



3Com® X Family Concepts Guide



X5 (25-user license) – 3CRTPX5-25-96
X5 (unlimited license) – 3CRTPX5-U-96
X506 – 3CRX506-96

Version 3.0

Part Number 10016442
Published November 2007
<http://www.3com.com/>



3Com Corporation
350 Campus Drive
Marlborough, MA
01752-3064

Copyright © 2005–2007, 3Com Corporation. All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from 3Com Corporation.

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms, or conditions of merchantability, satisfactory quality, and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hardcopy documentation, or on the removable media in a directory file named LICENSE.TXT or !LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

UNITED STATES GOVERNMENT LEGENDS:

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following:

United States Government Legend: All technical data and computer software is commercial in nature and developed solely at private expense. Software is delivered as Commercial Computer Software as defined in DFARS 252.227-7014 (June 1995) or as a commercial item as defined in FAR 2.101(a) and as such is provided with only such rights as are provided in 3Com's standard commercial license for the Software. Technical data is provided with limited rights only as provided in DFAR 252.227-7015 (Nov 1995) or FAR 52.227-14 (June 1987), whichever is applicable. You agree not to remove or deface any portion of any legend provided on any licensed program or documentation contained in, or delivered to you in conjunction with guide.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, the 3Com logo, TippingPoint, the TippingPoint logo, and Digital Vaccine are registered trademarks of 3Com Corporation or one of its subsidiaries.

OpenView is a trademark of Hewlett-Packard Development Company. Microsoft and Windows are registered trademarks or trademarks of Microsoft Corporation in the United States and other countries. Oracle is a registered trademark of Oracle Corporation.

Other brand and product names may be registered trademarks or trademarks of their respective holders.

Contents

About This Guide	ix
Welcome to the X Family Concepts Guide	ix
Target Audience	ix
Organization	x
Conventions	xi
Cross References	xi
Typeface	xii
Procedures	xii
Screen Captures	xii
Messages	xii
Related Documentation	xiii
Customer Support	xiv
Contact Information	xiv
Chapter 1 System Overview	1
Chapter 2 Key Concepts	3
Understanding Security Zones	3
Default Zones	4
Applying Security Zones	6
Firewall Configuration for Security Zones	8
IPS Filter Configuration for Security Zones	9
IP Interfaces	10
Types of IP Interfaces	10
Deployment Modes Implementation Methods	11
Deployment Examples	12
Using VLANs	17
Traffic Shaping and Bandwidth Management	17
Multicast and Dynamic Routing	19
Multicast Routing	19
Dynamic Routing	20
High Availability	21
How High Availability Works	21
Sample Configuration	23
Enforcing a Web Access Policy	24

Chapter 3 Network	25
Introduction	25
Security Zones	25
Dividing the Network into Security Zones	25
Default Security Zones	27
Advanced Security Zone Configuration	27
Configuring IP Interfaces	27
External Interface	28
Internal (LAN) Interface	28
GRE Interface	28
Configuring IP Interfaces	29
Configuring DHCP Server Settings	29
Using the DHCP Server	30
Using DHCP Relay	30
Using DHCP Relay over VPN	31
DHCP Client	32
Assigning a Static IP Address	32
Configuring IP Address Groups	32
Configuring Network Address Translation (NAT)	33
Setting Up Many-to-One NAT	34
Setting Up One-to-One NAT	34
Setting up NAT Within a VPN Tunnel	35
Configuring Routing	37
Default Route	37
Static Routing	37
Dynamic Routing	37
Multicast Routing	38
Network Tools	39
Chapter 4 Virtual Private Networks	41
Introduction to Virtual Private Networks	41
What is a VPN?	41
Benefits Of VPNs	42
Types of VPN Connections	42
X Family VPN Security Features	42
Tunnel Security Zones	42
VPN Connection Security Features	43
Site-To-Site VPNs	44
IPSec Modes	44
Summary of Site-To-Site VPN Methods	45
Security Association	46
IPSec Security Mechanisms	46
Keys and Keying Modes	46
Encryption and Data Integrity	48

IKE Proposals	49
IPSec Tunnel Setup	49
Site-to-Site VPN Operation	50
Client-to-Site VPNs	50
Supported Protocols	50
Summary of Client-To-Site VPN Methods	52
Using the Default Security Association	52
Configuring Client-to-Site User Authentication	53
Client-to-Site VPN Operation	53
Compatibility	53
Advanced VPN Configuration	54
Setting up NAT Within a VPN Tunnel	54
Setting Up a VPN Supernet	55
Configuring Your Network	57

Chapter 5 Firewall Rules 59

Overview	59
Firewall Rules Rank	60
How Firewall Rule Enforcement Works	60
X Family Firewall Components	63
Services	63
Service Groups	63
Schedules	63
Source and Destination Addresses	63
Firewall Actions	64
Creating a VPN Firewall Rule	64
Managing Bandwidth	65
Defining the Maximum Bandwidth	66
User Authentication	66
Other Firewall Options	67
Providing Access to Internal Servers	68
Firewall Rule Example	68
Usage	68
Setup	68
Implementation	69

Chapter 6 Web Content Filtering 71

Overview	71
Web Filter Profiles	71
Web Content Filtering Service	72
Manual Filtering	72
Web Content Filtering Components	72
Filtering Actions	72
Default Rule	72

Custom Response Page	73
URL Lists	73
How Web Content Filtering Works	74
Web Content Filtering Configuration Example	76
Background	76
Setup	76
Implementation	77
Chapter 7 Anti-Spam Filtering	79
Manual Filtering	80
Chapter 8 Intrusion Prevention System 81	
Overview	81
Configuring IPS	82
Security Profiles	82
IPS Filters	84
Action Sets	88
Notification Contacts	89
IPS Services	91
Chapter 9 User Authentication	93
Overview	93
How Local User Authentication Works	94
Applications of Authentication	95
Authentication and Firewall Privileges	96
Types of Firewall Privileges	96
Privilege Groups	96
Methods of Authenticating Local Users	97
Using the Local Device Database	97
Using an LDAP Authentication Server	98
Using a RADIUS Authentication Server	98
Configuration Example	98
Implementation Example	99
Chapter 10 Certificates	101
Overview	101
What are X.509 Certificates?	101
Certificates and Public Key Cryptography	102
Public Key Infrastructure	103
Digital Signatures	104
Contents of a Certificate	104

Certificate Revocation List (CRL)	105
Setting Up Your Certificate Infrastructure	106
Installing a CA Server	106
Installing Local Certificates on Your VPN Clients	106
Methods of Obtaining Certificates	106
Sending a Certificate Request to a CA	106
Directly Importing a Certificate	107
Creating Your Own Self-Signed Certificate	107
Installing and Managing Certificates	107
Certificate Configuration Example	107
Certificate Setup	107
VPN Setup	108
Implementation	109

Chapter 11 Events: Logs, Traffic Streams, and Reports **111**

Logs	112
Alert Log	112
Audit Log	113
IPS Block Log	113
Firewall Block Log	113
Firewall Session Log	113
VPN Log	114
System Log	114
Configuring Remote System Logs	114
Managed Streams	115
Health	115
Reports	116

Chapter 12 Deployment Scenarios **121**

Introduction	121
Examples of VPN Network Topologies	123
X Family Device in a Hub and Spoke Deployment	123
X Family Device in a Meshed Deployment	124
X Family Device in a Tree Deployment	125
X Family Device in a Mixed Deployment	126
Deployment Scenario Example	127
Description	127
Application	127
Network Topology	128
Security Zone Configuration	129
VPN Configuration	132
Creating Virtual VPN Tunnel Zones	133

Table of Contents

Firewall Rule Configuration	133
Multicast NBX Conference Calling	133
NBX Setup	133
Using Certificates to Ensure Security	
over a Public Network	135
Description	135
User Access Controlled by Firewall Policies	136
Description	136
Overview of the Web Content Filter Service	139
Core Categories	140
Productivity Categories	142
Purchasing a Web Filter License	148
Overview of the Anti-Spam Service	150
Purchasing an Anti-Spam License	150

About This Guide

Explains who this guide is intended for, how the information is organized, where information updates can be found, and how to obtain customer support if you cannot resolve a problem.

Welcome to the X Family Concepts Guide

This guide provides an overview of the terminology, concepts, setup, and configuration procedures that apply to the X Family of Unified Security Platforms and provides background information on each area of configuration. For more information about configuring the device using a web interface, see the *Local Security Manager User's Guide* or the online help. For information about a command-line interface, see the *Command Line Interface Reference*.

For more about this guide and how to use the X family documentation, see the following topics:

- [“Target Audience” on page ix](#)
- [“Conventions” on page xi](#)
- [“Related Documentation” on page xiii](#)
- [“Customer Support” on page xiv](#)

Target Audience

This guide is intended for users who manage one or more X Family devices.

Knowledge, Skills, and Abilities

This guide is written assuming that you are familiar with general networking concepts and the following standards and protocols:

- Transmission Control Protocol/Internet Protocol (TCP/IP)
- User Datagram Protocol (UDP)
- Internet Control Message Protocol (ICMP)
- Ethernet
- Network Time Protocol (NTP)
- Simple Mail Transport Protocol (SMTP)
- Simple Network Management Protocol (SNMP)

Organization

The *Concepts Guide* is organized as follows:

About This Guide

Explains who this book is intended for, how the information is organized, where information updates can be found, and how to obtain customer support if you cannot resolve a problem.

Key Architecture Concepts

Introduces key concepts for the X Family of Unified Security Platforms and their application in various deployments including security zones, IP interfaces, traffic shaping and management, and multicast and dynamic routing.

Network

Describes interfaces, security zones, DHCP functionality, routing, and IP address groups, and explains how to enable, disable, and modify their various features. It also describes the network management tools provided.

Virtual Private Networks

Provides an overview of Virtual Private Networks (VPNs) and describes how they are implemented using X family devices, including security features, site-to-site VPNs, client-to-site VPNs, and advanced VPN configuration.

Firewall Rules

Provides an overview of the X family firewall rules and describes the steps required to configure them. The following topics are included: firewall components, providing access to internal servers, setting up firewall policies, and an implementation example.

Intrusion Prevention System

Provides an overview of the Intrusion Prevention System (IPS) software.

User Authentication

Provides an overview of user authentication and describes how to configure user authentication on the X family device. The following topics are included: authentication and firewall privileges, methods of authentication methods for local users, user authentication setup, and an example implementation.

Certificates

Introduces certificates and public key cryptography concepts, describes methods of obtaining certificates, and provides an example illustrating how to configure certificates.

Web and Spam Filtering

Describes filtering on X family devices and the steps required to configure it including the following topics: Web content filtering components, spam filtering components, and purchasing filtering licenses.

Logging

Describes logging on X family devices and provides descriptions and examples of the available logs.

Deployment Scenarios

The X family device is multi-functional and can be deployed in a variety of scenarios. This chapter provides examples of common deployment scenarios. The following topics are included: VPN network topologies, multicast conference calling, ensuring security over a public network, and controlling user access with firewall policies.

Conventions

This guide follows several procedural and typographical conventions to better provide clear and understandable instructions and descriptions.

These conventions are described in the following sections:

- [“Cross References” on page xi](#)
- [“Typeface” on page xii](#)
- [“Procedures” on page xii](#)
- [“Messages” on page xii](#)

Cross References

When a topic is covered in depth elsewhere in this guide, or in another guide in this series, a cross reference to the additional information is provided. Cross references help you find related topics and information quickly.

Internal Cross References

This guide is designed to be used as an electronic document. It contains cross references to other sections of the document that act as hyperlinks when you view the document online. The following text is a hyperlink: [Procedures](#).

External Cross References

Cross references to other publications are not hyperlinked. These cross references will take the form: see <chapter name > in the *Publication Name*.

Typeface

This guide uses the following typeface conventions:

Bold	used for the names of screen elements like buttons, drop-down lists, or fields. For example, when you are done with a dialog, you would click the OK button. See “Procedures” below for an example.
Code	used for text a user must type to use the product
<i>Italic</i>	used for guide titles, variables, and important terms
Hyperlink	used for cross references in a document or links to a Web site

Procedures

This guide contains several step-by-step procedures that tell you how to perform a specific task. These procedures always begin with a phrase that describes the task goal, followed by numbered steps that describe what you must do to complete the task.

The beginning of every chapter has cross references to the procedures that it contains. These cross references, like all cross references in this guide, are hyperlinked.

Screen Captures

The instructions and descriptions in this document include images of screens. These screen captures may be cropped, focusing on specific sections of the application, such as a pane, list, or tab. See the application for full displays of the screen.

Messages

Messages are special text that are emphasized by font, format, and icons. There are four types of messages in this guide:

- [Warning](#)
- [Caution](#)
- [Note](#)
- [Tip](#)

A description of each message type with an example message follows.

Warning

Warnings tell you how to avoid physical injury to people or equipment. You should carefully consider this information prior to enacting actions or procedures that could potentially harm your staff, data, or security.



WARNING: The push-button on/off power switch on the front panel does not turn off the AC power. To remove AC power, you must unplug the AC power cord from either the power supply or the wall outlet.

Caution

Cautions tell you how to avoid a serious loss that could cause physical damage such as the loss of data, time, or security. You should carefully consider this information when determining a course of action or procedure.



CAUTION: You should disable password caching in the browser you use to access the LSM. If you do not disable password caching in your browser, and your workstation is not secured, your system security may be compromised.

Note

Notes tell you about information that might not be obvious or that does not relate directly to the current topic, but that may affect relevant behavior.



Note: If the device is not currently under SMS control, you can find out the IP address of the last SMS that was in control by checking your Audit log from the Logs page.

Tip

Tips are suggestions about how you can perform a task more easily or more efficiently.



Tip: You can see what percentage of disk space you are using by checking the Monitor page.

Related Documentation

The X Family of Unified Security Platforms has a full set of documentation. These publications are available in electronic format. For the most recent updates, check the 3Com Web site at <https://www.3Com.com>.

Customer Support

We are committed to providing quality customer support to all customers. A customer is provided with detailed customer and support contact information. For the most efficient resolution of your problem, please take a moment to gather some basic information from your records and your device before contacting customer support.

Information	Location
Your X family device serial number	You can find this number using the Local Security Manager (LSM) in the System Summary page, on the shipping invoice that came with the device, or on the bottom of the device.
Your TOS version number	You can find this information using the LSM in the System Summary page, or by using the CLI <code>show version</code> command.
Your X family device system boot time	You can find this information using the LSM in the System Summary page.

Contact Information

Please address all questions regarding the software to your authorized representative.

1 System Overview

The X Family of Unified Security Platforms enhances organizational flexibility and market responsiveness by addressing the network connectivity, security, and manageability needs of a multi-site, distributed enterprise. This unique solution provides rich user and application connectivity across the wide area network without compromising enterprise security or manageability.

X family devices provide the following functionality:

- **Stateful packet inspection firewall** — Flexible configuration of object-based firewall rules and unified control of multiple services, virtual servers, network address translation (NAT), and routing.
- **Security Zones** — Logically section your network for the purposes of applying firewall rules and IPS filters between internal sections of your network, between your network and the Internet, and between your network and remote office locations. The X family device can function as a software bridge to transparently connect security zones assigned to the same virtual interface.
- **Support for standards-based IPSec Virtual Private Networks (VPNs)** including:
 - o Hardware-accelerated encryption using the DES, 3DES, and AES encryption protocols
 - o Feature-rich client VPN capability using the PPTP or L2TP protocols
 - o Inspection and control of traffic both inside and outside of all VPN tunnel types using firewalls or the IPS software to ensure secure VPN connectivity
 - o Alternate VPN peers
- **Flexible user authentication** — Control access to the device and the Internet, authenticating via the device itself or through an external LDAP or RADIUS database.
- **Web content filtering** — URL filtering with configurable permit/block lists and regular-expression URL matching as well as a Web Filtering subscription service to enforce network security and usage policy by prohibiting the download of non-work related Web sites and offensive or illegal Web content.
- **Spam filtering** — Anti-Spam subscription filtering service that provides email filtering based on pattern recognition.
- **Bandwidth management** — Enforce network usage policy by rate-limiting applications such as peer-to-peer file sharing and instant messaging.

- **Zone-based Rate Limiting** — Prioritization of traffic inside and outside VPN tunnels with flexible, policy-based controls.
- **IP multicast routing (Protocol Independent Multicast — Dense Mode (PIM-DIM)) over IPSec, supporting next-generation IP conferencing applications** — Prioritizes real-time traffic and provides secure connectivity for IP multicast traffic.
- **Local Management** — Option to configure, monitor, and manage the device using either a web-based client application (the Local Security Manager) or a command-line interface (CLI).
- **Centralized Management** — Option to configure, monitor, and manage individual or multiple devices using the Security Management System (SMS) or third-party network management systems (NMSs).
- **Intrusion Prevention System (IPS)** — Identify and stop malicious traffic on the edge of the network using filters that detect and block malicious traffic. Customize default filters to meet the specific needs of your enterprise.
- **Digital Vaccine real-time protection** — The Threat Management Center monitors global network security threats and continually develops new attack filters which are automatically distributed on a subscription basis to preemptively protect against the exploit of new and zero-day vulnerabilities.
- **High Availability** — Prevent service interruptions by configuring a fallback device that will continue to pass network traffic in the event of any internal hardware or software failure on the primary device. High availability configurations can be synchronized across pairs of devices.

This guide explains concepts in the X family architecture and provides background information on each area of configuration. For more information about configuring, managing, and monitoring the device using the Local Security Manager, see the *Local Security Manager User's Guide* or the online help. For information about the command-line interface, see the *Command Line Interface Reference*.

2 Key Concepts

This chapter introduces concepts required to understand, deploy, and use X family devices. It covers the following topics:

- [“Understanding Security Zones” on page 3](#)
- [“IP Interfaces” on page 10](#)
- [“Traffic Shaping and Bandwidth Management” on page 17](#)
- [“Multicast and Dynamic Routing” on page 19](#)
- [“High Availability” on page 21](#)

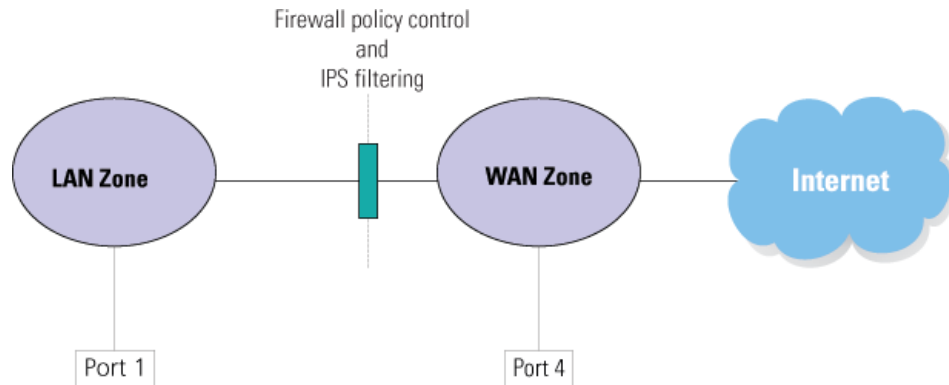
Understanding Security Zones

Security zones let you logically section your network so that the X family device recognizes each network section as a separate zone. Devices within a security zone are able to communicate freely with each other; policy (firewall rules, IPS filters, and Web content filtering) is not applied. After setting up the security zones required for your network, you configure the firewall rules, IPS filters, and Web content filters to apply to traffic passing between the zones. In addition, the device supports bridge mode, in which it functions as a software bridge to transparently connect security zones assigned to the same IP subnet.

Security zones let you logically segment your networks so that the device can apply firewall rules and IPS filters to control the traffic passing between the zones. Typically, each Ethernet port and VPN tunnel on the device is associated with one security zone, unless you use VLANs. If you configure VLANs, then a port can be in more than one security zone. Any traffic originating from or destined to devices in a zone will be directed through the device and policed by firewall policies if the traffic passes through to another zone. However, traffic moving between devices within a given zone that you have defined (intra-zone traffic) will not be subject to firewalling or IPS filtering (for example, a user on the LAN zone, accessing the local LAN printer) and will not pass through the device.

Figure 2–1 illustrates the point at which monitoring occurs for traffic passing between two security zones, the LAN zone and the WAN zone:

Figure 2–1: Simple Zone Configuration



Security zones can be associated with a physical port, or they can exist virtually by logical definition.

- Physical security zones are mapped to a single Ethernet port unless you use VLAN tagging, which allows you to assign different VLANs on the same port to different security zones. For details, see [“Using VLANs” on page 17](#).
- Virtual security zones are not mapped to a port. The default configuration defines two virtual zones: *this-device* and VPN. The *this-device* zone is used to control access to the X family device and to administer the device from a secure web interface. The VPN zone is used to apply policy to traffic coming in and out of a VPN tunnel.

Default Zones

The X family device is configured with the following default security zones:

Table 2–1: Preconfigured Zones

Preconfigured Zones	Ports
LAN	Port 1 (or “LAN”)
WAN	Port 6 (or “WAN”)
VPN	None (Virtual)

The *this-device* Zone

The *this-device* zone is preconfigured with firewall policy permissions that let you control the access of traffic to and from the device and administer the device. This zone is not configurable. Firewall rules can use *this-device* as a destination or source zone to allow you to control the traffic to and from the device. For example, to allow secure device management from the LAN, you must have a firewall rule that allows the LAN security zone access to the *this-device zone* for secure web management. See [Chapter 5, “Firewall Rules”](#) for more information about using firewall rules to enforce policies.

VPN Tunnel Zones


A VPN tunnel must be associated with a security zone. This can be either an existing zone, such as the LAN zone, or a virtual tunnel zone, which is not mapped to any ports on the X family device.

If you want to apply firewall policies to your VPN traffic, the VPN tunnel must be placed in a separate virtual tunnel zone. You can use the preconfigured VPN tunnel zone or create your own. You can also configure the VPN tunnel to use NAT so that multiple remote VPN sites can use the same IP subnet. You can then apply firewall policy control between the remote site traffic tunneling into the virtual zone and other zones, such as your LAN zone.

If the VPN tunnel terminates within an existing LAN zone, then no firewall policies can be applied to the VPN traffic, as it is in the same zone. This is shown in [Figure 2-2](#):

Figure 2-2: VPN Tunnel Terminating in LAN Zone

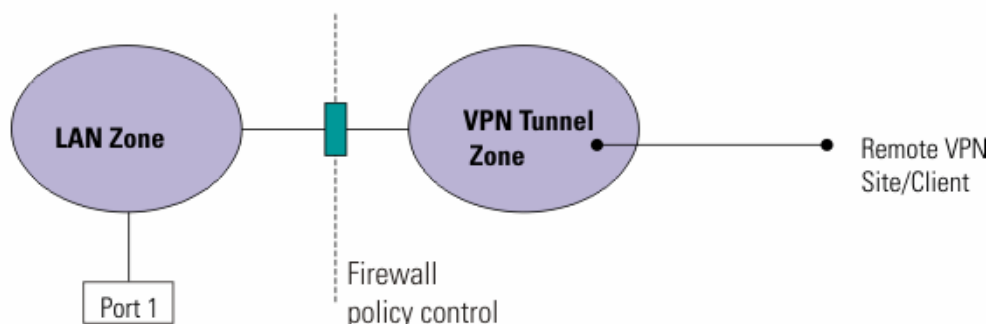
VPN tunnel terminates in existing zone, without policy control:



The use of a special VPN tunnel zone, which has no Ethernet ports assigned to it, allows firewall policy control between zones, as shown in [Figure 2-3](#):

Figure 2-3: VPN Tunnel Terminating in Dedicated VPN Zone

VPN tunnel terminates in virtual tunnel zone, with firewall policy control:



See [Chapter 3, "Network"](#) for more information on security zone configuration.

Applying Security Zones


You can use security zones to provide security partitioning to protect network data and users, both within and outside the network. This section describes three typical security zone applications:

- [“Policing the Network Perimeter” on page 6](#)
- [“Policing the Internal Network” on page 6](#)
- [“Supplementing a LAN Switch” on page 7](#)

Policing the Network Perimeter

In a typical deployment scenario, the device functions as a firewall and/or VPN gateway, placed on the edge of a network. The device applies firewall policy rules, IPS, and Web content filtering to all traffic between the LAN and WAN zones. It polices outbound traffic requests from the LAN zone on the network (for example, to apply Web content filtering to HTTP/HTTPS traffic or deny requests for restricted services), as well as inbound requests from remote devices. [Figure 2–4](#) shows this typical scenario:

Figure 2–4: Typical Zone Deployment — X Family Device as a Perimeter Firewall

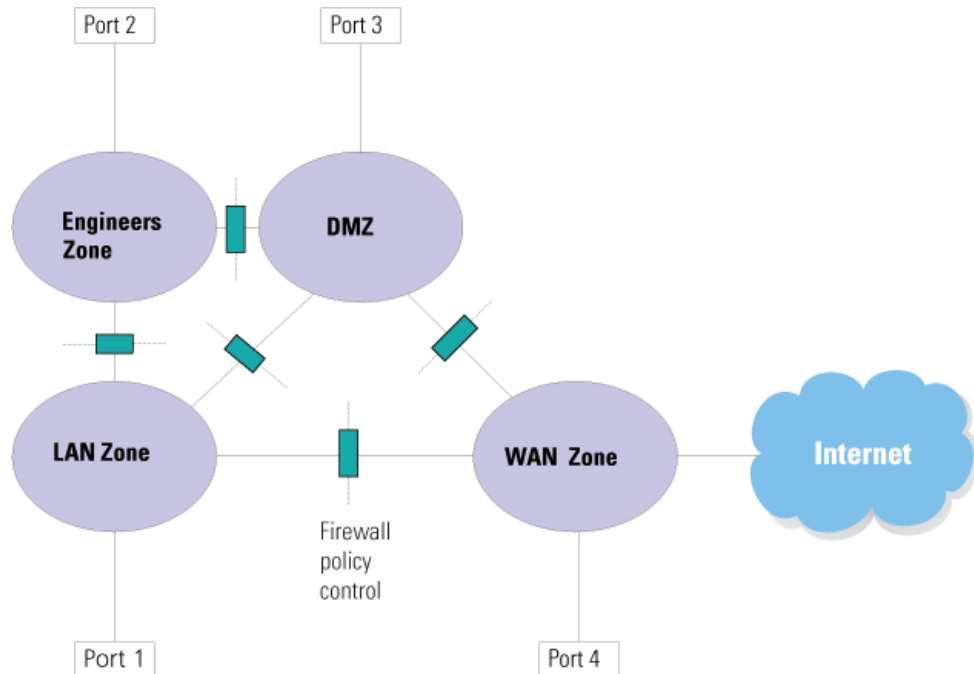


Policing the Internal Network

In some cases, you may want to monitor traffic within your network using firewall policy rules. For example, you may want to control access to a section of your network (such as a secure server section) or to services available to specific types of users or devices on your network (such as guest users and wireless devices). The X family device enables you to use security zones to restrict access, while at the same time allowing users or devices with authenticated access privileges (such as network

administrators) to access the secure sections of your network. A network where the device provides policy control within the LAN as well as between the LAN and the Internet is shown in [Figure 2-5](#):

Figure 2-5: Typical Zone Deployment Scenario — X Family Device for Policing Intranet Traffic



IPS Filtering Between Internal Network Sections

In addition to applying firewall policy rules to your internal network, you can also apply IPS filtering. In the network previously shown ([Figure 2-5](#)), IPS filtering can be configured to protect devices in the network from Internet security threats by setting up IPS filters to operate at every firewall policy control point shown. IPS filtering is most important between an internal zone and the WAN zone; however, it can also be applied between internal zones for extra protection from security threats.

Supplementing a LAN Switch

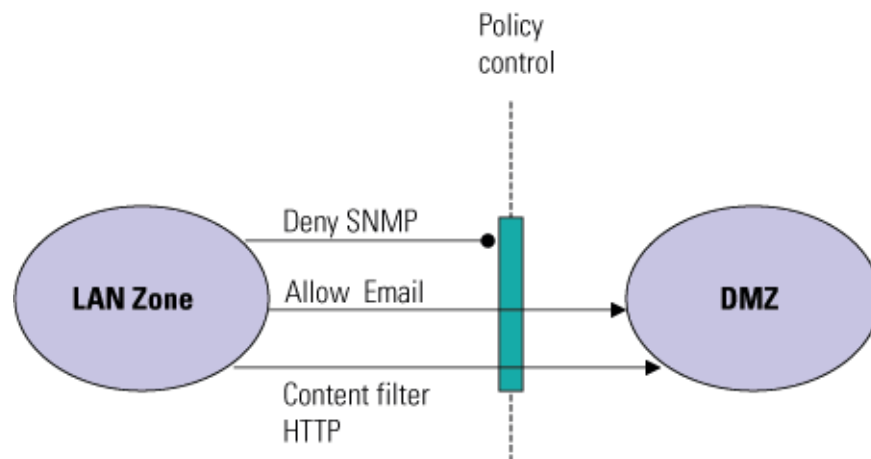
The X family device can be used to supplement an existing LAN switch without adding additional equipment. Using security zones, you can interconnect sections of the network and then monitor and direct traffic between the zones based on firewall rules and IPS filtering. This configuration provides policy control and management between sections of your internal network without requiring a separate infrastructure for each control point, improving security and management.

Firewall Configuration for Security Zones

Configuration of firewall rules for the security zones within the network enables you to control the types of services available to users and devices within any given zone.

[Figure 2–6](#) illustrates traffic shaping through firewall policy control between the LAN zone and the DMZ. Certain types of traffic are allowed, such as email and Web access, while access by other types of applications (for example, SNMP) are blocked.

Figure 2–6: Security Zones and Firewall Policies



Firewall rules can also be configured to prioritize traffic passing between zones to ensure quality of service for time-critical applications such as Voice over IP. While allowing email and Web traffic, subscription-service filters can block incoming email spam and access to various categories of Web sites.

To optimize network performance, place devices that need to communicate frequently, without firewall policy control, in the same zone. Note also that firewalling is only effective for IP traffic — all other types of traffic will be denied, if passed through the device.

See [Chapter 5, "Firewall Rules"](#) for more details on firewall policy configuration.

IPS Filter Configuration for Security Zones

Configuration of IPS filtering for network security zones enables you to protect the network from Internet security risks. IPS filtering is set up by establishing security profiles.

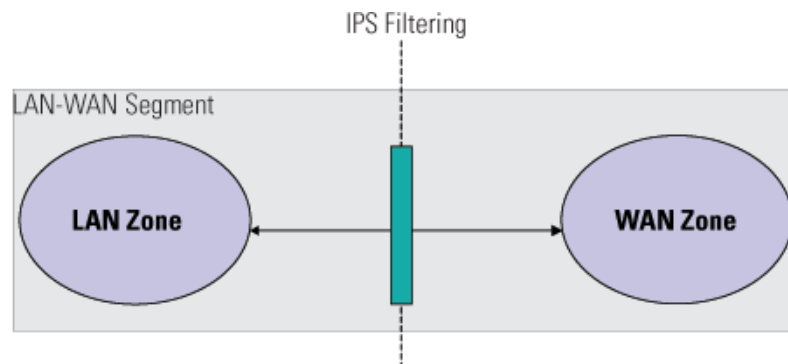
A security profile specifies the following:

- The traffic and direction of traffic to be monitored based on security zone pairs. For example, if a profile specifies the zone pair LAN ==> WAN, IPS filtering will be applied to all traffic going from the LAN to the WAN.
- The IPS filters to apply to the traffic.

Because the security zone pairs are directional, you can apply different IPS filtering rules to incoming and outgoing traffic. To apply the same rules, the security profile would have to include both pairs, for example LAN ==> WAN and WAN ==> LAN.

[Figure 2-7](#) illustrates how a device with a security profile configured for the LAN==>WAN and WAN==>LAN security zone pairs applies IPS filtering to all traffic passing between the LAN and WAN zones. All IPS filtering is performed on traffic that has been permitted through the X family firewall.

Figure 2-7: Security Zones and the IPS



Default Security Profile

Each device is configured with a default security profile of ANY ==> ANY. This profile applies the default IPS filter configuration to all traffic passing between any security zones configured on the device. If only the default zones are defined, the filters will operate on all traffic going from the LAN to the WAN and traffic going from the WAN to the LAN.

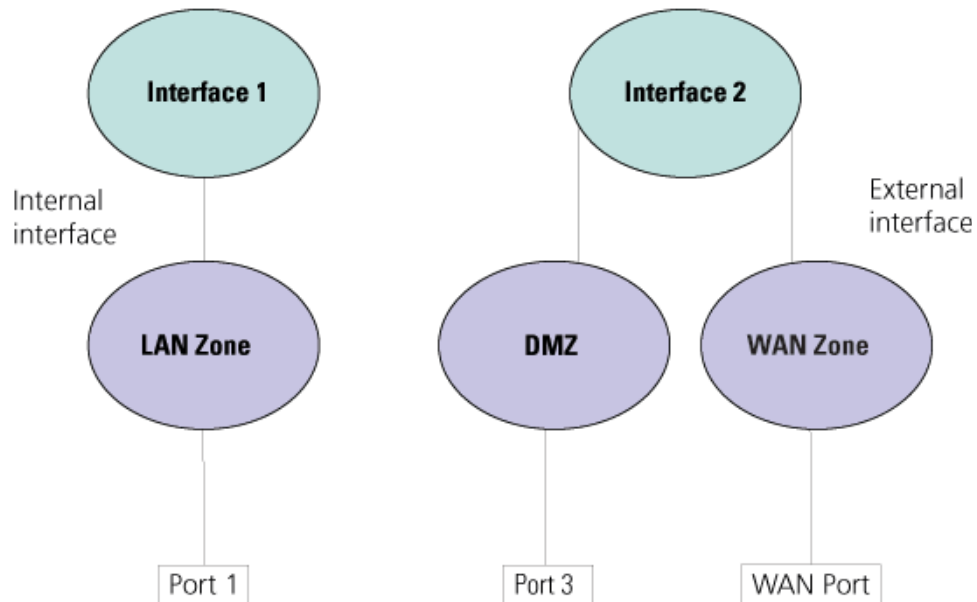
If you delete the default security profile and do not define any other security profiles, the device will not apply IPS filtering to any traffic. In addition, if you delete the ANY ==> ANY zone pair from a profile, any traffic that is not matched by another security zone pair will not be filtered by the IPS. If you create new security zones, or want IPS filtering applied to traffic to other zones such as LAN2 and LAN3 zones, you must create a new security profile, or add the zone pairs to the existing one.

See [Chapter 5, "Firewall Rules"](#) for more details on IPS filter types configuration.

IP Interfaces

IP interfaces (virtual interfaces) define how the device integrates with the IP addresses on your network (that is, the layer 3 network). [Figure 2–8](#) provides an overview of interfaces, zones, and ports.

Figure 2–8: Security Zones and IP Interfaces



You must configure one IP interface for every IP subnet that is directly connected to the device. For example, you need one for an Internet connection (external interface) and one for every directly connected network subnet (internal interfaces). If you are using site-to-site dynamic routing or multicasting, this also requires an IP interface (a Generic Route Encapsulation (GRE) tunnel interface).

Each IP interface is identified on the device by a number; for example: 1 (Internal), 2 (External), and so on.

Types of IP Interfaces

There are two types of IP interface that you can configure on an X family device:

- External
- Internal

External Interfaces

An external interface is a connection between the device and the Internet. It typically requires a public IP address. There are a number of methods by which an external interface can be allocated its IP address, usually dictated by your ISP (Internet Service Provider), including Static IP Address, DHCP, PPPoE client, PPTP client, L2TP client, DNS, and dynamic DNS.

The device supports up to two external interfaces. Using multiple external interfaces allows for load balancing and failover. Load balancing uses two external interfaces simultaneously to route traffic. Failover uses the second external interface as a backup link in case the primary link fails.

Internal (LAN) Interfaces

An internal (LAN) interface of the X family device is a connection between the device and a LAN. Internal interfaces of the device use static IP addressing. When changing an internal (LAN) interface IP address, choose an address that will be unique in your network and in your network's subnet.

You can also configure the internal (LAN) interface to share the IP address of the external interface. In this case, you will only have one IP interface (the external interface); your network and connections are operating in transparent mode, and direct connections from your network are on the same subnet.

If your network is routing internally (that is, there is more than one IP subnet directly connected to the device), you must configure one internal interface for each subnet. You can also use network address translation (NAT) to hide an IP subnet.

See [Chapter 3, "Network"](#) for more information on setting up and configuring interfaces.

Deployment Modes Implementation Methods

In an X family deployment, there are normally three ways in which you can implement zoning, depending on your current network configuration:

- Transparent deployment
- Routed deployment
- Full routed/NAT deployment

Transparent Deployment

In transparent deployment, the device acts like a Layer 2 switch. Effectively, the same IP subnet is used by all zones within a single IP interface. All devices in these zones share the same IP address space, which means that you only have a single IP interface for all zones that are in the same transparent group. Firewall policy can be applied between these zones.

Routed Deployment

In a routed deployment, the network is divided into multiple IP subnets. In this case, each security zone is associated with a unique IP interface so that the devices within each zone have a unique IP address space.

X family security zones can be configured on networks that use a combination of both transparent and routed deployments. In this scenario, security zones that are in the same IP interface would use transparent mode, while those security zones that are in different IP interfaces would use routed mode.

Full Routed/NAT

Network Address Translation (NAT) is used in typical network deployments to map LAN addresses to the external address of the device. In a full routed/NAT deployment, all security zones have unique IP

addresses and addresses going to the WAN zone may be translated. Each security zone is in a separate broadcast zone.

You can also use NAT within a VPN tunnel. See [Chapter 12, "Deployment Scenarios"](#) for details.

For deployment examples, see the following:

- [“Example 1: Transparent DMZ Deployment” on page 12](#)
- [“Example 2: Routed Deployment” on page 14](#)
- [“Example 3: Fully Transparent Deployment” on page 15](#)
- [“Example 4: NAT Deployment on a Switched LAN” on page 16](#)

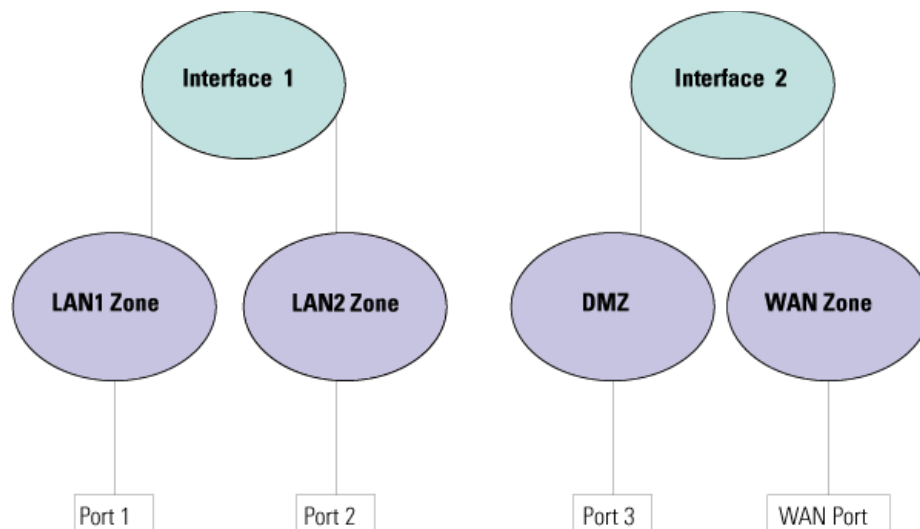
Deployment Examples

Example 1: Transparent DMZ Deployment

In this example, shown in [Figure 2–9](#), the X family device is:

- firewalling between the LAN and the Internet
- policing traffic from the Internet to the web servers in the DMZ
- firewalling between two LAN security zones
- IPS filtering between the two LAN security zones, and between all internal zones and the Internet

Figure 2–9: Transparent DMZ Deployment



Any traffic between ports will be firewallled (that is, policy will be applied). The two LAN security zones are assigned to one IP interface because they are using the same IP subnet. The DMZ and the WAN are also using one IP subnet (a different one to the LANs). Interface 2 is the external interface and has a static IP address. Any Web servers in the DMZ will have static IP addresses and these will be

visible from the Internet. Firewall rules forbid direct traffic between the two LAN zones and the DMZ. The device will protect the Web servers from hacker attacks and so on.

In this example, the segments shown in Table 2–2 have been configured for IPS filtering:

Table 2–2: Segments for IPS Filtering

Segment Name	Security Zone A	Security Zone B	Description
LAN1-WAN	LAN1	WAN	To protect LAN1 from security threats from the Internet
LAN2-WAN	LAN2	WAN	To protect LAN2 from security threats from the Internet
DMZ-WAN	DMZ	WAN	To protect the DMZ from security threats from the Internet
LAN1-LAN2	LAN1	LAN2	To protect LAN1 from security threats from LAN2 and vice-versa (for example if an unknown laptop that is infected with viruses is attached to an internal network)

The IPS segments ensure that all traffic passing from the Internet to the internal network will be filtered for security threats. The LAN1-LAN2 segment is optional, but if configured it will ensure that each LAN zone is protected from any threat that affects the other LAN zone.

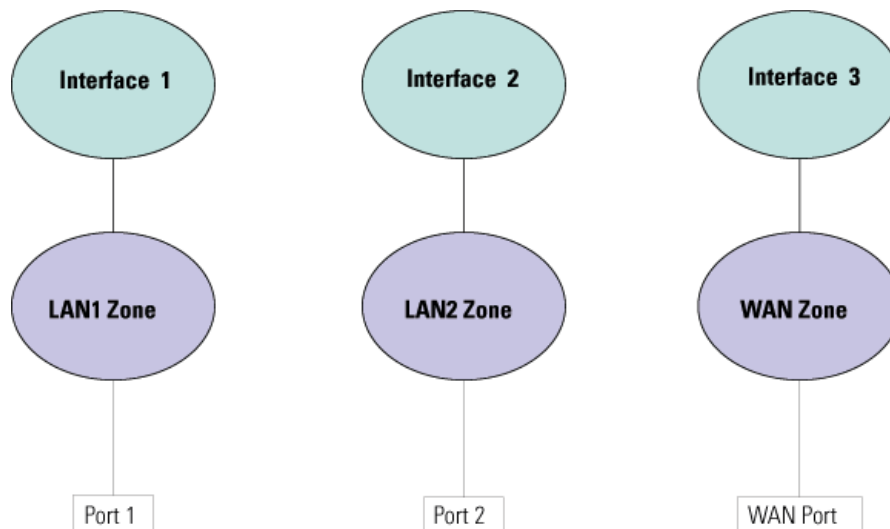
This configuration is typical in a medium-size network that has Web servers with static IP addresses. For example, an X family device is installed in the premises of an online book shop. The Web servers are in the DMZ. Customer records, human resources records, and other secure information are in LAN1. Store employees are in LAN2.

Example 2: Routed Deployment

In this example, shown in [Figure 2-10](#), the X family device is:

- firewalling between the LAN and the Internet
- firewalling between two LAN security zones
- IPS filtering between each LAN security zone and the Internet

Figure 2-10: Routed Deployment



LAN1 and LAN2 are using different IP subnets and therefore each require an IP interface. The external interface again is using a separate IP subnet with a static IP address. The device routes and applies firewall policy between each security zone.

In this example, the following security profile has been configured to protect the LAN1 and LAN2 security zones is shown in Table 2-3 has been configured for IPS filtering:

Table 2-3: Security Profile for IPS Filtering

Security Zone Pair	Description
WAN ==> LAN1	Protect LAN1 from security threats from the Internet
WAN ==>LAN2	To protect LAN2 from security threats from the Internet
LAN1 ==> LAN2 LAN2 <== LAN1	To protect LAN1 from security threats from LAN2 and vice-versa (for example if an unknown laptop that is infected with viruses is attached to an internal network)

The security profile ensures that all traffic passing from the Internet to the internal network will be filtered for security threats. The LAN1==>LAN2 and LAN <== LAN1 security zone pairs is optional,

but if configured it will ensure that each LAN zone is protected from any threat that affects the other LAN zone.

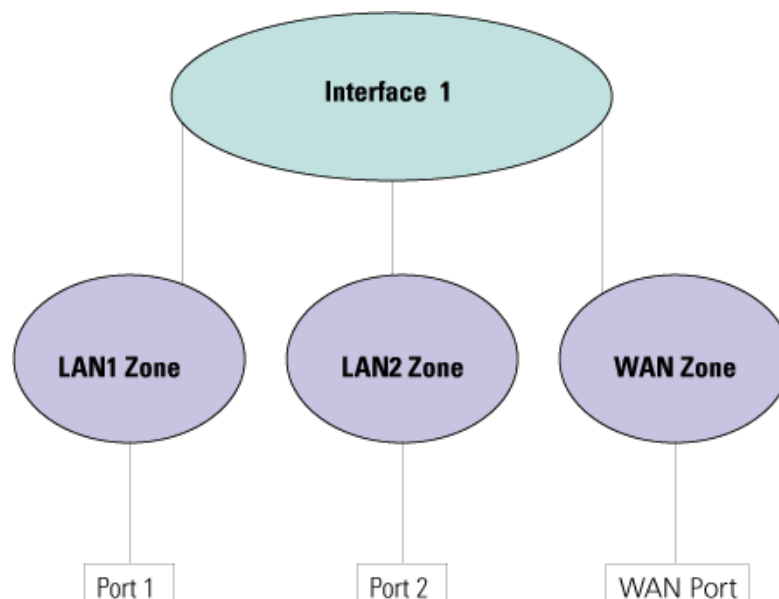
This configuration is typical in a large routed network where policy is required between network sections and can be scaled up to provide more zones. For example, an X family device is installed in a university. Staff and students are on different subnets and can be allocated separate security zones.

Example 3: Fully Transparent Deployment

In this example, shown in [Figure 2–11](#), the X family device is:

- firewalling between the LAN and the Internet
- firewalling between two LAN security zones
- IPS filtering between the two LAN security zones and the Internet

Figure 2–11: Fully Transparent Deployment



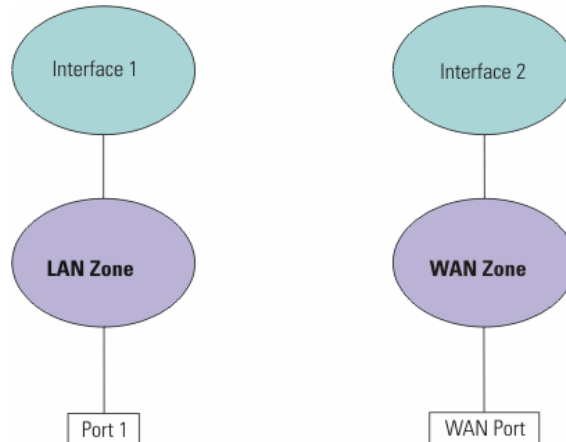
There is only one interface because only one external interface is used, as all security zones are accessible using public IP addresses. In this example, the security profile as shown in Table 2–3 has been configured for IPS filtering.

This configuration is typical in networks where the LAN and WAN are using the same IP subnet — that is, there is a public IP address allocated to every device on the network. For example, an X family device is installed in an archive. The archivists are in LAN1 and are allowed access to the Internet and the archive’s databases. There is a wireless access point in the archive which allows guest users into LAN2 to access the Family History database from their laptops, but not the Internet.

Example 4: NAT Deployment on a Switched LAN

In this example, shown in [Figure 2-12](#), all devices on the LAN network are in the same subnet and there is no requirement for firewalling policy between these devices. The X family device is being used as a firewall device to police traffic passing between the network and the Internet.

Figure 2-12: NAT Deployment on a Switched LAN



Interface 1 is an internal interface and Interface 2 is an external IP interface. LAN port 1 is in the LAN security zone. There is no firewalling between devices on the internal network. Firewall policy can only be applied between security zones. Many-to-one NAT is used between the internal and external interfaces.


This configuration is typical in a small network where there is a dynamically allocated IP address on the WAN. For example, an X family device is installed in a medical center. All doctors have access to all network resources. Client-to-site VPN access is provided to doctors who are not in the medical center. Site-to-site VPNs provide access to other medical centers. All VPNs terminate in the LAN security zone.

See [Chapter 3, "Network"](#) for more information on NAT configuration.

Using VLANs

If your network is divided into completely isolated VLANs, you can use the X family device to allow these network sections to communicate with each other in a secure way, and enable traffic to pass between the VLANs. [Figure 2–13](#) illustrates trunking of multiple VLANs on a single port:

Figure 2–13: VLAN Configuration



Mapping Zones to Physical Ports with VLANs

More than one zone can share the same port by using VLAN tagging and having multiple interfaces. You would need to do this if you have configured more zones than there are available ports on the X family device.

Using the VLANs configured on your network, each security zone has its own VLAN and the device ports can be in more than one security zone. If you create the security zones in this way, the ports on the device are known as tagged ports. Tagging describes the packet — a tagged packet has a VLAN ID appended to it. The VLAN tag is used to identify the network traffic to and from each security zone.

The only restriction on the number of security zones on the same port is the bandwidth of the physical connection to the port.

See [Chapter 3, "Network"](#) for more information on VLAN configuration.

Traffic Shaping and Bandwidth Management

The X family device allows you to prioritize traffic and control services and bandwidth allocation on your network. Traffic shaping is applied using firewall policies. The device allows you to shape traffic based on any of the following criteria:

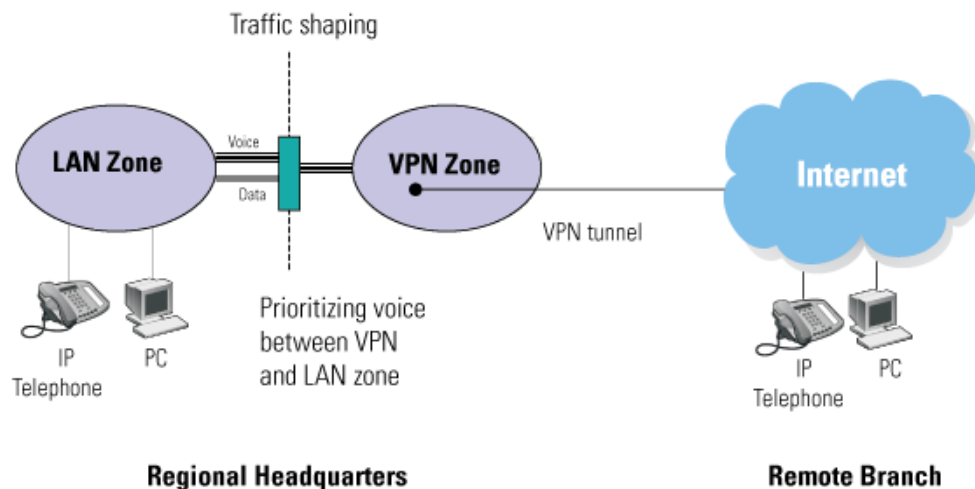
- Type of application — such as voice, email, HTTP, or video conferencing
- Source and destination security zone of the traffic
- IP address of the source or destination
- Time of day

Traffic shaping is applied to traffic passing between security zones, as well as to traffic inside and outside a VPN tunnel, and to both inbound and outbound traffic.

You can also rate limit the total traffic inbound to and outbound from a security zone. To do this, you must configure the inbound and outbound limits per security zone.

[Figure 2–14](#) shows how the device controls traffic passing between zones. Voice traffic, from the IP telephones, is prioritized over other traffic passing across the VPN tunnel.

Figure 2–14: Traffic Shaping Example



You can use the firewall policies to prioritize traffic as follows:

- You can prioritize traffic within a specified zone over the traffic in other zones.
- VPN tunnel traffic (terminating in a separate virtual zone) can also be prioritized above other zone traffic.
- Within the VPN tunnel, traffic from applications such as voice can be prioritized above other traffic, such as Web traffic.

You can also limit the bandwidth rate of both inbound and outbound zone traffic. By limiting the rate of traffic from any specific zone, you can prioritize one zone over another or prevent a zone from overutilizing network resources. To do this, you must configure the inbound and outbound limits per security zone.

Firewall rules used to shape traffic can be applied in one of two ways: per session or per rule. Per-rule shaping lets you prioritize the total bandwidth for a service, such as FTP. Per-session shaping lets you assign a fixed bandwidth to each session, such as a voice session. See [Chapter 5, "Firewall Rules"](#) for more information on traffic-shaping configuration using policies.

Bandwidth Rate Limiting

You can limit the bandwidth to conform to the specifications of the external link. For more information, see ["Managing Bandwidth" on page 65](#).

Multicast and Dynamic Routing

X family support for multicast and dynamic routing enables support of advanced applications, such as voice conferencing.

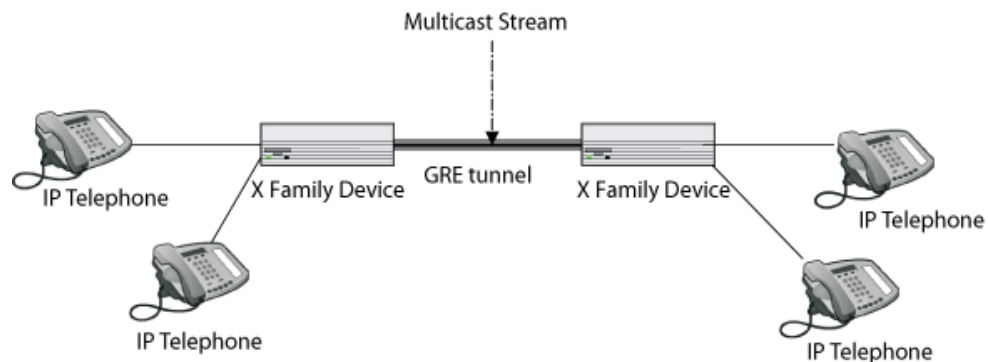
Multicast Routing

You can configure the X family device as an IP multicast router. Multicasting solves the problems associated with simultaneously sending the same information to multiple users by providing a mechanism for routers to forward a one-to-many transmission.

Using IP multicast, any host can join a multicast group, and any host can send a packet with the multicast group's destination address and have it delivered to all members of that group. The sender does not need to be a member of the group. This is particularly useful if you need to support voice or video conferencing.

The device supports IGMP v2 and Protocol Independent Multicast — Dense Mode (PIM-DM). [Figure 2-15](#) shows how multicast packets are encapsulated using GRE:

Figure 2-15: Multicast Stream Encapsulated within GRE



The device combines multiple streams into a single stream, which is transmitted across to the remote sites. At the remote end, the device splits the unicast stream into a multicast stream and forwards the streams to the target recipients.

Since multicast traffic cannot be sent across a VPN tunnel, GRE is used to encapsulate multicast packets. However, GRE alone does not provide security, hence the need for securing the GRE packets with IPsec.

Native IPsec does not support dynamic routing. So multicast traffic is encapsulated within a GRE tunnel as unicast traffic and then encrypted using IPsec.

The benefits of multicast routing are:

- You can deploy multicast applications across a global network.
- Multicast makes more efficient use of available WAN bandwidth.
- Multicast can be used for a variety of applications, including voice, video conferencing, data collaboration, and data distribution.

See [Chapter 3, "Network"](#) for more information on multicast routing configuration.

Dynamic Routing

The X family device supports dynamic routing using two protocols: OSPF and RIP. See [Chapter 3, "Network"](#) for more information on dynamic routing configuration.

OSPF Support

OSPF (Open Shortest Path First, RFC 2328) is an interior gateway protocol used within larger autonomous system networks.

OSPF is more efficient than RIP, an older routing protocol. Unlike RIP, in which the entire routing table is sent to a neighbor host every 30 seconds, the host using OSPF sends only the part that has changed, and then only when a change has taken place. Using OSPF, a host that obtains a change to a routing table or detects a change in the network immediately multicasts the information to all other hosts in the network so that all will have the same routing table information.

OSPF also allows for sophisticated routing schemes. Rather than simply counting the number of hops, OSPF bases its path descriptions on “link states” that take into account additional network information. OSPF calculates cost metrics to a given host router so that some paths can be given preference. OSPF also supports a variable network subnet mask so that a network can be subdivided.

RIP Support

The Routing Information Protocol (RIP) allows the X family device to determine the layer 3 routing required to send information across the network. The device provides support for RIPv1, RIPv2, or both for any interface, for sending and receiving data.

RIP uses multicast to exchange routing information with other routers or devices on the network. Since RIP uses multicast, this can also be encapsulated within the GRE tunnel.

With RIP configured, the device automatically adjusts to physical changes in the network’s layout. The RIP protocol regularly broadcasts routing information to other routers on the network.

High Availability

X family devices support High Availability configuration to provide a failover mechanism to minimize network downtime due to device failure. High availability allows two X family devices with the same configuration and licensing to be configured as a high availability pair. One device is the active device, forwarding packets; the other is a standby device, constantly monitoring the active device. The standby device automatically shifts from standby to active mode if the active device fails.

Once defined, a high-availability configuration can be synchronized between the primary and secondary X family devices.

How High Availability Works

The following sections describe how high availability works in failover and standby mode, how polling works to monitor the state of the active device, and how to synchronize device configurations. For details on configuring high availability, see the *LSM User's Guide* or the online help.

Failover Operation

After a pair of devices has been configured for high availability, the standby device only monitors the active device's HA state and does not route any network packets or monitor the dynamic behavior of the active device. If the standby device detects that the active device has failed, it assumes control of the IP interfaces used to route the packets on the network. When a device becomes active it sends an SNMP trap to any configured NMS trap destinations.

When a standby device takes over, it will not be aware of the final network state of the previous active device before it failed. This affects the device's network operation as follows:

- If dynamic routing is enabled, the new active device will start advertising its initial routing state and will need to relearn the network topology.
- TCP sessions that existed through the previously active device will be unknown to the new device and will be blocked. IPS and firewalling will only be performed on newly created sessions after the HA state transition.
- Site-to-site VPN tunnels that terminated on the previously active device will fail and will need to be re-established by the local device or its peer VPN terminator. To ensure that secondary devices recognize a HA state transition and quickly re-establish tunnels, enable the Dead Peer Detection (DPD) option on IKE proposals.
- Client VPN connections (PPTP, L2TP, and IPSec) will be closed and users will need to re-establish their VPN connection to the new active device using the same VPN IP address as before.
- The new active device will also be unaware of quarantined network equipment. However it will immediately establish quarantine for equipment that continues to transmit prohibited traffic.

When a device changes high availability states, it generates messages in the system log. For a list of these messages, see [“High Availability Log Messages” on page 323](#).

Standby Operation

You can ping the HA management IP addresses from a network device such as a PC to check network connectivity to the standby device. However, the following network tool will not function properly from the console when a device is in standby mode:

- Traffic Capture

As long as the device in standby mode has the appropriate Digital Vaccine (DV) license, the device can automatically retrieve the latest DV updates to ensure the up-to-date protection when the device switches to Active mode. To enable this functionality, the DV Web site must be accessible directly from an external interface through a static route.

Polling

The high availability function provides a polling feature. Polling is used to determine the regular heartbeat mechanism between the standby device and the active device. This function provides the following configuration parameters:

- A **polling interval** determines the period in seconds that the standby device waits before polling the active device. This determines how quickly the standby device will detect that the active device has failed.
- The active device should immediately respond to a poll from the standby device. If it does not, the standby device transmits the heartbeat message after a **retransmission interval** (specified in milliseconds).

Setting these values too low increases the load on the network and can cause the standby device to become active due to lost poll requests or responses.

- A **retry count** determines how many polls the standby device sends before it determines that the active device is not responding. If the active device does not respond on any of the IP interfaces, the standby device becomes the active device.

Synchronizing High Availability Configurations

You can synchronize the high availability configuration of an X family device and its peer. One device, identified by its serial number, is designated the primary HA device; the other device is the secondary HA device. Once synchronization is enabled, any configuration change made to the active device is replicated on the secondary device. 3Com recommends enabling synchronization.

High availability configuration synchronization uses a secure (SSL) channel that is automatically established between the pair of devices. The first super-user name and password pair (which must be the same on both devices) is used to authenticate the link. If the communication link between device pairs is down, configuration changes are queued; if the link is up, replication is immediate.



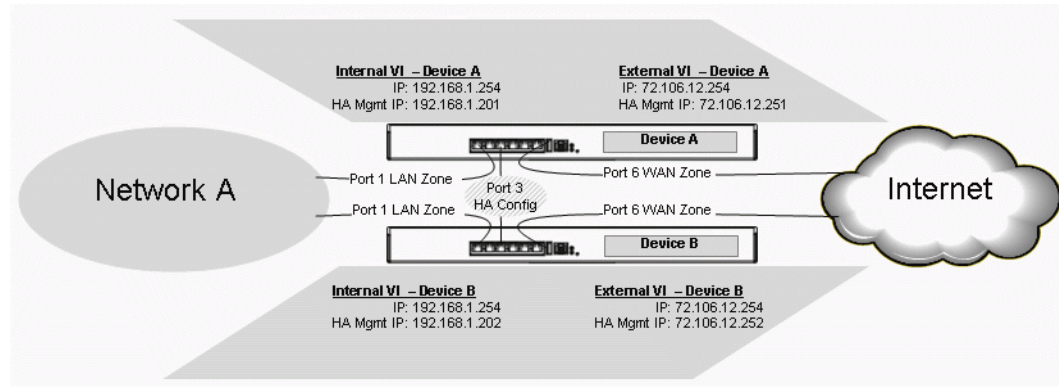
Note: Synchronization is not automatic; it must be manually enabled, and the two devices must be synchronized at least once. Then, once configuration synchronization is enabled, any configuration change to the primary device is automatically replicated to the secondary device.

Certificate Authority (CA), signed, and local certificates are synchronized. However, TOS updates or DV packages are not.

Sample Configuration

Figure 2–16 illustrates a simple HA deployment configured with a single internal IP interface and external IP interface:

Figure 2–16: High Availability Configuration



To configure and manage high availability, see the *LSM User's Guide* or the online help.

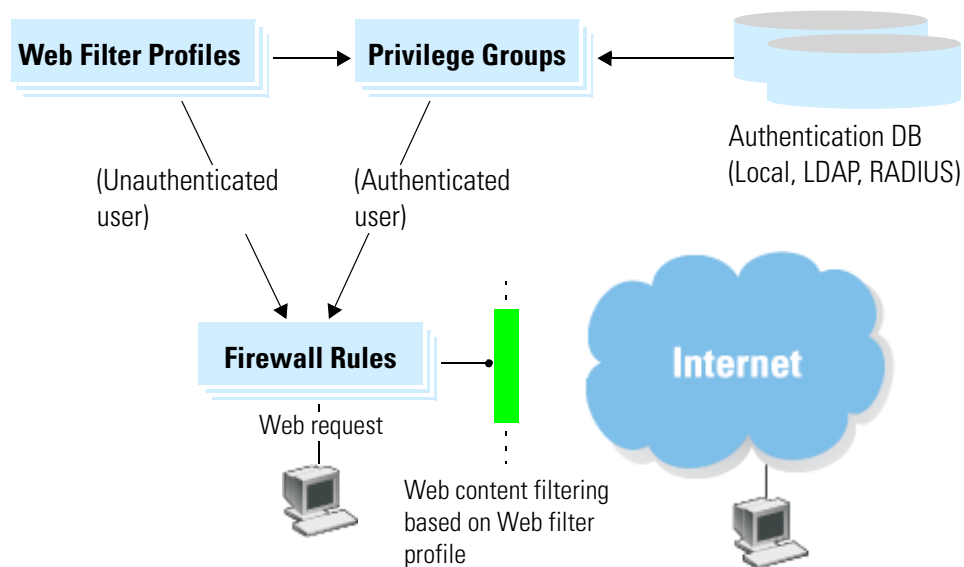
Enforcing a Web Access Policy

The X family device lets you implement a Web access policy. You can completely block Web access, permit unlimited Web access, or control access by user, group, Web site category, and URL.

- You can create firewall rules that block or permit access to the HTTP and HTTPS services.
- You can create Web filter profiles that specify Web access. You can block or permit access to predefined categories of Web sites, and you can further block or permit access to specific Web sites, domains, or pages by URL.
- You can create privilege groups that specify various levels of service access to authenticated users. You can assign a Web filter profile that is applied by a privilege group.
- You can also assign a Web filter profile that is applied by a firewall rule to unauthenticated users. The profile specifies what, if any, Web sites an unauthenticated user can reach.
- The device uses information from the TOS database, an LDAP server, or a RADIUS server to authenticate users and assign a user to a privilege group.

[Figure 2–17](#) illustrates the interaction of firewall rules, Web filter profiles, and privilege groups to control Web access for both authenticated and unauthenticated users:

Figure 2–17: Profile-Based Web Content Filtering



See [Chapter 5, "Firewall Rules"](#) for more information on firewalls and the *LSM User's Guide* or the online help for procedures to create and manage firewalls, privilege groups, and Web filter profiles. See [Appendix A, "Web Content Filter and Anti-Spam Services"](#) for a description of the Web Content Filtering Service.

3 Network

This section describes interfaces, security zones, DHCP functionality, routing, and IP address groups and explains how to enable, disable, and modify their various features. It also describes the network tools provided by the LSM.

Introduction

The information in this chapter relates to the options on the Network page of the LSM.

This chapter includes the following topics:

- [“Security Zones” on page 25](#)
- [“Configuring IP Interfaces” on page 27](#)
- [“Configuring DHCP Server Settings” on page 29](#)
- [“Configuring IP Address Groups” on page 32](#)
- [“Configuring Network Address Translation \(NAT\)” on page 33](#)
- [“Configuring Routing” on page 37](#)
- [“Network Tools” on page 39](#)

Security Zones

This section describes the process to setup security zones for your network. For an overview of security zones and security zone concepts, [“Key Concepts” on page 3](#) describes the concept of security zones and gives examples of how they might be applied in a network.

Dividing the Network into Security Zones

You can divide your network into security zones using physical ports or using VLANs. Each method is described below.

Using Physical Ports to Create Security Zones

If you divide your network physically into security zones, all the network segments that are in an individual security zone are attached to one physical port on the X family device. In a small network, this could mean an individual PC in its own security zone plugged into an individual port on the device. However, this method scales up to a hierarchy of switches and routers all in the same security zone and all connected into one port on the device. If you create the security zones in this way, the ports on the device are known as *untagged* ports.

Using VLANs to Create Security Zones

If you divide your network into security zones using VLANs currently configured on your network, each security zone has its own VLAN. In this way, you can configure ports to be in more than one security zone. If you create the security zones in this way, the ports on the device are known as *tagged* ports.

If your network is divided into completely isolated VLANs, you can use the device to allow these segments to communicate with each other in a secure way, and enable traffic to pass between the VLANs.

Security zones can be used in conjunction with VLANs and VLAN tagging, to let you use the device to provide appropriate policy control between the zones.

Partitioning LANs into different segments using VLANs helps control access and manage bandwidth. This enables you to logically, rather than physically, segment the network, and then apply security policies for each VLAN. When using VLANs for security, managing the switching infrastructure becomes as important as managing the device.

When using VLAN tagging, one port can be associated with multiple security zones. (You would need to use VLANs if you have configured more zones than there are available ports on the device.)

Tagging enables you to have multiple security zones on the same port. The only restriction is the bandwidth of the physical connection to the port.

Default Security Zones

Each X family device is configured with the following predefined, default security zones:

Table 3–1: Default Security Zones

Preconfigured Zones	Ports
LAN	Port 1 (or “LAN”)
WAN	Port 6 (or “WAN”)
VPN	VLAN ID = 4

You can modify the default zones or create your own security zones, with associated security policies and traffic shaping rules, according to the needs of your users and the topology of your network. For details on creating security zones, see the *LSM User’s Guide* or the online help.

Advanced Security Zone Configuration

You can configure a security zone to restrict bandwidth usage and/or add network protection.

- **Bandwidth Management** settings can prevent packet queuing on a WAN device in order to provide lower end-to-end latency on latency-sensitive traffic such as Voice over IP (VoIP). When configuring limits for inbound and outbound traffic, specify a rate slightly less than the actual WAN rate that is available. By testing the latency-sensitive application, you can determine the exact value to use. Bandwidth rate limiting for a security zone is used in conjunction with bandwidth management on firewall rules. You must enable a firewall rule to allow latency-sensitive traffic with priority 0 with the required guaranteed bandwidth. For more information about configuring policies to control bandwidth, see [“Managing Bandwidth” on page 65](#).
- **Network Protection** settings provide additional security for the zone. You can restrict the security zone so that it is only accessible to a subset of devices. You can also restrict outgoing traffic so it cannot be sent from the security zone to any VPN tunnel.

Configuring IP Interfaces

This section describes the three types of IP interfaces available on the X family device and how to configure them. For details, see the following:

- [“External Interface” on page 28](#)
- [“Internal \(LAN\) Interface” on page 28](#)
- [“GRE Interface” on page 28](#)
- [“Configuring IP Interfaces” on page 29](#)

See [Chapter 2, “Key Concepts”](#) for additional information on IP interfaces along with examples of how to apply them in a network.

External Interface

An external interface is typically a connection between the device and the Internet. It requires a public IP address. There are a number of methods by which an external interface can be allocated its IP address, usually dictated by your ISP (Internet Service Provider):

- Static IP Address — your ISP has allocated you a static IP address.
- DHCP — your ISP has told you to use DHCP, or you are connecting to a device that provides IP configuration using DHCP. In this case, the device is a DHCP client.
- PPPoE client — your ISP is using PPPoE to provide your IP configuration.
- PPTP client — your ISP is using PPTP to provide your IP configuration.
- L2TP client — your ISP is using L2TP to provide your IP configuration.
- DNS — your ISP is using a Distributed Name Service server to provide your IP configuration.
- Dynamic DNS — your ISP is using a dynamic DNS server to provide your IP configuration.

Internal (LAN) Interface

An internal interface of the X family device is a connection between the device and your own network. Internal interfaces of the device use static IP addressing. When changing an internal (LAN) interface IP address of the device, choose an address that will be unique in your network and in one of your network's IP subnets.

If your network is routing internally (that is, there is more than one IP subnet directly connected to the device), you must configure one internal interface for each IP subnet.

GRE Interface

Generic Route Encapsulation (GRE) is an interface to a remote site you configure to allow multicasting and dynamic routing between sites. This is necessary if, for example, you need to support voice conference calling (that is, networked IP telephony) between sites. A GRE interface uses a VPN tunnel that you have configured to perform GRE encapsulation.


You will need one GRE interface for each remote location to which you want to support multicast or dynamic routing. The GRE tunnel requires a security zone to function.

Because IPSec tunneling does not support multicast or dynamic IP routing, in order to transmit these packets over a site-to-site VPN, IPSec transport mode is used in combination with the GRE protocol. This allows IPSec to treat multicast and dynamic IP routing traffic as unicast traffic.

The data or payload that is going to pass through the tunnel is given a header and then placed inside a GRE packet; see [Figure 3-1](#). The GRE packet carries the data between the two tunnel endpoints. After

the GRE packet has arrived at the final destination (the endpoint of the tunnel), it is discarded and the encapsulated packet is then transmitted to its final destination.

Figure 3–1: GRE Packet



Configuring IP Interfaces

When you create an interface, there are a number of configurable settings:

- **IP allocation method** — the method by which this interface is allocated an IP configuration. For internal interfaces, this includes whether you require many-to-one NAT. For more information about many-to-one NAT, see [“Configuring Network Address Translation \(NAT\)” on page 33](#).
- **Dynamic routing method** — whether you require this interface to support OSPF or RIP (and related settings). For more information, see [“Dynamic Routing” on page 37](#).
- **Multicast** — whether you require this interface to support IGMP or PIM-DM multicast routing. For more information, see [“Multicast Routing” on page 38](#).
- **Zones** — the security zones that are associated with this interface. For more information, see [“Configuring IP Interfaces” on page 27](#).



Note Before you create a new interface, ensure that you have a security zone to allocate to the interface. A security zone must be present that is not currently allocated to any other interface.

For details on creating, configuring, and editing IP interfaces, see the *LSM User’s Guide* or the online help.

Configuring DHCP Server Settings

A DHCP server allows computers on your network to obtain an IP address and subnet mask automatically. DHCP assigns a temporary IP address and subnet mask to any PC that requires IP configuration information and reallocates that address when the PC disconnects from the network.

You can configure the X family device to act as a DHCP server for devices on its LAN-side interfaces that require IP address configuration.

This section describes the DHCP settings of the X family device. All DHCP settings are optional and are configured based on your network requirements.

The following DHCP functionality is provided by the device (and each feature is described below):

- DHCP server
- DHCP relay
- DHCP relay over VPN
- DHCP client
- Static mapping

Using the DHCP Server

Table 3–2 shows examples of policies to allow DHCP clients in the LAN security zone

Table 3–2: Example Policies to allow DHCP Clients in the LAN Security Zone

Action	Service	Source zone	Destination zone	Destination IP
Permit	dhcp-server	LAN	this-device	ANY
Permit	dhcp-client	this-device	LAN	ANY



Note When creating firewall policies to allow the DHCP or BOOTP functionality that you require, note that **All services** includes DHCP and BOOTP, so you might not need to specify these services individually.

The DHCP server also supports BOOTP requests for older PC clients.

Using NBX or VCX telephones

If you are using the X family device as your network’s DHCP server and you want to allow NBX or VCX telephones to retrieve the network call processor (NCP) IP address, provide the device with an NCP IP address (otherwise you will need to manually configure each phone).

See [Chapter 12, “Deployment Scenarios”](#) for more information about using the device in a network that contains NBX or VCX telephones.


Using DHCP Relay

DHCP relay allows DHCP to operate between a DHCP client on one subnet and a DHCP server on another. To use DHCP relay, you configure the X family device to act as a DHCP relay agent. The device will relay DHCP packets to the destination DHCP server and back to the client across network segments. This enables DHCP clients on different networks to use the same DHCP server.

You can enable and configure DHCP relay from the DHCP Server page in the LSM (**Network > DHCP Server**).


Using DHCP Relay over VPN

DHCP relay over VPN allows DHCP to operate between a DHCP client at a remote site and a DHCP server at a different site via a VPN tunnel. If you want to allow a DHCP server at one site to provide IP configuration to clients attached to a remote LAN via VPN, use DHCP relay over VPN. For example, a DHCP server might reside at the headquarters of a company. You can configure the devices in the remote branches to relay DHCP requests to the headquarter's device, which in turn relays the requests to the DHCP server.

 **Note** You can only use DHCP relay over VPN when the VPN between the two devices is set up to use Internet Key Exchange (IKE).

[Figure 3–2](#) shows how you can configure the device to act as a central or remote relay agent.

Figure 3–2: DHCP Relay: Device Configuration for Central and Remote Agent



- A **central relay agent** is connected to the network that contains the DHCP server. It receives requests from a remote agent and forwards them to the DHCP server on its LAN. You can configure this option to work over VPN so that the device allows a DHCP server at one site to provide IP configuration to clients attached to a remote LAN. In this configuration, the device acts as a DHCP Relay agent and supports DHCP over VPN tunnels using IKE.
- A **remote relay agent** is connected to a client network that requests a DHCP lease. It listens for DHCP requests from its LAN. When a client request is received, the agent inserts the Interface IP of the requestor into the DHCP request before it is relayed to the central DHCP server. This address, which is not contained in a DHCP address range, determines the scope of addresses used by the central DHCP server to allocate the address to the remote client.

For configuration instructions, see the *LSM User's Guide* or the online help.

DHCP Client

The X family device can be a DHCP client for the purpose of getting IP configuration for an external interface from a DHCP server via an Internet router if this is the method you require. This is described in [“External Interface” on page 28](#).

Assigning a Static IP Address

Static mapping allows you to configure devices with a static IP address. You can assign a static IP address to any device on your network using the X family DHCP server. This is useful for devices such as printers, where the IP address needs to be constant.

You can also assign static addresses when using DHCP relay over VPN if your remote site is using devices with fixed IP addresses that do not support DHCP.

To configure static mapping for DHCP, use the Static Reservations page in the LSM (**Network > DHCP Server**). Then, select the Static Reservations tab.

Configuring IP Address Groups

IP address groups allow you to specify and name a collection or range of IP addresses that will share the same configuration settings. Use IP address groups to:

- Simplify the process of defining security zones, firewall rules, and IPS filter IP address limits and exceptions (especially if you want to apply firewall rules to a subset of the devices within a security zone).
- Provide an IP address pool for devices that may be added to a security zone in the future.
- Specify a group of IP address ranges or subnets for the DHCP server to support DHCP clients in different interfaces.

You can add multiple addresses, ranges and hosts into an IP address group. To configure IP addresses, use the IP Address Groups page in the LSM (**Network > Configuration > IP Address Groups**).

Configuring Network Address Translation (NAT)

This section describes Network Address Translation (NAT) on the X family device and how to configure it. All NAT settings are optional and you must configure them depending on the requirements of your network.

See [Chapter 2, “Key Concepts”](#) for descriptions of network scenarios that might use NAT.

The X family device is able to perform NAT in two modes, as shown in [Figure 3–3](#) and [Figure 3–4](#).

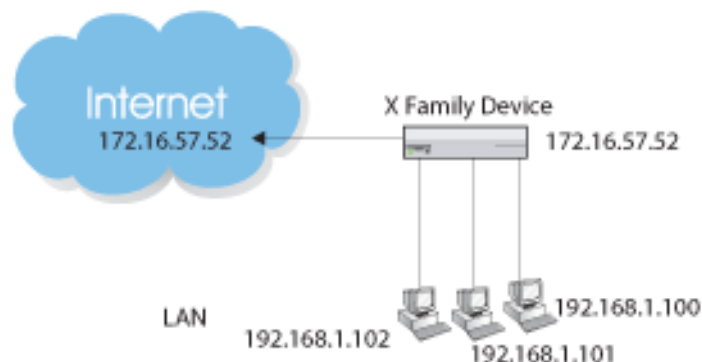
- **Many-to-one NAT** — The device only uses one Internet IP address for NAT. In many-to-one NAT, all the addresses on your LAN are mapped to the Internet address of the device. This means that:
 - the internal IP addresses on your network are not exposed externally.
 - you do not have to configure every PC on your network with a public IP address.
 - you only need one registered IP address.
- **One-to-one NAT** — The device uses a pool of Internet IP addresses for NAT. Each Internet IP address is associated with one LAN IP address. Effectively, each of these LAN IP addresses has its own public IP address. You can apply one-to-one NAT only where it is required. Typically, you only use one-to-one NAT on a network where you are also using many-to-one NAT. By using one-to-one NAT:
 - You can allow servers on your LAN, which are protected by the device firewall, to be accessed from the Internet, while the internal IP addresses of these hosts on your network are not exposed to the Internet.
 - Individual PCs (for example, Web servers) can have a fixed public IP address associated with them and therefore be accessible from the Internet with the correct device policy.



Note You can also configure the X family device to perform NAT over a VPN tunnel. For additional information, see [“Advanced VPN Configuration” on page 54](#).


Each figure below shows an X family device where the external interface has IP address 172.16.57.52. In [Figure 3–3](#), many-to-one NAT is configured. This means that the PCs or servers on the LAN have their private IP addresses mapped to the IP address of the external interface.

Figure 3–3: Many-to-One NAT



In [Figure 3-4](#), one-to-one NAT is configured. This means that the PCs on the LAN each have their IP address mapped to a public routable IP address associated with the external interface.

Figure 3-4: One-to-One NAT



Setting Up Many-to-One NAT

Many-to-one NAT can only operate between an internal and external interface. You configure it when you configure an internal interface. Many-to-one NAT works with any IP allocation mode and will map all the addresses associated with the internal interface (for which you have configured many-to-one NAT) to the address of the external interface.

Setting Up One-to-One NAT

When the X family device is set up to provide NAT, then internal servers (behind the firewall) cannot be accessed directly by external devices, because the internal network is private, and not exposed to external devices. If you are using NAT, the Virtual Server option (configured using the Virtual Servers page) enables you to configure firewall rules that allow external devices to access internal servers.

You can use the virtual server option to define a private LAN server IP address for each service passing through the firewall. Any external request for a service, directed at the device's WAN IP address, is forwarded to the virtual server. (When the default service is configured as a private LAN server, all incoming sessions, not otherwise intercepted as other private LAN servers for other services, are directed to the server's IP address.)

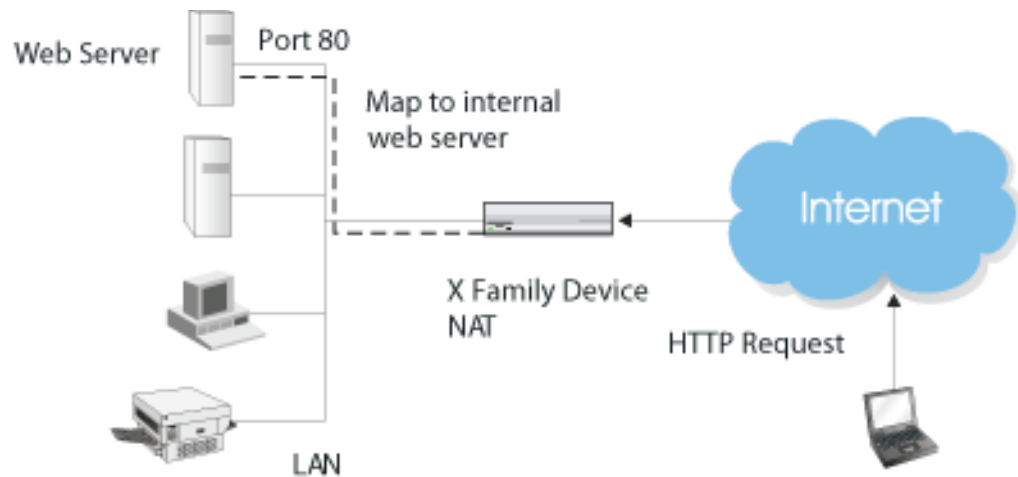
Outgoing sessions from the private server or device to the public network will use the public IP address configured for the virtual server. This allows one private IP address to be mapped to one public IP address. If you select **ALL** for the service, this provides one-to-one NAT for devices on the private LAN.



Note One-to-one NAT traffic is subject to firewall rules. You must set up a firewall rule to allow the traffic for the desired services through the firewall. To allow incoming traffic, use the IP address of the LAN device as the destination address of the firewall rule.

[Figure 3–5](#) illustrates how the virtual server works.

Figure 3–5: Virtual Server Configuration



To be able to use one-to-one NAT, you must have a static Internet IP address for every computer on your network that requires one, and one for the device itself.

Defining the Public IP Address

Two options are available for mapping to the X family device's public IP address:

- **Use External Interface IP address.** Use the preconfigured or automatically retrieved external IP address of the device.
- **Manually enter the external IP address.** You can configure an IP address that is part of the device's WAN IP subnet, but different from the one that the device is currently using.



Note The address that you define must be one that has been allocated to you by your Internet Service Provider.

Port Address Translation

Port Address Translation (PAT) maps a service to a different local port, so that you can run multiple instances of the same service on a single server. For example, an HTTP/Web server can be run on ports 80, 90 and 100 — effectively providing three separate Web servers, accessible by three separate public IP addresses.

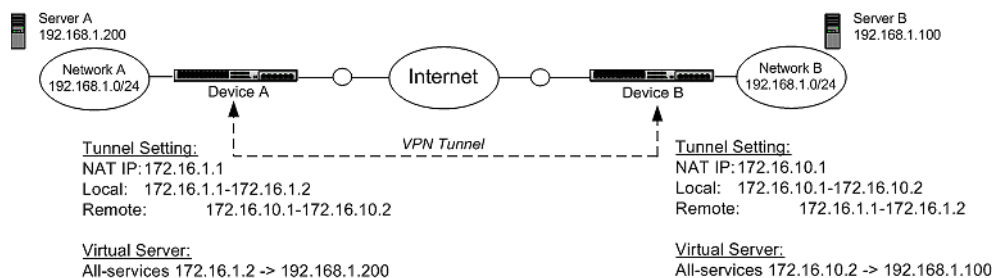
See [Chapter 5, “Firewall Rules”](#) for more information about firewall rules.

Setting up NAT Within a VPN Tunnel

One-to-one NAT for VPN tunnel allows you to perform NAT on traffic sent over the VPN tunnel. This allows multiple remote VPN sites to use the same IP subnet.

VPN NAT is a useful feature when the local and remote networks of the VPN tunnel overlap; VPN NAT can be used to translate the networks onto another neutral address space. In this scenario, both sides of the VPN tunnel must perform NAT on traffic sent over the tunnel.

The diagram below shows an example configuration where the local networks of the two devices are identical. Network A is able to access Server B via the virtual server and Network B is able to access Server A via the virtual server.



The following procedure describes the steps required to configure VPN NAT for this scenario.

- STEP 1** On Device A, configure the IPSec security association for the VPN tunnel for Device B:
- STEP A** Enable IPSec tunnel connections.
 - STEP B** Configure the Local ID (Local Networks IP address range).
Using this example, the Local ID (Local Networks IP address range) would be 172.16.1.1–172.16.1.2.
 - STEP C** Configure the Remote ID.
Using the example, the Remote ID (Remote Networks IP address range) would be 172.16.10.1–172.16.10.2.
 - STEP D** Enable NAT of local network addresses.
 - STEP E** Specify the NAT IP address.
Using the example, the NAT IP address would be 172.16.1.2.
 - STEP F** Make sure that the termination zone for the Security Association is set to a virtual zone that contains no physical ports.
- STEP 2** Configure the Virtual Server to be used by the tunnel:
- STEP A** From the LSM menu, select **Firewall > Virtual Servers**. Click **Create**.
 - STEP B** Select **All** from the **Services** drop-down list.
 - STEP C** Specify the **Local IP Address** which traffic will be redirected to.
Based on the example, specify 192.168.1.200.
 - STEP D** Specify the **Public IP address** which users or devices will use to access services.
(This address must be part of the Local ID IP subnet specified for the VPN tunnel).
Based on the example, specify 172.16.1.2.
- STEP 3** Configure firewall rules to allow traffic over the VPN tunnel.

See [Chapter 5, “Firewall Rules”](#) for more information about firewall rules.

STEP 4 Repeat Steps 1 through 3 for Device B.

Configuring Routing

This section describes how to configure the X family device for static, dynamic, and multicast routing.



Note The list of all routes (static and dynamic) are listed in the Routing page (**Network > Routing**) of the LSM.

Default Route

The default gateway is the route to which the device forwards a packet with a destination address not recognized by the device. Use the Default Gateway page in the LSM (**Network > Configuration > Default Gateway**) to configure this setting. For details, see the *LSM User's Guide* or the online help.



Note If you are using PPPoE, L2TP, PPTP or DHCP, then the default route will be automatically configured by your ISP and you cannot configure it yourself.

Static Routing

A static route is a manually added route, which you configure, to allow the X family device to transmit packets to the appropriate router.

The device supports the use of static routes to forward traffic:

- between the device and any external router.
- between the device and any GRE interface.



Note Static routes configured on the device are not used to route traffic to subnets at the other end of an IPSec VPN tunnel. The destination network's configuration in the Security Association associated with a VPN tunnel is used for this.

Dynamic Routing

The X family device supports both OSPF and RIP for dynamic routing.

OSPF for IP Interfaces

OSPF (Open Shortest Path First, RFC 2328) is an interior gateway protocol used within larger autonomous system networks.

OSPF is more efficient than RIP, an older routing protocol. Unlike RIP, in which the entire routing table is sent to a neighbor host every 30 seconds, the host using OSPF sends only the part that has changed, and then only when a change has taken place. Using OSPF, a host that obtains a change to a routing table or detects a change in the network immediately multicasts the information to all other hosts in the network so that all will have the same routing table information.

OSPF also allows more sophisticated routing schemes. Rather than simply counting the number of hops, OSPF bases its path descriptions on “link states” that take into account additional network information. OSPF also lets you assign cost metrics to a given host router so that some paths are given preference. OSPF supports a variable network subnet mask so that a network can be subdivided.

RIP for IP Interfaces

The Routing Information Protocol (RIP) is an interior gateway and dynamic routing protocol. RIP exchanges routing information between routers and adopts hop count to measure the distance from the destination and calculate the fewest number of hops between the source and the destination.

With RIP configured, the device automatically adjusts to physical changes in the networks layout. The RIP protocol regularly broadcasts routing information to other routers on the network.

RIP has two versions: RIP-1 and RIP-2. RIP-2 supports simple text authentication and MD5 authentication, as well as the variable-length subnet masks.

To improve performance and prevent route loops, the device supports split horizon and poison reverse:

- Split horizon — reduces convergence time by not allowing routers to advertise networks in the direction from which those networks were learned. The announcements only include networks in the opposite direction. If you enable this on an interface, routes that the device learns via that interface will not be advertised on that interface.
- Poison reverse — routes learned from a neighbor are advertised back to it with metric 16 (unreachable). This has a similar effect as split horizon, and is also called split horizon with poison reverse. In a single-path network this has no advantage over split horizon. However, in multi-path networks this greatly reduces loops. If you enable poison reverse on an interface, routes that the device learns from a neighboring router will be advertised back to that router as unreachable.

By default, split horizon and poison reverse are enabled on an interface. Typically, you would only disable these features for compatibility with other products.

The device provides support for RIPv1, RIPv2, or both for any interface, for sending and receiving data. For details on setting up dynamic routing, see the *LSM User's Guide* or the online help.

Set Up Dynamic Routing over a VPN

To set up the device so that it allows the use of OSPF or RIP over VPN, you need to create a GRE tunnel. To do this, see the *LSM User's Guide* or the online help.

Multicast Routing

[Chapter 2, “Key Concepts”](#) describes multicast routing. The following procedures describe how to configure the device to support multicast routing.

Setting Up Single-Site Multicast Support

To set up multicast groups that will only operate within one site (that is, there is no VPN connection in the multicast group), ensure IGMP is globally enabled and is enabled on every interface that will require it. By default, IGMP is disabled on all interfaces. This is the protocol that the device uses to manage multicast groups within a Layer 2 network. For details, see the *LSM User's Guide* or the online help.

Setting Up Site-to-Site Multicast Support

Site-to-Site multicast routing allows multicast groups to operate between sites (that is, there is a VPN within the multicast group). To enable this feature, you must configure IGMP and PIM-DM. For details, see the *LSM User's Guide* or the online help.

Network Tools

The LSM provides the following network tools:

- **DNS Lookup** — a network tool that displays the IP Address for a given DNS name.
- **Find Network Path**— a network tool that displays the physical interface/security zone (and router IP address if appropriate) that the X family device would use to reach a given location.
- **Traffic Capture** — a network tool that allows you to capture network packets into a file. This is useful for analyzing the type of traffic flowing through the device.
- **Ping** — a network tool that allows you to send out Ping requests to test whether devices on an IP network are accessible and functioning correctly. This feature is useful to diagnose connectivity problems such as a failed network device between the device and the web server being accessed, or to help diagnose DNS setup problems.
- **Traceroute** — a network tool that allows you to display the network hops from the device to another device on an IP network. This is a useful tool for network troubleshooting.

For details on these tools, see the *LSM User's Guide* or the online help.

4 Virtual Private Networks

This chapter provides an overview of Virtual Private Networks and describes how they are implemented using X family devices. It covers the following topics:

- [“Introduction to Virtual Private Networks” on page 41](#)
- [“X Family VPN Security Features” on page 42](#)
- [“Site-To-Site VPNs” on page 44](#)
- [“Client-to-Site VPNs” on page 50](#)
- [“Advanced VPN Configuration” on page 54](#)

Introduction to Virtual Private Networks

This section provides a general overview of the purpose and benefits of using a VPN. If you are familiar with these concepts, you can go directly to [“VPN Connection Security Features” on page 43](#).

What is a VPN?

A Virtual Private Network (VPN) is a means of establishing a secure connection between two points across a public network, for example, the Internet. A VPN is a private connection between two points, which in reality uses tunneling across a public connection from the Initiation Point to the Termination Point.

Initiation

Initiation occurs when the remote user or device requests access to the company LAN. Tunnel initiation is usually accomplished using VPN client software on a PC, or through VPN support in an access router or firewall, such as an X family device.

Termination

Termination is the point in the network at which the identity of the remote party is validated, the VPN tunnel is created, and the remote party enters the network. VPN termination is typically supported in routers, secure gateways, Internet firewalls, or in software residing on a network server.

Benefits Of VPNs

VPNs provide an easy, affordable, and secure means for businesses to conduct operations and provide network connectivity to all offices and partners across a public infrastructure. Users can exchange information from any location that has access to the Internet.

Data that is intended for delivery to a remotely connected site is automatically encrypted. The data is delivered via the Internet and decrypted at the intended destination.

Types of VPN Connections

In general, for the purpose of configuring the X family device, VPNs can be broadly grouped into two main types:

- **Site-to-site** — a VPN tunnel established between two firewall devices, typically used for office-to-office connectivity. For more information about site-to-site connections, see [“Site-To-Site VPNs” on page 44](#).
- **Client-to-site** — a VPN tunnel established between the device and a VPN client application, typically used to connect off-site users to the office network. For more information about client-to-site connections, see [“Client-to-Site VPNs” on page 50](#).

You can also use advanced VPN configuration options based on your network deployment:

- **NAT within a VPN tunnel** — this configuration is useful in environments where the local and remote networks of the VPN tunnel overlap. VPN NAT can be used to translate the networks onto another neutral address space.
- **VPN supernets** — Use this configuration for hub and spoke network topologies.

X Family VPN Security Features

Besides the usual security features of VPNs described below, the X family device also offers firewall-based VPN security, where the firewall features of the device can be used to improve the level of security afforded to the connection.

Tunnel Security Zones

Firewall security features are implemented through the use of firewall rules, which can be configured to permit or block access to specific services, based on the source and/or destination zone of the device, the type of service being requested, or the identity of the user.

By defining virtual tunnel zones for your VPNs (for example, security zones that have no physical ports associated with them), you can firewall traffic that flows between your VPN connections and the other

security zones in your network. By defining virtual tunnel zones, you will be able to configure IPS segments to provide IPS filtering between VPN traffic and the rest of your network.

For more information on security zones, refer to [“Key Concepts” on page 3](#). For information on firewall rules, see [“Firewall Rules” on page 59](#).

VPN Connection Security Features

The X family device uses three main VPN features to ensure the security of the network connection: tunneling, authentication, and encryption. These features work together to protect your resources and guarantee that your VPN connections are secure — even across the Internet.

Tunneling

Tunneling is the term used to describe the link created between the two endpoints in a VPN connection; for instance, between an employee's home-office computer and the company network.

The connection is termed a “tunnel” because the information exchanged across the link is “encapsulated.” This means that it is wrapped by protocols and data encryption methods that protect the data as it travels over the Internet, making unauthorized tapping into the information impossible. The three main secure VPN tunneling protocols used by the X family device are:

- IPSec
- L2TP over IPSec (recommended) or L2TP
- PPTP

These protocols are discussed in more depth later on in this chapter.

Authentication

Establishing the identity of a remote user is an essential element of VPN. This ensures that the person accessing the information is actually the person who has permission to access it; then, based on who the person is, limit where the remote user can go (LAN, intranet, or Internet); and finally what files and services the remote user can access once authenticated. There are two main types of authentication used in VPN:

- **User Authentication** — Using username/password verification methods, user authentication ensures that only authorized users have access to company information. X family VPNs can use a PKI (public-key infrastructure), with support for X.509 certificates, or an external RADIUS server. Access privileges are used to control different access levels for each employee or customer. Users are given permission to access only those areas of your Web site or network that you want them to access.
- **Packet Authentication** — Although interception and viewing of data on a shared network is the primary security concern, data integrity is also an issue. Packet authentication checks the electronic signature appended to the packet by the terminating device, to ensure that the packet has not been tampered with.

Encryption

Encryption is applied to the tunneled connection to scramble data, thus making data legible only to recipients with the correct key. Using cryptographic algorithms, information is scrambled (encrypted) by the initiator and then unscrambled (decrypted) when it reaches the recipient.

Recipients of encrypted data must have access privileges and hold specific keys in order to read the data.

Site-To-Site VPNs

Setting up a site-to-site VPN will let you connect your remote sites. For site-to-site VPNs, the X family device supports the IP security (IPSec) set of protocols for securing and authenticating IP traffic for a VPN.

A site-to-site connection refers to the establishment of a VPN tunnel between remote offices. This could be a connection between your headquarters and regional branch offices/distribution centers, or between branch offices and small home offices. Devices situated at each site are used to establish the VPN tunnel and pass traffic between the two sites. Once the VPN tunnel is created, information can be exchanged securely between both LANs. Users on each LAN communicate across the tunnel, making the two LANs appear as one.

For site-to-site connections, the device uses the IPSec protocol.

Figure 4-1: Site-to-Site Connection



IPSec Modes

Depending on how it is being used, IPSec can run in two modes:

- Tunnel mode
- Transport mode

IPSec Tunnel Mode

IPSec tunnel mode encapsulates and secures complete IP packets. It is typically used to provide site-to-site VPN connections where there are IPSec gateways, such as an X family device, at either end of the connection. The device tunnels the packets from one host on a private LAN across the Internet to the device at the other end, which terminates the VPN tunnel and forwards the unencapsulated packets to a host on the remote private LAN.

Since complete packets are encapsulated, the IP addresses of the LAN devices are hidden within the IPSec tunnel mode packets. Therefore, there is no need for these devices to be assigned with public (Internet routable) IP addresses.

IPSec tunnels always terminate inside a security zone. For information on terminating IPSec tunnels in security zones see [“IPSec Tunnel Setup” on page 49](#).

IPSec Transport Mode

IPSec transport mode encapsulates and secures only the data or payload of IP packets. As only the data portion of the packet is secured (instead of the full IP packet), the two X family devices (or another gateway device) require publicly routable IP addresses to communicate with each other over the Internet.

IPSec tunneling does not support multicast or dynamic IP routing. In order to transmit these packets over a site-to-site VPN, IPSec transport mode is used in combination with Generic Route Encapsulation (GRE) protocol. This allows IPSec to treat multicast and dynamic IP routing traffic as unicast traffic.

For more information on setting up your network to support site-to-site multicast, dynamic routing, and GRE interfaces, see [“Network” on page 25](#).

IPSec transport mode can also be used for a client-to-site VPN setup with L2TP. For more information, see [“Client-to-Site VPNs” on page 50](#).

Dynamic VPN Peer and Alternate VPN Peer

You can specify VPN peers either by IP address or by DNS name. Specifying DNS names is more convenient and allows more flexibility. In addition, the device supports use of an alternate VPN peer if connection to the primary VPN peer is lost.

Summary of Site-To-Site VPN Methods

The two methods of setting up a site-to-site VPN connection are summarized in [Table 4–1](#):

Table 4–1: VPN Site-to-Site Setup Methods

Method	Description
IPSec	The standard X family method of setting up a network-to-network VPN connection.
GRE over IPSec	Used to support applications such as dynamic routing and multicast, in which the data is encapsulated within a GRE packet. IPSec then treats the packets as unicast traffic.

Security Association

The security association (SA) defines the parameters that are used to set up a secure VPN connection. An SA is a group of security settings associated with a specific site-to-site VPN connection or set of client-to-site VPN connections.

A security association includes the following security features:

- Encryption
- Authentication of data integrity
- Sender authentication and non-repudiation (if using certificates)

The security features of your SA will depend on your network requirements. You can authenticate the packet source, to ensure content integrity, and you can apply encryption, to preserve the privacy of information.

The X family device is preconfigured with a default SA that can support multiple concurrent client-to-site VPN connections. The default SA has to be enabled before it can be used. If you require client-to-site VPNs that have different individual parameters, then configure them as if they were site-to-site connections, each with its own SA.



Note: Only the default SA can support multiple client-to-site connections.

IPSec Security Mechanisms

IPSec protects the header information and data within a packet through three mechanisms:

- Authentication Header (AH) — provides security by adding authentication information to an IP packet, protecting the entire TCP/IP packet.
- Encapsulation Security Payload (ESP) protocol — provides full protection of the data contents within a packet. It uses both encryption and authentication methods. For more information, see [“Encryption and Data Integrity” on page 48](#).
- Security Parameter Index (SPI) — identifies the cryptographic keys and algorithms to be used to establish a VPN tunnel. For more information, see [“Security Parameter Index \(SPI\)” on page 47](#).

Keys and Keying Modes

Keys are alphanumeric strings that can be used to encode data for encryption and authentication. Keys used in VPN communications can vary in length. The longer the key, the more difficult it is to break the encryption.

The X family device uses symmetric cryptography to encrypt and decrypt the data. As a result, the key on both ends of the VPN tunnel must match exactly.

The device supports two methods of creating keys for IPSec:

- Manual Key (low level of security)
- Internet key Exchange (IKE) with a preshared key or certificate (higher level of security)



Note: The key mode is selected via the security association for the VPN.

Manual Key

Using a manual key requires configuring security parameters at both ends of the tunnel. This may be more suitable for small, simple networks that are easy to configure and manage; use a manual key only if you require it for interoperability. Because manual keys are not regenerated on a regular basis, using a manual key provides a lower level of security than a key created through IKE. If you are using manual keying, you must configure the inbound and outbound Security Parameter Indexes (SPI) in the SA for the VPN.



CAUTION: Take care when delivering/exchanging this manual key to ensure that a third party cannot compromise the security of a VPN tunnel.

Security Parameter Index (SPI)

The SPI is used to establish a VPN tunnel. It is automatically generated if IKE is used, but must be manually created by the administrator if manual keying is used. The same SPI values must be configured on the X family devices/firewalls on both sides of the tunnel. The SPI is transmitted from the remote device to the local device. The local device uses the network, encryption, and key values that the administrator associated with the SPI to establish the tunnel. The SPI must be unique. For more information on creating an SPI, see the online help.



Note: The range from '0' to 'ff' inclusive is reserved by the Internet Engineering Task Force (IETF) and is not allowed for use as an SPI. This range is not accepted by the device when entered as an SPI.

Internet Key Exchange

Keys can be distributed between the VPN initiator and terminator using a key distribution protocol such as Internet Key Exchange (IKE). IKE is used to automatically generate the keys, the SPI, and security association used for encryption and authentication, and to negotiate the connection. IKE uses UDP port number 500 and precedes the actual IPSec data flow.

This method can be used if you need to create and manage multiple tunnels, as you do not need to configure each element manually.

The X family device supports the following IKE modes:

- *IKE with Pre-shared Key.* This mode uses a secret pre-shared key, and is the default keying mode. It offers more security than a manual key. A pre-shared key is a predefined value that the two endpoints

of a VPN tunnel use to set up an IKE SA. For more information on selecting a value for a pre-shared key, see the online help.



CAUTION: Take care when delivering/exchanging this shared key to ensure that a third party cannot compromise the security of a VPN tunnel.

- *IKE with X.509 Certificates.* This mode allows you to specify the local certificate you want to use to prove the identity of the remote device requesting to establish the VPN. This mode offers the highest level of security, but you must use a Certificate Authority (CA) before this mode can be used. You can specify the particular CA certificate used for authenticating incoming VPN requests. If you do not specify a CA certificate, the device will by default use any of the valid certificates currently configured on the device. See [Chapter 10, “Certificates”](#) for more information on using certificates.



Note: If you are using certificates to validate the VPN, then you do not need a shared key. If you have access to a certificate authority server, you should use the certificate method. 3Com recommends that if you are planning to use certificates, you should first set up the VPN using shared secrets to test the VPN and then implement certificates.

Note: Only valid, trusted local certificates are presented for use in the IPsec IKE configuration. If a local certificate in use with IKE is deleted, then this will disable the associated SA.

Encryption and Data Integrity

IPsec uses the Encapsulation Security Payload (ESP) standard for both encryption and authentication of data integrity. ESP uses encryption methods such as DES, 3DES, and AES to secure the traffic, and message digest methods such as SHA and MD5 to authenticate the integrity of the traffic.

Encryption

- Data Encryption Standard (DES) — The original Data Encryption Standard as defined by a U.S. government body. The X family device implementation of DES uses a 56-bit key.
- Strong Encryption (Triple DES or 3DES) — Strong Encryption, or Triple DES (3DES) is a variation on DES that uses a 168-bit key. 3DES is much more secure than DES. However, it uses more processing power, resulting in increased latency and decreased throughput.
- Advanced Encryption Standard (AES) — Advanced Encryption Standard (AES) is a standard that offers less latency and improved security over 3DES. AES uses 128, 192 and 256-bit keys.



Note: 3Com recommends that you use 3DES for encryption.

By default all new X family devices are supplied with 56-bit DES encryption only. To enable strong encryption functionality (3DES, 128-AES, 192-AES, 256-AES), you need to install the correct Strong Encryption Service Pack for your device, available from the TMC Web site.

Data Integrity

- Message Digest 5 (MD5) — This method produces a 128-bit digital signature (a 16-byte digital key), which is used to verify the content, source, and integrity of the data.
- Secure Hash Algorithm-1 (SHA1) — This method produces a 160-bit digital signature (a 20-byte key), and is considered to be more secure than MD5.



Note: 3Com recommends that you use SHA1 for authentication.

IKE Proposals

If you have decided to use a Keying Mode that uses IKE, you need to select an IKE proposal. An IKE proposal selects the security attributes which will be used to protect the VPN connection between VPN initiator and terminator.

IKE operation is divided into two phases:

- Initially, the device negotiates Phase 1 of the IKE and establishes a shared, secure channel.
- In Phase 2, the device establishes keying material for the VPN. Phase 2 is much quicker than Phase 1, since it can rely on the checks established during Phase 1, without needing to re-establish a shared, secure channel.

The device has the following preconfigured IKE proposals:

- DES - SHA1-PSK

You can select the IKE proposal through the IKE Proposals page of the LSM. Alternatively, you can configure your own IKE proposal. For more information, see the *LSM User's Guide* or the online help.

IPSec Tunnel Setup

All IPSec tunnels terminate in a security zone. The traffic received over the VPN has unrestricted access to all devices within the tunnel security zone, and all devices within this security zone have unrestricted access to the VPN.



Note: A VPN tunnel can also be used as the default route for traffic by checking **Use tunnel as default route**, on the Add page (**VPN > IPSec/IKE > Add**). This overrides the default route configured on the external security zone.

DHCP Relay over VPN

Use DHCP relay over VPN if you have a centralized DHCP server, in order to obtain IP addresses from a central site, over a VPN tunnel.

For more information on configuring DHCP over VPN Relay on your network, see [“Using DHCP Relay over VPN” on page 31](#).

Site-to-Site VPN Operation

The X family device either immediately attempts to establish a VPN tunnel to a remote site, or establishes a tunnel on demand (depending on your configuration). If you have configured the SA for IPSec tunnel mode, you need to manually specify the list of IP subnets in the destination networks that are accessed using this tunnel. If you have configured the SA for GRE and L2TP use, you need to configure a GRE Virtual Interface and use static or dynamic routing. For more information, see [“Configuring Your Network” on page 57](#).

For IPSec transport mode, data is automatically GRE encapsulated; the device supports IP multicast and dynamic IP routing in this mode. When a GRE IPSec transport mode security association is configured, you must specify the IP address associated with the GRE interface, when configuring the GRE virtual interface.

Client-to-Site VPNs

A client-to-site VPN is established between an office site and a remote mobile or telecommuter, using a PC connected to a public network such as the Internet. Users connect to the X family device over the Internet, using a VPN client (included with most operating systems).

This type of connection is a many-to-one connection, in that many client devices can share the same method of setting up a VPN connection with a single device.

These VPNs connect telecommuters, mobile workers, employees requiring remote access after hours, contractors, and external business partners to company resources.

Remote users gain access to the company LAN using whatever access method they have at their location and create a secure 'VPN tunnel' to the network across the Internet.

Figure 4-2: Client-to-Site Connection



Supported Protocols

For client-to-site VPN connections, the X family device supports the following protocols:

- IPSec tunnel mode
- L2TP/IPSec
- PPTP (with up to 128-bit MPPE) (useful for Linux or Mac OS X clients)

Back-end user authentication is performed using either a RADIUS server or the local device user database.

IPSec Tunnel Mode

When the remote user requests access to the company LAN, the VPN client on the remote worker's PC initiates the tunnel across the internet to the company LAN. An X family device on the company LAN terminates the tunnel. The VPN client encapsulates and secures complete IP packets and tunnels them across the Internet to the device, which forwards the unencapsulated packets to the destination address on the company LAN.

The remote user's PC and the destination address on the company LAN do not need to use Internet routable IP addresses because IPSec encapsulates complete packets, so the private IP addresses are hidden within the IPSec tunnel mode packets.



Note: Note that proprietary security applications are usually used to implement IPSec tunnel mode. These additional VPN gateway applications allow VPN pass-through over firewalls between the VPN client and the Internet.

L2TP/IPSec

L2TP over IPSec is the recommended method for setting up client-to-site VPNs.

Layer 2 Tunneling Protocol (L2TP) is a mechanism that provides a PPP (Point-to-Point Tunneling Protocol) connection between a user and a terminating device over an IP network. PPTP is the protocol that is typically used to allow a dial-up user to connect to the Internet, authenticate and obtain their IP configuration in order to access the private LAN behind the L2TP terminator.

Although PPTP provides a level of security and authentication, when L2TP is run over IPSec (L2TP/IPSec) much better security and authentication is provided.



Note: When secured with IPSec, L2TP uses IPSec transport mode, as both the VPN client and L2TP terminator have a public IP address.

After completing the IPSec connection between the client and the device, an L2TP tunnel is established and the PPP connection within L2TP authenticates the VPN client user and provides the VPN client with an appropriate *local* IP address.

The PPTP connection encapsulates complete IP packets and hence the VPN client can communicate with the hosts on the private LAN that are not using public IP addresses.

Authentication — For PPP-based authentication, the following authentication protocols are supported:

- PAP
- CHAP
- MS-CHAP
- MS-CHAPv2

Using L2TP Clients — Microsoft includes L2TP/IPSec as their standard VPN client. Windows 2000 includes IPSec and L2TP as separate components. Windows XP provides an integrated L2TP/IPSec client.

Microsoft also has an integrated L2TP/IPSec client for clients prior to Windows XP.

For more information on Microsoft Windows support for VPN clients, see your Windows documentation and to the Microsoft site: <http://support.microsoft.com>.

PPTP with MPPE

Point-to-Point Tunneling Protocol (PPTP) provides a Point-to-Point connection between a user and a terminating device over an IP network. PPTP is a legacy protocol, found in older versions of Microsoft Windows. It uses the Microsoft MPPE standard for encryption. You may want to use this method if your network has large numbers of PC clients with older versions of Windows (prior to Windows 2000 and XP).

Summary of Client-To-Site VPN Methods

[Table 4-2](#) summarizes the alternative methods of setting up a client-to-site VPN connection.

Table 4-2: VPN Client-to-Site Setup Methods

Method	Description
IPSec	Provides both encryption and tunneling. IPSec is installed on Microsoft Windows 2000 and XP. The X family device supports the widely used SafeNet VPN client
L2TP over IPSec	Most Windows clients now support this method and it is the recommended method for client VPNs. IPSec provides the encryption; L2TP provides the tunnel
PPTP	PPTP is a legacy protocol, used in older versions of Microsoft Windows. It uses MPPE as the standard Microsoft encryption method. You may want to use PPTP if your network has large numbers of PC clients with older versions, supporting only PPTP

Using the Default Security Association

The X family device provides a single Security Association, the *Default SA*, which you can use for all client-to-site VPN connections. This enables multiple client devices to concurrently connect to the device, without you needing to configure a separate VPN connection for each device.

During the establishment of a VPN tunnel using IKE, the device will run through the list of SAs and try to authenticate, using the gateway IP address or the IKE local identifier to determine the correct SA. If it cannot find a suitable SA, it will use the *Default SA* for authentication.

The *Default SA* is a simplified SA that uses IKE.

Destination networks are not configured in the Default SA. If a remote device establishes a VPN and negotiates an unknown destination network, the device will associate this network with the Default SA.



Note: Although the Default SA is preconfigured, you must enable it before it can be used.

If you need to, you can edit some of the parameters of the Default SA, but you cannot delete it. Manual Key mode is not supported with the Default SA. For more information on the Default SA, see the online help.

Configuring Client-to-Site User Authentication

The X family device supports several user authentication options. They require a user to log in to the firewall and authenticate, before accessing the VPN. These are:

- *RADIUS authentication.* The preferred method, especially if you have large numbers of users.
- *LDAP authentication.*
- *X family local database authentication.* More suited for small sites, with known and trusted users.

See [Chapter 9, “User Authentication”](#) for details on configuring user authentication.

Client-to-Site VPN Operation

A remote user, connecting to the corporate site over the Internet, typically receives a global IP address from their ISP. They use a VPN client to automatically (for example, via dial-on-demand routing) or manually (using dial-up VPN networking) connect to the corporate X family device's public IP address, using the appropriate VPN protocol.

If using IPSec, the client initiates IKE Phase 1 and authenticates with the device, using a shared secret or certificate. IKE Phase 2 establishes an IPSec SA, which is used to obtain the encryption and authentication keys, needed to encrypt traffic between the client and the X family device.

If using PPTP or L2TP, PPP user authentication is performed. PPP provides the client with their IP address, DNS and WINS information. PPTP negotiates MPPE to encrypt the link.

Compatibility

The X family client VPN termination is interoperable with the following VPN clients:

- Microsoft native PPTP client on Windows 95, 98, ME, NT, 2000 and XP
- Microsoft native L2TP/IPSec client on Windows XP

Advanced VPN Configuration

Advanced configuration features are described in this section. These include:

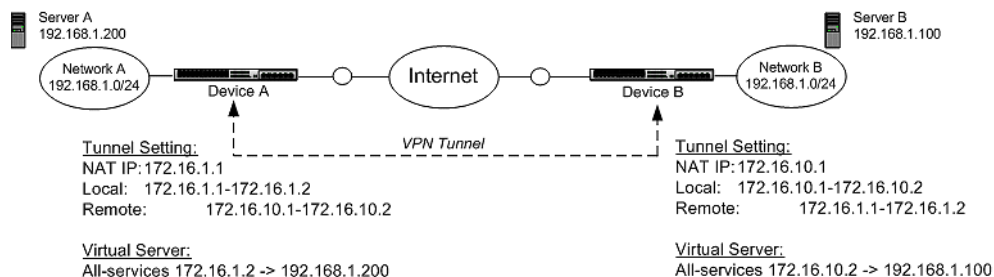
- Setting up NAT within a VPN tunnel
- Setting up a VPN supernet
- Configuring your network

Setting up NAT Within a VPN Tunnel

One-to-one NAT for VPN tunnel allows you to perform NAT on traffic sent over the VPN tunnel. This allows multiple remote VPN sites to use the same IP subnet.

VPN NAT is a useful feature when the local and remote networks of the VPN tunnel overlap; VPN NAT can be used to translate the networks onto another neutral address space. In this scenario, both sides of the VPN tunnel must perform NAT on traffic sent over the tunnel.

The diagram below shows an example configuration where the local networks of the two devices are identical. Network A is able to access Server B via the virtual server and Network B is able to access Server A via the virtual server.




The following procedure describes the steps required to configure VPN NAT for this scenario.

- STEP 1** On Device A, configure the IPsec security association for the VPN tunnel for Device B:
- STEP A** Enable IPsec tunnel connections.
- STEP B** Configure the Local ID (Local Networks IP address range).
 Using this example, the Local ID (Local Networks IP address range) would be 172.16.1.1-172.16.1.2.
- STEP C** Configure the Remote ID.
 Using the example, the Remote ID (Remote Networks IP address range) would be 172.16.10.1-172.16.10.2.
- STEP D** Enable NAT of local network addresses.
- STEP E** Specify the NAT IP address.
 Using the example, the NAT IP address would be 172.16.1.2.
- STEP F** Make sure that the termination zone for the Security Association is set to a virtual zone that contains no physical ports.

- STEP 2** Configure the Virtual Server to be used by the tunnel.
- STEP A** From the LSM menu, select **Firewall > Virtual Servers**. Click **Create**.
- STEP B** Select **All** from the **Services** drop-down list.
- STEP C** Specify the **Local IP Address** which traffic will be redirected to.
Based on the example, specify 192.168.1.200.
- STEP D** Specify the Public IP address which users or devices will use to access services.
(This address must be part of the Local ID ip subnet specified for the VPN tunnel).
Based on the example, specify 172.16.1.2.
- STEP 3** Configure firewall rules to allow traffic over the VPN tunnel.
See [Chapter 5, “Firewall Rules”](#) for more information about firewall rules.
- STEP 4** Repeat Steps 1 through 3 for Device B.

Setting Up a VPN Supernet

VPN supernet support makes configuration and management of hub and spoke VPN deployments easier and more efficient by allowing a single IP subnet to be used for central and remote sites. This configuration is useful if you have a central headquarters site and many branch offices as illustrated in the following figure:



With VPN Supernet:

- All sites use parts of a single destination subnet (central IP subnet).
- Only the device at the Headquarters needs to be updated when a new remote site is added.
- Same IP subnet can be shared between the Headquarters and branch offices.
- All remote sites do not have to use the supernet configuration.

Although all branch offices are configured with destination subnet as central IP subnet and use a portion of this address for themselves, each branch still negotiates the full IP subnet as the remote ID. The device determines whether to send the traffic locally or over the VPN tunnel only when routing the traffic.


Packet Routing

When VPN supernet is configured, the device uses the following priorities to route packets (listed from highest to lowest priority):

- Destined to local IP subnet
- Over VPN tunnel to specific IP remote network
- If the source IP does not match local ID, the packet is dropped
- To specific IP destination using a static route
- Over VPN tunnel used as Default route
- Default gateway

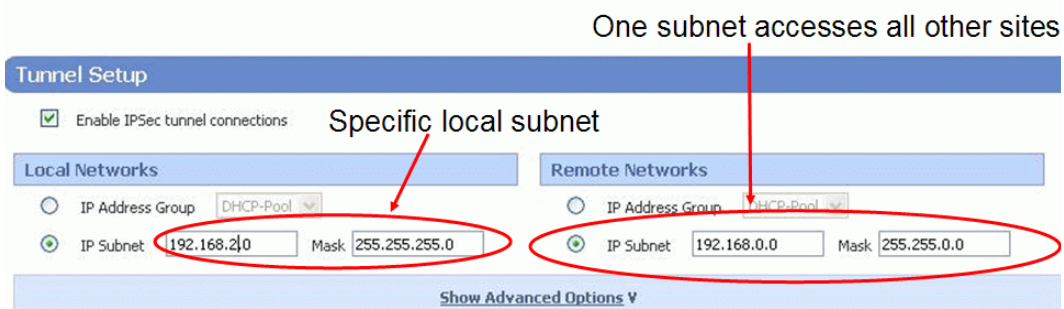
How to Configure VPN Supernet

To configure a VPN supernet define the local and remote networks of the IPsec tunnel appropriately. Any virtual interface subnet which is a strict subset of the remote network will automatically receive traffic in preference to the tunnel. The following figure provides an example showing the IP subnet shared by the central site and branch offices.



When you define the IPsec configuration for each branch office, configure the IPsec tunnel connection to use a specific local subnet for the Local Network and the shared subnet as the Remote Network address, as illustrated below:

Figure 4-3: Branch Office IPsec Tunnel Connection for VPN Supernet

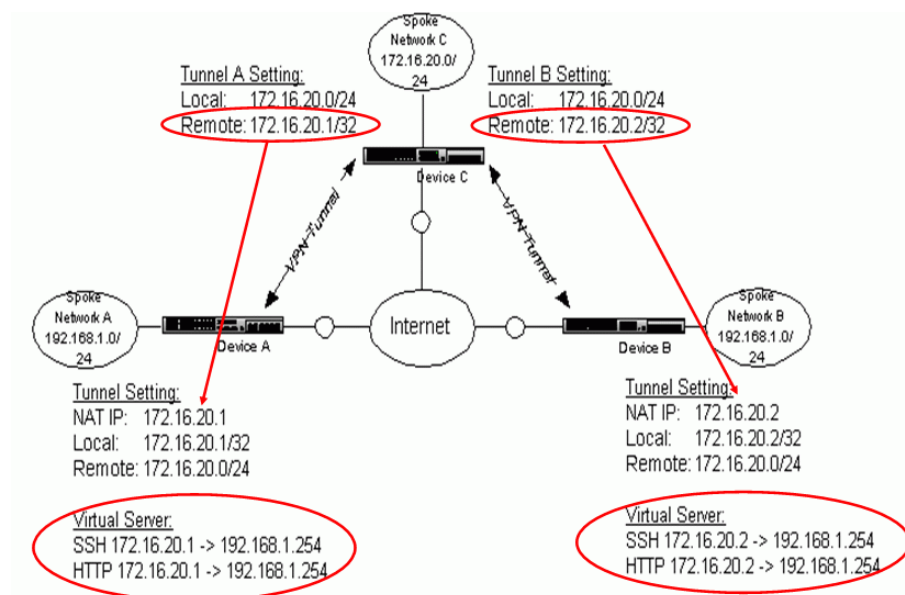


Deployment Example

This example describes how to configure a VPN supernet to set up a central office with a number of branch offices.

- The Central office uses a class C subnet of 172.16.20.0/24
- Each branch office has the same internal IP subnet
 - 192.168.1.0/24
 - X family device on each site has internal IP address 192.168.1.254
- Branch office VPN tunnel to HQ
 - VPN NAT allows each branch office to use only one VPN IP address
 - VPN Supernetting allows each branch office to use one of the HQ IP addresses
 - Branch offices use VPN IP addresses 172.16.20.1, 172.16.20.2, etcetera
- HQ needs to manage device at each branch office
 - Virtual server can map branch office VPN IP address to device internal IP address

The following figure illustrates the device deployment and the configuration settings for this scenario.



Configuring Your Network

You can also configure a VPN to:

- support site-to-site multicast routing.
- support site-to-site dynamic routing.
- forward DHCP requests to a DHCP server via a VPN.

[Chapter 3, “Network”](#) describes these features.

5 Firewall Rules

This chapter provides an overview of the X family firewall rules, and describes the steps required to configure them. It covers the following topics:

- [Overview](#)
- [X Family Firewall Components](#)
- [Providing Access to Internal Servers](#)
- [Firewall Rule Example](#)



Note: For detailed information about the options available on the Web Interface, see the **LSM User's Guide** or the online help.

Overview

The X family device is a Stateful Packet Inspection firewall, providing dynamic packet filtering. The device examines not only the packet header, but also checks the contents of the packet and monitors the status of the connection. To increase security, the device only opens TCP or UDP ports when an authorized device requests a connection to a specific TCP or UDP port number. The implementation of firewall rules enables you to use the X family dynamic packet filtering features.

Firewall rules control the flow of traffic between security zones (described in [Chapter 2, “Key Concepts”](#)), provide bandwidth management, and ensure quality of service.

You can use firewall rules to:

- Determine when and how traffic will be classified and controlled by the X family device.
- Prioritize specific types of network traffic.
- Permit or block a session request.
- Apply Web content filtering to specific categories of Web site, whether or not users are authenticated.
- Schedule when a service will be denied or permitted.
- Allocate bandwidth resources to a service and ensure a service has available bandwidth.
- Limit bandwidth resources to certain services.
- Time out idle sessions.
- Monitor network traffic.

Firewall rules can be used to screen both incoming and outgoing traffic. The following examples illustrate how the device could apply a firewall rule to inbound or outbound traffic:

- **Address inspection** — The device can block traffic that does not come from a known/trusted sender or is not going to an allowed destination.
- **Content inspection** — The device can apply Web content filtering to Web traffic (that is, HTTP or HTTPS GET or POST traffic). The device also screens data for known hacker attacks.
- **User authentication** — The device can apply user authentication such that it only allows traffic from users who have logged in to the device.

Firewall rules are applied based on connection requests. All permissions and restrictions governing both directions of a connection, for the life of the connection, are based on the network parameters (for example: source zone, destination zone, addresses) with which the connection is requested.

Firewall Rules Rank

Firewall rules are processed in order of precedence. The X family device applies the first rule that matches the category of the traffic in the request. If the traffic does not match any of the rules, the default behavior is to block the traffic. When you are defining firewall rules, ensure that you position the rule correctly in the list. For details on creating and ordering firewall rules, see the *LSM User's Guide* or the online help.

How Firewall Rule Enforcement Works

The following is an example of how the device enforces firewall rules for a session request, for example, when an unauthenticated (unknown) user requests a Web page using a browser.

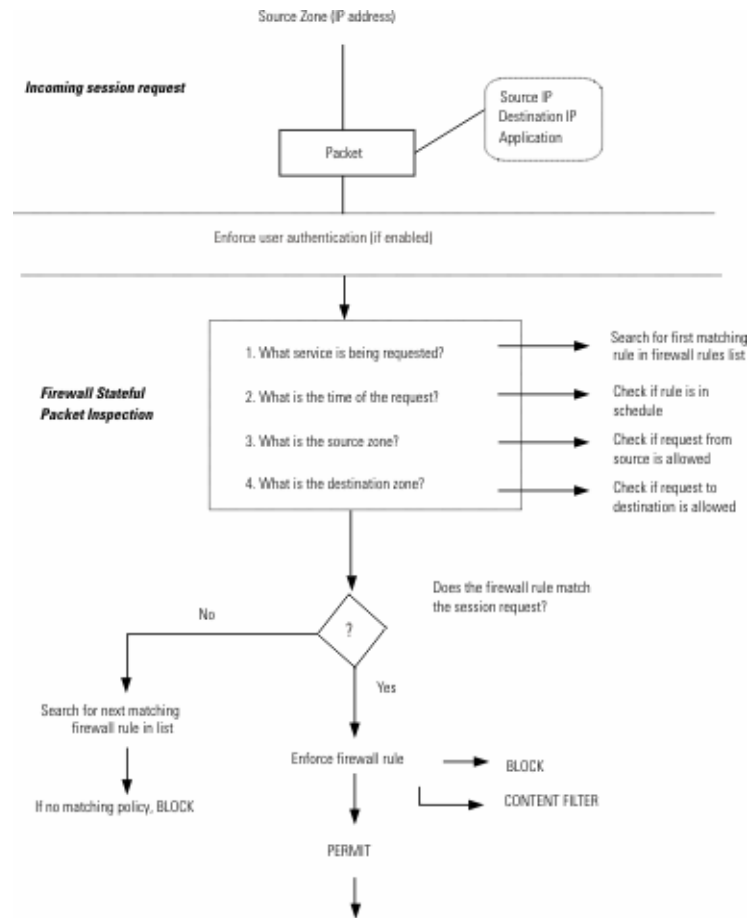
1. The user starts a Web browser. The Web browser resolves the DNS name for the URL and initiates a TCP connection to the target Web server via the X family device.
2. The device inspects the session header and identifies the following information about the request:
 - Source IP — The address of the device that initiated the request.
 - Destination IP — The address of the device for which the request is intended.
 - Application — Type of service/content and authenticated user (if any).

3. Using its routing table, the device decides which security zone the session has come from and which zone it is going to.
4. The device searches for the first firewall rule in its list that matches the session request. Rules are evaluated based on what options are configured:
 - User authentication
 - IP protocol service
 - Schedule
 - Source zone
 - Destination zone
 - Web content filtering
 - Anti-spam filtering

The firewall rule table is searched from the top of the table to the end (if necessary) looking for the first firewall rule that will match the session. Thus, it is important to put the most specific rules (for example, those configured with user authentication, IP address groups/ranges, or Web

content filtering) towards the top of the table. An illustration of how session requests are evaluated is shown in [Figure 5-1](#):

Figure 5-1: Handling Firewall Session Requests



5. When a rule is matched, the device enforces the firewall rule based on the action and logging configuration for the rule: Since the user is unknown, the firewall rule's Web profile is used; traffic is either permitted or blocked; the event is entered in the local log, entered locally and sent to a remote syslog server, or not logged at all.
6. If no matching firewall rule is found in the firewall rules list, the device denies the request using the implicit block rule preconfigured on the device.

X Family Firewall Components

This section provides more information on the components the X family device uses for policing traffic.

Services

Firewall rules are applied to services, which are applications that are known to the device. Services are associated with specific port numbers. The device supports a variety of common services types:

- *Predefined Services* — Applications that are included in the device's list of known applications.
- *Custom Services* — These services identify TCP/UDP ports or ICMP types used by protocols not known to the device, or specific IP protocols. You can define these services.

Custom Services

The device allows you to define new services that are not included in the predefined list. The device handles custom services in the same way as predefined services, verifying the validity of the port and protocol, based on the information you have defined. You need to specify:

- Service name
- Protocol (TCP, UDP, ICMP, ESP, AH, GRE, IGMP, or IPCOMP)
- Port number ranges or ICMP type
- For the IP protocol, the protocol number

Once defined, a custom service can be used within firewall rules or service groups.

Service Groups

You can combine services together to represent a service group. You can then apply a firewall rule to a service group, rather than to each predefined or custom service individually.

Schedules

The X family device allows you to create schedules that contain intervals of days and hours when the firewall rule applies. For example, Monday to Thursday, 8am to 5pm could be a “Work Hours” schedule. The **Always** (default) option can be used if you want the firewall rule to always be applied. Your schedule could include multiple sets of time intervals.

As an example, you may want to limit the use of a firewall rule defining a service such as FTP to office hours and weekdays only, to restrict the use of FTP access to this time. On the other hand, you may prefer to have a rule applying to your “Internet Access” Service Group enabled **Always**.

Source and Destination Addresses

A key component of a firewall rule is defining the source and destination addresses of the devices to which the firewall rule applies. This is specified using security zones, or specific IP address ranges within a security zone.

Security Zones

All firewall rules must define a source security zone and a destination security zone. Since firewall restrictions are only applied to traffic flowing between security zones (and not to traffic within a zone), you cannot define a firewall rule that has the same source and destination security zone. See [Chapter 2, “Key Concepts”](#) for more information about security zones.

For example, the email service could be only allowed for users in the LAN zone. This means that email traffic from any other zone will be blocked.

Providing Management Access

A preconfigured zone, called *this-device*, is used to provide management access to the X family device. If you want to set up a firewall rule for controlling access to the device for configuration purposes, you can select *this-device* for the source or destination zone. You can restrict management access to the device, or limit other types of traffic from or to the device.

IP Address Ranges

The device allows you to create groups of IP address ranges. IP address groups let you simplify the configuration process, and provide greater flexibility when defining policies. You can apply firewall rules to IP address groups, or use IP address groups with other firewall features.

IP address groups are used in conjunction with the source and destination security zones, to further define areas within a zone to which a firewall rule applies. For example, within the zone ‘LAN’, you may want to have a firewall rule for a subset of devices, defined by an IP address group. One example of this could be creating an IP address subset for the Administrator machines in your department. Although part of a larger zone, this subset or address range has more specific management and configuration access rights.



Note: Remember that firewall rules can only be applied if the traffic between IP address groups is also going between zones.

Firewall Actions

You can add, delete, or edit each firewall rule to perform one of the following actions:

- **Permit** — The device permits traffic of the specified service type.
- **Block** — The device blocks traffic of the specified service type.

In addition to the action applied, there are a number of general settings that are relevant to firewall rules. These are described later in this chapter.

Creating a VPN Firewall Rule

VPN tunnel traffic is allowed by default. However, you can create a firewall rule to restrict VPN tunnelled traffic to authenticated users from specific privilege groups, or provide bandwidth management over VPNs for specific applications (such as voice VPN). To achieve this, ensure that the

SA security zone is in a different security zone than the LAN physical ports and use firewall rules to define the traffic flow between these security zones.

When using VPNs, each type of VPN terminates within a predefined security zone. For example:

- PPTP security zone
- L2TP security zone
- IPSec SA security zone

Devices that are connected to the X family device by physical ports that are in the security zone where a VPN terminates have unrestricted access to the VPN tunnel. (Traffic over the VPN tunnel also has unrestricted access to these devices.)

If you wish to police the traffic between a LAN device and a VPN tunnel, ensure that the physical ports and the VPN security zone are separate and apply the appropriate firewall rules between these zones. (A security zone can be used purely for VPNs without being associated with any physical ports.)

Managing Bandwidth

The X family device provides a number of options for bandwidth management, enabling traffic shaping, ensuring Quality of Service and providing flexible control over network resources.



Note: If you enable bandwidth management, it may significantly impact device throughput.

Enabling Bandwidth Management

The **Enable bandwidth management** option can be used if you want bandwidth management applied to the firewall rule (this option is disabled by default).

For example, for voice VPN traffic, which is affected by network delays, causing jitter, it is important to prioritize this traffic over non-time critical applications, such as HTTP. One way of doing this is by guaranteeing bandwidth to the firewall rule that applies to voice traffic.

Selecting the Type of Bandwidth Management

The X family device offers two types of bandwidth management:

- **Per Session** — Bandwidth definitions are applied per matching session. The device assigns a fixed bandwidth allocation to a session. Use this option for services such as voice, to ensure a constant bandwidth (for example, of 200kbps) for the duration of the session.
- **Per Rule** — Bandwidth definitions are applied per rule. The device divides up the available bandwidth by the number of concurrent sessions. Because the available bandwidth is divided by the number of sessions, if there are 100 concurrent sessions, and you have defined a total of 15,000Kbps guaranteed bandwidth, then this would mean, on average, that 150 Kbps is assigned to each session.

Defining the Guaranteed Bandwidth

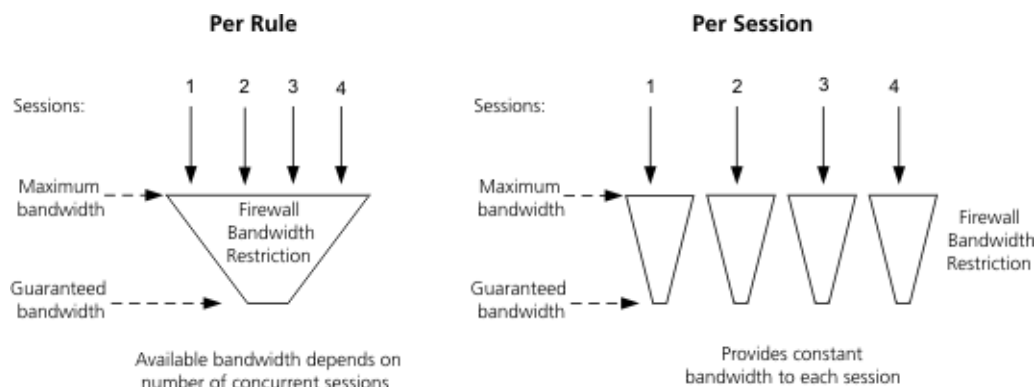
Guaranteed bandwidth can be applied to policies, to ensure that your mission and time-critical services (such as voice) receive the bandwidth allocation they require.

Defining the Maximum Bandwidth

The maximum bandwidth sets a limit to the amount of available bandwidth that the X family device can allocate to an application. Use this option to set an upper ceiling for bandwidth use, to ensure that this service does not drain resources from other network traffic (for example, by limiting the bandwidth available for FTP you ensure that file transfers do not block or slow down other access to the Internet).

An illustration of how the X family device applies bandwidth management to network sessions is shown in [Figure 5-2](#):

Figure 5-2: Applying Bandwidth Management



Selecting the Bandwidth Priority

The X family device uses priority queuing to assign traffic to one of four priority levels. It decreases the flow of lower priority traffic in order to guarantee bandwidth to an application. The top priority queue minimizes latency for latency-sensitive traffic (such as voice).

User Authentication

User authentication can be used to apply individual policies to individual (or groups of) users. You might configure user authentication in conjunction with policies to:

- Restrict an application only to particular users.
- Limit Internet access to specific groups of users or individual users.
- Allow some users to bypass Web content filtering.

User authentication is applied when a user accesses the X family device by entering the device's URL or IP address in their web browser. They then log in with a username and password.

If a firewall rule has authentication enabled, it can only be used if the user has previously authenticated with the device with a username that is associated with the privilege group specified within the firewall rule. Otherwise the rule is ignored and the device will search for the next matching firewall rule.

You can configure any firewall rule (applicable to any type of service) to require user authentication; the purpose of the authentication is to allow the device to apply the firewall rule to a group of users.

Enabling User Authentication

The authentication feature determines whether the user who is requesting a session has logged in to the X family device. The device's local database, an LDAP server, or a RADIUS server is used to authenticate the user.

When user authentication is used, you have two options:

- Any privilege group with firewall rule authentication
- Select a specific privilege group

To be able to apply privileges to individual users, you must enable user authentication on a user-by-user basis. See [Chapter 9, “User Authentication”](#) for more information about user authentication.

Allowing All Privileged Users

The option **Any privilege group with firewall rule authentication** allows access to authenticated users from all privilege groups which have the firewall privilege enabled.

Use this option to allow access to all authenticated users. For example, students and teachers at a school who log into the X family device for HTTPS access to class modules. This service would be open to all user privilege groups.

Allowing only Users in a Particular Privilege Group

The option to select a particular privilege group enables you to restrict access to authenticated users belonging to a specific privilege group. For example, teachers at a school who belong to the “staff” privilege group, would have access to additional services, such as FTP, for uploading class modules to the Web site.



Note: For information on creating privilege groups, see the **LSM User's Guide** or the online help.

Other Firewall Options

Monitoring Network Traffic

You can use the **Enable Logging** option to generate an event each time a firewall rule is enforced (that is, each time a new session starts or finishes). This can be used for diagnostic or testing purposes.

Timing Out Idle Sessions

Use the **Inactivity timeout** option to define when the device should end idle sessions. For example, an Internet session that has been idle for more than 30 minutes could be timed out.

Providing Access to Internal Servers

When the X family device is set up to provide Network Address Translation (NAT), then internal servers (behind the firewall) cannot be accessed directly by external devices, because the internal network is private and not exposed to external devices. If you are using NAT, the *Virtual Server* option (configured using the Virtual Servers page) enables external devices to access internal servers; that is, one-to-one NAT. You need to configure appropriate firewall rules to allow appropriate access. The device also allows you to configure Port Address Translation (PAT).

For more information on using NAT, see [“Configuring Network Address Translation \(NAT\)” on page 33](#).

Firewall Rule Example

This example illustrates a typical scenario for the application of a firewall rule.

Usage

An administrator for a school is requested to implement a set of firewall policies that restrict and allow access to the school’s Web server, depending on whether students or teachers need to access the site. Teachers need to use FTP to upload modules to the site. Students need to be able to access course modules on the Web site (using standard HTTPS browsing). Both groups need authentication (via log-in) for access to these services.

Setup

The administrator sets up two authentication privilege groups, *Teachers* and *Students*., defining a firewall rule for HTTPS access and a firewall rule for FTP access and ensures that authentication is enabled for each rule. The administrator defines a set of schedules which determine when students and staff are able to access their services ([Table 5–1](#)):

Table 5–1: Firewall Rule Setup

Action	Service	Source zone	Source IP	Destination zone	Destination IP	Privilege Group
Permit	HTTPS	LAN	ANY	WAN	Web server IP address	Any
Permit	FTP	LAN	ANY	WAN	Web server IP address	Teachers

Implementation

The X family device implements requests for FTP and HTTPS services as follows:

Access to FTP:

1. A teacher logs in to the device and is assigned the privileges of the *Teachers* privilege group. The teacher goes to the school's FTP server to upload lesson modules to the school's secure Web site.
2. The device first checks to determine that the teacher has logged in and then proceeds to process the request. Note that a student making a similar request would be denied, because the student does not have the privileges of a member of the *Teachers* privilege group.
3. The device searches for the first matching rule relating to the FTP service. When the matching rule is found, the device enforces that rule.
4. The device checks that the source IP (the teacher's computer) and destination IP (the FTP server) are within the permitted range defined by the firewall rule.
5. The device verifies that the request is within the hours of service defined for the firewall rule.
6. If the above criteria match, then the device allows the session setup to proceed.

Access to HTTPS:

1. A student browses to the school's Web server to download lesson modules from the school's secure Web site.
2. The device first checks to determine that the student has logged in and then processes the request.
3. The device searches for the first matching rule relating to the HTTPS service. When the matching rule is found, the device enforces the rule.
4. The device checks that the source IP (the student's computer) and destination IP (the Web server) are within the permitted range defined by the rule.
5. The device verifies that the request is within the hours of service defined for the firewall rule.
6. The student's privilege group permits HTTPS service to the requested Web site.
7. If the above criteria match, then the device allows the session setup to proceed.

6

Web Content Filtering

This chapter provides an overview of Web content filtering on X family devices, and describes the steps in configuring Web content filtering. It covers the following topics:

- [Overview](#)
- [Web Content Filtering Configuration Example](#)
- [Web Content Filtering Components](#)
- [Web Content Filtering Configuration Example](#)



Note: For detailed information about the options available on the Web Interface, see the **LSM User's Guide** or the online help.

Overview

Web content filtering allows you to control access to Web sites on the Internet, based on the user making the request, the security zone where the request originates, or the URL of the site requested. You can do this by creating Web filter profiles that call the Web Content Filtering Service as well as using manual filtering. If you apply both Web content filtering modes, manual filtering takes precedence over the Web Content Filtering Service. You can use manual filtering to identify specific Web sites to be allowed or denied, which overrides the categorization of those Web sites in the Web Content Filtering Service. Both filtering modes are described below.

Web Filter Profiles

You specify Web access using Web filter profiles. Web filter profiles block or permit access to predefined categories of Web sites defined by the Web Content Filtering Service, and you can further block or permit access to specific Web sites, domains, or pages on the basis of their URLs.

Privilege groups specify various levels of service access to authenticated users. You can assign a Web filter profile that is applied by a privilege group. You can also assign a Web filter profile to firewall rules

for application to unauthenticated users. The profile specifies what, if any, Web sites an unauthenticated user can reach.

Web Content Filtering Service

The Web Content Filtering Service is a content filtering service based on Web site classifications. This service is operated in partnership with SurfControl, a leading provider of content filtering solutions. The URL database has over 10 million entries and contains URLs in a variety of languages (65 languages) from over 200 countries. Web sites are classified into two main categories:

- **Core Categories** cover Web sites that contain offensive, potentially dangerous, or criminal content. The Web Content Filtering Service blocks URLs that are included in any Core category by default. If necessary, you can change the default setting for any category to allow access. For a list of the category types, see [“Core Categories” on page 140](#).
- **Productivity Categories** cover Web sites that could impair productivity when used in the work environment. The Web Content Filtering Service allows URLs that are included in any Productivity category by default. If necessary, you can change the default setting for any category to block access. For a list of the category types, see [“Productivity Categories” on page 142](#).

The Web Content Filtering Service is a subscription-based service that requires the purchase of the correct license for your product from a reseller. For details, see [“Purchasing a Web Filter License” on page 148](#).

Manual Filtering

You can enter a combination of URLs, domain names, IP addresses, keywords, and regular expressions to determine which Web requests are allowed or blocked.

Web Content Filtering Components

This section provides more information on the components that the X family device uses for Web content filtering.

Filtering Actions

Filtering actions allow you to specify what happens when a Web request is filtered. You can choose to permit Web requests or block them. You can also log Web access in the System Log. See [Chapter 11, “Events: Logs, Traffic Streams, and Reports”](#) for more information about the System Log.

Default Rule

The Default rule is implemented when a Web request is not a member of a currently blocked Web Content Filtering Service category or covered by a manual filtering rule, if the Web Content Filtering

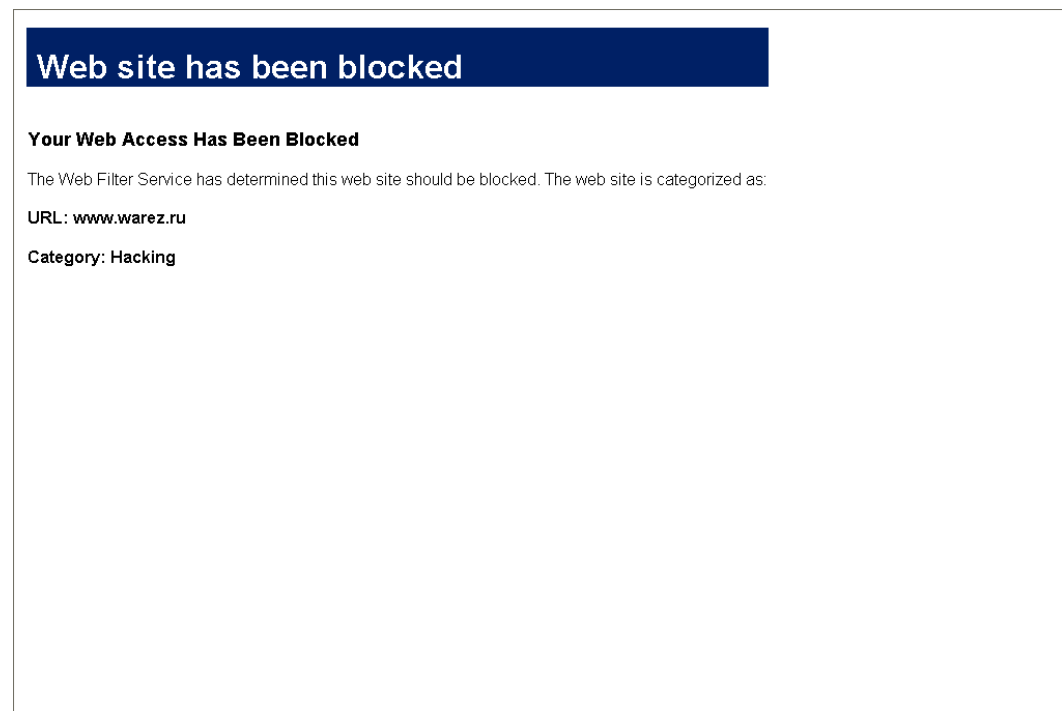
Service is not licensed, or the Content Portal Authority (CPA) Server cannot be contacted by the X family device. The Default Rule can be set to one of the following:

- **Allow unclassified or unknown sites** — the request is served and access is allowed.
- **Filter unclassified or unknown sites** — the filtering action is applied. You can set the filtering action to any of the following: **Block**, **Log**, or **Block and Log**.

Custom Response Page

The Custom Response page is displayed in the web browser via the standard HTTP 403 Access Denied response page when a Web request is blocked by the X family device. It comprises a fixed header and an HTML-formatted message of up to 1024 characters in length. You can configure the message appended to the header. An example is shown in [Figure 6-1](#):

Figure 6-1: Custom Response Page



URL Lists

Requests by client machines for connections to Web sites are checked against information in the *URL Permit and Block Lists* to see if they are permitted to view the site.

The URL Permit and Block Lists are simply lists of domain names, URLs, IP addresses, simple keywords or regular expressions entered by the administrator. Each entry in the list is referred to as a pattern. The device permits or blocks access if the Web request matches any of the patterns in the URL Permit and Block Lists.

The URL Permit List is checked first and therefore takes precedence over the URL Block List. Any request to a Web site that matches a URL Permit List pattern is allowed regardless of the Web Content Filtering Service categorization (if enabled). Any request to a Web site that matches a URL Block List pattern is blocked regardless of the Web Content Filtering Service categorization (if enabled).

URL Block List Example

If `192.168.2.100`, `www.Acme.com` (192.136.32.249) and `www.bbc.co.uk` were entered into the *URL Block List*, the X family device would permit and block the following sites:

Table 6–1: URL Block List Example

Site Accessed	Denied?
192.168.2.100	Yes
www.Acme.com	Yes
Cache.Acme.com	No
192.136.32.249	No
www.Acme.co.uk	No
www.bbc.co.uk	Yes
news.bbc.co.uk	No
www.bbc.co.uk/*	Yes

The maximum number of patterns that you can enter is 1000.

For details on creating and managing URL lists for manual Web content filtering, see the *LSM User's Guide* or the online help.

How Web Content Filtering Works

The following description describes the process flow of how a client Web site request is handled:

1. A Web browser on a PC issues an HTTP GET (or POST) request. If the first pertinent firewall rule permits HTTP (or HTTPS) service to the Internet, a connection is allowed.
2. If the firewall rule enables Web content filtering, the request is inspected; otherwise, the request is blocked.
3. The device inspects the session header of the request and identifies the IP address of the Web server. (This enables Web sites to be identified by domain or IP address, and prevents the Block List from being undermined by Web requests for its IP address.)
4. The device checks whether there is a user logged in to the device from this PC. If so, it checks the privilege group associated with that user to determine the appropriate Web filter profile to use, and whether Bypass Web Filtering is also enabled. (If the user is not logged in, the device uses the Web filter profile associated with the firewall rule itself.)

5. If the user is logged in and has Bypass Web Filtering specified in the associated privilege group, the request is passed on to the Web server without further inspection.
6. The device checks whether manual filtering is enabled in the Web filter profile. If so, it checks the Custom Filter URL Permit List for a pattern match. If there is a match, the request is passed on to the Web server without further inspection.
7. If there is no match in the URL Permit List, the device checks the URL Block List for a pattern match. If there is a match, the filter blocks the request.
8. If there is no pattern match in the URL Block List, the device checks to see if the Web Content Filtering Service is licensed and enabled. If it is enabled, the device contacts the Web Content Filtering Service server to determine if the URL matches a category included in the Web Content Filtering Service database. If the URL matches a blocked category, the request is blocked, executing the action configured for the Web Content Filtering Service: block only, log only, or block and log. If the URL matches a permitted category, the request is passed on to the Web server without further inspection.



Note: A local cache is available for the Web Content Filtering Service to speed up the filtering process. This cache can be configured using the Command-Line Interface.

9. If the request is not a member of a currently filtered category or covered by an entry in the Custom Filter List, the Web Content Filtering Service is not licensed, or the CPA Server cannot be contacted by the device, the Web filtering default rule is applied.
10. If the firewall rule has logging enabled, and the device denies access to a Web site based on the Custom Filter List or the Web Content Filtering Service, a “Warning” level Security event is sent to the Firewall Block Log. For allowed requests, “Informational” events are logged in the Firewall Session Log.
11. If access is permitted, the Web page is served to the user; if not, a customizable block response page is presented instead.

Web Content Filtering Configuration Example

The following example illustrates a typical scenario for the application of Web content filtering. For details on configuring Web content filtering, see the *LSM User's Guide* or the online help.

Background

An X family device administrator for a school wishes to restrict student access to illegal and productivity-draining Web sites and Web content, while ensuring that teachers have broad access to legally permissible Web sites. The administrator will not be subject to Web content filtering. Both groups are expected to log in for Internet access; the master registration file is on an LDAP server.

Setup

The administrator sets up a firewall rule allowing access to the service group web (consisting of HTTP and HTTPS), enabling Web content filtering, and assigning the Web filter profile *student* to unauthenticated users:

Table 6–2: Firewall Rule Setup

Action	Service Group	Source Zone	Destination Zone	Profile
Permit	web (HTTP, HTTPS)	LAN	WAN	student

The administrator sets up two Web filter profiles, *student* and *teacher*.

- The *student* profile defines the Web sites that students are prevented from accessing via the Web Content Filtering Service. Both Core Categories and Productivity Categories are blocked; unclassified sites are also blocked. In addition, the Web site YouTube.com is specifically blocked by a custom filter entry. This profile is also used, as defined in the controlling firewall rule, for Web traffic from users not authenticated by login.
- The *teacher* profile defines the Web sites that teachers are prevented from accessing via the Web Content Filtering Service. By default, Core Categories are blocked, and Productivity Categories are permitted; if a site is not classified, it is permitted.

Table 6–3: Web Filter Profile Setup

Name	General Configuration	Default Rule	Filtering Action	Custom Filter List
student	Enable Web Content Filtering Service Enable Manual URL Filtering Create default firewall rule	Filter unknown or unclassified sites	Block and Log	(Block) www.YouTube.com/*
teacher	Enable Web Content Filtering Service Create default firewall rule	Permit unknown or unclassified sites	Log	None

The administrator sets up three privilege groups: *student*, *teacher*, and *admin*. The privilege group names match the membership groups defined in the LDAP database. Privilege groups are assigned to users when they log in. The *admin* privilege group bypasses Web content filtering.

Table 6–4: Privilege Group Setup

Name	Web Filtering Bypass	Profile
student	No	student
teacher	No	teacher
admin	Yes	teacher

Implementation

Teacher Access to a Web site:

1. A teacher logs in to the device and is assigned the *teachers* privilege group. The teacher requests access to a Web site.
2. The device searches for the first matching (pertinent) firewall rule relating to the HTTP service. When the matching rule is found, the device enforces that rule.
3. The rule requires user authorization, so the device checks to determine that the teacher has logged in and then proceeds to process the request using the *teachers* privilege group.
4. The privilege group uses the Web filter profile *teachers*.
5. The Web filter profile does not include manual filtering, and the requested Web site is not listed in either the Core or Productivity Categories.
6. The default rule permits access, so the device allows the Web site to serve the page to the teacher.

Student Access to a Web site:

1. A student logs in to the device and is assigned the *students* privilege group. The student requests access to a Web site.
2. The device searches for the first matching (pertinent) firewall rule relating to the HTTP service. When the matching rule is found, the device enforces that rule.
3. The rule requires user authorization, so the device checks to determine that the student has logged in and then proceeds to process the request using the *students* privilege group.
4. The privilege group uses the Web filter profile *students*.
5. The Web filter profile *students* includes manual filtering, but the requested Web site does not match; however, the site is listed in one of the Productivity Categories.
6. The device blocks access to the site, displays a block page, and logs the event.

Administrator Access to a Web site:

1. The administrator logs in to the device and is assigned to the *admin* privilege group. The administrator requests access to a Web site.
2. The device searches for the first matching (pertinent) firewall rule relating to the HTTP service. When the matching rule is found, the device enforces that rule.
3. The rule requires user authorization, so the device checks to determine that the administrator has logged in and then proceeds to process the request using the *admin* privilege group.
4. The privilege group bypasses Web content filtering, so the device allows the Web site to serve the page to the administrator.

7 Anti-Spam Filtering

This chapter provides an overview of the Anti-Spam subscription service on X family devices, and describes the steps in configuring anti-spam filtering.

The X family device supports spam filtering to reduce the volume of unwanted and malicious email reaching users. Spam filtering is available both as a subscription service and through manual filtering. Spam filtering is applied based on firewall rules.

Anti-Spam IP Reputation Service

The **anti-spam IP reputation service** is a licensed service that provides spam filtering based on classifications of the sending IP address. The service assigns senders a risk based on factors such as the type and volume of email originating there and how long it has been sending the email. (Changes in these parameters often identify an activation of a botnet for sending spam.) The service includes both whitelists (known good addresses) and blacklists (known bad addresses). You can tune the filter, for example by changing the risk threshold above which email is not accepted.

The Anti-Spam IP Reputation Service provides email filtering based on weighted IP address classification. This service is operated in partnership with Commtouch Software Ltd., “dedicated to protecting the integrity of the world’s most widespread form of communication, e-mail.” Based on an analysis of billions of email messages each month, Commtouch classifies traffic based on factors such as the following:

- Statistical analysis of averages over time and recent changes of
 - Mail volume
 - Spam ratio
 - Valid bulk ratio
- Real-time zombie/botnet detection
- Continuously refined proprietary recurrent-pattern score
- Use of IP DNS and whois attributes

The Anti-Spam IP Reputation Service requires the purchase of the correct license for your product from a reseller. For details, see [“Purchasing an Anti-Spam License” on page 150](#).

Manual Filtering

Manual filtering lets you override the IP reputation service. You can create a custom whitelist or blacklist, permitting or blocking incoming email based on the sender’s IP address, IP address group, or IP host name. IP reputation information is not used to determine whether email from a given source is blocked or not.

8

Intrusion Prevention System

This chapter provides an overview of the Intrusion Prevention System (IPS) software.

Overview

The X family provides the TippingPoint Intrusion Prevention System (IPS) with Digital Vaccine (DV) filters that can be used to police your network to screen out malicious or unwanted traffic such as:

- Vulnerability Attacks and Exploits
- Worms
- Spyware
- Peer-to-Peer applications

In addition to the Digital Vaccine filters, IPS also provides Traffic Threshold filters you can use to profile and shape network bandwidth.

All IPS filtering occurs inline on traffic that has been permitted through the X family device firewall. Filtering is performed by the Threat Suppression Engine, custom software designed to detect and block a broad range of attacks at high speed. When a packet matches an IPS filter, the device handles the packets based on the Action configured on the filter. For example, if the action set is Block, then the packet is dropped. The device provides default actions to block or permit traffic with options to quarantine or rate-limit traffic and to notify users or devices when an action executes. Logging options are also available so you can review the types of traffic being filtered by the device. You can customize the default Actions, or create your own based on your network requirements.

A security profile defines the traffic to be monitored and the DV filters to be applied. Traffic monitoring is based on security zone pairs. For example, to create a security profile to monitor traffic coming from the WAN zone to the LAN zone, you select the security zone pair WAN ==> LAN. Then, you can

configure the DV filters to apply to that zone. The security zone pair specifies both the zone and the traffic direction which allows you to define separate security profiles for traffic in and out of a zone.

The default security profile is set to the ANY ==> ANY security zone pair with all IPS filters configured with the default Digital Vaccine settings. With the default profile in place, all incoming and outgoing traffic in any security zone configured on the device is monitored according to the recommended IPS filter configuration. You can edit the default security profile to customize the security zones that it applies to and create custom filter settings, or create your own security profiles as required.

Configuring IPS

You can monitor and configure IPS settings from the IPS menu pages available in the LSM. The following menu options are available:

- **Security Profiles** —View and manage the security profiles available on the device, view the security profile coverage by security zone.
- **Traffic Threshold** —View, manage, and create Traffic Threshold filters to monitor network traffic levels. These filters can be configured to trigger when traffic is either above or below normal levels.
- **Action Sets** — View, manage, and create actions that define the operations a filter performs when a traffic match occurs.
- **IPS Services** —Add and manage non-standard ports supported by the IPS device. Use this feature to configure additional ports associated with specific applications, services, and protocols to expand scanning of traffic. When filters scan traffic against the standard ports for listed services, the engine then accesses and scans traffic against the list of additional ports.
- **Preferences** —Reset IPS filters to the factory default values, configure timeout, logging, and congestion threshold settings to manage performance of the Threat Suppression Engine, configure the Adaptive Filter feature used to protect IPS performance from the effects of over-active filters.

For additional information, see the following topics:

- [“Security Profiles” on page 82](#)
- [“IPS Filters” on page 84](#)
- [“Action Sets” on page 88](#)
- [“Notification Contacts” on page 89](#)
- [“IPS Services” on page 91](#)
- [“Adaptive Filter Configuration” on page 91](#)

Security Profiles

On the X family device, a security profile defines the traffic to be monitored based on security zones (for example, ANY ==> ANY, LAN ==> WAN, or WAN ==> LAN) and the DV filters to be applied.

A security profile consists of the following components:

- **Identification** — Profile name and description.
- **Security Zones**— Specifies the incoming and outgoing security zones to which the security profile applies.
- **IPS Filter Category Settings**— determines the State and Action that applies to all filters within a given Filter Category group.
- **Filter overrides**—configure filter-level settings that override the Category Settings (optional).
- **Global Limits and Exceptions**—configure settings to apply filters differently based on IP address. You can limit filters to apply only to traffic between a source and destination IP address or address range, or apply filters to all traffic except the traffic between specified source and destination IP addresses or address ranges.

When a security profile is initially created, the recommended settings for all filter categories are set.

Default Security Profile

The default security profile is set to the ANY ==> ANY security zone pair with all IPS filters configured with the default Digital Vaccine settings. With the default profile in place, all incoming and outgoing traffic in any security zone configured on the device is monitored according to the recommended DV filter configuration. You can edit the default security profile to customize the security zones that it applies to and create custom filter settings, or create your own security profiles as required. 3Com recommends that you keep the default security profile with the default security zone pair ANY ==> ANY. This configuration ensures that all traffic will be inspected by the IPS using the default security profile, if the traffic does not match a more specific security zone configuration.

Applying Security Profiles to Traffic

In the IPS, it is possible for a packet to match more than one security profile depending how the security zone pairs are configured within each profile. As a general rule, the device will apply the filtering rules specified in the security profile that has the most specific security zone pair defined. To determine specificity, the device always considers the incoming zone first. See the following examples to see how the device applies filtering rules when a packet matches more than one security profile.

Example 1: Security Profile Zone Configuration

Security Profile	Applies To Security Zone Pair
#1	ANY ==> ANY
#2	LAN ==> WAN

In Example 1, a packet going from the LAN zone to the WAN zone matches both security profile #1 and #2. The X Family device applies the filtering rules from security profile #2 to the packet because the LAN zone is more specific than the ANY zone.

Example 2: Security Profile Zone Configuration

Security Profile	Applies To Security Zone Pair
#4	ANY ==> ANY

Example 2: Security Profile Zone Configuration (Continued)

Security Profile	Applies To Security Zone Pair
#5	ANY ==> WAN
#6	LAN ==> WAN

In Example 2, a packet going from the LAN zone to the WAN zone matches security profiles #4, #5 and #6. However, the device applies filtering rules from security profile #6 to the packet because the LAN zone is more specific than the ANY zone.

For details on creating and managing security profiles, see the online help.

IPS Filters

The IPS provides different types of filters to protect your network against attacks. For details, see the following topics:

- [“Digital Vaccine Filters” on page 84](#)
- [“Port Scan/Host Sweep Filters” on page 86](#)
- [“Traffic Threshold Filters” on page 87](#)

Digital Vaccine Filters

IPS Digital Vaccine (DV) Filters are used to monitor traffic passing between network security zones. Based on the security profiles configured, X family device applies the filters to traffic passing between network security zones. Each security profile has its own filter settings. Within a security profile, you can modify the filter (recommended) settings for a filter category and, if necessary, customize individual filters based on your network environment and security needs. The following sections provide an overview of the DV filters and the components used to configure them:

- [“About the Digital Vaccine Package” on page 84](#)
- [“Filter Components” on page 85](#)
- [“Categories and Category Settings” on page 85](#)

Categories and category settings are used to configure global settings for all filters within a specified category group.

- [“Filter Override Settings” on page 86](#)

Filter settings are used to override the global settings for individual filters within a category group.

About the Digital Vaccine Package

DV filters are contained in a Digital Vaccine (DV) package. All X family devices have a DV package installed and configured to provide out-of-the-box IPS protection for the network. After setting up the device, you can customize the filters in the DV through the LSM.

The filters within the DV package are developed by the Security Team to protect the network from specific exploits as well as potential attack permutations to address for Zero-Day threats. These filters include traffic anomaly filters and vulnerability-based filters. Vulnerability-based filters are designed to protect the network from an attack that takes advantage of a weakness in application software. For viruses that are not based on a specific vulnerability in software, the DV provides signature filters.

TippingPoint delivers weekly Digital Vaccine updates which can be automatically installed on the device (**System > Update**). If a critical vulnerability or threat is discovered, Digital Vaccine Updates are immediately distributed to customers. For details on automatic updates, see the *LSM User's Guide* or the online help.

Updates to Digital Vaccine Packages

Most customer choose to configure the AutoDV service on their X family device. When enabled, autoDV will automatically check for and download any Digital Vaccine upgrades on the Threat Management Center (TMC).

Alternatively, you can download Digital Vaccine updates manually from the TMC Web site (<https://tmc.tippingpoint.com>). You will need to create an account on the TMC before you can login and download the Digital Vaccine updates. When you register your device on <http://eSupport.3com.com> you will be automatically sent an email confirming the successful registration. The email provides a URL which you should click to access the TMC and create a login account. That URL will pre-populate most of the necessary fields on the TMC registration Web page. Once the registration page is completed you can register and create a new account to access updates.

Filter Components

IPS filters have the following components which determine the identity the filter type, global and customized settings, and how the device will respond when the Threat Suppression Engine finds traffic matching the filter:

- **Category** — defines the type of network protection provided by the filter. The category is also used to locate the filter in the LSM and to control the global filter settings using the Category Setting configuration.
- **Action set** — defines the actions that execute when the filter is matched.
- **Adaptive Filter Configuration State** — allows you to override the global Adaptive Filter configuration settings so that the filter is not affected by adaptive filtering (see [“Adaptive Filter Configuration” on page 91](#) for additional information).
- **State** — Indicates if the filter is enabled, disabled, or invalid. If the filter is disabled, the Threat Suppression Engine does not use the filter to evaluate traffic.

Categories and Category Settings

Categories and category settings are used to configure global settings for all filters within a specified category group.

DV Filters are organized into Categories and groups based on the type of protection provided:

- **Application Protection Filters** — defend against known exploits and exploits that may take advantage of known vulnerabilities targeting applications and operating systems. This filter type includes the following sub-categories: *Exploits*, *Identity Theft*, *Reconnaissance* (includes Port Scan/ Host Sweep filters), *Security Policy*, *Spyware*, *Virus*, and *Vulnerabilities*.
- **Infrastructure Protection Filters** — protect network bandwidth and network infrastructure elements such as routers and firewalls from attack by using protocols and detecting statistical anomalies. These filter types includes the sub-categories *Network Equipment* and *Traffic Normalization*.
- **Performance Protection Filters** — block or rate-limit traffic from applications that can consume excessive bandwidth, leaving network resources available for use by key applications. This filter type includes the following sub-categories: *IM*, *P2P*, and *Streaming Media*.

These categories are used to locate filters. *Category Settings* are used to assign global configuration settings to filters within a category. For example, if you don't want to use any filters to monitor P2P traffic, you can disable the P2P group in the Performance Protection category. You can configure the following global parameters:

- **State** — determines whether filters within the category are enabled or disabled. If a category is disabled, all filters in the category are disabled.
- **Action Set** — determines the action set that filters within a category will execute when a filter match occurs. If the *Recommended* action set is configured, filters within the category are configured with the settings recommended by the Digital Vaccine team, the group can have different settings.

For the best system performance, 3Com recommends that you use global Category Settings and the Recommended action set for all DV filters. However, in some cases, you may need to override the category settings and recommended action for individual filters due to specific network requirements, or in cases where the recommended settings for a filter interact poorly with your network.

Filter Override Settings

For the best system performance, 3Com recommends that you use global Category Settings and the Recommended action set for all DV filters. However, in some cases, you may need to override the category settings and recommended action for individual filters due to specific network requirements, or in cases where the recommended settings for a filter interact poorly with your network.

Filter override settings specify custom settings to be applied to the filter in the security profile. Once a filter has been customized, it is not affected by the global Category Settings that specify the filter State and Action. For details, see the *LSM User's Guide* or the online help.

Port Scan/Host Sweep Filters

A port scan attack scans a host looking for any open ports that can be used to infiltrate the network. A host sweep scans multiple hosts on the network looking for a specific listening port that can be used to infiltrate the network.

In the DV filters, the Port Scan/Host Sweep Filters (Filter numbers 7000- 7004) available in the *Application Protection Category - Reconnaissance* group are designed to protect the network against these types of attacks. These filters monitor the rate of connections generated by hosts on the network. The filter triggers when the connection rate during a specified interval goes above a given threshold.

The Port Scan/Host Sweep Attacks can only be used to monitor traffic on security zones that include physical ports. That is, you cannot run Port Scan/Host Sweep filters on VLANs or zones configured with a Virtual Server.

In the Category Settings, all Port Scan/Host Sweeps are disabled. If you want to apply these filters to the security profile, you can enable the filters, tune the *threshold* and *timeout* interval settings, and assign an action set to meet your network requirements. Because the *Recommended* setting for Port Scan/Host Sweep filters is disabled, you have to assign a specific action to the filter in order to enable it.

Filter Tuning

You can tune the sensitivity of Port Scan/Host Sweep filters by adjusting their *Timeout* and *Threshold* parameters. The timeout value is used in combination with the threshold value to determine whether or not an alert is sent.

For example, if the time interval is 300 seconds (5 minutes) and the connection threshold is 100 hits, then the filter is triggered every time the rate of connections exceeds 100, or a multiple of the threshold (101, 201, 301...) within the 300 second (five minute) time period.

The filters support any of the configured action sets available on the device. You can also configure IP address exceptions.

For details on configuring Port Scan/Host Sweep filters, see the *LSM User's Guide* or the online help.

Traffic Threshold Filters

Traffic threshold filters alert you and the device when network traffic varies from the norm. The device determines normal traffic patterns based on the network statistics over time. You can set four types of thresholds for each filter:

- **Major increase** — Traffic is greatly over the set threshold.
- **Minor increase** — Traffic is slightly over the set threshold.
- **Minor decrease** — Traffic is slightly below the set threshold.
- **Major decrease** — Traffic is greatly under the set threshold.

Thresholds are expressed as a “% of normal” traffic. For example, a threshold of 150% would fire if traffic exceeded the “normal” amount by 50%. A threshold of 60% would fire if the level of traffic dropped by 40% from “normal” amount of traffic.



Note: Network traffic rates are inherently erratic and can vary as much as 50% above or below the normal level on a regular basis. When you set up Traffic Threshold filters, avoid setting small variation percentages for minor and major thresholds to prevent the Traffic Threshold filter from triggering too often.

You can configure an action set for each threshold level the Traffic Threshold filter. When the filter triggers, the device executes the action specified for the threshold setting that triggered the filter. You can also configure traffic thresholds so that they only monitor traffic on the network without taking any action. All traffic threshold activity is recorded in the Traffic Threshold report (**Events > Reports > Traffic Threshold**).

Thresholds trigger when the traffic flow is above the *Above Normal* threshold percentage specified, or below the *Below Normal* threshold percentage specified by the set amounts. When traffic exceeds a threshold and returns to normal levels, the device executes the action specified for the threshold that triggered the filter and generates an alert. These alerts inform you of the triggered filter, when the thresholds are exceeded and return to normal, and the exceeded amount. After the filter triggers, you must reset it to re-establish it for use in the device. The filter is not disabled, but it does require resetting.



Note: A triggered Traffic Threshold filter will not perform functions until you manually reset it.

Traffic Threshold filter events are recorded in the Alert and Block logs (**Events > Logs**), based on the action set specified for the filter. Information on traffic threshold events is also available in the Traffic Thresholds report (**Events > Reports > Traffic Threshold**).



Note: Traffic Threshold filters are not included in the DV filter service. You must create and configure these filters based on your network requirements.

Action Sets

Action Sets determine what the X family device does when a packet triggers a filter. An action set can contain more than one action, and can contain more than one type of action. The types of action that can be specified include the following:

- **Flow Control** — determines where a packet is sent after it is inspected. A *permit* action allows a packet to reach its intended destination. A *block* action discards a packet. A block action can also be configured to *quarantine* the host and/or perform a *TCP reset*. A *rate limit* action enables you to define the maximum bandwidth available for the traffic stream.
- **Packet Trace** — allows you to capture all or part of a suspicious packet for analysis. You can set the packet trace priority and packet trace verbosity for action sets.
 - **Priority** — sets the relative importance of the information captured. Low priority items will be discarded before medium priority items if there is a resource shortage.
 - **Verbosity** — determines how much of a suspicious packet will be logged for analysis. If you choose *full* verbosity, the whole packet will be recorded. If you choose *partial* verbosity, you can choose how many bytes of the packet (from 64 to 1600 bytes) the packet trace log records.
- **Notification Contacts** — indicate the contacts to notify about the event. These contacts can be systems, individuals, or groups.

TCP Reset and Quarantine Actions

For Block action sets, you can configure TCP Reset and Quarantine options.

- **TCP reset** allows the device to reset the TCP connection for the source or destination IP when the Block action executes.



Note: Globally enabling the TCP Reset option may negatively impact your system performance. 3Com recommends using this option for issues related to mail clients and servers on email related filters.

- **Quarantine** allows the device to block packets based on the IP addresses in the packet that triggers the filter. When a filter with a quarantine option triggers, the device installs two blocks: one for the flow (as is normally done with Block actions) and another for the quarantined IP address. In addition to installing the two blocks, the device quarantines the IP address based on the instructions in the action set. For example, the user can display a Quarantine web page to notify the user of the problem and optionally provide instructions for fixing it, or the action may redirect all traffic from the quarantined IP address to a quarantine server that provides instructions to correct the problem.

Default Action Sets

The X family device is pre-configured with a collection of default Action Sets. You can edit the default settings for an action set, or create a new one. You cannot delete a default action set. The following actions sets are available:

- Recommended
- Block
- Block + Notify
- Block + Notify Trace
- Permit + Notify
- Permit + Notify + Trace

Rate Limit Action Set

A Rate Limit action set defines a maximum bandwidth that can be used by traffic that matches filters assigned to that action set. Incoming traffic in excess of this bandwidth is dropped. If two or more filters use the same rate limiting action set, then all packets matching these filters share the bandwidth. For example, if filters 164 (ICMP Echo Request) and 161 (ICMP Redirect Undefined Code) use the same 10 Mbps action set, then both “Echo Requests” and “Redirect Undefined Codes” filters share the 10 Mbps “pipe” as opposed to each filter getting a dedicated 10 Mbps pipe.

The supported rates are subject to restrictions based on the device model. Any of these listed rates can be used as long as it does not exceed 25% percent of the total bandwidth of the product.

Rate Limiting and Bandwidth Management for Firewall Policies

You can specify rate limiting for an individual filter or category setting. You can also specify bandwidth management for firewall rules between security zones. If there is a conflict between rules, the lower limit will be the value that will throttle the traffic first.

Quarantine Action Set

Quarantine Action Sets are Block action sets configured to block or redirect traffic from the host IP address for the filtered traffic. By enabling quarantine with a Block action set, you reduce the exposure of your network to internal and external threats.

When a filter with a quarantine option triggers, the device installs two blocks: one for the flow (as is normally done with Block actions) and another for the quarantined IP address. In addition to installing the two blocks, the device quarantines the IP address based on the instructions in the action set. For example, the user can display a Quarantine web page to notify the user of the problem and optionally provide instructions for fixing it, or the action may redirect all traffic from the quarantined IP address to a quarantine server that provides instructions to correct the problem.

You can review the list of currently quarantined IP addresses from the Quarantined Streams page (**Events > Managed Streams > Quarantined Streams**) in the LSM. You can also force an address into quarantine, or release a quarantined address. For details, see the *LSM User's Guide* or the online help.

Notification Contacts

Configuring notification contacts allows you to send messages to a recipient (either human or machine) in response to a traffic-related event that occurs on the X family device. The traffic-related

event can be the result of triggering an IPS filter configured with an action set that specifies a notification contact, or by triggering a Firewall Block rule with syslog logging enabled. A notification contact can be any of the following:

- **Remote System Log** — Sends messages to a syslog server on your network. This is a default contact available in all IPS action sets. Before using this contact, configure the IP address and port for the syslog server (**System > Configuration > Syslog Servers**). The Remote System Log is also the destination for all messages from Firewall Block rules with the *enable syslog logging option* turned on.
- **Management Console** — Sends messages to the LSM or the SMS device management application. This default contact is available in all action sets. If this contact is selected messages are sent to the Alert or IPS Block Log in the LSM, depending on whether a permit or block action has executed. When the device is under SMS management, messages are also sent to the SMS client application. This notification contact does not require any configuration, although you can change the default name and aggregation period.
- **Email or SNMP** — Sends messages to the email address or specified SNMP. All email or SNMP contacts must be added from the Notification Contacts page. If the default email server is not configured on the device, you will be prompted to configure it before adding a contact.

After configuring notification contacts, you can select them for IPS filter events when you create or edit the action set assigned to the filter. For Firewall Block rules, you can specify that messages be sent to the Remote System Log contact by selecting the *enable syslog logging* option when you edit the rule.

Alert Aggregation and the Aggregation Period

The X family device uses Alert Aggregation to protect system performance. Because a single packet can trigger an alert, attacks with large numbers of packets could potentially flood the alert mechanism used to send out notifications. Alert aggregation allows you to receive alert notifications at intervals to prevent this flooding. For example, if the aggregation interval is five minutes, the device sends an alert at the first IPS filter trigger, collects subsequent alerts and sends them out every five minutes.

On device, alert aggregation is controlled by the *aggregation period* that you configure when you create a notification contact. This setting is required for all notification contacts. For Email contacts, the aggregation period works in conjunction with the *Email Threshold* setting configured for the Email Server. By default, the device allows ten email alerts per minute. On the first email alert, a one-minute timer starts. The device sends e-mail notifications until the threshold is reached. Any notifications received after the threshold is reached are blocked. After one minute, the device resumes sending email alerts. The device generates a message in the system log whenever email notifications are blocked.



CAUTION: Short aggregation periods can significantly affect system performance. The shorter the aggregation period, the higher the system load. In the event of a flood attack, a short aggregation period can lead to system performance problems.

In addition to the user-configured aggregation period, the device also provides alert aggregation services to protect the device from over-active filters that can lower performance.

For details on configuring Notification Contacts, see the *LSM User's Guide* or the online help.

IPS Services

The TSE uses different filters to monitor different types of network traffic. For example, for http traffic, the TSE will monitor the default ports 80, 3128, 8000, and 8080 using filters that apply to http traffic. By default, the IPS port configuration uses the default port configuration for IP applications such as POP3, SNMP, SSH, and HTTP (for a list of all ports, see the IPS Services page in the LSM). If you know your network has traffic for a particular application, HTTP for example, on a non-standard port, you can add the port number to the http port configuration. Then, the TSE will also apply the http filters to traffic from the non-standard port.

When a non-standard port is configured, the TSE first applies filters to traffic from the standard ports for listed services, the engine then accesses and scans traffic against the list of additional ports. Each service supports 16 additional ports. However, 3Com does not recommend non-standard port configuration because it slows system performance.

Adaptive Filter Configuration

Adaptive Filtering is a mechanism to configure the TSE to automatically manage filter behavior when the X family device is under extreme load conditions. This feature protects your network against the potential adverse affects of a filter that interacts poorly with the network environment by preventing the device from entering High Availability mode.

The TSE Adaptive filtering mechanism monitors each filter to identify any suspected of causing congestion. When it identifies a filter, it manages the filter using one of the following methods, depending on how the global or filter-level Adaptive Filtering is configured:

- **Automatic Mode** — This setting enables the X family device to automatically disable and generate a system message regarding the defective filter.
- **Manual** — This setting enables the device to generate a system message regarding the defective filter. However, the filter is not disabled.

You can configure global settings for the Adaptive Filter from the IPS Preferences page (**IPS > IPS Preferences**) and the Configure Adaptive Filter Events page (**Events > Reports > Adaptive Filter**). At the filter level, you have the option to disable Adaptive Filter configuration so that a filter is never impacted by Adaptive Filter settings on the device. For details, see the *LSM User's Guide* or the online help.

9 User Authentication

This chapter provides an overview of user authentication and how to configure it on the X family device. It covers the following topics:

- [Overview](#)
- [Authentication and Firewall Privileges](#)
- [Methods of Authenticating Local Users](#)
- [Configuration Example](#)



Note: For detailed information about the options available on the Web Interface, see the **LSM User's Guide** or the online help.

Overview

User authentication on the X family device is a method of verifying the identity of a local user and associating the user with privilege rights configured on the firewall. You can use user authentication to:

- Enable VPN client access for remote users, over a secure VPN tunnel.
- Provide firewall rule authentication, ensuring secure access to network resources between security zones.
- Restrict Web site access of certain users.
- Permit certain users to bypass Web content filtering.

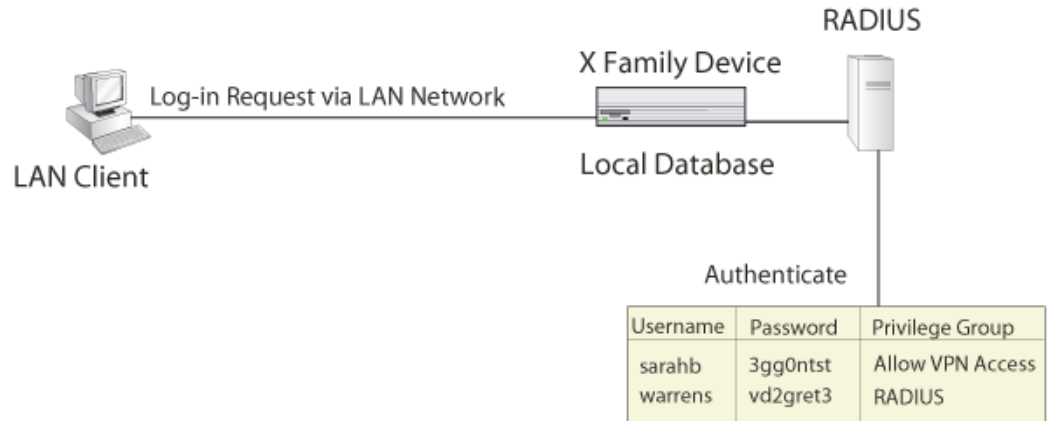
User authentication on the device can be implemented in conjunction with firewall rules, to restrict access to network services and applications. Firewall rules are described in more detail in [Chapter 5, “Firewall Rules”](#).

How Local User Authentication Works

The following explanation provides an overview of the authentication process as implemented by the X family device.

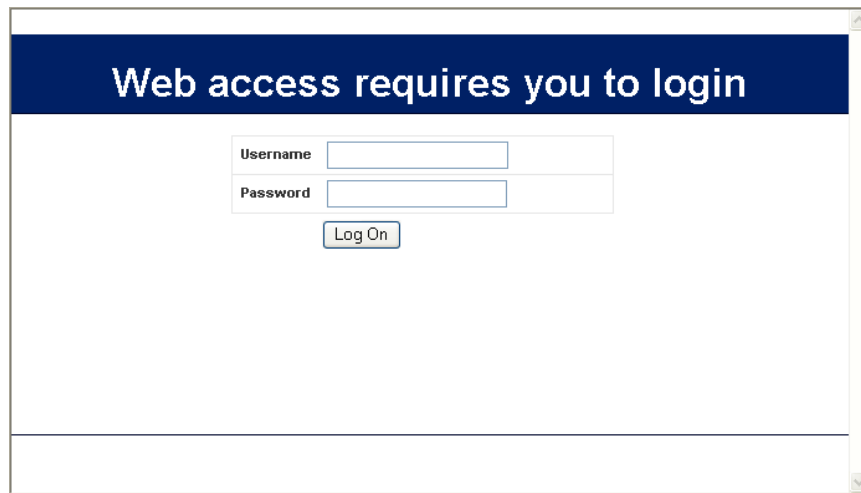
1. A local user logs on to the corporate network to gain access to a workstation or to corporate resources.

Figure 9–1: User Authentication



If username and password match:

1. Allow log-in
 2. Apply privileges of privilege group.
 3. Apply relevant firewall rules.
2. To access applications through the firewall, the user opens up a standard web browser and logs in using HTTPS and the LAN IP address of the device. A window is displayed, requesting a username and password. The user enters a username and password.



3. The device authenticates the user (checks that the user is listed in the database and that the username and password are correct). The methods available for user authentication are:
 - Using a RADIUS authentication server. The preferred method for large networks.
 - Using an LDAP authentication server.
 - Using the local device database. This can be used if no RADIUS or LDAP server is available, typically for small networks.
4. If no matching username and password can be located in the database, the device denies the login request.

If a matching user is found, the device applies the specific privileges associated with the privilege group to which the user belongs.

5. When a user requests a network service in another security zone, the device applies the relevant firewall rule for the type of service or application being requested:
 - If the firewall rule has user authentication enabled, then the device will allow access to this user, since he/she has authenticated, but will deny users who have not authenticated.
 - If the firewall rule is restricted to authenticated users belonging to a specific privilege group, then the device will apply the firewall rule only if the user belongs to the privilege group. If the user does not belong to the privilege group, the session request is denied.

Applications of Authentication

This section describes some scenarios in which you could apply authentication.

The use of firewall rules that utilize user authentication is optional. The main applications of user authentication are:

- If there are unknown devices/users within the LAN that you want to verify first, before providing access to resources. For example, wireless devices or “guest” users can be placed in a separate security zone, which requires user authentication to grant access to the main network services.
- If you need to provide differentiated access to users, based on their user profile/privilege group (for example, providing different privileges to teachers and students at a college).
- Once a VPN tunnel is established, additional firewall rule user authentication can optionally be used to provide differentiated access to services, based on the privileges of the user.
- The preconfigured privilege group *Allow_VPN_access* can be used to ensure that remote VPN clients (using L2TP/IPSec or PPTP) can authenticate with the device to access the secure private network.

Other Authentication Methods

The X family device also supports the use of certificates for secure VPN authentication. See [Chapter 4, “Virtual Private Networks”](#) for information on providing secure, authenticated remote VPN access to the network using X.509 certificates with IPSec.

Authentication and Firewall Privileges

Privileges are access rights to specific services on the network. Firewall privileges can be provided to users who have authenticated with the X family device. This section provides more information on using firewall privileges.

Types of Firewall Privileges

The following types of global privileges can be enabled for users:

- VPN client access
- Firewall rule authentication
- Web site access
- Web filter bypass

These global privileges are applied when a user logs in and is authenticated. They enable a user to bypass certain firewall restrictions that are applied using firewall rules.

VPN Client Access

Use the *VPN Client Access* option to enable authenticated VPN clients to access the network. For example, remote VPN users must log into the X family device using their VPN client and have this privilege before they are allowed to access the private LAN network.

This privilege only applies to L2TP/IPSec or PPTP VPNs.

Firewall Rule Authentication

Use the *Firewall Rule Authentication* option to grant the firewall rule authentication privilege to this user. This enables an authenticated user with this privilege to access services and devices in other security zones, if supported by the appropriate firewall rule.

You can disable this option if you want to restrict users in a privilege group to their own zone. See [Chapter 5, “Firewall Rules”](#) for more information.

Web Filter Bypass

Use the Web Filter Bypass option if you want users to be able to bypass a firewall rule enforcing Web content filtering.

For example, you may have a firewall rule for a school, enforcing Web content filtering on all Internet traffic. However, users belonging to the Privilege group *Teachers* could be allowed free access to the Internet, without Web content filtering. Although the firewall rule would apply Web content filtering to all Internet traffic, it would not enforce this for users belonging to the *Teachers* privilege group.

Privilege Groups

A privilege group enables you to define a set of firewall privileges that can be shared by multiple users, belonging to the group. All users in a privilege group share the same firewall privileges.

Users who log on to the X family device are authenticated and allowed to access all the privileges of the privilege group to which they belong.

The device has the following default privilege groups:

- *Allow_VPN_Access* — this can be associated with remote users that require client VPN access.
- *RADIUS* — this contains the privileges that a user who has been authenticated by RADIUS obtains. You can also configure the privileges on the RADIUS server itself.

You need to create additional privilege groups for firewall authentication and web filter bypass, or if you want a user to be associated with additional privileges, not covered by the default groups.

Once you have created your privilege groups and associated users with an appropriate group, you then use these groups in conjunction with firewall rules that enforce user authentication.

For a firewall rule that has authentication enabled, you can apply the firewall rule in one of two ways:

- To all local users who have authenticated with a privilege group that has firewall rule authentication enabled.
- Only to those local users who belong to a specific privilege group. This group must also have firewall rule authentication enabled.

See [Chapter 5, “Firewall Rules”](#) for more information.

Methods of Authenticating Local Users

These methods are available for authenticating users:

- Using the local X family device database
- Using an LDAP server
- Using a RADIUS authentication server

Using the Local Device Database

The local user database is used for authentication in the following circumstances:

- If you do not have a RADIUS server and want to authenticate on a small network, with a relatively small number of users.
- If remote authentication (RADIUS) fails or is disabled.

You can define a unique username and password for each user, and associate each user with a privilege group.

You can also delete users from the local database.

Using an LDAP Authentication Server

The X family device supports user authentication via **Lightweight Directory Access Protocol (LDAP)**. The following activities can be authenticated using LDAP:

- Web site access
- Web filtering bypass

Using a RADIUS Authentication Server

Using RADIUS is the recommended option, and requires that you have a RADIUS (Remote Authentication Dial-in User Service) server installed on your network. RADIUS authentication must be used if you have a large network with more than 100 users. Authentication and privilege group settings must be configured on the RADIUS server.

The X family device supports the following RADIUS servers:

- Microsoft ISA RADIUS Server (Microsoft Internet Authentication Service)
- Funk Steel Belted RADIUS

To set up a connection to the RADIUS server, see the *LSM User's Guide* or the online help.

Configuration Example

This example illustrates a typical scenario for the application of user authentication.

An administrator for a distributed organization needs to establish differentiated access to LAN applications/services for both on-site and off-site staff, based on their user profile.

The administrator sets up several privilege groups, including the *Sales* group. Users in the *Sales* group need to access the corporate intranet, to download email and to view confidential sales information located on a secure server.

The *Sales* privilege group has the following privileges:

Table 9-1: Privileges Assigned to Sales Group

Privilege	Value
VPN Client Access	Yes
Firewall Rule Authentication	Yes
Web Filter Bypass	Yes

To provide users in the sales group with access to the appropriate network services, the administrator defines a set of firewall rules and ensures that authentication is enabled based on the Sales privilege group.

Table 9–2: Firewall Rule Setup

Action	Service Group	Source zone	Source IP	Destination zone	Destination IP	Privilege Group
Permit	FTP, HTTPS	ANY	ANY	Flex	Intranet/Web server IP address	Sales
Permit	VPN	ANY	ANY	Flex	ANY	Sales
Permit	Email	ANY	ANY	Flex	Email server IP address	Sales

Implementation Example

Any user, such as a sales representative, who wants to access a secure service through the firewall would carry out a procedure similar to the following:

1. The remote sales representative dials in to the local ISP and establishes a secure VPN connection with the device. The sales representative is now able to access network resources which have user authentication enabled for the Sales privilege group as long as the resources are not restricted by firewall rules.
2. To access services that require prior X family user authentication, the sales representative must log into the device. Once authenticated, the sales representative is allowed to use all the applications and privileges that are available to the *Sales* privilege group.
3. When the sales representative attempts to access the email server to download email, the device searches for the first matching firewall rule relating to the email service. When the matching firewall rule is found, the device enforces the rules of the firewall rule (that is, allows the user to access the email server).

10 Certificates

This chapter provides an overview of certificates, and describes how to configure certificates on the X family device. It covers the following topics:

- [Overview](#)
- [Certificates and Public Key Cryptography](#)
- [Setting Up Your Certificate Infrastructure](#)
- [Methods of Obtaining Certificates](#)
- [Installing and Managing Certificates](#)
- [Certificate Configuration Example](#)



Note: For detailed information about the options available on the Web Interface, see the **LSM User's Guide** or the online help.

Overview

This section provides an overview of X.509 certificates as used by X family devices.

What are X.509 Certificates?

The X family device supports the use of X.509 certificates:

- for VPN authentication
- for secure management of the device using HTTPS authentication

A certificate is a data file that is used to verify the identity of a device. The file contains unique information about the device, such as a Distinguished Name (DN), email address, or domain name. This information is used to verify the identity of the device by validating it against information held by the Certificate Authority (CA). The certificate links this identity to a public key value, which is also contained within the certificate.

Authentication depends on the integrity of the public key value in the certificate. The role of a certificate is to guarantee that the public key bound to the certificate can be used to verify the identity contained in the certificate.

To prevent users from tampering with public keys, all certificates must be signed by the certification authority (CA). A CA is a trusted source that confirms the integrity of the public key value in a certificate. This could be a CA server within an organization, or a public company like Verisign.

Certificates and Public Key Cryptography

Public key cryptography is the mechanism behind the use of X.509 certificates.

In public key cryptography, data is encrypted and decrypted using public/private key pairs. Data encrypted with the public key may only be decrypted with the private key. The public key is made available to the public; the private key is always hidden, even from the user.

As an example of how public key cryptography is implemented, if a client device wants to establish a secure channel and send a message to the server, then the server can provide the client device with its public key. The client device encrypts the message using the public key. The server then decrypts the message using the private key.

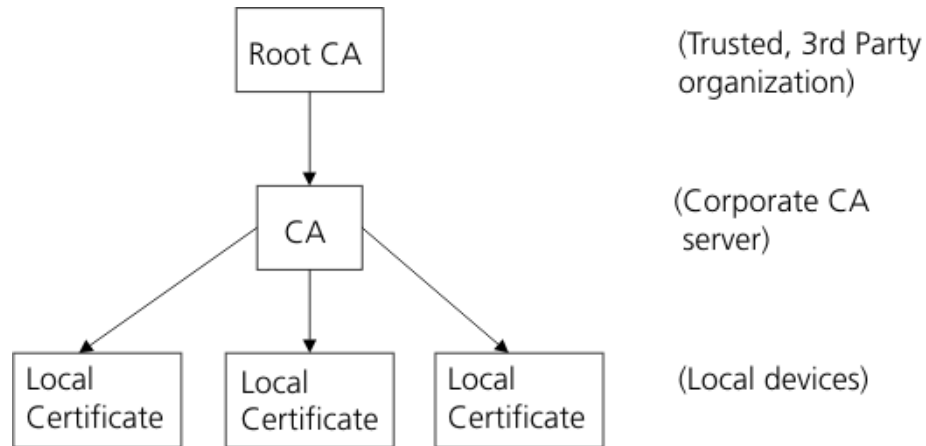
The reverse can also be used to verify the identity of the sender of the encrypted message. For example, the remote client encrypts the message with its private key, to identify itself as the sender. The server can then decrypt the message using the client's public key, and thus verify that it was the sender.

Public and private keys are the backbone of X.509 certificates.

Public Key Infrastructure

A hierarchical structure of trust is used for public key cryptography, as shown in [Figure 10-1](#):

Figure 10-1: PKI Hierarchy — CA Domain



Each device in this hierarchy has its own local certificate, which is signed by its parent CA. The root CA signs its own certificate.

In order to verify a certificate, you must be able to track it back from the local certificate to a root CA of a CA domain. Organizations that use certificates for internal purposes only may have their own CA domain and root CA server. The root CA certificate allows you to trust certificates lower down in the chain.

On the X family device, the trusted CA certificate that was used to sign a local certificate must also be installed before that local certificate can be used with IKE for VPN setup.

CA and Local Certificates

The following explains the difference between CA and local certificates:

- The CA certificate is an organizational certificate, installed on the CA server. The CA certificate verifies the local certificates by signing them.
- The local certificate is a personal certificate, installed on the X family device or remote device. Each device has a unique local certificate. The local certificate refers to the CA certificate for validation.

When setting up a VPN, the device first checks that the received local certificate is valid (Peer ID matches what is configured on the device), and then checks that the certificate was signed by one of the CA certificates.

Digital Signatures

A CA signs a certificate by adding its digital signature to the certificate. A digital signature is a message encoded with the CA's private key. The CA's public key is made available to applications by distributing a certificate for the CA. Applications verify that certificates are validly signed by decoding the CA's digital signature with the CA's public key.

Contents of a Certificate

In addition to the signature, all X.509 certificates have the following data:

- **Version** — This identifies which version of the X.509 standard applies to this certificate, which effects what information can be specified in it.
- **Serial Number** — The serial number is a unique number, issued by the entity that created the certificate. This information is used in numerous ways, for example when a certificate is revoked its serial number is placed in a Certificate Revocation List (CRL).
- **Signature Algorithm Identifier** — This identifies the algorithm used by the CA to sign the certificate.
- **Issuer Name** — The X.500 name of the entity that signed the certificate. This is normally a CA. Using this certificate implies that the entity that signed this certificate is trusted. (Note that in some cases, such as root or top-level CA certificates, the issuer signs its own certificate.)
- **Validity Period** — Each certificate is valid only for a limited amount of time. This period is described by a start date and time and an end date and time, and can vary between seconds to years. The validity period chosen depends on a number of factors, such as the strength of the private key used to sign the certificate or the amount one is willing to pay for a certificate. This is the expected period that entities can rely on the public value, if the associated private key has not been compromised.
- **Subject Distinguished Name** — The name of the entity whose public key the certificate identifies. This name uses the X.500 standard, so it is intended to be unique across the Internet. This is the Distinguished Name (DN) of the entity, for example,


```
CN=John Smith, OU=Software Division, O=Acme, C=US
```

 (These see the subject's Common Name, Organizational Unit, Organization, and Country.)
- **Subject Public Key Information** — This is the public key of the entity being named, together with an algorithm identifier which specifies which public key cryptographic system this key belongs to and any associated key parameters.

Certificates can only be manually configured.

Certificate requests include:

- Distinguished name (DN)
- Public key
- Set of attributes (optional)

This request is signed by the requestor's private key so the CA may verify authenticity. The CA transforms this request into a X.509 certificate.

Certificate Revocation List (CRL)

The CRL is a time-stamped list (maintained by the issuing CA) of issued certificates, which have later been revoked.

A certificate may be revoked for a number of reasons:

- Suspected compromise of the private part of a public / private key pair; this invalidates the public part
- Change of user details
- Certificate has expired

The CRL is periodically updated, as required, by the CA.

Setting Up Your Certificate Infrastructure

To support certificates on your network you will need to install a CA server and set up a mechanism for distributing signed local certificates to the devices in your network.

Installing a CA Server

You must set up a local CA server, such as Microsoft Windows 2003 Certificate Server, or use a third-party CA server. The X family device authenticates local certificates against the CA certificates distributed by the CA server.

In a typical certificate scenario, you would normally request a root CA certificate from a trusted organization such as Verisign. This certificate would be installed on your own CA servers. The CA servers would then typically sign and distribute the local certificates to all the devices in your network. This establishes the hierarchy of trust.

During a VPN setup, the local certificate is verified against the digital signature of the CA server, which is in turn verified against the signature of the root CA certificate. If this chain is broken — for example, the device does not recognize the root CA — the VPN will fail.

Installing Local Certificates on Your VPN Clients

You need to provide all your VPN clients with their local certificate. For VPN clients, the local certificate is commonly known as the “personal certificate.” Once installed, the personal certificate can be used by the VPN client to verify its identity. During a VPN connection between the client and the X family device, the device validates the client’s personal certificate against the CA certificate installed on the device.

Methods of Obtaining Certificates

There are three basic methods used to obtain certificates:

- Sending a certificate request to a Certificate Authority
- Directly importing a certificate with its private key
- Creating your own self-signed certificate

Sending a Certificate Request to a CA

If you are sending a request to a CA to issue you a certificate, you must provide your public key and some information about your device. This normally includes information such as its name and organizational address. If you ask a CA to issue a certificate for you, you will normally need to provide proof to show correctness of the information.

See the Web site of your CA for their certificate practise statement (CPS), defining the requirements and guidelines they follow for issuing certificates.

The X family device X.509 certificate implementation is interoperable with the following public CA servers:

- Verisign (<http://www.verisign.com/>)
- Entrust (<http://www.entrust.com>)
- Microsoft (<http://www.microsoft.com>)

Directly Importing a Certificate

You can export and import certificates along with their private keys. Be careful to ensure that you transfer your certificate files securely, as the private keys are the foundation for the certificate's security. A management password is used to protect this secret information.

Creating Your Own Self-Signed Certificate

The X family device only uses self-signed certificates for secure administration (HTTPS). Alternatively, a local certificate from a CA server can be used for secure administration. This avoids the initial warning message you will see when you first manage the device with its self-signed certificate.

If you are generating the certificate yourself, you must provide your public key and some information about your device, add any additional information (dates during which the certificate is valid, a serial number), and create the certificate.

Not everyone will accept self-signed certificates; one part of the value provided by a CA is to serve as a neutral and trusted introduction service, based in part on their verification requirements, which are published in their Certification Practice Statement (CPS).

Installing and Managing Certificates

For instructions on installing and managing certificates, see the *LSM User's Guide*.

Certificate Configuration Example

The following example illustrates a typical scenario for the application of certificates.

An administrator for a distributed organization wants to set up VPN access for all employees using certificates for authentication.

Certificate Setup

The administrator sets up a CA server on the local LAN, using Microsoft 2003 Certificate Server, and then creates and sends a certificate request to Verisign, a trusted third party certification organization. Verisign returns the signed certificate.

Using the signed certificate, the administrator creates the company's own CA certificates and signed local certificates. The administrator ensures that these are distributed to all devices which require them.

The administrator configures each of the devices, using the **Create Certificate Request** and **Import Signed Certificate** options available in the Local Security Manager (LSM) web application.

The administrator ensures that the connection to the appropriate Certificate Revocation List on the CA server is configured on each device.

VPN Setup

The administrator ensures that all VPN clients install their local certificates from the CA server. All local certificates are signed with the CA server's digital signature.

The administrator configures the *Default SA* (which is the Security Association for multiple concurrent client-to-site VPNs) to use local certificates for authentication. He also configures remote branch devices to use certificates for authentication with a unique SA for each site-to-site connection. The administrator sets up the Default SA IKE proposal as in [Table 10-1](#):

Table 10-1: Example X.509 IKE Proposal for Default SA

Field	Value
Name	default SA proposal
Authentication Type	X.509 Certificates
Local Certificate	Firewall_certificate_1
Peer ID Type	Distinguished Name

The above setup ensures that the local certificate called 'Firewall_certificate_1' is used to authenticate the VPN.

The administrator sets up a site-to-site IKE proposal as in [Table 10-2](#):

Table 10-2: Example X.509 IKE Proposal for Site-to-Site SA

Field	Value
Name	site to site proposal
Authentication Type	X.509 Certificates
Local Certificate	Firewall_certificate_2
Peer ID Type	Distinguished Name

Implementation

Any VPN client or remote device that wants to establish a secure VPN link with the X family device, would carry out a procedure similar to the following:

- 1 The VPN device sends its local certificate to the X family device, when requesting a VPN setup.
- 2 The X family device first checks that the certificate is valid (contains a matching peer ID type and value). It then verifies the certificate against the CA certificate.
- 3 Once authenticated, the X family device allows the VPN tunnel to be established.
- 4 The remote client/ VPN device also verifies the X family device's certificate and enables the tunnel to be established on its side.



Events: Logs, Traffic Streams, and Reports

This chapter describes the logs, views, and reports available to monitor system performance and traffic-related events triggered by firewall rules, Web content filters, IPS filters, and traffic threshold policies. In this section, you will review the information presented in the Events pages and learn how to manage the logs and reports. Only users with Super-user access may view all of the logs and reports available.

Overview

The LSM application provides the following information to view and monitor activity on the device:

- **Logs** — View information on system events and traffic-related events triggered by firewall, IPS, and traffic threshold security policies.
- **Managed Streams** — Review and manage traffic streams that have been blocked, rate-limited, or quarantined by IPS policies. You can also manually quarantine or release a quarantined IP address.
- **Health** — Review the current status and network performance of the device. Information includes memory and disk usage statistics, status of the Threat Suppression Engine and the Ethernet ports, and throughput performance.
- **Reports** — View graphs showing information on traffic flow, traffic-related events, and statistics on firewall rule hit counts and triggered filters (attack, rate limit, traffic threshold, quarantine, and adaptive filter).

For details, see the following sections:

- [“Logs” on page 112](#)
- [“Managed Streams” on page 115](#)
- [“Health” on page 115](#)
- [“Reports” on page 116](#)

Logs

The Logs menu pages provide information on system events and traffic-related events triggered by firewall, IPS, and traffic threshold security policies. Each menu page also provides functions to manage the log files.

When you review logs, you may also see the following type of administrator user levels. These users denote the type of account according to the interface they used in the device:


- **SMS** — Indicates the administrator used the SMS when the messages saved to the logs
- **LSM** — Indicates the administrator used the LSM when the messages saved to the logs
- **CLI** — Indicates the administrator used the CLI when the messages saved to the logs



Note: Users with any access level can view and print the system log, but only Administrator and Super-user level users can reset this log.

Log Maintenance

The X family device maintains two files for each log: a historical log file and a current log file. When the current log file reaches the default size (4MB), the log is de-activated and saved as the historical file, and a new log file is started as the current log. If a historical file already exists, that file is deleted. When the log is rolled over, the device generates a message in the Audit log. If you want to save log all data and create a backup, you can configure the device to offload log messages to a remote system log.

You can reset a log from its menu page, or use the Reset  function available on the System Summary page. For details, see the *LSM User's Guide* or the online help.

For details, see the following sections:

- [“Alert Log” on page 112](#)
- [“Audit Log” on page 113](#)
- [“IPS Block Log” on page 113](#)
- [“Firewall Block Log” on page 113](#)
- [“Firewall Session Log” on page 113](#)
- [“VPN Log” on page 114](#)
- [“System Log” on page 114](#)
- [“Configuring Remote System Logs” on page 114](#)

Alert Log

The Alert log contains information about network traffic that triggers IPS filters configured with a Permit + Notify or Permit+Notify+Trace action set. Any user can view the log, but only administrator and super-user level users can print the log.

To maintain a complete history of entries and provide a backup, you can configure the device to send Alert Log entries to a remote syslog server from the Notification Contacts page. For details, see the *LSM User's Guide* or the online help.

Audit Log

The audit log tracks user activity that may have security implications, including user attempts (successful and unsuccessful) to do the following:

- Change user information
- Change IPS, firewall, routing, or network configuration
- Gain access to controlled areas (including the audit log)
- Update system software and attack protection filter packages
- Change filter settings

To maintain a complete history of entries and provide a backup, you can configure the device to send Audit Block Log entries to a remote syslog server from the Syslog Servers page. For details, see the *LSM User's Guide* or the online help.

IPS Block Log

The IPS Block log contains information about packets that have triggered an IPS filter configured with a Block + Notify action set.

To maintain a complete history of entries and provide a backup, you can configure the X family device to send IPS Block Log entries to a remote syslog server from the Notification Contacts page. For details, see the *LSM User's Guide* or the online help.

Firewall Block Log

The Firewall Block Log captures information about events that have triggered a firewall rule that blocks matching traffic and has logging enabled.

A log entry is generated for each of the following events.

- Block web request event: occurs when the device blocks a Web request due to content filtering.
- Block event: occurs when a firewall rule with Block action is triggered.

To maintain a complete history of entries and provide a backup, you can configure the device to send Firewall Block Log entries to a remote syslog server from the Notification Contacts page. For details, see the *LSM User's Guide* or the online help.

Firewall Session Log

For firewall and Web content filter Permit rules with logging enabled, this log captures information on session creation and termination, including the time the session started, and the URL being accessed (for web requests). When a session terminates the Firewall Session Log shows how many bytes were transferred through the session.

A log entry is generated for each of the following events if the firewall rule had logging enabled.

- Web Request event: occurs when the device permits a web request to pass through.
- Session Started event: occurs when a firewall rule is triggered.
- Session Close event: occurs when the network connection is ended or closed due to inactivity.

To maintain a complete history of entries and provide a backup, you can configure the device to send Firewall Session Log entries to a syslog server from the Syslog Servers page. For details, see the *LSM User's Guide* or the online help.

Each log entry is tab-delimited. The log fields are populated based on the type of event being logged. If a field is not used, a tab is inserted to properly position the data in the next field.

VPN Log

The VPN log captures diagnostic messages relating to VPN tunnels to help troubleshoot and monitor VPN configurations. Each log entry is tab-delimited. The log fields are populated based on the type of event being logged. If a field is not used, a tab is inserted to properly position the data in the next field.

To maintain a complete history of entries and provide a backup, you can configure the device to send VPN Log entries to a syslog server from the Syslog Servers page.

Configuration

The logging level for the VPN log can be configured to provide more/less detailed information by configuring the **Enable Verbose** messages in the VPN Log option available on the IPSec Configuration page in the LSM.

System Log

The System Log contains information about the software processes that control the X family device, including startup routines, run levels, and maintenance routines. System log entries can provide useful troubleshooting information if you encounter problems with your device.

To maintain a complete history of entries and provide a backup, you can configure the device to send System Log entries to a syslog server from the Syslog Servers page. For details, see the online help.

Configuring Remote System Logs

All information logged by the LSM can be offloaded to a remote syslog server. Options to configure logging behavior for traffic-related events are available from the Edit Action Sets page (**IPS > Action Sets > Edit**) and the Edit Firewall Rule page. In order to use remote logging options, you must configure the contact information for the remote syslog servers.

For details on configuring, remote system logs, see the *LSM User's Guide* or the online help.

Managed Streams

From the LSM, you can view and manage traffic streams that are being monitored by the X family device.

The traffic streams include the following:

- **Blocked streams**— Traffic streams detected and blocked based on filters configured with a Block action set.
- **Rate-Limited streams** — Traffic streams detected and rate limited based on filters configured with a Rate-Limit action set.
- **Quarantined streams** — Traffic streams detected and blocked based on filters configured with a Quarantine action set, or quarantined manually.

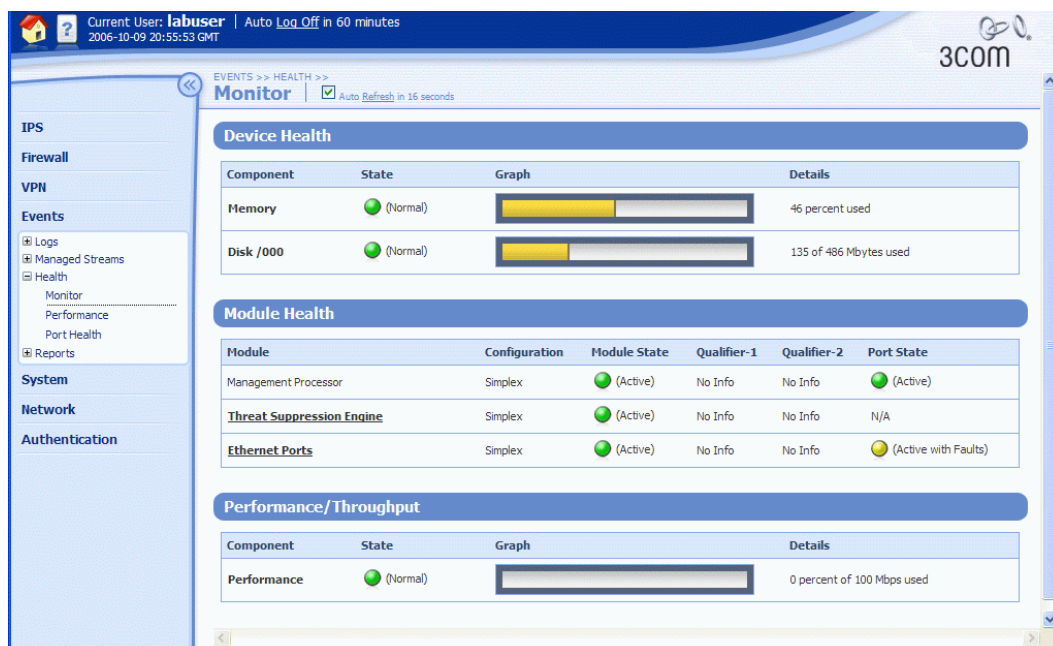
For details on viewing managed traffic streams, see the *LSM User's Guide* or the online help.

Health

The Health menu pages show the current status and network performance of the X family device. From the Monitor page you can review:

- Device health indicated by memory and disk usage statistics
- Module health including the Threat Suppression Engine and Ethernet ports
- Performance/Throughput

Figure 11–1: Monitor Page



For details on the Health menu pages, see the *LSM User's Guide* or the online help.

Reports

X family device **Reports** provides access to detailed information about the LSM system alert and traffic activity. Data for each report is gathered real-time. You can use the **Refresh** option on each report page to get the most current report information.

The following Reports menu options are available:

- **Attacks** — displays data on traffic that has been filtered by the device based on the IPS filter and firewall rule configuration in a security profile.
- **Rate Limits** — Displays a bar graph showing the percentage of rate limit bandwidth used for each action set configured with a rate limit.
- **Traffic** — Displays traffic flow data categorized by transmission type, protocol, frame size, and port.
- **Traffic Thresholds** — Displays a bar graph of traffic that has triggered a traffic threshold filter. The report graphs the amount of incoming traffic as a function of time.
- **Quarantine** — Displays a bar graph showing quarantine activity as a function of time.
- **Adaptive Filter Events** — Displays the global Adaptive Filter settings and a list of the 10 most recent filters impacted by adaptive filtering. You can also edit the Adaptive Filter settings from this report page.
- **Firewall** — Displays a bar graph showing the hit counts for each firewall rule as a percentage of total traffic based on firewall sessions.

For additional information, see the following:

- [“Viewing a Report” on page 116](#)
- [“Attack Reports” on page 117](#)
- [“Rate Limit Reports” on page 118](#)
- [“Traffic Reports” on page 118](#)
- [“Traffic Threshold Report” on page 119](#)
- [“Quarantine Report” on page 119](#)
- [“Configure Adaptive Filter Events Report” on page 119](#)
- [“Firewall Reports” on page 120](#)

Viewing a Report

- STEP 1** From the **Events > Reports** menu, select the desired Report menu option.
- STEP 2** On the selected reports page, click any available view options to update the report data.
- STEP 3** To update the report data, use the **Refresh** option. On some reports, an **Animate Charts** option is available to update the data in real time.

Attack Reports

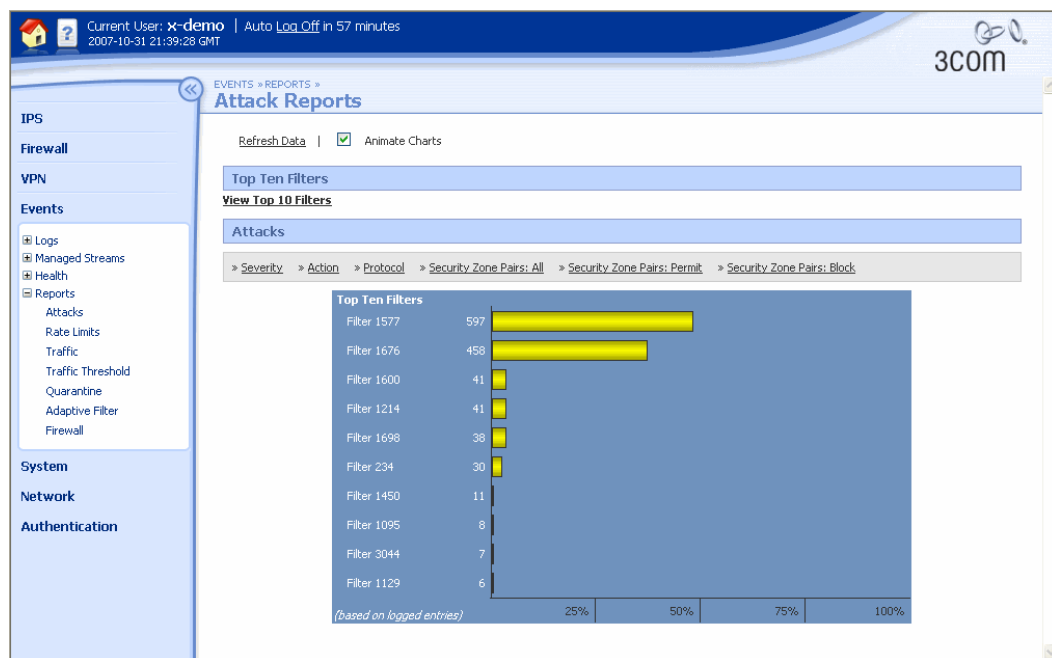
The **Attack Reports** allow you to view data on traffic that has been filtered by the device based on the IPS filter and firewall rule configuration. Firewall rules display as filter IDs in the 7400 to 7410 range. For example, filter ID 7400 is the default DENY ANY ANY rule implicitly added to the end of the Firewall Rule table.

Traffic data is reported based on the view options you select:

- **Top Ten Filters** — displays a bar graph of the top 10 attack filters by packet count, and the percentage of total traffic affected by the filter.
- **Severity** — displays the number of attacks categorized as Low, Minor, Major, and Critical. The graph also shows the percentage of total traffic for each severity level. The severity levels are assigned by the TippingPoint Digital Vaccine team and are included as part of the filter definition.
- **Action** — displays the actions taken on filtered traffic: traffic can be dropped (Invalid), blocked, or permitted. The report includes the number of packets processed by each action and the percentage of total traffic the number represents.
- **Protocol** — displays attack traffic categorized by protocol. The report includes the number of filtered packets for each protocol and the percentage of total traffic the number represents. Protocols include: ICMP, UDP, TCP, AND IP-Other.
- **By Port: All** — displays amount of all attack traffic reported by the security zone where the traffic was filtered, number of packets is reported as a percentage of total traffic.
- **By Port: Permit** — displays amount of attack traffic permitted reported by security zone. Number of packets is reported as a percentage of total traffic.
- **By Port: Block** — displays amount of attack traffic blocked reported by security zone. Number of packets is reported as a percentage of total traffic.

A sample Attack Reports page is shown in [Figure 11-2](#):

Figure 11-2: Attack Reports Page



Rate Limit Reports

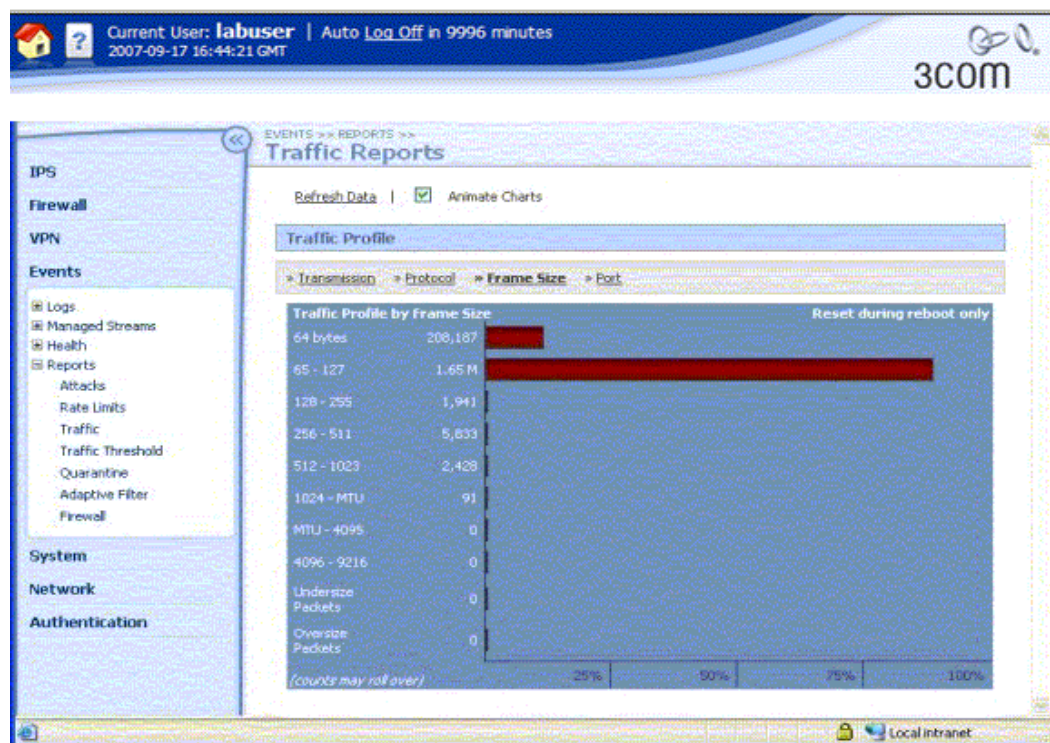
In the LSM, you can configure a rate limit action set to define the maximum amount of bandwidth available for traffic matching IPS filters that have a rate limit action set assigned. If two or more IPS filters use the same rate limit action set, then all packets matching these filters share the bandwidth. For each rate limit action set, the Rate Limit Reports page allows you to view the percentage of bandwidth consumed by rate-limited traffic graphed as a function of time.

For additional information on rate limit action sets and traffic streams, see [“Action Sets” on page 88](#). For details on rate-limited traffic streams, see the *LSM User’s Guide* or the online help.

Traffic Reports

The traffic report provides profile data on the packets flowing through the device (permitted packets only). A sample Traffic Profile Report page is shown in [Figure 11–3](#):

Figure 11–3: Traffic Profile Report — by Protocol



Traffic data is reported based on the view option you select on the Traffic Reports page:

- **Transmission Types** — graphs the number of packets transmitted for each of the following transmission categories: Unicast, Broadcast, MultiCast, MAC control, FCS Errors, Align Errors
- **Protocol** — graphs the number of packets transmitted by ICMP, UDP, TCP, IP-other, ARP, and Ethernet-Other
- **Frame size** — Traffic profile by framesize, by specified byte ranges
- **By Port** — Traffic profile by port, includes all security zones/ports

Updating Report Data

To update the traffic statistics in real time information, select the **Animate Charts** option. If this option is not selected, click the **Refresh Data** link to view the most current information.

Traffic Threshold Report

In the LSM, traffic threshold filters track statistical changes in network traffic patterns. You can specify the amount of traffic that triggers a Traffic Threshold filter from the Traffic Threshold Reports page. The units used in the report (packets/hour, bytes/minute, connections/second, etcetera) is determined by the units configured in the Traffic Threshold filter.



Note: The Traffic Threshold report is only available if an IPS Traffic Threshold filter has been configured for the device.

Traffic data is reported based on the viewing options you select on the Traffic Threshold Reports page:

- **Traffic Threshold filter name** — in the dropdown list under the Traffic Thresholds table heading, select the filter name to generate the traffic data for that filter.
- **Reporting time interval** — select the time interval for the reporting period: **Last 35 Days**, **Last 24 hours**, **Last 60 Minutes**, **Last 60 seconds**.

For additional information on Traffic Threshold filters, see the *LSM User's Guide* or the online help.

Quarantine Report

In the LSM, you can configure a filter with a quarantine action set. When a host computer triggers the filter, the host is quarantined according to the settings configured in the action set. You can monitor quarantine activity from the Quarantine Reports page.

Quarantine data is reported based on the viewing options you select on the Quarantine Reports page:

- **Total Hosts** — displays the total number of quarantined hosts as a function of time.
- **Packets Blocked** — displays the total number of packets blocked as a function of time.
- **Src Pages** — displays the number of LSM quarantine pages served to quarantined hosts as a function of time. The quarantine source pages are generated based on the configuration specified in the Quarantine action set.
- **Redirect Pages** — displays the number of times hosts have been redirected as a result of a quarantine action as a function of time.
- **Reporting time interval** — select the time interval for the reporting period: **Days** (last 35), **Hours** (last 24), **Minutes** (last 60), **Seconds** (last 60).



Note: Detailed information on quarantined hosts is available from the **Managed Streams > Quarantined Addresses** menu option in the LSM. For details, see the *Online Help*.

Configure Adaptive Filter Events Report

From the Configure Adaptive Filter Events Report page, you can:

- Review and modify the global Adaptive Filter configuration
- View a list of the 10 most recent filters managed by adaptive filtering
- Disable adaptive filter settings for an individual filter

For additional information on the Adaptive Filters, see [“Adaptive Filter Configuration” on page 91](#).


Firewall Reports

The Firewall Reports page provides links to the following graphs:

- Top Web sites — The most visited external Web sites by bandwidth
- Top clients — The clients on the internal virtual interfaces that are generating the most traffic by bandwidth
- Top services — The services that are consuming the most bandwidth
- Firewall rule hits — The firewall rules that are hitting the most traffic

A sample Firewall Reports page is shown in [Figure 11-4](#):

Figure 11-4: Firewall Reports Page



12 Deployment Scenarios

The X family device is multi-functional and can be deployed in a variety of scenarios. This chapter provides examples of common deployment scenarios. It covers the following topics:

- [Introduction](#)
- [Examples of VPN Network Topologies](#)
- [Deployment Scenario Example](#)
- [Multicast NBX Conference Calling](#)
- [Using Certificates to Ensure Security over a Public Network](#)
- [User Access Controlled by Firewall Policies](#)

Introduction

Some of the key requirements of organizations deploying a solution that includes firewall security and traffic prioritization, include:

- Partitioned approach to security — protection of network data and users through sectioning of the network and users.
- Security applications that integrate into the existing infrastructure.
- Ability to combine security services, such as authentication and authorization, with security technologies, such as VLANs and firewalls.
- Protection of the private network from external users and control over access to the Internet.
- Ability to control different type of network traffic.
- 24/7 availability and access to network resources and information.

The X family device can help meet these requirements, by providing the following applications:

- Protection, based on firewall rules and security zones for systems, applications and users within a dispersed network.
- Centralized control and configuration of your network segments, users and systems.
- Seamless connection between remote systems and users over the Internet, through VPN.
- Network traffic policing and prioritization.
- Firewall protection at the network perimeter.

Examples of common deployment requirements and the corresponding features on the device that are used to implement these requirements are shown in the following figure.

Table 12–1: Security Applications

Security Application	Function	X Family Implementation
Application Security	Prevents unauthorized access to specific applications.	Firewall rules; user authentication
Authentication	Verifies a user’s identity and ensures that the user is permitted to access network services.	User and firewall authentication
Authorization, Access Rights	Controls who has access to services, systems and subnets, and what information they can access.	User authentication; privilege groups
Intrusion Prevention	Prevents security threats from the Internet from causing damage to your network and compromising the security of your data.	IPS filtering
Availability, Business Continuity	Guarantees timely, reliable access to network resources for authorized users and high priority traffic, and enforces perimeter firewall protection as a first line of defense.	Firewall rules
Confidentiality, Privacy	Prevents disclosure of sensitive information to unauthorized persons or devices. Includes ability to encrypt messages across an unsecured network.	VPN; security zones; certificates
Content Management	Ensures that data is efficiently delivered, processed and stored. Provides internet filtering.	Web content filtering; firewall rules; Anti-Spam filtering
Integrity, Non-repudiation	Integrity ensures that messages are not altered during transmission.	VPN encryption; certificates
Security management	Provides secure network management and firewall-based networking. Simplifies IT management and administrative complexities of configuring and deploying secure networks.	Security zones; firewall rules

The next sections provide examples of common deployment scenarios that implement the above security applications.

Examples of VPN Network Topologies

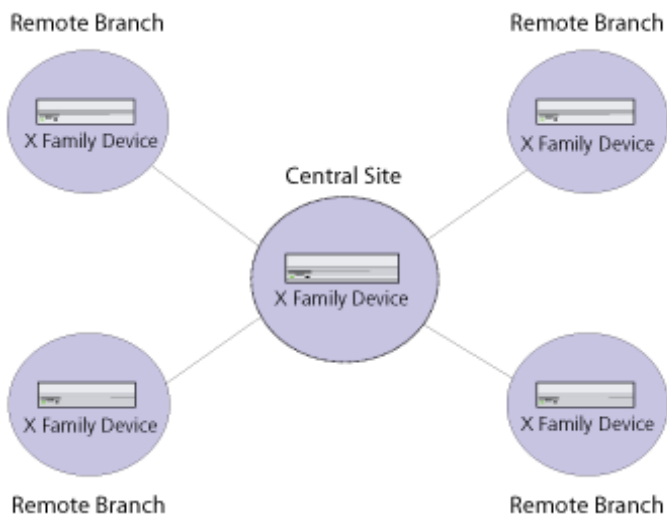
This section provides some examples of X family device VPN deployments in various types of network topologies, including:

- Hub and spoke deployment
- Meshed deployment
- Tree deployment

X Family Device in a Hub and Spoke Deployment

This example is typical of many corporate and public organizations in which a number of remote sites need to establish VPN connections with a central headquarters, to access servers and resources on the corporate network.

Figure 12–1: Hub and Spoke Deployment



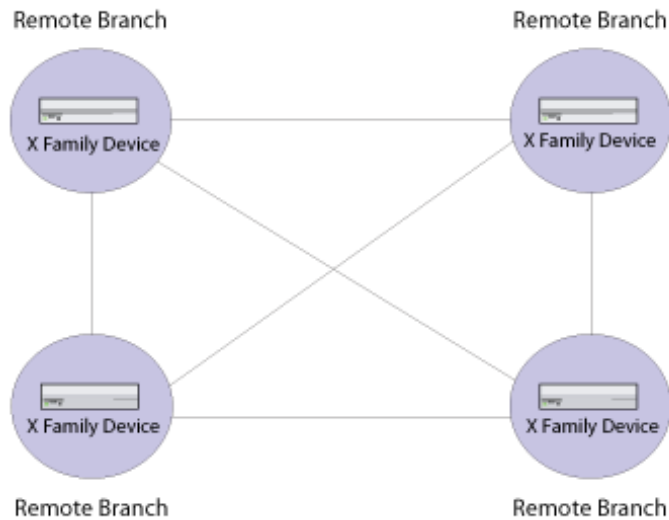
Since all sites are connected to a central location, this enables the central headquarters to maintain security and control the communication between the dispersed branches. The branches communicate with the central hub. Traffic can also be routed through the central hub to each of the branches.

A VPN security association is configured for each unique site-to-site connection. For deployments with a small number of remote sites, IKE with shared secrets can be used for authenticating the VPN connections. For larger deployments, involving many remote sites, IKE with certificates is recommended.

X Family Device in a Meshed Deployment

This example is typical of public organizations with no hierarchy between sites. Any site can connect directly to any other site to exchange information.

Figure 12–2: Meshed Deployment



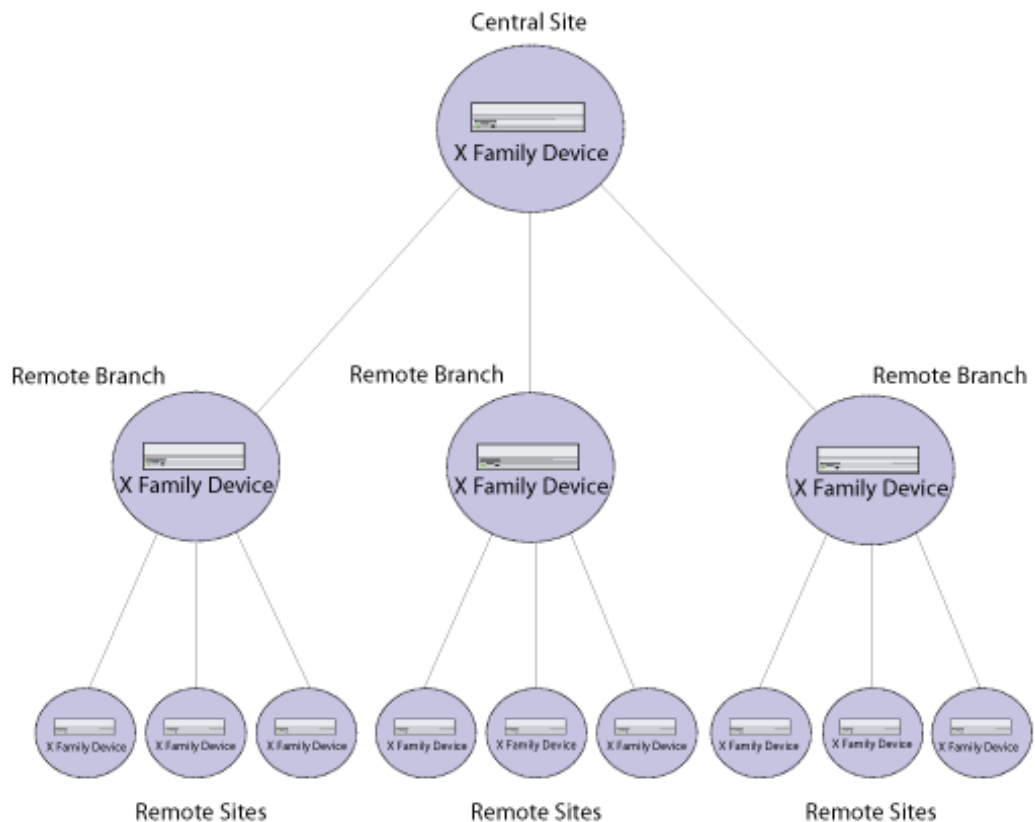
A meshed deployment does not offer a single, centralized point of control, since each branch has the same status as other branches. In this scenario, each branch is responsible for setting up the appropriate security zones and firewall policies for access to other branches.

A VPN security association is configured for each unique site-to-site connection. The use of IKE with shared secrets can be used for authenticating the VPN connections for a small number of remote sites. For larger numbers of remote sites, IKE with certificates is recommended.

X Family Device in a Tree Deployment

This example is typical of organizations in which a hierarchy exists. For example, a high street retail chain may have a central manufacturing site and several distribution centers across the country, which in turn connect to numerous sales outlets. The flow of information is from the central site out to the distribution centers and down through the chain and vice versa.

Figure 12–3: Tree Deployment



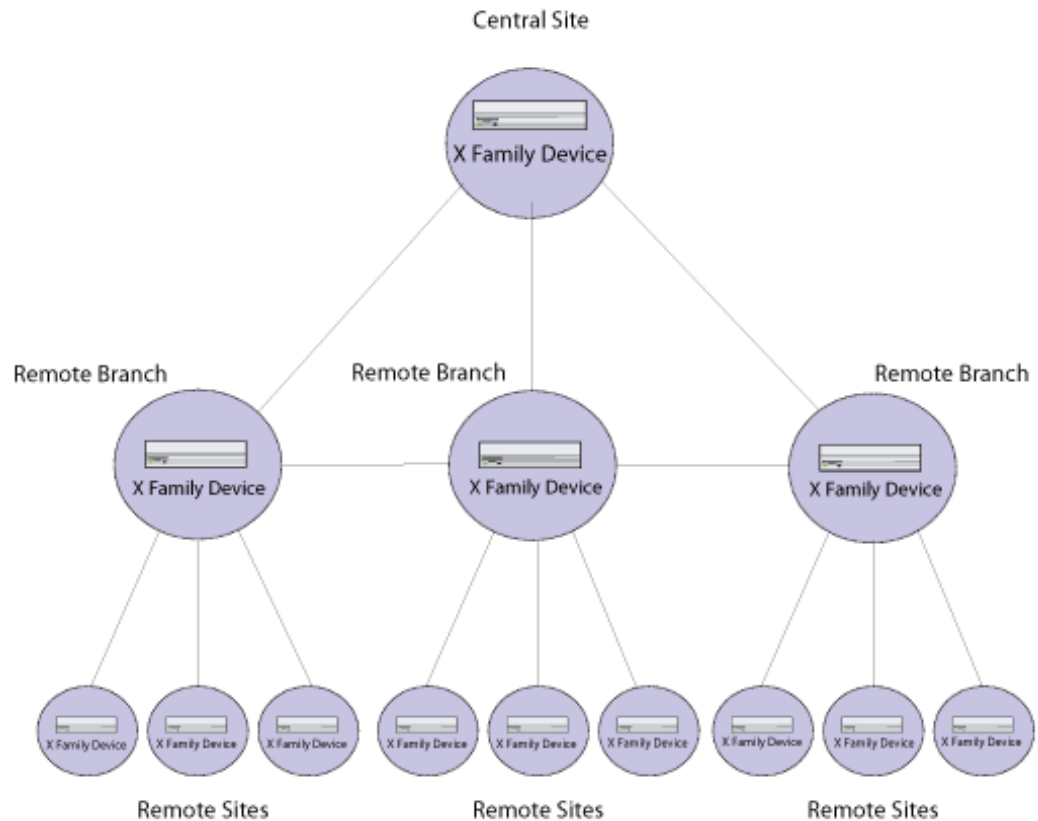
In a tree deployment, information flows down the hierarchy to the lowest nodes on the branch. Each level of the hierarchy is responsible for controlling the access of sites lower down in the hierarchy.

Sites that are highest in the hierarchy may have a unique site-to-site VPN connection with the central site, and a unique security association configured for each VPN connection. The use of X.509 certificates is the recommended option for authenticating the VPN connections, as this increases the security of the connection and add to the ease of configuring and maintaining the network.

X Family Device in a Mixed Deployment

Most typical deployments are a combination of tree, hub and spoke and meshed topologies.

Figure 12–4: Mixed Deployment



In this example, although there is still a hierarchy of sites, and a central site, as in a tree deployment, all the secondary level sites are interconnected.

The advantages of a mixed deployment are greater flexibility, to meet the needs of your organization. It provides improved redundancy and enables direct communication between the branches.

Each site is responsible for its own VPN security associations.

Deployment Scenario Example

The following scenario illustrates the implementation of security zones and VPN connectivity on a large, dispersed network.

Description

A successful New York high street institute has regional offices in several countries. Each region has hundreds of local branches and sales outlets. Each local branch and sales outlet reports to a regional branch headquarters.

The two main problem areas faced by the organization are communication and security requirements. As the company expands, their network faces an increasing demand for communication. Traffic flowing between the branches cannot be controlled or adequately supported with their current infrastructure. The company handles sensitive client data, so this needs to be kept secure on-site, as well as in any communications between the branches.

The nature of the business requires dealing with customers and quick communication between the branches. To support this, the business wants to run next generation applications, such as IP telephony, which provide a cost-effective, scalable solution to their expanding voice and video communication needs.

The branches are currently using a dialup connection with firewall software. The company decides that they are going to replace their current branch network infrastructure with X family devices, using VPN encryption for ensuring security and integrity of data passing over the open Internet, and firewall policies to manage traffic and prioritize voice and video traffic.

Application


Applications of the deployment solution being implemented by the company include:

- VPN — for secure communication between branches.
- GRE/multicast — to encapsulate and encrypt multicast traffic over IPSec.
- RIP — for dynamic routing.
- Voice, video, data— the network needs to support a variety of traffic, including IP telephony, conference calls, streaming video, and data streaming/data collaboration.
- Traffic shaping — to prioritize voice traffic and control other traffic passing over the network.
- Multiple security zones — these need to be set up to manage a diverse range of branch and home office staff, each with different levels of access to network resources.

Network Topology

The type of network deployment is a meshed topology between the regional branches and the main data centers. Lower down the chain, a hub and spoke topology is used, with sales outlets connecting to their local/regional branch, which in turn connect to the data centers ([Figure 12–5](#)).

Figure 12–5: Network Deployment Example



This topology enables each regional site to maintain centralized control over the local branches. Since each branch can only connect directly to its regional headquarters, this provides improved security and control. The meshed topology between the regional sites and data centers provides improved inter-site communication, as well as site redundancy.

The administrator deploys the following the X family devices at each site:

- One device at each of the central data centers, placed at the network perimeter and connected to the Internet via a WAN router.
- One device at each of the remote regional headquarters. Each device is placed at the network perimeter and connected to the internet via an ISDN link.
- One device at each local branch, placed at the network perimeter and connected to the internet via a dial-up modem connection.
- VPN client software installed on the workstations of each of the home office and travelling sales staff.

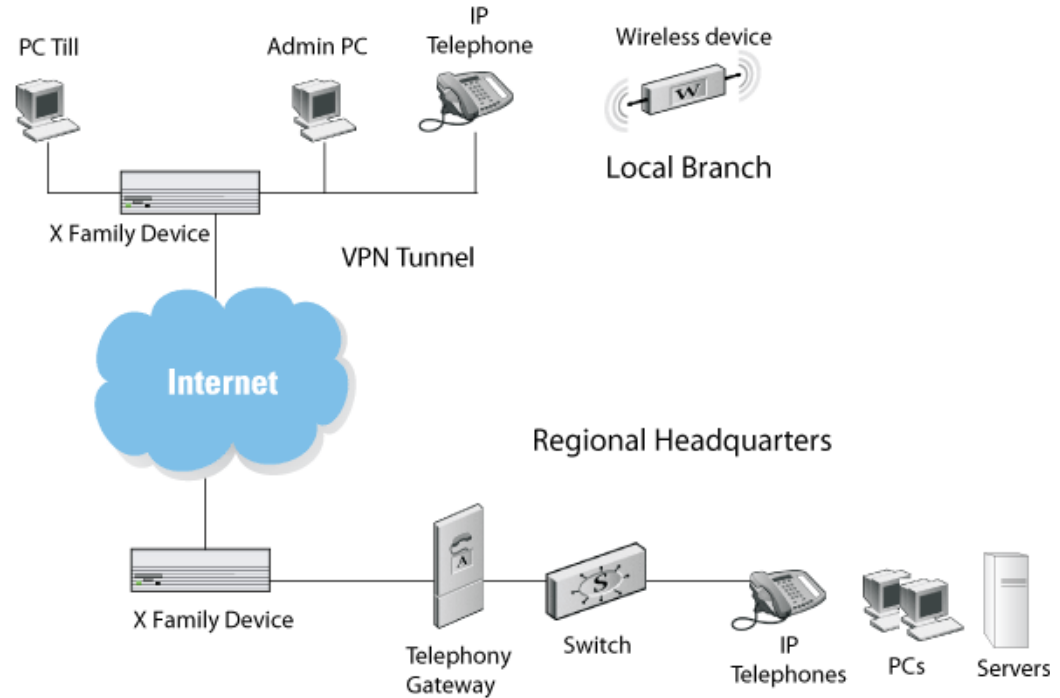
Benefits

The main benefits of deploying X family devices at each of the branches are:

- Provides flexibility on the WAN.
- Enables next-generation applications such as voice (IP telephony) and streaming video to remote branches.
- Implements security features, to conform to data protection legislation. All WAN traffic must be encrypted.

[Figure 12–6](#) provides a simplified description of the network deployment. Refer also to [Figure 12–7](#), which illustrates the corresponding security zones configured for this network.

Figure 12–6: Regional and Local Branch Setup



Local Branches

Each local branch has about a dozen employees. A store manager controls stock, employer and client records and order details. Sales staff manage the store's PC-based tills, which connect, over a VPN link, to the regional headquarters. Sales staff also have several hand-held wireless barcode reader devices, which are used in peak hours to scan goods and reduce queues.

Security Zone Configuration

The administrator configures security zones across the entire network:

- Local branch zones
- Home office zones for teleworkers and sales agents
- Regional headquarters zones
- Data center zones

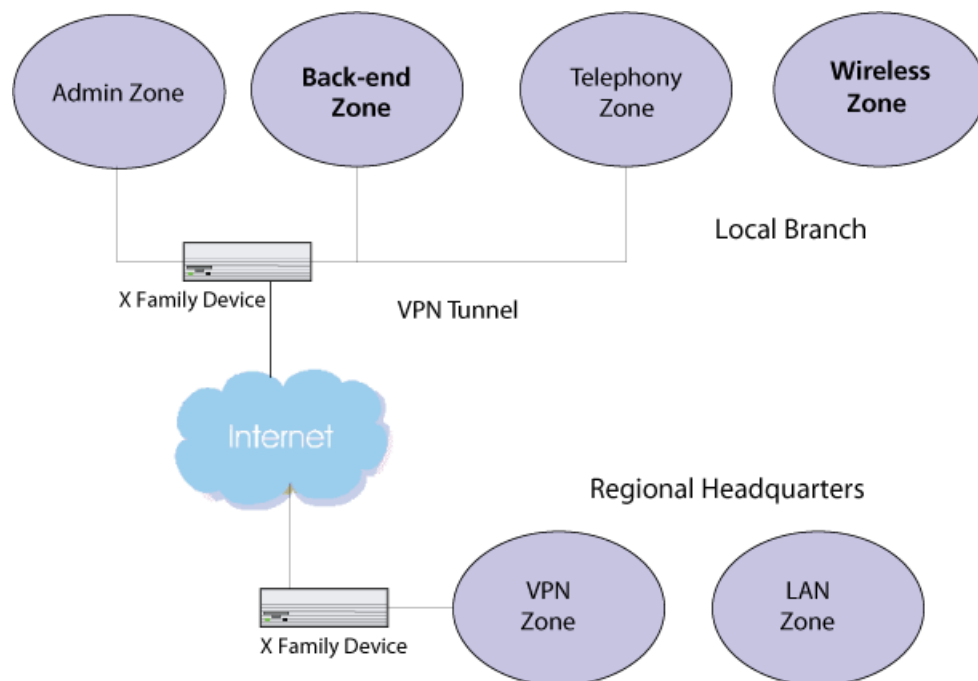
Local Branch Zones

The administrator creates the following zones on each of the local branches:

- Admin zone — contains the manager's PC.
- Wireless zone — contains the wireless barcode reader devices.
- Telephony zone — contains voice (IP telephony) and video devices.
- Back-end zone — contains the PC tills, used to process customer transactions.

[Figure 12-7](#) illustrates the security zones configured for the network. Refer also to [Figure 12-6](#), which provides a simplified description of the network deployment.

Figure 12-7: Security Zone Configuration



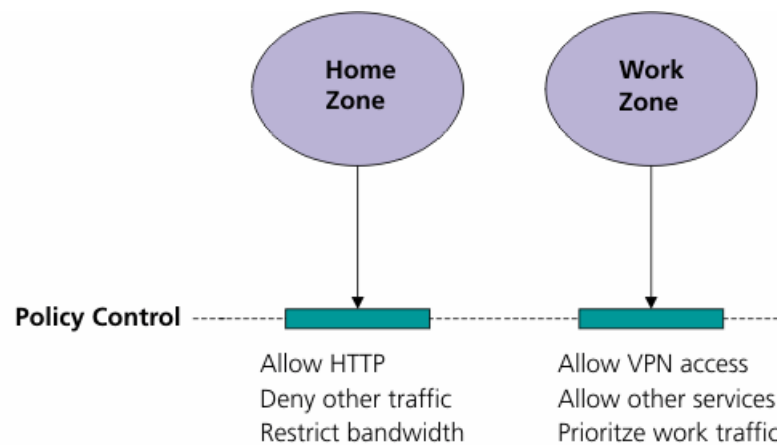
The back-end office till and administrator PC needs to connect, via a VPN link, to the regional branches/data centers, to transfer important transaction information and customer details across a secure link.

Home Office Zones

Teleworkers and sales agents, working from small home offices, have an X family device installed in the home office and can connect to their regional branch using a site-to-site VPN connection. Since the small home office of the employee may also contain devices that are used for personal use, not related to business, the administrator creates two zones:

- Home zone — allows HTTP browsing, but not allowing any other connection to the branch.
- Work zone — secure zone, supporting the VPN link to the branch and enabling all services applicable for this type of user.

Figure 12–8: Home and Work Zones



The administrator can increase security and prevent access to corporate servers from the home zone.

Bandwidth management of the home and work zones ensure that use of bandwidth is restricted for the home zone, while traffic from the work zone is prioritized.

Regional Headquarters Zones

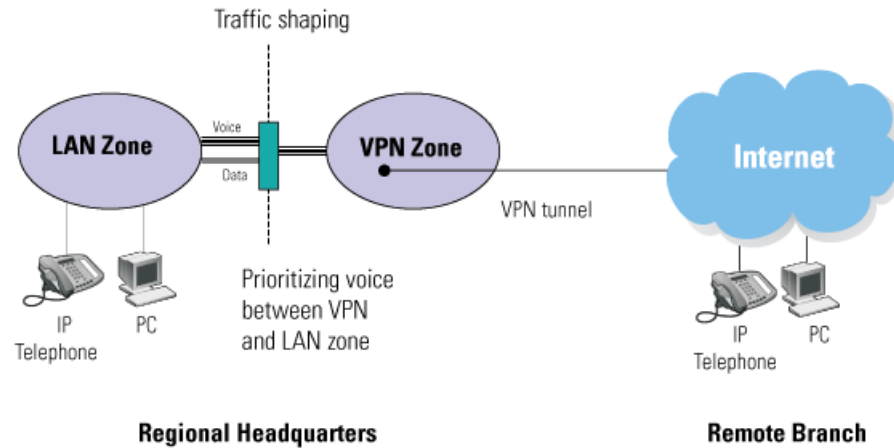
The administrator creates the following zones:

- LAN zone for all PCs and servers in the regional headquarters.
- WAN zone — for internet access devices (routers and switches).
- Virtual VPN tunnel zones — for tunneling VPN traffic.

Telephony and Virtual Zones

VPN traffic from the remote branches is set up with a target virtual voice VPN tunnel zone. This enables the administrator to set up firewall control for traffic passing across the VPN tunnel zone on to the LAN zone of the regional headquarters ([Figure 12-9](#)).

Figure 12-9: VPN Tunnel Zone Configuration



VPN Configuration

The network administrator sets up secure VPN connections between the headquarters and remote branches, to provide the remote branches with access to the servers on the corporate data centers.

A separate security association (SA) is set up for each remote site-to-site connection between branch headquarters and the data centers, using IPSec. VPN site-to-site connections are set up for all sales branches, for connecting to their regional headquarters.

To enable branch-to-branch communications in the mixed topology, the administrator configures routing.

Since all sales branches must go through their regional headquarters and cannot connect directly to the data centers, this provides an additional layer of security.

VPN Setup

The VPN is set up as follows:

- IPSec site-to-site security association for each VPN connection.
- Use of shared secrets for authenticating VPN access.

Creating Virtual VPN Tunnel Zones

In order to prioritize voice traffic, the administrator creates a Telephony zone in the branches, from which to tunnel the VPN traffic. 3Com NBX or VCX telephones are configured as part of this zone ([Figure 12-9](#)).

Firewall Rule Configuration

Firewall rules for traffic flowing between zones are set up as follows:

- Wireless to Back-end zone — Allow application data.
- Admin to Back-end zone — Allow HTTP, application data, SNMP.
- Admin to Remote Admin— Allow HTTP, FTP, POP3.

Multicast NBX Conference Calling

This scenario describes how IP telephony services are implemented in the example discussed in this chapter.

NBX Setup

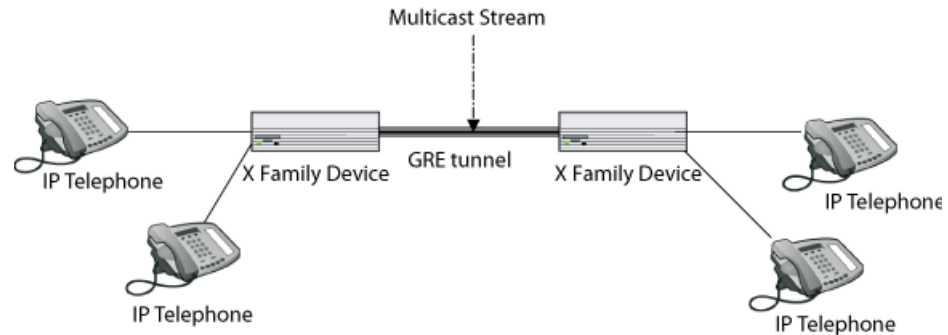
NBX telephones provide support for IP telephony on each of the branches. At the regional headquarters, NBX telephones are connected to the internet, via a Telephony gateway, for IP telephony sessions with remote branches ([Figure 12-6](#)).

To support voice across the VPN link, an NBX multicast service is required, in which packets are broadcast in real-time to all end stations involved in the communication. However, VPN is in effect a one-to-one, unicast connection. To enable multicast, the NBX packets need to be encapsulated within a GRE tunnel. The GRE packet is then broadcast to all other VPN endpoints to which an X family device is attached and opened at the remote end stations that it is intended for ([Figure 12-10](#)).

Traffic prioritization of NBX traffic is essential, to avoid problems of latency, which can effect the Quality of Service (QoS) levels over the network, during peak hour traffic. QoS is provided through the

creation of firewall policies that prioritize NBX traffic and guarantee a minimum bandwidth to this service.

Figure 12–10: Multicast Encapsulated within GRE



Multicast packets, enclosed within the GRE tunnel, are passed by the X family device across the Internet to remote sites, where the packets are distributed to the target IP telephones.

Network Setup

To support multicast, the administrator configures the network as follows:

- Selects GRE as the virtual interface for the NBX network segment; Enables RIP on this interface.
- Ensures that GRE is enabled for the site-to-site security associations configured for each VPN connection.
- Enables PIM-DM for dynamic multicast routing support between sites and IP subnets.
- Configures the X family device as a DHCP server: provides the NCP (network call processor) IP address, which avoids having to manually configure each NBX telephone.

Firewall Setup

To support NBX, the administrator configures firewall rules as follows:

- NBX traffic is assigned to the highest priority queue.
- Voice firewall rule: LAN to VPN; bandwidth control — guarantee 100Kbps, per session.
- WAN zone: 2Mbps — for all devices on zone.

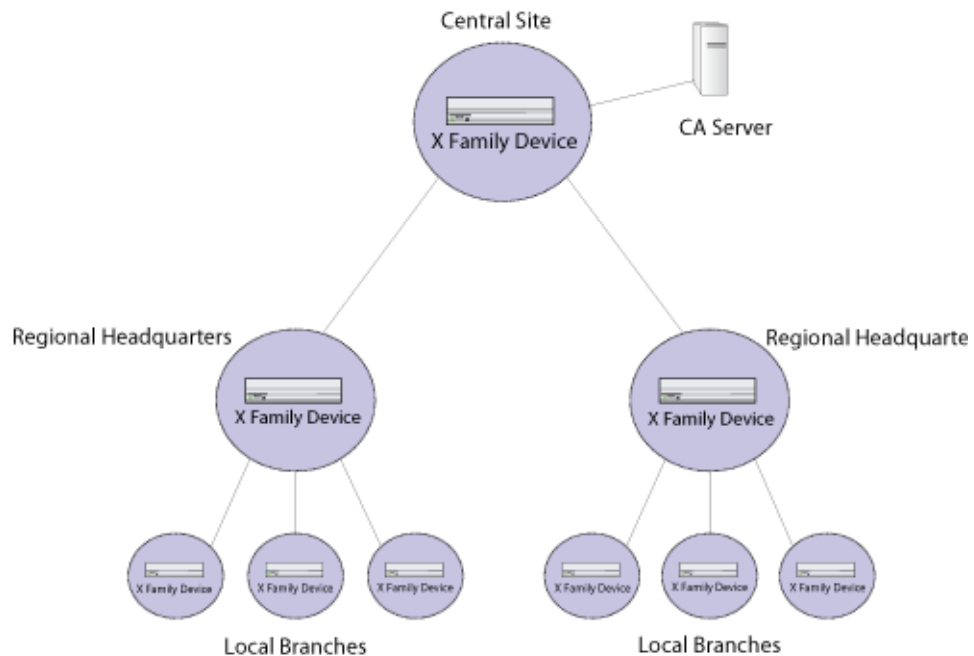
Using Certificates to Ensure Security over a Public Network

The following scenario illustrates the use of X.509 certificates in VPN, to connect remote sites over a public network, using the example discussed in this chapter.

Description

The network administrator sets up the network as described in the previous section. For additional authentication and non-repudiation, the administrator decides to use X.509 certificates as the means of authenticating the VPN.

Figure 12–11: Network Deployment Example Using Certificates



To implement X.509 certificates for VPN authentication, the administrator installs a CA server (Microsoft Certificate Server 2000) at the headquarters.

Certificates

To set up certificates, the administrator completes a procedure similar to the following:

- 1 The administrator submits a request to a third-party Certificate Authority, and is sent a signed CA certificate.
- 2 Using the certificate server, the administrator creates the company's own certificate, signed by the third-party CA.
- 3 The administrator exports the CA certificate from the CA server.

4 The administrator imports the certificate into each X family device.

On each device, the administrator performs the following:

- 1 Creates a certificate request and sends this to the CA (by email/FTP or other means).
- 2 Gets back the local certificate and installs it on the local machine.
- 3 Configures the CRL (Certificate Revocation List).

VPN

When configuring the VPNs for each remote site, the administrator:

- 1 Creates an IKE proposal and selects X.509 certificates as the method of authentication and chooses the Local Certificate installed on the X family device.
- 2 Creates an IPSec Security Association using IKE and selects the new IKE proposal.
- 3 Configures the Peer Distinguished Name in the SA to match the DN of the remote device's certificate.

User Access Controlled by Firewall Policies

The following scenario illustrates the implementation of firewall policies to control user access to network services, using the example discussed in this chapter.

Description

The administrator wants to provide differentiated access to the network's services, based on the user's identity.

The administrator installs a RADIUS server, which will be used to authenticate user names and passwords and assign privileges.

User Authentication for Local Branch and Headquarters

The administrator creates two privilege groups:

- *Staff* — for general staff, Web Filtering and Firewall Authentication privileges are enabled.
- *Admin* — for system administrators and IT support staff, all privileges are enabled.

User Authentication for Home Office Workers

The administrator creates two privilege groups:

- *Home* — Web Filtering and Firewall Authentication privileges are enabled.
- *Work* — VPN Client Access and Firewall Authentication privileges are enabled.

When configuring policies, the administrator ensures that the *Firewall Rule Authentication* option is enforced for all users who wish to access branch on-site services.

Firewall Configuration

Firewall rules are set up as follows for home office workers:

- *Home* — Allow HTTP; deny all other services; this restricts the user from accessing all other on-site staff services.
- *Work* — Allow SNMP, FTP, HTTP; enables staff to access on-site staff services.

Policies are set up as follows for on-site employees:

- *Staff* — Allow SNMP, FTP, HTTP; enables staff to access remote branches.
- *Admin* — Allow all services and access to all branches.

A Web Content Filter and Anti-Spam Services

Detailed information about the Web Content Filter Service and Anti-Spam subscription services used to control access to Web sites by categories and to control spam email. These services are offered in partnership, respectively, with SurfControl, a market-leading content filtering product, and Commtouch Software Ltd., dedicated to protecting the integrity of the world's most widespread form of communication, e-mail.

Overview of the Web Content Filter Service

The Web Content Filter Service is a subscription content filtering service that provides Web content filtering based on Web site category classifications. This service is operated in partnership with SurfControl, a provider of content filtering services.

On the X family device, all requests for Web sites within a particular category are allowed or blocked depending on how the category is configured for the Web Content Filter Service. You can configure Web Content Filter Service category settings from the Web Filtering Service page (**Firewall > Web Filtering**). For details, see the *LSM User's Guide* or the online help.

The Web Content Filter Service provides two main categories for filters:

- **Core Categories**

For details on what types of Web sites are included in each core category, see [“Core Categories” on page 140](#).

- **Productivity Categories**

For details on what types of Web sites are included in each productivity category, see [“Productivity Categories” on page 142](#).

In order to use the Web Content Filter Service, you need to purchase a license for the service. For details, see [“Purchasing a Web Filter License” on page 148](#).

Core Categories

Core Categories are used to classify Web sites that contain offensive, potentially dangerous, or criminal content. On the X family device, all Core Categories are blocked by default. For information on the type of Web sites included in each category, see the following topics:

- [“Adult/Sexually Explicit” on page 140](#)
- [“Criminal Skills” on page 141](#)
- [“Drugs, Alcohol & Tobacco” on page 141](#)
- [“Gambling” on page 141](#)
- [“Hacking” on page 141](#)
- [“Hate Speech” on page 141](#)
- [“Violence” on page 142](#)
- [“Weapons” on page 142](#)

Adult/Sexually Explicit

This Core category includes sites on the following topics:

- Sexually-oriented or erotic full or partial nudity depictions or images of sexual acts, including animals or other inanimate objects used in a sexual manner.
- Erotic stories and textual descriptions of sexual acts.
- Sexually exploitative, or sexually violent text or graphics.
- Bondage, fetishes, and genital piercing.
- Adult products including sex toys, CD-ROMs, and videos.
- Adult services including video conferencing, escort services, and strip clubs.

Sexual health, breast cancer, or sexually transmitted diseases (except in graphic examples) are not considered sexually explicit.

Criminal Skills

This Core category includes sites on the following topics:

- Advocating, instructing, or giving advice on performing illegal acts
- Tips on evading law enforcement
- Lock-picking and burglary techniques

Drugs, Alcohol & Tobacco

This Core category includes sites on the following topics:

- Recipes, instructions, or kits for manufacturing or growing illicit substances including alcohol. These include purposes other than industrial usage sites that glamorize, encourage, or instruct on the use of or masking the use of alcohol, tobacco, illegal drugs, or other substances that are illegal to minors.
- Alcohol and tobacco manufacturers' commercial Web sites.
- Sites detailing how to achieve “legal highs,” glue sniffing, misuse of prescription drugs, or abuse of other legal substances.
- Sites that make available alcohol, illegal drugs, or tobacco free or for a charge displaying, selling, or detailing use of drug paraphernalia.

Web sites discussing medicinal drug use, industrial hemp use, or public debate on the issue of legalizing certain drugs are not included in this Category set; nor are Web sites that are sponsored by a public or private agency that provides educational information on drug use.

Gambling

This Core category includes sites on the following topics:

- Online gambling or lottery sites that invite the use of real money.
This also includes Web sites that provide phone numbers, online contacts, or advice for placing wagers, participating in lotteries, or gambling real money, newsgroups or sites discussing number running, virtual casinos and offshore gambling ventures, sports picks, and betting pools.

Hacking

This Core category includes sites on the following topics:

- Promotion, instruction, or advice on the questionable or illegal use of equipment and/or software for purpose of hacking passwords, creating viruses, gaining access to other computers, and/or computerized communication systems
- Sites that provide instruction or work-arounds for filtering software
- Cracked software and information sites; Warez
- Pirated software and multimedia download sites
- Computer crime

Hate Speech

This core category includes sites on the following topics:

- Web sites advocating or inciting degradation or attack of specified populations or institutions based on associations such as religion, race, nationality, gender, age, disability, or sexual orientation; sites

that promote a political or social agenda which is supremacist in nature and exclusionary of others based on their race, religion, nationality, gender, age, disability, or sexual orientation.

- Holocaust revision/denial sites.
- Coercion or recruitment for membership in a gang or cult. A gang is defined as a group whose primary activities are the commission of felonious criminal acts, which has a common name or identifying sign or symbol, and whose members individually or collectively engage in criminal activity in the name of the group. A cult is defined as a group whose followers have been deceptively and manipulatively recruited and retained through undue influence such that followers' personalities and behavior are altered. Leadership is all-powerful, ideology is totalistic, and the will of the individual is subordinate to the group. A cult sets itself outside of society.

News, historical, or press incidents that may include the above criteria (except in graphic examples) are not blocked.

Violence

This Core Category includes sites on the following topics:

- Web sites portraying, describing, or advocating physical assault against humans, animals, or institutions.
- Depictions of torture, mutilation, gore, or horrific death.
- Web sites advocating suicide or self-mutilation.

News, historical, or press incidents that may include the above criteria (except in graphic examples) are not blocked.

Weapons

This Core Category includes sites on the following topics:

- Instructions, recipes, or kits for making bombs or other harmful or destructive devices.
- Web sites that primarily sell guns, weapons, ammunition, or poisonous substances.
- Web sites that allow online purchasing or ordering information, including lists of prices and dealer locations.

Productivity Categories

Productivity Categories are used to classify Web sites that could impair productivity when used in the work environment. On the X family device, all Productivity Categories are allowed by default.

Available Productivity Categories

This section provides a listing of the Productivity Categories available for the Web Content Filter Service. A description of the types of Web sites included is provided for each category. Use the cross-references in the following table to locate information on a specific category.

The following table provides a list of the available Productivity Categories.

Table A–1: Web Content Filtering Service: Available Productivity Categories

Advertisement (see page 143)	Arts & Entertainment (see page 144)	Chat (see page 144)	Computing & Internet (see page 144)
Education (see page 144)	Finance & Investment (see page 144)	Food & Drink (see page 144)	Games (see page 145)
Glamour & Intimate Apparel (see page 145)	Government & Politics (see page 145)	Health & Medicine (see page 145)	Hobbies & Recreation (see page 145)
Hosting Sites (see page 146)	Job Search & Careers (see page 146)	Sites for Children (see page 146)	Lifestyle & Culture (see page 146)
Motor Vehicles (see page 146)	News (see page 146)	Personals & Dating (see page 146)	Photo Searches (see page 146)
Real Estate (see page 147)	Reference (see page 146)	Religion (see page 147)	Remote Proxies (see page 147)
Sex Education (see page 147)	Search Engines (see page 147)	Shopping (see page 148)	Sports (see page 148)
Streaming Media (see page 148)	Travel (see page 147)	Usenet News (see page 148)	Web-based Email (see page 148)

Advertisement

- Banner Ad Servers
- Pop-Up advertisements
- Adware

Arts & Entertainment

- Museums, galleries, artist sites (sculpture, photography, etc.)
- Performing arts (theatre, vaudeville, opera, symphonies, etc.)
- Dance companies, studios, and training
- Book reviews and promotions, variety magazines, and poetry
- Television, movies, music, and video programming guides
- Online magazines and reviews on the entertainment industry
- Celebrity fan sites
- Broadcasting firms and technologies (satellite, cable, etc.)
- Horoscopes
- Jokes, comics, comic books, comedians, or any site designed to be funny or satirical
- Online greeting cards
- Amusement/theme parks

Chat

- Web-based chat
- Instant Message servers



Note This category filters HTTP traffic only.

Computing & Internet

- Reviews, information, buyer's guides of computers, computer parts and accessories, and software
- Computer/software/Internet companies, industry news, and magazines
- Pay-to-Surf sites

Education

- Educational institutions, including pre-, elementary, secondary, and high schools; universities
- Educational sites: pre-, elementary, secondary, and high schools; universities
- Distance education and trade schools, including online courses
- Online teacher resources (lesson plans, etc.)

Finance & Investment

- Web sites that provide stock quotes, stock tickers, and fund rates
- Web sites that allow stock or equity trading online
- Investing advice or contacts for trading securities
- Money management/investment services or firms

Food & Drink

- Recipes, cooking instruction and tips, food products, and wine advisors
- Restaurants, cafes, eateries, pubs, and bars
- Food/drink magazines, reviews

Games

- Web sites that allow a user to download or play online games
- Tips and advice on playing computer and Internet-based games
- Journals and magazines dedicated to game playing
- Web sites hosting games and contests

Glamour & Intimate Apparel

- Lingerie, negligee, or swimwear modeling
- Supermodel fan pages
- Fashion, clothing, and glamour magazines or catalogues
- Beauty and cosmetics
- Fitness models and sports celebrities
- Modeling information and agencies

Government & Politics

- Local, state, federal, and international government sites
- Government services such as taxation, armed forces, customs bureaus, or emergency services
- Political parties
- Political debate, canvassing, election information, and results
- Conspiracy theorist & alternative government views that are not hate-based

Health & Medicine

- Prescription medicines
- Medical information and reference about ailments, conditions, and drugs
- General health such as fitness and well-being
- Medical procedures, including elective and cosmetic surgery
- Dentistry, optometry, and other medical-related sites
- General psychiatry and mental well-being sites
- Psychology, self-help books, and organizations
- Promoting self-healing of physical and mental abuses, ailments, and addictions
- Alternative and complementary therapies, including: yoga, chiropractic, and cranio-sacral
- Hospital, medical insurance

Hobbies & Recreation

- Recreational pastimes such as collecting, gardening, kit airplanes
- Outdoor recreational activities such as hiking, camping, rock climbing
- Tips or trends focused on a specific art, craft, or technique
- Online publications on a specific pastime or recreational activity
- Online clubs, associations, or forums dedicated to a hobby
- Traditional (board, card, etc.) games and their enthusiasts
- Animal/pet related sites, including breed-specific sites, training, shows, and humane societies

Hosting Sites

Web sites that host business and individuals' Web pages (i.e. GeoCities, earthlink.net, AOL)

Job Search & Careers

- Sites hosting job and resume listings
- Tips and strategies for job seekers and interviewees
- Online job finding services

Sites for Children

Child-oriented sites and sites published by children

Lifestyle & Culture

- Home life and family-related topics, including weddings, births and funerals
- Parenting tips and family planning
- Gay/lesbian/bisexual (non-pornographic) sites
- Foreign cultures, socio-cultural information
- Tattoo, piercing parlors (non-explicit)

Motor Vehicles

- Car reviews, vehicle purchasing, or sales tips and parts catalogues
- Auto trading, photos, discussion of vehicles, including motorcycles, boats, cars, trucks, and RVs
- Journals and magazines on vehicle modification, repair, or customization
- Online automotive enthusiast clubs

News

- Online newspapers
- Headline news sites
- News wire services
- Personalized news sources

Personals & Dating

- Web sites that provide singles listings
- Matchmaking and dating services
- Advice for dating or relationships
- Romance tips and suggestions

Photo Searches

- Sites that provide resources for photo and image searches
- Online photo albums/digital photo exchange
- Image hosting

Real Estate

- Home, apartment, and land listings
- Rental or relocation services
- Tips on buying or selling a home
- Mortgage and home loan information
- Home improvement
- Real estate agents and agencies

Reference

- Personal, professional, or educational reference
- Online dictionaries, maps, and language translation sites
- Census, almanacs, and library catalogues
- Topic-specific search engines

Religion

- Churches, synagogues, and other houses of worship
- Any faith or religious beliefs, including non-traditional religions such as Wicca and witchcraft

Remote Proxies

- Remote proxies or anonymous surfing
- Search engine caches that circumvent filtering
- Web-based translation sites that circumvent filtering

Sex Education

- Pictures or text advocating the proper use of contraceptives
- Sites relating to discussion about the use of the Pill, IUDs, and other types of contraceptives
- Discussion sites on how to talk to your partner about diseases, pregnancy, and respecting boundaries



Note Not included in the category are commercial sites that sell sexual paraphernalia. These sites are typically found in the Adult category.

Search Engines

General search engines (Yahoo, AltaVista, Google)

Shopping

This Productivity category includes sites on the following topics:

- Internet malls and online auctions
- Department stores, retail stores, company catalogs online
- Online downloadable product warehouses; specialty items for sale
- Companies online dedicated to freebies or merchandise giveaways

Sports

This Productivity category includes sites on the following topics:

- Official team or conference Web sites
- National, international, college, and professional scores and schedules
- Virtual sports leagues and teams
- Sports-related online magazines or newsletters

Streaming Media

This Productivity category includes sites on the following topics:

- Streaming media files or events (any live or archived audio or video file)
- Internet TV and radio
- Personal (non-explicit) webcam sites
- Telephony sites that allow users to make calls via the Internet
- VoIP services

Travel

This Productivity category includes sites on the following topics:

- Airlines and online flight booking agencies
- Accommodation, information, and weather bureaus
- Leisure travel package listings
- Tourist information and maps

Usenet News

This blocks access to newsgroups accessed through the HTTP protocol.

Web-based Email

- Web-based email accounts
- Messaging sites (SMS, etc.)

Purchasing a Web Filter License

The Web Content Filter Service is a subscription-based service which requires the purchase of the correct license for your product from a reseller.



Note: You do not have to purchase a license to filter web sites using the Custom Filter List.

Each license allows one year of filtering for a specific X family product. Licenses cannot be transferred between base products, except through the standard Return Materials Authorization (RMA) process.

When you purchase a Web Content Filter Service subscription, you receive a License pack which includes a unique License Key. To enable the Web Content Filter Service for your product, register the License Key at <http://eSupport.3com.com>. You also need to provide the serial number of the specific X Family device for which you are enabling the service.



Note: The purchase of a Web Content Filter Service license does not extend any warranties or support contracts on the base product.

Free 14-day Trial Period

When you receive a new X family device, you can sign up for a 14-day trial period for the Web Content Filter Service. During the trial period, you do not need a license for the service. The trial period is activated when you register the device.

Overview of the Anti-Spam Service

The Anti-Spam Service is a subscription filtering service that provides email filtering based on pattern recognition. This service is operated in partnership with Commtouch Software Ltd.

[Purchasing an Anti-Spam License](#)

Purchasing an Anti-Spam License

The Anti-Spam Service is a subscription-based service which requires the purchase of the correct license for your product from a reseller.



Note: You do not have to purchase a license to filter spam using the manual spam filter list.

Each license allows one year of filtering for a specific X family product. Licenses cannot be transferred between base products, except through the standard Return Materials Authorization (RMA) process.

When you purchase an Anti-Spam Service subscription, you receive a License pack which includes a unique License Key. To enable the Anti-Spam Service for your product, register the License Key at <http://eSupport.3com.com>. You also need to provide the serial number of the specific X Family device for which you are enabling the service.



Note: The purchase of an Anti-Spam Service license does not extend any warranties or support contracts on the base product.

Free 14-day Trial Period

When you receive a new X family device, you can sign up for a 14-day trial period for the Anti-Spam Service. During the trial period, you do not need a license for the service. The trial period is activated when you register the device.

Glossary

action set

An integral part of an attack or peer-to-peer filter, action sets determine what the X family device does when a packet triggers a filter. An action set can contain more than one action, and can contain more than one type of action. The types of action that can be specified include the following:

- **Flow Control actions** — Determines where a packet is sent after it is inspected. *Permit* allows a packet to reach its intended destination. *Block* discards a packet. A block action can also be configured to *quarantine* the host and/or perform a *TCP reset*. *Rate limit* lets you define the maximum bandwidth available for the traffic stream.
- **Packet Trace action** — Captures all or part of a suspicious packet for analysis depending on how the packet trace options are configured.

The device comes with a set of default action sets that are applied to groups of filters based on a category setting recommended by the Threat Management Center. For details, see [“category settings” on page 152](#). The default action sets can be customized for individual filters or groups of filters. You can also create new action sets. For additional details, see [“Action Sets” on page 88](#).

Adaptive Filter Configuration

This function lets you configure the device to protect against potential adverse affects of a defective filter. When Adaptive Configuration is turned on and the network is experiencing heavy loads, the device automatically disables any filter that may be causing the congestion to prevent the device from entering High Availability mode and going offline. AFC settings are set to either Auto or Manual for the entire system. The default is Auto, which means that AFC is on. AFC can also be turned on or off for specific filters.

aggregation period

The length of time during which multiple instances of a specific attack can occur before notification is sent to a contact.

Application Protection

Category of filter types that defend against known and unknown exploits that target applications and operating systems of workstations and servers on a network. These filters include a variety of attack protection and security policy filters. These filters detect specific recognition data to recognize an attempted attack and take specific courses of action that you define when an attempt is detected.

attack filter package

See [“Digital Vaccine Package” on page 153](#).

attack traffic

Packets traversing a network that match at least one [Application Protection \(see page 152\)](#) filter.

Category

Digital Vaccine filters are organized into three main Categories based on the type of protection provided: [Application Protection \(see page 152\)](#), [Infrastructure Protection \(see page 154\)](#), and [Performance Protection \(see page 155\)](#). These categories are used to organize and locate filters in the LSM.

category settings

Category settings are used to assign global configuration settings to filters within a category. For example, a *Vulnerability* filter responds to attack traffic based on the category settings for the *Application Protection* category, while a *Network Equipment* filter would respond based on the category settings for the *Infrastructure Protection* category. You can edit individual filters within a sub-category to override the category settings for the filter. Category settings consist of the following global parameters:

- **State** — Determines whether filters within the sub-category are enabled or disabled. If a category is disabled, all filters in the Category are disabled.
- **Action Set** — Determines the action set that filters within a Category will execute when a filter match occurs. If the *Recommended* action set is configured, filters within the category are configured with the settings recommended by the Digital Vaccine team. If required, you can override the category setting on individual filters by editing the filter to define custom settings.

Classless Inter-Domain Routing (CIDR)

An address format similar to an IP address except that it is followed by a slash (/) and a specified number of bits. The number of bits indicates the significant bits in the address. In the following example, the IP source address of a packet must match all 32 bits of the IP address specified:

10.3.4.5/32

Custom Shield Writer (CSW)

An optional, stand-alone, TippingPoint application to write custom filters that can be imported for use on X family devices.

Digital Vaccine Filters

Digital Vaccine Filters block attacks and other malicious traffic from the network. Filters come with a set of recommended (default) settings that specify the filter status (enabled or disabled), the type of action to be taken when the filter is triggered (action set defined to permit or block traffic and/or send a notification), and the [Adaptive Filter Configuration \(see page 91\)](#) setting (on or off). You can accept the default settings or override them based on network security needs. Digital Vaccine filters are categorized in the following groups: [Application Protection \(see page 152\)](#), [Infrastructure Protection \(see page 154\)](#), and [Performance Protection \(see page 155\)](#).

Digital Vaccine filters are created by the Threat Management Center team, which monitors global network security threats and continually develops new attack filters that are automatically distributed to pre-emptively protect against the exploit of new and zero-day vulnerabilities. Updates are distributed using Digital Vaccine Packages.

Digital Vaccine Package

Downloadable software update that includes Digital Vaccine filters that provide the most current protection for your network. The Digital Vaccine Package is available from the [Threat Management Center \(TMC\) \(https://tmc.tippingpoint.com\)](#). Devices can also be configured to download and install the Digital Vaccine packages automatically.

DDoS filters

Group of infrastructure protection filters that detect denial of service attacks which flood a network with requests, including traditional SYN floods, DNS request floods against nameservers, and attempts to use protected systems as reflectors or amplifiers in attacks against third parties. These filters detect direct flood attacks and attacks hidden within larger packets and requests. DDoS filters include the following filters: SYN Proxy, Connection Flood, and CPS Flood filters.

Exploit filters

Filters that protect software from malicious attacks across a network by detecting and blocking the request. Exploits are attacks against a network using weaknesses in software such as operating systems and applications. These attacks usually take the form of intrusion attempts and attempts to destroy or capture data. These filters are part of the [Application Protection \(see page 152\)](#) filter category.

filter

Policy consisting of rules and conditions used to detect and manage malicious traffic on a network. Each filter includes an [action set](#) with instructions for managing data when the filter is triggered and [category settings](#). The LSM includes various types of filters, including Digital Vaccine filters in the [Performance Protection \(see page 155\)](#), [Application Protection \(see page 152\)](#), and [Infrastructure Protection \(see page 154\)](#) categories, along with traffic management, traffic threshold, and DDoS filters.

Identity Theft filters

Filters protect end users from phishing attacks by detecting and blocking connections to known phishing sites and attacks. A phishing attack is typically an email or Web site that has been spoofed to appear as if it is from a well-known financial or transaction institution. The attacks are usually geared to obtain account information from the end user. These filters are part of the [Application Protection \(see page 152\)](#) filter category.

IKE (Internet Key Exchange)

Internet Key Exchange (IKE) is used to negotiate the keying material that is used by the VPN encryption and integrity algorithms. IKE is a two-stage mechanism for automatically establishing IPSec tunnels with dynamically generated keying material. IKE uses UDP port number 500 and precedes the actual IPSec data flow.

IM filters

IM filters detect and control traffic from instant messaging applications such as Yahoo Messenger or MSN Messenger, chat, file transfer, and photo sharing. These filters can be used to block the operation of the instant messaging application. Many of the IM filters can also be used to rate-limit traffic from IM applications. These filters are part of the [Performance Protection \(see page 155\)](#) filter category.

Infrastructure Protection

Category of filter types that protect network bandwidth and network infrastructure elements such as routers and firewalls from attack using a combination of traffic normalization, DDoS protection, and application, protocol, and network equipment protection. These filters include DDoS, network equipment protection, and traffic normalization filters.

Intrusion Prevention System (IPS)

The TippingPoint Intrusion Prevention System in the X family device is an active network defense system that provides true intrusion prevention. Unlike intrusion detection systems, the IPS continually cleanses Internet and intranet traffic, identifying and preventing attacks before damage to critical resources occurs, ensuring network integrity and ultimately improving return on investment.

IP filter

A filter that blocks traffic based on the source, destination, port, protocol, and other parameters of the traffic.

IP interface

An IP interface is the Layer 3 configuration; that is, the IP configuration for its set of security zones (and hence Ethernet ports within the security zone). IP interfaces provide the X family device with the IP interfaces that it needs for the network connections you require.

IPSec

A protocol used to create secure VPNs by encrypting and authenticating all IP packets. It uses the IKE protocol for key exchange and authentication. IPSec provides security at the network layer.

L2TP

Layer 2 Tunneling protocol, a protocol for tunneling VPN (Virtual Private Network) traffic. L2TP is an extension to the Point-to-Point Tunneling Protocol (PPTP). L2TP supports multiple protocols and unregistered and privately administered IP addresses over the Internet. L2TP provides a more secure connection than the PPTP protocol.

Link State Advertisement (LSA)

A packet sent by an OSPF router describing routes within a given link.

Local Security Manager (LSM)

A browser-based management application that provides on-the-box administration, configuration, and reporting for a single X family device.

Network Equipment filters

Filters that detect and block the malicious attacks that target equipment accessible through a network. Network attacks can broadly or specifically seek access and data to corrupt on a network. These filters are part of the [Infrastructure Protection](#) filter category.

notification contacts

Recipients of alert messages. These contacts receive an email alert when a filter with the proper notification contacts settings triggers. Contacts include staff with email accounts and the SMS application.

OSPF (Open Shortest Path First, RFC 2328)

A routing protocol that determines the best path for routing IP traffic over a TCP/IP network. It calculate routes based on the number of routers, transmission speed, delays, and route cost. It is intended to replace the RIP protocol.

P2P filters

Filters that use the same algorithms as attack filters, but that block peer-to-peer protocol traffic. These protocols are primarily used to share music and video files. They essentially turn a personal computer into a file server which make its resources as well as those of its host network available to the peer-to-peer community. These filters are part of the [Performance Protection](#) filter category.

packet trace

Packet trace lets you capture all or part of a suspicious packet for analysis. You can set the packet trace priority and packet trace verbosity for action sets.

Performance Protection

Category of filter types that allow key applications to have prioritized access to bandwidth, ensuring that mission-critical applications perform adequately during times of high congestion. These filters include misuse and abuse, IP, and congestion/mitigation filters.

Port Scan/Host Sweep filters

Filters that perform port scans and host sweeps to prevent any malicious code, attacks, and exceeded threshold limits for traffic. Each filter scans a specific type of port and protocol to block attacks against ports and hosts. These filters are part of the [Application Protection](#) filter category.

PPTP (Point-to-Point Tunneling Protocol)

Point-to-Point Tunneling Protocol (PPTP) is a network protocol that enables the secure transfer of data from a remote client to a private enterprise server by creating a virtual private network (VPN) across TCP/IP-based data networks. PPTP supports on-demand, multi-protocol, virtual private networking over public networks, such as the Internet.

rate limiting

Setting in an action set that defines a maximum bandwidth that can be used by traffic that matches filters assigned to that action set. Incoming traffic in excess of this bandwidth is dropped. If two or more filters use the same rate limiting action set, then all packets matching these filters share the bandwidth.

Reconnaissance filters

Reconnaissance filters monitor traffic for events that indicate network activity usually associated with common information gathering techniques used by attackers to launch more sophisticated attacks. These attacks search through your network using various methods to locate vulnerabilities. After the attack has gathered data by probing your system and scanning your network, it continues with pointed attacks against those vulnerabilities. Reconnaissance filters look for these patterns and alert either the LSM or the SMS when an attack is detected. [Port Scan/Host Sweep filters \(see page 155\)](#) filters are included in this category. These filters are part of the [Application Protection](#) filter category.

RIP (Routing Information Protocol, RFC 2453)

RIP (Routing Information Protocol, RFC 2453) is a dynamic protocol that uses a distance vector algorithm to communicate route information with other routers in the network. RIP is well suited to small networks and uses the single metric *hop count* to determine distances. RIP sends route advertisements every 30 seconds using UDP broadcast or multicast packets. The best route to a destination will be the one that passes through the fewest number of routers (lowest hop count) to reach its destination. A destination with a metric of 16 hops or more is considered to be unreachable or of infinite distance.

Security Management System (SMS)

A Linux management server and Java-based client application for managing multiple X family devices. It provides coordination across your system for administration, configuration, monitoring, attack filter customization, centralized distribution of upgrades, and enterprise-wide reporting and trend analysis.

Security Profiles

A security profile is used to set up Digital Vaccine filters to monitor traffic passing on one or more virtual segments. The profile consists of category settings for the DV filters along with any user-defined filter overrides and IP address limits/exceptions. After a security profile is created, the device will begin monitoring traffic on the segments included in the profile using the specified filter settings.

Security Policy

Security Policy refers to all of the mechanisms available on the device to protect and manage network traffic including traffic management profiles, security profiles (Digital Vaccine Filters), DDoS, and Traffic Threshold filters. These profiles and filters are configured based on your network deployment and operational policy.

security zone

A security zone is a section of the network that is associated with a port or VLAN. Security zones let you logically segment your networks so that the X family device can apply policy rules and IPS filters to control the traffic passing between the zones.

SNMP Server

Provides access to interface counters and other statistics, configuration data, and general system information via the Simple Network Management Protocol (SNMP). The SNMP server must be enabled to use SMS management or to allow access by a third-party network management station.

spyware filters

Spyware filters detect and block downloads, communications, and popups sent via spyware.

Streaming Media filters

Streaming Media filters detect and control traffic from Streaming Media applications that deliver audio and video content utilizing IP protocols, typically UDP. Because these streaming media applications demand high bandwidth, the use of these applications can have a large negative impact on network performance. These filters can be used to block the operation of the instant messaging application. Many of the IM filters can also be used to rate-limit traffic from IM applications. These filters are part of the [Performance Protection \(see page 155\)](#) filter category.

Traffic Normalization filters

Filters that block network traffic when the traffic is considered malicious. These filters let you set alerts to trigger when the system recognizes this traffic. Traffic pattern filters alert when network traffic varies from normal. These filters are part of the [Infrastructure Protection](#) filter category.

Threat Management Center (TMC)

A 3Com service center that monitors sensors around the world for the latest attack information and builds and distributes attack filters. The TMC is available at the following URL: <https://tmc.tippingpoint.com>

Threat Suppression Engine (TSE)

A blend of Application-Specific Integrated Circuits (ASICs) and network processors that detect threats and anomalies in your network traffic. The TSE scans and reacts to malicious attacks or anomalous traffic based on the configuration of the IPS security profiles, traffic management, and traffic thresholds filters using the latest Digital Vaccine package updates.

Virus filters

Virus filters detect and block events triggered by viruses, worms, Trojans, backdoors, and other blended malware threats. These filters are part of the [Application Protection \(see page 152\)](#) filter category.

Vulnerabilities filters

Filters that detect any attempt to exploit a vulnerability in any application, operating system, or networked hardware device. These filters determine whether a vulnerability exists based on traffic requests and reaction by services. These filters are part of the [Application Protection \(see page 152\)](#) filter category.

Index

A

- action sets 151
 - category 152
 - flow control 88, 151
 - notification contacts 88
 - packet trace 88, 151
 - quarantine 89
 - rate-limiting 89
- adaptive filter events 91, 119
- address groups 32
- Advanced Encryption Standard (AES) 48
- aggregation period 151
- alternate VPN peer 45
- Anti-Spam Service 79, 150
 - manual filtering 80
 - purchasing a license 150
- Application Protection - Reconnaissance filters
 - filter tuning 86
 - port scans, host sweeps 86
- Attack Reports page 117
- authentication 94
 - and VPN client access 96
 - applications of 95
 - Authentication Header (AH) 46
 - deployment example 136
 - LDAP 1
 - local database 97
 - of local users 97
 - privileges 96
 - RADIUS server 24, 98
 - VPN 43, 95

B

- bandwidth
 - management 1, 17, 27
 - priority 66
 - rate limiting 27
- blacklist 79
- BOOTP 30
- bridge mode 3

C

- category, Digital Vaccine 152
- Certificate Revocation List (CRL) 105
- certificates. See X.509 certificates
- CHAP authentication protocol 51
- Classless Inter-Domain Routing (CIDR) 152
- client-to-site VPN 50
- command-line interface (CLI) 2

- configuring
 - bandwidth management 65
 - client-to-site user authentication 53
 - firewall rules 59
 - NAT 34
 - networks, deployment scenarios 121
 - rate limiting 27
 - virtual interfaces 29
- Content Portal Authority (CPA) server 73
- Custom Response page 73
- custom services 63
- Custom Shield Writer (CSW) 152
- customer support xiv

D

- Data Encryption Standard (DES) 48
- database, user authentication 97
- DDoS (distributed denial of service) attacks 154
- default
 - IKE proposal 49
 - privilege groups 97
 - route 37
 - security association 52
 - security profile 9
 - security zones 4, 27
 - Web content filtering rule 72
- deployment scenarios 121
- devices, restricting access to security zones 27
- DHCP
 - and NBX 30
 - and VCX 30
 - client 29
 - relay 29, 30
 - relay over VPN 29, 31, 49
 - server 29, 30
- Digital Vaccine (DV) 2
 - filters 81
- Digital Vaccine (DV) filters
 - overriding 86
- DMZ deployment 12
- DNS Lookup tool 39
- dynamic routing 19, 20, 37, 38
 - OSPF 20
 - RIP 20
- dynamic VPN peer 45

E

- email filtering. See Anti-Spam Service
- Encapsulation Security Payload (ESP) 46
 - standard 48
- encryption, VPN 44
- Ethernet port, and security zones 3
- events 111
 - adaptive filter 91, 119

- events, adaptive filter 91, 119
- external interface 10, 28

F

- failover 11
- filtering actions 72
- filtering, content. See content filtering
- filters 153
 - Application Protection - Reconnaissance filters
 - filter tuning 86
 - port scans, host sweeps 86
 - category 152
 - DDoS 154
 - Infrastructure Protection
 - traffic threshold filters 87
 - managing 84
 - misuse and abuse 155
 - notification contacts 89
 - rate-limiting 89
- Find Network Path tools 39
- firewall 1
 - actions 64
 - and applications 6
 - and security zones 64
 - policies 59
 - policies, example 136
 - rule enforcement 60
 - rule precedence 60
 - rules 59
 - zones and physical ports 26
- full routed/NAT deployment 11

G

- Generic Route Encapsulation (GRE) 10, 28
 - deployment example 133
 - interface 28
 - packet 28
- GRE. See Generic Route Encapsulation (GRE) 28
- guide
 - audience ix
 - conventions xi
 - caution xiii
 - note xiii
 - tip xiii
 - warning xiii
 - organization x
 - related documentation xiii
 - screen captures xii

H

- High Availability 2, 21
 - failover 21
 - polling 22
 - standby operation 22
- host sweeps filters 86
- hub and spoke deployment 123

I

- IGMP 39
- IGMP v2 19
- IKE (Internet Key Exchange) 154
- IKE. See Internet Key Exchange (IKE) 47
- Infrastructure Protection, traffic threshold filters 87
- interface
 - external 28
 - GRE 28
 - internal 28
 - virtual 10, 27
- internal
 - interface 28
 - servers and NAT 68
- Internet Key Exchange (IKE) 47
 - and DHCP relay 31
 - default proposal 49
 - proposals 49
- Intrusion Prevention System (IPS) 2, 81, 154
 - filtering 7
- IP
 - address 10, 27
 - address groups 32
 - addresses, static mapping 32
 - configuration 10, 27
 - filter 154
 - interfaces 10
 - configuring 27
 - external 10
 - internal 11
 - multicast routing 2
 - subnet 10
 - telephony 133
- IP security (IPSec). See IPSec
- IPSec 154
 - and Encapsulation Security Payload (ESP) 48
 - modes 44
 - protocols
 - security mechanisms 46
 - transport mode 45
 - tunnel mode 44, 51
 - tunnel setup 49

K

- keying modes 46

L

- L2TP/IPSec 51
- LAN interface 11
- LAN security zone 4

- LAN switch 7
- Layer 2 Tunneling Protocol (L2TP) 51
- Lightweight Directory Access Protocol (LDAP) server 24, 98
- load balancing 11
- Local Security Manager (LSM) 2, 155
 - Add page 49
 - Attack Reports page 117
 - Configure Adaptive Filter Events page 91
 - Configure Adaptive Filter Events Report page 119
 - Default Gateway page 37
 - DHCP Server page 30
 - Edit Action Sets page 114
 - Edit Firewall Rule page 114
 - Firewall Reports page 120
 - Health menu pages 115
 - IKE Proposals page 49
 - IP Address Groups page 32
 - IPS menu pages 82
 - IPS Preferences page 91
 - IPS Services page 91
 - IPSec Configuration page 114
 - Logs menu pages 112
 - Logs page xiii
 - Monitor page xiii, 115
 - Network page 25
 - Notification Contacts page 90, 112, 113
 - Quarantine Reports page 119
 - Quarantined Streams page 89
 - Rate Limit Reports page 118
 - Routing page 37
 - Static Reservations page 32
 - Syslog Servers page 113, 114
 - System Summary page xiv, 112
 - Traffic Profile Report page 118
 - Traffic Threshold Reports page 119
 - Virtual Servers page 34, 68
 - Web Filtering Service page 139
- local user database 97
- logs 111, 112
 - reports, rate limit 119
 - reset 112

M

- managing filters 84
- manual key 47
- meshed deployment 124
- Message Digest 5 (MD5) 49
- misuse and abuse, filtering out 155
- Monitor page 115
- MPPE encryption standard 52
- MS-CHAP authentication protocol 51
- MS-CHAPv2 authentication protocol 51
- multicast routing 19, 38, 133

N

- NAT 11, 33
 - and internal servers 68
 - and VPN tunnels
 - many-to-one 33
 - one-to-one 33
 - setting up 34
 - within VPN tunnel 35
- NBX 133
 - and DHCP 30
 - conference calling 133
 - multicast routing 38
 - setup example 133
- network
 - configuration and deployment 121
 - performance 8
 - tools 39
 - topologies 121
- Network Address Translation (NAT). See NAT
- network call processor (NCP) 30
- network management system (NMS) 2
- notification contacts 89, 155

O

- Open Shortest Path First (OSPF) 37
 - support 20
- OSPF. See Open Shortest Path First (OSPF)

P

- P2P (peer-to-peer) protocol, filtering 155
- packet
 - authentication 43
 - trace 155
- PAP authentication protocol 51
- Peer-to-Peer filter 155
- physical security zones 4
- Ping tool 39
- Point-to-Point Tunneling Protocol (PPTP)
 - with MPPE 52
- Point-to-Point Tunneling Protocol (PPTP) 52
- poison reverse 38
- policy
 - actions 64
 - authentication 96
 - bandwidth priority 17, 66
 - rank 60
 - schedules 63
 - service groups 63
 - services 63
 - source and destination addresses 63
 - user authentication 66
 - VPN 64
- Port Address Translation (PAT) 35
- ports
 - scan filters 86
 - tagged 17, 26
- PPTP. See Point-to-Point Tunneling Protocol (PPTP)
- predefined services 63
- prioritization of traffic. See rate limiting 2

- privilege groups 96
 - and Web content filtering 71
- privileges 96
- Protocol Independent Multicast — Dense Mode (PIM-DM) 19
- public key cryptography 102
- public-key infrastructure (PKI) 43

Q

- quarantine action set 89

R

- RADIUS server 98
- rate limiting 2, 89, 156
- reconnaissance filters
 - filter tuning 86
- related documentation xiii
- Remote Authentication Dial-in User Service (RADIUS) 98
- reports
 - attack 117
 - rate limit 119
 - top ten filters 117
- RFC 2453 156
- RIP. See Routing Information Protocol (RIP)
- routed deployment 11
- routing
 - default route 37
 - dynamic routing 20, 37
 - IP multicast 19
 - multicast 38, 133
 - poison reverse 38
 - split horizon 38
 - static routes 37
- Routing Information Protocol (RIP) 38
 - support 20
- rules, firewall 59

S

- Secure Hash Algorithm-1 (SHA1) 49
- security
 - applications 122
 - intranet 7
 - network perimeter 6
 - partitioning 6
 - profiles 9
 - profiles, default 9
- security association (SA) 46
 - default SA 52
- Security Management System (SMS) 2, 156
- Security Parameter Index (SPI) 46, 47

- security zones 1, 156
 - and applications 6
 - and Ethernet ports 3
 - and VLANs 3, 26
 - and VPN tunnels 3
 - concepts 3
 - configuration example 129
 - configuration of firewall policies 8
 - configuration of IPS filters 9
 - default 4, 27
 - implementation methods 11
 - LAN 4
 - physical 4
 - rate limiting 27
 - restricting devices 27
 - this-device 4
 - tunnel zones 42
 - virtual 4
 - VLANs 26
 - VPN 4
 - WAN 4
- segments 9
- self-signed certificates 107
- service groups 63
- services 63
- site-to-site VPN 44
- SMS. See Security Management System
- SNMP 157
- source and destination addresses 63
- spam filtering. See Anti-Spam Service
- split horizon 38
- static mapping 29, 32
- static routes 37
- static routing 37
- Strong Encryption (Triple DES or 3DES) 48
- Strong Encryption Service Pack 48
- switch replacement 7

T

- tagged ports 17, 26
- tech support xiv
- this-device security zone 4
- Threat Management Center (TMC) xiv, 2, 157
 - registration 85
- Threat Suppression Engine (TSE) 81, 157
- TMC 157
- Top Ten reports 117
- TOS
 - user database 24
 - version number xiv
- Traceroute tool 39
- traffic 116
 - shaping 8, 17
 - threshold filters 87
- Traffic Capture tool 39
- Traffic Threshold filters 81
- transparent deployment 11
- tree deployment 125
- troubleshooting 139
- tunnel security zones 42
- tunneling 43

U

- URL Permit and Block lists 73
- user authentication 1, 43, 93
 - and privileges 96
 - and VPN 96
 - example 98
 - local database 97
 - privilege groups 96
 - privileges 96
 - using LDAP 98
 - using RADIUS 98

V

- VCX 133
 - and DHCP 30
- virtual interface 10, 27, 154
 - configuring 29
- Virtual Private Network. See VPN
- virtual security zones 4
- virtual server 34
- virus filters 157
- VLANs 17, 26
 - and security zones 3
 - trunking 17
- Voice over IP (VoIP) 8
- VPN 41
 - advanced configuration 54
 - alternate peer 45
 - authentication 43, 95
 - benefits of 42
 - client access 96
 - client compatibility 53
 - client-to-site 42, 50
 - concurrent VPNs 52
 - connection types 42
 - definition 41
 - DHCP relay over 31, 49
 - dynamic peer 45
 - encryption 44
 - firewall rules 64
 - hub and spoke deployment 123
 - initiation 41
 - keying modes 46
 - L2TP/IPSec 51
 - meshed deployment 124
 - multiple VPNs 52
 - NAT 36, 42, 54
 - network topology examples 123
 - packet authentication 43
 - policy rules 64
 - security features 42
 - security zone 4
 - site-to-site 42, 44
 - supernets 42, 55
 - tree deployment 125
 - tunnel 3, 43
 - user authentication 43

W

- WAN security zone 4

- Web Content Filter Service 139
 - Core Categories 140
 - Productivity Categories 142
 - purchasing a license 148
- Web content filtering 1
 - actions 72
 - and Custom Response page 73
 - example 74
 - manual filtering 72
 - URL Permit and Block lists 73
- Web filter profiles 71
- whitelist 79

X

- X family device
 - and bandwidth management 65
 - and encryption 48
 - and NAT 33
 - and printers 32
 - and X.509 certificates 101
 - as central relay agent 31
 - as DHCP server 32
 - as remote relay agent 31
 - as router 37
 - deployment scenario 127
 - management access to 64
 - public IP address 35
 - serial number xiv
- X.509 certificates 43, 101
 - deployment example 135
 - example 135
 - supported CA servers 107

Z

- zones. See security zones