



R3.4 Beta Draft – Cisco Confidential Information



Cisco IOS XR Security Guide

Cisco IOS XR Software Release 3.4

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-??-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco IOS XR Security Guide

Copyright © 2006 Cisco Systems, Inc. All rights reserved.



Preface	ix
Changes to This Document	ix
Obtaining Documentation	x
Cisco.com	x
Documentation DVD	x
Ordering Documentation	x
Documentation Feedback	xi
Cisco Product Security Overview	xi
Reporting Security Problems in Cisco Products	xi
Obtaining Technical Assistance	xii
Cisco Technical Support Website	xii
Submitting a Service Request	xii
Definitions of Service Request Severity	xiii
Obtaining Additional Publications and Information	xiii
Introduction	15
Definition of Security	15
Document Modules	16
Fundamentals of Cisco IOS XR System Security	19
Hierarchy of the Layered Defense	19
Application-specific Integrated Circuits	20
Operating System and Infrastructure	20
High Availability	20
Security Applications	21
Management Access Security	21
Control Plane Protocols	21
Modularity and Compartmentalization in the Planes	22
Operating System Infrastructure and Security	23
uKernel	23
Scheduler	24
Memory Protection	24
CLI Availability	24

Restartable Processes	24
Resource Monitoring Through wdsysmon	24
Fault Manager	24
In-Service Software Upgrade	25
Fault Manager Usage	25
Memory Threshold Specifications	25
Keys and Key Chain Management	25
Node Access and Management Plane Security	27
Secure Shell and Other Access Protocols	27
Protecting Software Images	28
Authentication, Authorization, and Accounting	28
Task Groups and User Groups	28
Simple Network Management Protocol	29
Command-Line Interface	29
Data Path Security	31
Introduction to Data Plane Security Measures	31
Packet Walk	31
Packets Arriving at the Ingress	32
Transit Packets	32
For-us Packet Walk	32
Packet Paths	32
Summary of Data Plane Security Mechanisms	34
Access Control Lists	35
Cisco IOS XR Access Lists and Prefix Lists Feature Highlights	35
Purpose of IP Access Lists	36
Contents of an ACL	36
ACL Processes and Rules	36
Suggestions for Creating an ACL	37
Source and Destination Addresses	37
Wildcard Mask and Implicit Wildcard Mask	37
Transport Layer Information	38
Access Control Entry Sequence Numbering	38
Benefits	38
Sequence Numbering Characteristics	38

IP Access List Logging Messages	38
Access Lists with Fragment Control	39
Comments About Entries in Access Lists	40
Access Control List Counters	41
BGP Filtering Using Prefix Lists	41
Quality of Service	41
Benefits of Implementing Cisco IOS XR QoS Features	42
QoS Techniques in Cisco IOS XR Software	42
Packet Classification	43
Congestion Management	43
Congestion Avoidance	43
Traffic Policing	44
Traffic Shaping	44
Unicast Reverse Path Forwarding	44
Examples of Network Attacks	47
The Routing System	47
Types of Attacks Against a Routing System	48
Disrupting Peering	48
Falsifying Routing Information	48
Abusing Routing Stability Features to Reduce Network Availability	48
Forcing BGP Peer Damping by Injecting Flapping Routing Information	49
Forcing the Routing Protocol to Converge More Slowly by Injecting Flapping Routing Information	50
Attacking a Routing System	51
Port Flooding	52
Protocol Semantics Peering Attacks	52
Compromising a Legitimate Member of the Routing Domain	54
Masquerading as a Member of the Routing Domain	54
Attack Detection and Response	55
Always-Available CLI Commands	55
NetFlow	56
NetFlow Version 9 Template Format	58
Watchdog System Monitor	58
Fault Manager	59
System Event Detection	60

System Event Processing	61
Alarms	61
Alarm Logging and Debugging Event Management System	62
Correlator	62
System Logging Process	63
Alarm Logger	63
Logging Correlation	63
Correlation Rules	63
Root Message and Correlated Messages	64
Alarm Severity Level and Filtering	64
Implementing Network Security Mechanisms	65
The Six Phases of a Security Strategy	66
Phase 1: Preparation	66
Phase 2: Identification	66
Phase 3: Classification	66
Phase 4: Trace-back	66
Phase 5: Reaction	66
Phase 6: Post-Mortem	67
Default Settings	67
Password Encryption Service	68
Time Stamp on Debug	68
Disable IP Gratuitous ARPs	68
TCP Synwait Time	68
Disable IP Redirects	69
Disable IP Proxy ARP	69
Disable IP Unreachables (Host)	69
Disable IP Mask Reply	69
Disable IP Unreachables on Null Interface	70
Enable uRPF on Outside Interfaces	70
SYN Cache Default Value	70
SYN Rate Limit Default	70
General Recommendations	70
Password Management	71
Controlling Interactive Access	72
Console Ports	72
General Interactive Access	72

Controlling TTYs	73
Controlling VTYS and Ensuring VTY Availability	73
Warning Banners	73
Management Services Without Interactive Login	74
SNMP	74
Comparison of SNMPv1, SNMPv2c, and SNMPv3	75
Security Models and Levels for SNMPv1, SNMPv2, and SNMPv3	76
SNMPv3 Benefits	77
SNMPv3 Costs	77
User-based Security Model	78
Management and Interactive Access over the Internet and Other Untrusted Networks	78
HTTP and HTTPS	78
Defeating a Packet Sniffer	79
Other Internet Access Dangers	79
Logging	80
Saving Log Information	80
Recording Access List Violations	81
Securing IP Routing	81
Antispoofing	81
Antispoofing with ACLs	81
Antispoofing with Unicast Reverse Path Forwarding	82
Controlling Directed Broadcasts	82
Path Integrity	82
IP Source Routing	83
ICMP Redirects	83
Routing Protocol Filtering and Authentication	83
Flood Management	83
Transit Floods	84
Switching Modes and Cisco Express Forwarding	84
Cisco Discovery Protocol	84
Protecting the Legitimacy of the Routing Domain	85
Protecting Routers from Being Compromised	85
Protecting Routing Information on the Wire	85
Cisco Discovery Protocol	85
Protecting the Legitimacy of the Routing Domain	85
Protecting Routers from Being Compromised	85

Protecting Routing Information on the Wire 85

Protecting Against Illegitimate Devices Joining the Routing Domain 87

MD5 and Peer Authentication 88

TTL-Based Peering Session Protection (BGP TTL Security Hack) 89

Protecting the Routing Information 89

- Extranet Connections 90
- Using an Exterior Gateway Protocol for All Extranet Connections 91
- Filtering Routes Aggressively at the Extranet Edge 91
- Dampening Prefixes Aggressively at the Extranet Edge 92
- Limiting Route Count at the Extranet Edge 92
- Sample Extranet BGP Configuration 93
- Connections to the Internet 93
- Route Filtering 93
- Protecting Against Becoming a Transit 94
- Route Dampening 95
- Connections Within a Network 95
- Route Filtering Within a Network 95
- Traffic Segregation and Routing Protocol Traffic Filtering 96

INDEX

Preface

This guide is an overview of the security provisions on the Cisco XR 12000 Series Router and Cisco CRS-1 router. It describes the security measures built into the control plane, management plane, data planes, other areas of security interest, such as the types of attack that a router might experience and the intelligent security strategies you can use in a Cisco IOS XR software-based network.

The preface contains the following sections:

- [Changes to This Document, page ix](#)
- [Obtaining Documentation, page x](#)
- [Documentation Feedback, page xi](#)
- [Cisco Product Security Overview, page xi](#)
- [Obtaining Technical Assistance, page xii](#)
- [Obtaining Additional Publications and Information, page xiii](#)

Changes to This Document

[Table 1](#) lists the technical changes made to this document since it was first printed

Table 1 *Changes to This Document*

Revision	Date	Change Summary
OL-??	April 2006	Initial release of this document.
	July 2006	The feature descriptions were added: <ul style="list-style-type: none">• Support for the RADIUS per VRF AAA feature was added.• Key chain management was added, with new cryptographic algorithms for SHA1, HMAC-MD5, and HMAC-SHA1.• Cryptographic algorithms MD5 and SHA1 were added for RSVP.• Packet-length specification in an ACL was added for the Cisco XR 12000 Series Router.• TTL specification in an ACL was added for the Cisco CRS-1 router.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation DVD

Cisco documentation and additional literature are available in a Documentation DVD package, which may have shipped with your product. The Documentation DVD is updated regularly and may be more current than printed documentation. The Documentation DVD package is available as a single unit.

Registered Cisco.com users (Cisco direct customers) can order a Cisco Documentation DVD (product number DOC-DOCDVD=) from the Ordering tool or Cisco Marketplace.

Cisco Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/>

Cisco Marketplace:

<http://www.cisco.com/go/marketplace/>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:
<http://www.cisco.com/en/US/partner/ordering/>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 1 800 553-NETS (6387).

Documentation Feedback

You can send comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com
- Nonemergencies—psirt@cisco.com



Tip

We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.x through 8.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one that has the most recent creation date in this public key server list:

<http://pgp.mit.edu:11371/pks/lookup?search=psirt%40cisco.com&op=index&exact=on>

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, 365 days a year, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support Website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>



Note

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support Website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco TAC engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

<http://www.cisco.com/go/marketplace/>

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

<http://www.cisco.com/en/US/learning/index.html>

Introduction

This introduction contains the following:

- Definition of security and the role of security in the larger context of high network availability
- A preview of the modules in the *Cisco IOS XR Security Guide*

This guide describes many aspects of the system security that Cisco IOS XR software provides on the Cisco CRS-1 router and Cisco 12000 Series Router.

Definition of Security

Internet attacks occur every day, and these attacks can include core routers as targets—usually with some form of denial of service (DoS) attack. Fortunately, the Cisco CRS-1 router and Cisco 12000 Series Router come with a highly robust, built-in security structure. These routers can withstand attacks and remain in service until the necessary responses have been completed. This guide describes the software architecture and the hardware support for addressing these attacks. It also provides guidance on how you can further protect your investment in these products.

Our security goal for the Cisco CRS-1 router and the Cisco 12000 Series Router is to neutralize attacks and recover quickly from any degradation of service. The point is to keep the router always operating, regardless of even a distributed denial of service (DDoS) attack. We developed the systems that run Cisco IOS XR software so that an attack on one process, service, or plane does not compromise any other process, service, or plane. For example, a process that is shut down automatically restarts, and when needed, a patch can be applied without restarting the router.

Essentially, these mechanisms are:

- Modularity of services
- Separation and isolation of management, control, and data (or forwarding) planes
- Multiple layers of increasingly stringent defense
- Convenient upgradability

A primary goal of the security mechanisms is to ensure that the CPU on the route processor (RP) processes only what it must process and to exclude anything else, such as the garbage data that an attack attempts to bring. The increasing stringency of the layered defense makes a successful attack on the RP CPU unlikely, and at every stage along the packet paths, a provision exists to address an attack. The Cisco CRS-1 router and Cisco 12000 Series Router go far beyond a single ASIC, a cluster of ASICs, or a software security monitor that is attempting to protect the system. The notion of defensive layers and increasing stringency is introduced in the *Fundamentals of Cisco IOS XR System Security* module and is elaborated throughout this guide.

Security is not an add-on feature. Security is designed into the router, so the router is shipped as a secure system as the default. Each component of the system must participate in the security of the system. For example, applications are designed so that they protect themselves and can terminate themselves when attacked. The operating system image itself and all the software images are protected during production, shipment, download, and activation. The integrity of the software is ensured by the Software Authentication Manager (SAM). In addition to the contribution to high availability made by SAM, software patches can be applied to a service or process through software module updates (SMUs). SMUs are also a part of the high availability effort. Convenient upgradability is provided by in-service software upgrades (ISSUs), and this support includes SMUs.

Document Modules

The following sections summarize the modules in this guide.

Fundamentals of Cisco IOS XR System Security

This module describes issues and operating system infrastructure as they apply to router security. The subjects include:

- [Hierarchy of the Layered Defense, page SG-19](#)
- [Modularity and Compartmentalization in the Planes, page SG-22](#)
- [Operating System Infrastructure and Security, page SG-23](#)
- [In-Service Software Upgrade, page SG-25](#)
- [Fault Manager Usage, page SG-25](#)
- [Memory Threshold Specifications, page SG-25](#)
- [Keys and Key Chain Management, page SG-25](#)

Node Access and Management Plane Security

This module describes the issues of, and provisions for, the management plane. The subjects include:

- [Secure Shell and Other Access Protocols, page SG-27](#)
- [Protecting Software Images, page SG-28](#)
- [Authentication, Authorization, and Accounting, page SG-28](#)
- [Task Groups and User Groups, page SG-28](#)
- [Simple Network Management Protocol, page SG-29](#)
- [Command-Line Interface, page SG-29](#)

Data Path Security

This module describes the security measures implemented in the data plane. The subjects include:

- [Introduction to Data Plane Security Measures, page SG-31](#)
- [Access Control Lists, page SG-35](#)
- [Quality of Service, page SG-41](#)
- [Unicast Reverse Path Forwarding, page SG-44](#)

Examples of Network Attacks

This module presents examples of different types of attacks. The subjects include:

- [The Routing System, page SG-47](#)
- [Types of Attacks Against a Routing System, page SG-48](#)

Attack Detection and Response

This module describes how Cisco IOS XR software detects an attack and how it—and you—can respond. The subjects include:

- [Always-Available CLI Commands, page SG-55](#)
- [NetFlow, page SG-56](#)
- [Watchdog System Monitor, page SG-58](#)
- [Fault Manager, page SG-59](#)
- [Alarms, page SG-61](#)
- [Logging Correlation, page SG-63](#)

Implementing Network Security Mechanisms

This module describes a set of best practices and how to implement the security mechanisms. The subjects include:

- [The Six Phases of a Security Strategy, page SG-66](#)
- [Default Settings, page SG-67](#)
- [General Recommendations, page SG-70](#)
- [Protecting the Legitimacy of the Routing Domain, page SG-85](#)

Fundamentals of Cisco IOS XR System Security

This module describes the fundamental concepts of Cisco IOS XR software security that are based on the operating system and the Cisco IOS XR software infrastructure. The topics in this module include:

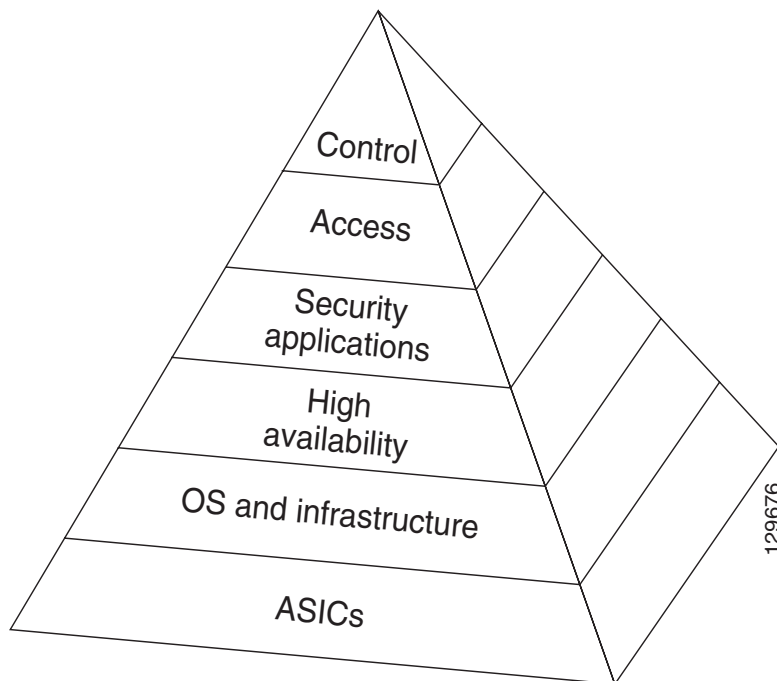
- [Hierarchy of the Layered Defense, page SG-19](#)
- [Modularity and Compartmentalization in the Planes, page SG-22](#)
- [Operating System Infrastructure and Security, page SG-23](#)
- [In-Service Software Upgrade, page SG-25](#)
- [Fault Manager Usage, page SG-25](#)
- [Memory Threshold Specifications, page SG-25](#)
- [Keys and Key Chain Management, page SG-25](#)

Hierarchy of the Layered Defense

This section introduces the security layers from the bottom up. The pyramid structure in [Figure 1](#) illustrates the increasing stringency of the router's defense. The topic at each layer of the pyramid are discussed in the sections of this module.

The peak of the pyramid is the router itself. Going up each layer in the pyramid, an attack is more likely to fail or at least be mitigated.

Figure 1 The Hierarchy of Defenses



Application-specific Integrated Circuits

Ideally, unwanted packets are dropped in the hardware on the line card (LC) or modular services card (MSC) ingress. The ASICs on these cards are the first defense against attack packets. Other parts of this hardware line of defense are the policers and microcode.

Operating System and Infrastructure

The operating system design and overall software infrastructure provide critical aspects of router security. These subjects are discussed in detail in the sections titled [Modularity and Compartmentalization in the Planes](#), page SG-22 and [Operating System Infrastructure and Security](#), page SG-23.

High Availability

In the context of this security guide, high availability (sometimes called “HA”) refers to the resiliency of the router and how HA is an antidote to a denial of service (DoS) attack. In Cisco IOS XR software, high availability contributes to security with code modularity, process isolation, and process resiliency. Support for HA includes:

- The support that software has for the patching of a software module. (See [In-Service Software Upgrade](#), page SG-25 for a description of a more complex dynamic software upgrade service.)
- Fault isolation:
 - Failure of one process does not compromise other processes.
 - The micro kernel (uKernel) architecture maintains protected memory space for each process.

- The uKernel is isolated and protected. (See [Operating System Infrastructure and Security, page SG-23](#) for more details.)
- The OS enables software developers to assign a relative priority to a process.
- Process management
 - Most system processes can restart without having an impact on packet forwarding.
 - Mirrored checkpoint servers support faster recovery.
 - The uKernel can limit the total number of processes, and it can limit the number of processes that an individual process can spawn. This capacity can mitigate the type of attack that attempts to spawn a flood of processes.
- Extensions (at the application level) for nonstop forwarding (NSF), which support graceful restart of routing and signaling protocols.

Security Applications

A variety of applications apply strictly to router security or include security benefits as a part of their operation. These applications are based either in the ASICs or in software:

- Unicast Reverse Path Forwarding (uRPF)
- Access control lists (ACLs)
- Local packet transport services (LPTS)
- Quality of service (QoS)
- Keys and key chain management

Management Access Security

The access facilities illustrated in [Figure 1](#) include:

- Access to the system resources requested by any running software application
- User access programs, such as Secure Shell (SSH)
- Task IDs and task groups
- User authentication through authentication, authorization, and accounting (AAA)—the protocols in TACACS+ or RADIUS)
- Command (task) authorization through AAA

For more details on user access, see “[Node Access and Management Plane Security](#).”

Control Plane Protocols

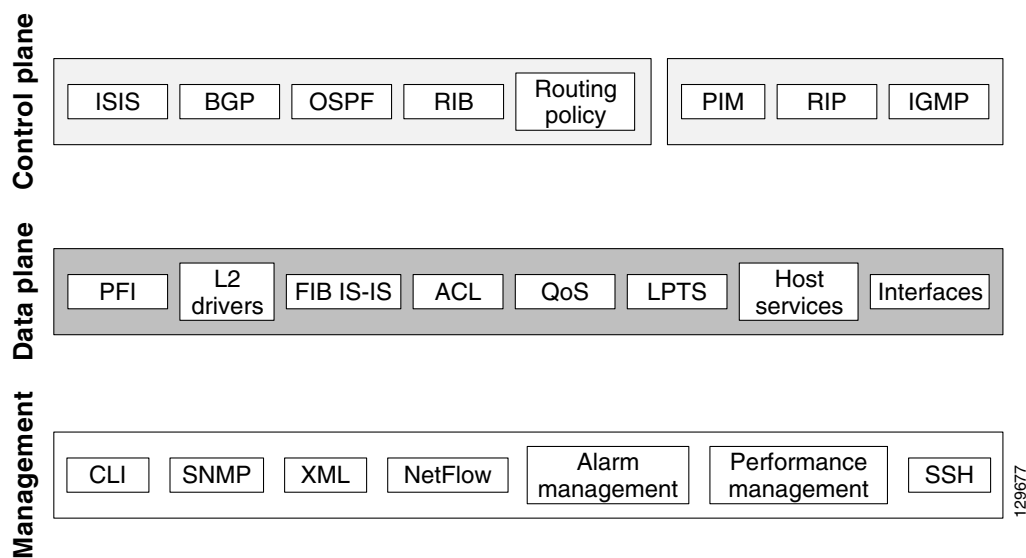
The control protocols and general-purpose routing applications must provide for security. For example, BGP must filter the peers so that BGP is not communicating with entities that it should not. The bases of control plane filtering reside in the ACLs, in the policies, and on the autonomous segment and the source from which it comes. In addition, the built-in feature LPTS automatically processes all control plane packets to provide protection against CPU and router resource attacks.

Modularity and Compartmentalization in the Planes

Modularity and compartmentalization are among the strongest foundations of security. In Cisco IOS XR software, the management plane, control plane, and data plane are fundamentally separated. For the management and control planes, software provides modularity and compartmentalization. Data plane security is based primarily in hardware, but software also reinforces data plane security. The planes are isolated so that the processes are protected from an attack on another plane. If one process fails due to an attack, other modules are not affected, and the compromised module automatically restarts without depending on other modules. For example, if TCP or the Network Time Protocol (NTP) were to fail, it would automatically restart.

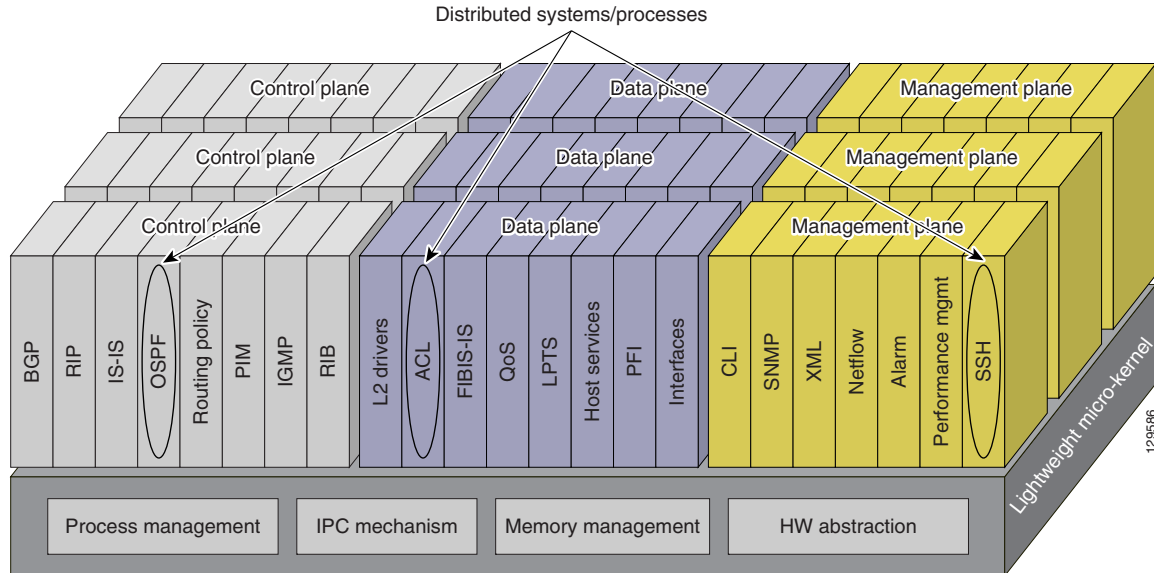
For a general view of the software modules operating within each plane, see [Figure 2](#).

Figure 2 Software Component Layering



Modularity and compartmentalization greatly contribute to the ability of Cisco IOS XR software to keep the Cisco CRS-1 router or Cisco 12000 Series Routers available during an attack. With the security scheme, no gating point of failure exists on a router. Services might be slowed, but the router continues to run during the attack, the attack response, and the service restoration process.

A comprehensive illustration of modularity appears in [Figure 3](#). This figure shows the management, control, and data planes and the processes and protocols that run within them. It also illustrates the layered core processes of the uKernel (bottom). Within [Figure 3](#), an example of the distributed system processes shows that a user is relying on SSH (management plane) to enter commands that relate to OSPF (control plane). An ACL (data plane) validates this action.

Figure 3 *Compartmentalization in the Planes and Lower Layers*

The system separates the planes in the following ways:

- In the data plane, all traffic forwarding or switching occurs in the ASICs.
- The CPU on an MSC or LC processes a packet without sending it to the route processor (RP), on which the control plane resides. This isolation of control from the line cards enhances security.
- Exception traffic that is punted to the CPU on the MSC or LC is separated from control traffic by dedicated queues.
- Exception packets are processed by different processes than control packets.

Operating System Infrastructure and Security

This section describes the aspects of the operating system and the overall software architecture that have been designed to enhance security.

uKernel

All essential services run outside the uKernel. Examples of these services are TCP, UDP, and hardware drivers. This architecture makes an attack on the uKernel highly unlikely. In addition, the uKernel is preemptive. It uses a mechanism of round robin and priority scheduling to protect the system against infinite loop attacks and flooding attacks.

The uKernel itself is very small, and everything around it runs on a process level. This process-oriented mechanism is an advantage during an attack. For example, if BGP fails, the rest of the system keeps running. For more information on processes, see [Restartable Processes](#).

Scheduler

The uKernel uses round-robin scheduling to protect against CPU starvation and lockups. (A lockup is a condition in which lower-priority processes are never serviced.) This aspect of the architecture helps with response time on the busiest system—even during an attack.

Memory Protection

Individual processes run inside their own memory space and do not corrupt other processes. Reserved memory is provided for critical processes that require protection during an attack. For example, memory is reserved for some always-available CLI commands that remain available during worst-case attacks.

CLI Availability

The CLI has a reserved memory pool to keep the CLI available during an attack. Usually, all CLI commands remain available during an attack. During an attack that causes low memory or some other out-of-resource (OOR) condition, a set of essential commands remains available. These always-available commands allow you to see the state of interface queues and processes and restart processes or, in extreme cases, shut down a process. This assortment of always-available commands ensures that you always have a way to mitigate attacks.

Restartable Processes

If an attack is able to shut down a process, the router continues running, and the process usually restarts itself. You can restart the process or apply a patch, if needed, without restarting the whole router. Furthermore, even if an attack on the uKernel is successful, the router or particular card on which the uKernel has stalled would restart on its own. You can patch the uKernel without restarting it. (However, if you must undo a patch to the uKernel, it must be restarted.)

Resource Monitoring Through wdsysmon

The watchdog system monitor (wdsysmon) monitors the CPU and memory. You can check the system log for potential low memory availability or high CPU usage. The Fault Manager is linked to wdsysmon. This linkage enables you to create customized Fault Manager responses to wdsysmon resource alerts.

Fault Manager

The Fault Manager can feed into the alarm subsystem to signal a possible attack. It can also run cleanup scripts or scripts that show the captured data that resulted from various attacks. This tool is flexible and depends on the topology and customer needs. The role of Fault Manager in security is described in more detail in the [Attack Detection and Response](#) module.

In-Service Software Upgrade

In-service software upgrade (ISSU) is the ability of the system to receive a software upgrade without requiring a restart of one or more processes. The ISSU contribution to high availability can apply from the highest-level software packages to the lowest-level modules.

For security in particular, ISSU lets you patch a package or even an individual process in the context of a software module upgrade (SMU). An SMU is the type of product installation envelope (PIE) most likely to apply to software that has been the object of an attack. In this case, only the affected processes are restarted, and no other parts of the system are disrupted.

The nonstop forwarding capability supports ISSU in that an LC or MSC continues to forward packets during the upgrade for a configurable period of time.

Fault Manager Usage

You can configure the Fault Manager to use scripts to respond to specific events in the system log. Conditions that need a response include:

- Unauthorized access attempts
- Potential low memory
- High CPU usage

Through the use of fault management scripts, you can provide for automatic detection and response to attacks. For more information on the Fault Manager, see *Cisco IOS XR System Management Configuration Guide* and *Cisco IOS XR System Management Command Reference*.

Memory Threshold Specifications

Memory usage can exceed tolerable limits during an attack, so memory thresholds can be configured to trigger an alert or remedial action.

If free memory drops below a configurable amount, a message to syslog is generated. (Expressed conversely, you could say if memory utilization exceeds a configurable threshold, a message is generated.) Configurable thresholds apply to memory usage by processes and memory usage by I/O.

Memory thresholds for applications can also be configured. Therefore, if an application is attacked, the attack cannot consume more memory than what you have allotted. In any case, sufficient memory is protected to support packets during their lifetime.

Keys and Key Chain Management

Routing protocols and network management applications frequently use software authentication capabilities for greater security when they communicate with each other. A common method of authentication is the use of shared secrets on all of the entities. These secrets typically are keys. Network entities exchange these keys before establishing trust with each other.

The collection of keys for authenticating a peer or peer group is managed in a logical container called a *key chain*. A key chain is a sequence of keys that is managed collectively. (Keys are used individually for authenticating a peer or peer group.) Key chain services do not affect the system-wide performance for control, data, or management traffic.

Key chain management lets you group a sequence of keys together under a key chain and assign a lifetime to each key. For task and command descriptions related to keys, key chains, and key chain management, see *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*.

In the current release, key chains apply to IP Service Level Agreements (IP-SLAs), Border Gateway Protocol (BGP), and Resource Reservation Protocol (RSVP).

The supported cryptographic algorithms from which you can choose are:

- MD5
- SHA1
- HMAC-MD5
- HMAC-SHA1

**Note**

In the current release, the cryptographic algorithms you can specify for RSVP are MD5 and SHA1.

**Note**

The key chain management feature is always enabled, so be aware that because keys have a user-configured lifetime, the validity of keys can be affected if the system clock changes.

Node Access and Management Plane Security

This module outlines the security measures implemented for the management plane in Release 3.4. For information on how to respond to attacks in the management plane, see the [Attack Detection and Response](#) module.

The management plane contains all issues related to user access to the router:

- [Secure Shell and Other Access Protocols](#), page SG-27
- [Protecting Software Images](#), page SG-28
- [Task Groups and User Groups](#), page SG-28
- [Simple Network Management Protocol](#), page SG-29
- [Command-Line Interface](#), page SG-29
- [Authentication, Authorization, and Accounting](#), page SG-28

Secure Shell and Other Access Protocols

This section lists the secure methods for accessing the node:

- Secure shell (SSH)
- Secure socket layer (SSL)
- Secure FTP (SFTP)
- Secure SNMP (version 3 is the most secure)

SSH is described in *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*.

SNMP security topics are described in the [Implementing Network Security Mechanisms](#) module.



Note

In general, we strongly advise against using Telnet as a method for accessing a router unless the crypto package is not available in your location. If you must use Telnet, best practices for passwords and other security measures should be used.

Protecting Software Images

This section introduces the methods used by Cisco to protect the software. The methods consist of:

- In-house procedures for creating security for a software image
- Authentication of software images when they are downloaded to the router
- Authentication of each executable when it starts up

Cisco Systems follows strict internal security procedures to ensure the integrity of software images. An image intended for installation on a router is shipped inside a file called a product installation envelope (PIE). The software is delivered with an encoded signature and a certificate attached to the PIE. The program that authenticates the image is the Cisco Software Authentication Manager (SAM).

SAM authenticates software during two different activities. First, it checks the code signature before the router can complete an image download. If SAM detects that an image is not valid, it blocks the download. Afterward, during run-time operation, SAM checks the authenticity of an application each time that executable starts up. For more details on SAM, see *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*.

Authentication, Authorization, and Accounting

The protocols for authentication, authorization, and accounting (AAA) operate on the router and have the following functions:

- Authentication confirms the identity of a user and thus, for example, oversees a login.
- Authorization determines whether a user can run the command that the user has entered.
- Accounting creates records of events associated with a user session. The recorded information can apply to billing, monitoring, auditing, and so on. Accounting requires an external server.

A variety of more complex AAA-related operations can be supported when an external server is used. The communication between resident AAA and the external servers relies on either the TACACS+ or RADIUS protocols.

On a smaller scale, the Per VRF AAA feature enables AAA services to be based on virtual private network (VPN) routing and forwarding (VRF) instances. The provider edge (PE) or virtual home gateway (VHG) communicates directly with the customer RADIUS server. This server is associated with the customer VPN without going through a RADIUS proxy.

The tasks and commands for AAA are described in *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*. Each of these documents has a module dedicated to AAA.

Task Groups and User Groups

To support the current method of controlling access to the router, Cisco IOS XR software uses a scheme of task groups, task identifiers (one task ID for each CLI command), and user groups:

- A task group is a collection of tasks (operations) that can be performed by commands on the CLI. An individual task group entry consists of a task ID. (Each command available in the CLI has a task ID associated with it. The task ID and possible action—read, write, execute, or debug—can be viewed in each command description in the command references.)
- Multiple task groups can be made available to a user group.

- A user group consists of the user-members and the task groups available to that user group. Thus, a user group is characterized by its associated task groups, but the *members* of the group are users. We provide a robust set of predefined user groups, and you can create custom user groups.
- A user can belong to multiple user groups. When you create a user account, you can associate the user with one or more of the user groups we provide or to a user group that you create.

Simple Network Management Protocol

Although SNMP is a part of the management plane, the description for this protocol appears in the [SNMP](#) section in the *Implementing Network Security Mechanisms* module.

Command-Line Interface

To access the CLI on a router, you can use SSH or Telnet. We strongly recommend that you use SSH if your routers have the crypto package. The crypto package is not available in all countries, so your only choice might be Telnet. Telnet is not secure because it transmits clear text.

Data Path Security

This module describes the security measures implemented for the data plane. The module sections are:

- [Introduction to Data Plane Security Measures, page SG-31](#)
- [Access Control Lists, page SG-35](#)
- [Quality of Service, page SG-41](#)
- [Unicast Reverse Path Forwarding, page SG-44](#)

Introduction to Data Plane Security Measures

This section focuses on:

- Paths that a packet can take on the router
- Summary of security actions in the data plane

We recommend that you become familiar with the high-level view of the packet paths by reading the “[Packet Walk](#)” section before you examine the other elements of data plane security measures. After the packet walk, the section titled “[Packet Paths](#)” gives details about where various types of packets travel on the router and where those packets are policed.

Packet Walk

A packet walk describes where a packet goes, beginning with its arrival at the ingress. This walk involves a variety of decisions that are based on tests—for example, is it a transit packet or a for-us (receive) packet, and does the packet have a known destination IP address on this router? A transit packet goes to an egress interface. A valid for-us packet is destined to the route processor (RP) or distributed route processor (DRP).

The first parts of the packet walk are:

- Initial processing upon arrival of a packet at the ingress (which could include invalid packet discard)
- Processing of a valid transit packet (if initial processing detects a transit packet)

Packets Arriving at the Ingress

When a packet arrives at the ingress, the receive ASICs do initial tests for validity:

- Initial validity checks apply to L2 framing, certain MPLS packet checks, and L3 sanity checks. Examples of L3 sanity checks are checksums, checks for illegal source or destination address, packet length checks, and so on.
- If the packet is valid, the interface then determines whether it is a transit packet or a for-us packet.
- If the packet is transit, it receives further checks before it passes to the egress interface. Broadly speaking, these checks are ACL checks and a lookup in the forwarding information base (FIB).
- If the packet is for-us, it begins its ACL check and lookup in the internal forwarding information base (iFIB).

Transit Packets

A transit packet takes the following course:

- The FIB within the receive ASIC checks it for validity, packet type (for example, multicast), and input features as determined by ACLs.
- The packet goes to the switch fabric queues. Broadly speaking, this queuing reflects packet priority.
- The switch fabric directs the packet to the egress interface.

For-us Packet Walk

After the receive ASIC determines that the packet has the current router as its final destination (and is therefore “for us”), a complex scheme of decisions and security checks begins. These actions are:

- The packet enters the iFIB to determine the packet type (for example, a fragment or an exception packet). If the packet is a fragment, it must be reassembled on the RP and returned to the iFIB to determine its next stage of processing.
- A check by the Generalized Time-to-Live (TTL) Security Mechanism (GTSM) determines if the packet should be dropped.
- Control plane policing also determines whether a packet should be dropped.
- Having passed initial security checks, the packet goes to RP queues that reside on the modular services card (MSC) or line card (LC). These queues perform rate limiting in case the packet stream is part of an attack that has progressed past the preceding security checks.
- After leaving the MSC or LC, the packet enters a queue on the RP before the CPU processes it. This queue also performs rate limiting to prevent the CPU from being overwhelmed in a flooding attack.

Packet Paths

For packets at the ingress of the MSC or LC, data plane protection initially depends on the ASICs and the access control lists that have been downloaded to these ASICs.

After the ASICs, an array of queues can shape the traffic. The queues on the MSC keep the local CPU from being overwhelmed by attack packets. Similarly, the switching fabric also has queues to prevent an attack from flooding the fabric.

Data plane exception packets go to the CPU on the RP or distributed route processors (DRPs). The punted packets are again policed on the RP, and the RP also performs some sanity checking on these exception packets. For example, sanity checking ensures that the IP address is valid and drops the packet if the address is invalid.

The columns in [Table 2](#) show where packets are policed on the LC, MSC, and RP.

The table shows where the different types of packets are sent and what part of the router can act on these packet types. The entries in the destination columns show policing or another action taken on the packet; an X shows other points where a packet can go; and an empty cell shows that a particular type of packet does not reach that point. For example, valid transit packets do not go to the CPU on the MSC, LC, or RP.

The locations where packets can go on the cards are:

- MSC/LC PSE—Packet processor on the MSC or LC
- MSC/LC CPU—CPU on the MSC or LC
- RP CPU—CPU on the route processor

Table 2 Where Packets Are Redirected, Policed, or Punted

Packet Type	MSC/LC PSE	MSC/LC CPU	RP CPU
For-us Packets			
For-us packets (TTL=255, 254, . . .)	GTSM	GTSM	X
Transit Traffic			
Transit packets	Full feature		
Transit packets, IP options	Policed	X	
Transit packets, IP option “router alert”	Policed	X	X
Source route			X
Packets in which TTL=1			
Packets that require ARP resolution	Policed	X	
Unicast Receive Traffic			
Internet Control Message Protocol (ICMP) echo request, packets that require logging	Policed	X	
Any other ICMP (including ICMP with options)	Policed	X	
Management traffic (SSH, SNMP, XML, and Telnet)	Policed		X
Management traffic (NetFlow and Cisco Discovery Protocol)	Policed	X	
Routing (BGP, OSPF, IS-IS, and so on)	Policed		X
Multicast, Broadcast			
Multicast control traffic (OSPF, PIM, HSRP, and so on)	Policed		X
First packet of multicast stream	Policed	X	
Broadcasts	Policed	X	X
Special			
Traffic needing fragmentation	Policed	X	
MPLS traffic needing fragmentation	Policed	X	
L2 packets (keepalives and so on)	Policed	X	

Summary of Data Plane Security Mechanisms

The following points summarize the mechanisms of data plane security on the Cisco CRS-1 router and Cisco XR 12000 Series Router:

- Most data plane packets are processed by hardware. Only exception packets, option packets, or locally destined packets (for-us packets) are processed by software. Some examples are ARP requests, Layer 2 items, encapsulations, and so on.
- The packets that are punted to the RP bypass the CPU on the MSC or LC and go directly to the RP.
- Filters can be static or dynamic. For example, an ACL is a static filter, but uRPF is dynamic.
- Dynamic filters and Cisco CRS-1 routerd policers are configured for packets that are to be processed by any of the CPUs. In the current release, the rates for the policers are not user-configurable.
- uRPF filtering of IPv4 and IPv6 addresses can be strict or loose. Strict means the return path to the source IP address of the packet must traverse the interface on which the packet arrived. Loose means the source IP address of the packet can simply be reachable through a valid prefix in the FIB. With a valid packet in the FIB, the source can be reachable through any interface in the router. For more information on uRPF, see “[Unicast Reverse Path Forwarding](#).”

Support for strict and loose filtering of IPv4 and IPv6 addresses is as follows:

- The Cisco CRS-1 router supports loose and strict filtering for IPv4 and IPv6.
- The Cisco XR 12000 Series Router supports strict and loose filtering for IPv6.
- In the current release, the Cisco XR 12000 Series Router supports only loose filtering for IPv4.
- The fabric responds to either of two packet priorities to ensure fast service for high-priority packets.
- Packet Accounting Statistics
 - NetFlow, for ingress or egress
 - Interface statistics



Note

To determine what has occurred at the ingress or egress of an interface, you can use NetFlow version 9. (For information on the use of NetFlow, see the [Attack Detection and Response](#) module.)

- Congestion avoidance mechanisms
 - Ingress and egress policing
 - Ingress and egress traffic shaping
 - Weighted random early discard (WRED)

The following is a list of additional details that apply to the access lists:

- The number of access controls elements (ACEs) has no impact on router performance. However, for IPv4, if hardware counters are enabled, a hit to performance occurs on the Cisco CRS-1 router but not on the Cisco XR 12000 Series Router.
- ACLs are applied on the interface level or the subinterface level.
- Matching on packet header supports prioritization.
- For ingress and egress ACLs, the following are supported:
 - IP source and destination addresses and masks
 - L4 source and destination ports (and ranges)

- TCP flags
 - IP DSCP/precedence
 - IPv6 option headers (routing, destination, and AH)
 - Non-initial fragments
 - Packet length
 - TTL value—supported on the Cisco CRS-1 router only in the current release
- Specific response actions can permit or deny forwarding. The router can log each such decision.

Access Control Lists

The access control lists (ACLs) function as the first line of defense against an attack. This section describes the benefits and operational characteristics of an ACL. It also introduces the advantages of using an address prefix list.

For a complete description of the access list and prefix list commands, see the *Access List Commands on Cisco IOS XR Software* and *Prefix List Commands on Cisco IOS XR Software* modules in *Cisco IOS XR IP Addresses and Services Command Reference*. For descriptions of the tasks that apply to ACLs and prefix lists, see *Cisco IOS XR IP Addresses and Services Configuration Guide*.

An ACL consists of one or more filter elements and an action that results from each filter test. The filter element consists of a criterion, such as a source address, destination address, protocol, packet length, or protocol-specific parameter. The action is either to permit or deny some action. Each combination of filter and action is called an access control entry (ACE). Together, the ACEs in an ACL define a network traffic profile. Many features (and therefore many commands) within Cisco IOS XR software can refer to the profile embodied in an ACL. Some examples of these features are traffic filtering, priority or custom queueing, and dynamic access control.



Note

Specification of packet length in an ACL helps to protect against certain DDoS attacks or worms in which the signature of the attack is well-known. The ACL can match on the length of the attack packet and drop it. A well-known signature is a common occurrence for many DoS attacks and worms.

A prefix list can serve as an alternative to an ACL in many Border Gateway Protocol (BGP) route filtering commands. A prefix list is used by route maps and route filtering operations. A prefix is a portion of an IP address, starting from the far left bit of the far left octet. By specifying how many bits of an address belong to a prefix, you can use a prefix to aggregate addresses and perform a function on them, such as redistribution (filter routing updates).

Cisco IOS XR Access Lists and Prefix Lists Feature Highlights

This section summarizes the high-level user control over ACLs and prefix lists. Cisco IOS XR software:

- Lets you clear counters for an ACL or prefix list by using a specific sequence number.
- Lets you copy the contents of an existing ACL or prefix list to another ACL or prefix list.
- Does not differentiate between standard and extended ACLs (for backward compatibility).
- Lets you apply sequence numbers to permit or deny statements and to resequence, add, or remove such statements from an ACL or prefix list.



Note In the current release, Cisco IOS XR software does not support resequencing of IPv6 prefix lists.

Purpose of IP Access Lists

In general, an ACL filters packets to control which packets move through the network and where packets can move. This control can help to limit network traffic and restrict the access of users and devices to a network. ACLs have many uses, so many commands accept a reference to an access list in their command syntax. More particularly, ACLs can do the following:

- Filter incoming packets on an interface. Ideally, attacking packets are dropped through ACL action.
- Filter outgoing packets on an interface. For example, bad packets can be kept from spreading.
- Restrict the contents of routing updates.
- Limit debug output based on an address or a protocol.
- Control VTY access to block access by unauthorized persons.
- Identify or classify traffic for advanced features, such as congestion avoidance, congestion management, and priority and custom queueing.

Contents of an ACL

An ACL is a sequential list consisting of permit and deny statements that apply to IP addresses and possibly to upper-layer IP protocols. An ACL is a standalone file: You can create an ACL, but it has no meaning or effect until the ACL is referenced by a command that accepts an ACL. Multiple commands can reference the same ACL. An ACL can control traffic arriving at the router or leaving the router, but not traffic originating at the router. Many software commands accept an ACL as part of their syntax.

ACL Processes and Rules

Note the following processes and rules before you configure an ACL:

- The software tests the source or destination address or the protocol of each packet being filtered against the conditions in the ACL, one condition (permit or deny statement) at a time.
- If a packet does not match an ACL statement, the packet is then tested against the next statement in the list.
- If a packet and an ACL statement match, the remaining statements in the list are skipped and the packet is permitted or denied as specified in the matched statement. The first entry that the packet matches determines whether the software permits or denies the packet.
- If the ACL denies the address or protocol, the software discards the packet and returns an ICMP Host Unreachable message. The generation of ICMP error messages is configurable in the Cisco IOS XR software.
- If no conditions match, the software drops the packet because each ACL ends with an unwritten deny statement. (If the packet has not been permitted by the time it was tested against each statement, it is denied.)
- The ACL should contain at least one permit statement; otherwise, all packets would be denied.

- Because the software stops testing conditions after the first match, the order of the conditions is critical. The same permit or deny statements specified in a different order could result in a packet being passed in one circumstance and denied in another circumstance.
- If an ACL is referenced by its name in a command but the ACL does not exist, all packets pass.
- One ACL is allowed for each interface, protocol, and direction.
- Transit packets are processed by an ACL before they are routed to an egress interface. An inbound ACL is efficient because it saves the overhead of routing lookups when the packet is discarded. If the packet is permitted by the tests, it is processed for routing. For an ingress ACL, **permit** means continue to process the packet after receiving it; **deny** means discard the packet.
- Outbound ACLs process transit packets before they leave the router. Transit packets are routed to the outbound interface and then processed by the egress ACL. For outbound ACLs, **permit** means send it to the output buffer; **deny** means discard the packet.
- An ACL cannot be removed if that list is associated with an access group that is in use. Remove the access group from the interface or service before you remove the ACL.
- An ACL must exist before you can use the **ipv4 access group** command for that ACL.

Suggestions for Creating an ACL

This section contains information you should understand before creating an ACL.

- Create the ACL before you apply it to an interface for the following reasons:
 - An interface that receives an empty ACL permits all traffic.
 - If you apply a nonexistent ACL to an interface and then configure the ACL, the first statement in that ACL goes into effect, but the implicit deny statement at the end could cause immediate access problems.
- Organize the ACL so that more specific references in a network or subnet appear before more general ones.
- To make the purpose of individual statements more easily understood at a glance, you can write a helpful remark before or after any statement.

Source and Destination Addresses

Source and destination addresses are two of the most typical fields in an IP packet on which to base an ACL. Specify source addresses to control packets from certain networking devices or hosts. Specify destination addresses to control packets being sent to certain networking devices or hosts.

Wildcard Mask and Implicit Wildcard Mask

Address filtering uses wildcard masks to indicate whether the software checks or ignores corresponding IP address bits when comparing the address bits in an ACE to a packet being submitted to the ACL. By carefully setting wildcard masks, you can select a single or several IP addresses for permit or deny tests.

Wildcard masking for IP address bits uses the number 1 and the number 0 to specify how the software treats the corresponding IP address bits. A wildcard mask is sometimes referred to as an *inverted mask*, because a 1 and 0 mean the opposite of what they mean in a subnet (network) mask.

- A wildcard mask bit 0 means *check* the corresponding bit value.
- A wildcard mask bit 1 means *ignore* that corresponding bit value.

You are not required to supply a wildcard mask with a source or destination address in an ACL statement. If you use the **host** keyword, the software assumes a wildcard mask of 0.0.0.0.

Unlike subnet masks, which require contiguous bits to indicate a network and subnet to be 1s, wildcard masks allow noncontiguous bits in the mask. However, in the current release of IPv6 ACLs, only contiguous bits are supported.

Transport Layer Information

You can filter packets based on transport layer information, such as whether the packet is a TCP, UDP, SCTP, ICMP, or IGMP packet, or any L4 type (0–255).

Access Control Entry Sequence Numbering

This section describes the benefits and characteristics of having a sequence number for each ACE.

Benefits

The IP Access List Entry Sequence Numbering feature lets you add sequence numbers to ACEs and to resequence them as needed. Using sequence numbers with ACEs simplifies the work of modifying an ACL. When adding an entry, you can choose the sequence number so that the ACE position in the ACL is optimal. If necessary, you can resequence current entries to create room for a new entry.

**Note**

In the current release, resequencing of an IPv6 prefix list is not supported.

Sequence Numbering Characteristics

This section describes the characteristics of the ACE sequence numbering within an ACL.

- If you start adding ACEs to a new ACL but do not include sequence numbers, the system assigns the number 10 to the first entry and then adds a 10 to each successive entry.
- If you add an ACE without a sequence number, the system gives the entry a number whose value is 10 higher than the last sequence number in that list and places the new ACE at the end of the list.
- You can add ACEs without affecting traffic flow or hardware performance.
- Distributed support ensures that the sequence numbers are synchronized in the RP and LC or MSC.
- This feature works with named standard and extended ACLs. An ACL name can also be a number.

IP Access List Logging Messages

Cisco IOS XR software can provide logging messages about packets that have been permitted or denied by a standard ACL. Any packet with a match in the ACL filters causes an informational message about the packet to go to the console. You can specify the level of logged messages by using the **logging console** command.

The first packet with a match in an ACL immediately triggers a logging message. Subsequent packets are logged over 5-minute intervals before the system logs or displays them. A message includes:

- ACL number
- Whether the packet was permitted or denied
- Source IP address of the packet
- Number of packets from the source that were permitted or denied in the prior 5-minute interval

You can use the {**ipv4 | ipv6**} **access-list log-update threshold** command to specify the number of packets that, when matched to an ACE (and are permitted or denied), can trigger a log message.

The logging facility might drop some logging message packets if it receives too many or if more than one logging message per second is processed. This restriction prevents the CPU from using too many cycles for logging ACL message packets. For the same reasons, you should not use the logging facility as a billing tool or an accurate source for the number of matches to an ACE.

Access Lists with Fragment Control

The IP Access Lists with Fragment Control feature significantly increases the granularity of control over non-initial fragments. (Before the fragment control feature became available, nonfragmented packets and the initial fragment of a packet were processed by IP extended access lists—if such an ACL had been applied—but non-initial fragments that were permitted by default.) You can specify whether the system examines non-initial IP fragments of packets when it applies an IP extended ACL.

(Before the fragment control feature became available, nonfragmented packets and the initial fragment of a packet were processed by IP extended access lists if such an ACL had been applied. In contrast, noninitial fragments were permitted by default.) You can specify whether the system examines noninitial IP fragments of packets when it applies an IP extended ACL.

Because non-initial fragments contain only Layer 3 information, an ACE that contains only Layer 3 information can be applied to non-initial fragments. The fragment has all the information the system requires for filtering, so the entry is applied to the fragments.

Fragment control depends on the optional **fragments** keyword in the following ACL commands:

- **deny (IPv4)**
- **permit (IPv4)**
- **deny (IPv6)**
- **permit (IPv6)**

By specifying the **fragments** keyword in an ACE, that particular entry applies only to non-initial fragments of packets, and the fragment is then permitted or denied.

The behavior of ACEs with respect to the **fragments** keyword is summarized in the table that follows.

If the Access-List Entry has . . .	Then . .
. . . no fragments keyword and all of the ACE information matches,	For an ACE containing Layer 3 and Layer 4 information: <ul style="list-style-type: none"> • An entry applies to nonfragmented packets and initial fragments. <ul style="list-style-type: none"> – If the entry matches and is a permit statement, the packet or fragment is permitted. – If the entry matches and is a deny statement, the packet or fragment is denied. • The entry is also applied to non-initial fragments as described here. Because non-initial fragments contain only Layer 3 information, only the Layer 3 portion of an ACE can be applied. <p>If the Layer 3 portion of the ACE matches, and</p> <ul style="list-style-type: none"> – If the entry is a permit statement, the non-initial fragment is permitted. – If the entry is a deny statement, the next access-list entry is processed.
. . . the fragments keyword and all the ACE information matches,	The ACE is applied only to non-initial fragments. Note The fragments keyword cannot be configured for an ACE that contains Layer 4 information.

ACEs that contain the **fragments** keyword are applied to only IP packets with a fragment offset greater than 0 (F0>0). Thus, both nonfragmented packets and initial fragments of IP packets will never match ACEs that contain the **fragments** keyword. Nonfragmented packets and initial fragments of IP packets will be compared only against ACEs that do not contain the **fragments** keyword. In some cases, therefore, you might have to configure two ACEs to control affected IP packets: One ACE would contain the **fragments** keyword to cover noninitial fragments, where F0>0. The second ACE would not have the 'fragments' keyword and would cover nonfragments and initial fragments, where F0=0.

Multiple ACEs with Layer 4 port information are neither necessary nor useful when you use the **fragments** keyword because noninitial fragments do not contain Layer 4 port information. Only one ACE per Layer 4 protocol is required when you use the **fragments** keyword with such entries.

You can also apply the **fragments** keyword to dynamic access lists.

Comments About Entries in Access Lists

You can include comments (remarks) about entries in an ACL by using the **remark** ACL configuration command. Comments can make an ACL easier to understand and scan. A comment line can have up to 255 characters.

A remark can go before or after a permit or deny statement, but you should be consistent about where remarks go. For example, having some remarks *before* the associated permit or deny statements and some remarks *after* the associated statements can be confusing. You can sequence remarks as needed.

Access Control List Counters

ACL counters are maintained in hardware and software. Hardware counters are used for packet-filtering applications, such as when an access group is applied to an interface. Software counters are used by all the applications and mainly pertain to software packet processing.

To display the hardware counters for an access group, use the **show {ipv4 | ipv6} access-lists** command.

To clear the hardware counters, use the **clear {ipv4 | ipv6} access-list** command in EXEC mode.

Hardware counting is disabled by default for IPv4 ACLs because of a small performance penalty. To enable hardware counting, use the **ipv4 access-group access-list-name {in | out} [hw-count]** command. This command is optional, and counting is enabled only on the specified interface.

Software counters are updated for the packets that software processes—for example, exception packets that are punted to the LC CPU for processing, or an ACL that is used by routing protocols, and so on. The counters that are maintained are an aggregate of all the software applications that are using that ACL. To display software-only ACL counters, use the **show ipv4 access-lists access-list-name [sequence number]** command in EXEC mode.

**Note**

All the information in this section applies to IPv6, with one exception: hardware counting is always enabled. (No *hw-count* option exists in the IPv6 access-group mode.)

BGP Filtering Using Prefix Lists

Prefix lists can be used as an alternative to access lists in many BGP route-filtering commands. The advantages of using prefix lists are as follows:

- Significant performance improvement in loading and route lookup of large lists.
- Incremental updates are supported.
- More user friendly CLI. The CLI for using access lists to filter BGP updates is difficult to understand and use because it uses the packet filtering format.
- Greater flexibility.

Before using a prefix list in a command, you must set up a prefix list. You should also consider the use of sequence numbers for the entries in the prefix list.

Quality of Service

Quality of service (QoS) is a technique for prioritizing flows of transit traffic and providing preferential forwarding for higher-priority transit packets. (In this QoS description, the words *traffic flow* and *packet* are used interchangeably.) The role of QoS in security matters is to let you determine the traffic that should have the highest priority during an attack in particular or during any congested situation. This section introduces the benefits and techniques of the QoS feature.

For descriptions of the tasks and commands that apply to QoS, see the following documents:

- *Cisco IOS XR Modular Quality of Service Configuration Guide*
- *Cisco IOS XR Modular Quality of Service Command Reference*
- *Cisco IOS XR IP Addresses and Services Configuration Guide*
- *Cisco IOS XR IP Addresses and Services Command Reference*

The fundamental reason for implementing QoS in a network is to provide better service for a specific traffic flow. Broadly speaking, a traffic flow is a packet that is moving from an ingress interface and is destined either to an egress interface or to the RP. A traffic flow must be identified, classified, and prioritized on all routers and then passed along the data path throughout the network. Specifically, a traffic flow is a combination of:

- Source and destination addresses
- Source and destination socket numbers
- A session identifier

QoS helps to provide better and more predictable network service by supporting:

- Bandwidth allocation
- Improved loss characteristics
- Avoidance and management of network congestion
- Metering of network traffic
- Setting of traffic flow priorities across the network

Benefits of Implementing Cisco IOS XR QoS Features

The QoS features let networks control and service a variety of networked applications and traffic types with predictability. (In all cases, the traffic types refer to transiting packets—not sent to the RP.) QoS in a network helps to achieve:

- Control over resources: You have control over the use of resources such as bandwidth, equipment, wide-area facilities, and so on. For example, you can limit bandwidth consumed over a backbone link by FTP transfers or give priority access to an important database.
- Tailored services: If you are an Internet service provider (ISP), the control and visibility provided by QoS enables you to offer carefully tailored grades of service differentiation to your customers.
- Coexistence of mission-critical applications ensures that:
 - The WAN is used efficiently by mission-critical applications.
 - The required bandwidth and minimal delays are ensured for time-sensitive applications.
 - All applications that are using a link get their fair share of service but without interfering in the mission-critical traffic.

QoS Techniques in Cisco IOS XR Software

This section describes the following techniques for supporting QoS across a heterogeneous network:

- Packet classification
- Congestion management
- Congestion avoidance

Before implementing QoS, you should identify and evaluate the characteristics of the network traffic because not all of these techniques are appropriate for every network.

Packet Classification

Packet classification techniques let you identify the traffic flows and partition traffic into multiple priority levels or classes of service. After a traffic flow is identified, it can be marked as a traffic class.

Identifying a traffic flow can be done through a variety of methods, such as:

- Access control lists (ACLs)
- Protocol match
- IP precedence
- IP differentiated service code point (DSCP)

Marking of a traffic flow is performed by one of two methods:

- Setting IP precedence bits
- Setting DSCP in the IP type of service (ToS) byte

For conceptual and configuration information about packet classification, see *Configuring Modular Quality of Service Packet Classification on Cisco IOS XR Software* module in *Cisco IOS XR Modular Quality of Service Configuration Guide*.

Congestion Management

Congestion management techniques control congestion after congestion has occurred. One way that network elements respond to an overflow of incoming traffic is by using a queueing algorithm to sort the traffic and then determine some servicing method to prioritize it on an output link.

Cisco IOS XR QoS software implements the Low Latency Queueing (LLQ) feature. LLQ brings strict priority queueing (PQ) to the Modified Deficit Round Robin (MDRR) scheduling mechanism. This mechanism guarantees a minimum bandwidth for each traffic class. LLQ with strict PQ allows delay-sensitive data, such as voice, to be dequeued and transmitted before packets in other queues. For conceptual and configuration information about these congestion management concepts, see the *Configuring Modular Quality of Service Packet Classification on Cisco IOS XR Software* module in *Cisco IOS XR Modular Quality of Service Configuration Guide*.

Congestion Avoidance

Congestion avoidance techniques monitor network traffic flows to help you anticipate and avoid congestion before a problem occurs at likely bottlenecks. These techniques can simultaneously:

- Provide preferential treatment for traffic that has been classified as realtime-critical under congestion situations (a video stream, for example)
- Maximize the network throughput and the utilization of capacity and resource
- Minimize packet loss and delay

To meet these objectives, Cisco IOS XR software supports Random Early Detection (RED), Weighted RED (WRED), and tail-drop QoS congestion avoidance features.

For detailed conceptual and configuration information about congestion avoidance techniques, see the *Configuring Modular Quality of Service Packet Classification on Cisco IOS XR Software* module in *Cisco IOS XR Modular Quality of Service Configuration Guide*.

Traffic Policing

The traffic policing feature limits the input or output transmission rate of a class of traffic (based on user-defined criteria) and marks the packets by setting the IP precedence value, QoS group, or DSCP value. Cisco IOS XR QoS software implements traffic policing through its rate-limiting and class-based traffic shaping capabilities.

For conceptual and configuration information about traffic policing, see the *Configuring Modular Quality of Service Congestion Management on Cisco IOS XR Software* module in *Cisco IOS XR Modular Quality of Service Configuration Guide*.

Traffic Shaping

Traffic shaping allows control over the traffic such that:

- When traffic leaves an interface, it matches the rate at the egress to the rate of the remote interface.
- The traffic complies with the policies that were contracted for it.

Traffic that complies with a profile can be shaped to meet downstream requirements and, therefore, eliminate bottlenecks in topologies that have data-rate mismatches.

Cisco IOS XR QoS software supports a class-based traffic shaping method through a CLI mechanism in which parameters are applied on a per-class basis.

For conceptual and configuration information about traffic shaping, see *Configuring Modular Quality of Service Congestion Management on Cisco IOS XR Software* module in *Cisco IOS XR Modular Quality of Service Configuration Guide*.

Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (uRPF) is a security feature that helps the router thwart an attack. uRPF verifies that the packet source IP address is actually reachable and thus limits the ability of an attacker to spoof the source IP addresses of network packets.

Unicast Reverse Path Forwarding (uRPF) is a security feature that helps mitigate problems caused by the introduction of packets with malformed or spoofed IP source addresses into the network. uRPF mitigates damage by discarding packets that lack a verifiable entry (route back) to the source IP address of each packet in the Forwarding Information Base (FIB).

Deploying this feature limits the ability of an attacker to use spoofed source IP addresses, and, at a minimum, provides traceback capabilities.

The forms of uRPF that are available in the current release are strict mode and loose mode:

- When uRPF strict mode is applied to an interface, the router examines all received packets to verify that the source IP address and interface upon which the packet was received `_match_` the entry in the FIB table.

uRPF strict mode is commonly deployed on point-to-point interfaces where the same interface is used for both directions of packet flows.

- When uRPF loose mode is applied to an interface, the router examines all received packets and verifies that only the source IP address of the packet `_matches_` a valid entry in the FIB table. The (stricter) verification of received interface is not checked.

uRPF loose mode is the foundation of triggered black hole filtering techniques that many ISPs use. In loose mode, uRPF can effectively drop DoS and DDoS attack packets based on a source IP address and very quickly push the scheme out to hundreds of routers.

**Note**

In the current release, the Cisco XR 12000 Series Router does not support uRPF strict mode for IPv4 packets on Engine 3 and Engine 5 line cards.

The interface types that support uRPF are 1GE, 10 GE, POS, Bundled-POS, Bundled-Ethernet, and VLAN-over-Bundled-Ethernet.

The IP address types that support uRPF are IPv6 and IPv4.

Beyond the basic functionality already described, uRPF does the following:

- Allows packets with the combination of 0.0.0.0 and 255.255.255.255 source and destination so that the BOOTP and DHCP functionality work.
- Drops any packets that resolve to a NULL0 route.
- Drops any packets that resolve to a default route.
- Drops any packets that have a source IP on the current router (commonly referred to as *self-ping*).

Examples of Network Attacks

This module starts with a high-level description of our approach to security. The remainder of the module consists of descriptions of attacks. The module sections are:

- [The Routing System](#)
- [Types of Attacks Against a Routing System](#)
 - [Disrupting Peering](#)
 - [Falsifying Routing Information](#)
 - [Abusing Routing Stability Features to Reduce Network Availability](#)
 - [Forcing BGP Peer Damping by Injecting Flapping Routing Information](#)
 - [Attacking a Routing System](#)

The Routing System

The first consideration for routing security usually is to secure routing protocols, but this view is too narrow. The focus must be on the routing system as a whole. The routing system consists of the three major categories of routing *protocols*, *devices*, and *topology information*.

- Each routing protocol consists of the following parts:
 - Semantics that transport topology information across a network or an internetwork
 - Algorithms that determine the shortest path to any given destination within the network
- The devices run the routing protocols and switch packets along the paths that the routing protocols have chosen as the best paths to each reachable destination within the network.
- The topology information represents the topology of the network and reachable destinations within the network. Topology information is carried within the routing protocol.

To protect routing within a network, each part of the routing system must be protected, rather than just the routing protocol and its semantics and algorithms.

Types of Attacks Against a Routing System

This section describes the types of possible threats to a routing system. Generally, the point of an attack against a routing system falls into one of two categories:

- Disrupting peering or neighbor associations
- Falsifying routing information

Each of these topics is covered for a variety of attacks within each category.

Disrupting Peering

Attacks on peering generally attempt to deny the use of network resources to authorized users of the network. These attacks usually are not as effective as they appear from outside the network because most routing protocols rebuild the peering session after the attack has stopped. For example, if BGP peering or the adjacency of two OSPF neighbors is disrupted, the routers rebuild their adjacency after the attack has stopped, and the routing system continues to forward traffic.

Redundant links and other high-availability techniques can also defeat most peering disruptions, and the routing protocols quickly adjust to the disruption and route around the disruption in the network. Some examples of attacks against peering sessions are provided in the section “[Attacking a Routing System](#).”

Falsifying Routing Information

The primary protection against the injection of false routing information from outside the network is Message Digest Algorithm 5 (MD5). With pervasive use of MD5 in your network, the only way that false routing information can be injected is through a compromised peering session, which would require access to a router.

A more subtle class of attacks seeks to falsify the information carried within the routing protocol, including the topology and reachability information. The purpose of this falsification is to misdirect traffic. Falsified routing information generally might cause a denial of service, falsify information being carried between two systems, or cause traffic to follow a path the traffic would not normally follow. The next sections contain examples of attacks that misdirect traffic.

Abusing Routing Stability Features to Reduce Network Availability

The abuse of routing stability features requires a thorough knowledge of routing protocol semantics and the implementation of the routing protocol that has been deployed. (The abuse of a routing stability feature is a class of attacks within the category of falsifying routing information, but this type of attack is described in this section for clarity.)

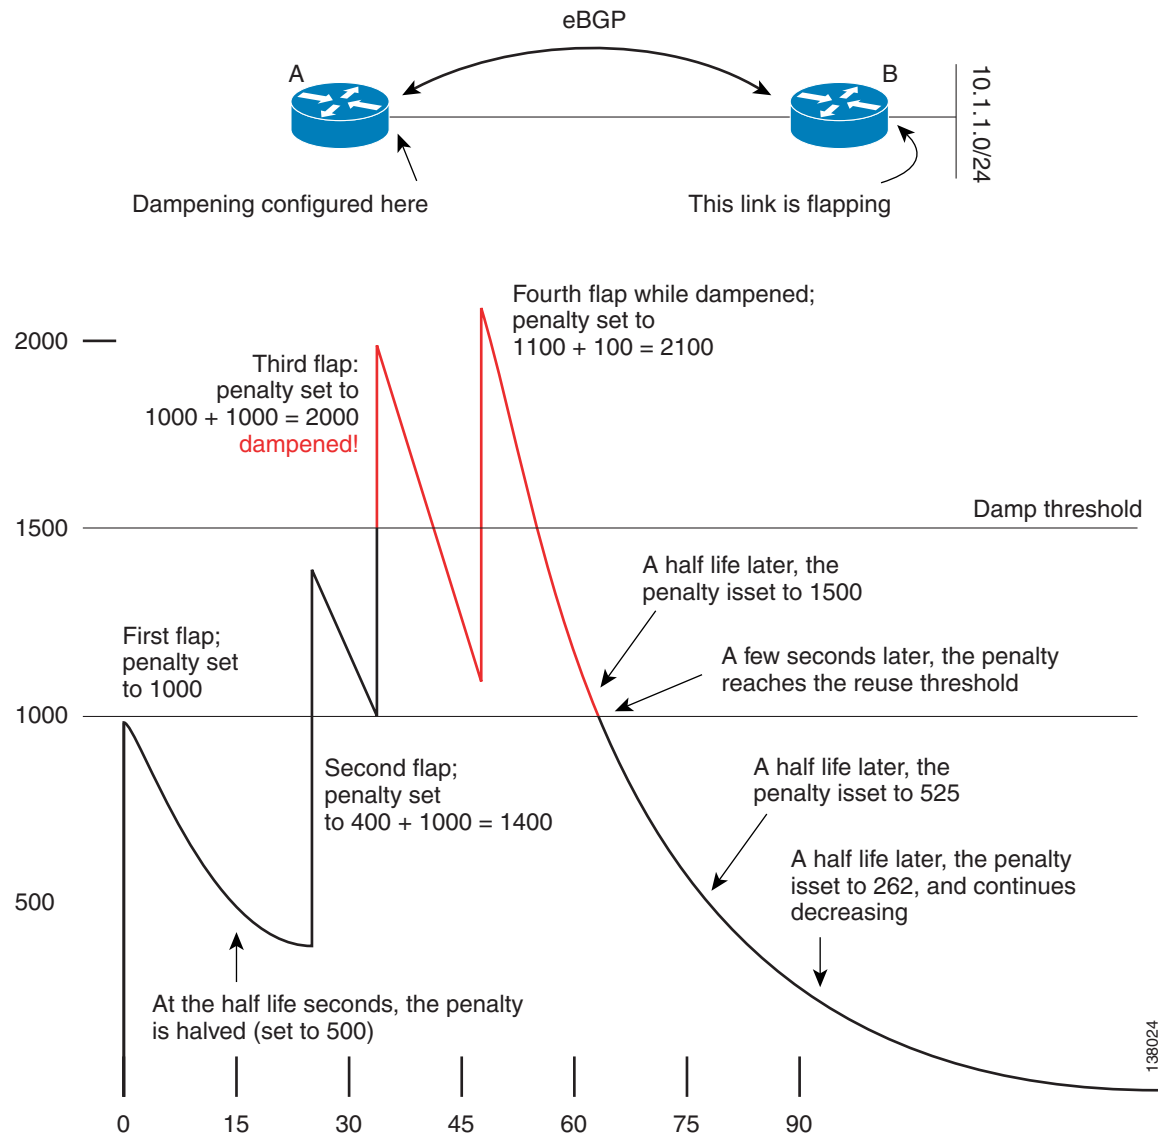
Most routing protocols provide some form of stabilization to prevent rapid changes in the network topology from overwhelming the processors, memory, and other resources on the devices that are running the routing protocol. An attacker might be able to exploit stabilizing features.

An abuse of routing stabilization normally falls within the realm of DoS attacks because it either removes legitimate routing information from the routing tables (forcing a router to understate its connectivity) or at least slows the convergence of the network so that any real changes in the network topology cause all current sessions across the network to drop. Two possible attacks using routing protocol stability features to reduce network availability are discussed in this section.

Forcing BGP Peer Dampening by Injecting Flapping Routing Information

Border Gateway Protocol (BGP) route dampening (or damping) can reduce the rate of change in the Internet routing tables. In general, nearly all BGP sessions between ISPs are aggressively dampened by the ISP to prevent flapping advertisements of routing information from wasting resources on the provider network. Figure 4 illustrates the concept of route dampening in BGP.

Figure 4 Peer Dampening Through Injected Flapping Data



The route dampening specification consists of a penalty, damp threshold, reuse threshold, and half-life. For example, a dampening specification could be:

- Penalty of 1000
- Dampening threshold of 1500
- Reuse threshold of 1000
- Half-life of 15 seconds

The following series of events explains the route dampening operation (see [Figure 4](#)):

- The first time 10.1.1.0/24 is withdrawn by Router B and advertised again, Router A sets the penalty on this route to 1000.
- Fifteen seconds after this event, the penalty is reduced by half to 500.
- The second time 10.1.1.0/24 is withdrawn by Router B and advertised again, Router A adds 1000 to the current penalty for the route, making it 1400.
- The penalty resumes its decrease until it reaches 1000.
- When Router B withdraws and advertises 10.1.1.0/24 a third time, Router A adds 1000 to the current penalty for the route, making it 2000.
- When the penalty jumps above 1500, Router A removes the route from the local routing table and stops advertising it to its peers. This action makes 10.1.1.0/24 unreachable through this path.
- When the penalty on the route decreases to 1100, Router B again withdraws and advertises 10.1.1.0/24 once more. This event causes Router A to add 100 to the penalty, making the total penalty 2100. The consequence is that if the route is being repeatedly withdrawn and advertised to Router A, 10.1.1.0/24 remains unreachable.
- When the penalty falls below 1000, Router A again places the route to 10.1.1.0/24 in its local routing table and begins advertising the destination as reachable to its peers.

An attacker could exploit the dampening mechanism in Router A by making a given destination route appear to be flapping. If an attacker could fake a set of withdrawals and updates that are advertising a given route and cause Router A to respond to these state changes when they are transmitted, the attacker could cause the ISP edge router to dampen the route for a long time. Dampening the route in this way would make the destination unreachable. This attack is worse than simply causing Router B to stop advertising 10.1.1.0/24 because the effects of the attack carry forward for a significant time after the attacker has left the network (possibly leaving no trace of how the attack was accomplished). Subsequently, the only way to make 10.1.1.0/24 reachable again is for the ISP to take corrective action, and this correction could take more time than what the client-business would find acceptable.

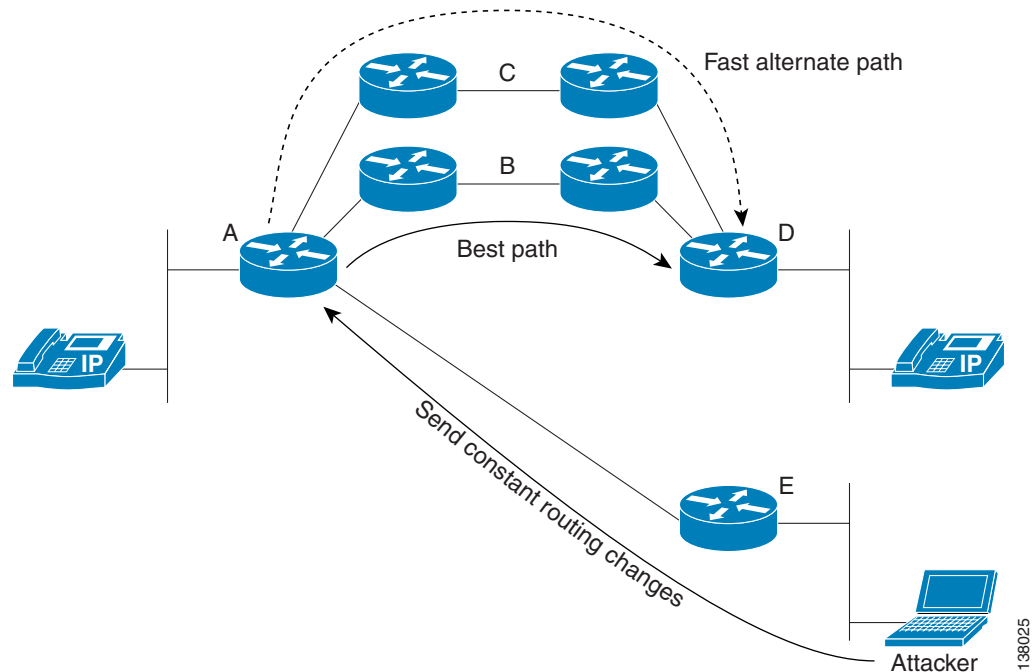
Forcing the Routing Protocol to Converge More Slowly by Injecting Flapping Routing Information

Routing protocols usually converge in 7 to 10 seconds in most networks, but this convergence is too slow for the newer applications used in networks. For example, if a link fails and traffic rerouting takes 7 to 10 seconds, voice calls are dropped, video streams are seriously degraded, and many other types of sessions fail. Improving the speed of routing convergence in a network (and within large internetworks, such as the Internet) has been an area of development by routing protocol designers.

The main challenge for a routing protocol that is attempting to converge quickly is information overload, so the question is: How much information is necessary to overwhelm the smallest or slowest devices in the network? If a network changes so quickly that it becomes overwhelmed, the routing systems quickly becomes unstable and could eventually fall into a condition called a *network melt*. In a network melt, the network does not converge again without human intervention—even when the original source of the information overload is removed from the network. Tools exist that can help prevent a network melt.

A primary tool for preventing information overload from causing a network to melt, but which still lets the network react quickly to a low or moderate number of topology or reachability changes, is the *exponential backoff*. Exponential backoff allows the routers to converge quickly when a single topology change occurs in the network, but it slows down the convergence as more and more events occur. Unfortunately, an attacker could take advantage of exponential backoff characteristics to force the network to converge slowly on a constant basis. Excessively slow convergence makes the network fail to meet the requirements of the applications running in it. See [Figure 5](#) for an illustration of this attack.

Figure 5 Attacking Through the Exponential Backoff Stability Feature



In a network with IP phones (Figure 5), two IP phones normally communicate over the path seen as {A,B,D} but have an alternate path of {A,C,D}. The routing protocol has been tuned to switch to the alternate path {A,C,D} very quickly by relying on exponential backoff. However, if the exponential backoff is forced to its longest possible wait before convergence is allowed, the Voice-over-IP (VoIP) sessions across the network fail.

With knowledge of the exponential backoff characteristics, an attacker sitting behind Router E could send constantly changing routing information to Router A. This trick forces the exponential backoff algorithm on Router A to maintain the highest (longest) timers possible because constant changes appear in the network topology. When the {A,B,D} link fails, the convergence time is long enough to force the call between the two IP phones in the network to fail.

This attack could be a coordinated attack to take down the primary path, but it could also be a *chance attack*. A chance attack is kept up for a long time with the intention that some link eventually fails and causes a service outage across the network.

Attacking a Routing System

This section describes some techniques an attacker could use to disrupt peering or to inject false routing information. It details specific actions an attacker could take to accomplish either goal, including:

- Flooding a port (thwarted in the iFIB)
- Using semantics of the protocol to attack the protocol (which are thwarted by GTSM for BGP only in the current release)
- Compromising a legitimate member of the routing domain


Port Flooding

Several well-known port-flooding attacks are commonly used against hosts, including:

- Various denial of service (DoS) attacks, which include a distributed denial of service (DDoS) attack.
- Synchronization (SYN) flooding.
- Other sorts of attacks that overrun the buffer or other resources on a device, causing the device to stop responding to legitimate traffic. These attacks could be used against routers to disrupt peering between two routers.

In the network shown in [Figure 6](#), an attacker is flooding Router A with TCP SYN packets. If Router A is not protected against it, this attack can eventually exhaust all TCP resources on the router and cause it to refuse new BGP connections or possibly to drop existing connections. A successful SYN attack could affect Internet connectivity for any network that depends on the routing information transferred across BGP sessions with Router A. The impact of this attack manifests only during the attack. After the attack ends, the BGP sessions are rebuilt, and routing information is exchanged, so packets can again be routed across the link.

Figure 6 A TCP SYN Flood Attack on BGP Peering




In another case, an attacker might try to flood the port that a protocol is using with the goal of consuming the processing resources and preventing the protocol from operating properly. In this case, the attacker could simply flood TCP port 179 with a large number of packets to starve out the TCP connection between the two routers.

Protocol Semantics Peering Attacks

Abusing the semantics of a protocol itself is another way to attack the routing protocol peering sessions. With an understanding of the routing protocol peering semantics, an attacker can cause a routing protocol peering session to be reset and thus disrupt the network services. Attacking a peering session between routers by using an attack against the routing protocol semantics is a DoS attack. This section illustrates three attacks that are possible against a peering session.

BGP uses TCP as its transport, so BGP is vulnerable to any attacks that causes the TCP session to fail. Because either end of a TCP session can reset the session by issuing a TCP reset, injection of a false TCP RST into the session can reset the BGP peering session. See [Figure 7](#). This attack is an example of a *single packet attack*. (A flood is an example of a multipacket attack.)


Figure 7 Breaking a BGP Peering Session Through a TCP Reset



To reset the BGP peering session in Figure 7, the attacker first listens to the TCP session between Router A and Router C and then, by using the information gathered, generates a TCP reset that *seems* to originate from Router C. Theoretically, an attacker could generate a TCP reset without first listening in on the TCP session between Router A and Router C, but this method is very difficult.

Similar attacks can be made against various interior gateway protocols (IGPs). Generally, these attacks are directed against the neighbor discovery and maintenance state machines. See Figure 8.

Figure 8 Breaking IGP Peering Through False Packet Injection



An attacker might transmit a packet that resets the neighbor state between Router A and Router B from any place in the network and do so with no prior knowledge of the state between these two routers. For the following descriptions of neighbor state resets, refer to Figure 8:

- If Router A and Router B are exchanging routing information by using Open Shortest Path First (OSPF), an attacker could unicast a packet to Router A. The attacker could send an empty hello packet that appears to have originated at Router B. (In OSPF, two-way connectivity is guaranteed by each router that lists the neighbors from which it has received hellos in the hello packets. Thus, to guarantee two-way connectivity, Router B would list the router ID of Router A in each hello it transmits.) If Router A receives a hello packet with an empty neighbor list, the router responds by resetting the neighbor adjacency based on false information. The false information is that the empty neighbor list shows that the two-way connectivity with the neighbor that sourced the hello is gone.
- If Router A and Router B are exchanging routing information by using the Enhanced Interior Gateway Routing Protocol (EIGRP), an attacker could transmit a unicast packet to Router A that appears to be an empty update with the initialization bit set. (When a router that is running EIGRP attempts to restart its neighbor relationship with another router, it sends an empty update with the initialization bit set.) The active initialization bit directs the other router to restart the relationship.

Neither of these reset-type attacks requires a knowledge of the state between Router A and Router B. The results of the attack are temporary, as are all attacks against peering relationships, so the damage to the network would probably be small. However, repeated attacks could bring routing across a link down (or the network down) until the source of the attack is discovered and blocked.

Compromising a Legitimate Member of the Routing Domain

A straightforward way to attack the routing system is to attack the routers running the routing protocols, gaining access to the routers themselves, and using that access to compromise the routing system in some way. For this scenario, the attacker might inject false information into the routing system. Successfully injecting false information could let the attacker accomplish any of the goals just listed. Access to a router could let an attacker bring peering relationships down between the compromised router and its peers, or it could let the attacker use the compromised router as a base for other types of peering disruption attacks.

Masquerading as a Member of the Routing Domain

An attacker does not necessarily need to take over a valid member of the routing domain by attacking a router within the network. An attacker could also inject false routing information by simply attaching a device that runs routing protocols to the network through an open port. With this access, the attacker would directly inject false routing information into the network. (For this reason, all ports are closed by default in Cisco IOS XR software.) By attaching a device to the network, the attacks that can be launched against a network are similar to the attacks that are possible by taking over a legitimate member of the routing domain.

Attack Detection and Response

This module introduces tools that you can use to detect an attack and help you respond to it. These tools have a wide range of uses in network operation, but they also are applicable to security.

The tools introduced in this module are:

- [Always-Available CLI Commands](#)
- [NetFlow](#)
- [Watchdog System Monitor](#)
- [Fault Manager](#)
- [Alarms](#)

Cisco IOS XR software can detect and respond to attacks at many points throughout the routing system. If an attack is intended to overwhelm the route processor (RP) to the point of making the router useless, the attack would have to pass many barriers. If an attack were to bypass one barrier, its chances of getting past the next barrier are—by design—fewer. If an attack can get closer to the CPU on the RP, its chances of success are reduced with each security barrier. If an attack reaches the RP, the uKernel is highly insulated by its own architecture.

A router that is running Cisco IOS XR software continues to run during an attack, and the set of always-available CLI commands are designed to help you bring an attack to an end and restore the router to full operation. These always-available CLI commands, along with support for software module updates (SMUs), make the system highly resilient.

Always-Available CLI Commands

The commands that are always available during an attack (if the management connection is up) come in two levels of importance. For the higher priority:

- The higher-priority commands let you debug or recover from an out-of-resources (OOR) attack. These commands relate to the operating system and allow you to:
 - Show processes
 - Show memory
 - Shut down or restart a process

The next priority of always-available commands apply to the interfaces and the IP layer. They can:

- Show an interface
- Show the layer

- Shut an interface
- Show ACLs in general or show a particular ACL

NetFlow

NetFlow is a built-in service that monitors and exports data for sampled IP traffic flows. For every flow on a router, a Netflow sampling contains the following information:

- Source and destination IP address
- Source and destination TCP/UDP ports
- Port utilization numbers
- Packet counts and bytes per packet
- Start time and stop time of data-gathering events and sampling windows
- Type of service (TOS)
- Type of protocol
- TCP flags
- In the category of routing and peering:
 - Next-hop address
 - Source AS number
 - Destination AS number
 - Source prefix mask
 - Destination prefix mask

Netflow output can be used for many purposes, but this section introduces only the insights into security concerns that NetFlow can give. Broadly speaking, all the routers with NetFlow are acting as sensors that provide you with data to help you:

- Detect and classify security incidents
- Understand the impact of network changes and services
- Improve network usage and application performance
- Reduce IP service and application costs
- Optimize network costs

For descriptions of the NetFlow configuration tasks and CLI commands, see the following modules:

- *Configuring NetFlow on Cisco IOS XR Software* module in *Cisco IOS XR Interface and Hardware Component Configuration Guide*
- *NetFlow Commands on Cisco IOS XR Software* module in *Cisco IOS XR Interface and Hardware Component Command Reference*

Key concepts in the security application of NetFlow are those of an anomaly and the characterization of an anomaly as an attack. An anomaly is an event or condition in the network that is recognized to be a statistical abnormality when compared to known traffic patterns. Patterns are based on established profiles and the all-important traffic baselines that are created during the phase we call *preparation*.

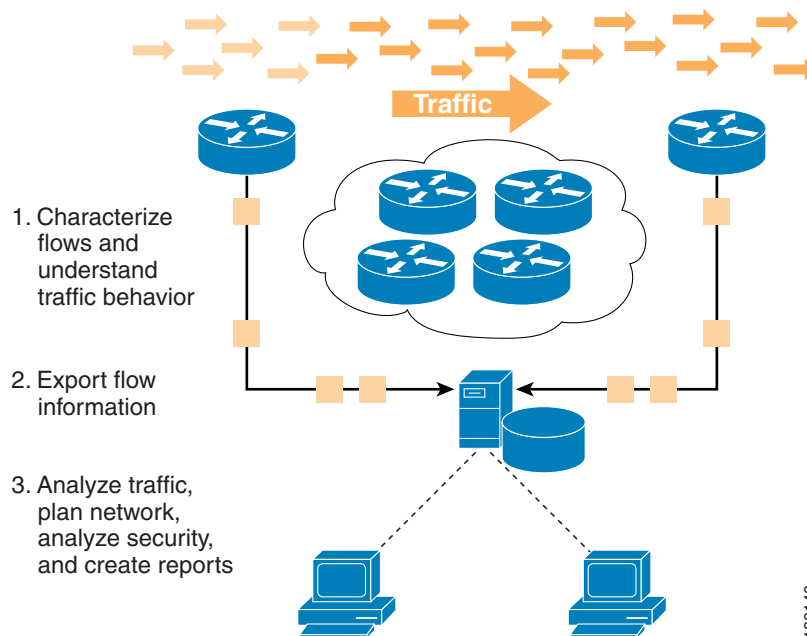
An important distinction between traditional telemetry gathering, such as that provided by Simple Network Management Protocol (SNMP), and that of NetFlow is that NetFlow gives you the ability to characterize a flow. In characterizing a flow, you can spot deviations from the intended configuration on, for example, an interface, but also in the sense of behavioral comparisons, such as peak flows at different times of the day or in comparison to other interfaces. Based on your knowledge of the network's baseline, you might characterize an anomaly as an attack. From another perspective, NetFlow also differs from intrusion detection systems (IDS) and packet capture. Packet capture gives a much more granular view of IP traffic, but the NetFlow security role really is to characterize the flow for the purpose of identifying an attack. The difference between packet capture and NetFlow telemetry can be illustrated as follows:

- Packet capture is like a wiretap in which the actual words of a conversation are recorded.
- Netflow is like a telephone bill, which shows both ends of the call, the date and time of the call, the duration of the call, and so on.

If necessary, upon determining that an attack is taking or has taken place, packet capture could be used to look into the contents of the attack packets. When the consideration is high network availability, however, the point is to prevent an attack through good planning or to end an attack quickly.

The connections between NetFlow data flow, data acquisition, and information access appear in [Figure 9](#). This figure includes in points 1, 2, and 3 the actions you can take using NetFlow.

Figure 9 NetFlow Data Acquisition Overview

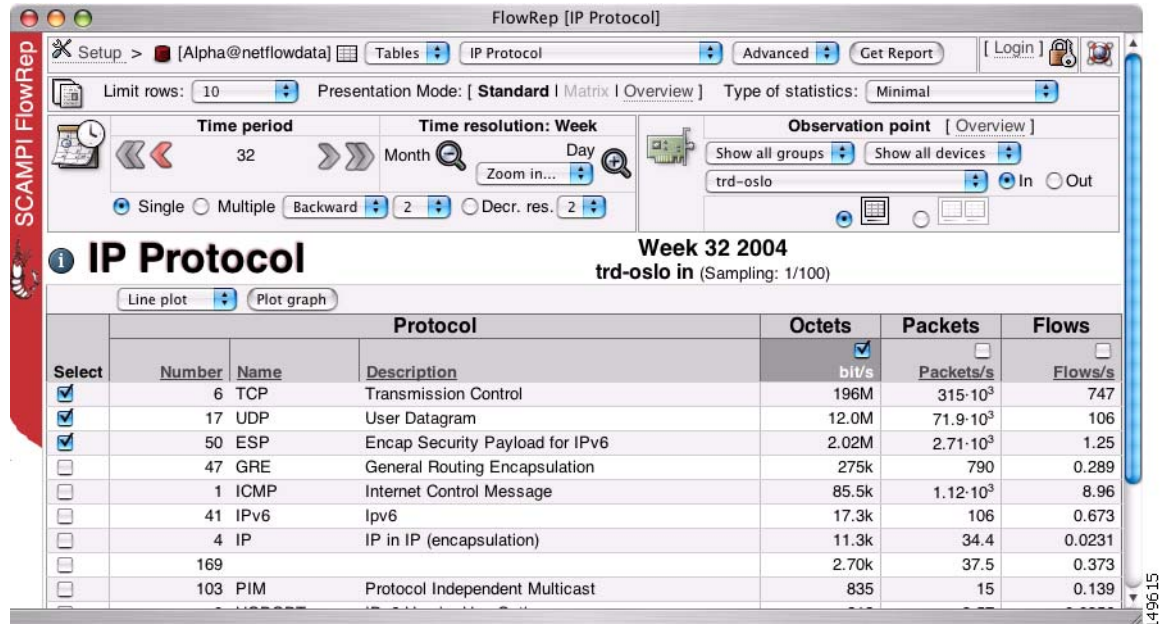


The amount of information that NetFlow generates can become very large, so NetFlow lets you control the parameters for the information that it exports. For example, you can configure the amount of time between exports or the high and low watermarks of packets on a port that would appear NetFlow output.

You can view the sampled information on the CLI by using the **show flow monitor** command or through a wide variety of front-end products designed to present the information in graphical and textual formats. Cisco offers various tools for the high-level viewing of information, such as the Cisco CS-MARS or NetQoS ReporterAnalyser. Open-source products are also readily available. These products include

Ethereal, Stager, and CU Flow. In addition, Cisco Systems has partnerships with other vendors, such as Arbor Networks, to work with the Arbor Peakflow tools. A sample screen from Stager appears in Figure 10.

Figure 10 Stager Display



NetFlow Version 9 Template Format

The basic output of NetFlow is a flow record. The record format that is best suited for security usage is supported by NetFlow Version 9. The most significant difference of the NetFlow Version 9 format is that it is template-based.

With the template format, new features can be added to NetFlow more quickly and without breaking current implementations. Third-party business partners who produce applications to collect or display data from NetFlow do not need to recompile applications each time a new NetFlow feature is added. Instead, our partners can use an external data file that documents the known template formats.

Key concepts in the security application of NetFlow are those of an anomaly and the characterization of an anomaly as an attack. An anomaly is an event or condition in the network that is recognized as a statistical abnormality when compared to typical traffic patterns. (Patterns are based on previously collected profiles and the all-important baselines.) You can identify anomalies through NetFlow after having created a detailed baseline of traffic flows during the phase we call preparation.

Watchdog System Monitor

The watchdog system monitor (wdsysmon) checks resource usage and reports to the Fault Manager or other destinations available to you. The role of wdsysmon for Fault Manager is crucial. For information on Fault Manager, see “[Fault Manager](#)” (a section that also contains references to the appropriate configuration guide and command reference).

Examples of the information that wdsysmon tracks are:

- High CPU usage
- Low memory levels
- Reserved packet memory pools (for ensuring that critical services, such as routing Hellos, are not getting lost during packet resource attacks)
- Kernel limitations on resources, such as the user-configurable number of processes that can be spawned
- Memory thresholds for applications, so that no application can exceed a user-configurable amount of memory usage that skyrockets because of an attack

You can use Fault Manager to respond to these resource issues and alarms of possible attack. Through FM scripts, you can set up automatic mitigation of attacks. The wdsysmon facility monitors resource utilization activity and prioritizes messages to send to the:

- Fault manager
- Console terminal
- Local (syslog) local buffer
- Syslog server

Fault Manager

The Cisco IOS XR Fault Manager (FM) is the central clearinghouse for events detected by any portion of the RP failover services. Among the large numbers of events, many events can pertain to security. For a detailed description of Fault Manager, see the *Configuring and Managing Fault Management Policies on Cisco IOS XR Software* module in *Cisco IOS XR System Management Configuration Guide*.

Essentially, FM manages system events. An event can be any significant occurrence within the system. FM runs on the RP and is responsible for:

- Detecting fault events (not limited to security breaches)
- Recovering from faults
- Managing process reliability statistics

FM *events* are notifications that something significant has occurred within the system. A small list of examples are:

- Operating or performance statistics that go outside allowable values (for example, free memory dropping below a critical threshold)
- Online insertion or removal (OIR) of a line card (LC) or a modular services card (MSC)
- Termination of a process

You can specify corrective actions in FM based on the current state of the system. For example, you can use FM to request notification by e-mail when a disk drive is running low on space.

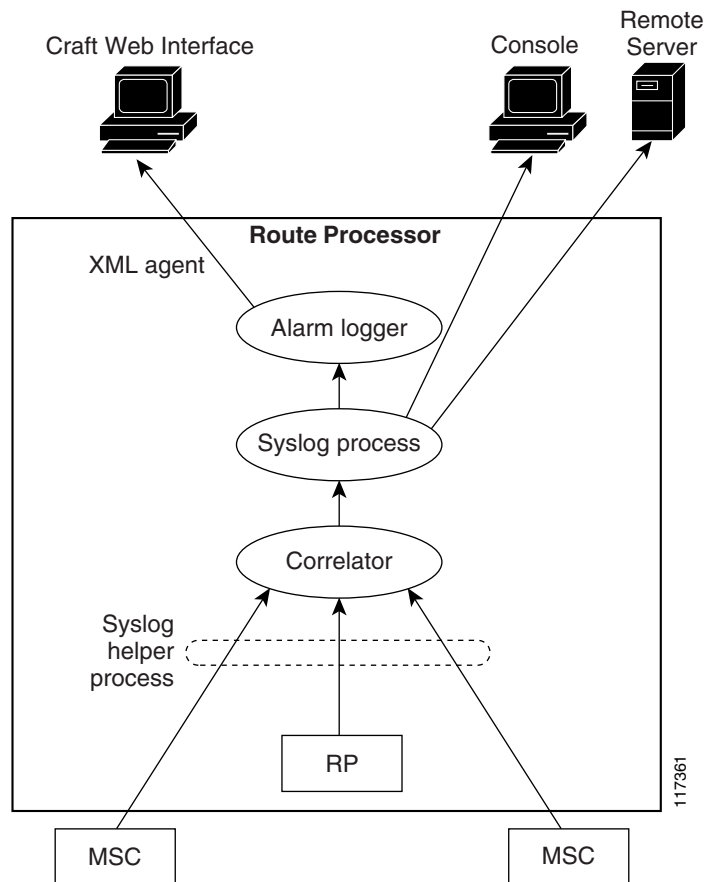
FM relies on software agents called *fault detectors* to notify it when the specified system events occur. After FM has detected an event, it can initiate corrective action. An action is prescribed in a routine called a *policy*. When the specified event is detected, FM implements the policy.

System Event Detection

FM uses fault detectors that it has provided to the syslog facility to detect certain system events. The fault detectors monitor the system for events and use a pattern match with the syslog messages. FM also relies on a timer-fault detector to detect when a certain date and time occur.

The FM *correlator* generates alerts according to the rules that you specify. The correlator ties information together to determine relations between events. For example, a rule could specify that two patterns together might indicate an attack—if event A is true and event B is true, a certain type of attack might be happening, so send message C. For a more specific example, if a link is flapping but the link has no hardware problem, then an attack might be under way. The syslog helper determines where the conclusion or distilled message should go. To see where the correlator and the syslog process fit in the larger context of fault and alarm monitoring, see [Figure 11](#).

Figure 11 Correlator and Syslog



System Event Processing

Upon receiving event notification, Fault Manager performs some preliminary tasks before it transmits messages or takes other actions. When FM receives an event notification, it:

- Checks for established policy handlers:
 - If a policy handler exists, FM initiates callback routines (*Fault Manager handlers*) or runs Tool Command Language (Tcl) scripts (*Fault Manager scripts*) that implement policies. The policies can include built-in FM actions.
 - If a policy handler does not exist, FM does nothing.
- Notifies the processes that have *subscribed* for event notification
- Records reliability data for each process in the system
- Provides access to FM-maintained system information through an API

FM can send you the messages that you have configured. FM communicates with wdsysmon, so it can get messages from wdsysmon and generate specific syslog messages for viewing. If you want the message to go to networking personnel in the form of a page or as e-mail, an external device should be used because the router should not be burdened with e-mail or a paging service.

Messaging can also be tied to an EMS system, such as HP OpenView. Even in an environment in which syslog messages are coming from all the routers in the network, you can configure HP OpenView to respond to a specific bit of information in a particular message and in a way that you specify.

You can configure FM to generate specific syslog messages, and thereafter the syslog server checks for those messages and alerts you when it detects a message that you have directed it to flag. You could also configure automatic responses to certain issues. For example, if a peer seems to be sending too many routes—exceeding a limit of 1000 that you have set—FM could shut down the peer for an hour. Also, you could program a response when the free space on a disk drops below a certain threshold.

Alarms

This section introduces concepts related to:

- Alarm log correlation
- Monitoring of the alarm logs and the correlated event records

For a detailed description of the subjects in this section and other related subjects, see the *Implementing and Monitoring Alarm Logs and Logging Correlation on Cisco IOS XR Software* module in *Cisco IOS XR System Management Configuration Guide*.

Alarm log correlation extends the system logging facility to include the ability to group and filter messages generated by various applications and system servers and isolate root messages on the router.

Alarms are generated during attacks. In the current release, you can check these alarms through wdsysmon. Applicable alarms include:

- Authentication alarms, which include login failures
- Resource threshold alarms for CPU, packet memory, kernel resources, and processes
- Denial of service (DoS), distributed denial of service (DDoS), and various intrusion alarms

Alarm Logging and Debugging Event Management System

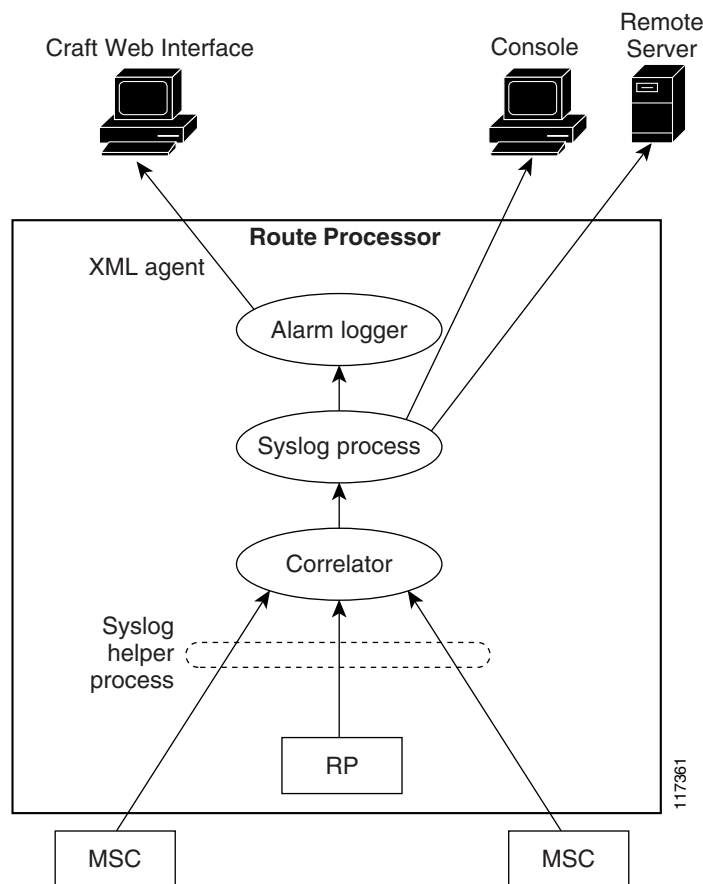
Alarm log correlation is one aspect of the Cisco IOS XR software Alarm Logging and Debugging Event Management System (ALDEMS), which:

- Monitors and stores alarm messages that are generated by system servers and applications
- Correlates alarm messages that are generated due to a single root cause

ALDEMS goes beyond the basic logging and monitoring functionality of Cisco IOS XR software by providing the level of alarm and event management necessary for a highly distributed system with hundreds of LCs and MSCs and thousands of interfaces. The Cisco IOS XR software achieves this level of alarm and event management by distributing logging applications across the nodes on the system.

For an illustration of the relationship among the components that constitute ALDEMS, see [Figure 12](#).

Figure 12 ALDEMS Component Communications



Correlator

The correlator receives messages from system logging (syslog) helper processes that are distributed across the nodes on the router. The correlator forwards syslog messages to the syslog process. If a logging correlation rule is configured, the correlator captures messages. It begins with the first occurrence of a

message from the set of messages specified in the correlation rule and continues for the interval specified in the timeout interval of the correlation rule. The logging correlator buffer resides inside the correlator and stores all correlated messages captured after a root message is received.

System Logging Process

The syslog helper gathers the syslog messages and distributes them across the nodes on the system. The system logging process controls the distribution of logging messages to all the possible destinations, such as the system logging buffer, the console, terminal lines, or a syslog server. The choices for these destinations are configurable but depend on the network device configuration.

By default, the Cisco CRS-1 router and Cisco XR 12000 Series Router are configured to send the system logging messages to a system logging (syslog) process.

Alarm Logger

The alarm logger is the final destination *on the router* for system logging messages that are generated on the router. The alarm logger stores alarm messages in the logging events buffer. This buffer is circular, so it writes over the oldest messages when the buffer becomes full. In addition, alarms are prioritized in the logging events buffer so that, when necessary, the buffer writes over messages in a particular order. You can examine messages in the logging events buffer to locate records that match a set of specific criteria.

Logging Correlation

The logging correlation facility groups the messages that are generated because of a shared, root cause. The correlator receives messages from the system logging helper processes that are distributed across each node on the system. With your correlation rules configured, the correlator scans the messages for the target fields that match the specifications in your correlation rules.

You can use logging correlation to isolate the most significant root messages for events affecting system security in particular, but you can also use it to monitor system performance in general. You can retrieve all correlated messages from the logging correlator buffer for display.

Correlation Rules

You can configure correlation rules to isolate root messages that might generate system alarms. Correlation rules prevent unnecessary overhead on ALDEMS caused by the accumulation of unnecessary messages. Each correlation rule depends on a message identification (ID). The message ID consists of a message category, message group name, and message code. The correlator scans the messages for occurrences of a particular message that matches a message ID.

The first message in the correlation rule configuration is considered the root message. If the correlator receives a root message, it stores the message in the logging correlator buffer and also forwards it to the syslog process on the RP. From there, the syslog process forwards the root message to the alarm logger to be stored in the logging events buffer. From the syslog process, the root message can also be forwarded to a destination, such as a console, a remote terminal, a remote server, the FM system, and SNMP.

If a message matches multiple correlation rules, all matching rules apply, and the message becomes a part of all matching correlation queues in the logging correlator buffer.

The message fields that define a message in a logging correlation rule are the:

- Message category
- Message group
- Message code

Root Message and Correlated Messages

When a correlation rule is applied, the timeout for the rule begins with the first occurrence of a message from the set of messages specified in a correlation rule. The first message (with category, group, and code triplet) configured in a correlation rule defines the root message. A root message is always forwarded to the logging events buffer. See the “[Correlation Rules](#)” section to learn how the root message is forward and stored.

When the timeout for the correlation rule expires, leaf messages that do not have an associated root message are forwarded to the logging events buffer. All subsequent messages from the message set in the correlation rule are stored in the logging correlation buffer.

Alarm Severity Level and Filtering

You can set up filters so that information is displayed based on the alarm severity. The alarm filter display indicates the severity level settings that report alarms, the number of records, and the current and maximum log sizes. Alarms can be filtered according to the severity level shown in [Table 3](#).

Table 3 Alarm Severity Levels for Event Logging

Severity Level	System Condition
0	Emergencies
1	Alerts
2	Critical condition
3	Errors
4	Warnings
5	Notifications
6	Informational entry

Implementing Network Security Mechanisms

Routers that run Cisco IOS XR software are designed to be secure by default—to arrive at the site ready to run securely. Nevertheless, this module describes issues you should consider for securing the network. It describes the implementation of security and has three stages of discussion: default secure states, general recommendations for improving security that apply to most IP networks, and more specific kinds of protection. For information on detecting attacks, see the [Attack Detection and Response](#) module. For examples of attacks, see [Examples of Network Attacks](#) module. The sections in this module are:

- [The Six Phases of a Security Strategy](#)
- [Default Settings](#)
 - [Password Encryption Service](#)
 - [Time Stamp on Debug](#)
 - [Disable IP Gratuitous ARPs](#)
 - [TCP Synwait Time](#)
 - [Disable IP Redirects](#)
 - [Disable IP Proxy ARP](#)
 - [Disable IP Unreachables \(Host\)](#)
 - [Disable IP Mask Reply](#)
 - [Disable IP Unreachables on Null Interface](#)
 - [Enable uRPF on Outside Interfaces](#)
 - [SYN Cache Default Value](#)
 - [SYN Rate Limit Default](#)
- [General Recommendations](#)
 - [Password Management](#)
 - [Controlling Interactive Access](#)
 - [Management Services Without Interactive Login](#)
 - [SNMP](#)
 - [Management and Interactive Access over the Internet and Other Untrusted Networks](#)
- [Protecting the Legitimacy of the Routing Domain](#)
 - [Protecting Routers from Being Compromised](#)
 - [Protecting Routing Information on the Wire](#)

- [Protecting Against Illegitimate Devices Joining the Routing Domain](#)
- [MD5 and Peer Authentication](#)
- [TTL-Based Peering Session Protection \(BGP TTL Security Hack\)](#)
- [Protecting the Routing Information](#)

The Six Phases of a Security Strategy

This section describes the six phases of security actions that we recommend as a part of best practices.

Phase 1: Preparation

A supremely important factor in a security strategy is preparation. For example, network analysis to generate the network baseline of normal traffic patterns and subsequent planning for the use of Netflow and analytic software ahead of an attack are vital steps in the preparation phase.

Phase 2: Identification

Equipped with a baseline of network traffic patterns, you can use a tool such as Arbor Peakflow to look through Netflow output for historical precedents that are based on traffic, bits per second, IP addresses, and so on. Knowledge of traffic patterns enables you to identify the nature of suspicious traffic in bits per second, duration, sources, and so on.

Phase 3: Classification

You can classify and scope the threat by using information from Netflow and Arbor Peakflow. The actual threat posed by anomalous traffic to network availability can vary, so in the classification phase you determine the seriousness of the threat. For this phase, you can use Arbor Peakflow or other tools to classify the threat.

Phase 4: Trace-back

Trace-back is the phase in which you identify the sources of an attack. With proper planning and the available tools, you can identify even indirect vectors, such as VPNs or laptop computers.

Phase 5: Reaction

After identifying potential vectors, you can use this information for effective action. For example, for a very serious attack you could close network access to the attack traffic by placing access control lists (ACLs) on the inbound and outbound traffic at all points of presence on the network and at strategic places in the WAN backbone.

Phase 6: Post-Mortem

After an attack has been thwarted, continued monitoring might be appropriate to confirm that the problem has been eradicated. Analysis of the network baseline might also be appropriate.

Default Settings

With Cisco IOS XR software, our approach has been to disable services by default, so most services need to be explicitly enabled. We call this approach *default secure*. For example, default secure means that you need to enable the address-family ipv4-unicast in some routing protocols (IS-IS and BGP); these items are disabled by default. It also means that external services and services that listen to TCP/UDP ports and IP protocol numbers (SSH and so on) are off by default.

This section contains the default settings for security-related components. In [Table 4](#), the default state, or sometimes the default value, is shown for a variety of features. The sections that follow this table explain the security-related reason for the default state or value.

Table 4 *Default Settings*

Feature	Default State or Value
Disable Cisco Discovery Protocol (CDP)	Disabled
Disable IP source route	Disabled
Password encryption service	Always on
Enable TCP keepalive	Enabled
Time stamp on debug	Enabled
Disable IP gratuitous ARPs	Disabled
Set TCP Synwait time	30 seconds
Disable IP redirects	On
Disable IP proxy ARPs	Off
Disable IP directed broadcast	Off
Disable IP unreachable (host)	On
Disable IP mask reply	Off
Disable IP unreachable on Null interface	Always off
Enable Unicast Reverse Path Forward (uRPF) on outside interface	Disabled
SYN cache	4000 packets
SYN rate limit	100 packets per second

Password Encryption Service

Passwords are encrypted by default. Cisco IOS XR software encrypts the following:

- Passwords
- Challenge Handshake Authentication Protocol (CHAP) secrets
- Similar data that is saved in its configuration file

These mechanisms are useful for preventing a casual observer from reading passwords by, for example, surreptitiously looking at the screen.

Time Stamp on Debug

We recommend that you keep time stamps enabled on all debug and log messages whenever possible. Time stamps on debug and log messages indicate the time and date that the message was generated. Knowing the time that messages are generated is an important tool in diagnosing potential attacks.

Disable IP Gratuitous ARPs

A *gratuitous* Address Resolution Protocol (ARP) is an ARP broadcast in which the source and destination MAC addresses are the same. Primarily, a host uses gratuitous ARPs to inform a network about its IP address. A spoofed gratuitous ARP message can cause network mapping information to be stored incorrectly and subsequently result in network malfunction. We recommend that, whenever possible, you leave disabled all IP gratuitous ARP requests.

TCP Synwait Time

The TCP synwait time is a value that is useful in defeating SYN flooding attacks, a form of denial of service (DoS) attack. The router ships with a TCP synwait time of 30 seconds. We recommend that you normally set the TCP synwait time to 10 seconds, if possible.

A TCP connection initially requires a three-phase handshake to establish the connection:

- A connection request is sent by the originator.
- An acknowledgement is sent by the receiver.
- An acceptance of that acknowledgement is sent by the originator.

After this three-phase handshake is done, the connection is established and data transfer can begin. A SYN flood attack sends repeated connection requests to a host but never sends the acceptance of acknowledgements that complete the connections, creating a growing number of incomplete connections at the host. Knowing that the buffer for incomplete connections is usually smaller than the buffer for complete connections, the attacker hopes to overwhelm and disable the host. Setting the TCP synwait time to 10 seconds causes the router to shut down an incomplete connection after 10 seconds and thus to prevent the buildup of incomplete connections at the host. For example, to set the TCP synwait time to 10 seconds, enter the following command:

```
ip tcp synwait-time 10
```

Disable IP Redirects

Internet Message Control Protocol (ICMP) supports IP traffic by relaying information about paths, routes, and network conditions. When a router receives power for the first time at a new installation, the only protocol that is enabled is ICMP.

ICMP redirect messages instruct an end node to use a specific router as its path to a particular destination. In a properly functioning IP network, the following rules are followed:

- A router sends redirects only to hosts on its own local subnets.
- No end node ever sends a redirect.
- No redirect ever traverses more than one network hop.

An attacker might attempt to violate these rules. Disabling ICMP redirects prevents an attacker from successfully exploiting these rules and causes no operational impact to the network.

We recommend that you leave disabled the redirect mechanism for ICMP messages.

Disable IP Proxy ARP

Address resolution protocol (ARP) is used by the network to convert IP addresses into MAC addresses. Normally, ARP is confined to a single LAN, but a router can act as a proxy for ARP requests, making ARP queries available across multiple LAN segments. Because it breaks the LAN security barrier, proxy ARP should be used only between two LANs with an equal security level and only when necessary.

We recommend that you leave proxy ARP disabled whenever possible.

Disable IP Unreachables (Host)

The router transmits ICMP *host unreachable* messages if:

- The router receives a nonbroadcast packet that uses an unknown protocol.
- The router receives a packet that it cannot deliver to the ultimate destination because it knows of no route to the destination address.

These messages can be exploited by an attacker who is attempting to gain network mapping information, so we recommend that you leave ICMP host unreachable messages disabled whenever possible.

Disable IP Mask Reply

ICMP *mask reply messages* are sent when a network device must have the subnet mask for a particular subnetwork in the internetwork. ICMP mask reply messages are sent to the device requesting the information by devices that have the requested information. An attacker might try to exploit these messages to gain network mapping information.

We recommend that you leave ICMP mask reply messages disabled whenever possible.

Disable IP Unreachables on Null Interface

ICMP supports IP traffic by relaying information about paths, routes, and network conditions.

The router transmits ICMP *host unreachable* messages if:

- The router receives a nonbroadcast packet that uses an unknown protocol.
- The router receives a packet that it cannot deliver to the ultimate destination because it has no route to the destination address.

Because the null interface is a packet sink, two things happen:

- Packets forwarded to it are discarded.
- Host unreachable messages are never generated.

If the null interface is being used to block a DoS attack, host unreachable messages would simply flood the local network. Disabling these messages prevents this flood of messages. In addition, because all blocked packets are forwarded to the null interface, an attacker receiving host unreachable messages could use the messages to determine security policy configurations.

Enable uRPF on Outside Interfaces

The Unicast Reverse Path Forwarding (uRPF) feature checks the source address of any packet against the interface through which the packet entered the router. This source address verification defeats IP spoofing and is recommended for at least the outside interfaces.

If the input interface is not a valid path to the source address according to the routing table, the packet is dropped. To enable uRPF on an outside interface, enter the interface configuration mode for that the interface and enter the appropriate command for IPv4 or IPv6. For details on uRPF, see “[Antispoofing with Unicast Reverse Path Forwarding](#).”

SYN Cache Default Value

The default SYN cache size is 4000 packets.

SYN Rate Limit Default

The default SYN rate is 100 packets per second to thwart SYN flood attacks.

General Recommendations

This section contains basic configuration suggestions that apply almost universally in IP networks.

The Cisco CRS-1 router and Cisco 12000 Series Router have many security features, such as packet-filtering access control lists; the Cisco firewall; TCP Intercept; authentication, authorization, and accounting (AAA); and encryption. Other features, such as packet logging and quality of service (QoS), help you improve network security. These features are discussed in the following Cisco publications:

- See *About Cisco IOS XR Software Documentation* for the documents that specifically apply to Cisco IOS XR software-based routers.
- See the Documentation area of our website for other Cisco products, such as firewalls.

In the most general terms, the fundamentals (beyond the default security) that you can utilize for locking down a Cisco CRS-1 router or Cisco XR 12000 Series Router include:

- Well-planned structure for task groups and user groups.
- Extensive use of ACLs.
- Use of QoS.
- User authentication through AAA.
- Use of cryptographic algorithms on all peering sessions. Supported algorithms are:
 - MD5
 - SHA1
 - HMAC-MD5
 - HMAC-SHA1
- Logging of any configuration changes (and following through by monitoring the logs).
- One-time passwords. (The password is discarded after a single use.)
- Keeping the console from being available to people who do not specifically need access.
- Not letting customers send routes they should not and using route prefix filtering to enforce this rule.

Task groups and user groups are very important features for controlling command access and helping you protect sensitive user accounts.

A user belongs to one or more user groups, and each of those user groups has an array of tasks that the task group lets a user perform. By being a member of multiple user groups, a user has access to multiple sets of tasks. User groups and task groups are configured through the Cisco IOS XR software command set used for authentication, authorization, and accounting (AAA) services. AAA is part of the base package and so is available by default.

Authentication verifies the identity of a user. Authorization verifies that an authenticated user (or principal) has permission to perform the task—in the form of a CLI command—that he or she entered at the CLI. The accounting function logs user-sessions and can create an audit trail by recording certain user or system-generated actions.

For descriptions of the AAA commands, see the *Authentication, Authorization, and Accounting Commands on Cisco IOS XR Software* module in *Cisco IOS XR System Security Command Reference*. For descriptions of the AAA tasks that use AAA commands and other related commands, see *Cisco IOS XR System Security Configuration Guide*.

The task group and user group mechanisms help prevent misuse by employees but also from an attacker who might gain access to the router. A highly sensitive account might, for example, belong to an individual with access to hundreds of business accounts around the world. If all users in this case had the same access and one account were compromised, then all accounts would be compromised. The capabilities of a user should be configured based on what that user actually needs to do. For example, an administrator should log in at the ADMIN level only when administrative tasks are required and log out after those tasks are completed.

Password Management

Passwords and similar secrets, such as SNMP community strings, are primary defenses against unauthorized access to a router. The best way to handle most passwords is to maintain them on a TACACS+ or RADIUS authentication server. However, routers still have a locally configured password for privileged access and can also have other password information in configuration files.

The following password strategies are well-known:

- All passwords are encrypted.
- Password strings that are not based on a language should be used to thwart dictionary attacks.
- Unique passwords should be created on each network device.
- One-time passwords should be used where feasible. (These passwords are generated by a broker such as SofToken.)

Cisco IOS XR software encrypts passwords automatically for the following reason: If a password is not encrypted and that password is configured for the console TTY line, the password might be intercepted and used to gain unauthorized, privileged access. This access also pertains to a remote VTY sessions.

Controlling Interactive Access

Any person who can log in to a Cisco router can display information that casual observers should not be able to see. An unauthorized user who can log in to the router might be able to use it to launch an attack against the network. In addition, anyone who can get privileged access to the router can also reconfigure it. To prevent intrusive access, you need to control interactive logins to the router.

Although most interactive access is disabled by default, some exceptions exist. The most obvious exception is an interactive session from an asynchronous terminal (such as the console terminal) that is directly connected to a router. Another example is an integrated modem line.

Console Ports

The console port on a router that is running Cisco IOS XR software has special privileges. In particular, if a send-BREAK signal goes to the console port during the first few seconds after a reboot, the password recovery procedure can easily be used to take control of the system. An attacker could exploit this fact if that attacker can access the console port through a hardwired terminal, a modem, a terminal server, or some other network device and then interrupt power or induce a system crash. Thereafter, the attacker might be able to take control of the system, even without physical access to it or the ability to log in to it during normal operation.



Note

Any modem or network device that provides access through the console port must be secured according to a standard that is comparable to the security used for privileged access to the router. At a minimum, any console-connected device should be configured to require a password from the dialup user, and the password should be carefully managed.

General Interactive Access

With the applicable software package in place and the correct configuration, Cisco IOS XR software can support connections through:

- SSH
- SSL
- Telnet
- Non-IP-based network protocols

Although interactive Telnet access is available on both the standard Telnet TCP port (port 23) and on a variety of higher-numbered ports, we strongly recommend SSH (if the crypto package is available).

All interactive access mechanisms use the Cisco IOS XR TTY abstraction regardless of the protocol: they all involve sessions on a line of some type. Local asynchronous terminals and dialup modems use TTYs. Remote network connections use virtual TTYs (VTYs). For user access, the best way to protect a system is to apply appropriate controls on all lines. By default, interactive login is prevented on a line.

Controlling TTYs

The use of local asynchronous terminals is now uncommon, but they still exist in some installations. All such terminals should be physically secured, but even if they are secure, the router should be configured to require users on local asynchronous terminals to log in. Most TTY ports in modern routers are either connected to external modems or implemented by integrated modems. Securing these ports is even more important than securing local terminal ports. By default, no TTY lines are configured.

Controlling VTYs and Ensuring VTY Availability

A VTY should be configured to accept connections that have only the protocols that are actually needed. If your software supports an encrypted access protocol, such as SSH, we recommend that you enable only SSH and leave clear-text Telnet disabled. In addition, we advise that you restrict the IP addresses from which the VTY accepts connections. By default, no VTYs are configured.

Cisco routers support a few VTY lines. When all of the VTYs are in use, no more remote interactive connections are allowed. This limit creates the opportunity for a DoS attack. If an attacker can open remote sessions to all the VTYs on the system, a legitimate user might not be able to log in.

One way of reducing this exposure caused by the VTY limit is to configure a more restrictive IP access class on the *last* VTY in the system that is different from the other VTYs. The last VTY can be restricted to accept connections only from a single, specific administrative workstation, whereas the other VTYs can accept connections from any address in a network domain.

Another useful tactic is to configure VTY timeouts to prevent an idle session from using a VTY indefinitely. Although the effectiveness of this tactic against deliberate attacks is relatively limited, it also provides some protection against sessions that are accidentally left idle for too long. Similarly, enabling TCP keepalives on incoming connections can help to guard against both malicious attacks and orphaned sessions caused by remote system crashes.

Complete VTY protection can be provided by leaving all non-IP-based remote access protocols disabled and using IPsec encryption for all remote interactive connections to the router. (IPsec is an extra-cost option. For information on IPsec tasks and commands, see *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*.)

Warning Banners

In some jurisdictions, civil or criminal prosecution of hackers who break into a system can be easier if you provide a banner informing unauthorized users that their use is in fact unauthorized. In some jurisdictions, you might be forbidden to monitor the activities of even unauthorized users unless you have taken steps to notify them of your intent to do so. One way of providing this notification is to put it into a banner message by using the **banner login** command.

Legal notification requirements are complex and vary between different jurisdictions or situations. Even within a jurisdiction, legal opinions vary, so this issue should be discussed with your legal counsel. In cooperation with counsel, you should consider which of the following information should be in a banner:

- Notice that the system is to be logged in to or used only by authorized personnel and perhaps information about who specifically can authorize its use
- Notice that any unauthorized use of the system is unlawful and might be subject to civil penalties, criminal penalties, or both
- Notice that any use of the system can be logged or monitored without further notice and that the resulting logs can be used as evidence in court
- Specific notices that local laws might require

From a security rather than a legal point of view, your login banner should not contain any specific information about your router. You should not include the router name, model, software, or who owns it, because a miscreant could use this information.

Management Services Without Interactive Login

Many companies manage their networks by using protocols other than interactive login. The most common of these protocols used for management are Simple Network Management Protocol (SNMP) and HTTP. These protocols are disabled by default. If you enable a protocol, use the security guidance in the next sections.

SNMP

This section contains introductory information about three versions of SNMP and some tips on using SNMP. The introductory material is followed by a more detailed comparison of the security strengths of each version.

SNMP is widely used for router monitoring and frequently used for router configuration changes. Unfortunately, version 1 of SNMP uses a weak authentication scheme based on a *community string*. In effect, SNMPv1 uses a fixed password transmitted over the network without encryption. SNMPv2c also uses a community string and supports an MD5-based digest authentication scheme and allows for restricted access to various management data. SNMPv3 has the strongest security and, therefore, is the preferable version. For descriptions of SNMP versions 1, 2, and 3, see the *Implementing SNMP on Cisco IOS XR Software* module in *Cisco IOS XR System Management Configuration Guide*. This module contains vital information about SNMP implementation in Cisco IOS XR software and the tasks related to SNMP. The command reference contains the details for the commands you use to implement the tasks.

The Cisco IOS XR software supports the following versions of SNMP:

- Simple Network Management Protocol Version 1 (SNMPv1)
- Simple Network Management Protocol Version 2c (SNMPv2c)
- Simple Network Management Protocol Version 3 (SNMPv3)

SNMPv1 and SNMPv2c use a community-based form of security. The community of network managers that are able to access the agent MIB is defined by an IP address access control list and password.

Support for SNMPv2c includes a bulk retrieval mechanism and more detailed error message reporting to management stations. Bulk retrieval supports the retrieval of tables and large amounts of information, minimizing the number of round trips required. The SNMPv2c-improved error handling support includes expanded error codes that distinguish different kinds of error conditions, whereas in SNMPv1 these conditions are reported through a single error code.

SNMPv3 is a security model that supports RFC 3411 through RFC 3418. A *security model* is an authentication strategy that is set up for a user and the group in which the user resides. A *security level* is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when an SNMP packet is handled. See [Table 6](#) for a list of security levels available in each version of SNMP. The benefits of SNMPv3 are further described in “[SNMPv3 Benefits](#).”

You must configure the SNMP agent to use the version of SNMP supported by the management station. An agent can communicate with multiple managers; for this reason, you can configure Cisco IOS XR software to support communications with one management station using the SNMPv1 protocol, one using the SNMPv2c protocol, and another using SMNPv3.

If you must use SNMP versions 1 or 2, carefully choose community strings that would be very difficult to guess—not, for example, “public” or “private.” If possible, avoid using the same community strings for all network devices. Use a different string or strings for each device, or at least for each area of the network. Do not make a read-only string the same as a read-write string. If possible, periodic SNMP version 1 or 2 polling should be done with a read-only community string, and read-write strings should be used only for actual write operations.

SNMP version 1 is poorly suited to use across the public Internet because:

- Version 1 uses clear-text authentication strings.
- Most SNMP implementations send those strings repeatedly as part of periodic polling.
- Version 1 is an easily spoofable, packet-based transaction protocol.

In most networks, legitimate SNMP messages come from only specific management stations. If your network has this arrangement, you should probably use the access list name option on the **snmp-server community** command to restrict SNMP access to only the IP addresses of the management stations.

For SNMP version 2 or 3, assign a different MD5 secret value for each router.

SNMP management stations often have large databases of authentication information, such as community strings. This information might provide access to many routers and other network devices. This concentration of information makes the SNMP management station a natural target for attack, so it should have an appropriate level of security.

Comparison of SNMPv1, SNMPv2c, and SNMPv3

SNMPv1, SNMPv2c, and SNMPv3 support the following operations:

- **get-request**—Retrieves a value from a specific variable.
- **get-next-request**—Retrieves the value following the named variable; this operation is often used to retrieve variables from within a table. With this operation, an SNMP manager does not need to know the exact variable name. The SNMP manager searches sequentially to find the needed variable from within the MIB.
- **get-response**—Replies to a get-request, get-next-request, or set-request that was sent by a network management system (NMS).
- **set-request**—Stores a value in a specific variable.
- **trap**—Sends an unsolicited message from an SNMP agent to an SNMP manager after an event.

Table 5 identifies other key SNMP features supported by the SNMPv1, SNMPv2c, and SNMPv3.

Table 5 *SNMPv1, v2c, and v3 Feature Support*

Feature	SNMP v1	SNMP v2c	SNMP v3
Get-bulk operation	No	Yes	Yes
Inform operation	No	Yes (no in Cisco IOS XR software)	Yes (no in Cisco IOS XR software)
64-bit counter	No	Yes	Yes
Textual conventions	No	Yes	Yes
Authentication	No	No	Yes
Privacy (encryption)	No	No	Yes
Authorization and access controls (views)	No	No	Yes

Security Models and Levels for SNMPv1, SNMPv2, and SNMPv3

The security level determines if an SNMP message needs to be protected from disclosure and if the message needs to be authenticated. The security levels that exist within a security model are as follows:

- noAuthNoPriv—Security level that does not provide authentication or encryption.
- authNoPriv—Security level that provides authentication but does not provide encryption.
- authPriv—Security level that provides both authentication and encryption.

Three security models are available: SNMPv1, SNMPv2c, and SNMPv3. The security model combined with the security level determine the security mechanism applied when the SNMP message is processed.

The combinations of security models and levels are defined in Table 6.

Table 6 *SNMP Security Models and Levels*

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.

Table 6 *SNMP Security Models and Levels (continued)*

Model	Level	Authentication	Encryption	What Happens
v3	authNoPriv	HMAC-MD5 or HMAC-SHA	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC) Message Digest 5 (MD5) algorithm or the HMAC Secure Hash Algorithm (SHA).
v3	authPriv	HMAC-MD5 or HMAC-SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard.

SNMPv3 Benefits

SNMPv3 provides secure access to devices by providing authentication, encryption, and access control. These added security benefits secure SNMP against the following security threats:

- masquerade—An SNMP user might assume the identity of another SNMP user to perform management operations for which that SNMP user does not have authorization.
- message stream modification—Messages might be maliciously reordered, delayed, or replayed (to an extent that is greater than can occur through the natural operation of a subnetwork service) to cause SNMP to perform unauthorized management operations.
- disclosure—Unauthorized monitoring of exchanges between SNMP engines might occur. Protecting against this threat might be required as a matter of local policy.

In addition, SNMPv3 provides access control over protocol operations on SNMP-managed objects.

SNMPv3 Costs

SNMPv3 authentication and encryption contribute to a slight increase in the response time when SNMP operations on MIB objects are performed, but the security advantages provided by SNMPv3 far outweigh the small performance hit.

[Table 7](#) shows the order of response time for the various security model and security level combinations. The order of entries is from least (the fastest) response time to greatest (the slowest) response time.

Table 7 *Order of Response Times from Least to Greatest*

Security Model	Security Level
SNMPv2c	noAuthNoPriv
SNMPv3	noAuthNoPriv
SNMPv3	authNoPriv
SNMPv3	authPriv

User-based Security Model

The SNMPv3 user-based security model (USM) applies to SNMP message-level security and offers the following services:

- Message integrity—Ensures that messages have not been altered or destroyed in an unauthorized manner and that data sequences have not been altered to an extent greater than can occur harmlessly.
- Message origin authentication—Ensures that the claimed identity of the user on whose behalf received data was originated is confirmed.
- Message confidentiality—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMPv3 authorizes management operations only by configured users and encrypts SNMP messages.

USM uses two authentication protocols:

- HMAC-MD5-96 authentication protocol
- HMAC-SHA-96 authentication protocol

USM uses CBC-DES (DES-56) as the privacy protocol for message encryption.

Management and Interactive Access over the Internet and Other Untrusted Networks

Many users manage routers remotely and sometimes do so over the Internet. Any unencrypted remote access is risky, but access over a public network, such as the Internet, is especially dangerous. All remote management schemes, including interactive access, HTTP, and SNMP, are vulnerable.

HTTP and HTTPS

Cisco IOS XR software supports remote configuration and monitoring through HTTP protocol and relies on Secure Socket Layer (SSL) for security. When HTTP can use SSL, it becomes the secured version—HTTPS.

In general, HTTP access is equivalent to interactive access to the router. The authentication protocol used for HTTP is equivalent to sending a clear-text password across the network. Unfortunately, HTTP has no effective provision for challenge-based or one-time passwords and is therefore a relatively risky choice for use across the public Internet.

If you use HTTP for management, you should restrict access to appropriate IP addresses and use authentication for users. As with interactive logins, the best choice for HTTP authentication is to use authentication, authorization, and accounting (AAA) services with a TACACS+ or RADIUS server.

If at all possible, avoid using an unencrypted protocol to log in to a router over any untrusted network. If a router has the encrypted login protocol SSH, you can configure HTTP to run over a secure socket by using the **http server** command and including the **ssl** optional keyword. This command enables the HTTP server to run SSL when accessing web pages or files provided by the HTTP server of the router and disables access through the regular HTTP port. Another possibility is to use IPSec encryption for router management traffic, including Telnet, SNMP, and HTTP.

The attacks discussed in this section are relatively sophisticated. These attacks can often be thwarted if the public network providers involved have taken proper security measures. We recommend that you evaluate the trustworthiness of the security measures used by all the providers that are carrying your management traffic. Even if the network security of an ISP seems substantial, we recommend that you remain diligent in protecting your networking investment.

The cautions described in this section apply to both the hosts and the routers. These cautions apply to protecting the router login sessions, but you should also consider applicable mechanisms for protecting host systems if you administer those hosts remotely. Remote Internet administration is useful but requires careful attention to security.

Defeating a Packet Sniffer

Some attackers install *packet sniffer* programs when they break into computers owned by ISPs or in other large networks. A packet sniffer program monitors traffic in the network and steals data, such as passwords and SNMP community strings. Although packet sniffing is becoming more difficult as network operators improve security, it still happens. In addition to the risk from outside attackers, rogue employees have also installed sniffers. Any password sent over an unencrypted channel is at risk, and this risk includes the login and enable passwords for a router.

If at all possible, avoid using an unencrypted protocol to log in to a router over any untrusted network. If a router has the encrypted login protocol SSH, you can configure HTTP to run over a secure socket by using the **http server** command and including the **ssl** optional keyword. This command enables the HTTP server to run SSL when accessing web pages or files provided by the HTTP server of the router and disables access through the regular HTTP port. Another possibility is to use IPsec encryption for router management traffic, such as Telnet, SNMP, and HTTP.



Note

These encryption features are subject to certain export restrictions imposed by the United States Government and are special-order, extra-cost items.

Other Internet Access Dangers

This section describes security risks (other than sniffer programs) to the management of routers over the Internet. In addition to packet sniffers, remote Internet management of routers presents the following security risks:

- To manage a router over the Internet, you must permit at least some Internet hosts to have access to the router. The host addresses could be spoofed, or the hosts could be compromised in other ways. If you permit interactive access over the Internet, security then depends not only on your own antispoofing measures but also on those of the ISPs involved.

These dangers can be reduced by making sure that all the hosts that are permitted to log in to your router are under your control and by using encrypted login protocols with strong authentication.

- *Hijacking* of an unencrypted TCP connection (such as a Telnet session) is a possibility. With connection hijacking, the attacker takes control of a user session. Although such hijacking attacks are both more complicated and less common than simple packet sniffing, an attacker who wants to target a specific network might try this method. The only real prevention against session hijacking is to use an encrypted management protocol that is strongly authenticated.

- DoS attacks are relatively common on the Internet. If a network is experiencing a DoS attack, you might not be able to reach your router to collect information or take defensive action. Although you can take steps to make your network more resistant to DoS attacks, the only real defense against this risk is to have a separate, out-of-band management channel, such as the current Management Ethernet interface.

Logging

Cisco routers can log information about a variety of events, and many of these events are significant for security. Logs can help greatly in characterizing security incidents and in informing your response. The main types of logging on Cisco routers are as follows:

- *AAA logging* collects information about user logins, logouts, HTTP accesses, commands executed, and similar events. AAA log entries are sent to authentication servers that run the TACACS+ or RADIUS protocol. If you use a TACACS+ or RADIUS server, consider enabling AAA logging of various types by using AAA configuration commands. For important information on AAA tasks and commands, see *Cisco IOS XR System Security Configuration Guide* and *Cisco IOS XR System Security Command Reference*.
- *SNMP trap logging* sends notifications of significant changes in system status to SNMP management stations.
- *System logging* records a large, configurable variety of events. System logging events can be reported to a variety of destinations:
 - System console port (**logging console**)
 - Servers using the UNIX “syslog” protocol (**logging trap**)
 - Remote sessions on VTYS and local sessions on TTYs (**logging monitor** and **terminal monitor**)
 - Local logging buffer in router RAM (**logging buffered**)

From a security point of view, the most important events usually recorded in the log are:

- Interface status changes
- Changes to the system configuration
- Access list matches
- Events detected by the optional firewall and intrusion detection features

Each system logging event receives a value for its urgency level. The levels range from debugging information (the lowest urgency) to system emergencies. You can configure each logging destination with an urgency threshold such that an event is logged if it comes in at or above the threshold.

Saving Log Information

By default, system logging information goes to the asynchronous console port only. Because many console ports are either not monitored or connect to terminals that lack any historical memory and have relatively small displays, this information might not be available when you need it. This deficiency can be a greater problem for debug sessions conducted remotely, across the network.

Even a moderately sized logging buffer is very useful. You can use the **show memory** command to ensure that your router has enough free memory to support a logging buffer. Create the buffer by using the **logging buffered** command.

The Cisco CRS-1 router and Cisco 12000 Series Router support syslog servers. Logged information can go to a server after the server has been identified through the **logging** command. You can also:

- Control the threshold of message severities sent to the server by using the **logging trap** command.
- Configure the time stamp format for log entries by using the **service timestamps** command.
- Set the size of the log file for the local logging service through the **logging localfilesize** command.

Recording Access List Violations

If a router uses ACLs to filter traffic, it should log packets that violate the filtering criteria. ACL logging has the capability of characterizing traffic associated with network attacks by logging the suspect traffic.

Securing IP Routing

This section describes some basic security measures related to IP packet forwarding.

Antispoofing

Many network attacks rely on *spoofing* (falsifying) of the source addresses of IP packets. An attack is more difficult to trace if the attacker uses somebody else's address.

Antispoofing measures should be enabled at every point in the network that supports antispoofing, but usually these measures are easiest and most effective at the borders between large address blocks or between domains of network administration. To implement antispoofing measures on every router in the entire network is less practical because of the difficulty of determining which source addresses could legitimately appear on any particular interface. ISPs should be careful to apply antispoofing controls at end-user connection points (see also RFC 2267).

Many administrators of corporate firewalls or perimeter routers install antispoofing measures to prevent hosts on the Internet from taking on the addresses of internal hosts. However, they might overlook steps that prevent internal hosts from taking on the addresses of hosts on the Internet. The prevention of spoofing in both directions is good practice.

Some reasons for setting up antispoofing in both directions at an organizational firewall are:

- An internal user is less likely to attempt an attack and is less likely to succeed.
- Accidentally misconfigured internal hosts are less likely to cause trouble for remote sites.
- External attackers often attempt to break into a network to use it as a base for further attacks.

Antispoofing with ACLs

The goal of spoofing protection is to discard packets that arrive on interfaces that are not authentic paths from the supposed source addresses of those packets. For example, on a router that connects a corporate network to the Internet, a packet should be discarded if it arrives on an interface with a source address field that seems to be from a device inside the corporate network. Similarly, any packet arriving on the interface connected to the corporate network, but whose source address field seems to be from a machine outside the corporate network, should be discarded. In routers that run Cisco IOS XR software, the ACLs operate at the ingress of the line cards (LCs) and modular service cards (MSCs).



Note

For external network purposes, the better protection against spoofing is uRPF. For protecting the external network interfaces, ACLs are adequate. Configuring ACLs on the interior is practical because you know the topology of your own network.

An ISP that carries transit traffic might have limited opportunities to configure antispoofing ACLs, but at least it can usually filter outside traffic that seems to originate within the address space of that ISP.

In general, antispoofing filters are built with input ACLs because packets are filtered at the interface ingress instead of the egress. ACLs are built with the **ipv4 access-list** or **ipv6 access-list** command.

When a router has antispoofing ACLs, it should always reject packets with broadcast or multicast source addresses and packets with the reserved loopback address as a source address. An antispoofing ACL should also filter out all ICMP redirects regardless of the source or destination address.

Antispoofing with Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (uRPF) defeats IP address spoofing. With uRPF, the router checks the source address of any packet against the interface through which the packet entered the router. This source address verification is a dynamic method to defeat spoofing and is much more efficient than an ACL to protect against spoofing.

uRPF can support *loose* or *strict* filtering. Loose filtering means that the interface can accept packets from a path to the source address that is merely a valid path. Strict filtering means that the interface accepts packets from only one valid path.

In the current release of the Cisco 12000 Series Router, certain variations exist in uRPF support:

- IPv4: Only uRPF loose mode is supported on Engine3 and Engine 5 line cards.
- IPv6: uRPF strict mode and uRPF loose mode are supported on Engine3 and Engine 5 line cards.

In addition to a local inside interface, you can enable uRPF on an outside interface. Enter the interface configuration mode for the interface, and then enter the applicable command. The applicable command is either the **ipv4 verify unicast source reachable-via** command or **ipv6 verify unicast source reachable-via** command.

Controlling Directed Broadcasts

IP-directed broadcasts can be exploited by the common *smurf* DoS attack and similar attacks. An IP-directed broadcast is a packet that is sent to the broadcast address of a subnet to which the sending device is not directly attached. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet. At the target subnet, the packet is converted into a link-layer broadcast. Because of the nature of the IP addressing architecture, only the last router in the chain, the one that is connected directly to the target subnet, can conclusively identify a directed broadcast. Directed broadcasts are occasionally used for legitimate purposes, but such use is not common outside the financial services industry.

The smurf attacker sends ICMP echo requests from a spoofed (falsified) source address to a directed broadcast address. The requests cause all the hosts on the target subnet to send replies to the falsified source. By sending a continuous stream of such requests, the attacker can create a much larger stream of replies. The intended result is for the replies to overwhelm the host whose address has been spoofed.

Path Integrity

Many attackers depend on their ability to influence the path that a packet takes over a network. If they can control routing, attackers might be able to spoof the address of another machine and have the return traffic sent to themselves, or they might be able to intercept and surreptitiously read data. Routing could also be disrupted to deny service.

IP Source Routing

The IP protocol supports source routing options that allow the packet sender to:

- Control the route that packet takes to its destination.
- Influence the route that any reply is likely to take.

IP source routing is disabled by default and is rarely used for legitimate purposes in real networks.

ICMP Redirects

An ICMP redirect message instructs an end node to use a specific router as its path to a particular destination. In a properly functioning IP network:

- A router sends redirects to hosts on its own local subnets only.
- No end node ever sends a redirect.
- No redirect ever traverses more than one network hop.

An attacker might attempt to violate these rules to start an attack. Filtering out the incoming ICMP redirects at the ingress is a good practice on any router that exists at the border between administrative domains. Note, however, that only remote attackers are thwarted by this type of filtering. An attacker can still cause significant trouble by using ICMP redirects if the host that is used by an attacker directly connects to the same segment as the host that is targeted in an attack.

Routing Protocol Filtering and Authentication

If a router uses a dynamic routing protocol that supports authentication, that authentication should always be used. Authentication prevents some malicious attacks on the routing infrastructure and can also help to prevent damage caused by misconfigured rogue devices on the network.

We strongly advise that service providers and other operators of large networks use route filtering to prevent their routers from accepting obviously incorrect routing information.

Although excessive use of route filtering can greatly reduce the advantages of dynamic routing, proper filtering often helps to prevent unwanted consequences. For example, if a router uses a dynamic routing protocol to communicate with a customer network, the router should not accept any routes from that customer other than routes that have actually been defined for that customer. The customer provides a list of its routes to you, and your network filters out everything that is not on that list.

Flood Management

Many DoS attacks consist of the flooding of useless packets. These floods are intended to congest network links, slow down hosts, or overload routers. Cisco IOS XR software was designed to stop or greatly reduce the impact of any flood.

An important part of flood management is to be aware of the locations of performance bottlenecks. If a flood is overloading a particular line, filtering out the flood on the router at the source end of the line is effective, but filtering at the destination end provides little or no help.

Transit Floods

Using the Cisco QoS features can protect hosts and links against certain kinds of floods. Unfortunately, a general treatment of flood management is beyond the scope of this guide, and the protection depends largely on the attack. In general, we recommend that you use weighted fair queueing (WFQ) to protect against flooding attacks.

If you plan to use QoS features to control floods, we recommend that your organization have a robust understanding of QoS features and how common flooding attacks work. For example, WFQ is much more effective against ping floods than against SYN floods because the typical ping flood appears to WFQ as a single traffic flow, whereas each packet in a SYN flood generally appears as a separate flow. A smurf reply stream falls somewhere between the two. For information on QoS, see *Cisco IOS XR Modular Quality of Service Configuration Guide*, *Cisco IOS XR Modular Quality of Service Command Reference*, and *Cisco IOS XR IP Addresses and Services Command Reference*.

We provide features intended specifically to reduce the impact of SYN flood attacks on hosts. The TCP intercept feature and SYN flood protection are two such features. SYN flood protection can be complex, and results can vary with the flood rate, router speed, memory size, and hosts that are involved.

Switching Modes and Cisco Express Forwarding

Although most flooding DoS attacks send their traffic to either one or a few targets—and, therefore, do not overwhelm the standard cache maintenance subsystem—many common SYN flooding attacks use randomized source addresses. The host under attack replies to some fraction of the SYN flood packets and consequently creates traffic for a large number of destinations. Routers configured for Cisco Express Forwarding (CEF), therefore, perform well during SYN floods that are directed at hosts, not at the routers themselves. CEF is always enabled in Cisco IOS XR software.

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) can serve some network management needs. CDP is useful when you are trying to configure something at the far end but you do not know what precisely is on the other end. Using CDP is not necessary if the network is properly diagrammed and the ports and so on are properly labeled. CDP lets a device on a directly connected segment learn:

- That the router is a Cisco device
- Router model number
- Cisco IOS XR software version that is running
- Host names from a router (a bad idea for external ports)
- Remote port ID

Using CDP is dangerous on external ports because an attacker might be able to use this information to design an attack against the router. On internal ports, CDP is not dangerous. The CDP protocol is off by default but can be enabled through the CLI.

The Cisco CRS-1 router and Cisco XR 12000 Series Routers use CDP Version-2 (CDPv2). This version provides device-tracking features that include a reporting mechanism for rapid error tracking. Error messages can be sent to the console or a logging server and can include instances of unmatched native VLAN IDs (IEEE 802.1Q) on connecting ports and unmatched port duplex states between connecting devices. For detailed information on CDP, see the *Implementing CDP on Cisco IOS XR Software* module in *Cisco IOS XR System Management Configuration Guide*.

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) can serve some network management needs. CDP is useful when you are trying to configure something at the far end but you do not know what precisely is on the other end. Using CDP is not necessary if the network is properly diagrammed and the ports and so on are properly labeled. CDP lets a device on a directly connected segment learn:

- That the router is a Cisco device
- Router model number
- Cisco IOS XR software version that is running
- Host names from a router (a bad idea for external ports)
- Remote port ID

Using CDP is dangerous on external ports because an attacker might be able to use this information to design an attack against the router. On internal ports, CDP is not dangerous. The CDP protocol is off by default but can be enabled through the CLI.

The Cisco CRS-1 router and Cisco 12000 Series Router use CDP Version-2 (CDPv2). This version provides device-tracking features that include a reporting mechanism for rapid error tracking. Error messages can be sent to the console or a logging server and can include instances of unmatched native VLAN IDs (IEEE 802.1Q) on connecting ports and unmatched port duplex states between connecting devices. For detailed information on CDP, see the *Implementing CDP on Cisco IOS XR Software* module in *Cisco IOS XR System Management Configuration Guide*.

Protecting the Legitimacy of the Routing Domain

The information for the remainder of the module is deeper and more specific than the general description of security measures described previously in this module. The various classes of attacks and the possible types of attacks against a routing system are described in the [Examples of Network Attacks](#) module. This section describes the steps you can take to prevent the attacks from taking place.

Protecting Routers from Being Compromised

The first point of attack against a routing system is the router itself. Attacking the routing information and data flows on a router is easier than reaching the same objective on the wires between the routers. Therefore, the physical router and its site should be secure.

Protecting Routing Information on the Wire

Although an attack on a router itself is the easiest way to attack the routing *system*, attacking the routing information on the wire is difficult but possible. The best way to protect routing information is to use MD5 to authenticate the routing protocol packets and provide an integrity check.

A cryptographic signature combines three items, as seen in [Figure 13](#):

- Encryption algorithm, which is generally public knowledge
- Key used in the encryption algorithm, which is a secret shared by the routers that are authenticating their packets
- Contents of the packet itself

MD5 participates in signature computation in ways that are beyond its authentication of routing protocol updates through a password or passphrase. The MD5 key is used to compute a signature based on the data within the routing protocol update. MD5 signatures can be used to detect even a single-bit change in a packet. For more information on the MD5 signature algorithm, see RFC1321 (<http://www.ietf.org/rfc/rfc1321.txt?number=1321>).

IPSec provides another way to protect routing information that is transmitted between routers. You can think of an IPSec connection as a secure tunnel between two routers. For information on the IPSec tasks or commands, see *Cisco IOS XR System Security Configuration Guide* or *Cisco IOS XR System Security Command Reference*. IPSec supports two basic *protocols* and two *modes*.

The protocols within IPSec are as follows:

- The *authentication header* applies a signature (protection) to just the header of the packet.
- *Encapsulated Security Payload* (ESP) encrypts the entire payload.

The two IPSec modes operate as follows:

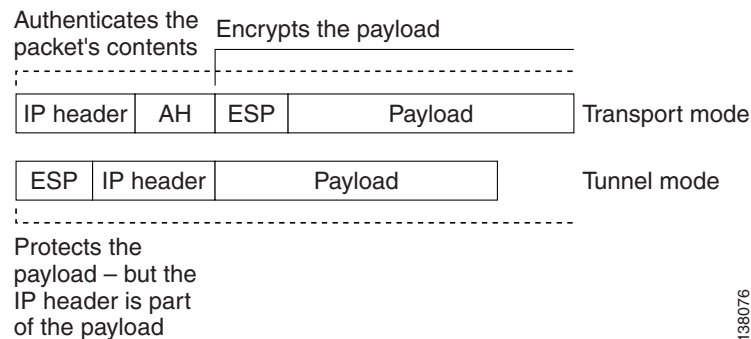
- In tunnel mode, the IP packet is tunneled through IPSec.
- In transport mode, the IPSec header is placed after the IP header, and authentication, encryption, or both of these occurs on just the payload of the packet.

With the protocols and transport modes, IPSec can operate in four possible ways:

- Tunnel and authentication header
- Transport and authentication header
- Tunnel and ESP with NULL encryption
- Transport and ESP with NULL encryption

The typical combinations that make up an IPSec implementation are illustrated in [Figure 14](#).

Figure 14 IP Security Modes of Operation



Of the possible combinations of protocols and modes, the most likely combination to be used for transporting a routing protocol session is tunnel mode and ESP with NULL encryption. ESP can be combined with different forms of encryption, but the *integrity* of the information within a routing protocol packet is a greater concern than the *confidentiality* of the information.

IPSec usually relies on the Internet Key Exchange (IKE) protocol to manage the distribution of keys between devices that are running IPSec. Generally, these keys are public or private key pairs exchanged between each pair of IPSec devices. IPSec is designed for unicast, point-to-point communications between devices rather than broadcast or multipoint models of communication. IPSec cannot be used to

protect data that is multicast between routers. For this reason, the best use of IPsec protection is between BGP peers, especially external BGP (eBGP) peers bordering between two different autonomous systems. OSPF or EIGRP can operate with IPsec if you take either of the following actions:

- Disable the broadcast capabilities of each protocol and thus force all neighbor adjacencies to be built through unicast hellos and routing updates only.
- Run the routing protocol over an encrypted Generic Routing Encapsulation (GRE) tunnel. (Two end points can build a logical IPsec or tunnel GRE for getting traffic from another site.)

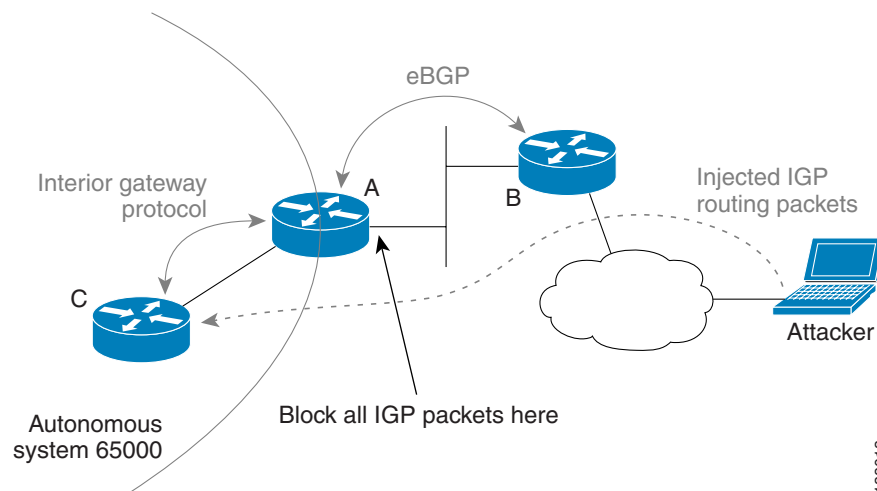
Protecting Against Illegitimate Devices Joining the Routing Domain

This section describes how a device can be prevented from illegitimately joining a routing domain. The mechanism for preventing such intrusion depends on the cryptographic signature that is built by using a key, the routing data being transmitted, and the encryption hash.

If the routing data changes during transit, the receiver is not able to compute the same signature that the transmitter sent, so the routing data is discarded. In the same way, if the *key* that the transmitter uses to build the signature is not the same as the key the receiver has, then the receiver cannot produce the same signature that the transmitter has sent along with the data, so the receiver discards the data. Thus, unless the transmitter has the right key with which it can build the signature, it cannot be part of the routing domain. Using MD5 or IPsec can be an effective way to prevent unauthorized devices from inserting false routing information into a routing domain.

When a network connects to devices outside the local administrative domain—whether the Internet, business partners, or others—the Interior Gateway Protocols (IGP) should be filtered at the edge of the network to prevent outsiders from injecting any false routing data into the network. For an illustration of how to protect the network through filtering at the edge, refer to [Figure 15](#) and the configuration description that follows it.

Figure 15 Preventing Interior Gateway Protocol Attacks Through a Domain Border



The attacker in [Figure 15](#) wants to inject false routing information into AS 65000 to disrupt routing. The easiest way to do this is to find the address of Router C and then unicast OSPF, EIGRP, or RIP packets directly to Router C itself. The network administrator should set up filters on Router A to prevent all interior gateway protocol packets from passing into the domain from outside.

MD5 and Peer Authentication

A critical issue for authentication of routing protocols is the choice of passwords. When you select the passwords for authenticating the routing protocols, four well-known tactics for meeting the security objectives are as follows:

- Create passwords that are hard to guess or break.
- Avoid using the same password on a large number of devices in the network.
- Change passwords on a regular basis.
- Inside the routing domain, do not use passwords that are shared outside the routing domain.

Risks exist when you use the same password on numerous devices in a network. For example:

- If a single device is compromised, all other devices in the network can also be compromised. Using multiple passwords to protect routing information creates a firewall between devices and helps prevent an attacker from taking down the entire network if one router is compromised.
- Every router using the same key to sign data provides more information that an attacker could use to feed into analysis and break the key.

Standard rules to help you pick passwords that are hard to guess or break have been widely published. For more information on this topic, see:

- *Key Management Considerations for the TCP MD5 Signature Option*, RFC 3562, by M. Leech <http://www.ietf.org/rfc/rfc3562.txt?number=3562>
- *Choosing a Safe Password*, at <http://www.passwordexperts.com/article01.shtml>
- *Choosing a Password*, at <http://www.unix-manuals.com/tutorials/unix/choose-password/password-choice.html>

Consider using several keys for MD5 authentication in the network. The most logical approach is likely to split the key usage along topological boundaries in the network. Consider changing keys at the same places where area boundaries exist and summarization is used. The benefits of following these guidelines are that they:

- Provide a natural boundary for keys
- Conveniently let operators know what keys are used where
- Provide breaks at logical points

Keys that are used for routing authentication should be changed on a regular basis, too. The frequency of change depends on many factors, including:

- Number of devices within the network that are using a specific key
- Whether a key is used over links that can be easily compromised (by internal or external attacks)
- Average number of packets that a device using a specific key transmits over a period of time

Any key that protects a session directed to outside routing domains should be different from keys that protect routing traffic and peering sessions within the routing domain.


When using MD5 or IPSec to protect routing data, you should consider the processing required to perform the required cryptographic operations. The performance of these cryptographic algorithms is constantly scrutinized, and various methods are developed to improve their performance. However, you should be aware that because of the processing overhead, a router that runs peering and route authentication can peer with fewer routers and process fewer outside tasks as the overhead grows.

TTL-Based Peering Session Protection (BGP TTL Security Hack)

In many of the attacks illustrated in the [Examples of Network Attacks](#) module, the examples show the attacker situated either multiple hops away from the router under attack or possibly outside the routing domain. To take advantage of this fact, the Generalized Time-to-Live Service Mechanism (GTSM) is designed to thwart all attackers who do not connect directly to the physical network that connects two routers. For this description of GTSM, refer to [Figure 16](#).

In [Figure 16](#), the object of an attacker is Router A (as shown by the line labeled “BGP attack packets”), and the attacker is connected to Router C.

Figure 16 TTL-Based Peering Session Protection



If Router A and Router B in [Figure 16](#) are configured with normal BGP peering, Router B generates BGP packets destined to Router A with a TTL of 1. As Router A receives the packets, the TTL is reduced to 0, and the packets are processed normally because they are destined to Router A. If the attack uses the injection of false TCP or BGP packets into the A-to-B peering session, it can generate these packets with a TTL of 3 and transmit them. Router C reduces the TTL on the packets to 2 and then forwards them to Router B. Router B reduces the TTL to 1 and then forwards the packets to Router A. When Router A receives these packets, it reduces the TTL to 0 and processes them normally. This example illustrates that an attacker can easily inject packets into the peering session between Router A and Router B.

If Router B is instead configured with the TTL to be 255 for BGP packets that it transmits, and if Router A is configured to ignore (or drop) all BGP packets with a TTL of 254 or lower, the attacker cannot send packets to Router A through Router C. With an incorrect TTL, attack packets are discarded.

Consider a situation in which an attacker attempts to direct packets to Router A with the intention of disrupting the peering session between Router A and Router B. The highest TTL the attacker can assign to the packet is 255 because of limitations in the IP header. When Router C receives this packet, it reduces the TTL by 1 and forwards the packet. Router B likewise reduces the TTL by 1, so the TTL is now 253, and forwards it. In this case, when Router A receives the packet, it discards the packet because it is a BGP packet with a TTL of less than 254.

GTSM keeps an attacker from injecting false routing information from anywhere other than on the link that directly connects the routers. It effectively prevents attackers from affecting the peering sessions between routers. Thus, GTSM dramatically reduces the scope of attacks against routing protocols.

Protecting the Routing Information

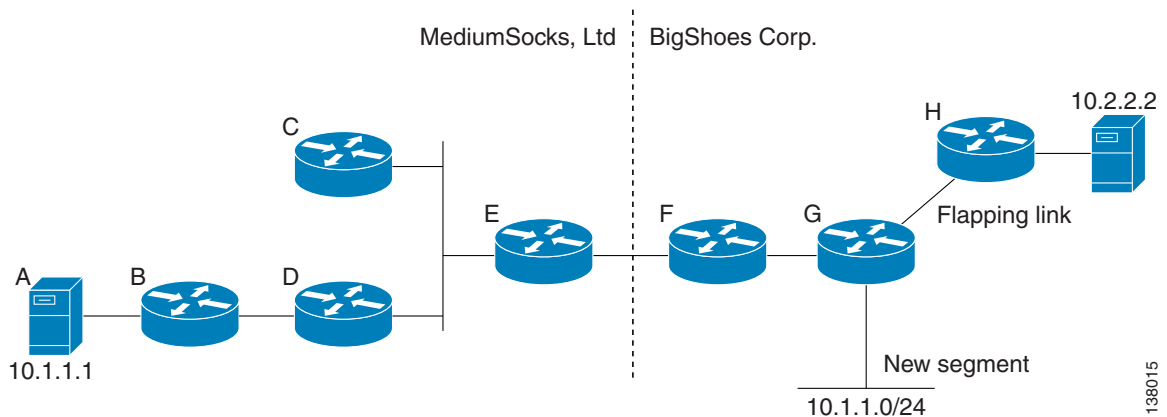
This section describes ways that you can protect routing information. This protection is provided in addition to what is provided to peering sessions between routers with techniques like filtering at the edge of your network and GTSM. The routing information that needs protection is:

- In the routing tables
- Carried by the routing protocols—Information that directs packets to the right destinations

Extranet Connections

A primary point at which an attacker might inject false routing information into a network is at an extranet connection—the place where a network connects to networks outside your administrative control (usually for business-to-business transactions or information sharing). An extranet connection tends to be a prime point for the injection of false routing information because it represents trusted relationships (which generally use an interior gateway protocol to exchange routing information). In fact, an extranet connection can be the source of many internal routing problems, whether or not they are sent with criminal intent. Two of the possible problems are illustrated in Figure 17.

Figure 17 Possible Impacts of Exchanging Routing Information with Outside Networks



MediumSocks, Ltd., has an inventory-sharing arrangement with BigShoes Corp. To facilitate this agreement, they have set up an extranet connection between their networks, shown here between Router E and Router F. MediumSocks has a server, 10.1.1.1, which has long been used for internal purposes and is not reachable to hosts in the BigShoes network. The administrators at BigShoes do not know MediumSocks is using the 10.1.1.0/24 range of addresses, and BigShoes sets up a new segment, 10.1.1.0/24, attached to Router G. When this segment is advertised into the BigShoes' routing protocol, it is also advertised into the MediumSocks network, becoming the preferred path, at Router C, to 10.1.1.0/24. This move cuts off internal access to the 10.1.1.1 server from behind Router C in the MediumSocks network. Although this problem can be resolved by simple route filtering at the Router E and Router F border between the two companies, finding this problem is not necessarily easy, nor are these types of problems easily foreseen when filtering is initially configured for Router E and Router F.

Another problem with sharing routing information in this way is the impact that changes in one network have on the another network. For instance, if the server that the MediumSocks hosts actually interact with is 10.2.2.2, behind Router H, the link between Router G and Router H link becomes unstable and, constantly flaps. Each time the link flaps, the MediumSocks network must converge to account for this change in network topology.

The unreachability and flapping examples demonstrate the impact of misconfigurations, miscoordination, or failing links. These events are normal events. However, these typical problems can also be used as attacks against the network routers. Therefore, you should take steps to prevent these network problems in another network from causing problems in your network. The sections that follow describe some strategies that can protect your network from attacks or prevent routing information from other routing domains from causing problems in your network.

Using an Exterior Gateway Protocol for All Extranet Connections

This section focuses on the use of the Exterior Gateway Protocol BGP (eBGP) to exchange routing information with an outside routing domain.

The essential rule for all connections to extranets is: Never use an Interior Gateway Protocol (IGP) to exchange routing information dynamically with an outside routing domain. Always use eGBPs, or use static routes, rather than dynamically exchanging routing information through an IGP. Many network administrators reject the use of static routes, primarily because of redundancy and resiliency concerns.

The use of BGP to exchange a small number of routes might seem overly protective, but the following considerations might make its use seem more logical:

- BGP is designed to tolerate well the rapid changes involved with the routing information that is outside the control of the local network administrator.
- Security for the BGP protocol continues to be the subject of dedicated research because BGP is such a widely used exterior gateway protocol. Exterior gateway protocols operate between routing domains, so they operate between mutually defensive parties. Interior gateway protocols operate within a trusted administrative domain.
- Because BGP is an exterior gateway protocol, it lets you build and act on policies rather than just on reachability information.

Based on this information, you should use BGP when exchanging routing information with any routing domain outside your administrative control.



Note

At the provider edge-customer edge (PE-CE) boundary, the service provider does not accept external routing information into its interior gateway protocol. Instead, this information is redistributed into BGP in a way that allows the routes to be redistributed back into the customer Interior Gateway Protocol (IGP) without losing the original routing information.

Filtering Routes Aggressively at the Extranet Edge

Filtering at the network edge should be based on an approach of allowing the minimum necessary routes rather than on accepting everything except a few specific networks. Therefore, the filters for limiting the routes that are learned through BGP at the network edge should deny all routes by default and permit just the routes that are necessary to reach the hosts and servers needed in each network.

Beyond this basic filtering, you should also be sure that the outside peer does not advertise your routes to other peers. Although no way exists for you to guarantee this restriction, you can mark the routes you advertise to the extranet peer with the `NO_EXPORT` community to instruct outside routers not to advertise the route to any of their external peers. For information on related tasks or commands, see *Cisco IOS XR Routing Configuration Guide* or *Cisco IOS XR Routing Command Reference*.

Similarly, your network should:

- Refuse to accept a route from an extranet peer that did not originate within the extranet itself.
- Refuse to propagate a route learned from an extranet peer to any other external BGP peers you may have configured.

The easiest way to accomplish both safeguards is to use AS path filters on every eBGP peering session in your network. The goal is to filter any routes that have an AS path length greater than 1 (or any AS path containing more than one autonomous system) that the network receives. Another goal is to prevent your edge routers from advertising any routes not originating within your autonomous system (any route with AS path length of greater than 1). For more information on AS path filters, see *Cisco IOS XR Routing Command Reference*.

Dampening Prefixes Aggressively at the Extranet Edge

Route dampening is designed to prevent constant changes in a single prefix from destabilizing the local routing system (in which such a prefix has been learned from an eBGP peer). Dampening can specifically address concerns about constantly changing routing information that affects the stability and convergence times of local routing.

We recommend that you dampen routes learned from an extranet much more quickly than routes learned from an ISP for two reasons.

- The AS path of routes learned from an extranet peer should be only 1. Therefore, convergence times should be quick, and little of the usual churn that results from large internetwork convergence should occur.
- The routing information learned from an extranet peer is typically redistributed back into an interior gateway protocol to facilitate reachability from the interior of your network to the extranet destinations (unless you are running BGP on all routers along the path to networks requiring reachability to the destinations in the extranet).

Dampening of BGP route flapping has the parameters shown in the following syntax:

```
bgp dampening [half-life [reuse suppress max-suppress-time] | route-policy name]
```

Before using the **bgp dampening** command, consider the following issues:

- How many times should a route flap before the router dampens it?
- How long should the route be dampened before it is reinstalled in the routing table?
- Most network configurations would call for the route to be suppressed rapidly but also advertised again in a fairly short time after it stops flapping.
- Leave the route out of the table if it is flapping.

Set the parameters as follows:

- Half-life period: 5 minutes
- Reuse threshold: 1500
- Suppress threshold: 3000
- Maximum suppress time: 120

Limiting Route Count at the Extranet Edge

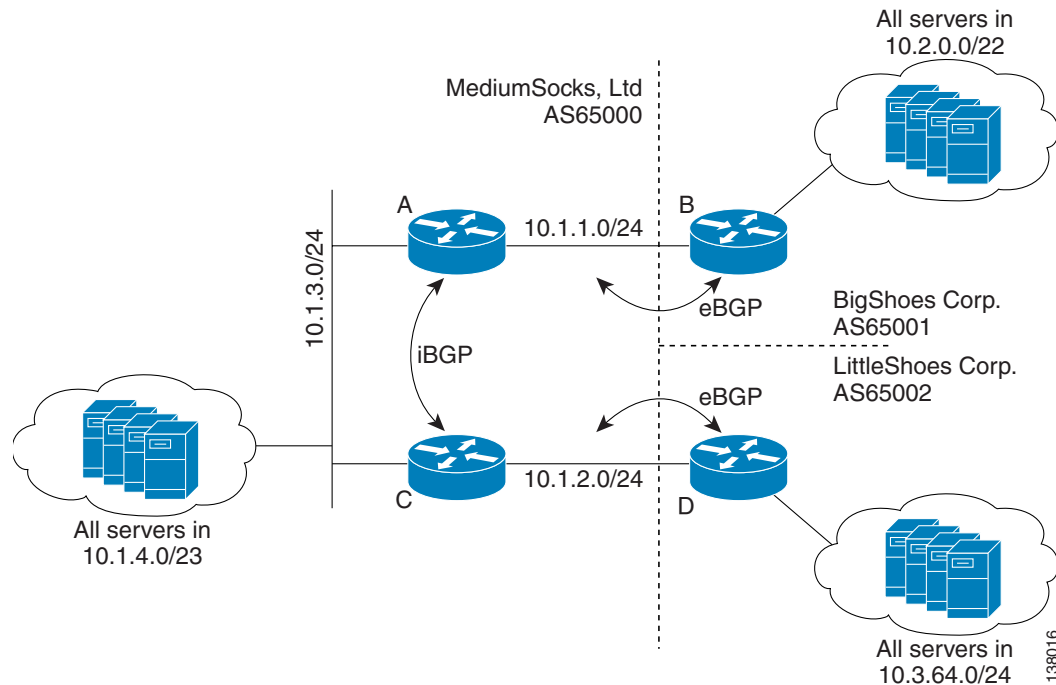
Extranet connections tend to have a very limited scope. A set of well-defined, accepted routes exists, so the number of routes you should be receiving over an extranet connection is relatively easy to know (with some margin of error). Another filtering strategy you can implement at the edge is to limit the number of routes that the routers can learn through BGP from the extranet. Limiting these routes can prevent an outside network from flooding a network with extraneous network information and possibly overloading that network.

Generally, the maximum number of prefixes that are accepted should be set to a low number. For example, if you would normally expect the number of prefixes received across the BGP session with an extranet to be less than 20, then you set the maximum number of prefixes allowed to 30. By setting the maximum higher, you allow for some change in the partnering network and, at the same time, provide some protection for your network.

Sample Extranet BGP Configuration

This section illustrates a pair of extranet connections in which a corporate network (MediumSocks) is sufficiently isolated from those of two of its partners. The connections are illustrated in Figure 18. The MediumSocks network is running iBGP inside the network and eBGP at the boundaries.

Figure 18 Sample Extranet BGP Configuration



Connections to the Internet

Connections to the Internet are often treated in the same way as connections to extranets, but several significant differences exist, including:

- Route filtering
- Protecting the local network from becoming a transit
- Different parameters in route dampening

Some of these issues do not directly relate to security, so this section merely introduces them to help build a comprehensive view of BGP when BGP connects the network to an ISP. For more recommendations for connections to an ISP, see *Recommended Internet Service Provider Security Services and Procedures*, RFC 3013 (<http://www.ietf.org/rfc/rfc3013.txt?number=3013>).

Route Filtering

Route filtering is very important for security when a network connects to the Internet through a service provider, but this type of filtering generally differs from what is applied to extranet connections. Instead of denying all routes by default and permitting relatively small number of routes, you should permit most routes and deny only specific routes. The routes you should deny follow:

- Internal routes—actually, anything you advertise to the service provider—which should be filtered so that the router cannot relearn those destinations from the service provider


- Private address space, defined in RFC 1918 at <http://www.ietf.org/rfc/rfc1918.txt?number=1918>
- Other addresses, called *bogons*, which are either not used at all or—for any reason—should not be accepted

For more information on bogons, see *The Bogon Reference Page* at <http://www.cymru.com/Bogons/>.

Protecting Against Becoming a Transit

If a certain enterprise network connects two ISPs, a poor network design could cause the two ISPs to see the enterprise network as a transit network that can carry their traffic. For example, the red line in [Figure 19](#) shows where BigShoes is functioning as a transit. The problem occurred because BigShoes advertised the AS65004 routes to AS65005.

Figure 19 Two Service Providers Transiting Traffic Through BigShoes, Ltd.



The blue line illustrates the correct traffic flow. The normal sequence for traffic flow follows:

1. The traffic between 10.1.1.0/24 and the destinations in AS65004 should travel to AS65000, an ISP.
2. From AS65000, the traffic flows through AS65003, another ISP, to AS65002.
3. The traffic flows to AS65004 (the blue traffic path in [Figure 19](#)).

An alternate path (in red) carries the traffic from AS65005 to AS65000, then through AS65001 to AS65002 to AS65004, using the BigShoes corporate network as a transit between AS65000 and AS65002. Although this path does not create a security problem for the BGP, it is a resource utilization issue for BigShoes, Ltd. The routers at the edge of AS65001 should be configured so that:

- They prevent routing information learned from AS65000 from being advertised to AS65002.
- They prevent routing information learned from AS65002 from being advertised to AS65000.

The simplest way to prevent the resource problem is to use an AS path filter at these edge routers so that only routes originating in AS65001 can be advertised to AS65002 and AS65000.

Route Dampening

The primary difference between route dampening on an extranet connection and dampening on an ISP connection is the aggressiveness of the dampening parameters. (Route dampening for connecting to extranets is discussed in the section “[Dampening Prefixes Aggressively at the Extranet Edge](#).”) For more information on the recommended route dampening parameters when connecting to an ISP, see RIPE Routing-WG Recommendations for Coordinated Route-flap Dampening Parameters (ripe-229) at <http://www.ripe.net/ripe/docs/ripe-229.html>.

Connections Within a Network

The largest threats to routing are likely to come from within a network. (Preceding sections focused on protecting the network routers from attacks through outside connections.) Knowing how to prevent attacks on a routing system through connections to open ports in wiring closets (especially at remote sites) is therefore very important. Many of the techniques describes in this module, such as peer authentication, can mitigate internal attacks. The forthcoming sections describes two other techniques:


- Route filtering within a network
- Traffic segregation and routing protocol traffic filtering

Route Filtering Within a Network

Route filtering should be employed at topological edges within a network to prevent false routing information from being injected. Typical approaches are to:

- Filter routing information that is coming from remote sites back to a central location (or data center)
- Filter routing information coming from any open areas of the network, such as a lab network

The configuration in [Figure 20](#) demonstrates the filtering of routes that come from remote sites. On Router A, the routes learned from each of the remote networks are filtered to prevent false routing information from being injected from a remote site.

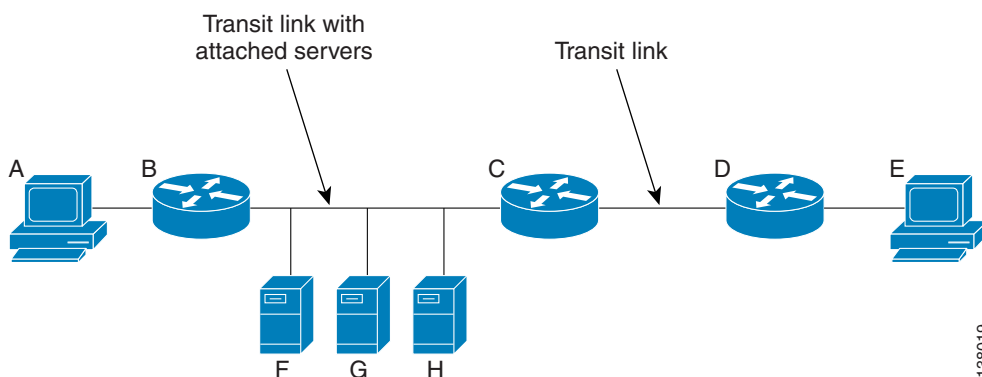
Figure 20 Filtering Routes Within a Network

This filtering is not strict but rather is more of a precaution that is intended to reduce the amount of administrative overhead associated with configuring and maintaining the filter.

For more information on filtering, see *Cisco IOS XR Routing Command Reference* and *Cisco IOS XR Routing Configuration Guide*.

Traffic Segregation and Routing Protocol Traffic Filtering

One way to protect the routing protocols and the routing system from internal attack is to segregate transit links from nontransit links in the network. See [Figure 21](#) and the explanation that follows it.

Figure 21 Transit Links and Traffic Segregation

In [Figure 21](#), note the following details about the sequence:

1. Traffic between Router A and Router E passes over the link between Router B and Router C and passes by Server F, Server G, and Server H. This link is a transit link with server connectivity.
2. Packets are forwarded by Router C to the Router C-to-Router D link, across another transit link.
3. The packets are forwarded by Router D to the final destination, Host E.

The Router A-to-Router B and Router D-to-Router E links are considered access links because various hosts are attached to the network on these links, and the links from Router C to Router D are considered transit links because traffic transits those links. In contrast, the link between Router B and Router C is a mixed link because both of the following are true:

- This link carries traffic between Router A and Router E.
- Servers are also attached to this link.

From a network design standpoint, this arrangement is a bad practice. From a network security standpoint, this arrangement is also bad for the following reasons:

- If Server F, Server G, or Server H is compromised, Router B and Router C (including the peering session between them) can be attacked directly from the compromised server. For example, even if a solution like GTSM is running on all the routers in this network, Server F, Server G, and Server H could attack Router B and Router C directly by injecting routing protocol packets from a segment.
- Although you can filter all routing protocol traffic from being accepted on the access links attached to Router B and Router D, you cannot filter routing protocol traffic on the link between Router B and Router C because these routers need to be able to exchange routing information.

In general, good practices are to:

- Segregate transit links and access links—A shared link should never be a transit link. For the example in [Figure 21](#), you could add a new, preferred physical link between Router B and Router C. Make the cost higher on the transit link that is common to the hosts, so that the traffic on the preferred link is not seen by Server F, Server G, or Server H.
- Filter out routing protocol traffic at points where hosts access the network.

For more information on filtering the routing protocols traffic, see *Cisco IOS XR Routing Command Reference* and *Cisco IOS XR Routing Configuration Guide*.

A

AAA [21, 71](#)
access control list, see ACL [35](#)
access controls elements, see ACE [34](#)
ACE [34, 39, 40](#)
ACL [32, 34, 35, 36, 38, 40](#)
 antispoofing [82](#)
 contents [36](#)
 counters [41](#)
 packet length [ix, 35](#)
 purpose [36](#)
 show ipv4 access-lists command [41](#)
 TTL specification [ix](#)
alarms
 severity level and filtering [64](#)
ALDEMS (Alarm Management and Debugging Event System), description [62](#)
always-available commands [55](#)
antispoofing [79, 81](#)
 filters [82](#)
attack, SYN flood [68](#)
attacks
 authentication [83](#)
 by falsified addresses [81](#)
 DoS [80](#)
 falsifying routing information [48](#)
 hijacking [79](#)
 peering disruption [48](#)
 spoofing [81](#)
attacks, DoS [82](#)
attacks, smurf [82](#)
attacks, SYN flood [84](#)
authentication, authorization, and accounting
 see AAA [21](#)

B

banner login command [73](#)
BGP
 prefix list [41](#)

C

CDP [84, 85](#)
CEF [84](#)
checkpoint [21](#)
Cisco Discovery Protocol, see CDP [84, 85](#)
Cisco Express Forwarding, see CEF [84](#)
classification
 summary [43](#)
code signing [28](#)
command
 logging console [38](#)
 show {ipv4 | ipv6} access-lists [41](#)
 show ipv4 access-lists [41](#)
commands
 always-available CLI commands during attack [55](#)
 banner login [73](#)
 http server [78](#)
 ipv4 access-list [82](#)
 ipv6 access-list [82](#)
 show flow monitor [57](#)
 snmp-server community [75](#)
commands, ip tcp synwait-time [68](#)
Commented IP Access List Entries [40](#)
compartmentalization [22](#)
congestion avoidance
 summary [43](#)
congestion management
 summary [43](#)

correlator [62](#)
 CPU, monitoring [24](#)
 CPU, starvation and lockup [24](#)
 cryptographic algorithms [ix, 26, 71](#)
 HMAC-MD5 [26, 71](#)
 HMAC-SHA1 [26, 71](#)
 MD5 [26, 71](#)
 SHA1 [26, 71](#)

D

denial of service, see DoS [20](#)
 DoS [20, 82](#)

F

falsifying routing information [48](#)
 FIB [32, 34](#)
 filters
 antispoofing [82](#)
 forwarding information base, see FIB [32](#)
 fragment [39](#)
 non-initial [39](#)

G

Generalized Time-to-Live Security Mechanism, see
 GTSM [32](#)
 get-next-request [75](#)
 get-request [75](#)
 get-response [75](#)
 GTSM [32, 33](#)

H

helper process, for syslog [62](#)
 hijacking [79](#)
 HMAC-MD5 [26, 71](#)
 HMAC-SHA1 [26, 71](#)

HTTP [78](#)
 http server command [78](#)

I

ICMP [36](#)
 iFIB [32](#)
 image authentication [28](#)
 in-service software upgrade, see ISSU [25](#)
 internal forwarding information base, see iFIB [32](#)
 IP Access Lists with Fragment Control feature [39](#)
 ip tcp synwait-time [68](#)
 ipv4 access-list command [82](#)
 ipv6 access-list command [82](#)
 ISSU [25](#)

L

logging console command [38](#)
 logging correlation [63](#)
 logging correlation rules [63](#)
 logging correlator buffer [63](#)
 logging events buffer [63](#)
 logging process [63](#)

M

MD5 [26, 48, 71, 74](#)
 memory [24](#)
 memory, monitoring [24](#)
 memory, protected space [24](#)
 memory space, protection [20](#)
 Message Digest Algorithm, see MD5 [48](#)
 micro kernel, see uKernel [20](#)
 modularity [22](#)

N

NetFlow [34, 56](#)

Network Time Protocol, see NTP [22](#)
 non-initial fragment [39](#)
 nonstop forwarding, see NFS [21](#)
 NSF [21](#)
 NTP [22](#)

O

one-time password [72](#)

P

packet length [35](#)
 packet length in ACLs [35](#)
 packet walk [31](#)
 passwords [72](#)
 passwords, guidance [72](#)
 passwords, one-time [72](#)
 peering disruption attack [48](#)
 peering sessions
 cryptographic algorithms [71](#)
 Per VRF AAA [28](#)
 PIE [28](#)
 prefix list [41](#)
 route filtering [41](#)
 process, limits [21](#)
 product installation envelope, see PIE [28](#)
 Protecting Software Images [28](#)

Q

QoS [41](#)
 QoS (Quality of Service)
 benefits [42](#)
 features
 traffic policing [44](#)
 traffic shaping [44](#)
 techniques
 congestion avoidance [43](#)

 congestion management [43](#)
 packet classification [43](#)

QoS (Quality of service)

 characteristics [41](#)

Quality of service, see QoS [41](#)

R

restart, automatic process [24](#)
 route cost [97](#)
 route prefix filtering [71](#)
 RSVP [ix](#)

S

SAM [28](#)
 scheduling [24](#)
 segregating traffic [97](#)
 set-request [75](#)
 SHA1 [26, 71](#)
 show {ipv4 | ipv6} access-lists commands [41](#)
 show flow monitor command [57](#)
 show ipv4 access-lists command [41](#)
 SMU [55](#)
 smurf attack [82, 84](#)
 SNMP [74](#)
 SNMP (Simple Network Management Protocol)
 versions
 security models and levels [76](#)
 SNMPv1,v2c, and v3 comparison [75](#)
 SNMPv3 benefits [77](#)
 SNMPv3 costs [77](#)
 snmp-server community command [75](#)
 SofToken [72](#)
 Software Authentication Manager, see SAM [28](#)
 software module update, see SMU [55](#)
 source route [67](#)
 spoofing [81](#)
 SSH [78](#)

ssl [78](#)
SYN flood [68, 84](#)
synwait time [68](#)
syslog [62](#)
syslog, helper processes [62](#)
system logging, see syslog [62](#)

T

task group [71](#)
TCP keepalives [73](#)
traffic, segregation [97](#)
traffic policing
 summary [44](#)
transit link [96, 97](#)
trap [75](#)

U

uKernel [20](#)
Unicast Reverse Path Forwarding, see uRPF [44](#)
upgradability [25](#)
uRPF [34, 44](#)
 strict [34](#)
 strict or loose [44](#)
 supported address types [45](#)
 supported interfaces [45](#)

W

watchdog system monitor, see wdsysmon [24, 58](#)
wdsysmon [24, 58, 61](#)
weighted fair queueing, see WFQ [84](#)
weighted random early discard, see WRED [34](#)
WFQ [84](#)
WRED [34](#)