduces the HA priority of all HA groups on the AX device, if a specific route is missing from the IPv4 or IPv6 route table.

You can configure this feature for individual IP routes. When you configure this feature for a route, you also specify the value to subtract from the HA priority of all HA groups, if the route is missing from the route table.

You can configure this option for up to 100 IPv4 routes and up to 100 IPv6 routes. This option is valid for all types of IP routes supported in this release (static and OSPF).

If the priority of an HA group falls below the priority for the same group on the other AX device in an HA pair, a failover can be triggered.

Notes

- This feature applies only to routes in the data route table. The feature does not apply to routes in the management route table.
- For failover to occur due to HA priority changes, the HA pre-emption option must be enabled.

For a configuration example, see "Route-Based Failover Example" on page 193.

Real Server or Port Health-based Failover

You can configure the AX device to decrease the HA priority of an HA group, if a real server or port's health status changes to Down.

You can configure this feature on individual real servers and ports. The feature is disabled by default. To enable the feature, assign an HA weight to the server or port. If the server or port’s health status changes to Down, the weight value is subtracted from the priority value of the HA group. You can specify a single HA group or allow the priority change to apply to all HA groups.

If the server or port’s status changes back to Up, the weight value is added back to the HA group’s priority value.

If the HA priority of a group falls below the priority of the same group on the other AX device, HA failover can be triggered.

Notes

- The lowest HA priority value a server or port can have is 1.
- If HA weights for an HA group are assigned to both the server and an individual port, and both health checks are unsuccessful, only the server weight is subtracted from the HA group’s priority.
- For failover to occur due to HA priority changes, the HA pre-emption option must be enabled.

How the Active AX Device Is Selected

In Active-Standby configurations, only one AX device is Active and the other is the Standby. After you configure HA, the Active AX device is selected using the process shown in Figure 25.

FIGURE 25 Initial Selection of Active AX Device

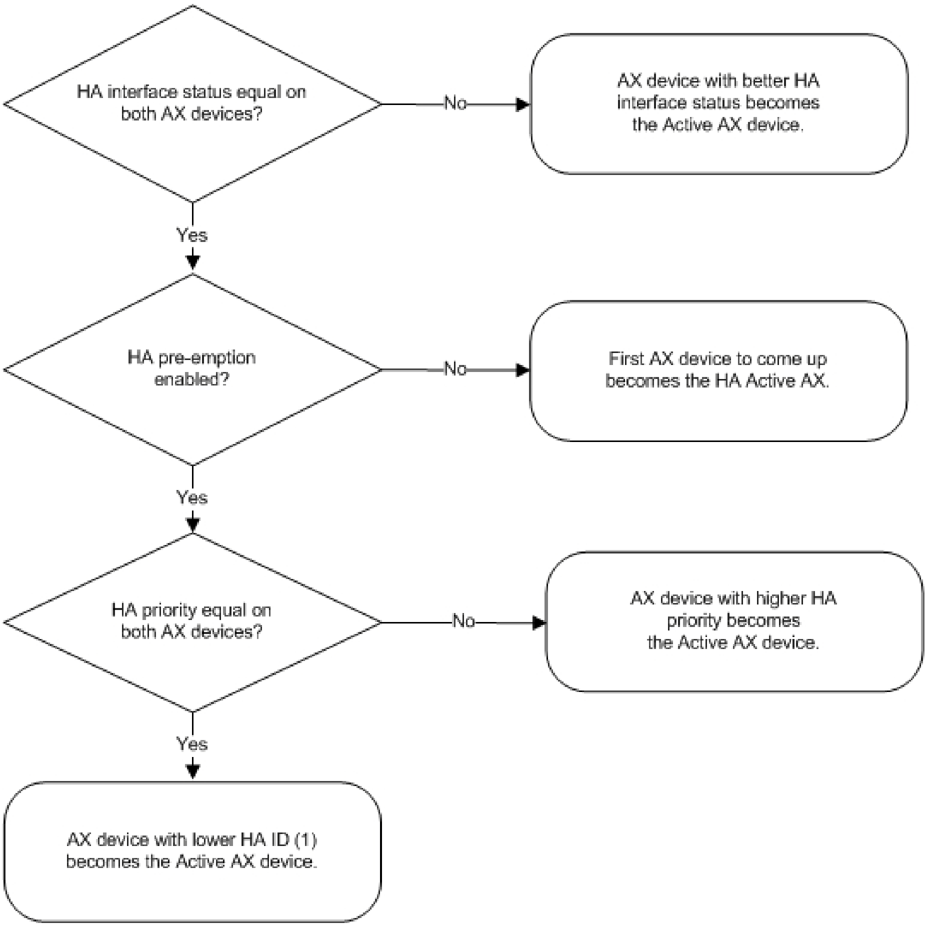
HA interface status equal on both AX devices?
No
AX device with better HA interface status becomes the Active AX device.
Yes
HA pre-emption enabled?
No
First AX device to come up becomes the HA Active AX.
Yes
HA priority equal on both AX devices?
No
AX device with higher HA priority becomes the Active AX device.
Yes
AX device with lower HA ID (1) becomes the Active AX device.

After initial selection of the Active AX device, that device remains the Active AX device unless one of the following events occurs:

- The Standby AX device stops receiving HA heartbeat messages from the Active AX device.
- The HA interface status of the Active AX device becomes lower than the HA interface status of the Standby AX device.
- VLAN-based failover is configured and the VLAN becomes inactive.

- Gateway-based failover is configured and the gateway becomes unavailable.
- HA pre-emption is enabled, and the configured HA priority is changed to be higher on the Standby AX device.

Figure 26 shows the events that can cause an HA failover.

FIGURE 26 HA Failover

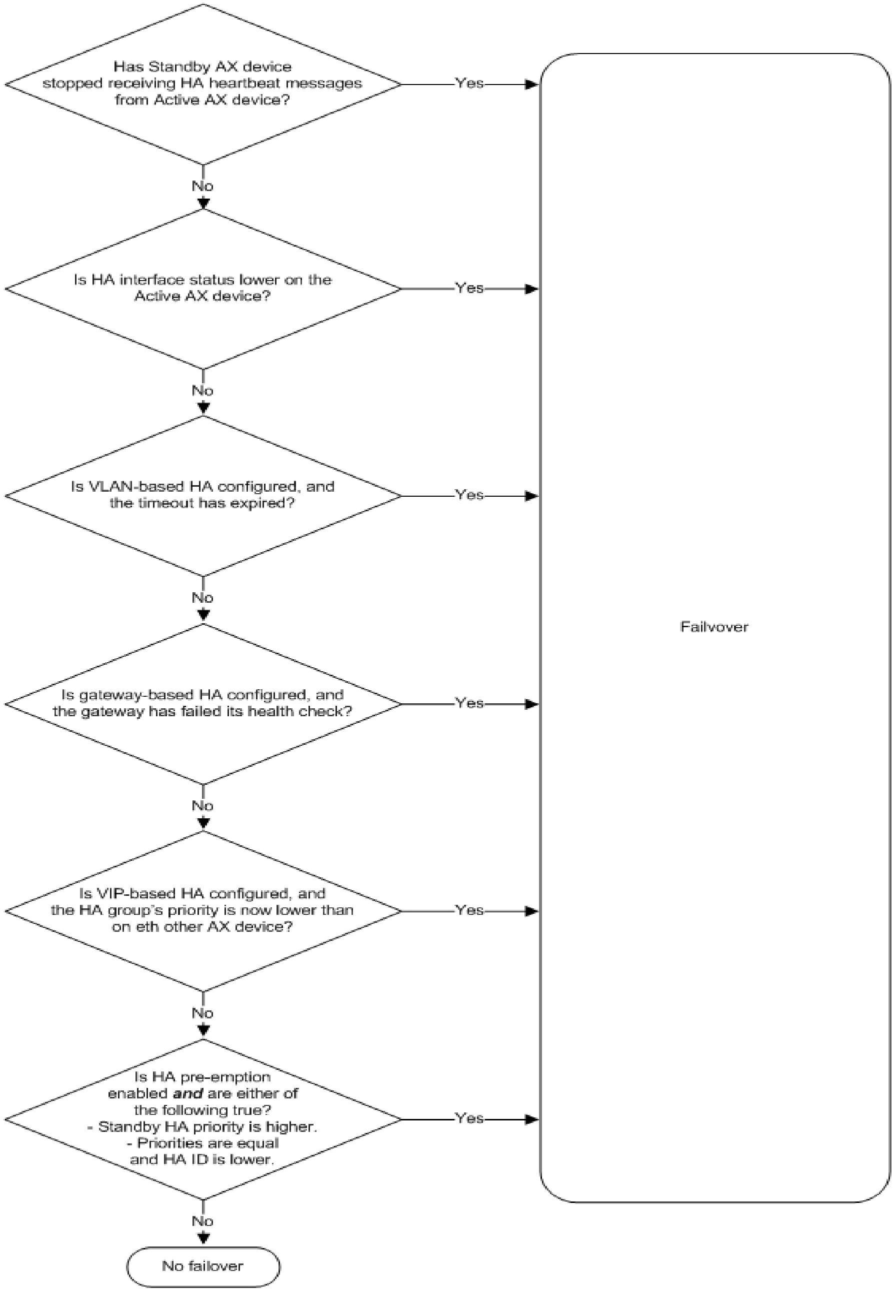
Has Standby AX device stopped receiving HA heartbeat messages from Active AX device?
Yes
No
Is HA interface status lower on the Active AX device?
Yes
No
Is VLAN-based HA configured, and the timeout has expired?
Yes
No
Is gateway-based HA configured, and the gateway has failed its health check?
Yes
No
Is VIP-based HA configured, and the HA group's priority is now lower than on eth other AX device?
Yes
No
Is HA pre-emption enabled and are either of the following true?
- Standby HA priority is higher.
- Priorities are equal and HA ID is lower.
Yes
No
Failvover
No failover

HA Pre-Emption

By default, a failover occurs only in the following cases:

- The Standby AX device stops receiving HA heartbeat messages from the other AX device in the HA pair.
- HA interface state changes give the Standby AX device a better HA state than the Active AX device. (See "HA Interfaces" on page 173.)
- VLAN-based failover is configured and the VLAN becomes inactive. (See "VLAN-based Failover" on page 180.)
- Gateway-based failover is configured and the gateway becomes unavailable. (See "Gateway-based Failover" on page 180.)

By default, failover ***does not*** occur due to HA configuration changes to the HA priority.

To enable the AX devices to failover in response to changes in priority, enable HA pre-emption. When pre-emption is enabled, the AX device with the higher HA priority becomes the Active AX device. If the HA priority is equal on both AX devices, then the AX device with the lower HA ID (1) becomes the Active AX device.

Note: To force Active groups to change to Standby status, without changing HA group priorities and enabling pre-emption, see "Forcing Active Groups to Change to Standby Status" on page 195.

HA Sets

Optionally, you can provide even more redundancy by configuring multiple sets of HA pairs.

FIGURE 27 Multiple HA Pairs

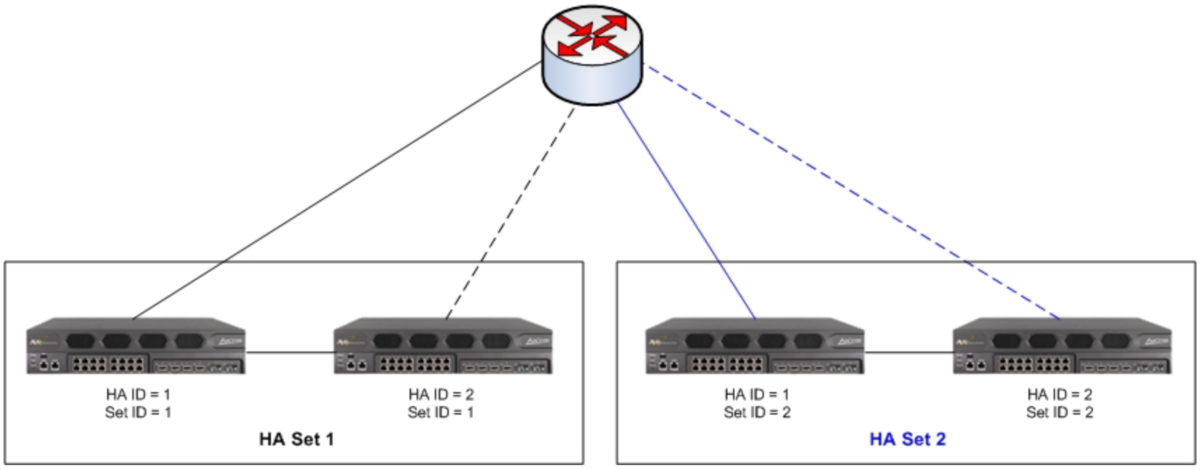

In this example, two HA pairs are configured. Each pair is distinguished by an HA set ID. Each HA pair can be used to handle a different set of real servers.

You can configure up to 7 HA sets. This feature is supported for Layer 2 and Layer 3 HA configurations. The set ID can be specified along with the HA ID.

HA Status Indicators

The HA status of an AX device is displayed in the GUI and CLI. The HA status indicators provide the following information:

- Current HA status of the AX device: Active or Standby
- Configuration status:
  - Most recent configuration update – The system time and date when the most recent configuration change was made.
  - Most recent configuration save – The system time and date when the configuration was saved to the startup-config.
  - Most recent config-sync – The system time and date when the most recent configuration change was made.

In the GUI

The current HA status is shown as one of the following:

- Active
- Standby
- Not Configured

The config-sync status is shown as one of the following:

- Sync
- Not-Sync

The GUI does not indicate when the most recent configuration update or save occurred. This information is available in the CLI. (See below.)

Note: In the current release, the configuration synchronization status is not updated from Not-Sync to Sync if the synchronization target is the running-config.

In the CLI

In the CLI, the HA the status is shown in the command prompt. The status can be one of the following:

- `AX-Active#`
- `AX-Standby#`
- `AX-Forced_Standby#`

Note: If HA is not configured, the prompt is simply the hostname ("AX" by default).

Configuration status is displayed in show running-config output. Here is an example:

```
AX-Active#show running-config
!Current configuration: 8134 bytes
!
!Configuration last updated at 08:11:05 IST Mon May 17 2010
!Configuration last saved at 15:16:49 IST Sat May 15 2010
!Configuration last synchronized at 08:15:02 IST Mon May 17 2010
```

Disabling HA Status Display in the CLI Prompt

Display of the HA status in the CLI prompt is enabled by default. To disable it, use the following command at the global configuration level of the CLI:

```
[no] terminal no-ha-prompt
```

Configuration Examples

The following sections provide HA deployment examples.

- "HA with DS-Lite" on page 188
- "VLAN-Based Failover Example" on page 190
- "Gateway-Based Failover Examples" on page 191
- "Route-Based Failover Example" on page 193

HA with DS-Lite

The following commands configure the HA deployment shown in Figure 24 on page 170.

Commands on AX1

The following commands configure the data interfaces:

```
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ipv6 address 3001::1/16
AX(config-if:ethernet4)#ipv6 address 5001::1/16
AX(config-if:ethernet4)#interface ethernet 5
AX(config-if:ethernet5)#ip address 9.9.9.9 255.255.255.0
AX(config-if:ethernet5)#interface ethernet 6
AX(config-if:ethernet6)#ip address 182.168.20.1 255.255.255.0
AX(config-if:ethernet6)#exit
```

The following commands configure global HA parameters:

```
AX(config)#ha group 1 priority 100
AX(config)#ha interface ethernet 4  no-heartbeat
AX(config)#ha interface ethernet 5
AX(config)#ha interface ethernet 6  no-heartbeat
AX(config)#ha conn-mirror ip 9.9.9.10
AX(config)#ha preemption-enable
```

The following commands configure the floating IP addresses:

```
AX(config)#floating-ip 182.168.20.5 ha-group 1
AX(config)#floating-ip 3001::5 ha-group 1
AX(config)#floating-ip 5001::5 ha-group 1
```

The following commands configure DS-Lite:

```
AX(config)#class-list dslite-1
AX(config-class list)#3001::/16 lsn-lid 1
AX(config-class list)#5001::/16 lsn-lid 1
AX(config-class list)#exit
AX(config)#ip nat pool dslite-1 182.168.20.100 182.168.20.105 netmask /24  ha-
group-id 1 lsn max-users-per-ip 2
AX(config)#lsn-lid 1
AX(config-lsn lid)#source-nat-pool dslite-1
AX(config-lsn lid)#exit
AX(config)#ds-lite fragmentation inbound ipv4
AX(config)#ds-lite fragmentation inbound df-set ipv6
AX(config)#ds-lite fragmentation outbound df-set ipv4
AX(config)#ds-lite inside source class-list dslite-1
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ipv6 nat inside
AX(config-if:ethernet4)#interface ethernet 6
AX(config-if:ethernet6)#ip nat outside
```

Commands on AX2

Here are the commands on device AX2. The following parameters are unique:

- IP interfaces
- Connection mirror address
- HA ID
- HA priority on group 1

The other parameters have the same values as on AX1.

```
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ipv6 address 3001::10/16
AX(config-if:ethernet4)#ipv6 address 5001::111/16
AX(config-if:ethernet4)#interface ethernet 5
AX(config-if:ethernet5)#ip address 9.9.9.10 255.255.255.0
AX(config-if:ethernet5)#interface ethernet 6
AX(config-if:ethernet6)#ip address 182.168.20.2 255.255.255.0
AX(config-if:ethernet6)#exit
AX(config)#ha group 1 priority 100
AX(config)#ha interface ethernet 4  no-heartbeat
AX(config)#ha interface ethernet 5
AX(config)#ha interface ethernet 6  no-heartbeat
AX(config)#ha conn-mirror ip 9.9.9.9
AX(config)#ha preemption-enable
AX(config)#floating-ip 182.168.20.5 ha-group 1
AX(config)#floating-ip 3001::5 ha-group 1
AX(config)#floating-ip 5001::5 ha-group 1
AX(config)#class-list dslite-1
AX(config-class list)#3001::/16 lsn-lid 1
AX(config-class list)#5001::/16 lsn-lid 1
AX(config-class list)#exit
AX(config)#ip nat pool dslite-1 182.168.20.100 182.168.20.105 netmask /24  ha-
group-id 1 lsn max-users-per-ip 2
AX(config)#lsn-lid 1
AX(config-lsn lid)#source-nat-pool dslite-1
AX(config-lsn lid)#exit
AX(config)#ds-lite fragmentation inbound ipv4
AX(config)#ds-lite fragmentation inbound df-set ipv6
AX(config)#ds-lite fragmentation outbound df-set ipv4
AX(config)#ds-lite inside source class-list dslite-1
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ipv6 nat inside
AX(config-if:ethernet4)#interface ethernet 6
AX(config-if:ethernet6)#ip nat outside
```

VLAN-Based Failover Example

To configure VLAN-based failover, use either of the following methods.

(For a description of the feature, see “VLAN-based Failover” on page 180.)

Using the GUI

1. Select Config > HA > Setting > HA Global.
2. In the Status Check section, enter the VLAN ID in the VLAN ID field.
3. Enter the timeout in the Timeout field.

   The timeout can be 2-600 seconds. You must specify the timeout. Although there is no default, A10 recommends trying 30 seconds.
4. Click Add.
5. Repeat step 2 through step 4 for each VLAN to be monitored for HA.
6. Click OK.

Using the CLI

To enable HA checking for a VLAN, use the following command at the global configuration level of the CLI:

```
[no] ha check vlan vlan-id timeout seconds
```

The timeout can be 2-600 seconds. You must specify the timeout. Although there is no default, A10 recommends trying 30 seconds.

The following command enables VLAN-based failover for VLAN 10 and sets the timeout to 30 seconds:

```
AX(config)#ha check vlan 10 timeout 30
```

Gateway-Based Failover Examples

To configure gateway-based failover, use either of the following methods.

(For a description of the feature, see “Gateway-based Failover” on page 180.)

Using the GUI

1. Configure a health monitor that uses the ICMP method:
   a. Select Config > Service > Health Monitor.
   b. Select Health Monitor on the menu bar.
   c. Click Add.
   d. In the Health Monitor section, enter a name for the monitor.
   e. In the Method section, select ICMP from the Type drop-down list.
   f. Click OK.

2. Create a server configuration for the gateway and apply the ICMP health monitor to it:
   a. Select Config > Service > SLB.
   b. Select Server on the menu bar.
   c. Click Add. The General section appears.
   d. In the General section, enter a name for the gateway in the Name field.
   e. In the IP Address field, enter the IP address of the gateway.

f. In the Health Monitor drop-down list, select the ICMP health monitor you configured in step 1.

g. Click OK.

3. Enable gateway-based failover:

   a. Select Config > HA > Setting > HA Global.

   b. In the Status Check section, enter the IP address of the gateway in the IP Address field.

   c. Click Add.

   d. Repeat step b and step c for each gateway to be monitored for HA.

   e. Click OK.

USING THE CLI

1. To configure a health monitor for a gateway, use the following commands.

   [no] **health monitor** *monitor-name*

   Enter this command at the global Config level.

   [no] **method icmp**

   Enter this command at the configuration level for the health monitor.

2. To create a server configuration for the gateway and apply the health monitor to it, use the following command.

   [no] **slb server** *server-name ipaddr*

   [no] **health-check** *monitor-name*

3. To enable HA health checking for the gateway, use the following command at the global configuration level.

   [no] **ha check gateway** *ipaddr*

CLI Example—IPv4

The following commands configure an ICMP health method:

```
AX(config)#health monitor gatewayhm1
AX(config-health:monitor)#method icmp
AX(config-health:monitor)#exit
```

The following commands configure a real server for the gateway and apply the health monitor to it:

```
AX(config)#slb server gateway1 10.10.10.1
AX(config-real server)#health-check gatewayhm1
AX(config-real server)#exit
```

The following command enables HA health checking for the gateway:

```
AX(config)#ha check gateway 10.10.10.1
```

CLI Example—IPv6

The following commands configure an ICMP health method:

```
AX(config)#health monitor v6-gw1
AX(config-health:monitor)#method icmp
AX(config-health:monitor)#exit
```

The following commands configure a real server for the IPv6 gateway and apply the health monitor to it:

```
AX(config)#slb server gateway1 2001::10
AX(config-real server)#health-check v6-gw1
AX(config-real server)#exit
```

The following command enables HA health checking for the gateway:

```
AX(config)#ha check gateway 2001::10
```

Route-Based Failover Example

You can configure HA route awareness for IPv4 routes and IPv6 routes.

Note: The current release does not support this feature in the GUI.

HA Route Awareness for IPv4 Routes

To configure HA route awareness for an IPv4 route, use the following command at the global configuration level of the CLI:

```
[no] ha check route destination-ipaddr /mask-length
priority-cost weight
[gateway ipaddr]
[protocol {static | dynamic}]
[distance num]
```

The *destination-ipaddr /mask-length* option specifies the destination IPv4 subnet of the route.

The **priority-cost** *weight* option specifies the value to subtract from the HA priority of each HA group, if the IP route table does not have a route to the destination subnet.

The **gateway** *addr* option specifies the next-hop gateway for the route.

The **protocol** option specifies the source of the route:

- **static** – The route was added by an administrator.
- **dynamic** – The route was added by a routing protocol. (This includes redistributed routes.)

The **distance** *num* option specifies the metric value (cost) of the route.

Omitting an optional parameter matches on all routes. For example, if you do not specify the next-hop gateway, routes that match based on the other parameters can have any next-hop gateway.

HA Route Awareness for IPv6 Routes

To configure HA route awareness for an IPv6 route, use the following command at the global configuration level of the CLI:

```
[no] ha check route
destination-ipv6addr/mask-length
priority-cost weight
[gateway ipv6addr]
[protocol {static | dynamic}]
[distance num]
```

The *destination-ipv6addr/mask-length* option specifies the destination IPv6 address. The other options are the same as those for IPv4 routes. (See above.)

CLI Examples

The following command configures HA route awareness for a default IPv4 route. If this route is not in the IP route table, 255 is subtracted from the HA priority of all HA groups.

```
AX(config)#ha check route 0.0.0.0 /0 priority-cost 255
```

Note: The lowest possible HA priority value is 1. Deleting 255 sets the HA priority value to 1, regardless of the original priority value.

The following command configures HA route awareness for a dynamic route to subnet 10.10.10.x with route cost 10. If the IP route table does not have a dynamic route to this destination with the specified cost, 10 is subtracted from the HA priority value for each HA group.

```
AX(config)#ha check route 10.10.10.0 /24 priority-cost 10 protocol dynamic distance 10
```

The following commands configure HA route awareness for an IPv6 route to 3000::/64. Based on the combination of these commands, if the IPv6 route table does not contain any routes to the destination, 105 is subtracted from the HA priority of each HA group.

If the IPv6 route table does contain a static route to the destination, but the next-hop gateway is not 2001::1, the AX device subtracts only 5 from the HA priority of each HA group.

```
AX(config)#ha check route 3000::/64 priority-cost 100
AX(config)#ha check route 3000::/64 priority-cost 5 protocol static gateway 2001::1
```

Forcing Active Groups to Change to Standby Status

To force HA groups to change from Active to Standby status, use the following command at the global configuration level of the CLI:

```
[no] ha force-self-standby [group-id]
```

If you specify a group ID, only the specified group is forced to change from Active to Standby. If you do not specify a group ID, all Active groups are forced to change to Standby status.

CLI Example

The following command forces HA group 1 to change from Active to Standby status:

```
AX(config)#ha force-self-standby 1
```

Enabling Session Synchronization

Session synchronization backs up live client sessions on the Backup AX device.

To enable session synchronization:

- Globally enable the feature, specifying the IP address of the other AX device in the HA pair.
- Enable the feature on individual virtual ports. Session synchronization is supported for Layer 4 sessions.

Note: HA session synchronization is required for persistent sessions (source-IP persistence, and so on), and is therefore automatically enabled for these sessions by the AX device. Persistent sessions are synchronized even if session synchronization is disabled in the configuration.

Using the GUI

To globally enable the feature:

1. Select Config > HA > Setting.
2. On the menu bar, select HA Global.
3. In the Mirror IP Address field, enter the IP address of a data interface on the other AX device in the HA pair.
4. Click OK or Apply.

To enable the feature on individual virtual ports:

1. Select Config > Service > Server.
2. On the menu bar, select Virtual Server.
3. Click on the virtual server name.
4. On the Port tab, select the port and click Edit.
5. Select Enabled next to HA Connection Mirror.

Note: If the HA Connection Mirror option is not displayed, session synchronization is not supported for this service type.

6. Click OK to redisplay the Port tab.
7. Click OK again.

Using the CLI

To globally enable session synchronization, use the following command at the global configuration level of the CLI:

```
[no] ha conn-mirror ip ipaddr
```

The *ipaddr* must be an IP address of a data interface on the other AX device.

To enable session synchronization on a virtual port, use the following command at the configuration level for the port:

```
[no] ha-conn-mirror
```

CLI Example

The following command sets the session synchronization address to 10.10.10.66, the IP address of the other AX in this HA pair:

```
AX(config)#ha conn-mirror ip 10.10.10.66
```

The following commands access the configuration level for a virtual port and enable connection mirroring on the port:

```
AX(config)#slb virtual-server vip1 10.10.10.100
AX(config-slb virtual server)#port 80 tcp
AX(config-slb virtual server-slb virtua...)#ha-conn-mirror
```

Configuring OSPF-Related HA Parameters

The following sections describe how to configure OSPF-related HA parameters.

OSPF Awareness of HA

The AX device uses HA-aware VIPs, floating IPs, IP NAT pools, and IP range lists with route redistribution to achieve HA-aware dynamic routing. However, by default, the OSPF protocol on the AX device is not aware of the HA state (Active or Standby) of the AX device. Consequently, following HA failover of an AX device, other OSPF routers might continue forwarding traffic to the Standby AX device (the former Active AX device), instead of the new Active AX device.

You can assign an additional cost to an AX device’s OSPF interfaces when the HA status for any group on the device is Standby. If failover of one or more HA groups from Active to Standby occurs, the AX device does the following:

- Updates the cost of all its OSPF interfaces
- Sends Link-State Advertisement (LSA) updates to its OSPF neighbors advertising the interface cost change

After an OSPF neighbor receives the LSA update, the neighbor updates its OSPF link-state database with the increased cost of the links. The increased cost biases route selection away from paths that use the Standby AX device.

Similarly, if a failover results in HA status Active for all HA groups on an AX device, the device removes the additional cost added for Standby status from all its OSPF interfaces and sends LSA updates to advertise the change. The reduced cost biases route selection toward paths that use the Active AX device and away from paths that use the Standby AX device.

Note: The additional cost for Standby status is removed only if the HA status for *all* HA groups on the device is Active. Otherwise, if the status of *any* of the groups is Standby, the additional cost remains in effect for all OSPF interfaces on the device.

Enabling OSPF Awareness of HA

To enable OSPF awareness of HA, use the following command at the OSPF configuration level.

```
[no] ha-standby-extra-cost num
```

The *num* option specifies the extra cost to add to the AX device's OSPF interfaces, if the HA status of one or more of the device's HA groups is Standby. You can specify 1-65535. If the resulting cost value is more than 65535, the cost is set to 65535.

Enter the command on each of the AX devices in the HA pair.

Manually Synchronizing Configuration Information

This section describes how to manually synchronize the configuration. You can use config-sync options to manually synchronize some or all of the following:

- Startup-config, to the other AX device's startup-config or running-config
- Running-config, to the other AX device's running-config or startup-config)
- Data files:
  - SSL certificates and private-key files
  - aFleX files
  - External health check files
  - Black/white-list files

Requirements

Session synchronization (connection mirroring) is required for manual config sync. Config sync uses the session synchronization link. To enable session synchronization, see “Enabling Session Synchronization” on page 195.

SSH management access must be enabled on both ends of the link. (See “Securing Admin Access by Ethernet” on page 250.)

The link must be on a data interface, not on the management interface.

Configuration Items That Are Backed Up

The following configuration items are backed up during HA configuration synchronization:

- Admin accounts and settings
- AAA settings
- DDoS protection settings
- ICMP rate limiting
- Floating IP addresses
- IP NAT configuration, including IPv6 migration features
- ACLs
- Health monitors
- Data Files:
  - External health check files
  - SSL certificates, private-key files, and CRLs
  - Class-list files

Note: For IP NAT configuration items to be backed up, you must specify an HA group ID as part of the NAT configuration.

Configuration Items That Are *Not* Backed Up

The following configuration items are *not* backed up during HA configuration synchronization:

- Interface-specific management access settings (the ones described in “Securing Admin Access by Ethernet” on page 250)
- AX Hostname

- MAC addresses
- Management IP addresses
- Trunks or VLANs
- Interface settings
- RIP, OSPF, IS-IS, or BGP settings
- ARP entries or settings

Reload of the Target AX Device

In certain cases, the target AX device is automatically reloaded, but in other cases, reload is either optional or is not allowed.

Table 10 lists the cases in which reload is automatic, optional, or not allowed.

TABLE 10 Reload of Target AX Device After Config-Sync

Admin Role	Status of Target AX*	Target Config	Reload?
Root or Super User (Read-Write)	Standby	startup-config	Automatic
		running-config	Automatic
	Active	startup-config	Optional† Not reloaded by default
		running-config	Automatic

*. “Active” means the AX device is currently the active device for at least one HA group.

†. If the target AX device is not reloaded, the GUI Save button on the Standby AX device does not blink to indicate unsaved changes. It is recommended to save the configuration if required to keep the running-config before the next reboot.

Caveats

Before synchronizing the Active and Standby AX devices, verify that both are running the same software version. HA configuration synchronization between two different software versions is not recommended, since some configuration commands in the newer version might not be supported in the older version.

The HA configuration synchronization process does not check user privileges on the Standby AX device and will synchronize to it using read-only privileges. However, you must be logged onto the Active AX with configuration (read-write) access.

Do not make other configuration changes to the Active or Standby AX device during synchronization.

Data that is synchronized from a Standby AX device to an Active AX device is not available on the Active AX device until that device is rebooted or the software is reloaded.

Performing HA Synchronization

To synchronize the AX devices in an HA configuration, use the CLI commands described below.

USING THE GUI

1. Select Config > HA > Config Sync.
2. In the User and Password fields, enter the admin username and password for logging onto the other AX device.
3. Next to Operation, select the information to be copied to the other AX device:
   - All – Copies all the following to the other AX device:
     - Floating IP addresses
     - IP NAT configuration
     - Access control lists (ACLs)
     - Health monitors
     - Data files (see below)

     The items listed above that appear in the configuration file are copied to the other AX device’s running-config.
   - Data Files – Copies only the SSL certificates and private-key files, aFleX files, External health check files, and black/white-list files to the other AX device
   - Running-config – Copies everything listed for the All option, ***except*** the data files, from this AX device’s running-config
   - Startup-config – Copies everything listed for the All option, ***except*** the data files, from this AX device’s startup-config
4. Next to Peer Option, select the target for the synchronization:
   - To Running-config – Copies the items selected in step 3 to the other AX device’s running-config
   - To Startup-config – Copies the items selected in step 3 to the other AX device’s startup-config
5. To reload the other AX device after synchronization, select With Reload. Otherwise, the other AX device is not reloaded following the synchronization.

Note: In some cases, reload of the other AX device either is automatic or is not allowed. See Table 10 on page 200.

6. Click OK.

Using the CLI

The **ha sync** commands are available at the global configuration level of the CLI.

To synchronize data files and the running-config, use the following command:

```
ha sync all
{to-startup-config [with-reload] |
  to-running-config} ipaddr
```

Note: The *ipaddr* option specifies the IP address of the other AX device.

In some cases, reload of the other AX device either is automatic or is not allowed. See Table 10 on page 200.

To synchronize the Active AX device's startup-config to the Standby AX device's startup-config or running-config, without also synchronizing the data files, use the following command:

```
ha sync startup-config
{to-startup-config [with-reload] |
  to-running-config} ipaddr
```

To synchronize the Active AX device's running-config to the Standby AX device's running-config or startup-config, without also synchronizing the data files, use the following command:

```
ha sync running-config
{to-startup-config [with-reload] |
  to-running-config} ipaddr
```

To synchronize the data files by copying the Active AX device's data files to the Standby AX device, use the following command:

```
ha sync data-files ipaddr
```

Tip for Ensuring Fast HA Failover

You can use health checking of the upstream and downstream routers to help ensure rapid HA failover.

The time it takes for traffic to reconverge following HA failover can vary based on the network environment, and depends on the following:

- How fast the ARPs (typically, ARPs of the default gateways) are learned on the newly active AX device
- How fast the MAC tables in the devices along the traffic paths are updated

To help reconvergence occur faster, you can create a real server configuration for each router, and use an ICMP health monitor for checking the health of the gateways. The health checks keep the ARP entries for the gateway routers active, which can help to reduce reconvergence time considerably.

In a typical configuration that includes a client-side router and a server-side router, configure a real server for each router.

To configure health checking of the gateway routers:

1. (Optional) Configure an ICMP health monitor.

   For Layer 3 inline deployments, it is recommended to use very short values (1 second) for the interval and timeout.

2. Create a server configuration for each gateway. If you plan to use a custom ICMP health monitor (previous step), apply the health monitor to the server.

Perform these steps on ***both*** AX devices in the HA pair.

Note: The AX device also has an HA gateway health checking feature. This feature also uses ICMP health monitors. However, if you use the HA gateway health checking feature, HA failover is triggered if a gateway fails a health check. If you use real server configurations instead, as shown in the following examples, HA failover is not triggered by a failed health check.

CLI Example – IPv4

```
AX(config)#health monitor gatewayhm1
AX(config-health:monitor)#method icmp interval 1 timeout 1
AX(config-health:monitor)#exit
AX(config)#slb server gateway-upstream 192.168.10.1
AX(config-real server)#health-check gatewayhm1
AX(config-real server)#exit
AX(config)#slb server gateway-downstream 10.10.10.1
AX(config-real server)#health-check gatewayhm1
AX(config-real server)#exit
```

To use the default ICMP health monitor instead, the configuration is even simpler:

```
AX(config)#slb server gateway-upstream 192.168.10.1
AX(config-real server)#exit
AX(config)#slb server gateway-downstream 10.10.10.1
AX(config-real server)#exit
```

CLI Example – IPv6

```
AX(config)#health monitor gatewayhm1
AX(config-health:monitor)#method icmp interval 1 timeout 1
AX(config-health:monitor)#exit
AX(config)#slb server up-router 2309:e90::1
AX(config-real server)#health-check gatewayhm1
AX(config-real server)#exit
AX(config)#slb server down-router 2309:e90::3
AX(config-real server)#health-check gatewayhm1
AX(config-real server)#exit
```

To use the default ICMP health monitor:

```
AX(config)#slb server up-router 2309:e90::1
AX(config-real server)#exit
AX(config)#slb server down-router 2309:e90::3
AX(config-real server)#exit
```

Network Address Translation

This chapter describes Network Address Translation (NAT) and how to configure it. NAT translates the source or destination IP address of a packet before forwarding the packet.

Note: This chapter does not include information about NAT features for Server Load Balancing (SLB) or for IPv6 migration.

Overview

The AX device supports traditional, Layer 3 IP source NAT. IP source NAT translates internal host addresses into routable addresses before sending the host’s traffic to the Internet. When reply traffic is received, the AX device then retranslates addresses back into internal addresses before sending the reply to the client.

You can configure dynamic or static IP source NAT:

- Dynamic source IP NAT – Internal addresses are dynamically translated into external addresses from a pool.
- Static source IP NAT – Internal addresses are explicitly mapped to external addresses.

Configuration Elements for Dynamic NAT

Dynamic NAT uses the following configuration elements:

- Access Control List (ACL) – to identify the inside host addresses to be translated
- Pool – to identify a contiguous range of external addresses into which to translate inside addresses
- Optionally, pool group – to use non-contiguous address ranges. To use a non-contiguous range of addresses, you can configure separate pools, then combine them in a pool group and map the ACL to the pool group. The addresses within an individual pool still must be contiguous, but you can have gaps between the ending address in one pool and the starting address in another pool. You also can use pools that are in different subnets.

  A pool group can contain up to 5 pools. Pool group members must belong to the same protocol family (IPv4 or IPv6) and must use the same HA ID. A pool can be a member of multiple pool groups. Up to 50 NAT pool groups are supported.

If a pool group contains pools in different subnets, the AX device selects the pool that matches the outbound subnet. For example, if there are two routes to a given destination, in different subnets, and the pool group has a pool for one of those subnets, the AX selects the pool that is in the subnet for the outbound route.

The AX device searches the pools beginning with the first one added to the group, and selects the first match. If none of the pools are in the destination subnet, the AX uses the first pool that has available addresses.

- Inside NAT setting on the interface connected to the inside host.
- Outside NAT setting on the interface connected to the Internet. Inside host addresses are translated into external addresses from a pool before the host traffic is sent to the Internet.

Note: The AX device enables you to specify the default gateway for an IP source NAT pool to use. However, the pool's default gateway can be used only if the data route table already has either a default route or a direct route to the destination of the NAT traffic. In this case, the pool's default gateway will override the route, for NAT traffic that uses the pool.

If the data route table does not have a default route or a direct route to the NAT traffic destination, the pool's default gateway can not be used. In this case, the NAT traffic can not reach its destination.

Configuration Elements for Static NAT

Static NAT uses the following configuration elements:

- Static mappings or an address range list – A static mapping is a one-to-one mapping of an inside address to an external address. An address range list is a contiguous range of inside addresses and external addresses to translate them into.
- Inside NAT setting on the interface connected to the inside host.
- Outside NAT setting on the interface connected to the Internet. Inside host addresses are translated into external addresses from a static mapping or a range list before the host traffic is sent to the Internet.

Configuring Dynamic IP Source NAT

To configure dynamic source NAT:

1. Configure an Access Control List (ACL) to identify the inside addresses that need to be translated.

2. Configure a pool of external addresses to use for translation. To use non-contiguous ranges of addresses, configure multiple pools and add them to a pool group.

3. Enable inside source NAT and map the ACL to the pool.

4. Enable inside NAT on the interfaces connected to the inside hosts.

5. Enable outside NAT on the interfaces connected to the Internet.

Note: In addition, on some AX models, if Layer 2 IP NAT is required, you also must enable CPU processing on the NAT interfaces. This applies to models AX 3200-12, AX 3400, AX 5100, AX 5200, AX 5200-11, and AX 5630. This additional step is performed at the configuration level for each NAT interface. The procedures below do not include this additional step.

Using the GUI

1. To configure an ACL to identify the inside addresses that need to be translated:
   a. Select Config > Network > ACL.
   b. Select the ACL type, Standard or Extended, on the menu bar.
   c. Click Add.
   d. Enter or select the values to filter.
   e. Click OK. The new ACL appears in the Standard ACL table or Extended ACL table.
   f. Click OK to commit the ACL change.

2. To configure a pool of external addresses to use for translation:
   a. Select Config > Service > IP Source NAT.
   b. Select IPv4 Pool or IPv6 Pool on the menu bar.
   c. Click Add.
   d. Enter a name for the pool.
   e. Enter the start and end addresses.

f. Enter the network mask.

g. If the AX device is deployed in transparent mode, enter the default gateway to use for NATted traffic.

h. To use session synchronization for NAT translations, select the HA group.

i. Click OK.

3. To enable inside source NAT and map the ACL to the pool:

a. Select Config > Service > IP Source NAT, if not already selected.

b. Select Binding on the menu bar.

c. Select the ACL number from the ACL drop-down list.

d. Select the pool ID from the NAT Pool drop-down list.

e. Click Add. The new binding appears in the ACL section.

f. Click OK.

4. To enable inside NAT on the interfaces connected to the inside hosts:

a. Select Config > Service > IP Source NAT, if not already selected.

b. Select Interface on the menu bar.

c. Select the interface connected to the internal hosts.

d. In the Direction drop-down list, select Inside.

e. Click Add.

f. Repeat for each interface connected to the internal hosts.

g. Do not click OK yet. Instead, go to the next step.

5. To enable outside NAT on the interfaces connected to the Internet:

a. Select the interface connected to the Internet.

b. In the Direction drop-down list, select Outside.

c. Click Add.

d. Repeat for each interface connected to the Internet.

e. Click OK.

FIGURE 28 Configure > Network > ACL > Standard ACL

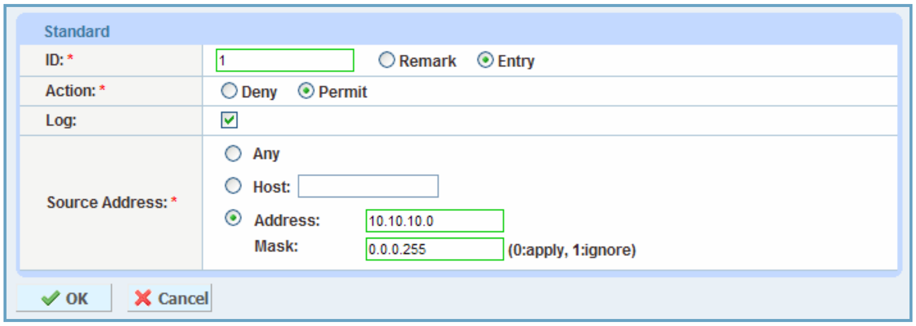

Standard

ID: * 1 ○ Remark ● Entry

Action: * ○ Deny ● Permit

Log: ☑

Source Address: *
○ Any
○ Host:
● Address: 10.10.10.0
Mask: 0.0.0.255 (0:apply, 1:ignore)

FIGURE 29 Configure > Service > IP Source NAT > IPv4 Pool

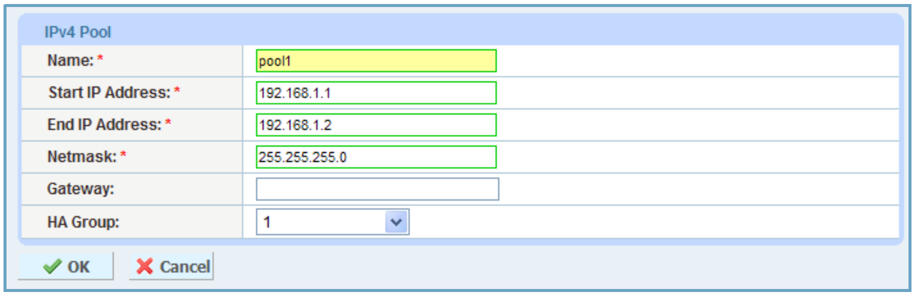

IPv4 Pool

Name: *	pool1
Start IP Address: *	192.168.1.1
End IP Address: *	192.168.1.2
Netmask: *	255.255.255.0
Gateway:	
HA Group:	1

FIGURE 30 Configure > Service > IP Source NAT > Binding

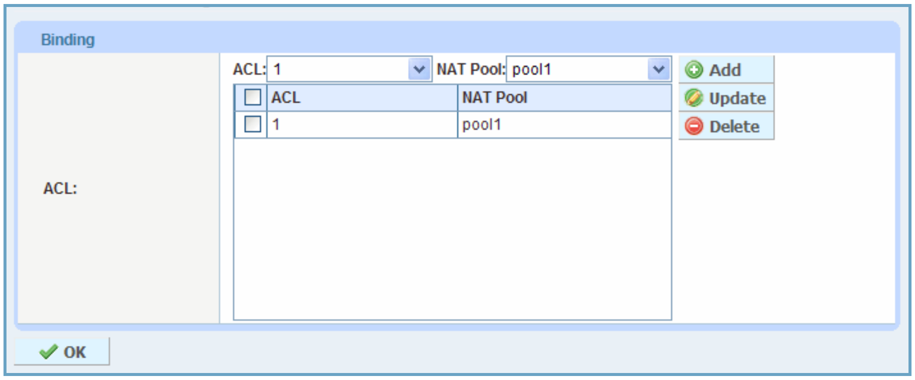

Binding

ACL: ACL: 1 NAT Pool: pool1

□	ACL	NAT Pool
□	1	pool1

FIGURE 31 Configure > Service > IP Source NAT > Interface

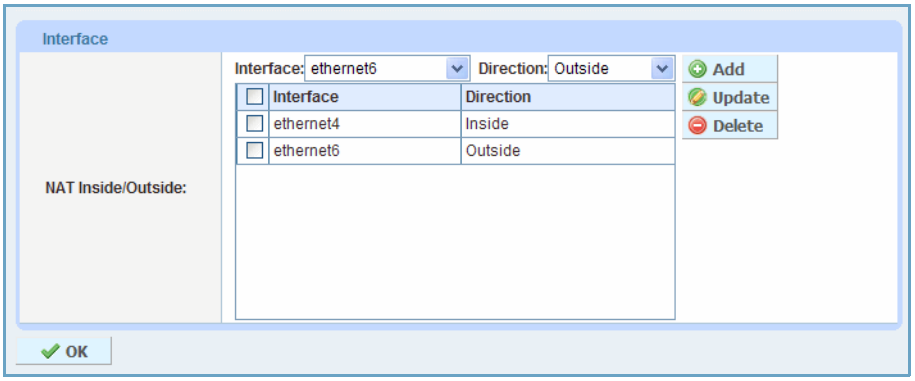

NAT Inside/Outside:

Interface: ethernet6 Direction: Outside

☐	Interface	Direction
☐	ethernet4	Inside
☐	ethernet6	Outside

USING THE CLI

1. To configure an ACL to identify the inside addresses that need to be translated, use either of the following commands at the global configuration level of the CLI.

   Use a standard ACL to specify the host IP addresses to translate. All host addresses that are permitted by the ACL are translated before traffic is sent to the Internet.

   To also specify other information including destination addresses and source and destination protocol ports, use an extended ACL.

Standard ACL Syntax

```
access-list acl-num {permit | deny}
source-ipaddr {filter-mask | /mask-length}
```

Extended ACL Syntax

```
access-list acl-num {permit | deny} {ip | icmp}

{any | host host-src-ipaddr |
  net-src-ipaddr {filter-mask | /mask-length}}

{any | host host-dst-ipaddr |
  net-dst-ipaddr {filter-mask | /mask-length}}
```

or

```
access-list acl-num {permit | deny} {tcp | udp}

{any | host host-src-ipaddr |
  net-src-ipaddr {filter-mask | /mask-length}}
  [eq src-port | gt src-port | lt src-port |
  range start-src-port end-src-port]

{any | host host-dst-ipaddr |
  net-dst-ipaddr {filter-mask | /mask-length}}
  [eq dst-port | gt dst-port | lt dst-port |
  range start-dst-port end-dst-port]
```

2. To configure a pool of external addresses to use for translation, use one of the following commands at the global configuration level of the CLI.

   To configure an IPv4 pool:

   ```
   ip nat pool pool-name start-ipaddr end-ipaddr
   netmask {subnet-mask | /mask-length}
   [gateway ipaddr]
   [ha-group-id group-id [ha-use-all-ports]]
   ```

Note: The **ha-use-all-ports** option applies only to DNS virtual ports. Using this option with other virtual port types is not valid. (For information about this option, see the *AX Series CLI Reference*.)

   To configure an IPv6 pool:

   ```
   ipv6 nat pool pool-name
   start-ipv6-addr end-ipv6-addr
   netmask mask-length
   [gateway ipaddr] [ha-group-id group-id]
   ```

   To configure a pool group:

   ```
   ip nat pool-group pool-group-name
   {pool-name ...}
   ```

3. To enable inside source NAT and map the ACL to the pool, use the following command:

   ```
   ip nat inside source list acl-name
   pool {pool-name | pool-group-name}
   ```

4. To enable inside NAT on the interfaces connected to the inside hosts, use the following commands:

   ```
   interface [ethernet port-num | ve ve-num]
   ip nat inside
   ```

   The interface command changes the CLI to the configuration level for the interface connected to the internal hosts. These are the hosts identified by the ACL configured in step 1 and used by the commands in step 2 and step 3.

5. To enable outside NAT on the interfaces connected to the Internet, use the following commands:

   ```
   interface [ethernet port-num | ve ve-num]
   ip nat outside
   ```

CLI Example

The following commands configure an ACL to specify the internal hosts to be NATted. In this example, all hosts in the 10.10.10.x subnet are to receive NAT service for traffic to the Internet.

```
AX(config)#access-list 1 permit 10.10.10.0 0.0.0.255
```

The following command configures an IPv4 pool of external addresses to use for the NAT translations. In this example, 10.10.10.x addresses will be translated into 192.168.1.1 or 192.168.1.2:

```
AX(config)#ip nat pool pool1 192.168.1.1 192.168.1.2 netmask /24
```

The following command enables inside source NAT and associates the ACL with the pool:

```
AX(config)#ip nat inside source list 1 pool pool1
```

The following commands enable inside source NAT on the interface connected to the internal hosts:

```
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ip nat inside
```

The following commands enable source NAT on the interface connected to the external hosts:

```
AX(config-if:ethernet4)#interface ethernet 6
AX(config-if:ethernet6)#ip nat outside
```

Configuring Static IP Source NAT

You can configure individual static source NAT mappings or configure a range of static mappings.

After configuring the static source NAT mappings, do the following:

- Enable inside NAT on the interfaces connected to the inside hosts.
- Enable outside NAT on the interfaces connected to the Internet.

Limitations for Static NAT Mappings

- Application Level Gateway (ALG) services other than FTP are not supported when the server is on the inside.
- HA session synchronization is not supported. However, sessions will not be interrupted by HA failovers.
- Syn-cookies are not supported.

Using the GUI

Note: The GUI supports configuring a static NAT range but does not support configuring individual mappings.

1. To configure the static translations of internal host addresses to external addresses:
   a. Select NAT Range on the menu bar.
   b. Click Add.
   c. Enter a name for the range.
   d. Select the address type (IPv4 or IPv6)
   e. In the From fields, enter the first (lowest numbered) address and network mask in the range of inside host addresses to be translated.
   f. In the To field, enter the first (lowest numbered) address and network mask in the range of external addresses into which to translate the inside host addresses.
   g. In the Count field, enter the number of addresses to be translated.
   h. To apply HA to the addresses, select the HA group.
   i. Click OK.

2. To enable inside NAT on the interfaces connected to the inside hosts:
   a. Select Interface on the menu bar.
   b. Select the interface from the Interface drop-down list.
   c. Select Inside in the Direction drop-down list.
   d. Click OK.
   e. Repeat for each inside interface.

3. To enable outside NAT on the interfaces connected to the Internet:
   a. Select Interface on the menu bar.
   b. Select the interface from the Interface drop-down list.
   c. Select Outside in the Direction drop-down list.
   d. Click OK.
   e. Repeat for each outside interface.

USING THE CLI

1. To configure the external addresses to use for translation, use one of the following commands.

   To configure individual address mappings, use the following command to configure each mapping:

   ```
   ip nat inside source
   static source-ipaddr nat-ipaddr
   [ha-group-id group-id]
   ```

   The *source-ipaddr* is the internal host that will send requests. The *nat-ipaddr* is the address into which the AX device will translate the *source-ipaddr* before forwarding the requests.

   Similarly, for inbound static NAT, the following syntax is supported:

   ```
   [no] ip nat inside source static destination-
   ipaddr nat-ipaddr
   ```

   The *destination-ipaddr* is the internal host to which external hosts send requests. The *nat-ipaddr* is the address into which the AX device will translate the *destination-ipaddr* before forwarding the requests.

   To configure a range list to use for static mappings:

   ```
   ip nat range-list list-name
   source-ipaddr /mask-length
   nat-ipaddr /mask-length
   count number [ha-group-id group-id]
   ```

The *source-ipaddr* specifies the starting address in the range of internal host addresses. The *nat-ipaddr* command specifies the first address in the range of external addresses to use for the translations.

The **count** option specifies how many mappings to create.

2. If you used the **ip nat inside source** command, enter the following command at the global configuration level of the CLI, to enable static NAT support:

```
ip nat allow-static-host
```

Note: This step is not required if you use a static source NAT range list instead.

3. To enable inside NAT on the interfaces connected to the inside hosts, use the following commands:

```
interface [ethernet port-num | ve ve-num]
ip nat inside
```

The **interface** command changes the CLI to the configuration level for the interface connected to the internal hosts.

4. To enable outside NAT on the interfaces connected to the Internet, use the following commands:

```
interface [ethernet port-num | ve ve-num]
ip nat outside
```

CLI Example

The following commands enable static NAT, configure an IP address range named “nat-list-1” that maps up to 100 local addresses starting from 10.10.10.97 to Internet addresses starting from 192.168.22.50, set Ethernet interface 2 as the inside NAT interface, and set Ethernet interface 4 as the outside NAT interface.

```
AX(config)#ip nat range-list nat-list-1 10.10.10.97 /16 192.168.22.50 /16 count
100
AX(config)#interface ethernet 2
AX(config-if:ethernet2)#ip nat inside
AX(config-if:ethernet2)#exit
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#ip nat outside
```

NAT ALG Support for PPTP

NAT Application Level Gateway (ALG) support for the Point-to-Point Tunneling Protocol (PPTP) enables clients and servers to exchange Point-to-Point (PPP) traffic through the AX device over a Generic Routing Encapsulation (GRE) tunnel.

PPTP is used to connect Microsoft Virtual Private Network (VPN) clients and VPN hosts. Figure 32 shows an example.

FIGURE 32 NAT ALG for PPTP

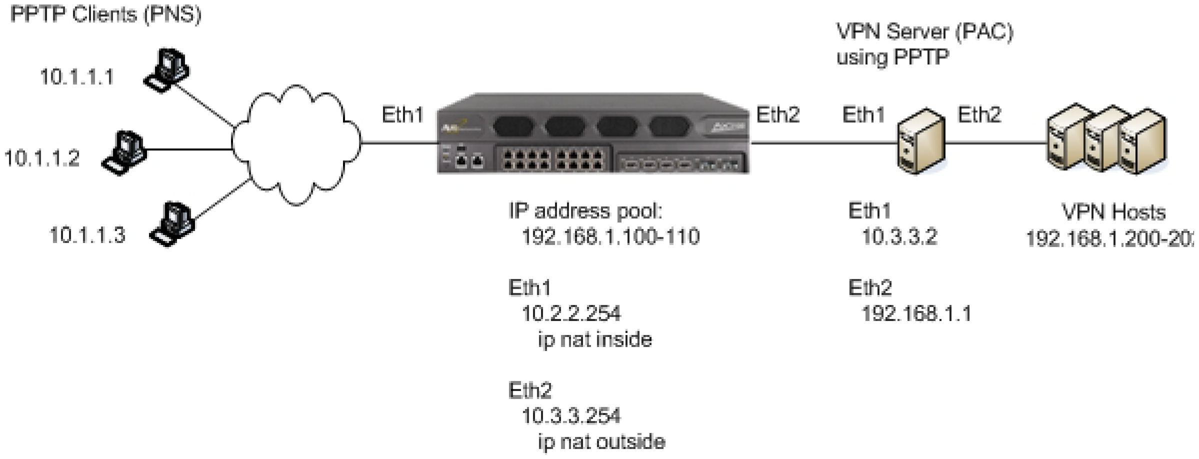

The AX device is deployed between PPTP clients and the VPN server (VPN Server using PPTP). The AX interface connected to the PPTP clients is enabled for inside source NAT. The AX interface connected to the VPN server is enabled for outside source NAT.

Each client runs a PPTP Network Server (PNS). To set up a VPN session, the PNS sends an Outgoing-Call-Request to the PPTP Access Concentrator (PAC), which is the VPN server. The destination TCP port is the PPTP port (1723 by default). The request includes a Call ID that the PNS chooses.

Because multiple clients may share the same NAT address, the AX device must ensure that clients do not share the same Call ID as well. Therefore, the AX device assigns to each client a NAT Call ID (analogous to a NAT source port for TCP) and modifies the Outgoing-Call-Request to use the NAT Call ID instead.

The PAC replies to the Outgoing-Call-Request with a Call ID of its own. This is like a TCP destination port. The AX device does not change the

PAC's Call ID. The PAC then assigns to the client an IP address belonging to the VPN subnet.

On the AX device, the GRE session is created after the PNS sends its reply. In the GRE session, the Call ID is used as the Layer 4 port, instead of a TCP/UDP port number. (See the example of **show session** output in "CLI Example" on page 219.)

In Figure 32 on page 216, client (PNS) 10.1.1.1 wants to connect to a VPN through the VPN Server (PAC) 10.3.3.2, which is using PPTP. Client 10.1.1.1 establishes a PPTP control session (on port 1723) with 10.3.3.2. When the client sends the Outgoing-Call-Request over that TCP session with its desired Call ID, the AX device will translate the Call ID into a unique Call ID for NAT. Once the VPN server replies with its own Call ID, the AX device will establish the GRE session.

After the Call IDs are exchanged, the client and server encapsulate VPN subnet traffic in a GRE tunnel. The GRE tunnel packets are sent under normal IP between 10.1.1.1 and 10.3.3.2. A GRE packet for PPTP uses a Call ID in the same way as a TCP or UDP destination port. Therefore, GRE packets from the server (10.3.3.2) will use the NAT Call ID. The AX device translates the NAT Call ID back into the client's original Call ID before sending the packet to the client.

Note: One GRE session is supported per control session, which means one call at a time is supported. In practice, PPTP is used only for VPNs, in which case multiple concurrent calls do not occur.

Configuring NAT ALG for PPTP

To configure an AX device to support NAT ALG for PPTP:

- Configure dynamic IP source NAT:
  - Configure an ACL that matches on the PPTP client subnet(s).
  - Configure an IP source NAT pool that contains the range of IP addresses into which to translate client addresses.
  - Configure an inside source NAT list, using the ACL and pool.
  - Enable inside IP source NAT on the AX interface connected to the VPN clients.
  - Enable outside IP source NAT on the AX interface connected to the VPN server.
- If NAT ALG support for PPTP is disabled, enable it. (The feature is enabled by default.)

Note: In the current release, NAT ALG support for PPTP is not supported with static NAT or NAT range lists.

Note: In the current release, NAT ALG support for PPTP can not be disabled or re-enabled using the GUI.

Using the CLI

To configure dynamic IP source NAT, use the following commands.

First, to configure the ACL, use the following command at the global configuration level of the CLI:

```
access-list acl-num permit
source-ipaddr {filter-mask | /mask-length}
```

Note: The ACL must permit IP traffic. The syntax above is for a standard ACL. If you plan to use an extended ACL instead, make sure to use the **ip** option, instead of **icmp**, **tcp**, or **udp**.

To configure the IP address pool, use the following command at the global configuration level of the CLI:

```
ip nat pool pool-name start-ipaddr end-ipaddr
netmask {subnet-mask | /mask-length}
[gateway ipaddr] [ha-group-id group-id]
```

To configure an IP source NAT list, use the following command at the global configuration level of the CLI:

```
ip nat inside source list acl-name
pool {pool-name | pool-group-name}
```

To enable inside source NAT on an interface, use the following command at the configuration level for the interface:

```
[no] ip nat inside
```

To enable outside source NAT on an interface, use the following command at the configuration level for the interface:

```
[no] ip nat outside
```

To enable or disable NAT ALG support for PPTP, use the following command at the global configuration level of the CLI:

```
ip nat alg pptp {enable | disable}
```

The feature is enabled by default. The default protocol port number is 1723 and can not be changed.

To display GRE sessions, use the following commands:

show session

To display or clear statistics, use the following commands:

show ip nat alg pptp statistics

clear ip nat alg pptp statistics

CLI Example

The commands in this section implement the NAT ALG for PPTP configuration shown in Figure 32 on page 216.

The following commands configure dynamic inside source NAT.

```
AX(config)#access-list 1 permit 10.1.1.0 0.0.0.255
AX(config)#ip nat pool pptp-pool 192.168.1.100 192.168.1.110 netmask /24
AX(config)#ip nat inside source list 1 pool pptp-pool
```

The following commands specify the inside NAT interface and the outside NAT interface.

```
AX(config)#interface ethernet 1
AX(config-if:ethernet1)#ip address 10.2.2.254 255.255.255.0
AX(config-if:ethernet1)#ip nat inside
AX(config-if:ethernet1)#interface ethernet 2
AX(config-if:ethernet2)#ip address 10.3.3.254 255.255.255.0
AX(config-if:ethernet2)#ip nat outside
```

The following command displays session information:

```
AX(config-if:ethernet2)#show session
Prot Forward Source       Forward Dest         Reverse Source       Reverse Dest
Age Hash

-----------------------------------------------------------------------------------------
--------------------

Gre  10.1.1.1:49152     10.3.3.2:32799     10.3.3.2:32799     192.168.1.100:2109
240 1

Tcp  10.1.1.1:2301      10.3.3.2:1723      10.3.3.2:1723      192.168.1.100:2109
240 2
```

This example shows the GRE session and the TCP session over which the GRE session is transported. For the GRE session, the number following each IP address is the PPTP Call ID. For the TCP session, the number is the TCP protocol port.

The following command displays PPTP NAT ALG statistics.

```
AX(config-if:ethernet2)#show ip nat alg pptp statistics
Statistics for PPTP NAT ALG:
-----------------------------
Calls In Progress:                   10
Call Creation Failure:               0
Truncated PNS Message:               0
Truncated PAC Message:               0
Mismatched PNS Call ID:              1
Mismatched PAC Call ID:              0
Retransmitted PAC Message:           3
Truncated GRE Packets:               0
Unknown GRE Packets:                 0
No Matching Session Drops:           4
```

Fast Aging for IP NATted ICMP and DNS Sessions

The AX device uses application-aware aging for IP NATted sessions, in cases where the AX device performs IP NAT translation of the internal client IP addresses.

The default timeout for IP NATted ICMP sessions, as well as UDP sessions on port 53 (DNS), is set to the SLB maximum session life (MSL), which is 2 seconds by default.

Note: Fast aging applies to sessions between internal clients and external resources, in cases where the AX device performs IP NAT translation of the client addresses. This type of traffic is not SLB traffic between clients and a VIP configured on the AX device. For SLB DNS traffic, short aging based on the MSL time is the default aging mechanism.

Table 11 summarizes the session timeouts and how to configure them.

TABLE 11 Session Timeout for IP NATted ICMP and UDP Sessions

Default Timeout for IP NATted ICMP DNS Sessions	Method To Change Timeout
SLB MSL timeout (2 seconds by default) **Note:** For DNS, this is the default only for the default DNS port (53).	You can use either of the following methods: • Change the SLB MSL timeout. • Change the IP NAT translation timeout: • ICMP – Change the IP NAT translation ICMP timeout, by specifying the number of seconds for the timeout, instead of “fast”. • DNS – Change the IP NAT translation UDP timeout for the DNS port, by specifying the number of seconds for the timeout, instead of “fast”. The timeout is configurable for individual UDP ports.

Using the CLI

To display the timeout that will be used for IP NATted sessions, use the following command:

```
show ip nat timeouts
```

To change the IP NAT translation timeout for ICMP, use the following command:

```
[no] ip nat translation icmp-timeout
        {seconds | fast}
```

To change the IP NAT translation timeout for a UDP port, use the following command:

```
[no] ip nat translation service-timeout
        udp port-num {seconds | fast}
```

The *port-num* option specifies the UDP port number. The **fast** option sets the timeout to the SLB MSL timeout, for the specified UDP port.

CLI Example

The following command displays the current IP NAT translation timeouts:

```
AX#show ip nat timeouts
NAT Timeout values in seconds:
SYN     TCP     UDP     ICMP
------------------------
60      300     300     fast
Service 53/udp has fast-aging configured
```

In this example, the output indicates that fast aging is used for IP NATted ICMP sessions, and for IP NATted DNS sessions on port 53.

The message at the bottom of the display indicates that the fast aging setting (SLB MSL timeout) will be used for IP NATted UDP sessions on port 53. If the message is not shown in the output, then the timeout shown under "UDP" will be used instead.

Client and Server TCP Resets for NATted TCP Sessions

By default, the AX device does not send TCP Resets to the client and server when a NATted TCP session becomes idle. To enable this option, use the following command at the global configuration level of the CLI:

```
ip nat reset-idle-tcp-conn
```

Requirements for Translation of DNS Traffic

If you plan to use IP NAT for DNS traffic, make sure the configuration includes the following:

- Both the DNS request from the inside client, and the response from the external DNS server, must pass through the IP NAT outside interface.
- If an ACL is configured on the interface that will receive the DNS responses (the IP NAT outside interface), the ACL must include a permit rule that allows traffic from the DNS server. Otherwise, the traffic will be denied by the implicit (non-visible) deny any any rule at the end of the ACL.

IP NAT Use in Transparent Mode in Multi-Netted Environment

If the AX device is deployed in transparent mode, the device uses NAT IP addresses to perform health monitoring on servers that are outside the IP subnet or VLAN of the AX device. If there are multiple IP addresses in the NAT pool, the AX device uses only the last IP address in the pool for the health checks. Also, the AX device only responds to control traffic (for example, management and ICMP traffic) on the last IP address in the pool.

In the following example, the AX device's IP address is on the 172.168.101.0/24 subnet. A NAT pool has been configured to reach servers outside of that subnet/VLAN.

```
AX#show ip
System is running in Transparent Mode
IP address:                  172.168.101.4 255.255.255.0
IP Gateway address:          172.168.101.251
SMTP Server address:         Not configured

AX#show ip nat pool
Total IP NAT Pools: 4
Pool Name     Start Address    End Address      Mask  Gateway         HA Group
----------------------------------------------------------------------------
Pool-A        173.168.10.20    173.168.10.25    /24   173.168.10.250 0
```

In this configuration, the AX device will initiate health checks using the last IP address in the pool as the source IP address. In this example, the AX device will use IP address 173.168.10.25. In addition, the AX device will only respond to control traffic directed to 173.168.10.25 from the 173.168.10.0/24 subnet.

NAT Range List Requires AX Interface or Route Within the Global Subnet

In an IP source NAT configuration, return UDP or ICMP traffic may not be able to reach the AX device. This can occur under the following circumstances:

- IP source NAT is configured using a NAT range list.
- The AX device does not have any data interfaces or routes that contain an address within the subnet of the range list's global address(es).

To work around this issue, configure an IP interface that is within the NAT range list’s global subnet. You can configure the address on any active data interface on the AX device.

This issue does not affect NATted traffic other than ICMP or UDP traffic, or use of an ACL with a NAT pool.

IP NAT in HA Configurations

If you are using IP source NAT or full NAT in an HA configuration, make sure to add the NAT pool or range list to an HA group. Doing so allows a newly Active AX device to properly continue management of NAT resources following a failover.

Using the GUI

In the GUI, you can select the HA group from the HA Group drop-down list on the following configuration tabs:

- Config > Service > IP Source NAT > IPv4 Pool
- Config > Service > IP Source NAT > IPv6 Pool
- Config > Service > IP Source NAT > NAT Range

Using the CLI

In the CLI, the **ha-group-id** option is supported with the following NAT commands:

```
[no] ip nat pool pool-name start-ipaddr end-ipaddr

netmask {subnet-mask | /mask-length} [gateway ipaddr]
[ha-group-id group-id]

[no] ipv6 nat pool pool-name start-ipv6-addr
end-ipv6-addr netmask mask-length [gateway ipaddr]
[ha-group-id group-id]

[no] ip nat range-list list-name
source-ipaddr   /mask-length   nat-ipaddr   /mask-length
count number [ha-group-id group-id]
```

Configuring Dynamic IP NAT with Many Pools

The AX device supports use of up to 1023 pools for dynamic IP NAT. If you plan to use a lot of pools (100 or more), use the procedure in this section.

To configure dynamic IP NAT with more than 100 pools:

1. Configure the NAT pools. Each pool can contain a single, contiguous range of public IP addresses.
2. For each pool, configure a Limit ID (LID) and add the pool to the LID.
3. Configure a class list that maps internal clients to the LIDs.
4. Enable inside source NAT.

Configure NAT Pools

To configure a NAT pool, use the following command at the global configuration level of the CLI:

```
[no] ip nat pool pool-name start-ipaddr end-ipaddr
netmask {subnet-mask | /mask-length}
```

The *start-ipaddr* and *end-ipaddr* options specify the beginning and ending public IP addresses in the range to be mapped to internal addresses. The **netmask** option specifies the subnet mask or mask length for the addresses.

Configure LIDs

A Limit ID (LID) assigns a numeric ID to a NAT pool. To configure a LID, use the following commands:

```
[no] lid num
```

Enter this command at the global configuration level of the CLI. The *num* specifies the LID number and can be 1-1024, for a maximum of 1024 LIDs. This command changes the CLI to the configuration level for the LID, where the following command is available:

```
[no] use-nat-pool pool-name
```

This command binds a NAT pool to the LID.

Configure a Class List

To bind internal addresses to LIDs (and thus to NAT pools), use a class list.

You can configure a class list in either of the following ways:

- Use a text editor on a PC or other device to create the list, then import it onto the AX device.
- Use CLI commands to create the list.

Class List Syntax

Each entry (row) in the class list defines a client class, and has the following format: `ipaddr /network-mask glid num`

Each entry consists of the following:

- *ipaddr* – Specifies the inside subnet that requires NAT. The *network-mask* specifies the network mask.

  To configure a wildcard IP address, specify 0.0.0.0 /0. The wildcard address matches on all addresses that do not match any entry in the class list.

- **glid** *num* – Specifies the global LID.

Importing a Class List

After the class list is configured, import it onto the AX device, using the following command at the Privileged EXEC or global configuration level of the CLI:.

```
import class-list file-name url
```

The *file-name* specifies the name the class list will have on the AX device. The *url* specifies the file transfer protocol, username (if required), and directory path.

You can enter the entire URL on the command line or press Enter to display a prompt for each part of the URL. If you enter the entire URL and a password is required, you will still be prompted for the password. To enter the entire URL:

- `tftp://host/file`
- `ftp://[user@]host[:port]/file`
- `scp://[user@]host/file`
- `rcp://[user@]host/file`

Configuring a Class List Using the CLI

To configure a class list in the CLI, use the following commands:

```
[no] class-list name [file]
```

Enter this command at the global configuration level of the CLI.

The **file** option saves the class list as a separate file. Without this option, the class list is instead saved in the startup-config. The **file** option is valid only when you create the class list. After you create the list, the list remains either in the startup-config or in a separate file, depending on whether you use the **file** option when you create the list.

Note: If the class list contains 100 or more entries, as is likely with this feature, it is recommended to use the **file** option.

A class list can be exported from the AX device only if you use the **file** option.

The **class-list** command creates the class list if it is not already configured, and changes the CLI to the configuration level for the list.

```
[no] ipaddr /network-mask glid num
```

- To add an entry to the class list, use the command without “**no**”.
- To modify an entry, use the command without “**no**”. Use the same source IP address as the entry to replace. Entries are keyed by source IP address.
- To delete an entry, use “**no**” followed by the source IP address.

Enable Inside Source NAT

To enable inside source NAT, use the following commands:

[**no**] **ip nat inside source class-list** *name*

Use this command at the global configuration level of the CLI.

[**no**] **ip nat inside**

Use this command on each interface connected to the inside clients.

[**no**] **ip nat outside**

Use this command on each interface connected to the external network or Internet.

Configuration Example

The commands in this example configure dynamic NAT with more than 100 pools.

Class List

This example uses the following class list:

```
10.10.10.1 /32 glid 1
10.10.10.2 /32 glid 2
10.10.10.3 /32 glid 3
10.10.10.4 /32 glid 4
10.10.10.5 /32 glid 5
10.10.10.6 /32 glid 6
...
10.10.10.254 /32 glid 254
```

For brevity, this example and the CLI example do not show LIDs or pools 7-253.

This example uses mask length /32 for each entry, which uses a separate LID (and therefore, a separate pool) for each individual host. The entries are not required to be for individual hosts. For example, the following entry assigns all hosts in subnet 10.10.20.x to LID 10:

```
10.10.10.0 /24 glid 10
```

AX Configuration Commands

The following commands configure the NAT pools:

```
AX(config)#ip nat pool p1 192.168.217.201 192.168.217.201 netmask /24
AX(config)#ip nat pool p2 192.168.217.202 192.168.217.202 netmask /24
AX(config)#ip nat pool p3 192.168.217.203 192.168.217.203 netmask /24
AX(config)#ip nat pool p4 192.168.217.204 192.168.217.204 netmask /24
AX(config)#ip nat pool p5 192.168.217.205 192.168.217.205 netmask /24
AX(config)#ip nat pool p6 192.168.217.206 192.168.217.206 netmask /24
...
AX(config)#ip nat pool p254 192.168.217.254 192.168.217.254 netmask /24
```

The following commands configure the LIDs:

```
AX(config)#lid 1
AX(config-global lid)#use-nat-pool p1
AX(config-global lid)#lid 2
AX(config-global lid)#use-nat-pool p2
AX(config-global lid)#lid 3
AX(config-global lid)#use-nat-pool p3
AX(config-global lid)#lid 4
AX(config-global lid)#use-nat-pool p4
AX(config-global lid)#lid 5
AX(config-global lid)#use-nat-pool p5
AX(config-global lid)#lid 6
AX(config-global lid)#use-nat-pool p6
...
AX(config-global lid)#lid 254
AX(config-global lid)#use-nat-pool p254
AX(config-global lid)#exit
```

The following command imports the class list:

```
AX(config)#import class-list revnat ftp:
Address or name of remote host []?1.1.1.2
User name []?axadmin
Password []?*********
File name [/]?revnat
```

The following command enables inside source NAT:

```
AX(config)#ip nat inside source class-list revnat
```

The following commands enable inside NAT on the interface connected to the internal clients:

```
AX(config)#interface ethernet 1
AX(config-if:ethernet1)#ip nat inside
```

The following commands enable outside NAT on the interface connected to the Internet:

```
AX(config)#interface ethernet 2
AX(config-if:ethernet2)#ip nat outside
AX(config-if:ethernet2)#exit
```

Stateful Firewall for Transparent Sessions

This release provides stateful firewall support for transparent sessions. In the context of this feature, a transparent session is a Layer 3 session that passes through the AX device without any NAT of any kind. By providing ALG services for transparent traffic, the AX device can perform some of the basic functions of a stateful firewall.

This ability can be helpful for scenarios in which an internal user has been assigned a public IP address. In such scenarios, instead of performing NAT, the AX device performs Layer 3 forwarding, which could leave the user exposed to attacks from the public network.

The stateful firewall feature protects the user from outside attacks by using access control lists to deny or reject traffic from unrecognized external sources. In addition, the AX device maintains state information for Application Layer Gateway protocol traffic, which can originate from either side of the firewall, enabling that traffic to pass through the firewall unimpeded.

As Figure 33 shows, the Stateful Firewall feature allows traffic from internal clients to pass through the firewall, while external traffic is filtered using an ACL or based on the state information for the ALG session.

FIGURE 33 AX device acting as Stateful Firewall

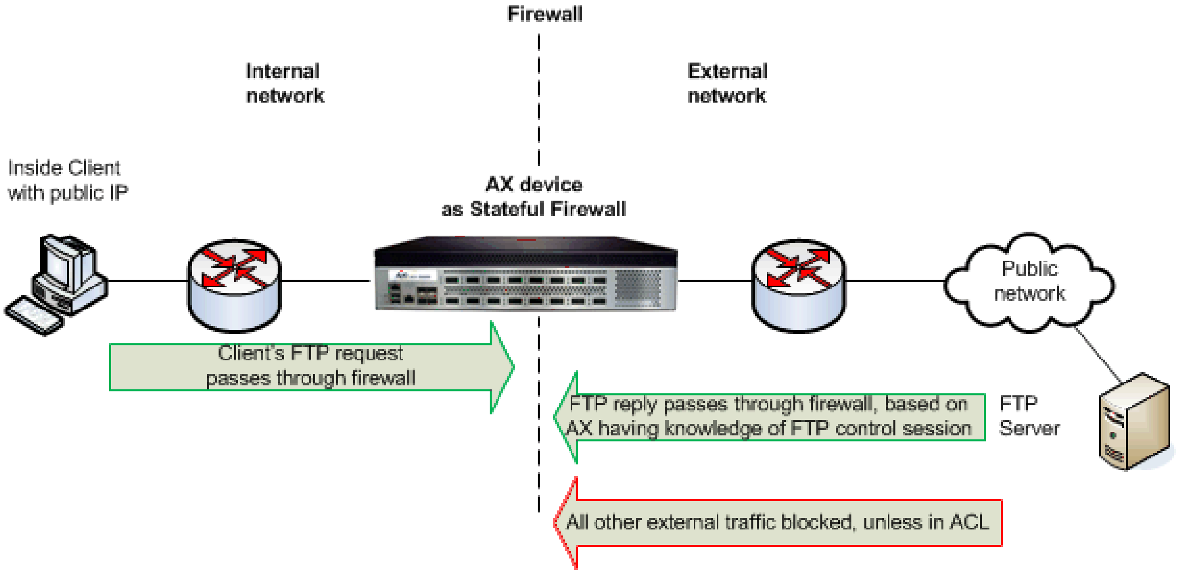
Firewall
Internal network
External network
Inside Client with public IP
AX device as Stateful Firewall
Public network
Client's FTP request passes through firewall
FTP reply passes through firewall, based on AX having knowledge of FTP control session
FTP Server
All other external traffic blocked, unless in ACL

You can enable stateful firewall support for one or more of the following ALG protocols:

- File Transfer Protocol (FTP)
- Trivial File Transfer Protocol (TFTP)
- Real Time Streaming Protocol (RTSP)
- Point-to-Point Tunneling Protocol (PPTP) Generic Routing Encapsulation (GRE)
- Session Initiation Protocol (SIP)

Stateful firewall support is disabled by default. If you enable stateful firewall support without specifying a particular port, then endpoint-independent filtering (EIF) is enabled on all ports (1-65535).

Stateful Firewall Timers

The following timers apply to stateful firewall sessions:

- Idle timeout – Number of seconds a stateful firewall session can remain idle before the AX device terminates the session. You can specify 60-15000 seconds. The default is 300 seconds.
- Session Traversal Utilities for NAT (STUN) timeout – Number of minutes for EIF. You can specify 0-60 minutes. The default is 2 minutes.
- SYN timeout – Amount of time the session stays alive before the TCP handshake is completed and the session is established. You can specify 2-30 seconds. The default is 4 seconds. (Note that the second session can remain in a half-open state before being deleted.)

Configuration

To configure stateful firewall support:

1. Globally enable the feature.
2. Enable the stateful firewall feature on the inside and outside ports.
3. Enable ALG support for individual protocols.
4. (If applicable) Enable EIF. Endpoint-independent-filtering (EIF) allows an inside user to circumvent the access-list by opening a specified TCP or UDP port for an outside user to respond.
5. (Optional) Change the Session Traversal Utilities for NAT (STUN) timeout for transparent sessions.

USING THE GUI

1. Select Config Mode > Service > Stateful Firewall. The Stateful Firewall configuration page appears, as shown in Figure 34.

FIGURE 34 Config Mode > Service > Stateful Firewall

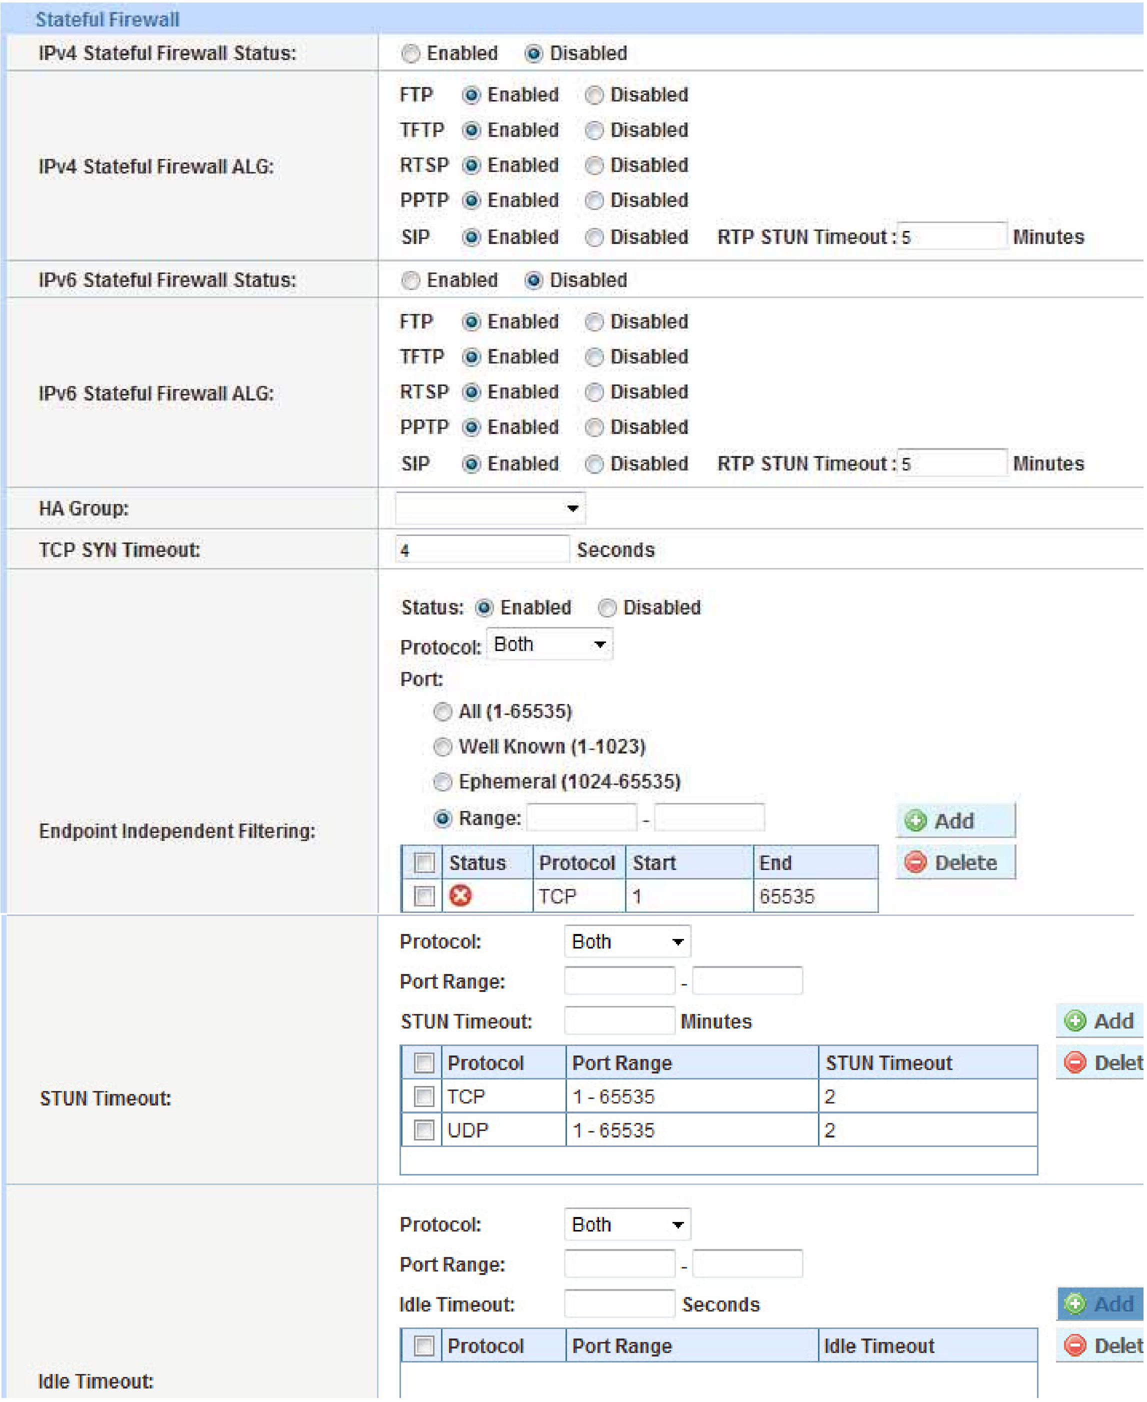

Stateful Firewall

IPv4 Stateful Firewall Status: ○ Enabled ● Disabled

IPv4 Stateful Firewall ALG:

- FTP ● Enabled ○ Disabled
- TFTP ● Enabled ○ Disabled
- RTSP ● Enabled ○ Disabled
- PPTP ● Enabled ○ Disabled
- SIP ● Enabled ○ Disabled RTP STUN Timeout : 5 Minutes

IPv6 Stateful Firewall Status: ○ Enabled ● Disabled

IPv6 Stateful Firewall ALG:

- FTP ● Enabled ○ Disabled
- TFTP ● Enabled ○ Disabled
- RTSP ● Enabled ○ Disabled
- PPTP ● Enabled ○ Disabled
- SIP ● Enabled ○ Disabled RTP STUN Timeout : 5 Minutes

HA Group:

TCP SYN Timeout: 4 Seconds

Endpoint Independent Filtering:

Status: ● Enabled ○ Disabled

Protocol: Both

Port:

- ○ All (1-65535)
- ○ Well Known (1-1023)
- ○ Ephemeral (1024-65535)
- ● Range: -

□	Status	Protocol	Start	End
□	☒	TCP	1	65535

STUN Timeout:

Protocol: Both

Port Range: -

STUN Timeout: Minutes

□	Protocol	Port Range	STUN Timeout
□	TCP	1 - 65535	2
□	UDP	1 - 65535	2

Idle Timeout:

Protocol: Both

Port Range: -

Idle Timeout: Seconds

□	Protocol	Port Range	Idle Timeout

2. Configure the applicable trigger options. (See “Global Stateful Firewall Options” on page 234.)

TABLE 12 Global Stateful Firewall Options

Parameter	Description	Default
IPv4 Stateful Firewall Status:	Enabled or Disabled	Disabled
IPv4 Stateful Firewall ALG:	• File Transfer Protocol (FTP) • Trivial File Transfer Protocol (TFTP) • Real Time Streaming Protocol (RTSP) • Point-to-Point Tunneling Protocol (PPTP) Generic Routing Encapsulation (GRE) • Session Initiation Protocol (SIP) RTP STUN Timeout : 2-10	All ALG protocols are enabled by default. Default: 5 minutes
IPv6 Stateful Firewall Status:	Enabled or Disabled	Disabled
IPv6 Stateful Firewall ALG:	• File Transfer Protocol (FTP) • Trivial File Transfer Protocol (TFTP) • Real Time Streaming Protocol (RTSP) • Point-to-Point Tunneling Protocol (PPTP) Generic Routing Encapsulation (GRE) • Session Initiation Protocol (SIP) RTP STUN Timeout : 2-10	All ALG protocols are enabled by default. Default: 5 minutes
HA Group:	HA Group ID for stateful firewall (IPv4 and IPv6).Range is 1-31.	None
TCP SYN Timeout:	2-30 seconds	Default: 4 seconds
Endpoint Independent Filtering:		
Status:	Enabled or Disabled	
Protocol:	• Both • TCP • UDP	Both
Port:	• All (1-65535) • Well Known (1-1023) • Ephemeral (1024-65535) • Range (1-65535)	
STUN Timeout:		
Protocol:	• Both • TCP • UDP	Both
Port Range:	1-65535	
STUN Timeout:	0-60 minutes	

TABLE 12 Global Stateful Firewall Options

Parameter	Description	Default
Idle Timeout:		
Protocol:	• Both • TCP • UDP	Both
Port Range:	1-65535	
Idle Timeout:	60-15000 seconds	

3. Next, select Config Mode > Network > Interface > LAN, and select the desired data interface on the AX device for which you would like to enable the stateful firewall feature (internal).

   A window similar to the one shown below appears.

FIGURE 35 Config Mode > Network > Interface > LAN

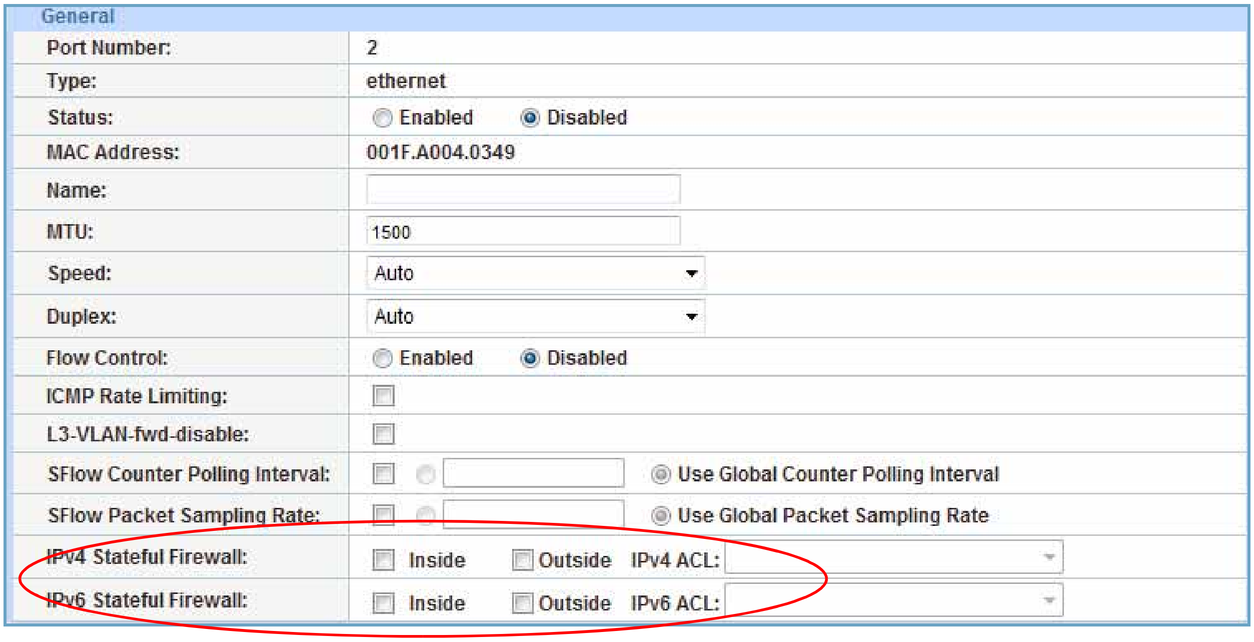

General	
Port Number:	2
Type:	ethernet
Status:	○ Enabled ● Disabled
MAC Address:	001F.A004.0349
Name:	
MTU:	1500
Speed:	Auto
Duplex:	Auto
Flow Control:	○ Enabled ● Disabled
ICMP Rate Limiting:	□
L3-VLAN-fwd-disable:	□
SFlow Counter Polling Interval:	□ ○ ● Use Global Counter Polling Interval
SFlow Packet Sampling Rate:	□ ○ ● Use Global Packet Sampling Rate
IPv4 Stateful Firewall:	□ Inside □ Outside IPv4 ACL:
IPv6 Stateful Firewall:	□ Inside □ Outside IPv6 ACL:

4. In the *IPv4/IPv6 Stateful Firewall* section, select the **Inside** checkbox to enable stateful firewall on the inside (private) port.

5. Navigate to the other data interface for which you would like to enable the stateful firewall feature (public), and then select the **Outside** checkbox (for IPv4 or IPv6) to enable stateful firewall on the outside port.

   The IPv4 or IPv6 ACL drop-down menu is no longer grayed-out. Click the drop-down menu and select the appropriate access control list. The

ACL is applied when someone from the outside tries to initiate a connection to the inside user.

6. Click OK to save your changes.

Using the CLI

All configuration commands for stateful firewall support are entered at the global configuration level of the CLI.

1. To globally enable stateful firewall support, use the following command:

   ```
   ip stateful-firewall enable
   ```

2. To enable stateful firewall support for a logical interface, use the following command:

   ```
   [no] {ip | ipv6} stateful-firewall
   {inside | outside [access-list num]}
   ```

   If you do not specify an access list for the outside interface, the behavior is to deny all Layer 3 traffic coming from the outside interface.

3. To enable or disable ALG support for individual protocols, use the following command:

   ```
   [no] ip stateful-firewall alg
   [ftp | tftp | rtsp | pptp | sip]
   {enable | disable}
   ```

   If you do not specify a protocol, ALG support is enabled for all the protocols.

4. (Optional) To enable EIF, use either of the following commands:

   ```
   [no] ip stateful-firewall
   endpoint-independent-filtering
   {enable | disable}
   ```

   Use this command to enable or disable EIF for ephemeral, well-known, or a range of ports.

   ```
   [no] ip stateful-firewall
   endpoint-independent-filtering
   {tcp | udp} {enable | disable}
   [ephemeral | well-known |
     port-num [to port-num]]
   ```

   The **ephemeral** option enables or enables EIF on ports 1024-65535.

   The **well-known** option enables or disables EIF on ports 1-1023.

The **tcp** | **udp** and *port-num* [**to** *port-num*] options enable or disable EIF on a specified port or ports 1-65535.

5. (Optional) To change stateful firewall timers, use the following commands. For valid ranges and defaults, see "Stateful Firewall Timers" on page 232.

   This command configures the idle timeout.

   ```
   [no] ip stateful-firewall {tcp | udp}
   idle-timeout [port portnum [to portnum]] seconds
   ```

   You can use either of the commands below to configure the STUN timeout for EIF sessions.

   ```
   [no] ip stateful-firewall {tcp | udp}
   stun-timeout [port portnum [to portnum]] minutes
   ```

   ```
   [no] ip stateful-firewall stun-timeout
   [port portnum [to portnum]]
   minutes
   ```

   The following command configures a 10-second SYN timeout for half-open TCP connections.

   ```
   [no] ip stateful-firewall tcp syn-timeout 10
   ```

Displaying Stateful Firewall Information

To display stateful firewall statistics, use the following command:

```
show ip stateful-firewall statistics
```

To display ALG statistics for stateful firewall support, use the following command:

```
show ip stateful-firewall alg
{ftp | pptp | rtsp | sip | tftp} statistics
```

Clearing Stateful Firewall Statistics

To clear stateful firewall statistics, use the following commands:

```
clear ip stateful-firewall statistics
```

```
clear ip stateful-firewall alg
{ftp | pptp | rtsp | sip | tftp} statistics
```

CLI Example

The following example globally enables the stateful firewall feature and sets up the access list. Then, an inside stateful firewall is enabled on private VE port 21, and an outside stateful firewall is enabled on public VE port 22, and access list “101” is applied.

```
AX(config)#ip stateful-firewall enable

AX(config)#access-list 101 permit tcp any any log
AX(config)#access-list 101 permit udp any any log

AX(config)#interface ve 21
AX(config-if:ve21)#ip address 10.10.10.33 255.255.255.0
AX(config-if:ve21)#ip stateful-firewall inside

AX(config)#interface ve 22
AX(config-if:ve22)#ip address 20.20.20.33 255.255.255.0
AX(config-if:ve22)#ip stateful-firewall outside access-list 101
AX(config-if:ve22)#exit
```

Management Security Features

In addition to basic security provided by login and enable passwords, AX Series devices support the following advanced management access security features:

- Multiple admin accounts with distinct levels of access – see "Configuring Additional Admin Accounts" on page 240
- Admin account lockout in response to excessive invalid passwords – see "Configuring Admin Lockout" on page 247
- Admin access control based on management interface (GUI or CLI) – see "Admin Access Control Based on the Management Interface" on page 249
- Interface-level access control for individual management access types (Telnet, SSH, and so on) – see "Securing Admin Access by Ethernet" on page 250
- Public key authentication for SSH management access – see "SSH Public Key Authentication for SSH Management Access" on page 255
- Web access features for securing access through the GUI – see "Changing Web Access Settings" on page 257
- Command auditing – see "Command Auditing" on page 259
- Authentication, Authorization, and Accounting (AAA) for remotely managed access security – see "Configuring AAA for Admin Access" on page 261

The following sections describe these features and show how to configure them.

Note: If you have not already changed the default "admin" password and the enable password, A10 Networks recommends that you do so now, before implementing security options described in this chapter.

Configuring Additional Admin Accounts

The AX device comes with one admin account, “admin”, by default. The “admin” account has global Read Write privileges.

The admin account, and other admin accounts with global Read Write privileges, can configure additional admin accounts. For each admin account, the following settings can be configured:

- Username and password
- IP host or subnet address from which the admin is allowed to log on
- Management interfaces the admin is allowed to use (CLI or GUI)
- GUI access Role (read-write privileges for GUI page access)
- Account state (enabled or disabled)

Configuring an Admin Account

To configure an admin account, use either of the following methods.

USING THE GUI

1. Select Config > System > Admin > Administrator.
2. Click Add. The Administrator section appears.
3. Enter the name in the Administrator Name field.
4. Enter the password for the new admin account in the Password and Confirm Password fields.
5. To restrict login access by the admin to a specific host or subnet:
   a. Enter the address in the Trusted Host IP Address field.
   b. To restrict access to just a single host, edit the value in the Netmask field to 255.255.255.255.
   c. To restrict access to a subnet, edit the value in the Netmask field to the subnet mask for the subnet.

Note: To allow access from any host, leave the Trusted Host IP Address and Netmask fields blank.

6. Select the role from the Role drop-down list. The role defines the read or write access allowed to the admin for each GUI page. (See “GUI Access Roles” on page 242.)

7. To restrict access to specific management interfaces, click the checkboxes next to Access Type.

8. Leave the Status enabled.

9. Click OK. The new admin appears in the Admin table.

Note: For information about the SSH Key File section, see “SSH Public Key Authentication for SSH Management Access” on page 255.

FIGURE 36 Config > Admin > Admin

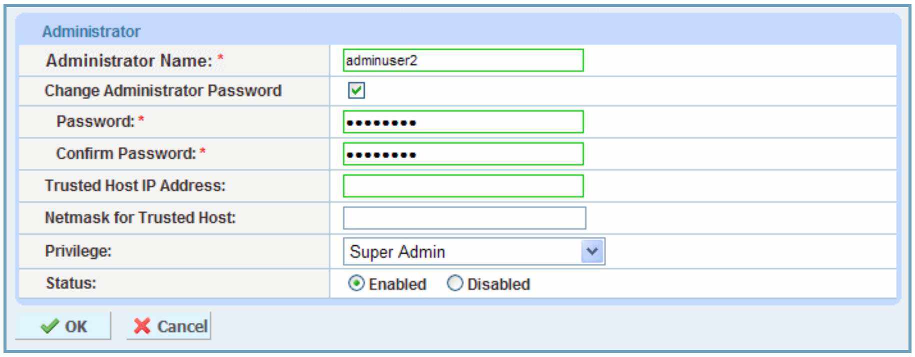

Administrator

Administrator Name: *	adminuser2
Change Administrator Password	☑
Password: *	•••••••••
Confirm Password: *	•••••••••
Trusted Host IP Address:	
Netmask for Trusted Host:	
Privilege:	Super Admin
Status:	● Enabled ○ Disabled

FIGURE 37 Config > Admin - new admin added

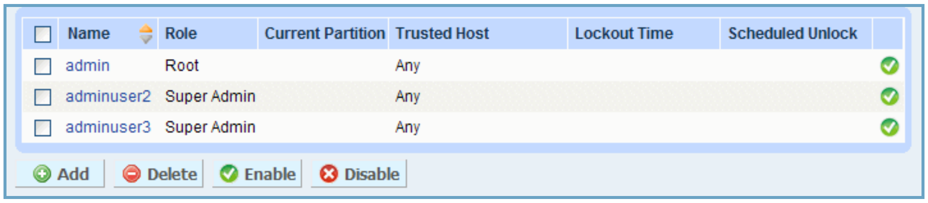

□	Name	Role	Current Partition	Trusted Host	Lockout Time	Scheduled Unlock	
□	admin	Root		Any			✓
□	adminuser2	Super Admin		Any			✓
□	adminuser3	Super Admin		Any			✓

GUI Access Roles

Admin roles enable you to restrict the GUI options an admin is authorized to use. For each GUI page, the admin role specifies whether the admin is allowed to access (view) the page. If the admin is allowed to access the page, the role specifies whether the admin has read-only or read-write privileges for the page.

You can assign an admin to a preconfigured role or a custom role that you configure. You also can customize the preconfigured roles. Table 13 on page 242 lists the preconfigured roles and the types of GUI page access allowed by each one.

Table Column Descriptions

In the Role and Access column, the numbers indicate the roles.

Note: If you configure GUI-based access in RADIUS or TACACS+, these are the numbers to use when specifying a preconfigured role.

- 1 – ReadOnlyAdmin
- 2 – ReadWriteAdmin
- 3 – SystemAdmin
- 4 – NetworkAdmin
- 5 – NetworkOperator
- 6 – SlbServiceAdmin
- 7 – SlbServiceOperator

The following letters indicate the access privileges for the GUI page:

- **R** – Read-only
- **W** – Read-write
- **H** – Hidden (page can not be viewed by the admin)

TABLE 13 Preconfigured GUI Access Roles

GUI Page*	Role and Access						
	1	2	3	4	5	6	7
Monitor Pages							
Monitor > Overview > Summary	R	R	R	R	R	R	R
Monitor > Overview > Status	R	R	H	H	H	R	R
Monitor > Overview > Statistics	R	R	H	H	H	R	R
Monitor > Overview > Performance	R	R	H	H	H	R	R

TABLE 13 Preconfigured GUI Access Roles (Continued)

GUI Page*	Role and Access						
	1	2	3	4	5	6	7
Monitor > Service > SLB	R	R	H	H	H	R	R
Monitor > Service > Health Monitor	R	R	H	H	H	R	R
Monitor > Service > IP Source NAT	R	R	H	H	H	R	R
Monitor > Service > LSN	R	R	H	H	H	R	R
Monitor > Service > NAT64	R	R	H	H	H	R	R
Monitor > Service > DS-Lite	R	R	H	H	H	R	R
Monitor > Service > Stateless NAT46	R	R	H	H	H	R	R
Monitor > Service > 6RD	R	R	H	H	H	R	R
Monitor > Service > Fixed NAT	R	R	H	H	H	R	R
Monitor > Service > Stateful Firewall	R	R	H	H	H	R	R
Monitor > Service > NetFlow Monitor	R	R	H	H	H	R	R
Monitor > Service > SFlow	R	R	H	H	H	R	R
Monitor > Service > Session	R	R	H	H	H	R	R
Monitor > Service > Application	R	R	H	H	H	R	R
Monitor > Network > Interface	R	R	H	R	R	H	H
Monitor > Network > Trunk	R	R	H	R	R	H	H
Monitor > Network > LACP	R	R	H	R	R	H	H
Monitor > Network > VLAN	R	R	H	R	R	H	H
Monitor > Network > ACL	R	R	H	R	R	H	H
Monitor > Network > ARP	R	R	H	R	R	H	H
Monitor > Network > Route	R	R	H	R	R	H	H
Monitor > System > Admin	R	R	R	H	H	H	H
Monitor > System > Logging	R	R	R	H	H	H	H
Monitor > HA > Group	R	R	H	H	H	R	R
Monitor > HA > Status	R	R	H	H	H	R	R
Config Pages							
Config > Get Started > Basic System	R	W	H	W	R	H	H
Config > Service > SLB	R	W	H	H	H	W	R
Config > Service > Template	R	W	H	H	H	W	R
Config > Service > Health Monitor	R	W	H	H	H	W	R
Config > Service > IP Source NAT	R	W	H	H	H	W	R
Config > Service > LSN	R	W	H	H	H	W	R
Config > Service > NAT64	R	W	H	H	H	W	R
Config > Service > DS-Lite	R	W	H	H	H	W	R
Config > Service > Stateless NAT46	R	W	H	H	H	W	R
Config > Service > 6RD	R	W	H	H	H	W	R
Config > Service > Stateful Firewall	R	W	H	H	H	W	R
Config > Service > NetFlow Monitor	R	W	H	H	H	W	R
Config > Service > SFlow	R	W	H	H	H	W	R
Config > Network > Interface	R	W	H	W	R	H	H

TABLE 13 Preconfigured GUI Access Roles (Continued)

GUI Page*	Role and Access						
	1	2	3	4	5	6	7
Config > Network > Trunk	R	W	H	W	R	H	H
Config > Network > LACP	R	W	H	W	R	H	H
Config > Network > VLAN	R	W	H	W	R	H	H
Config > Network > ACL	R	W	H	W	R	H	H
Config > Network > ARP	R	W	H	W	R	H	H
Config > Network > Route	R	W	H	W	R	H	H
Config > Network > DNS	R	W	H	W	R	H	H
Config > Network > ICMP Rate Limiting	R	W	H	W	R	H	H
Config > Network > BPDU-Fwd-Group	R	W	H	W	R	H	H
Config > System > Settings > Web	R	W	W	H	H	H	H
Config > System > Settings > Web Certificate	R	W	W	H	H	H	H
Config > System > Settings > Terminal	R	W	W	H	H	H	H
Config > System > Settings > Log	R	W	W	W	R	H	H
Config > System > Settings > General	R	W	W	H	H	H	H
Config > System > Settings > Boot	R	W	W	H	H	H	H
Config > System > Settings > Action	R	W	W	H	H	H	H
Config > System > Admin	R	W	W	H	H	H	H
Config > System > Access Control	R	W	W	W	R	H	H
Config > System > Time	R	W	W	W	R	H	H
Config > System > SNMP	R	W	W	W	R	H	H
Config > System > Maintenance	R	W	W	H	H	H	H
Config > System > Console	R	W	W	H	H	H	H
Config > System > Config File	R	W	W	H	H	H	H
Config > System > Diagnosis	R	W	W	H	H	H	H
Config > HA > Setting	R	W	H	H	H	W	R
Config > HA > Config Sync	R	W	H	H	H	W	R

*. In some cases where the same access privileges apply to all pages at a given GUI level, only the high-level page name is listed in this table. However, access is configurable on an individual page basis for all GUI pages.

Using the CLI

1. Log on through the CLI and access the global configuration level.

2. Enter the following command to create the new admin account:

   ```
   [no] admin admin-username
   ```

   This command changes the CLI to the configuration level for the new account.

3. Use the following commands to complete the configuration:

   ```
   password string
   ```

trusted-host *ipaddr* {*subnet-mask* | */mask-length*}

access {**cli** | **web**}

privilege *priv-level*

The **password** command assigns the password, which can be 1-63 characters. The default is “a10”.

The **trusted-host** command specifies the host or subnet from which the admin is allowed to log in. The default is 0.0.0.0 /0 (any host or subnet).

The **access** command specifies the management interfaces the admin is allowed to access. By default, access is allowed to all interfaces (CLI and GUI).

The **privilege** command specifies the privileges granted to the admin account:

- **read** – The admin can access the User EXEC and Privileged EXEC levels of the CLI only. This is the default.
- **write** – The admin can access all levels of the CLI.

Note: For information about the **ssh-pubkey** command, see “SSH Public Key Authentication for SSH Management Access” on page 255.

Note: To restrict write access to specific configuration areas, see “Configuring AAA for Admin Access” on page 261.

4. To verify the new admin account, enter the following command:
   show admin

CLI Examples

The following commands add admin “adminuser2” with password “12345678” and read-write privilege:

```
AX(config)#admin adminuser2
AX(config-admin:adminuser2)#password 12345678
AX(config-admin:adminuser2)#privilege write
AX(config-admin:adminuser2)#show admin
UserName                      Status    Privilege Partition
-------------------------------------------------------
admin                      Enabled      R/W
adminuser2                 Enabled      R/W
```

The following commands add admin “adminuser3” with password “abcdefgh” and read-write privilege, and restrict login access to the 10.10.10.x subnet only:

```
AX(config)#admin adminuser3
AX(config-admin:adminuser3)#password abcdefgh
AX(config-admin:adminuser3)#privilege write
AX(config-admin:adminuser3)#trusted-host 10.10.10.0 /24
AX(config-admin:adminuser3)#show admin
UserName                      Status    Privilege Partition
-------------------------------------------------------
admin                      Enabled      R/W
adminuser2                 Enabled      R/W
adminuser3                 Enabled      R/W

AX(config-admin:adminuser3)#show admin adminuser3 detail
  User Name             ...... adminuser3
  Status                ...... Enabled
  Privilege             ...... R/W
  Partition             ......
  Access type             .....cli web
  GUI role             ......
  Trusted Host(Netmask) ...... 10.10.10.0(255.255.255.0)
  Lock Status           ...... No
  Lock Time             ......
  Unlock Time           ......
  Password Type         ...... Encrypted
  Password              ...... $1$6334ba07$CKbWL/LuSNdY12kcE.KdS0
```

Deleting an Admin Account

An admin with Root privileges can delete other admin accounts. If you need to delete an admin account:

1. Display the admin session table to determine whether the admin has any active admin sessions.

2. Clear any sessions the admin has open.

3. Delete the admin account.

Note: To delete an admin account, you first must terminate any active sessions the admin account has open. The account is not deleted if there are any open sessions for the account.

USING THE GUI

1. To display the admin session table, select Monitor > System > Admin.
2. To clear an admin session, click on the checkbox next to the session to select it, then click Delete.
3. To delete the admin account:
   a. Select Config > System > Admin.
   b. Click on the checkbox next to the admin name.
   c. Click Delete.

USING THE CLI

1. To display the admin session table, use the following command at the Privileged EXEC level or any configuration level:

   `show admin session`

2. To clear an admin session, use the following command at the Privileged EXEC level or any configuration level:

   `clear admin session` *`session-id`*

   The *session-id* is the ID listed in the ID column of the show admin session output.

3. To delete the admin account, use the following command at the global configuration level:

   `no admin` *`admin-username`*

Configuring Admin Lockout

By default, there is no limit to the number of times an incorrect password can be entered with an admin account to attempt access. You can enable the AX device to lock admin accounts for a specific period of time following a specific number of invalid passwords entered for the account.

Table 14 lists the admin lockout parameters you can configure.

TABLE 14 Admin Lockout Parameters

Parameter	Description	Default
Feature state	Controls whether admin accounts can be locked.	Disabled
Threshold	Number of failed login attempts allowed for an admin account before it is locked.	5

TABLE 14 Admin Lockout Parameters (Continued)

Parameter	Description	Default
Reset time	Number of minutes the AX device remembers a failed login attempt. For an account to be locked, greater than the number of failed login attempts specified by the threshold must occur within the reset time.	10 minutes
Duration	Number of minutes a locked account remains locked. To keep accounts locked until you or another authorized administrator unlocks them, set the value to 0.	10 minutes

To configure admin lockout, use either of the following methods.

USING THE GUI

To enable the lockout feature:

1. Select Config > System > Admin.
2. Select Lockout Policy on the menu bar.
3. Select Enable Administrator lockout Feature.
4. Optionally, change lockout settings. (For descriptions, see Table 14 on page 247.)
5. Click OK.

To view lockout status or manually unlock a locked account:

1. Select Monitor > System > Admin.
2. Select the admin account.
3. Click Unlock.

USING THE CLI

1. Log on through the CLI and access the global configuration level.
2. Optionally, enter the following commands to change lockout settings:

   ```
   admin lockout threshold number
   admin lockout duration minutes
   admin lockout reset-time minutes
   ```

   (For descriptions, see Table 14 on page 247.)

3. Use the following command to enable admin lockout:

   `admin lockout enable`

To view lockout status or manually unlock a locked account:

1. Log on through the CLI and access the global configuration level.

2. Enter the following command to view the lockout status of an admin account:

   show admin *admin-username* **detail**

3. Enter the following command to access the configuration level for the admin account:

   admin *admin-username*

4. Use the following command to unlock the account:

   `unlock`

Admin Access Control Based on the Management Interface

You can specify the AX management interfaces individual admins are allowed to access. In this release, you can deny an admin from accessing the AX device through one or more of the following management interfaces:

- CLI
- GUI

By default, admins are allowed to use any of the management interfaces. To deny access through specific management interfaces, use either of the following methods.

Using the GUI

1. Select Config > Settings > Admin > Administrator.

2. Click on the admin name or click Add to add a new one.

3. If configuring a new admin, enter the username and password.

4. Next to Access Type, select the interfaces the admin is allowed to access.

5. Click OK.

Note: For information about the admin roles listed in the Role drop-down list, see “Admin Access Control Based on the Management Interface” on page 249. For information about the SSH Key File option, see “SSH Public Key Authentication for SSH Management Access” on page 255.

USING THE CLI

To deny or permit an admin to access the AX device through a specific management interface, use the following command at the configuration level for the admin account:

```
[no] access {cli | web}
```

The following commands deny management access by admin “admin2” using the CLI:

```
AX(config)#admin admin2
AX(config-admin:admin2)#no access cli
```

Securing Admin Access by Ethernet

By default, certain types of management access through the AX device’s Ethernet interfaces are blocked. Table 15 lists the default settings for each management service.

TABLE 15 Default Management Access

Management Service	Ethernet Management Interface	Ethernet and VE Data Interfaces
SSH	Enabled	Disabled
Telnet	Disabled	Disabled
HTTP	Enabled	Disabled
HTTPS	Enabled	Disabled
SNMP	Enabled	Disabled
Ping	Enabled	Enabled

You can enable or disable management access, for individual access types and interfaces. You also can use an Access Control List (ACL) to permit or deny management access through the interface by specific hosts or subnets.

To set management access through Ethernet interfaces, use either of the following methods.

Notes Regarding Use of ACLs

If you use an ACL to secure management access, the action in the ACL rule that matches the management traffic's source address is used to permit or deny access, regardless of other management access settings.

For example, if you disable Telnet access to a data interface, but you also enable access to the interface using an ACL with permit rules, the ACL permits Telnet (and all other) access to the interface, for traffic that matches the permit rules in the ACL.

If you want certain types of management access to be disabled on an interface, do not use a permit ACL to control management access to the interface.

Each ACL has an implicit **deny any any** rule at the end. If the management traffic's source address does not match a permit rule in the ACL, the implicit **deny any any** rule is used to deny access.

On data interfaces, you can disable or enable access to specific services and also use an ACL to control access. However, on the management interface, you can disable or enable access to specific services ***or*** control access using an ACL, but you can not do both.

USING THE GUI

To change management access settings for interfaces:

1. Select Config > System > Access Control.
2. For each interface (each row), select or de-select the checkboxes for the access types.
3. To use an ACL to control access, select the ACL from the ACL drop-down list in the row for the interface.
4. After selecting the settings for all the interfaces, click OK.

To reset the access settings to the defaults listed in Table 15, click Reset to Default.

USING THE CLI

Disabling Management Access

To disable management access, use either of the following commands at the global configuration level of the CLI:

```
disable-management service
 {all|ssh|telnet|http|https|snmp|ping}
 {management|ethernet port-num [to port-num]|
   ve ve-num [to ve-num]}
```

or

```
disable-management service acl acl-num
 {management|ethernet port-num [to port-num]|
   ve ve-num [to ve-num]}
```

In both commands, the following options specify the interfaces to protect:

- **management** – The out-of-band Ethernet management interface (MGMT)
- **ve** *ve-num* [**to** *ve-num*] – A VE data interface or range of VE data interfaces
- **ethernet** *port-nu*m [**to** *port-num*] – An Ethernet data interface or range of Ethernet data interfaces

In the first command, the following options specify the type of management access you are configuring:

- **all** – Disables access to all the management services listed below.
- **ssh** – Disables SSH access to the CLI.
- **telnet** – Disables Telnet access to the CLI.
- **http** – Disables HTTP access to the management GUI.
- **https** – Disables HTTPS access to the management GUI.
- **snmp** – Disables SNMP access to the AX device’s SNMP agent.
- **ping** – Disables ping replies from AX interfaces. This option does not affect the AX device’s ability to ping other devices.

Note: Disabling ping replies from being sent by the AX device does not affect the device’s ability to ping other devices.

In the second command, the **acl** *acl-id* option specifies an ACL. Management access from any host address that matches the ACL is either permitted or denied, depending on the action (permit or deny) used in the ACL.

CLI Examples:

The following command disables HTTP access to the out-of-band management interface:

```
AX(config)#disable-management service http management
You may lose connection by disabling the http service.
Continue? [yes/no]:yes
```

Enabling Management Access

To enable management access, use either of the following commands at the global configuration level of the CLI:

```
enable-management service
  {all | ssh | telnet | http | https | snmp | ping}
  {management | ethernet port-num [to port-num] |
    ve ve-num [to ve-num]}
```

or

```
enable-management service acl acl-num
  {management | ethernet port-num [to port-num] |
    ve ve-num [to ve-num]}
```

The options are the same as those for the **disable-management** command.

CLI Example:

The following command enables Telnet access to data interface 6:

```
AX(config)#enable-management service telnet ethernet 6
```

Displaying the Current Management Access Settings

To display the management access settings that are currently in effect, enter the following command at any level of the CLI:

```
show management
```

CLI Examples

Here is an example for an AX device that has 10 Ethernet data ports. In this example, all the access settings are set to their default values.

```
AX#show management
     PING   SSH    Telnet HTTP   HTTPS  SNMP   ACL
------------------------------------------------------
mgmt on     on     off    on     on     on     -
1    on     off    off    off    off    off    -
2    on     off    off    off    off    off    -
3    on     off    off    off    off    off    -
4    on     off    off    off    off    off    -
5    on     off    off    off    off    off    -
6    on     off    off    off    off    off    -
7    on     off    off    off    off    off    -
9    on     off    off    off    off    off    -
10   on     off    off    off    off    off    -
ve1  on     off    off    off    off    off    -
```

Here is an example after entering the commands used in the configuration examples above.

```
AX#show management
     PING   SSH    Telnet HTTP   HTTPS  SNMP   ACL
------------------------------------------------------
mgmt on     on     off    off    on     on     -
1    on     off    off    off    off    off    1
2    on     off    off    off    off    off    1
3    on     off    off    off    off    off    1
4    on     off    off    off    off    off    1
5    on     off    off    off    off    off    1
6    on     off    on     off    off    off    1
7    on     off    off    off    off    off    1
9    on     off    off    off    off    off    1
10   on     off    off    off    off    off    1
ve1  on     off    off    off    off    off    -
```

Regaining Access if You Accidentally Block All Access

If you disable the type of access you are using on the interface you are using at the time you enter a **disable-management** command, your management session will end. If you accidentally lock yourself out of the device altogether (for example, if you use the **all** option for all interfaces), you can still access the CLI by connecting a PC to the AX device's serial port.

SSH Public Key Authentication for SSH Management Access

Public key authentication allows an AX admin to log in through SSH without entering a password. When the admin enters their username and presses Enter, the SSH client on the admin's PC sends a signature file for the admin. The AX device compares the signature file to the admin's public key stored on the AX device. If they match, the admin is granted access.

To use public key authentication for management access to the AX device:

1. On the PC from which the admin will access the AX CLI, use the PC's SSH client to generate an RSA key pair for the admin. The key pair consists of a public key and a private key.

   Do not use a passphrase.

Note: In the current release, only the OpenSSH client is supported.

2. Log into the AX device with root or global read-write privileges.

3. Access the configuration level for the admin account.

4. Import ***only*** the public key onto the AX device. (***Do not import the private key onto the AX device.***)

   You can import public keys in separate files or grouped together into a single file.

Note: The "admin" account has root privileges and can manage the public certificates for all admins. Any other admin account can manage only the public key belonging to that admin account.

Using the CLI

To import an SSH public key onto the AX device, use the following command at the configuration level for the admin account:

ssh-pubkey import *url*

The *url* specifies the file transfer protocol, username (if required), and directory path for exporting the public key file.

You can enter the entire URL on the command line or press Enter to display a prompt for each part of the URL. If you enter the entire URL and a password is required, you will still be prompted for the password. To enter the entire URL:

- **tftp:**//*host*/*file*
- **ftp:**//[*user*@]*host*[:*port*]/*file*
- **scp:**//[*user*@]*host*/*file*
- **rcp:**//[*user*@]*host*/*file*

To delete an SSH public key from the AX device, use the following command:

ssh-pubkey delete *num*

The *num* option specifies the key number on the AX device. The key numbers are displayed along with the keys themselves by the **ssh-pubkey list** command. (See below.)

To verify installation of a public key, use the following command:

ssh-pubkey list

CLI Example

The following commands configure public key authentication for “admin2”.

Key Configuration on PC

On the admin’s PC, the following OpenSSH commands configure the public-private key pair:

```
OpenSSHclient$mkdir ~/.ssh
OpenSSHclient$chmod 700 ~/.ssh
OpenSSHclient$ssh-keygen -q -f ~/.ssh/ax_admin2 -t rsa
Enter passphrase (empty for no passphrase): …
Enter same passphrase again: …
```

Note: Do not type “empty” at the passphrase prompts. Just press Enter.

Configuration on AX Device

The following commands import admin2’s public key and verify its presence on the AX device:

```
AX(config)#admin admin2
AX(config-admin:admin2)#ssh-pubkey import scp:
Address or name of remote host []?10.10.10.69
User name []?axadmin2
```

```
Password []?*********
File name [/]?ax_admin2.pem
AX(config-admin:admin2)#ssh-pubkey list
```

Changing Web Access Settings

By default, access to the AX management GUI is enabled and is secure. A valid admin username and password are required to log in.

Table 16 lists the default settings for Web access.

TABLE 16 Default Web Access Settings

Parameter	Description	Default
Auto-redirect	Automatically redirects requests for the unsecured port (HTTP) to the secure port (HTTPS).	Enabled
HTTP server	HTTP server on the AX device.	Enabled
HTTP port	Protocol port number for the unsecured (HTTP) port.	80
HTTPS server	HTTPS server on the AX device.	Enabled
HTTPS port	Protocol port number for the secure (HTTPS) port.	443
Timeout	Number of minutes a Web management session can remain idle before it times out and is terminated by the AX device.	Range: 0-60 minutes To disable the timeout, specify 0. Default: 10 minutes
aXAPI Timeout	Not applicable.	

Note: If you disable HTTP or HTTPS access, any sessions on the management GUI are immediately terminated.

USING THE GUI

1. Select Config > System > Settings.
2. On the menu bar, select Web.
3. Edit the settings you want to change.
4. Click OK.

Note: The Preference section sets the default IP address type (IPv4 or IPv6) for GUI configuration fields that require an IP address. The Preference section does not affect access to the GUI itself.

USING THE CLI

At the global configuration level of the CLI, use the following command:

```
[no] web-service
{
auto-redir |
port protocol-port |
secure-port protocol-port |
server |
secure-server |
timeout-policy idle minutes
}
```

To view Web access settings, use the following command:

```
show web-service
```

CLI EXAMPLE

The following commands disable management access on HTTP and verifies the change:

```
AX(config)#no web-service server
AX(config)#show web-service
AX Web server:
         Idle time:                10 minutes
         Http port:                80
         Https port:               443
         Auto redirect:            Enabled
         Https:                    Enabled
         Http:                     Disabled
```

Command Auditing

You can enable command auditing to log the commands entered by AX admins. Command auditing logs the following types of system management events:

- Admin logins and logouts for CLI and GUI sessions
- Unsuccessful admin login attempts
- Configuration changes. All attempts to change the configuration are logged, even if they are unsuccessful.
- CLI commands at the Privileged EXEC level (if audit logging is enabled for this level)
- HA configuration synchronization

The audit log is maintained in a separate file, apart from the system log.

Note: Backups of the system log include the audit log.

Command auditing is disabled by default. To enable it, use either of the following methods.

Audit Log Examples

The following audit log indicates a change to the image to use for booting, performed using the CLI:

```
Jul 06 2010 23:27:25  admin cli: bootimage hd sec
```

The following audit logs indicate configuration and operational actions related to virtual server "vip1" performed using the GUI:

```
Jun 08 2010 09:06:04 [12] web: [admin] add virtual server [name:vip1, ip:1.1.1.1,
vport1:8001(TCP).] successfully.
Jun 08 2010 09:06:05 [12] web: [admin] edit virtual server [name:vip1, ip:1.1.1.1,
vport1:8001(TCP).] successfully.
Jun 08 2010 09:06:06 [12] web: [admin] disable virtual server [vip1] successfully.
Jun 08 2010 09:06:06 [12] web: [admin] enable virtual server [vip1] successfully.
Jun 08 2010 09:06:07 [12] web: [admin] delete virtual server [vip1] successfully.
```

Using the GUI

To enable command auditing:

1. Select Config > System > Settings > Log.
2. Click to expand the Audit section.

3. Select the audit level:
    - Disabled – Command auditing is disabled.
    - Enabled – Auditing of configuration commands is enabled.
    - Enable Privilege – Auditing of configuration commands ***and*** Privileged EXEC commands is enabled.

4. To modify the maximum number of entries the log can hold, edit the number in the Audit Buffer Size field. You can specify 1000-30000 entries. The default is 20000.

5. Click OK.

To show audit log entries, navigate to the following page:

Monitor > System > Audit > Logging

Using the CLI

To enable command auditing, use the following command at the global configuration level of the CLI:

```
[no] audit enable [privilege]
```

The **privilege** option enables logging of Privileged EXEC commands also. Without this option, only configuration commands are logged.

To specify the number of entries the audit log file can hold, use the following command at the global configuration level of the CLI:

```
[no] audit size num-entries
```

You can specify 1000-30000 entries. When the log is full, the oldest entries are removed to make room for new entries. The default is 20,000 entries.

To show audit log entries, use the following command:

```
show audit
```

Configuring AAA for Admin Access

You can configure the AX device to use remote servers for Authentication, Authorization, and Accounting (AAA) for admin sessions. The AX device supports RADIUS and TACACS+ servers.

Authentication

Authentication grants or denies access based on the credentials presented by the person who is attempting access. Authentication for management access to the AX device grants or denies access based on the admin username and password.

By default, when someone attempts to log into the AX device, the device checks its local admin database for the username and password entered by the person attempting to gain access.

Without additional configuration, the authentication process stops at this point. If the admin username and password are in the local database, the person is granted access. Otherwise, they are denied.

You can configure the AX device to also use external RADIUS or TACACS+ servers for authentication.

You can use TACACS+ *or* RADIUS for external authentication. Only one external authentication method can be used.

Authentication Process

You can specify whether to check the local database or the remote server first. Figure 38 and Figure 39 show the authentication processes used if the AX device is configured to check remote AAA servers (RADIUS or TACACS+) first.

If the RADIUS or TACACS+ server responds, the local database is not checked.

- If the admin name and password are found on the RADIUS or TACACS+ server, the admin is granted access.
- If the admin name and password are not found on the RADIUS or TACACS+ server, the admin is denied access.

Only if there is no response from any RADIUS or TACACS+ server, does the AX device check its local database for the admin name and password.

Username “admin” Always Authenticated Locally By Default

An exception is made for the “admin” account. By default, the AX device always uses local authentication for “admin”. Optionally, you can disable automatic local authentication for “admin”, in which case the authentication process is the same as for other admin accounts.

FIGURE 38 Authentication Process When Remote Authentication Is First (2 remote servers configured) – Example shown is for RADIUS

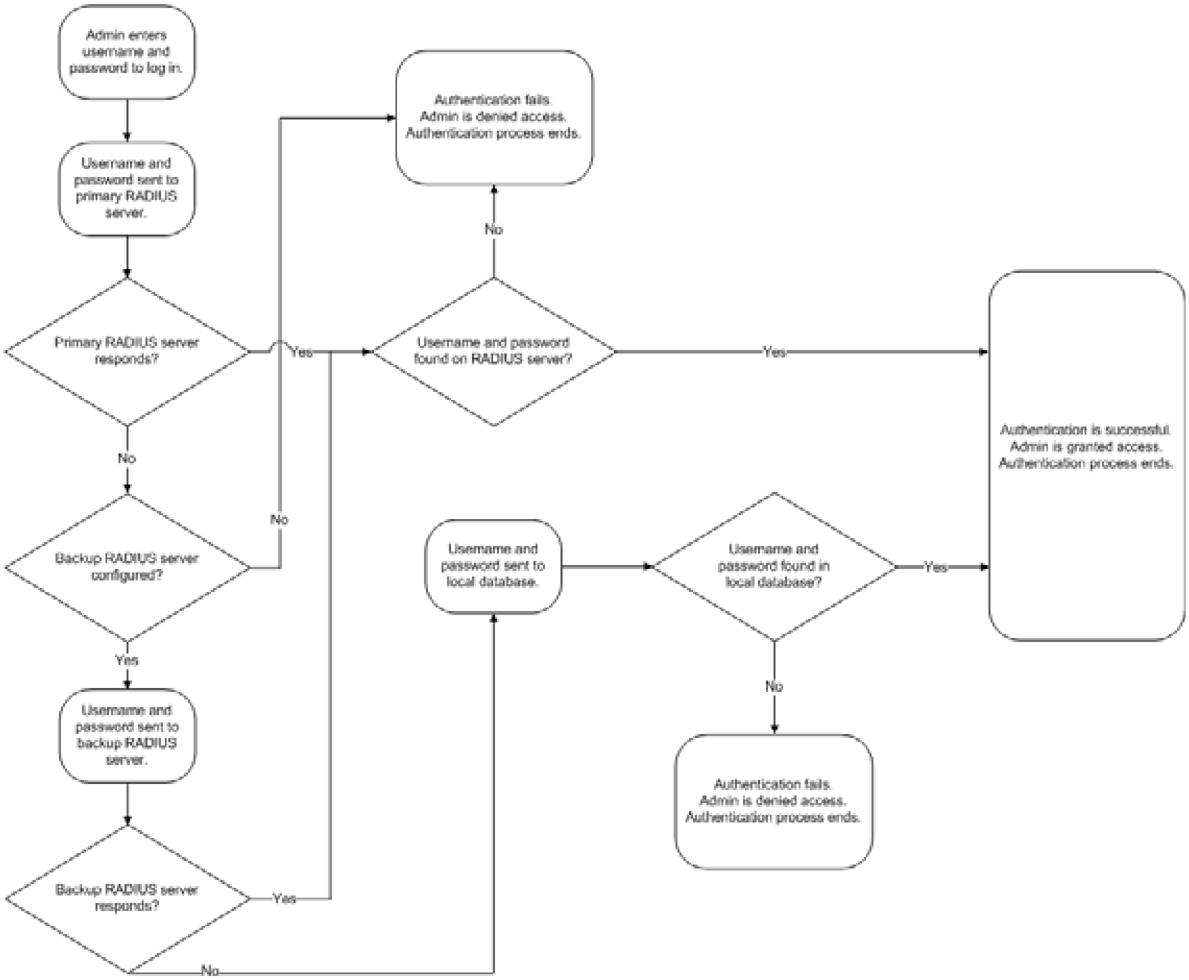
Admin enters username and password to log in.
Username and password sent to primary RADIUS server.
Primary RADIUS server responds?
Yes
No
Backup RADIUS server configured?
No
Yes
Username and password sent to backup RADIUS server.
Backup RADIUS server responds?
Yes
No
Authentication fails. Admin is denied access. Authentication process ends.
No
Username and password found on RADIUS server?
Yes
Authentication is successful. Admin is granted access. Authentication process ends.
Username and password sent to local database.
Username and password found in local database?
Yes
No
Authentication fails. Admin is denied access. Authentication process ends.

FIGURE 39 Authentication Process When Remote Authentication Is First (1 remote server configured) – Example shown is for TACACS+

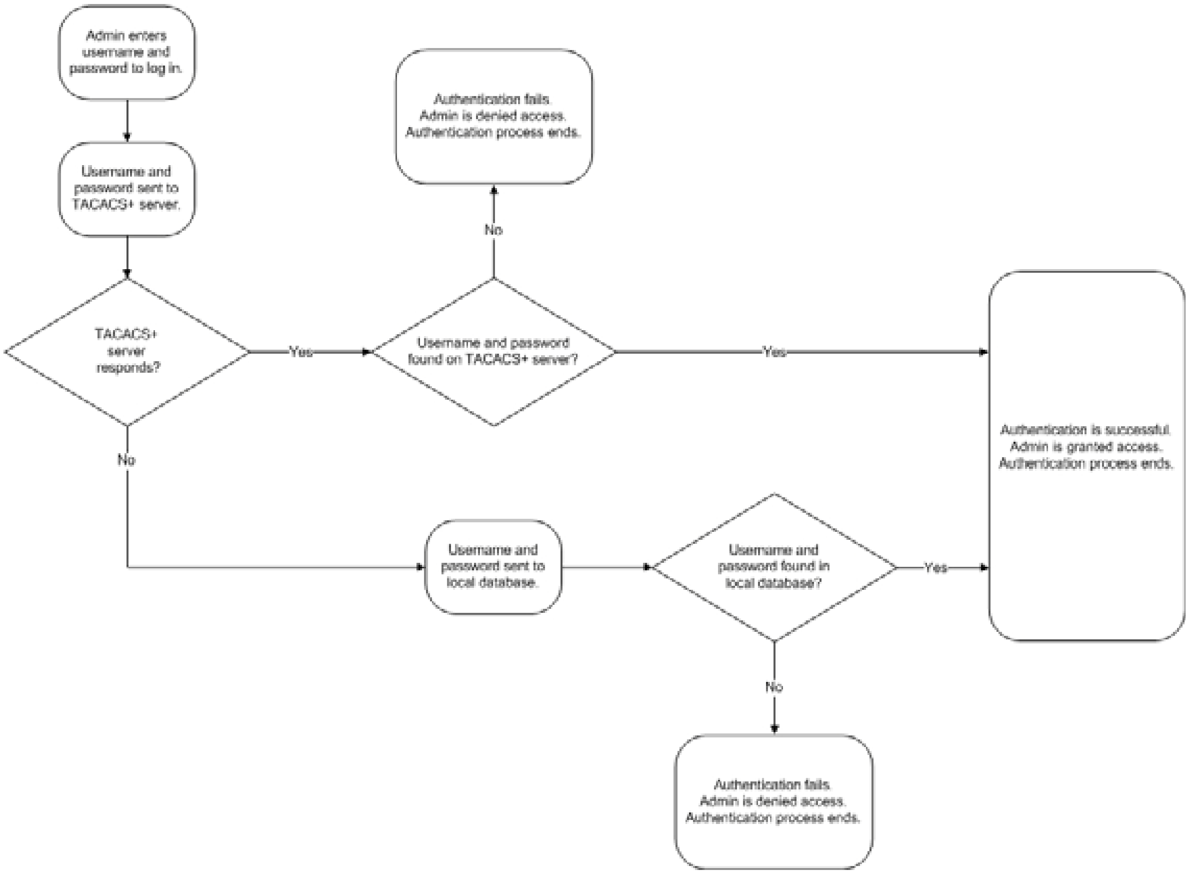
Admin enters username and password to log in.
Username and password sent to TACACS+ server.
TACACS+ server responds?
Yes
No
Username and password found on TACACS+ server?
Yes
No
Authentication fails. Admin is denied access. Authentication process ends.
Authentication is successful. Admin is granted access. Authentication process ends.
Username and password sent to local database.
Username and password found in local database?
Yes
No
Authentication fails. Admin is denied access. Authentication process ends.

Authorization

You can configure authorization based on the following:

- Management interface used by the admin (CLI or GUI)
- GUI page requested by the admin
- CLI command entered by the admin

Authorization Based on Management Interface

You can deny an admin from accessing the AX device through one or more of the following management interfaces:

- CLI
- GUI

By default, admins are allowed to use any of the management interfaces.

RADIUS Configuration for Management Interface Access

To configure authorization based on management interface, use the following A10-Admin-Access-Type values:

- `cli`
- `web`

To authorize access to more than one management interface, use a comma between each value. For example: `cli,web`

If you do not specify an A10-Admin-Access-Type value, access through all three interfaces is permitted.

TACACS+ Configuration for Management Interface Access

To configure authorization based on management interface, use the following AVP:

```
a10-access-type=mgmt-int
```

The *mgmt-int* can be one or more of the following:

- `cli`
- `web`

To authorize access to more than one management interface, use a comma between each value. For example: `a10-access-type=cli,web`

If you do not specify an A10-Admin-Access-Type value, access through all three interfaces is permitted.

Authorization for GUI Access

Each admin account configured on the AX device includes a GUI access role. The GUI access role specifies the GUI pages to which the admin has write privileges, the pages to which the admin has read-only privileges, and if applicable, the pages that are hidden from the admin.

For each GUI page, the admin role specifies whether the admin is allowed to access (view) the page. If the admin is allowed to access the page, the role specifies whether the admin has read-only or read-write privileges for the page.

You can assign an admin to a preconfigured role or a custom role that you configure. You also can customize the preconfigured roles. Table 13 on page 242 lists the preconfigured roles and the types of GUI page access allowed by each one.

Note: The GUI access roles do not apply to admins who log in through the CLI. (See "Authorization for CLI Access" on page 266.)

Note: Also see "RADIUS Authorization Based on Service-Type" on page 269.

RADIUS Configuration for GUI Access Roles

To configure role-based authorization for access to the GUI, use the A10-Admin-Privilege option. For example, to authorize access to the GUI pages associated with the ReadOnlyAdmin role, use the following statement in the admin definition:

```
A10-Admin-Role = "ReadOnlyAdmin"
```

Note: In the current release, the A10-Admin-Privilege option applies only to GUI access. It does not restrict CLI access.

TACACS+ Configuration for GUI Access Roles

To configure role-based authorization for access to the GUI, use the following Attribute Value Pair (AVP):

```
a10-admin-role=role-name
```

Note: In the current release, this AVP applies only to GUI access. It does not restrict CLI access.

Compatibility with Privilege Levels Assigned by RADIUS or TACACS+

It is not required to assign a privilege level to an AX admin on the RADIUS or TACACS+ server used to authenticate the admin. The AX device uses the GUI access role assigned to the admin in the admin's account on the AX device.

However, if a privilege level ***is assigned*** to the admin on the RADIUS or TACACS+ server, that privilege level ***must match*** the role assigned to the admin in the AX configuration. Otherwise, the admin will be denied access.

Table 17 lists the RADIUS and TACACS+ privilege levels that match the GUI access roles.

TABLE 17 RADIUS / TACACS+ Privilege Levels and Matching GUI Access Roles

	Privilege Level		
GUI Access Role	**RADIUS**	**TACACS+**	**Partition Role**
ReadWriteAdmin	2	15	N
SystemAdmin	3	14	N
NetworkAdmin	4	13	N
NetworkOperator	5	12	N
SlbServiceAdmin	6	11	N
SlbServiceOperator	7	10	N
ReadOnlyAdmin	1	0	N
PartitionReadWrite	8	9	Y
PartitionNetworkOperator	9	8	Y
PartitionSlbServiceAdmin	10	7	Y
PartitionSlbServiceOperator	11	6	Y
PartitionReadOnly	12	5	Y

Note: Partitions are not applicable to this release.

Authorization for CLI Access

You can configure the AX device to use external RADIUS or TACACS+ servers to authorize commands entered by admins who log in using the CLI.

Following successful Authentication, the authenticated party is granted access to specific system resources by Authorization. For an AX admin, authorization specifies the CLI levels they can access.

Operational Commands Disabled for Read-Only Admins

Admins who are authenticated by RADIUS or TACACS+, and authorized for read-only access directly to the Privileged EXEC level of the CLI, are not allowed to run certain operational commands.

For these admins, the following operational commands at the Privileged EXEC level of the CLI are disabled:

- **`backup`**
- **`config`**
- **`import`**
- **`locale`**
- **`reboot`**
- **`reload`**
- **`shutdown`**

This includes admins with the ReadOnlyAdmin or PartitionReadOnly role.

RADIUS CLI Authorization

To configure RADIUS CLI Authorization, use the following settings on the RADIUS server:

```
VALUE A10-Admin-Privilege Read-only-Admin 1
VALUE A10-Admin-Privilege Read-write-Admin 2
```

The first line grants access to the User EXEC level and Privileged EXEC level. The admin’s CLI session begins at the User EXEC level. The admin can access the Privileged EXEC level, ***without*** entering an enable password. Access to the configuration level is not allowed.

```
login as: admin3
Using keyboard-interactive authentication.
Password: ********
Last login: Fri Mar 26 20:03:39 2010 from 192.168.1.140

[type ? for help]

AX>enable
AX#
```

The second line grants access to all levels. The admin’s CLI session begins at the Privileged EXEC level.

```
login as: admin4
Using keyboard-interactive authentication.
Password: ********
Last login: Fri Mar 26 20:03:39 2010 from 192.168.1.140

[type ? for help]

AX#
```

Note: Also see “RADIUS Authorization Based on Service-Type” on page 269.

TACACS+ CLI Authorization

To configure TACACS+ CLI Authorization:

- Configure the TACACS+ server to authorize or deny execution of specific commands or command groups.
- Configure the AX device to send commands to the TACACS+ server for authorization before executing those commands.

Note: This authorization process does not apply to admins who log in through the GUI. (See “Authorization for GUI Access” on page 265.)

CLI Access Levels

You can use TACACS+ to authorize an admin to execute commands at one of the following CLI access levels:

- 15(admin) – This is the most extensive level of authorization. Commands at all CLI levels, including those used to configure admin accounts, are sent to TACACS+ for authorization.
- 14(config) – Commands at all CLI levels ***except*** those used to configure admin accounts are sent to TACACS+ for authorization. Commands for configuring admin accounts are automatically allowed.
- 1(priv EXEC) – Commands at the Privileged EXEC and User EXEC levels are sent to TACACS+ for authorization. Commands at other levels are automatically allowed.
- 0 (user EXEC) – Commands at the User EXEC level are sent to TACACS+ for authorization. Commands at other levels are automatically allowed.

Access levels 1-15 grant access to the Privileged EXEC level or higher, without challenging the admin for the enable password. Access level 0 grants access to the User EXEC level only.

Note: Command levels 2-13 are equivalent to command level 1.

Caution: The most secure option is 15(admin). If you select a lower option, for example, 1(priv EXEC), make sure to configure the TACACS+ server to deny any unmatched commands (commands that are not explicitly allowed by the server). Otherwise, unmatched commands, including commands at higher levels, will *automatically be authorized* to execute.

TACACS+ Authorization Debug Options

You can enable the following TACACS+ debug levels for troubleshooting:

- 0x1 – Common system events such as “trying to connect with TACACS+ servers” and “getting response from TACACS+ servers”. These events are recorded in the syslog.
- 0x2 – Packet fields sent out and received by the AX Series device, not including the length fields. These events are written to the terminal.
- 0x4 – Length fields of the TACACS+ packets will also be displayed on the terminal.
- 0x8 – Information about TACACS+ MD5 encryption will be sent to the syslog.

RADIUS Authorization Based on Service-Type

The AX device supports the RADIUS Service-Type attribute. The following attribute values are supported:

- `Service-Type=Login` – Allows access to the EXEC level of the CLI (AX>), and read-only access to the GUI
- `Service-Type=NAS Prompt` – Allows access to the Privileged EXEC level of the CLI (AX#), and read-only access to the GUI
- `Service-Type=Administrative` – Allows access to the configuration level of the CLI [`AX(config)#`], and read-write access to the GUI

By default, if the Service-Type attribute is not used, or the A10 vendor attribute is not used, successfully authenticated admins are authorized for read-only access. You can change the default privilege authorized by RADIUS from read-only to read-write. To change the default access level authorized by RADIUS, use the following command at the global configuration level of the CLI:

```
[no] radius-server default-privilege-read-write
```

Accounting

You can configure the AX device to use external RADIUS or TACACS+ servers for Accounting.

Accounting keeps track of user activities while the user is logged on. For AX admins, you can configure Accounting for the following:

- Login/logoff activity (start/stop accounting)
- Commands

Command Accounting (TACACS+ only)

You can use TACACS+ servers to track attempts to execute commands at one of the following CLI access levels:

- 15(admin) – This is the most extensive level of accounting. Commands at all CLI levels, including those used to configure admin accounts, are tracked.
- 14(config) – Commands at all CLI levels ***except*** those used to configure admin accounts are tracked. Commands for configuring admin accounts are not tracked.
- 1(priv EXEC) – Commands at the Privileged EXEC and User EXEC levels are tracked. Commands at other levels are not tracked.
- 0 (user EXEC) – Commands at the User EXEC level are tracked. Commands at other levels are not tracked.

Note: Command levels 2-13 are equivalent to command level 1.

TACACS+ Accounting Debug Options

The same debug levels that are available for TACACS+ Authorization are also available for TACACS+ Accounting. (See “TACACS+ Authorization Debug Options” on page 269.)

Configuring AAA for Admin Access

To configure AAA for admin access:

1. Prepare the AAA servers:
   - Add admin accounts (usernames and passwords).
   - Add the AX device as a client. For the client IP address, specify the AX IP address.
   - For authorization, configure the following settings for the admin accounts:
     - Specify the management interfaces the admin is allowed to access (CLI or GUI).
     - If using TACACS+, specify the CLI commands or command groups that are to be allowed or denied execution.
     - If using RADIUS, specify the access role for the GUI.

2. To use RADIUS or TACACS+ for Authentication:

   a. Add the RADIUS or TACACS+ server(s) to the AX device.

   b. Add RADIUS or TACACS+ as an authentication method to use along with the local database.

3. Configure Authorization:

   a. Add the TACACS+ or RADIUS servers, if not already added for authentication.

   b. Specify the access level:
      - If using TACACS+, specify the CLI command levels to be authorized.
      - If using RADIUS, specify the GUI access to be authorized.

4. Configure Accounting:

   a. Add the TACACS+ or RADIUS servers, if not already added for Authorization.

   b. Specify whether to track logon/logoff activity. You can track both logons and logoffs, logoffs only, or neither.

   c. Optionally, is using TACACS+, specify the command levels to track.

Configuring Authentication

To configure remote authentication, use either of the following methods.

Using the GUI

1. Select Config > System > Admin > External Authentication > General.

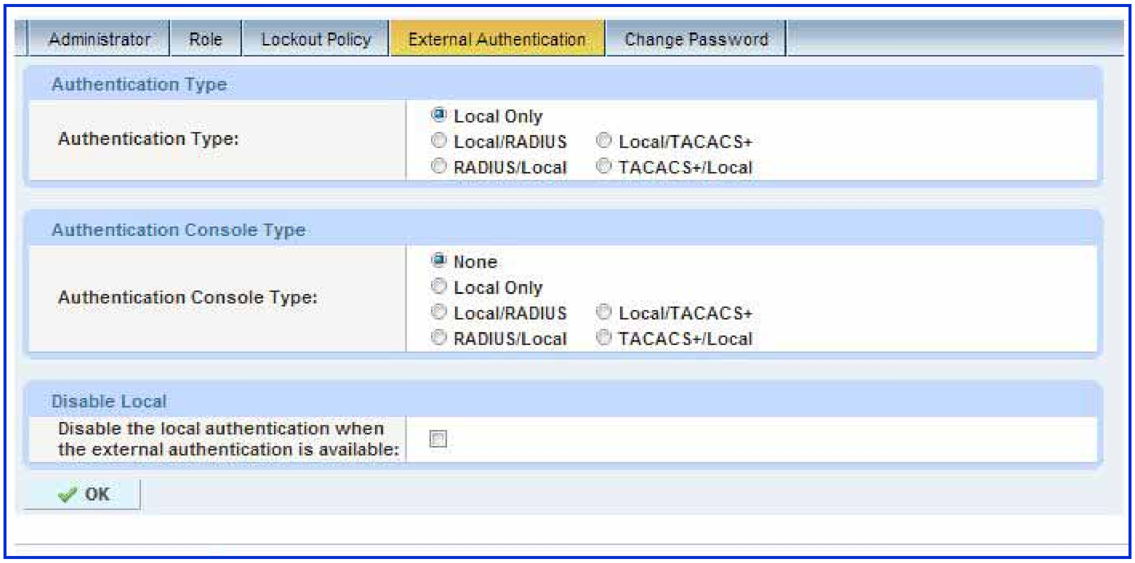

Authentication Type

Authentication Type: ● Local Only ○ Local/RADIUS ○ Local/TACACS+ ○ RADIUS/Local ○ TACACS+/Local

Authentication Console Type

Authentication Console Type: ● None ○ Local Only ○ Local/RADIUS ○ Local/TACACS+ ○ RADIUS/Local ○ TACACS+/Local

Disable Local

Disable the local authentication when the external authentication is available: □

2. Select the authentication methods an d the order in which to use them. You can select one of the following:
   - Local Only – The local admin database on the AX device is used. No external authentication servers are used.
   - Local/RADIUS – The local admin database is consulted first. If the local database does not contain an account for the username entered by the admin, the RADIUS server is consulted.
   - Local/TACACS+ – The local admin database is consulted first. If the local database does not contain an account for the username entered by the admin, the TACACS+ server is consulted.
   - RADIUS/Local – The RADIUS server is consulted first. If the RADIUS server does not respond, the local admin database is consulted.
   - TACACS+/Local – The TACACS+ server is consulted first. If the TACACS+ server does not respond, the local admin database is consulted.

3. Select the authentication console type:

   - Authentication Console Type – Specifies a separate authentication policy for the console (serial) port. The options are the same as Authentication Type.

4. Select the option to disable Local authentication:

   - Disable Local – Disables automatic local authentication of the "admin" account. Without this option, the "admin" account is always authenticated locally, regardless of the authentication configuration used for the other admin accounts.

5. Click OK.

6. Select one of the following options:

   - Config > System > Admin > External Authentication > RADIUS
   - Config > System > Admin > External Authentication > TACACS+

7. To add the primary server, click Server 1 to display the configuration fields for the server.

8. Enter the hostname or IP address of the server in the Hostname field.

9. In the Secret and Confirm Secret fields, enter the shared secret (password) expected by the server when it receives requests.

10. To add a backup server to use if the primary server can not be reached, click Server 2 and enter the configuration information for the server.

11. Click OK.

USING THE CLI

Note: The command syntax shown in this section is simplified to show the required or more frequently used options. For complete syntax information, see the *AX Series CLI Reference*.

1. Use one of the following commands at the global configuration level of the CLI to add the primary server:

   [no] **radius-server host** {*hostname* | *ipaddr*} **secret** *secret-string*

   [no] **tacacs-server host** {*hostname* | *ipaddr*} **secret** *secret-string*

   The *secret-string* is the shared secret (password) expected by the server when it receives requests.

2. To add a backup server to use if the primary server can not be reached, repeat the command, using the backup server's information.

3. Use one of the following commands to specify the order in which to use the authentication methods:

   authentication [**console**] **type** *method1* [*method2*]

   The **console** option applies the authentication settings only to access through the console (serial) port. Without this option, the settings apply to all types of admin access.

   (For more information, see "Authentication Process" on page 261.)

Separate Authentication Policy for the AX Console Port

This release provides a new option that allows a separate authentication policy to be configured for the console (serial) port. In previous releases, the authentication policy used for access over the network is also used for access through the console port.

By default, the authentication policy used for access over the network is also used for access through the console port.

Optionally, you can configure the AX device to use RADIUS or TACACS+ to authenticate access through the console port. In this case, the AX device tries to authenticate using the external server(s) first. If the external server(s) are unavailable, the AX device then uses the local admin database for authentication.

Note: Use of the local admin database is required, even if you plan to use RADIUS or TACACS+ as the primary authentication method. If the RADIUS or TACACS+ servers are unavailable, the AX device then uses the local admin database.

In addition to configuring the authentication policy, the RADIUS or TACACS+ servers must be added to the AX configuration. For information, see the "Management Security Features" chapter of the *AX Series System Configuration and Administration Guide*.

Using the GUI

1. Select Config > System > Admin > External Authentication > General.

2. In the Authentication Console Type section, select the authentication policy:

   - None – No separate policy is used for console authentication.
   - Local Only – The local admin database on the AX device is used. No external authentication servers are used.

- Local/RADIUS – The local admin database is consulted first. If the local database does not contain an account for the username entered by the admin, the RADIUS server is consulted.
- Local/TACACS+ – The local admin database is consulted first. If the local database does not contain an account for the username entered by the admin, the TACACS+ server is consulted.
- RADIUS/Local – The RADIUS server is consulted first. If the RADIUS server does not respond, the local admin database is consulted.
- TACACS+/Local – The TACACS+ server is consulted first. If the TACACS+ server does not respond, the local admin database is consulted.

3. Click OK.

Note: If you select a policy that includes RADIUS or TACACS+, you also need to configure the RADIUS or TACACS+ settings. Select Config > System > Admin > External Authentication > RADIUS or Config > System > Admin > External Authentication > TACACS+.

Using the CLI

To configure a separate authentication policy for the console port, use the following command at the global configuration level of the CLI:

```
[no] authentication console type
[radius | tacplus] local
```

CLI Example

The following command adds a RADIUS server to the AX configuration:

```
AX(config)#radius-server host 10.10.10.13 secret radp1
```

The following command configures a separate authentication policy for the console port. The policy uses the RADIUS server as the primary authentication method:

```
AX(config)#authentication console type radius local
```

Local Authentication Disable for “admin”

You can disable automatic local authentication for the “admin” account. By default, the AX device always locally authenticates “admin” even if RADIUS or TACACS+ is used as the primary authentication method.

Automatic local authentication of “admin:” is still the default behavior. To disable this behavior, use either of the following methods.

Note: If the RADIUS or TACACS+ server can not be reached, the AX device then uses local authentication for "admin". This is the same behavior as is used for other admin accounts when the remote AAA server can not be reached.

USING THE GUI

To disable automatic local authentication of the "admin" account using the GUI.

1. Select Config > System > Admin > External Authentication > General.

USING THE CLI

To disable automatic local authentication of the "admin" account:

1. Log in using the "admin" account.
2. Use the following command at the global configuration level of the CLI:

```
[no] authentication disable-local
```

Configuring Authorization

Note: The command syntax shown in this section is simplified to show the required or more frequently used options. For complete syntax information, see the *AX Series CLI Reference*.

Note: The configuration options described in this section are available only in the CLI.

1. Add the RADIUS or TACACS+ server(s), if not already added.

```
[no] tacacs-server host {hostname | ipaddr}
secret secret-string
```

```
[no] radius-server host {hostname | ipaddr}
secret secret-string
```

2. Optionally, if using TACACS+, specify the command levels the TACACS+ server will be used to authorize:

```
authorization commands cmd-level method tacplus
[none]
```

The *cmd-level* can be one of the following: 15, 14, 1, or 0.

The **none** option allows a command to execute if Authorization cannot be performed (for example, if all TACACS+ servers are down).

(For descriptions, see "Authorization for CLI Access" on page 266.)

Note: If using RADIUS, you can set the GUI access levels on the RADIUS server itself. See "Authorization for GUI Access" on page 265.

3. Optionally, if using TACACS+, enable Authorization debugging:

   authorization debug *debug-level*

   The *debug-level* can be one of the following: 0x1, 0x2, 0x4, or 0x8.

   (See "TACACS+ Authorization Debug Options" on page 269.)

Configuring Accounting

Note: The command syntax shown in this section is simplified to show the required or more frequently used options. For complete syntax information, see the *AX Series CLI Reference*.

Note: The configuration options described in this section are available only in the CLI.

1. Add the RADIUS or TACACS+ server(s), if not already added.

   [**no**] **tacacs-server host** {*hostname* | *ipaddr*} **secret** *secret-string*

   [**no**] **radius-server host** {*hostname* | *ipaddr*} **secret** *secret-string*

2. To configure Accounting for logon/logoff activity, use the following command:

   [**no**] **accounting exec** {**start-stop** | **stop-only**} {**radius** | **tacplus**}

3. Optionally, if using TACACS+, configure accounting for command execution:

   accounting commands *cmd-level* **stop-only tacplus**

4. Optionally, if using TACACS+, enable Accounting debugging:

   accounting debug *debug-level*

CLI Examples

RADIUS Authentication

The following commands configure a pair of RADIUS servers and configure the AX device to use them first, before using the local database. Since 10.10.10.12 is added first, this server will be used as the primary server. Server 10.10.10.13 will be used only if the primary server is unavailable.

```
AX(config)#radius-server host 10.10.10.12 secret radp1
AX(config)#radius-server host 10.10.10.13 secret radp2
AX(config)#authentication type radius local
```

TACACS+ Authorization

The following commands configure the AX device to use TACACS+ server 10.10.10.13 to authorize commands at all CLI levels. In this example, the **none** option is not used. As a result, if TACACS+ authorization cannot be performed (for example, due to server unavailability), the command is denied.

```
AX(config)#tacacs-server host 10.10.10.13 secret SharedSecret
AX(config)#authorization commands 15 method tacplus
```

TACACS+ Accounting

The following commands configure the AX device to use the same TACACS+ server for accounting of logon/logoff activity and of all command activity:

```
AX(config)#accounting exec start-stop tacplus
AX(config)#accounting commands 15 stop-only tacplus
```

Example Including RADIUS Server Setup

This example shows the AX commands to configure an AX device to use a RADIUS server, and also shows the changes to make on the RADIUS server itself.

The RADIUS server in this example is freeRADIUS. The IP address is 192.168.1.157, and the shared secret is “a10rad”.

To implement the solution, the following steps are required:

1. On the AX device:
   a. Add the RADIUS server.
   b. Enable RADIUS authentication.
2. On the freeRADIUS server:
   a. In the clients.conf file, add the AX device as a client.
   b. Add a dictionary file for vendor “a10networks”, and add the file to the dictionary.
   c. In the users file, add each AX admin as a user.

Configuration on the AX Device

Enter the following commands at the global configuration level of the CLI:

```
AX(config)#radius-server host 192.168.1.157 secret a10rad
AX(config)#authentication type local radius
```

Configuration on the freeRADIUS Server

Changes in clients.conf File

The AX device is added to the clients.conf file as a RADIUS client:

```
vi /usr/local/etc/raddb/clients.conf

client 192.168.1.0/24 {
        secret          = a10rad
        shortname       = private-network-1
}
```

Note: In this example, the AX device's subnet is added as the client.

Creation of dictionary.a10networks File

Here is a copy of the RADIUS dictionary file for the AX device. If you plan to use the AX device as a RADIUS server, copy-and-paste this file onto the external RADIUS servers.

The dictionary file contains objects for AX admin AAA. The file also contains objects for IPv6 migration and external logging features that use RADIUS. For information about these features, see the following documents:

- *AX Series Traffic Logging Guide for IPv6 Migration* for IPv6 migration
- *AX Series IPv4-to-IPv6 Transition Solutions Guide*

AX RADIUS Dictionary File

```
# A10 Networks dictionary.
#
# Version:      1.1  08-Nov-2012
#               $Id$
#
```

```
VENDOR         A10-Networks             22610


BEGIN-VENDOR   A10-Networks

#
#              Admin
#
ATTRIBUTE      A10-App-Name                           1                         string
ATTRIBUTE      A10-Admin-Privilege                    2                         integer
ATTRIBUTE      A10-Admin-Partition                    3                         string
ATTRIBUTE      A10-Admin-Access-Type                  4                         string
ATTRIBUTE      A10-Admin-Role                         5                         string

VALUE          A10-Admin-Privilege             Read-only-Admin                  1
VALUE          A10-Admin-Privilege             Read-write-Admin                 2
VALUE          A10-Admin-Privilege             System-Admin                     3
VALUE          A10-Admin-Privilege             Network-Admin                    4
VALUE          A10-Admin-Privilege             Network-Operator                 5
VALUE          A10-Admin-Privilege             Slb-Service-Admin                6
VALUE          A10-Admin-Privilege             Slb-Service-Operator             7
VALUE          A10-Admin-Privilege             Partition-Read_write             8
VALUE          A10-Admin-Privilege             Partition-Network-Operator       9
VALUE          A10-Admin-Privilege             Partition-SlbService-Admin       10
VALUE          A10-Admin-Privilege             Partition-SlbService-Operator    11
VALUE          A10-Admin-Privilege             Partition-Read-Only              12

#
#              CGN accounting
#
ATTRIBUTE      A10-CGN-Timestamp                      6                         date
ATTRIBUTE      A10-CGN-Protocol                       7                         integer
ATTRIBUTE      A10-CGN-Port-Batch-Size                8                         short
ATTRIBUTE      A10-CGN-Port-Batch-Step-Size           9                         short
ATTRIBUTE      A10-CGN-Inside-Addr                    10                        ipaddr
ATTRIBUTE      A10-CGN-Inside-Port                    11                        short
ATTRIBUTE      A10-CGN-NAT-Addr                       12                        ipaddr
ATTRIBUTE      A10-CGN-NAT-Port                       13                        short
ATTRIBUTE      A10-CGN-Dest-Addr                      14                        ipaddr
ATTRIBUTE      A10-CGN-Dest-Port                      15                        short
ATTRIBUTE      A10-CGN-NAT-Dest-Addr                  16                        ipaddr
ATTRIBUTE      A10-CGN-NAT-Dest-Port                  17                        short
ATTRIBUTE      A10-CGN-Fixed-NAT-Port-Start           18                        short
ATTRIBUTE      A10-CGN-Fixed-NAT-Port-End             19                        short

VALUE          A10-CGN-Protocol                TCP                              1
VALUE          A10-CGN-Protocol                UDP                              2
VALUE          A10-CGN-Protocol                ICMP                             3
VALUE          A10-CGN-Protocol                IP                               4
VALUE          A10-CGN-Protocol                GRE                              5
VALUE          A10-CGN-Protocol                RTSP                             6
VALUE          A10-CGN-Protocol                OTHER                            0

ATTRIBUTE      A10-CGN-Action                         20                        integer

VALUE          A10-CGN-Action                  Port-Allocated                   1
VALUE          A10-CGN-Action                  Port-Freed                       2
VALUE          A10-CGN-Action                  Session-Created                  3
VALUE          A10-CGN-Action                  Session-Deleted                  4
VALUE          A10-CGN-Action                  Port-Batch-Allocated             5
VALUE          A10-CGN-Action                  Port-Batch-Freed                 6
VALUE          A10-CGN-Action                  Fixed-NAT-Port-Range-Allocated   7
VALUE          A10-CGN-Action                  Fixed-NAT-Port-Range-Freed       8
VALUE          A10-CGN-Action                  MSISDN-Query                     9
VALUE          A10-CGN-Action                  HTTP-Request-Got                 10

ATTRIBUTE      A10-CGN-Response                       21                        integer

VALUE          A10-CGN-Response                       Success                   1
VALUE          A10-CGN-Response                       Failure                   2

ATTRIBUTE      A10-CGN-HTTP-Request-Method            22                        integer

VALUE          A10-CGN-HTTP-Request-Method            GET                       1
VALUE          A10-CGN-HTTP-Request-Method            HEAD                      2
VALUE          A10-CGN-HTTP-Request-Method            PUT                       3
VALUE          A10-CGN-HTTP-Request-Method            POST                      4
VALUE          A10-CGN-HTTP-Request-Method            OPTIONS                   5
VALUE          A10-CGN-HTTP-Request-Method            DELETE                    6
VALUE          A10-CGN-HTTP-Request-Method            TRACE                     7
VALUE          A10-CGN-HTTP-Request-Method            CONNECT                   8

ATTRIBUTE      A10-CGN-HTTP-Host                      23                        string
ATTRIBUTE      A10-CGN-HTTP-Url                       24                        string

ATTRIBUTE      A10-CGN-MSISDN                         25                        string
ATTRIBUTE      A10-CGN-IMEI                           26                        string
ATTRIBUTE      A10-CGN-IMSI                           27                        string

ATTRIBUTE      A10-CGN-Timestamp-Millisecond          28                        octets
ATTRIBUTE      A10-CGN-Inside-IPv6-Addr               29                      ipv6addr
ATTRIBUTE      A10-CGN-NAT-IPv6-Addr                  30                      ipv6addr
ATTRIBUTE      A10-CGN-Dest-IPv6-Addr                 31                      ipv6addr
ATTRIBUTE      A10-CGN-NAT-Dest-IPv6-Addr             32                      ipv6addr
ATTRIBUTE       A10-CGN-Inside-Tunnel-Addr          33                        ipaddr
ATTRIBUTE       A10-CGN-NAT-Tunnel-Addr             34                        ipaddr
ATTRIBUTE       A10-CGN-Dest-Tunnel-Addr            35                        ipaddr
ATTRIBUTE       A10-CGN-NAT-Dest-Tunnel-Addr        36                        ipaddr
ATTRIBUTE       A10-CGN-Inside-Tunnel-IPv6-Addr     37                       ipv6addr
ATTRIBUTE       A10-CGN-NAT-Tunnel-IPv6-Addr        38                       ipv6addr
ATTRIBUTE       A10-CGN-Dest-Tunnel-IPv6-Addr       39                       ipv6addr
ATTRIBUTE       A10-CGN-NAT-Dest-Tunnel-IPv6-Addr   40                       ipv6addr

END-VENDOR A10-Networks
```

Notes

The following notes apply to RADIUS logging for all types of traffic logging events.

- The Acct-Session-Id type is octets for all AX generated CGN traffic logging event when log-receiver is set to RADIUS.
- To enable parsing A10 CGN logging RADIUS attributes in Wireshark, we need the following three steps:
  1. Copy the dictionary file from the section above or from the *AX Series System Configuration and Administration Guide* and save it to a text file named “dictionary.a10networks”.
  2. Place the “dictionary.a10networks” file in the Wireshark RADIUS folder. The default path in Windows is “C:\Program Files\Wireshark\radius)”.

Edit the file “dictionary” in the Wireshark RADIUS folder, to add the following line:

```
$INCLUDE dictionary.a10networks
```

Changes in users File

```
vi /usr/local/etc/raddb/users
```

Here are some examples of AX admin definitions in a RADIUS users file on the RADIUS server.

```
###################################

#this is a read-write user
rw                          Cleartext-Password := "111111"
                            A10-Admin-Privilege = Read-write-Admin,
#this is a read-only user
```

```
ro                          Cleartext-Password := "111111"
                            A10-Admin-Privilege = Read-only-Admin,
```

Enabling Parsing of A10 RADIUS Attributes in Wireshark

To enable parsing of A10 RADIUS attributes in Wireshark, perform the following steps:

1. Copy the dictionary file from this document to a file named “dictionary.a10networks”.

2. Place the dictionary.a10networks file in the Wireshark RADIUS folder. (The default path in Windows is “C:\Program Files\Wireshark\radius”.)

3. Open the dictionary in the Wireshark RADIUS folder, and add the following line to it:

   ```
   $INCLUDE dictionary.a10networks
   ```

Configuring Windows IAS for AX RADIUS Authentication

This section describes how to configure Windows Server 2003 Internet Authentication Service (IAS) for use with AX RADIUS authentication. These steps assume that IAS ***and*** Active Directory (AD) are already installed on the Windows 2003 server.

Procedure Overview

To configure Windows IAS for AX RADIUS authentication:

1. On the IAS server, create the following access groups:
   - AX-Admin-Read-Only
   - AX-Admin-Read-Write

2. On the IAS server, configure a RADIUS client for the AX device.

3. On the IAS server, configure the following remote access policies:
   - AX-Admin-Read-Only-Policy
   - AX-Admin-Read-Write-Policy).

4. On the IAS server, add AD users to appropriate AX device access groups, either AX-Admin-Read-Only or AX-Admin-Read-Write.

5. Register the IAS server in AD.

6. On the AX device, configure RADIUS.

7. Test the configuration by attempting to log onto the AX device with AD users added in step 4.

The following sections provide detailed steps for each of these tasks.

Configure Access Groups

1. Select Start > All programs > Administrator tools > Active directory user and computers.

If Active Directory Is Not Installed

If AD is not installed on the IAS server, you can use the following steps to add the users and groups. However, the rest of this section assumes that AD will be used.

1. Open the Computer Management tool by selecting Start > Programs > Administrative Tools > Computer Management.

2. Open the System Tools and Local Users and Groups items, if they are not already open.

3. Right click on Group and select New Group.

4. Enter the following information for the first group:
   - Group Name – AX-Admin-Read-Only
   - Group Description – Read-Only Access to AX devices
   - Members – Add the members using the Add button.

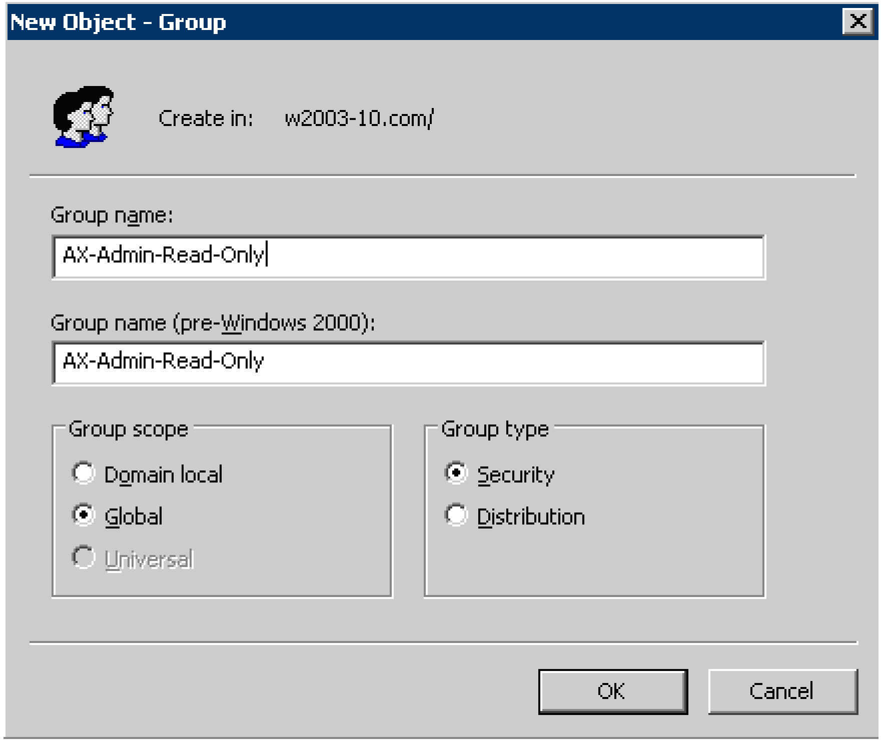

New Object - Group

Create in: w2003-10.com/

Group name:
AX-Admin-Read-Only

Group name (pre-Windows 2000):
AX-Admin-Read-Only

Group scope
○ Domain local
● Global
○ Universal

Group type
● Security
○ Distribution

5. Click Create.

6. Enter the following information for the second group:
   - Group Name – AX-Admin-Read-Write
   - Group Description – Read-Write to AX devices
   - Members – Add members as desired using the Add button

7. Click Create.

8. Click Close.

Configure RADIUS Client for AX Device

1. Open Internet Authentication Service, by selecting Start > Programs > Administrative Tools > Internet Authentication Service.

2. Right-click on Client and select New Client.

3. Enter the following information in the Add Client dialog box:
   - Friendly name – Useful name for the AX device; for example, ax2000_slb1
   - Protocol – RADIUS

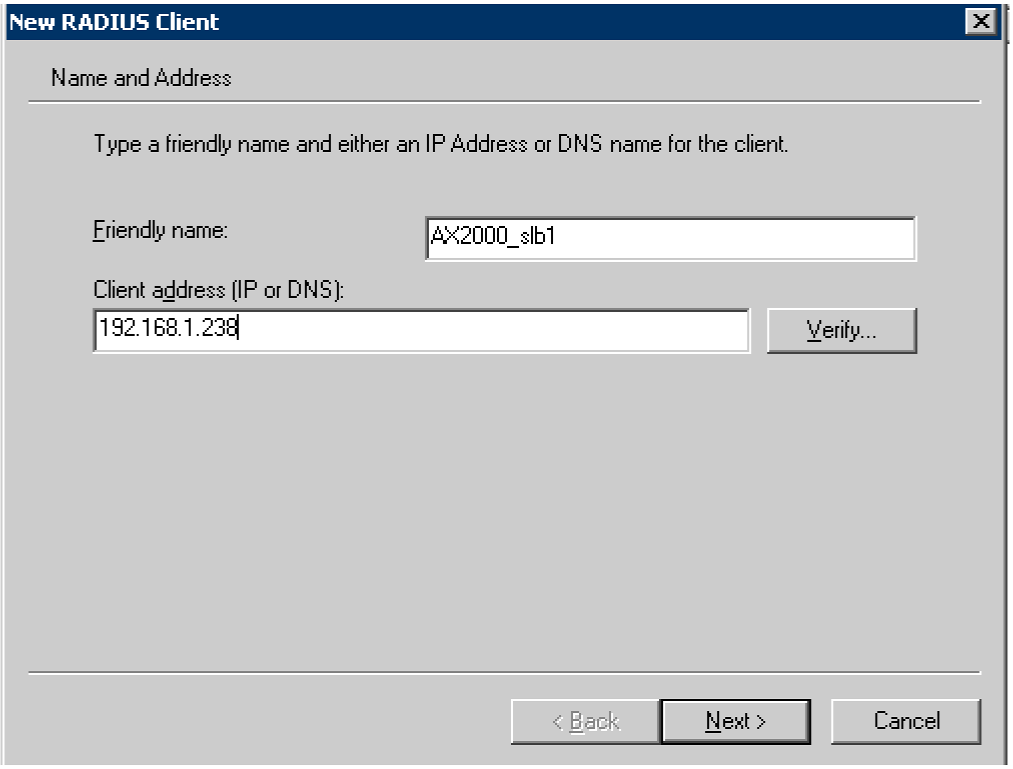

New RADIUS Client

Name and Address

Type a friendly name and either an IP Address or DNS name for the client.

Friendly name: AX2000_slb1

Client address (IP or DNS): 192.168.1.238

Note: 192.168.1.238 is the IP address of the AX device that will use the IAS server for external RADIUS authentication.

4. Click Next.

5. Enter the following information in the Add RADIUS Client dialog box:

   - Client address – IP address or domain name for the client (AX device)
   - Client-Vendor – RADIUS Standard
   - Shared secret – Secret to be shared between IAS and AX. You also will need to enter this in the RADIUS configuration on the AX device.
   - Confirm shared secret – Same as above

Note: Do not select “Request must contain the Message Authenticator attribute”. AX RADIUS authentication does not support this option.

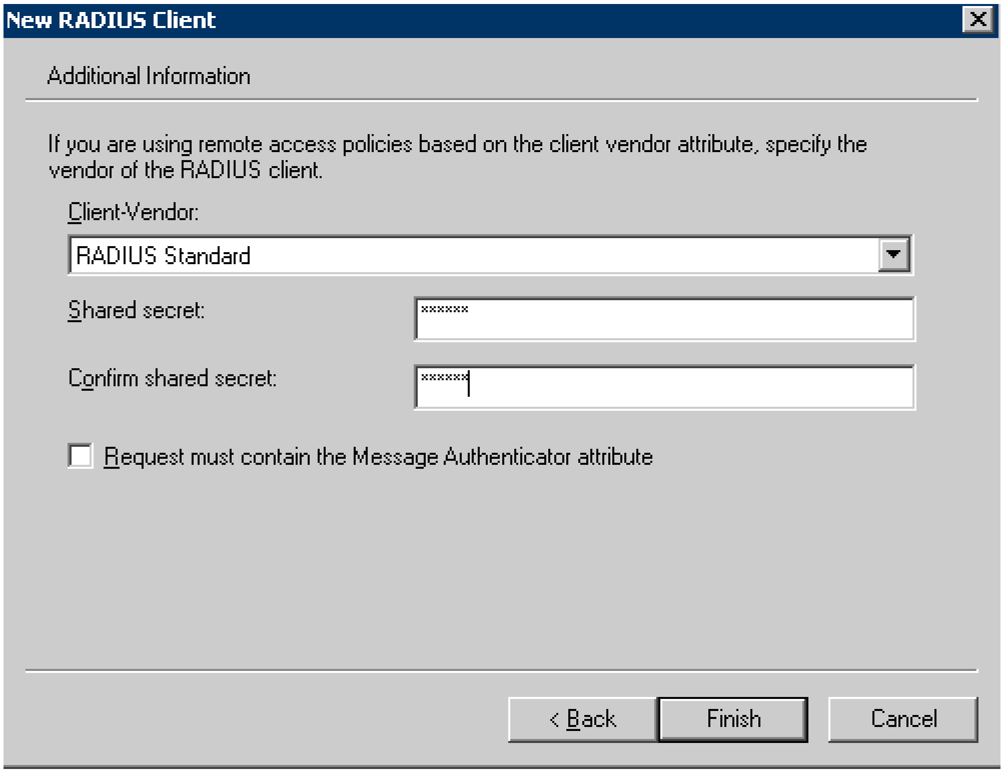

If you are using remote access policies based on the client vendor attribute, specify the vendor of the RADIUS client.

Client-Vendor: RADIUS Standard

Shared secret: ******

Confirm shared secret: ******

□ Request must contain the Message Authenticator attribute

6. Click Next.

Configure Remote Access Policies

1. Open the Internet Authentication Service, if not already open.
2. To create the first remote access policy, right-click on Remote Access Policies, select New Remote Access Policy, and enter the following information:

   Policy Friendly name – AX-Admin-Read-Only-Policy

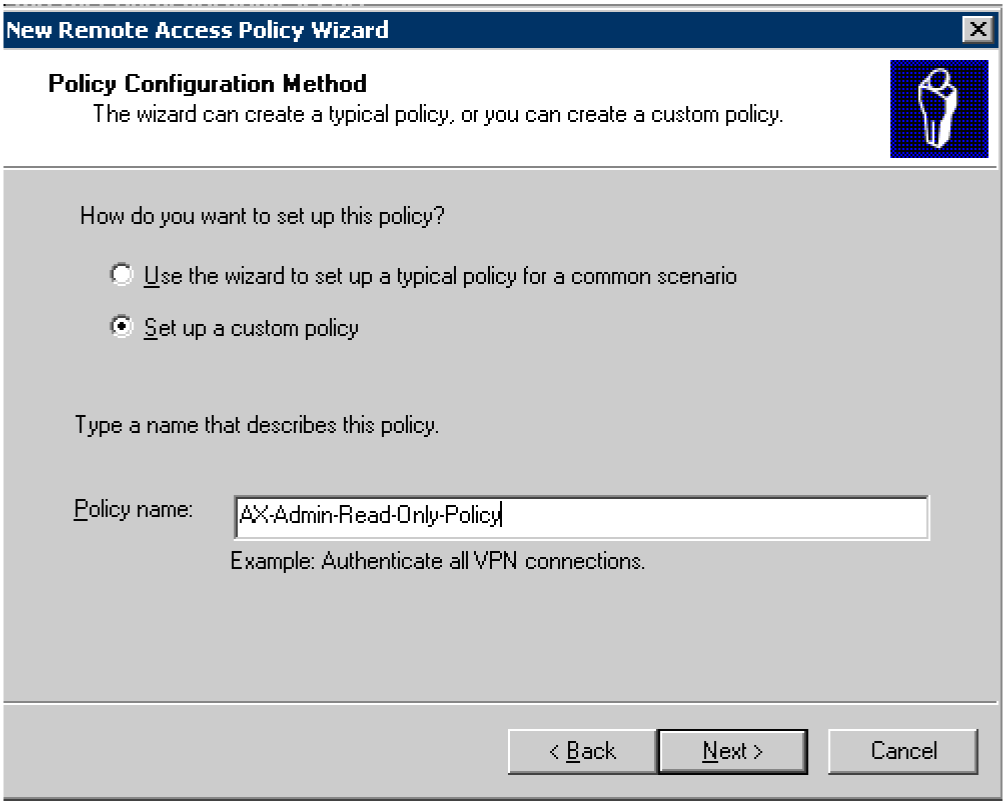

New Remote Access Policy Wizard

Policy Configuration Method

The wizard can create a typical policy, or you can create a custom policy.

How do you want to set up this policy?

○ Use the wizard to set up a typical policy for a common scenario

● Set up a custom policy

Type a name that describes this policy.

Policy name: AX-Admin-Read-Only-Policy

Example: Authenticate all VPN connections.

3. Click Next.
4. In the Add Remote Access Policy dialog box, click Add.
5. In the Select Attribute dialog box, double-click Client Friendly Name.
6. In the Client-Friendly-Name dialog box, enter the friendly name used to define the AX device (for example, AX-Admin-Read-Only-Policy) and click OK.
7. In the same Add Remote Access Policy dialog box as before, click Add again.

8. In the Select Attribute dialog box, double-click Windows-Groups.

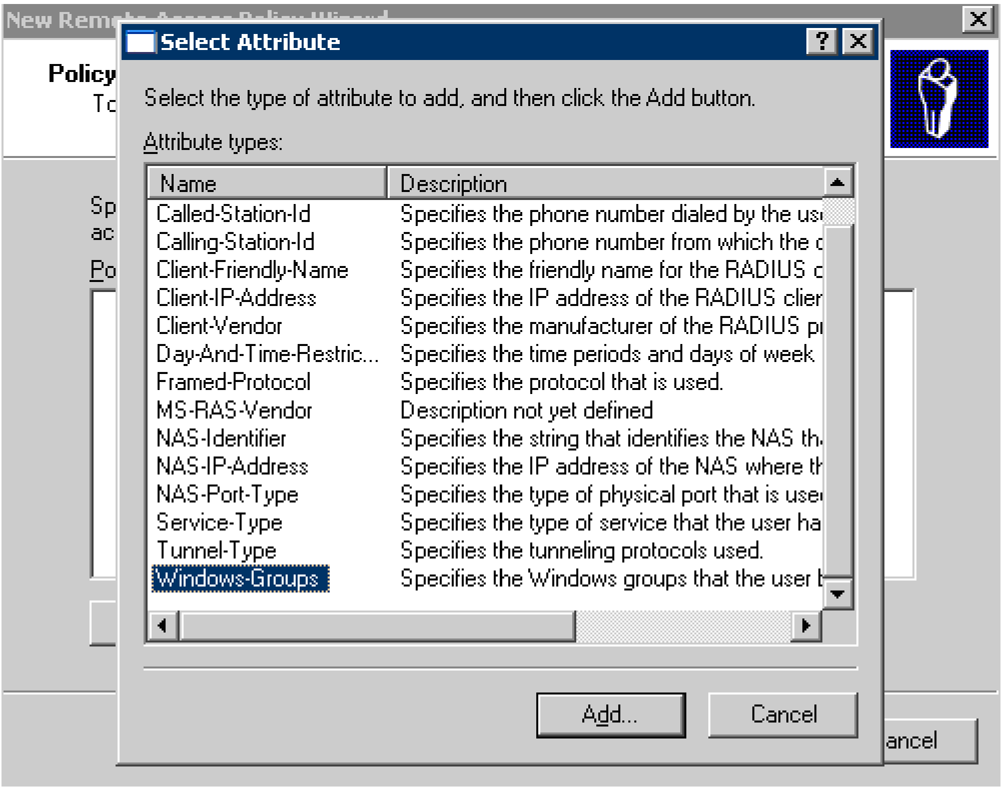

Select the type of attribute to add, and then click the Add button.

Attribute types:

Name	Description
Called-Station-Id	Specifies the phone number dialed by the us
Calling-Station-Id	Specifies the phone number from which the c
Client-Friendly-Name	Specifies the friendly name for the RADIUS c
Client-IP-Address	Specifies the IP address of the RADIUS clier
Client-Vendor	Specifies the manufacturer of the RADIUS pı
Day-And-Time-Restric...	Specifies the time periods and days of week
Framed-Protocol	Specifies the protocol that is used.
MS-RAS-Vendor	Description not yet defined
NAS-Identifier	Specifies the string that identifies the NAS th
NAS-IP-Address	Specifies the IP address of the NAS where th
NAS-Port-Type	Specifies the type of physical port that is use
Service-Type	Specifies the type of service that the user ha
Tunnel-Type	Specifies the tunneling protocols used.
Windows-Groups	Specifies the Windows groups that the user b

9. In the Groups dialog box, click Add, then double-click AX-Admin-Read-Only group, Click OK to add the group, then click OK once more to confirm the groups.

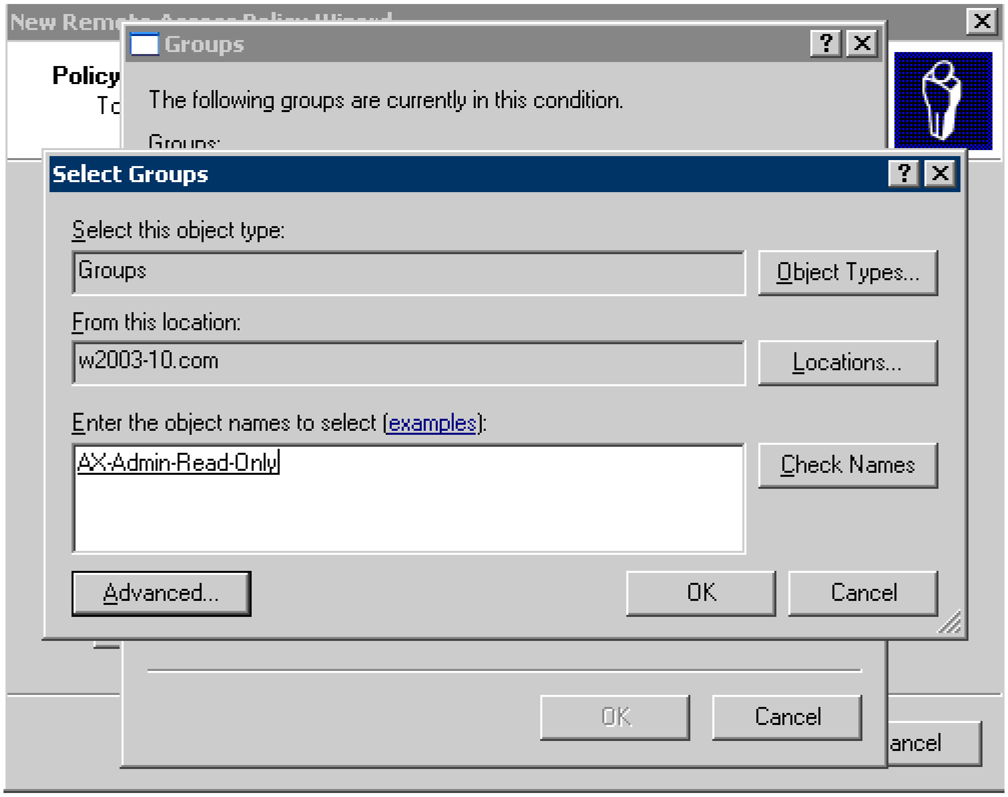

10. In the same Add Remote Access Policy dialog box as before, click Next.

11. Select Grant remote access permission, and click Next.

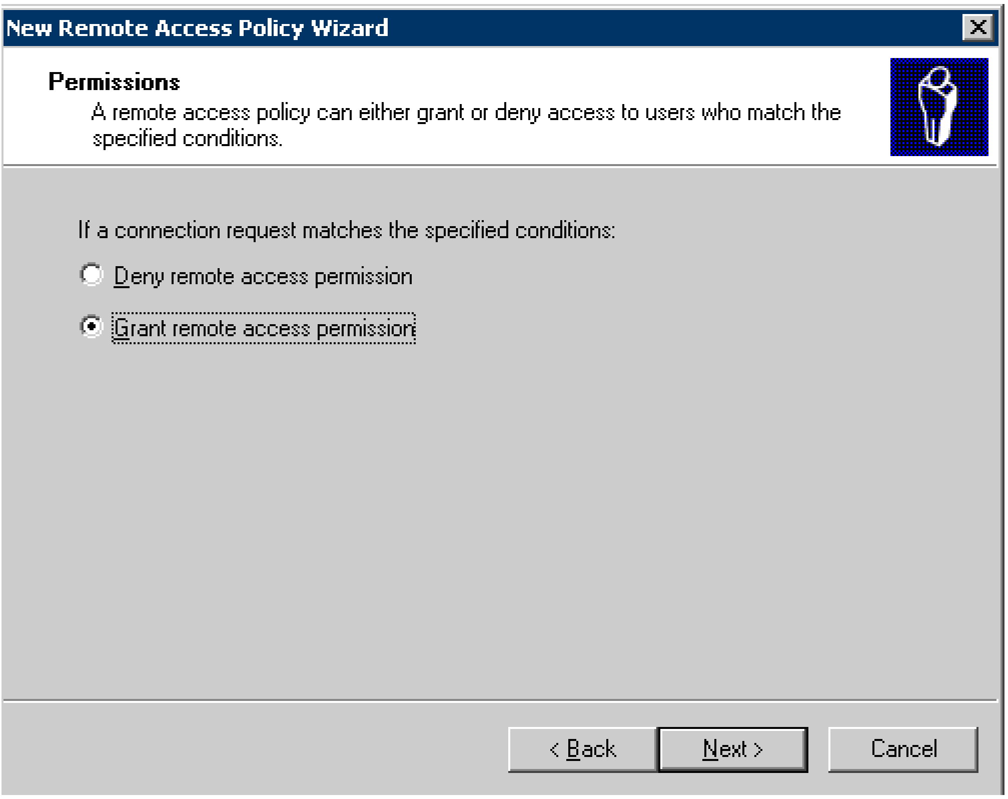

New Remote Access Policy Wizard

Permissions

A remote access policy can either grant or deny access to users who match the specified conditions.

If a connection request matches the specified conditions:

○ Deny remote access permission

● Grant remote access permission

12. Click Edit Profile.

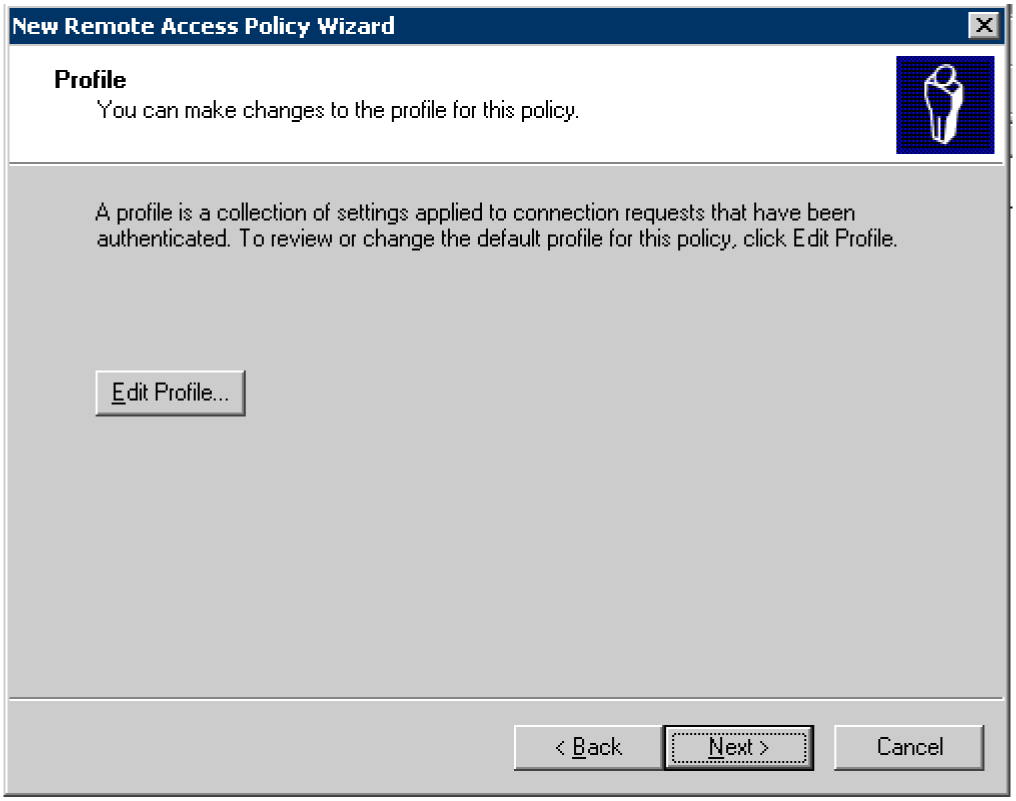

Profile

You can make changes to the profile for this policy.

A profile is a collection of settings applied to connection requests that have been authenticated. To review or change the default profile for this policy, click Edit Profile.

13. In the Edit Dial-in Profile dialog box, select the Authentication tab. Select the type of authentication you are using: CHAP and PAP.

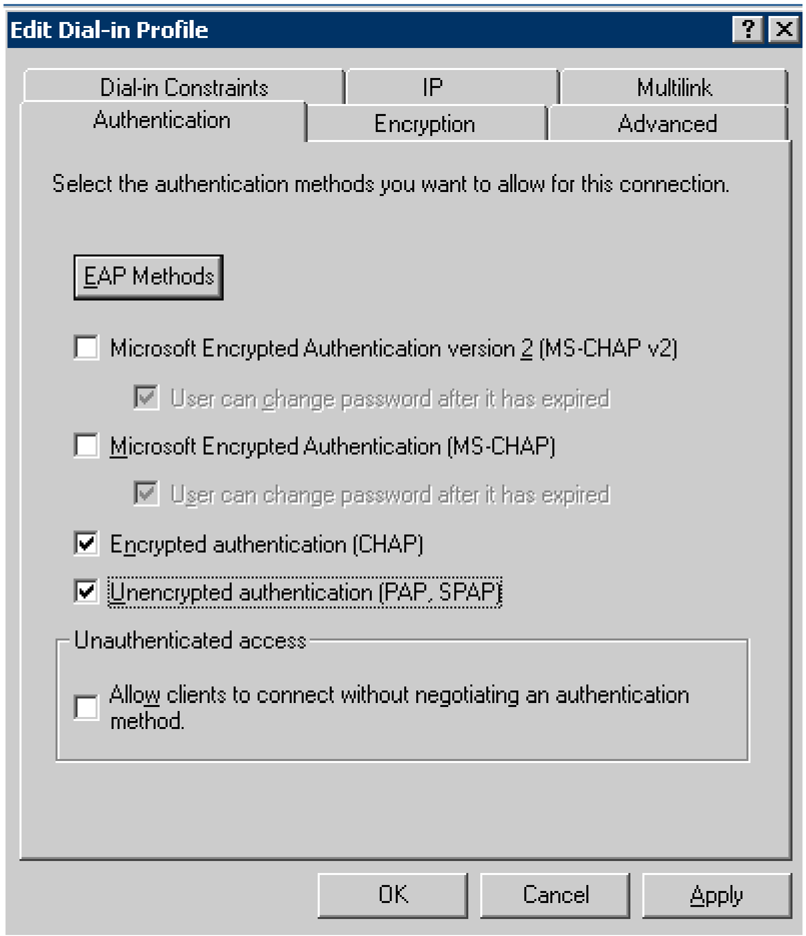

Edit Dial-in Profile

Select the authentication methods you want to allow for this connection.

□ Microsoft Encrypted Authentication version 2 (MS-CHAP v2)

☑ User can change password after it has expired

□ Microsoft Encrypted Authentication (MS-CHAP)

☑ User can change password after it has expired

☑ Encrypted authentication (CHAP)

☑ Unencrypted authentication (PAP, SPAP)

Unauthenticated access

□ Allow clients to connect without negotiating an authentication method.

14. Select the Advanced tab, and click Add.

15. In the RADIUS attributes list, find and double-click the line beginning with Vendor-Specific.

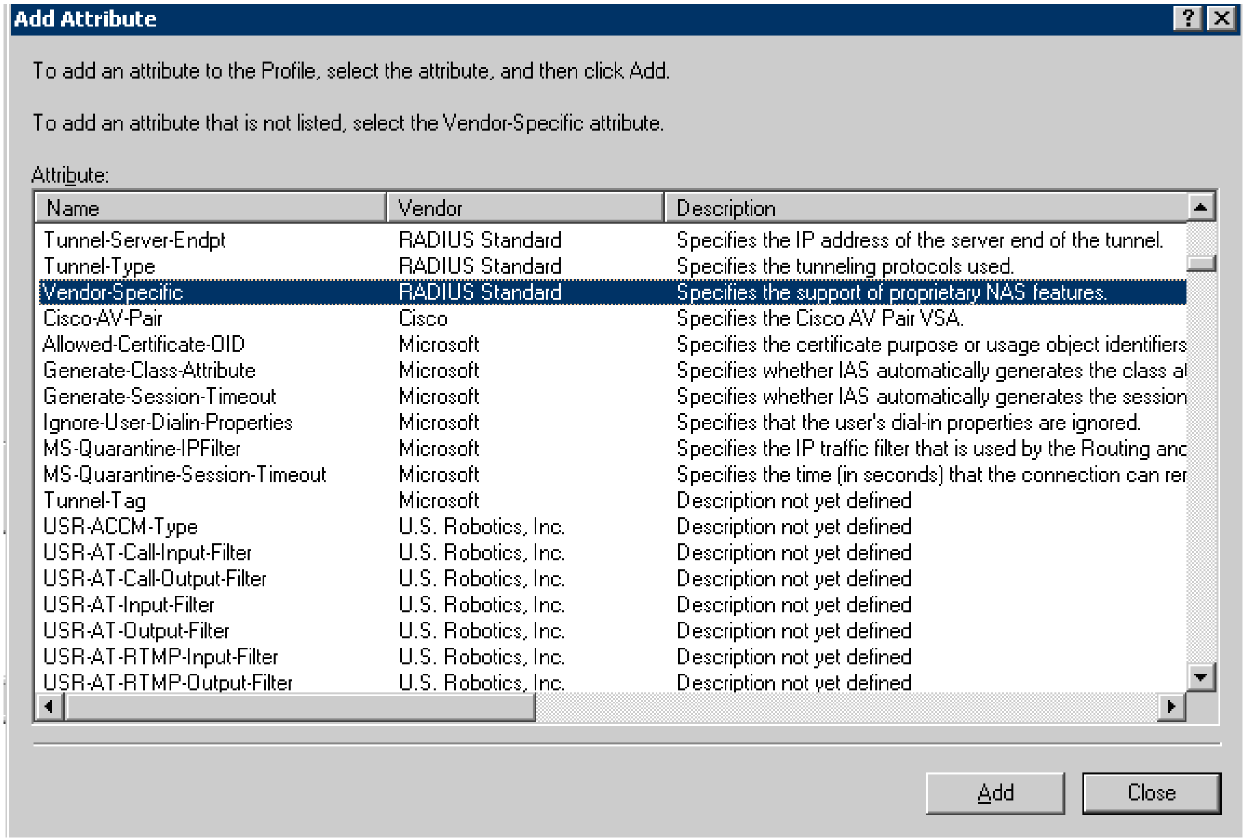

To add an attribute to the Profile, select the attribute, and then click Add.

To add an attribute that is not listed, select the Vendor-Specific attribute.

Attribute:

Name	Vendor	Description
Tunnel-Server-Endpt	RADIUS Standard	Specifies the IP address of the server end of the tunnel.
Tunnel-Type	RADIUS Standard	Specifies the tunneling protocols used.
Vendor-Specific	RADIUS Standard	Specifies the support of proprietary NAS features.
Cisco-AV-Pair	Cisco	Specifies the Cisco AV Pair VSA.
Allowed-Certificate-OID	Microsoft	Specifies the certificate purpose or usage object identifiers
Generate-Class-Attribute	Microsoft	Specifies whether IAS automatically generates the class at
Generate-Session-Timeout	Microsoft	Specifies whether IAS automatically generates the session
Ignore-User-Dialin-Properties	Microsoft	Specifies that the user's dial-in properties are ignored.
MS-Quarantine-IPFilter	Microsoft	Specifies the IP traffic filter that is used by the Routing anc
MS-Quarantine-Session-Timeout	Microsoft	Specifies the time (in seconds) that the connection can rer
Tunnel-Tag	Microsoft	Description not yet defined
USR-ACCM-Type	U.S. Robotics, Inc.	Description not yet defined
USR-AT-Call-Input-Filter	U.S. Robotics, Inc.	Description not yet defined
USR-AT-Call-Output-Filter	U.S. Robotics, Inc.	Description not yet defined
USR-AT-Input-Filter	U.S. Robotics, Inc.	Description not yet defined
USR-AT-Output-Filter	U.S. Robotics, Inc.	Description not yet defined
USR-AT-RTMP-Input-Filter	U.S. Robotics, Inc.	Description not yet defined
USR-AT-RTMP-Output-Filter	U.S. Robotics, Inc.	Description not yet defined

16. In the Multivalued Attribute Information dialog box, click Add and enter the following:

    - Enter vendor code – 22610 (for A10 Networks)
    - Conforms to RADIUS RFC – Yes

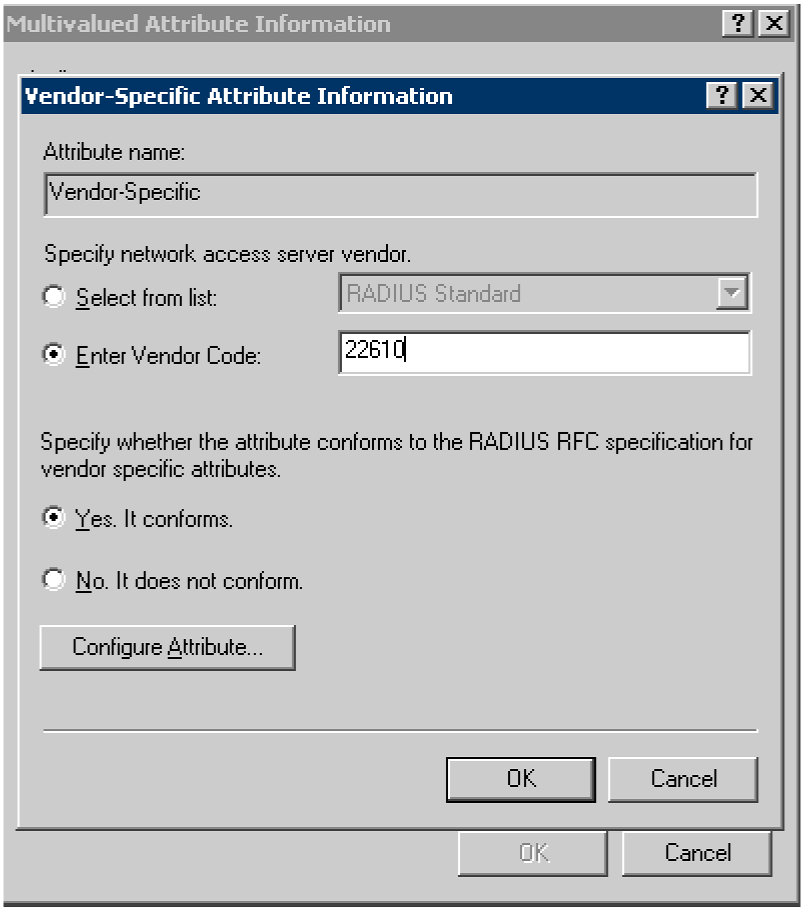
Multivalued Attribute Information
Vendor-Specific Attribute Information
Configure Attribute...
OK
Cancel
OK
Cancel

Attribute name:
Vendor-Specific

Specify network access server vendor.

○ Select from list: RADIUS Standard

● Enter Vendor Code: 22610

Specify whether the attribute conforms to the RADIUS RFC specification for vendor specific attributes.

● Yes. It conforms.

○ No. It does not conform.

17. Click Configure Attribute, and enter the following information:
    - Vendor-assigned attribute number – 2
    - Attribute format – Decimal
    - Attribute value – 1

Note: Attribute value 1 is read-only. Attribute value 2 is read-write.

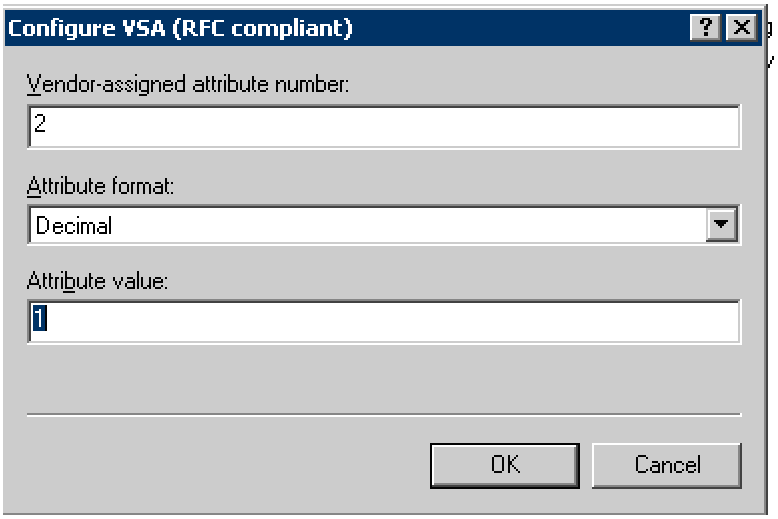

Vendor-assigned attribute number: 2

Attribute format: Decimal

Attribute value: 1

18. Click OK for the Configure VSA, Vendor-Specific Attribute Information, and Multivalued Attribute Information dialog boxes.

19. Click Close in the Add Attributes dialog box.

20. Click OK in the Edit Dial-In Profile dialog box. Optionally, read the suggested help by clicking OK.

21. Click Finish in the Add Remote Access Policy dialog box.

22. To create the second Remote Access Policy, repeat the above steps with the following changes:
    - Policy Friendly name – AX-Admin-Read-Write-Policy
    - Group to add – AX-Admin-Read-Write
    - Attribute value – 2

Add AD Users to AX Access Groups

1. In the Active Directory management console, add the AX access group to the user, tester1:

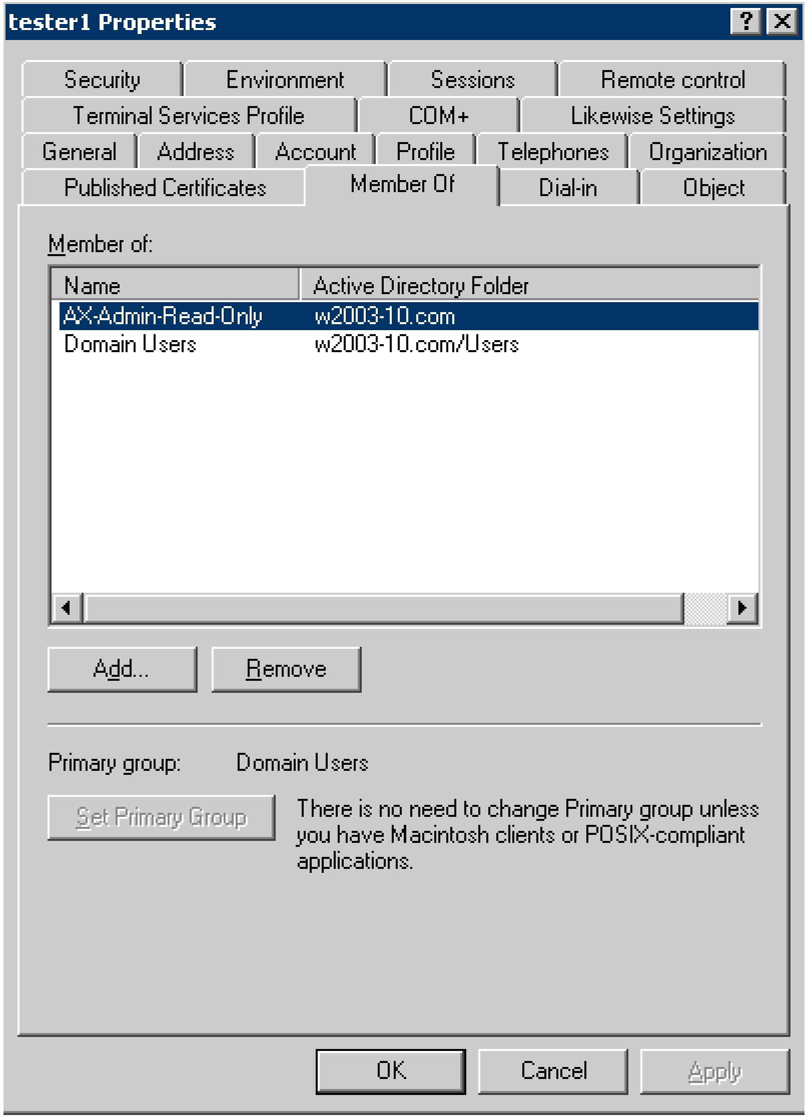

tester1 Properties

Member of:

Name	Active Directory Folder
AX-Admin-Read-Only	w2003-10.com
Domain Users	w2003-10.com/Users

Primary group: Domain Users

There is no need to change Primary group unless you have Macintosh clients or POSIX-compliant applications.

2. Make sure Remote Access Permission is enabled:

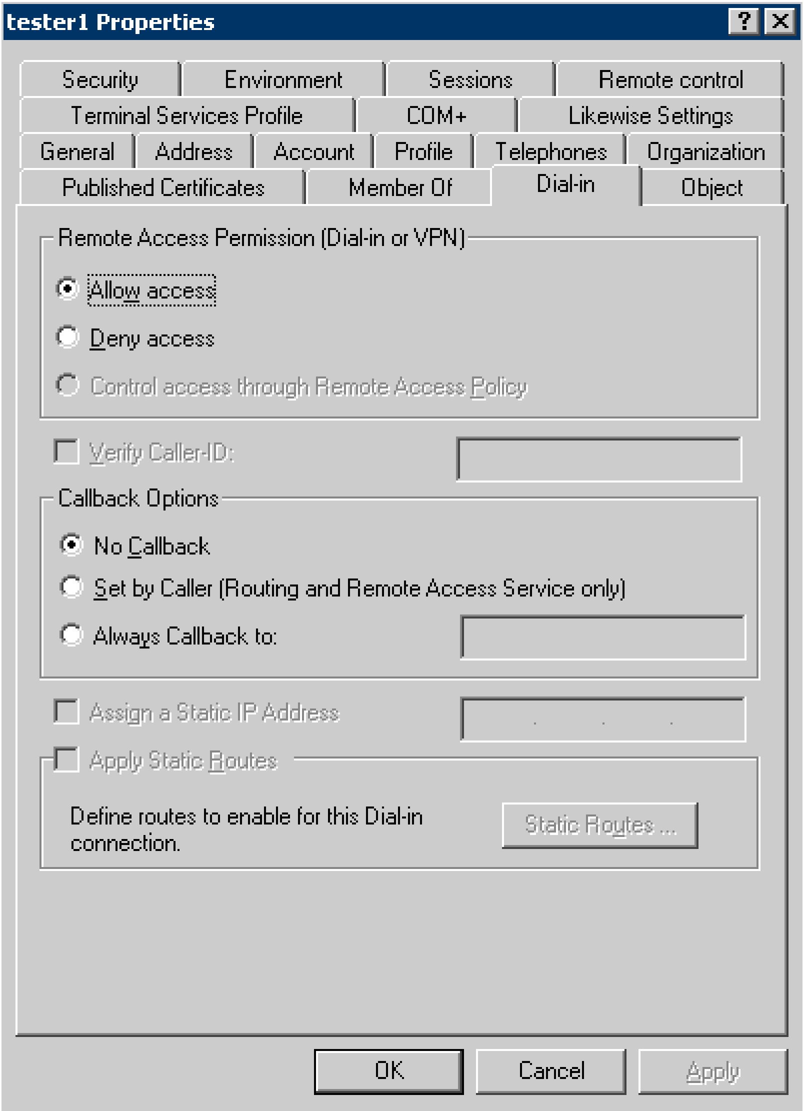

Remote Access Permission (Dial-in or VPN)

● Allow access

○ Deny access

○ Control access through Remote Access Policy

□ Verify Caller-ID:

Callback Options

● No Callback

○ Set by Caller (Routing and Remote Access Service only)

○ Always Callback to:

□ Assign a Static IP Address

□ Apply Static Routes

Define routes to enable for this Dial-in connection.

Register the IAS Server in Active Directory

The IAS RADIUS server must be registered with AD. Otherwise, RADIUS will use compatibility mode instead of AD to authenticate users.

1. Open the IAS main window.

2. Click Action on the menu bar, and click “register server on active directory”.

Configure RADIUS in the AX Device

Add the RADIUS server (IAS server) to the AX device. Make sure the shared secret is the same as the one specified for the RADIUS client configured for the AX server on the IAS server.

```
AX(config)#radius server 192.168.230.10 secret shared-secret
AX(config)#authentication type local radius
```

Note: 192.168.230.10 is the IP address of w2003-10.com, and *shared-secret* is the secret entered in the step 5 in "Configure RADIUS Client for AX Device" on page 285.

Test the Configuration

1. Access the AX CLI command prompt.
2. Enter the login name, in the following format:

   *user-name***@***AD-domain-name*

   In this example, use "tester1@w2003-10.com".
3. Enter the password.
4. Press Enter.

Traffic Security Features

AX Series devices support the following advanced security features, which are described in this chapter:

- “DDoS Protection” on page 301
- “ICMP Rate Limiting” on page 303
- “Access Control Lists (ACLs)” on page 306

DDoS Protection

AX Series devices provide enhanced protection against distributed denial-of-service (DDoS) attacks, with IP anomaly filters. The IP anomaly filters drop packets that contain common signatures of DDoS attacks.

Note: On models AX 3200-12, AX 3400, AX 5100, AX 5200, AX 5200-11, and AX 5630, DDoS protection is hardware-based. On models AX 5630, AX 5200-11, AX 5200, AX 5100, AX 3400 and AX 3200-12, DDoS protection is software-based.

DDoS detection applies only to Layer 3, Layer 4, and Layer 7 traffic. Layer 2 traffic is not affected by the feature. Layer 4 and Layer 7 DDoS applies only to software releases in which Server Load Balancing (SLB) is supported.

All IP anomaly filters except “IP-option” apply to IPv4 and IPv6. The “IP-option” filter applies only to IPv4.

You can enable the following DDoS filters. All filters are supported for IPv4. All filters except IP-option are supported for IPv6.

- Frag – Drops all IP fragments, which can be used to attack hosts running IP stacks that have known vulnerabilities in their fragment reassembly code
- IP-option – Drops all packets that contain any IP options
- Land-attack – Drops spoofed SYN packets containing the same IP address as the source and destination, which can be used to launch an “IP land attack”
- Ping-of-death – Drops all jumbo IP packets, known as “ping of death” packets

Note: On the following models, the ping-of-death option drops all IP packets longer than 32000 bytes: AX 3530, AX 3030, AX 3000-11, AX 3000, AX 2600, and AX 2500. On the following models the option drops IP packets longer than 65535 bytes: AX 5630, AX 5200-11, AX 5200, AX 5100, AX 3400 and AX 3200-12.

- TCP-no-flag – Drops all TCP packets that do not have any TCP flags set
- TCP-SYN-FIN – Drops all TCP packets in which both the SYN and FIN flags are set
- TCP-SYN-frag – Drops incomplete (fragmented) TCP Syn packets, which can be used to launch TCP Syn flood attacks
- Invalid HTTP or SSL payload
- Zero-length TCP window
- Out-of-sequence packet

Enabling DDoS Protection

To enable DDoS protection, use either of the following methods.

Using the GUI

1. Select Config > Service > SLB.
2. On the menu bar, select Global > DDoS Protection.
3. Select each type of DDoS protection filter to enable.

   To enable all of them, select Drop All.
4. Click OK.

Using the CLI

Use the following command at the global configuration level of the CLI:

```
ip anomaly-drop {drop-all | frag | ip-option |
land-attack | ping-of-death | tcp-no-flag | tcp-
syn-fin | tcp-syn-frag}
```

You can enable the following options individually or specify **drop-all** to enable all the options:

As an example, the following command enables DDoS protection against ping-of-death attacks:

```
AX(config)#ip anomaly-drop ping-of-death
```

Displaying and Clearing IP Anomaly Statistics

USING THE CLI\

Select Monitor > Service > Application > Switch.

USING THE CLI

To display IP anomaly statistics, use the following command:

```
show slb l4
```

To clear all Layer 4 SLB statistics, including the IP anomaly counters, use the following command:

```
clear slb l4
```

ICMP Rate Limiting

ICMP rate limiting protects the AX device against denial-of-service (DoS) attacks such as Smurf attacks, which consist of floods of spoofed broadcast ping messages.

ICMP rate limiting monitors the rate of ICMP traffic and drops ICMP packets when the configured thresholds are exceeded.

You can configure ICMP rate limiting filters globally, on individual Ethernet interfaces, and in virtual server templates. If you configure ICMP rate limiting filters at more than one of these levels, all filters are applicable.

ICMP Rate Limiting Parameters

ICMP rate limiting filters consist of the following parameters:

- Normal rate – The ICMP normal rate is the maximum number of ICMP packets allowed per second. If the AX device receives more than the normal rate of ICMP packets, the excess packets are dropped until the next one-second interval begins. The normal rate can be 1-65535 packets per second.
- Maximum rate – The ICMP maximum rate is the maximum number of ICMP packets allowed per second before the AX device locks up ICMP traffic. When ICMP traffic is locked up, all ICMP packets are dropped until the lockup expires. The maximum rate can be 1-65535 packets per second.

- Lockup time – The lockup time is the number of seconds for which the AX device drops all ICMP traffic, after the maximum rate is exceeded. The lockup time can be 1-16383 seconds.

Note: Specifying a maximum rate (lockup rate) and lockup time is optional. If you do not specify them, lockup does not occur.

Log messages are generated only if the lockup option is used and lockup occurs. Otherwise, the ICMP rate-limiting counters are still incremented but log messages are not generated.

Note: The maximum rate must be larger than the normal rate.

Using the GUI

To globally configure ICMP rate limiting:

1. Select Config > Network > ICMP Rate Limiting.
2. Select the ICMP Rate Limiting checkbox to activate the configuration fields.
3. Enter the normal rate in the Normal Rate field.
4. Enter the maximum rate in the Lockup Rate field.
5. Enter the lockup time in the Lockup Period field.
6. Click OK.

To configure ICMP rate limiting on an individual Ethernet interface:

1. Select Config > Network > Interface.
2. Click on the interface name to display the configuration sections for it.
3. Select the ICMP Rate Limiting checkbox to activate the configuration fields.
4. Enter the normal rate in the Normal Rate field.
5. Enter the maximum rate in the Lockup Rate field.
6. Enter the lockup time in the Lockup Period field.
7. Click OK.

To configure ICMP rate limiting in a virtual server template:

Note: This option is applicable only in software releases that support SLB.

1. Select Config > Service > SLB.

2. On the menu bar, select Template > Virtual Server.

3. To edit an existing template, click on the template name. To create a new template, click Add.

   The Virtual Server section appears.

4. Select the ICMP Rate Limit Status checkbox to enable the configuration fields.

5. Enter the normal rate in the Normal Rate field.

6. To configure the lockup time, click Lockup Status.

7. Enter the maximum rate in the Lockup Rate field.

8. Enter the lockup time in the Lockup Period field.

9. Click OK.

USING THE CLI

To configure an ICMP rate-limiting filter, use the following command. You can enter this command at the global configuration level, the configuration level for a physical or virtual Ethernet interface, or the configuration level for a virtual server template.

[**no**] **icmp-rate-limit** *normal-rate* **lockup** *max-rate* *lockup-time*

For descriptions of the parameters, see "ICMP Rate Limiting Parameters" on page 303.

To display ICMP rate limiting information, use the following commands:

show icmp

show interfaces

show slb *virtual-server server-name* **detail**

CLI Example

The following commands configure a virtual server template that sets ICMP rate limiting:

```
AX(config)#slb template virtual-server vip-tmplt
AX(config-vserver)#icmp-rate-limit 25000 lock 30000 60
```

Access Control Lists (ACLs)

You can use Access Control Lists (ACLs) to permit or deny packets based on address and protocol information in the packets. AX devices support the following types of ACLs:

- Standard IPv4 ACL – Standard IPv4 ACLs filter based on source IPv4 address.
- Extended IPv4 ACL – Extended IPv4 ACLs filter based on source and destination IPv4 addresses, IP protocol, and TCP/UDP port numbers.
- Extended IPv6 ACL – Extended IPv6 ACLs filter based on source and destination IPv6 addresses, IP protocol, and TCP/UDP port numbers.

How ACLs Are Used

You can use ACLs for the following tasks:

- Permit or block through traffic.
- Permit or block management access.
- Specify the internal host or subnet addresses to which to provide Network Address Translation (NAT).

An ACL can contain multiple rules. Each rule contains a single permit or deny statement. Rules are added to the ACL in the order you configure them. The first rule you add appears at the top of the ACL.

Rules are applied to the traffic in the order they appear in the ACL (from the top, which is the first rule, downward). The first rule that matches traffic is used to permit or deny that traffic. After the first rule match, no additional rules are compared against the traffic.

Access lists do not take effect until you apply them.

- To permit or block through traffic on an interface, apply the ACL to the interface. (See “Applying an ACL to an Interface” on page 319.)
- To permit or block management access, use the ACL with the **enable-management** command. (See “Securing Admin Access by Ethernet” on page 250.)
- To specify the internal host or subnet addresses to which to provide NAT, use the ACL when configuring the pool. (See “Network Address Translation” on page 205.)

- To permit or block through traffic on a virtual server port, apply the ACL to the virtual port. (See “Applying an ACL to a Virtual Server Port” on page 320.)

Note: ACL use on virtual ports is supported only in software releases that support SLB.

Configuring Standard IPv4 ACLs

To configure a standard IPv4 ACL, use either of the following methods.

Using the GUI

1. Select Config > Network > ACL.
2. Select Standard on the menu bar.
3. Click Add.
4. Enter or select the values to filter. (For descriptions, see the CLI syntax below.)
5. Click OK. The new ACL appears in the Standard ACL table.
6. Click OK to commit the change.

Using the CLI

To configure a standard ACL, use the following command:

```
access-list acl-num [seq-num]
{permit | deny | remark string}
source-ipaddr {filter-mask | /mask-length}
[log [transparent-session-only]]
```

The *acl-num* specifies the ACL number, from 1-99.

The *seq-num* option specifies the sequence number of this rule in the ACL. (See “Resequencing ACL Rules” on page 322.)

The **deny** | **permit** option specifies the action to perform on traffic that matches the ACL:

- **deny** – Drops the traffic.
- **permit** – Allows the traffic.

The **remark** option adds a remark to the ACL. (For more information, see "Adding a Remark to an ACL" on page 317.)

The source address to match on is specified by one of the following:

- **any** – The ACL matches on all source IP addresses.
- **host** *host-src-ipaddr* – The ACL matches only on the specified host IP address.
- *net-src-ipaddr* {*filter-mask* | */mask-length*} – The ACL matches on any host in the specified subnet. The filter-mask specifies the portion of the address to filter:
  - Use 0 to match.
  - Use 255 to ignore.

For example, the following filter-mask filters on a 24-bit subnet: 0.0.0.255

Alternatively, you can use *mask-length* to specify the portion of the address to filter. For example, you can specify "/24" instead "0.0.0.255" to filter on a 24-bit subnet.

The **log** option configures the AX device to generate log messages when traffic matches the ACL. This option is disabled by default. The **transparent-session-only** option limits logging for an ACL rule to creation and deletion of transparent sessions for traffic that matches the ACL rule. (See "Transparent Session Logging" on page 317.)

When ACL logging is enabled, the AX device writes the log messages to the local logging buffer. If you configure an external log server, the AX device also sends the messages to the server. For more information, see "Log Rate Limiting" on page 37.

Note: If you plan to use an external log server, the server must be attached to an AX data port in order for ACL logging messages to reach the server. They will not reach the server if the server is attached to the AX management port.

CLI Example

The following commands configure a standard ACL to deny traffic sent from subnet 10.10.10.x, and apply the ACL to inbound traffic received on Ethernet interface 4:

```
AX(config)#access-list 1 deny 10.10.10.0 0.0.0.255
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#access-list 1 in
```

Configuring Extended IPv4 ACLs

To configure an extended IPv4 ACL, use either of the following methods.

Using the GUI

1. Select Config > Network > ACL.
2. Select Extended on the menu bar.
3. Click Add.
4. Enter or select the values to filter. (For descriptions, see the CLI syntax below.)
5. Click OK. The new ACL appears in the Extended ACL table.
6. Click OK to commit the change.

Using the CLI

To configure an extended ACL, use the following commands.

Syntax for Filtering on Source and Destination IP Addresses

```
[no] access-list acl-num [seq-num]
{permit | deny | l3-vlan-fwd-disable |
  remark string} ip

{any | host host-src-ipaddr |
  net-src-ipaddr {filter-mask | /mask-length}}

{any | host host-dst-ipaddr |
  net-dst-ipaddr {filter-mask | /mask-length}}

[fragments] [vlan vlan-id] [dscp num]

[log [transparent-session-only]]
```

The *acl-num* specifies the ACL number, from 100-199.

The *seq-num* option specifies the sequence number of this rule in the ACL. (See “Resequencing ACL Rules” on page 322.)

The **deny** | **permit** option specifies the action to perform on traffic that matches the ACL:

- **deny** – Drops the traffic.
- **permit** – Allows the traffic.

The **remark** option adds a remark to the ACL. (For more information, see "Adding a Remark to an ACL" on page 317.)

The source address to match on is specified by one of the following:

- **any** – The ACL matches on all source IP addresses.
- **host** *host-src-ipaddr* – The ACL matches only on the specified host IP address.
- *net-src-ipaddr* {*filter-mask* | /*mask-length*} – The ACL matches on any host in the specified subnet. The filter-mask specifies the portion of the address to filter:
  - Use 0 to match.
  - Use 255 to ignore.

For example, the following filter-mask filters on a 24-bit subnet: 0.0.0.255

Alternatively, you can use *mask-length* to specify the portion of the address to filter. For example, you can specify "/24" instead "0.0.0.255" to filter on a 24-bit subnet.

The options for specifying the destination address are the same as those for specifying the source address.

The **fragments** option matches on packets in which the More bit in the header is set (1) or has a non-zero offset.

The **vlan** option matches on the specified VLAN. VLAN matching occurs for incoming traffic only.

The **dscp** option matches on the 6-bit Diffserv value in the IP header, 1-63.

The **established** option matches on TCP packets in which the ACK or RST bit is not set. This option is useful for protecting against attacks from outside. Since a TCP connection from the outside does not have the ACK bit set (SYN only), the connection is dropped. Similarly, a connection established from the inside always has the ACK bit set. (The first packet to the network from outside is a SYN/ACK.)

The **log** option configures the AX device to generate log messages when traffic matches the ACL. This option is disabled by default. The **transparent-session-only** option limits logging for an ACL rule to creation and deletion of transparent sessions for traffic that matches the ACL rule. (See "Transparent Session Logging" on page 317.)

When ACL logging is enabled, the AX device writes the log messages to the local logging buffer. If you configure an external log server, the AX device also sends the messages to the server. For more information, see "Log Rate Limiting" on page 37.

Note: If you plan to use an external log server, the server must be attached to an AX data port in order for ACL logging messages to reach the server. They will not reach the server if the server is attached to the AX management port.

Syntax for Filtering on ICMP Traffic

```
[no] access-list acl-num [seq-num]
{permit | deny | l3-vlan-fwd-disable |
  remark string} icmp

[type icmp-type [code icmp-code]]

{any | host host-src-ipaddr |
  net-src-ipaddr {filter-mask | /mask-length}}

{any | host host-dst-ipaddr |
  net-dst-ipaddr {filter-mask | /mask-length}}

[fragments] [vlan vlan-id] [dscp num]

[log [transparent-session-only]]
```

The **type** and **code** options enable you to filter on ICMP traffic.

The **type** *type-option* option matches based on the specified ICMP type. You can specify one of the following. Enter the type name or the type number (for example, **dest-unreachable** or **3**). The *type-option* can be one of the following:

- **any-type** – Matches on any ICMP type.
- **dest-unreachable** | **3** – Type 3, destination unreachable
- **echo-reply** | **0** – Type 0, echo reply
- **echo-request** | **8** – Type 8, echo request
- **info-reply** | **16** – Type 16, information reply
- **info-request** | **15** – Type 15, information request

- **mask-reply** | **18** – Type 18, address mask reply
- **mask-request** | **17** – Type 17, address mask request
- **parameter-problem** | **12** – Type 12, parameter problem
- **redirect** | **5** – Type 5, redirect message
- **source-quench** | **4** – Type 4, source quench
- **time-exceeded** | **11** – Type 11, time exceeded
- **timestamp** | **13** – Type 13, timestamp
- **timestamp-reply** | **14** – Type 14, timestamp reply
- *type-num* – ICMP type number, 0-254

The **code** *code-num* option matches based on the specified ICMP code. To match on any ICMP code, specify **any-code**. To match on a specific ICMP code, specify the code, 0-254.

Syntax for Filtering on Source and Destination IP Addresses *and* on TCP or UDP Protocol Port Numbers

```
[no] access-list acl-num [seq-num]
{permit | deny | l3-vlan-fwd-disable |
  remark string} {tcp | udp}

{any | host host-src-ipaddr |
  net-src-ipaddr {filter-mask | /mask-length}}
  [eq src-port | gt src-port | lt src-port |
  range start-src-port end-src-port]

{any | host host-dst-ipaddr |
  net-dst-ipaddr {filter-mask | /mask-length}}

  [eq dst-port | gt dst-port | lt dst-port |
  range start-dst-port end-dst-port]

[fragments] [vlan vlan-id] [dscp num] [established]

[log [transparent-session-only]]
```

The **tcp** and **udp** options enable you to filter on protocol port numbers. Use one of the following options to specify the source port(s) on which to filter:

- **eq** *src-port* – The ACL matches on traffic from the specified source port.
- **gt** *src-port* – The ACL matches on traffic from any source port with a higher number than the specified port.
- **lt** *src-port* – The ACL matches on traffic from any source port with a lower number than the specified port.
- **range** *start-src-port end-src-port* – The ACL matches on traffic from any source port within the specified range.

The same options can be used to specify the destination port(s) on which to filter.

CLI Example

The following commands configure an extended IPv4 ACL to deny traffic sent from subnet 10.10.10.x to 10.10.20.5:80, and apply the ACL to inbound traffic received on Ethernet interface 7:

```
AX(config)#access-list 100 deny tcp 10.10.10.0 0.0.0.255 10.10.20.5 /32 eq 80
AX(config)#interface ethernet 7
AX(config-if:ethernet7)#access-list 100 in
```

Configuring Extended IPv6 ACLs

To configure an extended IPv4 ACL, use either of the following methods.

Using the GUI

1. Select Config > Network > ACL.
2. Select IPv6 on the menu bar.
3. Click Add.
4. Enter or select the values to filter. (For descriptions, see the CLI syntax below.)
5. Click OK. The new ACL appears in the IPv6 ACL table.
6. Click OK to commit the change.

USING THE CLI

To configure an IPv6 ACL, use the following commands:

```
[no] ipv6 access-list name
```

Enter this command at the global configuration level of the CLI. The *name* can be a string up to 16 characters long. This command changes the CLI to the configuration level for the ACL, where the following ACL-related commands are available.

The permit | deny Command

This command specifies the action to take for traffic that matches the ACL, specifies the source and destination addresses upon which to perform the action, and optionally, enables logging.

```
[no] [seq-num] {permit | deny} {ipv6 | icmp}

{any | host host-src-ipv6addr |
  net-src-ipv6addr /mask-length}

{any | host host-dst-ipv6addr |
  net-dst-ipv6addr /mask-length}

[fragments] [vlan vlan-id] [dscp num]

[log [transparent-session-only]]
```

or

```
[no] {permit | deny} {tcp | udp}

{any | host host-src-ipv6addr |
  net-src-ipv6addr /mask-length}
  [eq src-port | gt src-port | lt src-port |
  range start-src-port end-src-port]

{any | host host-dst-ipv6addr |
  net-dst-ipv6addr /mask-length}
  [eq dst-port | gt dst-port | lt dst-port |
  range start-dst-port end-dst-port]

[fragments] [vlan vlan-id] [dscp num] [established]

[log [transparent-session-only]]
```

Parameter	Description
seq-num	Sequence number of this rule in the ACL. You can use this option to resequence the rules in the ACL.
deny | **permit**	Action to take for traffic that matches the ACL. **deny** – Drops the traffic. **permit** – Allows the traffic.
ipv6 | **icmp**	Filters on IPv6 or ICMP packets.
tcp | **udp**	Filters on TCP or UDP packets. The **tcp** and **udp** options enable you to filter on protocol port numbers.
any | **host** *host-src-ipv6addr* | *net-src-ipv6addr /mask-length*	Source IP address(es) to filter. **any** – The ACL matches on all source IP addresses. **host** *host-src-ipv6addr* – The ACL matches only on the specified host IPv6 address. *net-src-ipv6addr /mask-length* – The ACL matches on any host in the specified subnet. The *mask-length* specifies the portion of the address to filter.
eq *src-port* | **gt** *src-port* | **lt** *src-port* | **range** *start-src-port end-src-port*	For **tcp** or **udp**, the source protocol ports to filter. **eq** *src-port* – The ACL matches on traffic from the specified source port. **gt** *src-port* – The ACL matches on traffic from any source port with a higher number than the specified port. **lt** *src-port* – The ACL matches on traffic from any source port with a lower number than the specified port.

	range *start-src-port end-src-port* – The ACL matches on traffic from any source port within the specified range.
any | **host** *host-dst-ipv6addr* | *net-dst-ipv6addr /mask-length*	Destination IP address(es) to filter.
eq *dst-port* | **gt** *dst-port* | **lt** *dst-port* | **range** *start-dst-port end-dst-port*	For **tcp** or **udp**, the destination protocol ports to filter.
fragments	Matches on packets in which the More bit in the header is set (1) or has a non-zero offset.
vlan *vlan-id*	Matches on the specified VLAN. VLAN matching occurs for incoming traffic only.
dscp *num*	Matches on the 6-bit Diffserv value in the IP header, 1-63.
established	Matches on TCP packets in which the ACK or RST bit is not set. This option is useful for protecting against attacks from outside. Since a TCP connection from the outside does not have the ACK bit set (SYN only), the connection is dropped. Similarly, a connection established from the inside always has the ACK bit set. (The first packet to the network from outside is a SYN/ACK.)
log [**transparent-session-only**]	Configures the AX device to generate log messages when traffic matches the ACL. The **transparent-session-only** option limits logging for an ACL rule to creation and deletion of transparent sessions for traffic that matches the ACL rule. (See "Transparent Session Logging" on page 317.)

The remark Command

The **remark** command adds a remark to the ACL. The remark appears at the top of the ACL when you display it in the CLI. Here is the syntax:

```
[no] remark string
```

The *string* can be 1-63 characters. To use blank spaces in the remark, enclose the entire remark string in double quotes.

Adding a Remark to an ACL

You can add a remark to an ACL. The remark appears at the top of the ACL when you display it in the CLI, or next to the ACL in the ACL tables displayed in the GUI.

Here is a CLI example:

```
AX(config)#access-list 42 permit host 192.168.1.42
AX(config)#access-list 42 deny 192.168.1.0 /24
AX(config)#access-list 42 remark "The meaning of life"
AX(config)#show access-list ipv4 42
Access List 42 "The meaning of life"
access-list 42 10 permit host 192.168.1.42  Hits: 0
access-list 42 20 deny 192.168.1.0 0.0.0.255  Hits: 0
```

As shown in this example, the remark appears at the top of the ACL, above the first rule.

To use blank spaces in the remark, enclose the entire remark string in double quotes, as shown in the example. The ACL must already exist before you can configure a remark for it.

Transparent Session Logging

The **transparent-session-only** option limits logging for an ACL rule to creation and deletion of transparent sessions for traffic that matches the ACL rule.

A *transparent session* is a non-SLB Layer 2 or Layer 3 session that the AX device sets up for traffic that is transiting through the AX device, but is not initiated or terminated on the device.

Sample Log Messages for Transparent Sessions

The following sections show examples of the log messages generated for transparent sessions.

IPv4 Sessions

The following example shows the log messages for creation and deletion of an IPv4 transparent session:

```
Oct 29 2009 12:00:55 Notice  [AX]:[eth 1] TCP 200.200.200.100:32548 >
1.1.1.100:80  ACL rule transparent session expired (ACL 150)
Oct 29 2009 12:00:55 Notice  [AX]:[eth 1] TCP 200.200.200.100:32548 >
1.1.1.100:80  ACL rule transparent session created (ACL 150)
```

The interface on which the ACL matched traffic is indicated in brackets (in this example, “eth 1”). The addresses are shown as *src-ip***:***port* > *dst-ip***:***port*. The ACL number or ACL name is shown at the end of the message.

IPv6 Sessions

For successfully created TCP or UDP sessions, a separate message is generated when the session is created and when it expires:

```
Feb 24 2010 02:18:27 Notice  [AX]:[ve 21] UDP 2001:10::100:50213 >
2001:7::40:53  ACL rule transparent session expired (IPV6_LIST)
Feb 24 2010 02:18:12 Notice  [AX]:[ve 21] UDP 2001:10::100:50213 >
2001:7::40:53  ACL rule transparent session created (IPV6_LIST)
Feb 24 2010 02:15:12 Notice  [AX]:[ve 21] TCP 2001:10::100:4401 > 2001:7::40:22
ACL rule transparent session expired (IPV6_LIST)
Feb 24 2010 02:15:08 Notice  [AX]:[ve 21] TCP 2001:10::100:4401 > 2001:7::40:22
ACL rule transparent session created (IPV6_LIST)
```

For all other types of transparent IPv6 sessions, a message is generated if the packet is forwarded:

```
Feb 24 2010 02:18:07 Notice  [AX]:[ve 21] IP 2001:10::100 > 2001:7::40  ACL
rule permitted this packet (IPV6_LIST)
```

If a TCP or UDP packet is denied, a message such as the following is generated:

```
Feb 24 2010 02:18:07 Notice [AX]:[ve 21] TCP 2001:10::100:57373 > 2001:7::40:80
ACL rule transparent session denied (IPV6_LIST)
```

For all other types of transparent IPv6 sessions, a message such as the following is generated:

```
Feb 24 2010 02:18:07 Notice  [AX]:[ve 21] IP 2001:10::100 > 2001:7::40  ACL
rule denied this packet (IPV6_LIST)
```

Configuration

To configure session filtering for transparent IPv6 sessions on an interface:

1. Configure an IPv6 ACL that uses the **log transparent-session-only** option.

2. Apply the ACL to the interface that receives incoming traffic for the sessions.

3. For the following AX models only, enable the **cpu-process** option on the interface that receives incoming traffic for the sessions: AX 5630, AX 5200-11, AX 5200, AX 5100, AX 3400, and AX 3200-12

CLI Example

The following commands configure an IPv6 ACL for transparent session logging, and apply it to an IPv6 interface:

```
AX(config)#ipv6 access-list tran_sess_log1
AX(config-access-list:trans_sess_log1)#permit tcp any any log transparent-session-only
AX(config-access-list:trans_sess_log1)#exit
AX(config)#interface ve 21
AX(config-if:ve21)#ipv6 access-list tran_sess_log1 in
```

Applying an ACL to an Interface

To apply a configured ACL to an interface, use either of the following methods.

USING THE GUI

To apply an ACL to an Ethernet port:

1. Select Config > Network > Interface.

2. Select LAN on the menu bar.

3. Click on the port number.

4. In the IPv4 section, select the ACL from the Access List field.

5. Click OK.

To apply an ACL to a Virtual Ethernet (VE) interface:

1. Select Config > Network > Interface.

2. Select Virtual on the menu bar.

3. Click on the VE name.

4. Select IPv4.

5. Select the ACL from the Access List field.

6. Click OK.

USING THE CLI

Access the configuration level for the interface and use one of the following commands:

```
[no] access-list acl-num in
[no] ipv6 access-list name in
```

The following commands configure a standard IPv4 ACL to deny traffic from subnet 10.10.10.x, and apply the ACL to the inbound traffic direction on Ethernet interface 4:

```
AX(config)#access-list 1 deny 10.10.10.0 0.0.0.255
AX(config)#interface ethernet 4
AX(config-if:ethernet4)#access-list 1 in
```

Applying an ACL to a Virtual Server Port

You can apply an ACL to a virtual server port. An ACL applied to a virtual server port permits or denies traffic just as an ACL applied to a physical port or Virtual Ethernet (VE) interface does.

Note: ACL use on virtual ports is supported only in software releases that support SLB.

To apply a configured ACL to a virtual server port, use either of the following methods.

USING THE GUI

1. Select Config > Service > SLB.

2. Select Virtual Server on the menu bar.

3. Click Add or click on the name of a configured virtual server.

4. Enter or change information in the General section, if you are configuring a new virtual server.

5. In the Port section, click Add or select a port and click Edit.

6. In the Virtual Server Port section, select the ACL from the Access List drop-down list.

7. Click OK.

8. Click OK again to return to the virtual server table.

USING THE CLI

To apply an ACL to a virtual port in the CLI, use the following command at the configuration level for the virtual port:

```
[no] access-list {acl-num | name acl-name}
```

The *acl-num* option specifies an IPv4 ACL. The **name** *acl-name* option specifies the name of an IPv6 ACL.

Using an ACL to Control Management Access

To use an ACL to control management access, see "Securing Admin Access by Ethernet" on page 250.

Using an ACL for NAT

To use an ACL for NAT, configure the ACL, then use either of the following methods to bind the ACL to a NAT pool.

USING THE GUI

To bind an ACL to an IP source NAT pool:

1. Select Config > Service > IP Source NAT.

2. Select Binding on the menu bar.

3. Select the ACL number from the ACL drop-down list.

4. Select the pool ID from the NAT Pool drop-down list.

5. Click Add. The new binding appears in the ACL section.

6. Click OK.

USING THE CLI

To use a configured ACL in an IPv4 NAT pool, use the following command:

```
[no] ip nat inside source
{list acl-name
  {pool pool-name | pool-group pool-group-name}
static local-ipaddr global-ipaddr}
```

The **list** *acl-name* option specifies the ACL.

Resequencing ACL Rules

An ACL can contain multiple rules. Each **access-list** command configures one rule. Rules are added to the ACL in the order you configure them. The first rule you add appears at the top of the ACL.

Rules are applied to the traffic in the order they appear in the ACL (from the top rule, which is the first, downward). The first rule that matches traffic is used to permit or deny that traffic. After the first rule match, no additional rules are compared against the traffic.

Each ACL has an implicit deny any rule at the end of the ACL. This rule is applied to any traffic that does not match any of the configured rules in the ACL. The deny any rule at the end of the ACL is not displayed and cannot be removed.

You can resequence the rules in an ACL. When you create an ACL rule, the AX device assigns a sequence number to the rule and places the rule at the bottom of the ACL. Here is an example:

```
AX(config)#access-list 86 permit host 10.10.10.12
AX(config)#access-list 86 deny 10.10.10.0 /24
AX(config)#show access-list ipv4 86
access-list 86 10 permit host 10.10.10.12 log Hits: 0
access-list 86 20 deny 10.10.10.0 0.0.0.255 log Hits: 0
```

In this example, two rules are configured for ACL 86. The default sequence numbers are used. The first rule has sequence number 10, and each rule after that has a sequence number that is higher by 10.

The intent of this ACL is to deny all access from the 10.10.10.x subnet, except for access from specific host addresses. In this example, the permit rule for the host appears before the deny rule for the subnet the host is in, so the host will be permitted. However, suppose another permit rule is added for another host in the same subnet.

```
AX(config)#access-list 86 permit host 10.10.10.13
AX(config)#show access-list ipv4 86
access-list 86 10 permit host 10.10.10.12 log Hits: 0
access-list 86 20 deny 10.10.10.0 0.0.0.255 log Hits: 0
access-list 86 30 permit host 10.10.10.13 log Hits: 0
```

By default, since no sequence number was specified when the rule was configured, the rule is placed at the end of the ACL. Because the deny rule comes before the permit rule, host 10.10.10.13 will never be permitted.

To resequence the ACL to work as intended, the deny rule can be deleted, then re-added. Alternatively, either the deny rule or the second permit rule can be resequenced to appear in the right place. To change the sequence number of an ACL rule, delete the rule, then re-add it with the sequence number.

```
AX(config)#no access-list 86 30
AX(config)#access-list 86 11 permit host 10.10.10.13 log
AX(config)#show access-list ipv4 86
access-list 86 10 permit host 10.10.10.12 log Hits: 0
access-list 86 11 permit host 10.10.10.13 log Hits: 0
access-list 86 20 deny 10.10.10.0 0.0.0.255 log Hits: 0
```

In this example, rule 30 is deleted, then re-added with sequence number 11. The ACL will now work as intended, and permit hosts 10.10.10.12 and 10.10.10.13 while denying all other hosts in the 10.10.10.x subnet. To permit another host, another rule can be added, sequenced to come before the deny rule.

```
AX(config)#access-list 86 12 permit host 10.10.10.14 log
AX(config)#show access-list ipv4 86
access-list 86 10 permit host 10.10.10.12 log Hits: 0
access-list 86 11 permit host 10.10.10.13 log Hits: 0
access-list 86 12 permit host 10.10.10.14 log Hits: 0
access-list 86 20 deny 10.10.10.0 0.0.0.255 log Hits: 0
```

USING THE GUI

Each row in the Standard ACL and Extended ACL tables is a separate ACL rule. You can configure multiple rules in the same ACL. In this case, they still appear as separate rows, with the same ACL number.

The AX device applies the ACL rules in the order they are listed, starting at the top of the table. The first rule that matches traffic is used to permit or deny that traffic. After the first rule match, no additional rules are compared against the traffic.

If you need to re-order the ACL rules, you can do so by clicking the up or down arrows at the ends of the rows containing the ACL rules.

Click OK to commit the changes.

USING THE CLI

See the description above.

Using the Management Interface as the Source for Management Traffic

By default, the AX device attempts to use a route from the main route table for management connections originated on the AX device. You can enable the AX device to use the management route table to initiate management connections instead.

This chapter describes the AX device's two route tables, for data and management traffic, and how to configure the device to use the management route table.

Route Tables

The AX device uses separate route tables for management traffic and data traffic.

- Management route table – Contains all static routes whose next hops are connected to the management interface. The management route table also contains the route to the device configured as the management default gateway.
- Main route table – Contains all routes whose next hop is connected to a data interface. These routes are sometimes referred to as *data plane* routes. Entries in this table are used for load balancing and for Layer 3 forwarding on data ports.

  This route table also contains copies of all static routes in the management route table, excluding the management default gateway route.

You can configure the AX device to use the management interface as the source interface for automated management traffic. In addition, on a case-by-case basis, you can enable use of the management interface and management route table for various types of management connections to remote devices:

The AX device automatically will use the management route table for reply traffic on connections initiated by a remote host that reaches the AX device on the management port. For example, this occurs for SSH or HTTP connections from remote hosts to the AX device.

Note: Static routes whose next hop is the management interface are duplicated in the management route table.

Keep the Management and Data Interfaces in Separate Networks

It is recommended to keep the management interface and the data interfaces in separate networks. If both tables have routes to the same destination subnet, some operations such as pinging may have unexpected results. An exception is the default route (0.0.0.0/0), which can be in both tables.

To display the routes in each table, use the following commands:

- **show ip route mgmt** – This command displays the routes in the management route table.
- **show ip route** or **show ip fib** – These commands display data plane routes.

Management Routing Options

You can configure the AX device to use the management interface as the source interface for the following management protocols, used for automated management traffic:

- SYSLOG
- SNMPD
- NTP
- RADIUS
- TACACS+
- SMTP

For example, when use of the management interface as the source interface for control traffic is enabled, all log messages sent to remote log servers are sent through the management interface. Likewise, the management route table is used to find a route to the log server. The AX device does not attempt to use any routes from the main route table to reach the server, even if a route in the main route table could be used.

In addition, on a case-by-case basis, you can enable use of the management interface and management route table for the following types of management connections to remote devices:

- Upgrade of the AX software
- SSH or Telnet connection to a remote host
- Import or export of files
- Export of **show techsupport** output

- Reload of black/white lists
- SSL loads (keys, certificates, and Certificate Revocation Lists)
- Copy or restore of configurations
- Backups

Caution: **If you enable this feature, then downgrade to AX Release 1.2.4 or earlier, it is possible to lose access to the AX device after you downgrade. This can occur if you configure an external AAA server (TACACS+ server) to authorize CLI commands entered on the AX device, and the TACACS+ server is connected to the AX device through the management default gateway.**

If this is the case, before you downgrade, remove the TACACS+ configuration from the AX device. After you downgrade, you can re-add the configuration, but make sure the TACACS+ server can be reached using a route other than through the management default gateway.

Enabling Use of the Management Interface as the Source for Automated Management Traffic

By default, use of the management interface as the source interface for automated management traffic is disabled.

To enable it, use the following command at the configuration level for the management interface:

```
[no] ip control-apps-use-mgmt-port
```

Here is an example:

```
AX(config-if:management)#ip control-apps-use-mgmt-port
```

Using the Management Interface as the Source Interface for Manually Generated Management Traffic

To use the management interface as the source interface for manually generated management traffic, use the **use-mgmt-port** option.

In the GUI, this option is provided as a Use Management Port checkbox on the applicable pages.

In the CLI, this option is supported with the following commands.

Commands at the User EXEC Level

```
ssh [use-mgmt-port] {host-name | ipaddr)
login-name [protocol-port]
```

```
telnet [use-mgmt-port] {host-name | ipaddr)
[protocol-port]
```

Commands at the Privileged EXEC Level

```
export {ssl-cert |ssl-key | axdebug}
file-name [use-mgmt-port]url
```

```
ssh [use-mgmt-port] {host-name | ipaddr)
login-name [protocol-port]
```

```
telnet [use-mgmt-port] {host-name | ipaddr)
[protocol-port]
```

Commands at the Global Configuration Level

```
backup {config | log} [use-mgmt-port] url
```

```
copy {running-config | startup-config |
from-profile-name}
[use-mgmt-port]
{url | to-profile-name [cf]}
```

```
health external
{delete program-name |
import [use-mgmt-port] [description] url |
export [use-mgmt-port] program-name url}
```

```
[no] restore [use-mgmt-port] url
```

```
[no] slb ssl-load
{certificate file-name | private-key file-name}
[use-mgmt-port] url
```

```
upgrade {cf | hd} {pri | sec} [use-mgmt-port] url
```

Show Commands

```
show techsupport [[use-mgmt-port] export url]
[page]
```

Boot Options

This chapter describes how to display or change the storage area from which the AX device boots.

Note: This chapter does not describe how to upgrade the system image. For upgrade instructions, see the release notes for the release to which you plan to upgrade.

Storage Areas

The AX device has four storage areas (also called "image areas") that can contain software images and configuration files:

- Primary storage on the Solid State Drive (SSD) or disk
- Secondary storage on the SSD or disk
- Primary storage on the compact flash (CF)
- Secondary storage on the compact flash

These storage areas are depicted in Figure 40.

FIGURE 40 Software Image Locations on the AX Device

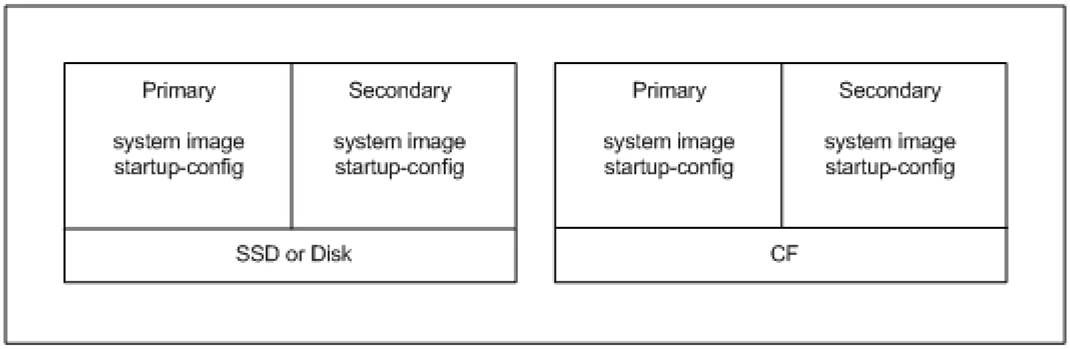

The SSD or disk storage areas are used for normal operation. The compact flash storage areas are used only for system recovery.

Normally, each time the AX device is rebooted, the device uses the same storage area that was used for the previous reboot. For example, if the primary storage area of the SSD or disk was used for the previous reboot, the system image and startup-config from the primary storage area are used for the next reboot.

Unless you change the storage area selection or interrupt the boot sequence to specify a different storage area, the AX device always uses the same storage area each time the device is rebooted.

Note: The AX device always tries to boot using the SSD or disk first. The compact flash is used only if the SSD or hard disk is unavailable. If you need to boot from compact flash for system recovery, contact A10 Networks.

Displaying Storage Information

To display the software images installed in the AX storage areas, and the currently running software version, use either of the following methods.

Using the GUI

The first page that is displayed when you log onto the GUI is the Summary page. This page lists the software image versions installed in each of the storage areas.

FIGURE 41 Monitor > Overview > Summary

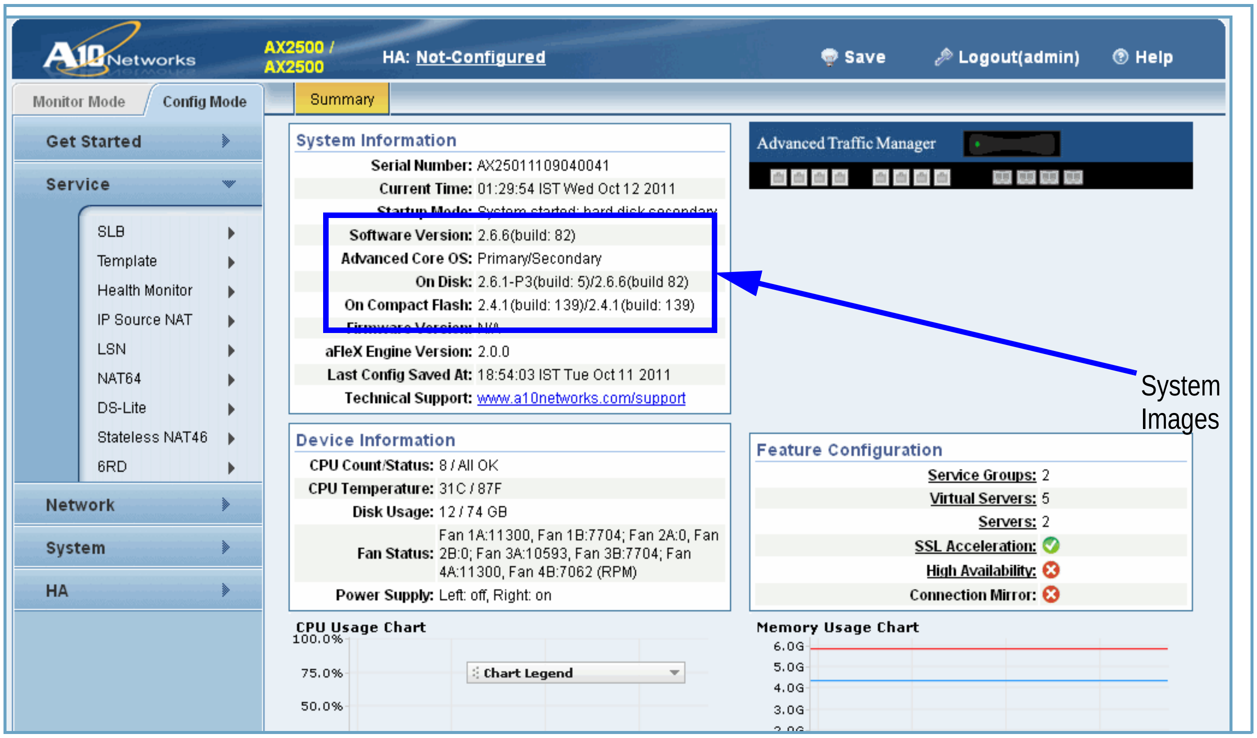

System Information

Serial Number: AX25011109040041
Current Time: 01:29:54 IST Wed Oct 12 2011
Startup Mode: System started; hard disk secondary
Software Version: 2.6.6(build: 82)
Advanced Core OS: Primary/Secondary
On Disk: 2.6.1-P3(build: 5)/2.6.6(build 82)
On Compact Flash: 2.4.1(build: 139)/2.4.1(build: 139)
Firmware Version: N/A
aFleX Engine Version: 2.0.0
Last Config Saved At: 18:54:03 IST Tue Oct 11 2011
Technical Support: www.a10networks.com/support

Device Information

CPU Count/Status: 8 / All OK
CPU Temperature: 31C / 87F
Disk Usage: 12 / 74 GB
Fan Status: Fan 1A:11300, Fan 1B:7704; Fan 2A:0, Fan 2B:0; Fan 3A:10593, Fan 3B:7704; Fan 4A:11300, Fan 4B:7062 (RPM)
Power Supply: Left: off, Right: on

Feature Configuration

Service Groups: 2
Virtual Servers: 5
Servers: 2
SSL Acceleration: ✓
High Availability: ☒
Connection Mirror: ☒

Above the highlighted fields, the Startup Mode field lists the storage area used for the most recent reboot. The Software Version field lists the currently running software version.

Using the CLI

The **show version** command shows storage area information. The command also lists other information, including the currently running software version.

```
AX#show version
AX Series Advanced Traffic Manager AX2500
  Copyright 2007-2011 by A10 Networks, Inc.  All A10 Networks products are
  protected by one or more of the following US patents and patents pending:
  7716378, 7675854, 7647635, 7552126, 20090049537, 20080229418, 20080040789,
  20070283429, 20070271598, 20070180101

      64-bit Advanced Core OS (ACOS) version 2.6.4-P1, build 120 (May-31-2011,03:22)
        Booted from Hard Disk primary image
      Serial Number: AX25011109040041
      aFleX version: 2.0.0
      Hard Disk primary image (default) version 2.6.4-P1, build 120
      Hard Disk secondary image version 2.6.0-P1, build 31
      Compact Flash primary image version 2.4.1, build 139
      Compact Flash secondary image (default) version 2.4.1, build 139
      Last configuration saved at May-31-2011, 18:29
      Hardware: 8 CPUs(Stepping 5), Single 74G Hard disk
      Memory 6123 Mbyte, Free Memory 1585 Mbyte
      Current time is Jun-3-2011, 22:54
      The system has been up 3 days, 4 hours, 24 minutes
```

Displaying the Storage Location for Future Reboots

To display the storage area that will be used for the future reboots, use either of the following methods.

Note: The AX device always tries to boot using the SSD or disk first. The compact flash is used only if the SSD or hard disk is unavailable. If you need to boot from compact flash for system recovery, contact A10 Networks.

Using the GUI

Select Config > System > Settings > Boot.

Using the CLI

Use the following command: **show bootimage**

In the following example, the AX device is configured to boot from the primary storage area on the SSD or disk:

```
AX(config)#show bootimage
                      (* = Default)
                          Version
-----------------------------------------------
Hard Disk primary         2.6.4-P1.120 (*)
Hard Disk secondary       2.6.0-P1.31
Compact Flash primary     2.4.1.139
Compact Flash secondary   2.4.1.139 (*)
```

Booting from a Different Storage Area

To reboot from a different storage area, do one of the following:

- Interrupt the boot sequence and use the bootloader menu to temporarily select the other storage area.
- Configure the AX device to use the other storage area for all future reboots, then reboot.

Temporarily Changing the Storage Location for the Next Reboot

To temporarily change the storage location to use for a reboot, interrupt the boot sequence to access the bootloader menu.

To access the bootloader menu, reboot the AX device, then press Esc within 3 seconds when prompted.

When the bootloader menu appears, use the Up and Down arrow keys to select the image area from which to boot, and press Enter. The menu does not automatically time out. You must press Enter to reboot using the selected image.

Caution: **Each storage area has its own version of the startup-config. When you save configuration changes, they are saved only to the startup-config in the storage area from which the AX device was booted.**

If you plan to reboot from a different storage area, but you want to use the same configuration, first save the configuration to the other storage area. (The procedures in “Permanently Changing the Storage Location for Future Reboots” on page 334 include steps for this.)

Note: The bootloader menu is available on new AX devices that are shipped with AX Release 2.6.1 or later. However, the bootloader menu is not automatically installed when you upgrade from a release earlier than 2.6.1. To install the bootloader menu on upgraded devices, see the AX Release 2.6.1 release notes, or the description of the **boot-block-fix** command in the *AX Series CLI Reference* for 2.6.1 or later.

```
AX#reboot
Rebooting System Now !!!
Proceed with reboot? [yes/no]:yes
INIT:

Shutting down........Restarting system.
Press `ESC' to enter the boot menu... 1
  Admin presses Esc within 3 seconds.
```

```
   #      #     ###     #     #
  # #    ##    #   #    ##    # ###### ##### #    #  ####  #####  #    #  ####
 #   #  # #   #     #   # #   # #        #   #    # #    # #    # #   #  #
#     #   #   #     #   #  #  # #####    #   #    # #    # #    # ####    ####
#######   #   #     #   #   # # #        #   # ## # #    # #####  #  #        #
#     #   #    #   #    #    ## #        #   ##  ## #    # #   #  #   #  #    #
#     # #####   ###     #     # ######   #   #    #  ####  #    # #    #  ####

Copyright 2005-2011 by A10 Networks, Inc.  All A10 Networks products are
protected by one or more of the following US patents and patents pending:
7716378, 7675854, 7647635, 7552126, 20090049537, 20080229418, 20080040789,
20070283429, 20070271598, 20070180101
-------------------------------------------------------------------
 0: AX ACOS (Primary Image)
 1: AX ACOS (Secondary Image)
-------------------------------------------------------------------

Use the Up and Down arrow keys to select the image from which to boot.
Press enter to boot the selected image.

  Admin presses down arrow to select 1.

    Highlighted entry is 1:

  Admin presses Enter to reboot using the selected image.

Booting 'AX ACOS (Secondary Image)'
Please wait while the system boots...

Booting........................[OK]

AX login:
```

Permanently Changing the Storage Location for Future Reboots

To change the storage area that will be used for future reboots, use either of the following methods.

Note: The procedures in this section change the storage area selection for all future reboots (unless you later change the selection again). If you only need to temporarily override the storage area selection for a single reboot, see "Temporarily Changing the Storage Location for the Next Reboot" on page 333.

Caution: **Each storage area has its own version of the startup-config. When you save configuration changes, they are saved only to the startup-config in the storage area from which the AX device was booted.**

If you plan to reboot from a different storage area, but you want to use the same configuration, first save the configuration to the other storage area. The procedures in this section include a step for this.

Using the GUI

1. Save the configuration, by clicking the Save icon at the top of the GUI window.

   This step copies any unsaved configuration changes from the running-config to the startup-config.

FIGURE 42 Save the Configuration

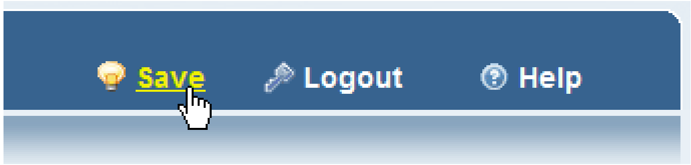

2. Copy the startup-config to the storage area you plan to use for the next reboot:

   a. Select Config > System > Config.

   b. Click one of the following:

      - Primary Startup – This option selects the startup-config in the primary storage area of the SSD or hard drive. Click this link if you plan to use the primary storage area for the next reboot.
      - Secondary Startup – This option selects the startup-config in the secondary storage area of the SSD or hard drive. Click this link if you plan to use the secondary storage area for the next reboot.

      The selected startup-config is displayed, in the Content field.

   c. Use the Copy From drop-down list to select the startup-config file that was used for the most recent reboot. This is the startup-config to which you saved configuration changes in step 1. The commands in the selected startup-config replace the commands that were displayed in the Content field.

   d. Click OK. The startup-config you selected in step b is replaced with the one selected in step c.

3. Change the storage area to use for booting:
   a. Select Config > System > Settings > Boot.
   b. Select Primary or Secondary.
   c. Click OK.

Note: You can select the storage area for the SSD or disk and for the compact flash. However, the AX device always tries to boot using the SSD or disk first. The compact flash is used only if the SSD or hard disk is unavailable.

Using the CLI

1. Save the configuration, by entering the following command:

   ```
   write memory
   ```

   This step copies any unsaved configuration changes from the running-config to the startup-config.

2. Copy the startup-config to the storage area you plan to use for the next reboot. Use the following command:

   ```
   write memory {primary | secondary}
   ```

   - **primary** – Select this option if the secondary storage area is the one the AX device was booted from, and you plan to boot from the primary storage area next time.
   - **secondary** – Select this option if the primary storage area is the one the AX device was booted from, and you plan to boot from the secondary storage area next time.

3. Select the storage area to use for future reboots. Use the following command.

   ```
   bootimage {hd | cf} {pri | sec}
   ```

Note: This command is available at the global configuration level of the CLI.

- The **hd** | **cf** option specifies whether you are selecting a storage area on the SSD or hard drive, or the compact flash.
- The **pri** | **sec** option specifies whether the AX first tries to boot using the image in the primary image area or the secondary image area.

To verify the setting, use the following command: `show bootimage`

CLI Example

In this example, the AX device was booted from the primary storage area, and will be configured to use the secondary image area for future reboots.

The following command displays the current setting for the storage area to use for reboots:

```
AX(config)#show bootimage
                      (* = Default)
                          Version
-----------------------------------------------
Hard Disk primary         2.6.4-P1.120 (*)
Hard Disk secondary       2.6.0-P1.31
Compact Flash primary     2.4.1.139
Compact Flash secondary   2.4.1.139 (*)
```

The following commands save the configuration and copy it to the other storage area:

```
AX(config)#write memory
Building configuration...
[OK]
AX(config)#write memory secondary
Building configuration...
[OK]
```

The following commands configure the AX device to use the secondary storage area on the SSD or hard drive for future reboots, and verify the setting:

```
AX(config)#bootimage hd sec
Secondary image will be used if AX is booted from hard disk
AX(config)#show bootimage
                      (* = Default)
                          Version
-----------------------------------------------
Hard Disk primary         2.6.4-P1.120
Hard Disk secondary       2.6.0-P1.31 (*)
Compact Flash primary     2.4.1.139
Compact Flash secondary   2.4.1.139 (*)
```

Configuration Management

By default, when you click the Save button in the GUI or enter the **write memory** command in the CLI, all unsaved configuration changes are saved to the startup-config. The next time the AX device is rebooted, the configuration is reloaded from this file.

In addition to these simple configuration management options, the AX device has advanced configuration management options that allow you to save multiple configuration files. You can save configuration files remotely on a server and locally on the AX device itself.

Note: For information about synchronizing configuration information between AX devices in a High Availability (HA) pair, see "Manually Synchronizing Configuration Information" on page 198.

Note: For upgrade instructions, see the release notes for the AX release to which you plan to upgrade.

Backing Up System Information

The AX device allows you to back up the system, individual configuration files, and even log entries onto remote servers. You can use any of the following file transfer protocols:

- Trivial File Transfer Protocol (TFTP)
- File Transfer Protocol (FTP)
- Secure Copy Protocol (SCP)
- Unix Remote Copy (RCP)

USING THE GUI

1. Select Config > System > Maintenance.
2. Select one of the following from the menu bar:
   - Backup > System – This option backs up the configuration file(s), aFleX scripts, and SSL certificates and keys.
   - Backup > Config – This option backs up only the specified configuration file.

- Backup > Syslog – This option backs up the log entries in the AX device’s syslog buffer. (If there are any core files on the system, this option backs them up as well.)

3. Select the backup location:
   - Local – Saves the backup on the PC or workstation where you are using the AX GUI.
   - Remote – Saves the backup onto another PC or workstation.

4. If you selected Local:
   a. Click OK.
   b. Click Save and navigate to the save location. Optionally, you can edit the filename.
   c. Click Save.

5. If you selected Remote:
   a. In the Protocol drop-down list, select the file transfer protocol: FTP, TFTP, RCP, or SCP.
   b. If using FTP and the remote device does not use the default FTP port, change the port.
   c. In the Host field, enter the hostname or IP address of the remote device.
   d. In the Location field, enter the pathname. To change the system backup file from the default name (“backup_system.tar”), specify the new name at the end of the path.
   e. In the User and Password fields, enter the username and password required for write access to the remote device.
   f. Click OK.

USING THE CLI

At the Privileged EXCE level of the CLI, use the following command:

```
backup {system | log} url
```

The **system** option backs up the startup-config file, aFleX scripts, and SSL certificates and keys.

The **log** option backs up the log entries in the AX device’s syslog buffer.

The *url* option specifies the file transfer protocol, username, and directory path. You can enter the entire URL on the command line or press Enter to display a prompt for each part of the URL. If you enter the entire URL and a password is required, you will still be prompted for the password. To enter the entire URL:

- **tftp:**//*host*/*file*
- **ftp:**//[*user*@]*host*[:*port*]/*file*
- **scp:**//[*user*@]*host*/*file*
- **rcp:**//[*user*@]*host*/*file*

Saving Multiple Configuration Files Locally

The AX device has CLI commands that enable you to store and manage multiple configurations on the AX device.

Note: Unless you plan to locally store multiple configurations, you do not need to use any of the advanced commands or options described in this section. Just click Save in the GUI or enter the **write memory** command in the CLI to save configuration changes. These simple options replace the commands in the startup-config stored in the image area the AX device booted from with the commands in the running-config.

Configuration Profiles

Configuration files are managed as configuration profiles. A *configuration profile* is simply a configuration file. You can locally save multiple configuration profiles on the AX device. The configuration management commands described in this section enable you to do the following:

- Save the startup-config or running-config to a configuration profile.
- Copy locally saved configuration profiles.
- Delete locally saved configuration profiles.
- Compare two configuration profiles side by side to see the differences between the configurations.
- Link the command option “startup-config” to a configuration profile other than the one stored in the image area used for the most recent reboot. (This is the profile that “startup-config” refers to by default.) This option makes it easier to test a configuration without altering the configuration stored in the image area.

Note: Although the enable and admin passwords are loaded as part of the system configuration, they are not saved in the configuration profiles. Changes to the enable password or to the admin username or password take effect globally, regardless of the values that were in effect when a given configuration profile was saved.

USING THE GUI

You can use the Config > System > Config File page to perform the following configuration management tasks:

- Display individual configuration files.
- Add, modify, and delete configuration files.
- Display side-by-side comparisons of configuration files.

Displaying a Configuration File

1. Select Config > System > Config File.
2. Click on the configuration file name.

Adding a Configuration File

1. Select Config > System > Config File.
2. Click Add. The Config File page appears.
3. Enter the name in the Name field.
4. To use another configuration file as a template, select the file from the Copy drop-down list.
5. Edit the file if required.
6. Click OK.

Modifying a Configuration File

1. Select Config > System > Config File.
2. Click on the configuration file name.
3. Edit the file.
4. Click OK.

Deleting a Configuration File

1. Select Config > System > Config File.
2. Select the checkbox next to each configuration file to delete.
3. Click Delete.

Comparing Configuration Files

1. Select Config > System > Config File.
2. Select the checkbox next to each of the 2 configuration files to compare.
3. Click Diff.

Note: You can compare a maximum of 2 files at a time.

The device configurations appear side-by-side in a new window. Differences between the two configurations are highlighted:

- Yellow – Indicates a configuration section that is present in each device’s configuration, but does not contain exactly the same configuration on both devices.
- Red – Indicates a configuration command that is present in the device configuration shown on the left, but is not present in the device configuration shown on the right.
- Green – Indicates a configuration command that is present in the device configuration shown on the right, but is not present in the device configuration shown on the left.

USING THE CLI

To manage multiple locally stored configurations, use the following commands. All commands are available at the global configuration level of the CLI.

```
write memory
[primary | secondary | profile-name] [cf] |
terminal
```

This command replaces the configuration commands in the specified configuration profile with the commands in the running-config.

If you enter **write memory** without additional options, the command replaces the configuration profile that is currently linked to by startup-config with the commands in the running-config. If startup-config is set to its default (linked to the configuration profile stored in the image area that was used for the last reboot), then **write memory** replaces the configuration profile in the image area with the running-config.

If you enter **write memory primary**, the command replaces the configuration profile stored in the primary image area with the running-config. Likewise, if you enter **write memory secondary**, the command replaces the configuration profile stored in the secondary image area with the running-config.

If you enter **write memory** *profile-name*, where *profile-name* is the name of a configuration profile, the AX device replaces the commands in the specified profile with the running-config.

The **cf** option replaces the configuration profile in the specified image area (primary or secondary) on the compact flash rather than the hard disk. If you omit this option, the configuration profile in the specified area on the hard disk is replaced.

The **terminal** option displays the running-config on the management terminal.

```
show startup-config [all | profile-name] [cf]
```

When entered without the **all** or *profile-name* option, this command displays the contents of the configuration profile that is currently linked to "startup-config". To display the contents of a different configuration profile, use the *profile-name* option. To display a list of the locally stored configuration profiles, use the **all** option.

The **cf** option displays the configuration profile in the specified image area (primary or secondary) on the compact flash rather than the hard disk. If you omit this option, the configuration profile in the specified area on the hard disk is displayed. If the **all** option is also used, the **cf** option displays all the configuration profiles stored on the compact flash.

```
copy {running-config | startup-config |
from-profile-name}
{url | to-profile-name [cf]}
```

The **copy startup-config** *to-profile-name* command copies the configuration profile that is currently linked to “startup-config” and saves the copy under the specified profile name.

The **copy running-config** *to-profile-name* command copies the running-config and saves the copy under the specified profile name.

The **cf** option copies the profile to the compact flash instead of the hard disk.

Note: Copying a profile from the compact flash to the hard disk is not supported.

(The *url* option backs up the configuration to a remote device. See “Backing Up System Information” on page 339.)

diff {**startup-config** | *profile-name*}
{**running-config** | *profile-name*}

Displays a side-by-side comparison of the commands in a pair of configurations.

The **diff startup-config running-config** command compares the configuration profile that is currently linked to "startup-config" with the running-config. Similarly, the **diff startup-config** *profile-name* command compares the configuration profile that is currently linked to "startup-config" with the specified configuration profile.

To compare a configuration profile other than the startup-config to the running-config, enter the configuration profile name instead of s**tartup-config**.

To compare any two configuration profiles, enter their profile names instead of **startup-config** or **running-config**.

In the CLI output, the commands in the first profile name you specify are listed on the left side of the terminal screen. The commands in the other profile that differ from the commands in the first profile are listed on the right side of the screen, across from the commands they differ from. The following flags indicate how the two profiles differ:

- – This command has different settings in the two profiles.
- – This command is in the second profile but not in the first one.
- – This command is in the first profile but not in the second one.

link startup-config {**default** | *profile-name*}
[**primary** | **secondary**] [**cf**]

This command links the "startup-config" token to the specified configuration profile. By default, "startup-config" is linked to "default", which means the configuration profile stored in the image area from which the AX device most recently rebooted.

This command enables you to easily test new configurations without replacing the configuration stored in the image area.

The **primary** | **secondary** option specifies the image area. If you omit this option, the image area last used to boot is selected.

The **cf** option links the profile to the specified image area in compact flash instead of the hard disk.

The profile you link to must be stored on the boot device you select. For example, if you use the default boot device selection (hard disk), the profile you link to must be stored on the hard disk. If you specify **cf**, the profile must be stored on the compact flash. (To display the profiles stored on the boot devices, use the **show startup-config all** and **show startup-config all cf** commands.)

After you link "startup-config" to a different configuration profile, configuration management commands that affect "startup-config" affect the linked profile instead of affecting the configuration stored in the image area. For example, if you enter the **write memory** command without specifying a profile name, the command saves the running-config to the linked profile instead of saving it to the configuration stored in the image area.

Likewise, the next time the AX device is rebooted, the linked configuration profile is loaded instead of the configuration that is in the image area.

To relink “startup-config” to the configuration profile stored in the image area, use the default option (**link startup-config default**).

delete startup-config *profile-name* [**cf**]

This command deletes the specified configuration profile. The **cf** option deletes the profile from compact flash instead of the hard disk.

Note: Although the command uses the **startup-config** option, the command only deletes the configuration profile linked to “startup-config” if you enter that profile’s name. The command deletes only the profile you specify.

Note: If the configuration profile you specify is linked to “startup-config”, “startup-config” is automatically relinked to the default. (The default is the configuration profile stored in the image area from which the AX device most recently rebooted).

CLI Examples

The following command saves the running-config to a configuration profile named “slbconfig2”:

```
AX(config)#write memory slbconfig2
```

The following command shows a list of the configuration profiles locally saved on the AX device. The first line of output lists the configuration profile that is currently linked to “startup-config”. If the profile name is “default”, then “startup-config” is linked to the configuration profile stored in the image area from which the AX device most recently rebooted.

```
AX(config)#show startup-config all
Current Startup-config Profile: slb-v6
Profile-Name                           Size    Time
------------------------------------------------------------
1210test                               1957    Jan 28  18:39
ipnat                                  1221    Jan 25  10:43
ipnat-l3                               1305    Jan 24  18:22
ipnat-phy                              1072    Jan 25  19:39
ipv6                                   2722    Jan 22  15:05
local-bwlist-123                       3277    Jan 23  14:41
mgmt                                   1318    Jan 28  10:51
slb                                    1354    Jan 23  18:12
slb-v4                                12944    Jan 23  19:32
slb-v6                                13414    Jan 23  19:19
```

The following command copies the configuration profile currently linked to “startup-config” to a profile named “slbconfig3”:

```
AX(config)#copy startup-config slbconfig3
```

The following command compares the configuration profile currently linked to “startup-config” with configuration profile “testcfg1”. This example is abbreviated for clarity. The differences between the profiles are shown in this example in bold type.

```
AX(config)#diff startup-config testcfg1
!Current configuration: 13378 bytes                           (
!Configuration last updated at 19:18:57 PST Wed Jan 23 2008   (
!Configuration last saved at 19:19:37 PST Wed Jan 23 2008     (
!version 1.2.1                                                (
!                                                             (
hostname AX                                                   (
!                                                             (
clock timezone America/Tijuana                                (
!                                                             (
ntp server 10.1.11.100                                        (
!                                                             (
...
!                                                             (
interface ve 30                                               (
 ip address 30.30.31.1 255.255.255.0                          |  ip address
10.10.20.1 255.255.255.0
 ipv6 address 2001:144:121:3::5/64                            |  ipv6 address
fc00:300::5/64
!                                                             (
!                                                             (
                                                              > ip nat range-
list v6-1 fc00:300::300/64 2001:144:121:1::900/6
!                                                             (
ipv6 nat pool p1 2001:144:121:3::996 2001:144:121:3::999 netm <
!                                                             <
slb server ss100 2001:144:121:1::100                          <
   port 22  tcp                                               <
--MORE--
```

The following command links configuration profile “slbconfig3” with “startup-config”:

```
AX(config)#link startup-config slbconfig3
```

The following command deletes configuration profile “slbconfig2”:

```
AX(config)#delete startup-config slbconfig2
```

VLAN-to-VLAN Bridging

VLAN-to-VLAN bridging allows an AX device to selectively bridge traffic among multiple VLANs. The AX device selectively forwards packets from one VLAN to another based on the VLAN-to-VLAN bridging configuration on the AX device. This feature allows the traffic flow between VLANs to be tightly controlled through the AX device without the need to reconfigure the hosts in the separate VLANs.

VLAN-to-VLAN bridging is useful in cases where reconfiguring the hosts on the network either into the same VLAN, or into different IP subnets, is not desired or is impractical.

You can configure a bridge VLAN group to forward one of the following types of traffic:

- IP traffic only (the default) – This option includes typical traffic between end hosts, such as ARP requests and responses.

  This option does not forward multicast packets.
- All traffic – This option forwards all types of traffic.

Configuration Notes

VLAN-to-VLAN bridging is supported on AX devices deployed in transparent mode (Layer 2) or in gateway mode (Layer 3).

Each VLAN to be bridged must be configured on the AX device. The normal rules for tagging apply:

- If an interface belongs to only one VLAN, the interface can be untagged.
- If the interface belongs to more than one VLAN, the interface must be tagged.

Each VLAN can belong to only a single bridge VLAN group.

Each bridge VLAN group can have a maximum of 8 member VLANs. Traffic from any VLAN in the group is bridged to all other VLANs in the group. Up to 64 bridge VLAN groups are supported.

If the AX device is deployed in gateway mode, a Virtual Ethernet (VE) interface is required in the bridge VLAN group.

Configuring VLAN-to-VLAN Bridging

To configure VLAN-to-VLAN bridging:

1. Configure each of the VLANs to be bridged. In each VLAN, add the AX device's interfaces to the VLAN.

2. Configure a bridge VLAN group. Add the VLANs to the group.

   If the AX device is deployed in gateway mode, add a Virtual Ethernet (VE) interface to the group.

   Optionally, you can assign a name to the group. You also can change the types of traffic to be bridged between VLANs in the group.

3. If the AX device is deployed in gateway mode, configure an IP address on the VE to place the AX device in the same subnet as the bridged VLANs.

USING THE CLI

To configure a bridge VLAN group, use the following commands.

```
[no] bridge-vlan-group group-num
```

Use this command at the global configuration level of the CLI to create the bridge VLAN group and enter the configuration mode for it, where the following commands are available. The *group-num* can be 1-64.

```
[no] name string
```

This command configures a name for the group. The string can be 1-63 characters long. If the string contains blank spaces, use double quotation marks around the entire string.

```
[no] vlan vlan-id
  [vlan vlan-id ... | to vlan vlan-id]
```

This command adds the VLANs to the group.

```
[no] router-interface ve num
```

On an AX device deployed in gateway mode, this command adds the VE to the group.

```
forward-all-traffic
```

This command configures the AX device to forward all types of traffic between the VLANs in the group. By default, only IP traffic is forwarded. If you change the traffic type but later want to change it back to the default, you can use the following command: **`forward-ip-traffic`**

Group Information and Statistics

To display information for a bridge VLAN group, use the following command:

show bridge-vlan-group [*group-id*]

CLI Example – Transparent Mode

The commands in this section configure an AX device deployed in transparent mode to forward IP traffic between VLANs 2 and 3.

The following commands configure the VLANs:

```
AX(config)#vlan 2
AX(config-vlan:2)#tagged ethernet 2
AX(config-vlan:2)#exit
AX(config)#vlan 3
AX(config-vlan:3)#tagged ethernet 3
AX(config-vlan:3)#exit
```

The following commands configure the bridge VLAN group:

```
AX(config)#bridge-vlan-group 1
AX(config-bridge-vlan-group:1)#vlan 2 to 3
AX(config-bridge-vlan-group:1)#exit
```

CLI Example – Gateway Mode

The commands in this section configure an AX device deployed in gateway mode to forward IP traffic between VLANs 2 and 3 on IP subnet 192.168.1.x.

The following commands configure the VLANs:

```
AX(config)#vlan 2
AX(config-vlan:2)#tagged ethernet 2
AX(config-vlan:2)#exit
AX(config)#vlan 3
AX(config-vlan:3)#tagged ethernet 3
AX(config-vlan:3)#exit
```

The following commands configure the bridge VLAN group, which includes a VE:

```
AX(config)#bridge-vlan-group 1
AX(config-bridge-vlan-group:1)#vlan 2 to 3
AX(config-bridge-vlan-group:1)#router-interface ve 1
AX(config-bridge-vlan-group:1)#exit
```

The following commands assign an IP address to the VE:

```
AX(config)#interface ve 1
AX(config-if:ve1)#ip address 192.168.1.100 /24
AX(config-if:ve1)#exit
```

A10 Networks

Customer Driven Innovation

Corporate Headquarters

A10 Networks, Inc.
3 West Plumeria Dr.
San Jose, CA 95134 USA

Tel: +1-408-325-8668 (main)
Tel: +1-408-325-8676 (support - worldwide)
Tel: +1-888-822-7210 (support - toll-free in USA)
Fax: +1-408-325-8666

www.a10networks.com