SITCS

Implementing Cisco Threat Control Solutions

Volume 2 **Student Guide**
Version 1.0

Part Number: 97-3338-01

https://t.me/learningnets

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned 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. (1110R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS" AND AS SUCH MAY INCLUDE TYPOGRAPHICAL, GRAPHICS, OR FORMATTING ERRORS. CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

© 2014 Cisco Systems, I

https://t.me/learningnets

Table of Contents

Cisco Email Security Appliance 4-1

Describing the Cisco Email Security Solutions 4-3

Cisco Modular Network Architecture and Cisco ESA 4-4
Cisco Hybrid Email Security Solution Overview 4-7
SMTP Terminologies 4-8
SMTP Flow 4-9
SMTP Conversation 4-10
Cisco ESA Services Overview 4-12
Cisco ESA GUI Management 4-14
Cisco ESA Committing the Configuration Changes 4-15
Cisco ESA Licensing 4-16
Incoming Mail Processing Overview 4-18
Outgoing Mail Processing Overview 4-20
Cisco ESA LDAP Integration Overview 4-21
Cisco Registered Envelope Service (CRES) Overview 4-23
Summary 4-25

Describing the Cisco Email Security Appliance Basic Setup Components 4-27

Cisco ESA Listener Overview 4-29
Cisco ESA Listener Type: Private and Public 4-30
Cisco ESA One Interface/One Listener Deployment Example 4-31
Cisco ESA Two Interfaces/Two Listeners Deployment Example 4-33
Cisco ESA Listener Major Components: HAT and RAT 4-34
Cisco ESA One Listener Deployment Scenario 4-35
One Listener Deployment Scenario: Interfaces and Listener 4-37
One Listener Deployment Scenario: LDAP Accept Query 4-41
One Listener Deployment Scenario: HAT 4-44
One Listener Deployment Scenario: HAT > Sender Group 4-46
One Listener Deployment Scenario: HAT > Sender Group SBRS 4-49
One Listener Deployment Scenario: HAT > BLACKLIST Sender Group 4-50
One Listener Deployment Scenario: HAT > RELAYLIST Sender Group 4-52
One Listener Deployment Scenario: HAT > Add Sender Group 4-53
One Listener Deployment Scenario: HAT > Mail Flow Policy 4-55
One Listener Deployment Scenario: HAT > Mail Flow Policy > Anti-Spam and Anti-Virus 4-58
One Listener Deployment Scenario: HAT > Mail Flow Policies Summary 4-59
One Listener Deployment Scenario: RAT 4-60
One Listener Deployment Scenario: SMTP Routes 4-62
One Listener Deployment Scenario: Email Relaying on Internal Mail Server 4-64
Summary 4-65

Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies 4-67

Cisco ESA Incoming and Outgoing Mail Policies Overview 4-69
Cisco ESA Mail Policies Matching 4-72

https://t.me/learningnets

Anti-Spam Overview 4-76
Anti-Spam Configuration 4-77
Spam Quarantine Configuration 4-82
Policy, Virus, Outbreak Quarantines Configuration 4-84
Anti-Virus Overview 4-85
Anti-Virus Configuration 4-87
Content Filters Overview 4-90
Content Filters Configuration 4-92
Outbreak Filters Overview 4-95
Outbreak Filters Configuration 4-97
Data Loss Prevention Overview 4-100
Data Loss Prevention Configuration 4-103
Reporting Overview 4-108
Message Tracking 4-111
Trace 4-114
Summary 4-117
Module Summary 4-119
References 4-120
Module Self-Check 4-121
Cisco Intrusion Prevention Systems 5-1
Describing Cisco IPS Solutions 5-3
Cisco Modular Network Architecture and Cisco IPS 5-4
Intrusion Detection Systems vs. Intrusion Prevention Systems 5-8
Intrusion Prevention Systems Terminologies 5-10
Traditional Network Intrusion Prevention Systems 5-13
Sourcefire NGIPS Benefits Overview 5-16
ASA NGFW Services with NGIPS Overview 5-18
ASA NGFW Services NGIPS Threat Profile 5-19
Apply NGIPS Threat Profile in NGFW Services Access Policy 5-22
NGFW Services NGIPS Global Settings 5-23
Viewing NGIPS Threat Information in PRSM Event Viewer 5-24
Summary 5-25
Integrating Cisco IPS Sensor into a Network 5-27
Cisco IPS Sensor Deployment Modes Overview 5-28
Deploy Sensors in Promiscuous Mode 5-29
Deploy Sensors in Inline Interface Pair Mode 5-33
Deploy Sensors in Inline VLAN Pair Mode 5-36
Deploy Sensors in Inline VLAN Group Mode 5-38
Cisco IPS Management Options Overview 5-40
Configure Interfaces on IPS Sensor 5-42
Configure Promiscuous Interfaces 5-43
Configure Inline Interface Pairs 5-44
Configure Inline VLAN Pairs 5-46
Configure Inline VLAN Groups 5-48

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Traffic Redirection to IPS Module Overview 5-50
Summary 5-53
Configuring Basic Cisco IPS Operations 5-55
IPS Signatures 5-56
IPS Signature Properties 5-59
IPS Actions 5-60
Configure a Virtual Sensor 5-61
IPS Threat Profiles Overview 5-64
Configure Basic Signature Properties 5-66
Risk Rating 5-69
Risk Rating Calculation 5-70
Threat Rating 5-73
Event Action Overrides Overview 5-74
Configure and Verify Event Action Overrides 5-76
Event Action Filters Overview 5-81
Configure and Verify Event Action Filters 5-82
Examine IPS Events using IDM 5-84
Summary 5-86
Tuning Cisco IPS Signatures 5-87
False Negatives 5-88
False Positives 5-89
False Positive Examples 5-90
Cisco IPS Tuning Approaches 5-91
Tune Cisco IPS to Reduce False Positives 5-92
Reduce False Positives: Narrow Search Context 5-93
Reduce False Positives: Narrow Header Values 5-94
Reduce False Positives: Limit Number of Matched Patterns 5-95
Reduce False Positives: Limit Number of Matched Patterns Example 5-96
Reduce False Positives: Decrease Attention Span 5-97
Reduce False Positives: Increase Number of Events 5-98
Tune Cisco IPS to Reduce False Negatives 5-99
Reduce False Negatives: IP Reassembly 5-101
Reduce False Negatives: TCP Reassembly 5-102
Reduce False Negatives: Deobfuscation 5-103
Summary 5-104
Configuring Custom Cisco IPS Signatures 5-105
Custom Signatures Overview 5-106
Signature Engines 5-107
Custom Signature Creation Procedure 5-108
Custom Signature Matching Strategies 5-109
Configure Custom Signatures 5-110
Custom Signature Wizard 5-111
Custom Signature Wizard Options 5-112
Verify Custom Signatures 5-113
Summary 5-114

https://t.me/learningnets

Configuring Cisco IPS Anomaly Detection **5-115**
Anomaly Detection 5-116
Worm Scanning Methods 5-117
Scanners and Histograms 5-118
Anomaly Detection Zones 5-122
Anomaly Detection Learning and Knowledge Base 5-123
Anomaly Detection and Actions 5-125
Anomaly Detection Scenario 5-126
Anomaly Detection Configuration Procedure 5-127
Verify Knowledge Base Creation 5-136
Verify Anomaly Detection Operational Mode 5-137
Verify Anomaly Detection Statistics 5-138
Summary 5-139
Configuring Cisco IPS Reputation-Based Features **5-141**
Global Correlation and Reputation-Based Filtering Overview 5-142
IPS Processing Flow with Global Correlation and Reputation Filter 5-144
Reputation Filter Operations 5-145
Global Correlation Operations 5-146
Dynamic Updates from Cisco SensorBase 5-148
IPS Sensor Network Participation to Cisco SensorBase 5-149
Global Correlation Configurations 5-150
Verify Global Correlation and Reputation Filter 5-155
Summary 5-158
Module Summary **5-159**
References 5-161
Module Self-Check **5-163**
Glossary ***G-1***

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module 4

Cisco Email Security Appliance

Introduction

Cisco Email Security Appliance (ESA) protects the email infrastructure and employees who use email at work by filtering unsolicited and malicious email before it reaches the user. It provides high-availability email protection against the constant, dynamic, and rapidly changing threats affecting email today. The appliance combine best-in-class technologies to prevent and respond to multilevel threats. Cisco ESA integrates into existing email infrastructures with a high degree of flexibility.

This module covers the basic of email security deployment and configurations using the Cisco Email Security Appliance.

Upon completing this module, you will be able to meet the objectives:

Illustrate the SMTP flows and conversations and provides a high level overview of the Cisco Email Security Appliance services

Describe the basic configuration components to setup the Cisco ESA, which includes the listener, LDAP queries, HAT, RAT, Mail Flow Policies and SMTP Routes table

Explain how to configure the different features within the incoming and outgoing mail policies (anti-spam, anti-virus, content filters, outbreak filters, data loss prevention)

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 1

Describing the Cisco Email Security Solutions

Overview

This lesson illustrates the SMTP flows and conversations and provides a high level overview of the Cisco Email Security Appliance services.

Upon completing this lesson, you will be able to meet these objectives:

- Identify how to place Cisco WSA into a typical modularized network design
- Describe the Cisco ESA Hybrid Solution
- Describe the SMTP terms
- Describe the email flow
- Describe the email conversation
- Describe the Cisco ESA services
- Describe the Cisco ESA GUI management
- Describe how to commit configuration changes on the Cisco Email Security Appliance
- Describe the Cisco Email Security Appliance licensing requirements
- Describe the incoming mail processing
- Describe the incoming mail processing
- Explain Cisco ESA to LDAP Integration functions
- Explain the Cisco Registered Envelope Service

https://t.me/learningnets

Cisco Modular Network Architecture and Cisco ESA

This topic identifies how to place Cisco Email Security Appliance into a typical modularized network design.

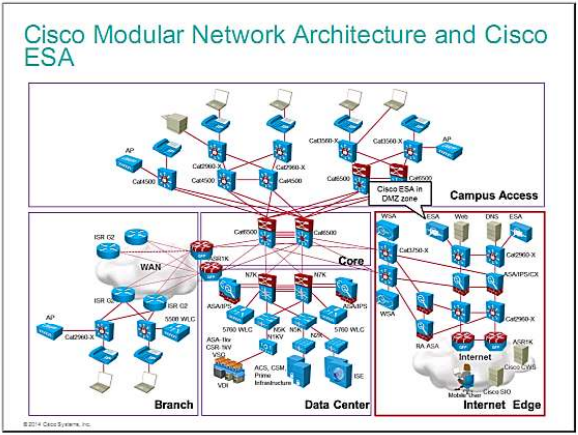

The figure illustrates where in the Cisco modular network architecture the Cisco ESA should be deployed. Ideally, the Cisco ESA should be implemented in a firewall demilitarized zone of Cisco modular network architecture.

E-mail is a medium through which spyware and viruses can be propagated. In addition to outside threats, E-mail spam and malicious malware can reduce employee productivity. The ESA is a type of firewall and threat monitoring appliance for SMTP traffic (TCP port 25).

Cisco ESA solution delivers:

- Fast, comprehensive email protection that can block spam and threats before they even hit your network
- Flexible cloud, virtual, and physical deployment options to meet your ever-changing business needs
- Outbound message control through on-device DLP and email encryption

Cisco's ESA solution offers simple, fast deployment, with few maintenance requirements, low latency, and low operating costs. The solution automatically forwards security updates to Cisco's cloud-based threat intelligence solution. This threat intelligence data is pulled by the Cisco ESAs every three to five minutes, providing you with threat defense.

The Cisco Email Security Virtual Appliance significantly lowers the cost of deploying email security, especially in highly distributed networks. The appliance lets your network manager create instances where and when they are needed, using your existing network infrastructure. The Cisco ESAV is a software version of the Cisco ESA and runs on top of a VMware ESXi hypervisor and Cisco Unified Computing System servers.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Logically speaking, the Cisco ESA acts as a MTA within the E-mail delivery chain. There are multiple deployment approaches for the security appliance depending on the number of interfaces used. The Cisco ESA may be deployed with a single physical interface to transfer emails to and from both the Internet and the internal mail servers. If desired, two physical interfaces may be used, one for email transfer to and from the Internet, and another one for email communications to the internal mail servers. In the former approach, the Cisco ESA would reside on the DMZ, while in the later, the Cisco ESA would have an interface connecting to the DMZ and the other one connecting to the inside network. In this case, the DMZ-based interface would send and receive E-mail to and from the Internet. The inside network interface would be used to deliver E-mail to the internal mail servers.

You can deploy Cisco Email Security solutions:

On premises: The Cisco ESA is an email gateway typically deployed in a firewall demilitarized zone. Incoming SMTP traffic is directed to the Cisco ESA data interface according to specifications set by your mail exchange records. The Cisco ESA filters it and redelivers it to your network mail server. Your mail server also directs outgoing mail to the Cisco ESA data interface, where it is filtered according to outgoing policies and then delivered to external destinations.

Virtual: With Cisco UCS running in your small branch office, you could host the Cisco ESAV with other Cisco products such as the Cisco Web Security Virtual Appliance. Together, they provide the same level of protection as their hardware equivalents but save you money on space and power resources.

Hybrid: Cloud Based Email Security for Inbound + On-Premise Email Security Appliance for Outbound

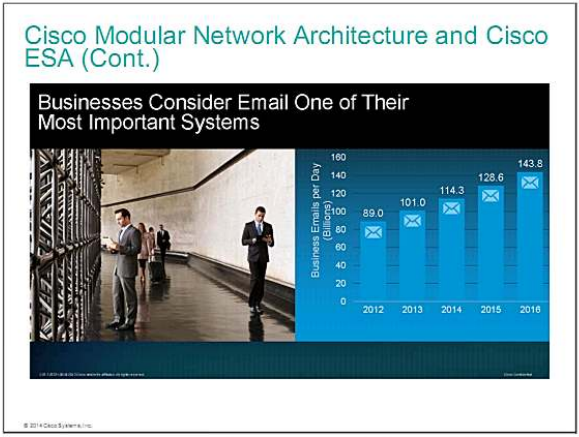

Over the past 20 years, email has evolved from a tool used primarily by technical and research professionals to become the backbone of corporate communications. Each day, more than 100 billion corporate email messages are exchanged. As the level of use rises, security becomes a greater priority. Mass spam campaigns are no longer the only concern. Today, spam and malware are just part of a complex picture that includes inbound threats and outbound risks.

The two major threats to your organization's email system are:

A flood of unsolicited and unwanted email, called spam, that wastes employee time through sheer volume and uses valuable resources like bandwidth and storage.

Malicious email, which comes in two basic forms: embedded attacks, which include viruses and malware that perform actions on the end device when clicked, and targeted or directed attacks, such as phishing attacks, which try to mislead employees into releasing sensitive information like credit card numbers, social security numbers, or intellectual property. Phishing attacks might direct employees to inadvertently browse malicious websites that distribute additional malware to computer endpoints.

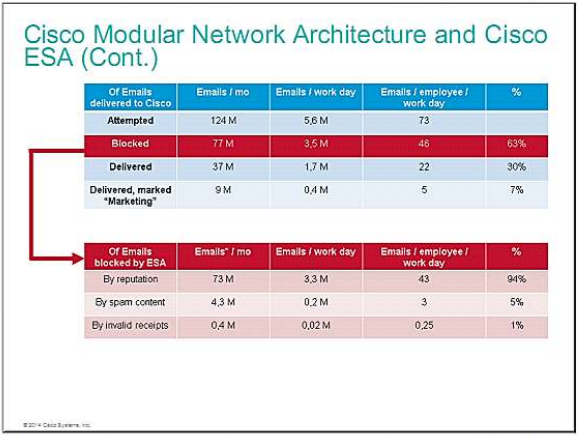

The Cisco ESA protects the email infrastructure and employees who use email at work by filtering unsolicited and malicious email before it reaches the user.

The Cisco ESA uses the Cisco Security Intelligence Operations. The Cisco SIO detects and correlates threats in real time using the largest threat detection network in the world. It monitors 100 TB of daily security intelligence, 1.6 million deployed security devices, 13 billion daily web requests, and 35 percent of worldwide email traffic. Cisco SIO prevents zero-hour attacks by continually generating new rules that feed updates to the Cisco ESA. These updates occur every 3-5 minutes to provide industry-leading threat defense.

The figure shows some sample statistics gathered from the Cisco ESA deployed at Cisco, showing that 63% of the emails delivered to Cisco were blocked by Cisco ESA and a vast majority of the emails were blocked based on the sender reputation. Reputation filters provide the first layer of defense by looking at the source IP address of the email server and comparing it to the reputation data downloaded from Cisco SIO. The Cisco SIO is one of the world's largest repository for security data, including sources of spam, botnets, and other malicious hosts. When hosts on the Internet engage in malicious activity, the Cisco SIO lowers the reputation of that host.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Hybrid Email Security Solution Overview

This topic describes the Cisco Email Security Appliance Hybrid Solution

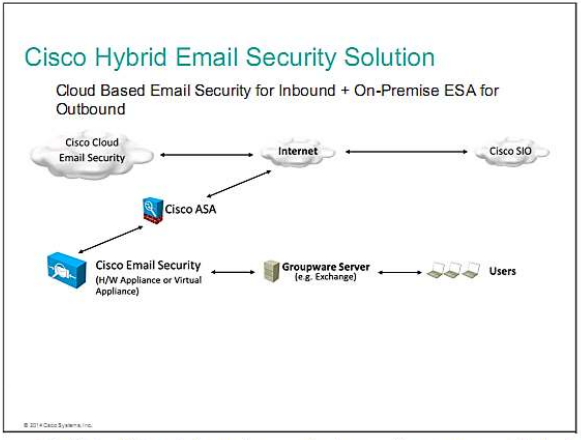

Cisco IronPort Hybrid Email Security is a unique service that provides customers with the choice to deploy email security in the way that best meets their business needs.

The Cisco IronPort Hybrid Email Security architecture includes email security infrastructure divided between cloud-based and on-premises form factors. It is a unique service offering that provides maximum flexibility by dividing control between the customer site and the Cisco cloud-based, SaaS infrastructure.

Organizations typically deploy inbound email security controls (anti-spam, anti-virus, etc...) through a hosted cloud infrastructure while protecting sensitive outbound email with encryption and data loss protection via on-premises email security appliances.

SMTP Terminologies

This topic describes the SMTP terms.

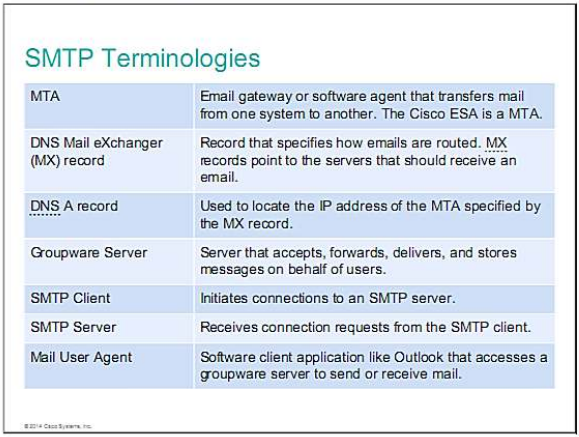

These terms are necessary for a discussion involving mail transfer:

MTA: The MTA also called a message transfer agent, or SMTPD, is a computer program or software agent that transfers electronic mail messages from one computer to another. Since personal computers do not send mail between themselves directly, they use client applications like MS Outlook which send mail to a groupware server, which then relay their mail through the MTA to some other mail domain. Another name for an MTA is an email gateway. The Cisco Email Security Appliance occupies this position in the network.

DNS MX Record: An MX record, or mail exchanger record, is a type of resource record in the DNS specifying how Internet email should be routed. MX records point to the servers that should receive an email, and their priority relative to each other. The DNS uses MX records to locate the IP address of an MTA.

DNS A Record: Used to locate the IP address of the MTA specified by the MX record.

Groupware Server: A server that accepts, forwards, delivers, and stores messages on behalf of users who only need to connect to the email infrastructure. Additionally, it manages collaborative schedules, maintains Calendars, and performs other services relating to members interacting within a group.

SMTP Client: Will initiate the connection request to an SMTP server that is located within the same enterprise or out on the Internet.

SMTP Server: Will receive the connection request from the SMTP client. Our discussion of transferring mail from a groupware server to an MTA and then on to another MTA represents an exchange of roles between SMTP client to SMTP server.

Mail User Agent: The Mail User Agent is a software client application like Outlook that accesses a groupware server, for example, an exchange server, to send or receive mail.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

SMTP Flow

This topic describes a SMTP flow.

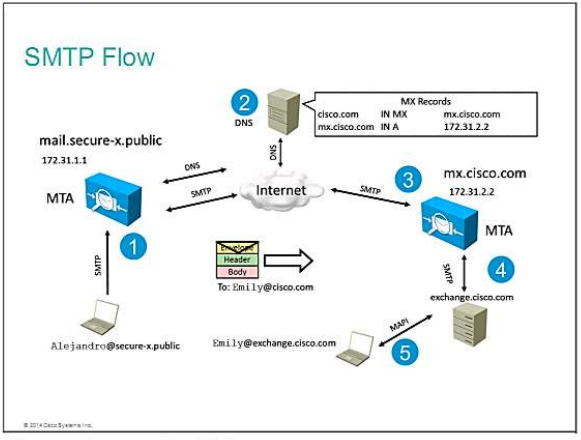

This figure illustrates the stages of mail delivery:

1. The MTA located at secure-x.public receives an email from joe@secure-x.public to Sam@cisco.com. For our discussion, this MTA is known as the sending MTA.
2. The sending MTA needs a destination IP address for the mail packet's IP header, so it examines the domain portion of the recipient email address (the part that follows the @) and performs a query of DNS MX records for www.cisco.com. The MX record points to the FQDN of the public interface of the MTA at the company Cisco. The lookup then proceeds to the A record for the server. This A record returns the public IP address of the SMTP Server.
3. The SMTP client forwards the packet to the SMTP server IP address. The receiving MTA acts as the SMTP server in this session. The receiving MTA receives the request to accept the session. It will check the host's address (smtp.secure-x.public) against a HAT to obtain this sender's reputation.
4. Assuming that the reputation of smtp.secure-x.public is good, the receiving MTA performs an LDAP lookup and forwards the email to the mail server.
5. The mail server then communicates with the mail user agent on the user's computer through POP, IMAP, and MAPI .

SMTP Conversation

This topic describes an SMTP conversation.

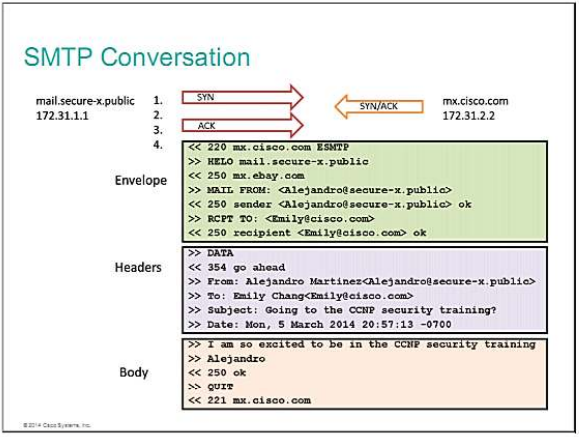

The figure shows the three parts of the SMTP conversation. The events occur sequentially beginning with the envelope and ending with the message content. The SMTP envelope specifies the recipient and the sender.

Headers

The headers are used by the user agents. They are sent after receiving a 354 status code from the SMTP server. They contain the following information: sender's display name and email, the recipient's display name and email, and the subject and date. A blank line separates headers from any message content.

Message Content

The body is an optional text region between the header and the one line containing a period (.). The terms "body," "message content," and "mail data" are used interchangeably. It is the material transmitted after the **DATA** command is accepted and before the end of data indication is transmitted. A period (.) on one line indicates the end of data transmission.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

SMTP Commands

SMTP commands transfer requests from the client to the SMTP Server. The most common commands are:

HELLO (HELO) or EHLO commands are used to identify the SMTP client to the SMTP server. The argument field contains the FQDN of the SMTP client if one is available.

MAIL FROM command is used to initiate a mail transaction in which the mail data is delivered to an SMTP server which may, in turn, deliver it to one or more mailboxes. In general, the **MAIL** command may be sent only when no mail transaction is in progress. It specifies who the mail originator is.

RCPT TO command is used to identify an individual recipient of the mail data. Multiple recipients are specified by multiple use of this command. The argument field contains a forward path and may contain optional parameters.

DATA. The receiver normally sends a 354 response to DATA, and then treats the lines (strings ending in **<CRLF>** sequences) as mail data from the sender. This command causes the mail data to be appended to the mail data buffer.

QUIT command specifies that the receiver must send an OK reply, and then close the transmission channel. The receiver must not intentionally close the transmission channel until it receives and replies to a **QUIT** command (even if there was an error). The sender must not intentionally close the transmission channel until it sends a **QUIT** command and should wait until it receives the reply (even if there was an error response to a previous command).

Three-Digit Response Codes

The three-digit response codes define server response to the SMTP client:

The first digit denotes whether the response is good, bad, or incomplete:

- 2 success (250 OK)
- 3 okay so far (354 Start mail input)
- 4 temporary failure (452 mailbox full)
- 5 permanent failure (550 user unknown)

The second digit categorizes the result described by the first digit:

- 0 syntax
- 1 connection
- 5 mail system

The third digit adds finer detail.

Reply Codes

Replies to SMTP commands serve to ensure the synchronization of requests and actions in the process of mail transfer and to guarantee that the SMTP client always knows the state of the SMTP server. Every command must generate exactly one reply. The following are examples of reply codes:

220 (FQDN of Server) Simple Mail Transfer Service Ready

250 OK

452 Too many recipients

Cisco ESA Services Overview

This topic describes the Cisco Email Security Appliance services.

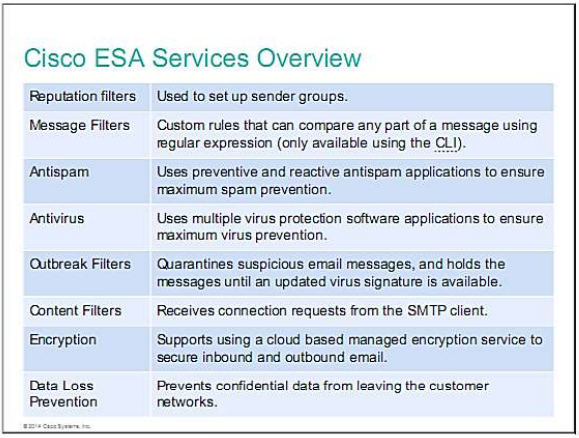

This table lists the main security services provided by the Cisco ESA.

Reputation Filters

The SBRS is similar to a credit reporting service for email. It protects legitimate senders from spam sources. SBRS allows identification and blocking of IP addresses from unsolicited commercial email or authentication of legitimate incoming email. Sender groups are used to associate one or more senders with a group, upon which you can apply message filters and other mail flow policies.

Message Filters

Message filters allow you to create access policies to handle messages and attachments as they are received. They identify messages based on message or attachment content, information about the network, message envelope, message headers, or message body. Message filters provide a script like interface for custom rules and can compare any part of a message with regular expression patterns. Message filters configuration is only available using the CLI, it is not available in the GUI. This course does not cover how to configure message filters using the CLI, please refer to the latest Cisco ESA User Guide for more information.

Antispam

Antispam scanning offers complete, Internetwide, server-side antispam protection. It stops spam attacks before they slow or overwhelm your network. You can remove unwanted mail before it reaches your users' inboxes, without violating their privacy. Antispam scanning can be configured to deliver mail to the Spam Quarantine (either on- or off-box). Messages released from the Spam Quarantine proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

Antivirus

The Cisco ESA has integrated virus scanning engines. You can configure it to scan messages and attachments for viruses on a per-mail policy basis.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Content Filters

Content filters are similar to message filters but are applied after the message has undergone message filters processing and anti-spam and anti-virus scanning. The Cisco ESA uses content filters to scan messages on a per-user (sender or recipient) basis. A content filter is limited to scanning either incoming or outgoing messages. You cannot define a filter that scans both types of messages.

Outbreak Filters

The Outbreak Filters feature includes special filters that act proactively to provide a critical first layer of defense against new outbreaks. Based on Outbreak Rules published by Cisco, messages with attachments of specific filter types can be sent to a quarantine named Outbreak. Messages in the Outbreak quarantine are processed according to newly updated rules. This processing may mean that they are released for delivery, or sent back to antivirus for rescanning. Outbreak filters can be configured to detect targeted phishing attacks, and can modify the URLs inside the message so that malicious sites are blocked to the recipients.

Data Loss Prevention

To ensure compliance with industry and government regulations worldwide and help prevent confidential data from leaving customer networks, DLP offers comprehensive protection, and unparalleled accuracy in identifying confidential data in outgoing mail traffic.

Encryption

To satisfy the compliance requirements to secure messaging. Meet the encryption requirements for regulatory requirements such as PCI, HIPAA, SOX, and GLBA — as well as state privacy regulations and European directives — without burdening the senders, recipients, or email administrators. Offer encryption not as a mandate, but as a service that's easy to use. Give the sender complete control of their content, even after it's been sent. With Cisco's email encryption, senders don't fear mistyped recipient addresses, mistakes in content, or time-sensitive emails because the sender always has the option to lock the message.

To take advantage of the most advanced cloud-based encryption key service available today, manage the recipient registration, authentication, and per-message/per-recipient encryption keys using the Cisco Registered Envelope Service. Cisco Registered Envelope Service provides all user registration and authentication as a highly available managed service. There's no additional infrastructure to deploy. For enhanced security and reduced risk, message content goes straight from your gateway to the recipient.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Cisco ESA GUI Management

This topic describes the primary management method of the Cisco Email Security Appliance.

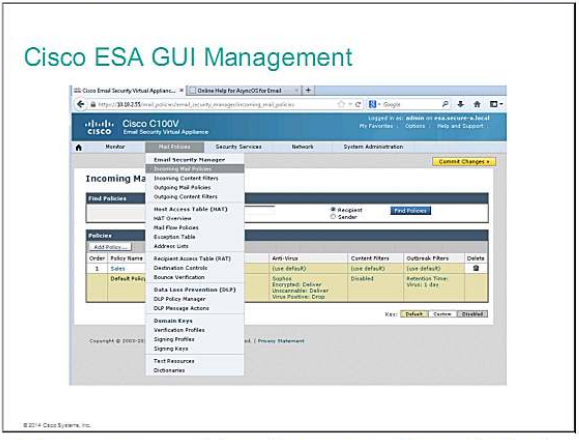

The Cisco ESA is an easy-to-manage platform with a management GUI organized into these main menus:

Monitor: The Monitor section contain pages for the Mail Flow Monitor feature (Overview, Incoming Mail, Outgoing Destinations, Outgoing Senders, Delivery Status, Internal Users, Content Filters, Virus Outbreaks, Virus Types, System Capacity, System Status), Local and External Quarantines, and Scheduled Reports features. You can also access message tracking from this menu.

Mail Policies: The Mail Policies section contains pages for the Email Security Manager feature (including Mail Policies and Content Filters), the HAT and RAT configuration, Destination Controls, Bounce Verification, Domain Keys, Text Resources, and Dictionaries.

Security Services: The Security Services section contains pages to set global settings for the Anti-Spam, Anti-Virus, Cisco Email Encryption, Outbreak Filters, and SenderBase Network Participation features. You also enable the following features from this menu: Reporting, Message Tracking, External Spam Quarantine.

Network: The Network section contains pages for creating and managing IP interfaces, Listeners, SMTP Routes, DNS, Routing, Bounce Profiles, SMTP Authentication, and Incoming Relays.

System Administration: The System Administration section contains pages for the Trace, Alerting, User Management, LDAP, Log Subscription, Return Addresses, System Time, Configuration File management, Feature Key Settings, Feature Keys, Shutdown/Reboot, Upgrades, and System Setup Wizard features.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA Committing the Configuration Changes

This topic describes how to commit configuration changes on the Cisco Email Security Appliance.

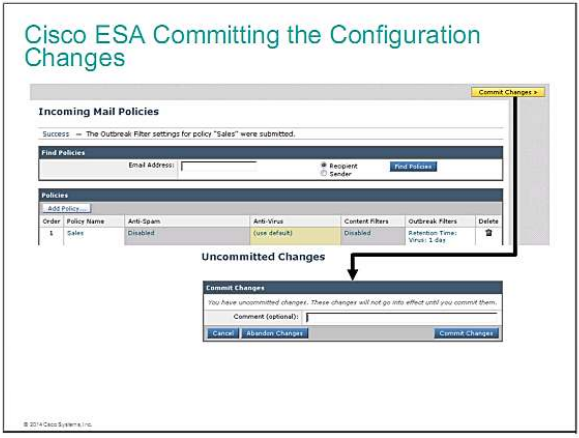

When you change the configuration of the Cisco ESA using the GUI, you must commit the changes before they go into effect. Or, you can choose to clear the changes you have made if you do not want to commit them.

Commit changes in the Web GUI using the **Commit Changes** button in the upper right corner of the Web interface. You can make multiple configuration changes before you commit all of them. When you make a change, the **Commit Changes** button color is yellow.

To commit changes made in the web interface:

1. Click the **Commit Changes** button.
2. The **Uncommitted Changes** page appears. Enter comments in the Comment field if you choose. Click the **Commit Changes** button.
3. When there are no changes to commit, the button color is gray and the button text is **No Changes Pending**.

Commit changes in the CLI using the **commit** command. Most configuration changes you make in the CLI are not effective until you issue the **commit** command.

Cisco ESA Licensing

This topic describes the Cisco Email Security Appliance licensing requirements.

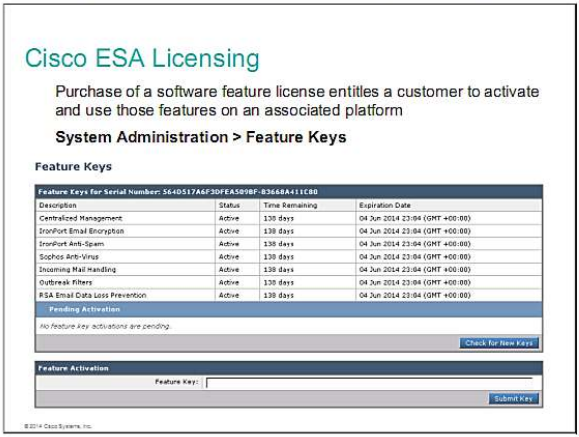

Purchase of a software feature license entitles a customer to activate and use those features on an associated platform. For example, customers who buy a Cisco Email Security Inbound Bundle license are entitled to activate all of the Inbound bundle software features on any of the Cisco ESA that they own.

All software feature licenses are term based subscriptions of length 1, 3, or 5 years. Each software feature or feature set has corresponding SKUs for the three term options.

The Cisco ESA uses tiered quantity based pricing, with the user count entered as the applicable quantity.

System Administration > Feature Keys is where the license keys for the different features on the Cisco ESA are displayed.

Three Cisco Email Security software licenses are available: Cisco Email Security Inbound, Cisco Email Security Outbound, and Cisco Email Security Premium. The major components of each software offering are provided below:

Cisco Email Security Inbound Bundle

- Antispam scanning
- Sophos Antivirus solution
- Virus outbreak filters
- Clustering

Cisco Email Security Outbound Bundle

- DLP compliance
- Email encryption
- Clustering

Cisco Email Security Premium Bundle

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

- Antispam scanning
- Sophos Antivirus solution
- Virus outbreak filters
- DLP compliance
- Email encryption
- Clustering

https://t.me/learningnets

Incoming Mail Processing Overview

This topic describes the incoming mail processing on the Cisco Email Security Appliance.

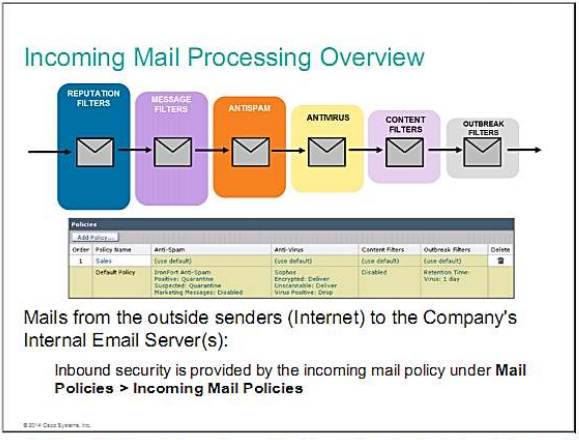

Inbound security is provided by the incoming mail policy. This policy includes six layers of filters. The following list shows the order in which they are applied to a message traversing the appliance from an external source.

1. **Threat prevention with Reputation Filters:** Reputation filtering is the first layer of spam protection, allowing you to control the messages that come through the email gateway based on senders' trustworthiness as determined by the Cisco SenderBase Reputation Service. When enabling reputation filtering, incoming mails from known bad senders is simply refused. Known good mail from global companies is automatically routed around the spam filters, reducing the chance of false positives. Unknown, or "grey" email is routed to the anti-spam scanning engine. Using this approach, reputation filters can reduce the load on the content filters. Reputation Filtering Cisco ESA service is controlled using the HAT Sender Group SenderBase Reputation Service setting. The HAT maintains a set of rules that control the incoming SMTP connections to the Cisco ESA.
2. **Policy enforcement with Message Filters:** Message filters are special rules describing how to process messages and attachments as they are received using a script like interface with regular expression (only available in the CLI).
3. **Spam detection with Anti-spam and the Context Adaptive Scanning Engine**:
   - **Email reputation:** Who is sending this message?
   - **Message content:** What content is included in this message?
   - **Message structure:** How was this message constructed?
   - **Web reputation:** Where does the call to action take you?
4. **Virus detection with Sophos and/or McAfee Antivirus:** Antivirus provides a virus detection engine that scans files for viruses, Trojan horses, and worms.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

5. **Content Filters:** These filters can be used to filter specific file types or content. Some common reasons for using content filters include: filtering attachments, adding disclaimers, rerouting messages to other systems and etc...
6. **Outbreak Filters:** Newly released viruses that do not have a published ID can be blocked by stopping files with the infected file's characteristics. This feature provides the day zero protections.

The Cisco ESA GUI is used to create the incoming mail policies that checks for virus outbreaks, spam, and content of email messages on a per-recipient or per-sender basis.

Antispam, antivirus, content filters and outbreak filters Cisco ESA services are controlled using incoming mail policies, which are configurable using the GUI **Mail Policies > Incoming Mail Policies** screen.

Specific user based incoming mail policies can be defined based on email address, email domains (recipient or sender), or LDAP group queries.

https://t.me/learningnets

Outgoing Mail Processing Overview

This topic describes the outgoing mail processing on the Cisco Email Security Appliance.

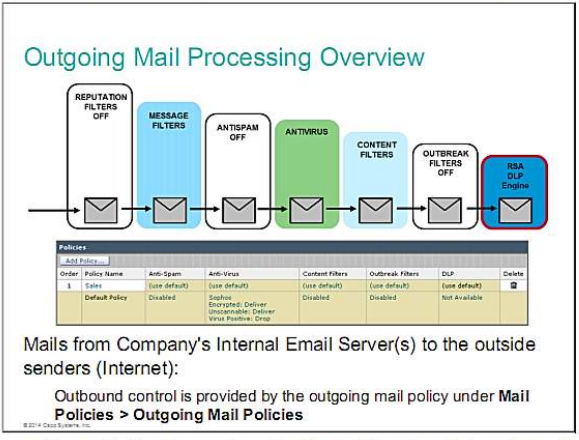

Outbound control is provided by the outgoing mail policy. It differs from the incoming mail processing by:

- The default outgoing mail policy has anti-spam, content filters and outbreak filters disabled by default. Anti-spam, content filters and outbreak filters can be enabled by editing the default outgoing mail policy.
- Provides an additional DLP function to the process to ensure unsuitable or unauthorized information does not leak out of the company. Data Loss Prevention scanning can only be performed on outgoing messages.
- Reputation Filtering is also not used for the outgoing mails, it is only used for the incoming mails.

Anti-spam, anti-virus, content filters, outbreak filters and DLP Cisco ESA services are controlled using the outgoing mail policies, which are configurable using the GUI **Mail Policies > Outgoing Mail Policies** screen.

Specific user based outgoing mail policies can be defined based on email address, email domains (recipient or sender), or LDAP group queries.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA LDAP Integration Overview

This topic explains the reasons for integration the Cisco Email Security Appliance to LDAP.

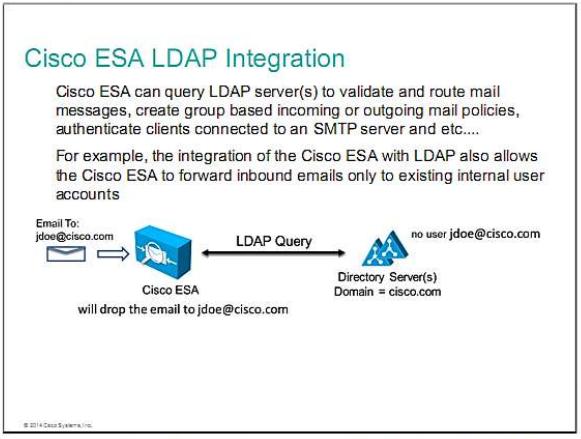

If the user information is stored within LDAP directories in the network infrastructure — for example, in Microsoft Active Directory or OpenLDAP directories, the Cisco ESA can use different types of LDAP queries to accept recipients, route messages, masquerade headers, autheticate the SMTP clients, create group based incoming or outgoing mail policies and etc... LDAP queries can also be used in conjunction with message filters to create rules for handling messages as they are received by the Cisco ESA.

The Cisco ESA can be configured to work with one or multiple LDAP servers.

For examples:

Use LDAP accept query to determine if a recipient or sender is a member of a group as defined by your LDAP directory. Incoming and outgoing mail policies then can reference the LDAP group to match the recipent or sender.

Configure the Cisco ESA to perform actions on messages based on the groups in the LDAP directory. You do this by associating a LDAP group query with a message filter. You can perform any message action available for message filters on messages that match the defined LDAP group.

Configure the Cisco ESA to use LDAP routing query to route messages to the appropriate address and/or mail host based upon the information available in LDAP directories on the network. The Cisco ESA supports alias expansion (LDAP routing with multiple target addresses) where it replaces the original email message with a new, separate message for each alias target (for example, recipient@yoursite.com might be replaced with new separate messages to recipient1@internal.company.com and recipient2@internal.company.com, etc. Routing queries are sometimes known as aliasing queries on other mail processing systems.

Support SMTP Authentication to authenticate clients connected to an SMTP server using LDAP authentication query. The practical use of this mechanism is that users at a given organization are able to send mail using that entity's mail servers even if they are connecting remotely (e.g. from home or while traveling). Remote users using remote MUAs can issue an authentication request (challenge/response) when attempting to send a piece of mail.

Use LDAP Masquerade query to masquerade the envelope senders (for outgoing mail) and message headers (for incoming mail, such as To:, Reply To:, From: or CC:). Masquerading is a feature that rewrites the Envelope Sender (also known as the sender, or MAIL FROM) and the To:, From:, and/or CC: headers on email based on queries you construct. A typical example implementation of this feature is "Virtual Gateway". The Virtual Gateway technology enables users to logically separate a single Cisco ESA into multiple Virtual Gateway addresses from which to send and receive email. Each Virtual Gateway address is given a distinct IP address, hostname and domain, and email queue. Another typical implementation is "hiding" your network infrastructure by "stripping" the subdomains from strings in email headers.

Please refer to the latest Cisco ESA User Guide for more information about all the other usages and types of LDAP queries.

The figure shows an incoming email to the jdoe@cisco.com user, by using LDAP accept query, the Cisco ESA knows that user does not exist in the Active Directory and will therefore drop the incoming email to the jdoe@cisco.com user and not forward it to the cisco.com internal email server.

Note	This course does not cover the details on how to implement all the different types of LDAP queries. Please refer to the latest Cisco ESA User Guide for more information. The Cisco ESA lab you will perform later will also step you through how to configure the Cisco ESA to integration with LDAP and to enable Accept Query which allows the Cisco ESA to forward inbound email only to existing internal AD user accounts.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Registered Envelope Service (CRES) Overview

This topic explains the Cisco Registered Envelope Service

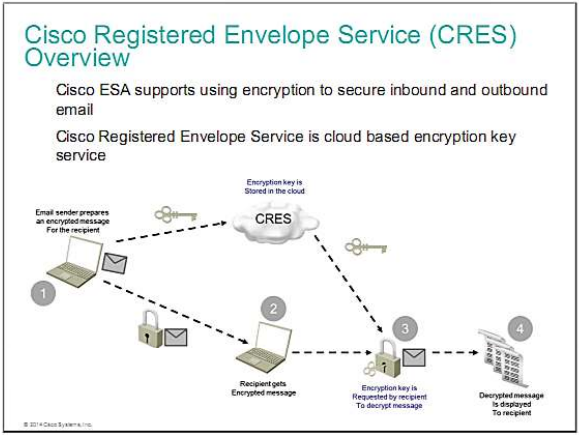

The Cisco ESA supports using encryption to secure inbound and outbound email. To use this feature, you create an encryption profile that specifies characteristics of the encrypted message and connectivity information for the key server. The key server may either be:

Cisco Registered Envelope Service - cloud based managed hosted service:

- The Cisco Registered Envelope Service is an advanced cloud based email encryption key service. Whether you need to meet compliance requirements, secure customer and partner communications, or protect intellectual property, the Cisco Registered Envelope Service provides a flexible and scalable solution to help you support your secure messaging requirements without having to invest in additional infrastructure. Although regular emails are not a secure communication medium, encryption and key management is often seen as too complex to be used in everyday communications. The Cisco Registered Envelope Service takes away the complexity behind encryption and makes it seamless for senders and recipients to exchange secure messages.
- Message management is provided as part of the business-class email features and includes message recall and expiration and guaranteed read-receipt. These features can be accessed through the Cisco Registered Envelope Service web interface or directly from the email client by installing dedicated plug-in.

Cisco Ironport Encryption Appliance - locally managed server:

- The Cisco Ironport Encryption Appliance has reached End of Sales.

The figure shows how simple it is for users to send and receive an encrypted e-mail using the cloud based Cisco Registered Envelope Service. The encrypted messages can be received by any email user - independent of the email client, the operating system, or the device used - without the need to install any software or requiring the sender to pre-exchange the encryption credentials with the recipients.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Note	This course does not cover the details on how to implement the encryption. Please refer to the latest Cisco ESA User Guide for more information.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

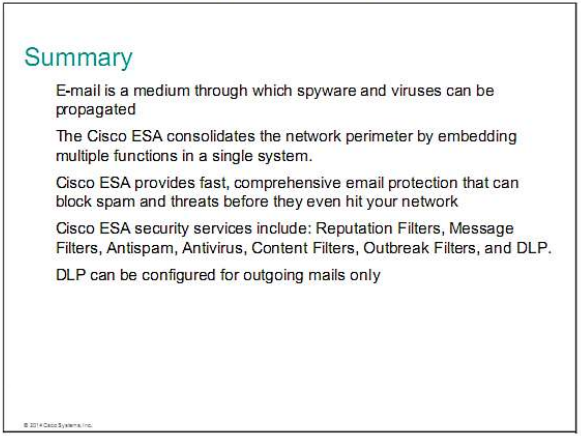

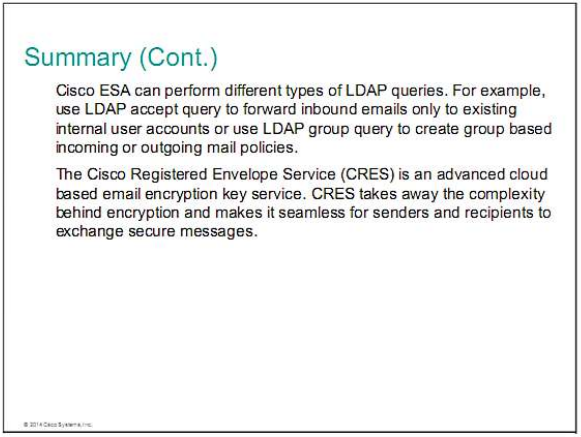

References

For additional information, refer to these references:

Email Security Using Cisco ESA Deployment Guide

http://www.cisco.com/en/US/docs/security/esa/esa7.5/ESA_7.5_Configuration_Guide.pdf

http://www.cisco.com/en/US/docs/solutions/SBA/February2013/Cisco_SBA_BN_EmailSecurityUsingCiscoESADeploymentGuide-Feb2013.pdf

Cisco IronPort AsyncOS for Email Configuration Guide

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 2

Describing the Cisco Email Security Appliance Basic Setup Components

Overview

The Cisco ESA easily integrates into existing email infrastructures with a high degree of flexibility. The Cisco ESA can be deployed with a single physical interface in order to filter email to and from an organization's mail server. The second deployment option is a two-interface configuration, one interface for email transfers to and from the Internet and the other for email transfers to and from the internal servers.

This lesson describe the basic configuration components to setup the Cisco ESA, which includes the listener, LDAP queries, HAT, RAT, Mail Flow Policies and SMTP Routes table.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the Cisco ESA listener
- Describe the two Cisco ESA listener types
- Describe a one interface one listener deployment example
- Describe a two interfaces two listeners deployment example
- Describe the Cisco Email Security Appliance Listener Components - HAT and RAT Overview
- Describe the Cisco Email Security Appliance scenario using one listener
- Describe how to set IP address on interface and how to create listener on the Cisco Email Security Appliance
- Provide an example of LDAP integration using LDAP Query to drop emails to a nonexisting users
- Describe the Cisco ESA HAT
- Describe the Cisco ESA Sender Group HAT component

https://t.me/learningnets

Describe the Cisco ESA Sender Group SBRS HAT component

Describe the Cisco ESA BLACKLIST Sender Group HAT component

Describe the Cisco ESA RELAYLIST Sender Group HAT component

Describe how to Add a Sender Group

Describe the Cisco ESA Mail Flow Policy

Describe the Cisco ESA Mail Flow Policy Examples

Describe the Cisco ESA Mail Flow Policies connection behavior, throttling and anti-spam and anti-vrius typical settings

Describe the Cisco ESA RAT

Describe the Cisco ESA SMTP Routes

Describe the Cisco ESA SMTP Routes

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA Listener Overview

This topic describes the Cisco Email Security Appliance listener.

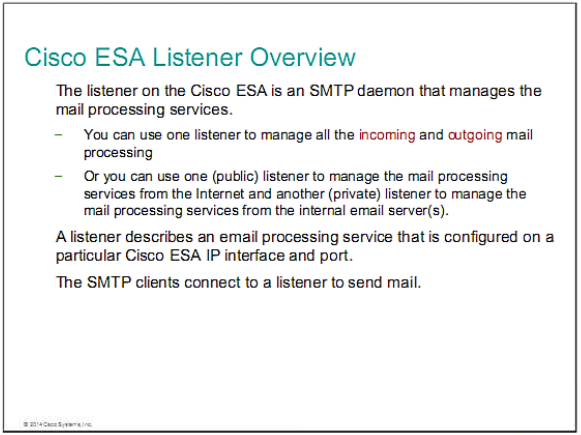

A listener, on the Cisco ESA, is an SMTP daemon that accepts incoming connections from SMTP clients on a particular interface and port. The listener waits for connection requests on the particular interface and port. Depending on the topology, a single listener can accept incoming SMTP requests from the Internet and from an internal mail server(s). Or you can configure listeners to specialize based on the direction; either from the Internet on a public listener or from an internal mail server on a private listener.

A Cisco ESA can be configured with one or more listeners. You can segregate the incoming and outgoing email traffic over separate listeners and on separate interfaces.

https://t.me/learningnets

Cisco ESA Listener Type: Private and Public

This topic describes the two Cisco Email Security Appliance listener types.

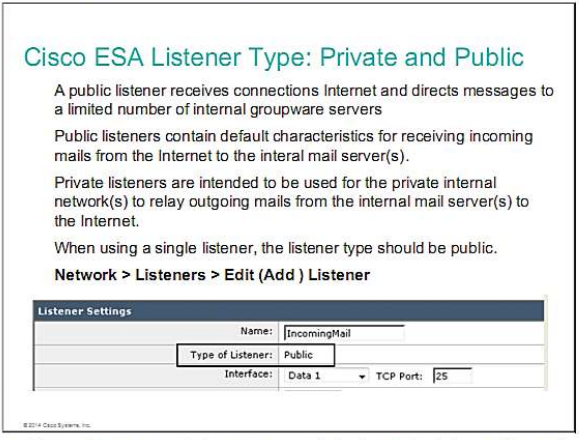

When you configure a listener, one of the parameter to define is whether it is a public or private listener.

Technically, only one of the three interfaces available on the Cisco ESA is required to send and receive email. This interface can have one listener that is facing both public and private SMTP servers. However, depending on the network requirements, more interfaces and more listeners may be required, as in the two-interface deployment. In the two-interface deployment, the public interface is in the DMZ and the private interface is facing the inside groupware server.

A public listener receives connections from many external hosts (Internet) and directs messages to a limited number of internal groupware servers. Public listener contains default characteristics for receiving email from the Internet.

A private listener receives connections from a limited number of internal groupware servers and directs messages to many external mail hosts. Private listener is intended to be used for private (internal) networks.

When using a single listener, the listener type should be public.

The figure shows a single listener example where the public listener is named **IncomingMail**. This single public listener will be used for both receiving incoming mail from the Internet and for relaying outgoing mails from the internal mail server(s).

https://t.me/learningnets

Cisco ESA One Interface/One Listener Deployment Example

This topic describes a one interface one listener deployment example.

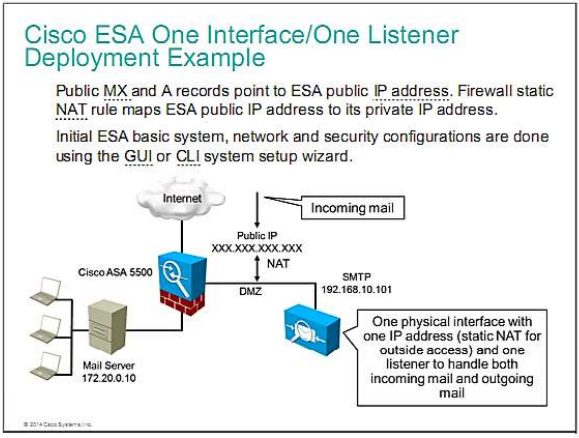

The figure shows a typical deployment where the Cisco Email Security Appliance is physically deployed on the DMZ of the Internet Edge firewall, and uses a single interface for simplicity. This interface handles all incoming and outgoing email and carries the management traffic (or you could use a separate management interface to implement an OOB management network). Usually, the customer's MX record points to a public address on the firewall. The firewall is configured with a static NAT rule to translate the Cisco ESA real (private) 192.168.10.101 IP address to the translated (public) IP address. The firewall's DMZ is a portion of the network where, typically, traffic to and from other parts of the network is tightly restricted. Organizations place network services in a DMZ for exposure to the Internet.

In this scenario, the typical firewall security policy should allow only SMTP traffic to Cisco ESA. The Cisco ESA is allowed to send SMTP traffic as well as make HTTP and HTTPS connections (needed for reputation updates) to the Cisco Security Intelligence Operations. Cisco ESA is allowed to make inbound SMTP connections to the corporate mail server as well as DNS requests to the organization's DNS server and requests to the LDAP servers if using LDAP integration.

In this deployment, one listener handles incoming and outgoing connections. It relays outgoing connections to external domains, and restricts incoming connections to internal mail domains. The incoming and outgoing connections to and from this one listener are identified by the session source IP address. Mail from internal source IP addresses will be processed under a different security criterion than mail arriving from public addresses in the outside networks. For mail arriving from the Internet, only messages with an **RCPT To** field that matches an internal mail domain will be processed by the filters and forwarded internally.

Instead of using one interface, you can configure two Ethernet interfaces and segregate your internal network from your external Internet network connections (not shown in the figure).

Note

The Cisco ESA ships with a default IP address of 192.168.42.42 with HTTP for the GUI access enabled on either the Management interface or the Data 1 interface depending on the Cisco ESA model. To access the GUI, browse to the http://192.168.42.42 URL. When the login page is displayed, log in to the system using the default **admin** username with the default **ironport** password.

The Cisco ESA GUI or CLI system setup wizard must be used to setup the Cisco ESA with the basic system, network and security configurations. Later, you can configure custom options not available in the system setup wizard.

Basic Cisco ESA configurations such as the system hostname, default gateway, NTP server, DNS and etc.. are not covered in this course. Please refer to the latest Cisco ESA user guide for more information.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA Two Interfaces/Two Listeners Deployment Example

This topic describes a two interfaces two listeners deployment example.

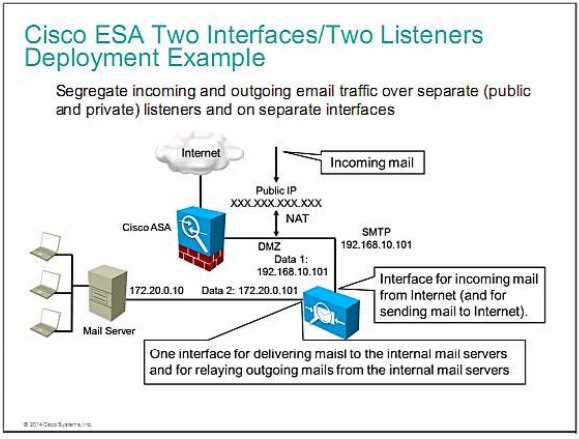

Depending on customer network needs, multiple data interfaces may be required. For example, a corporate security policy states that gateway security devices must reside in the DMZ with the restriction that internal servers and assets must be on separate networks (including internal mail servers) with a "security device" in between.

In the network diagram, the Data 1 interface is connected to the DMZ using the 192.168.10.101 IP address. The firewall is configured with a static NAT rule to translate the Cisco Email Security Appliance real (private) 192.168.10.101 IP address to the translated (public) IP address. This interface supports a public listener. It is a public listener because it is receiving email from the Internet. This public listener uses SMTP protocol on Port 25 of the Data1 Ethernet interface to accept messages from the Internet. This interface also sends messages to destination hosts on the Internet.

The Data 2 interface (172.20.0.101) is facing the internal mail server (172.20.0.10). It supports the private listener. The private listener will relay outgoing mail connections from the internal mail server. The private listener uses SMTP protocol on Port 25 of Data 2 Ethernet interface to accept messages from internal email server. This interface also sends the incoming messages to internal mail server.

Cisco ESA Listener Major Components: HAT and RAT

This topic describes the Cisco Email Security Appliance Listener Components - HAT and RAT.

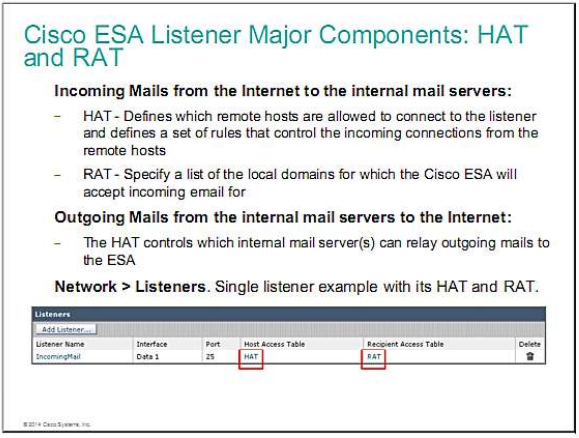

The figure shows an example where the Cisco ESA is configured with a single public listener on the Cisco ESA named "IncomingMail". This listener is configured on the Cisco ESA Data 1 Interface and is listening for SMTP traffic on TCP port 25.

The listener controls whether the Cisco ESA accepts or rejects the SMTP connections and at what rates the Cisco ESA will accept the SMTP connections using the HAT and the RAT.

HAT is used to specify which remote hosts that are allowed to connect to a listener (that is, the mail senders allowed to send incoming mails). This setting is done by determining the reputation of the sending host IP address, and applying the appropriate action of accepting or rejecting the connection.

RAT controls whether the Cisco ESA will accept the incoming mails for a given recipient or recipient domain. For inbound email only, the RAT allows you to specify a list of all local domains for which the Cisco ESA will accept incoming mails for.

The HAT and RAT together control the incoming mails from the Internet.

The HAT is also used to control the outgoing mail by controlling which internal mail server(s) can relay outgoing mails to the Cisco ESA. The Cisco ESA relays the outgoing mails to the Internet only if it is received from the internal mail server(s) specified in the HAT (using the **RELAYLIST** sender group configurations in the HAT).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA One Listener Deployment Scenario

This topic describes the Cisco Email Security Appliance scenario using one listener.

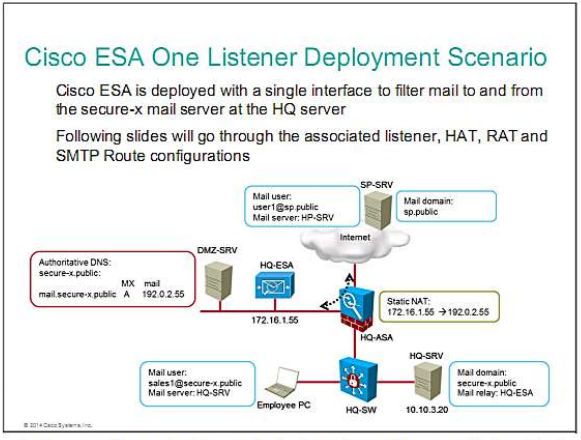

In the scenario, the access list on the Cisco Adaptive Security Appliance outside interface is configured to allow SMTP traffic from the outside to the Cisco ESA and the access list on the DMZ interface is configured to allow the required traffic from the Cisco ESA to the internal HQ Server which is the secure-x mail server. The Cisco ESA acting as the MTA is deployed using a single interface to filter mail to and from the secure-x mail server at the HQ server.

The figure shows how the Cisco ESA is deployed. It shows high level configuration of DNS and mail settings on mail servers. The DNS server on the DMZ server is configured as authoritative server for the **secure-x.public** domain. When the Internet mail server wants to send an email to the **secure-x.public** domain, it queries the DNS server on the DMZ server for MX record of the **secure-x.public** domain. The MX record points to the **mail.secure-x.public** host, for which the A record points to the translated 192.0.2.55 IP address of the Cisco ESA. The Cisco ASA is configured with a static NAT rule to translate the real 172.16.1.55 (private) IP address to the 192.0.2.55 translated (public) IP address. The internal mail server on the HQ server (10.10.3.20) is configured with a mail relay which points to the Cisco ESA. The internal mail server uses the mail relay (Cisco ESA) to send the outgoing mails from the secure-x domain.

The Cisco ESA also uses out-of band management interface with the 10.10.2.55 IP address, which is not shown in the figure for the sake of better visibility.

The following slides will go through the Cisco ESA listener, HAT, RAT and SMTP Route configurations associated with this scenario where:

The Cisco ESA is configured with a single listener on the Data 1 interface listening on TCP port 25.

This single listener will accept incoming mails only for the **secure-x.public** domain:

- The Cisco ESA is configured to only accept incoming mails for the **secure-x.public** domain using the RAT.
- The HAT is used to categorized the mail senders to the **secure-x.public** domain into the different sender groups.

- Each sender group has an associated mail flow policy to define whether the remote hosts in the sender group are allowed to connect to the listener and under what conditions.

This single listener will relay the outgoing mails only from the 10.10.3.20 internal mail server to the Internet:

- The Cisco ESA is configured to only allow the 10.10.3.20 internal mail server to establish SMTP connections to the listener using the **RELAYLIST** Sender Group in the HAT.

A SMTP route is configured on the Cisco ESA for routing the mails to the internal mail server.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: Interfaces and Listener

This topic describes how to set IP address on interface and how to create listener on the Cisco Email Security Appliance.

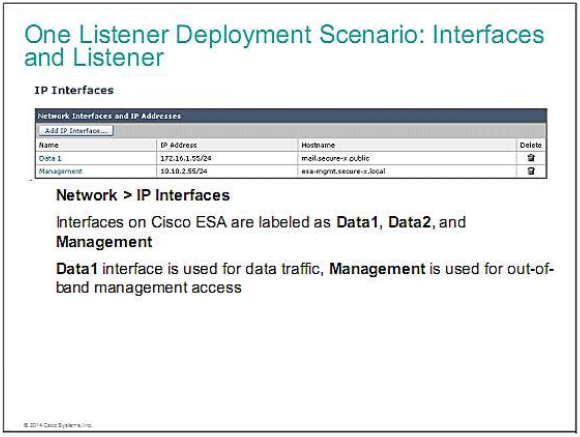

To add IP interface, navigate to **Network > IP Interfaces**. Interfaces on Cisco ESA are labeled as **Data1**, **Data2**, and **Management**. In our scenario, we will use two interfaces:

Data1: IP address of this interface is **172.16.1.55/24**. The Cisco ASA is configured with a static NAT rule to translate the real **172.16.1.55** (private) IP address to the **192.0.2.55** translated (public) IP address.

Management: This interface is used for out-of-band management. IP address of this interace is **10.10.2.55**.

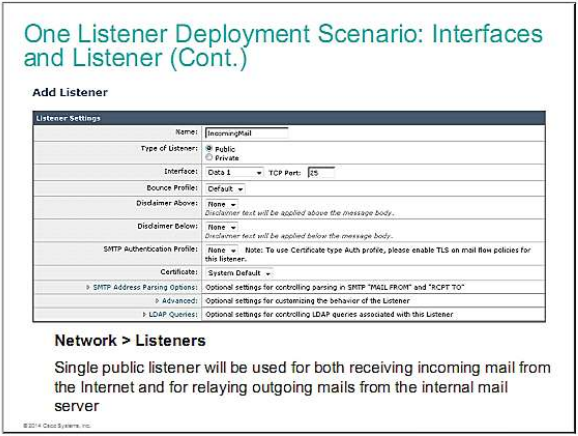

The figure shows a listener named **"IncomingMail"** which has the public listener type. In this scenario, the Cisco ESA is using a single listener on the Data 1 interface using TCP port 25. This single public listener will be used for both receiving incoming mail from the Internet and for relaying outgoing mails from the internal mail server.

Bounce Profile: The Cisco ESA classifies undeliverable email, or "bounced messages" into two categories: soft bounces and hard bounces. Soft bounces: A messages that is temporarily undeliverable. For example, a user's mailbox may be full. These messages can be retried at a later time. (e.g. An SMTP 4XX error code). Hard bounces: A message that is permanently undeliverable. For example, the user no longer exists for that domain. These messages will not be retried. (e.g. An SMTP 5XX error code).

Disclaimer Above and **Disclaimer Below**: Select a disclaimer to attach above or below emails (disclaimers created via the **Mail Policies > Text Resources** page or the **textconfig** command in the CLI.

SMTP Authentication Profile: Specify an SMTP Authentication profile. The Cisco ESA provides support for SMTP authentication. SMTP authentication is a mechanism for authenticating clients connected to an SMTP server. The Cisco ESA supports two methods to authenticate user credentials: use an LDAP directory or use a different SMTP server. The default is none as shown in the figure.

Certificate: Specify a certificate for TLS connections to the listener, certificates are added via the **Network > Certificates** page or the **certconfig** command in the CLI.

SMTP Address Parsing Options: Configure the settings for controlling parsing in SMTP "MAIL FROM" and "RCPT TO" commands.

Advanced: Configure the advanced settings for customizing the behavior of the listener.

LDAP Queries: Configure the settings for controlling LDAP queries associated with this listener. Use these settings to enable LDAP queries on the listener. You must create the LDAP query first, before using this option. Each type of query has a separate subsection to configure. Click the **type of query** to expand the subsection.

Note	Most of these configurations are beyond the scope of this course, please refer to the latest Cisco ESA User Guide for more information.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

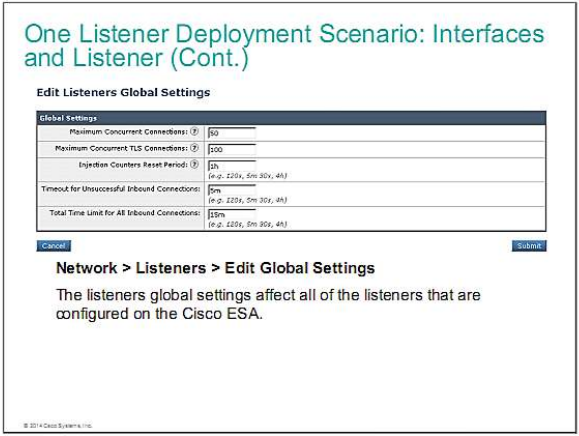

The global settings for the listeners affect all of the listeners that are configured on the Cisco ESA. If the listener uses an interface that has both IPv4 and IPv6 addresses, the listeners global settings apply to both the IPv4 and IPv6 traffic.

Maximum Concurrent Connection: Set the maximum number of concurrent connections for listeners. The default value is 50. If the listener accepts both IPv4 and IPv6 connections, the number of connections is divided between the two.

Maximum Concurrent TLS Connections: Set the maximum concurrent TLS connections across all listeners combined. The default value is 100. If the listener accepts both IPv4 and IPv6 TLS connections, the number of connections is divided between the two.

Injection Counters Reset Period: Allows you to adjust when the injection control counters are reset. For very busy systems maintaining counters for a very large number of different IP addresses, configuring the counters to be reset more frequently (for example, every 15 minutes instead of every 60 minutes) will ensure that the data does not grow to an unmanageable size and impact system performance. The current default value is 1 hour. You can specify periods ranging from as little as 1 minute (60 seconds) to as long as 4 hours (14,400 seconds).

Email injection is a security vulnerability that can occur in Internet applications that are used to send email messages. It is the email equivalent of HTTP Header Injection. Like SQL injection attacks, this vulnerability is one of a general class of vulnerabilities that occur when one programming language is embedded within another. When a form is added to a Web page that submits data to a Web application, a malicious user may exploit the MIME format to append additional information to the message being sent, such as a new list of recipients or a completely different message body.

Timeout Period for Unsuccessful Inbound Connections: Set the length of time the Cisco ESA will allow an unsuccessful inbound connection to remain intact before closing it. An unsuccessful connection can be an SMTP conversation in which SMTP or ESMTP commands continue to be issued without a successful message injection occurring. When the specified timeout is reached, the behavior is to send an error and disconnect: "421 Timed out waiting for successful message injection, disconnecting." A connection is considered unsuccessful until it successfully injects a message. Only available for SMTP connections on public listeners. The default value is 5 minutes.

Total Time Limit for All Inbound Connections: Set the length of time Cisco ESA will allow an inbound connection to remain intact before closing it. This setting is intended to preserve system resources by enforcing a maximum allowable connection time. Once about 80% of this maximum connection time is reached the following message is issued: "421 Exceeded allowable connection time, disconnecting." The appliance will attempt to disconnect when the connection exceeds 80% of the maximum connection time in order to prevent disconnecting mid-message. It is likely that a problem is occurring with the inbound connection if it is open long enough to reach 80% of the maximum connection time. Keep this threshold in mind when specifying the time limit. Only available for SMTP connections on public listeners. The default value is 15 minutes.

Note	Most of these configurations are beyond the scope of this course, please refer to the latest Cisco ESA User Guide for more information.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: LDAP Accept Query

This topic provides an example of the LDAP integration using LDAP Query to drop emails to a nonexisting users.

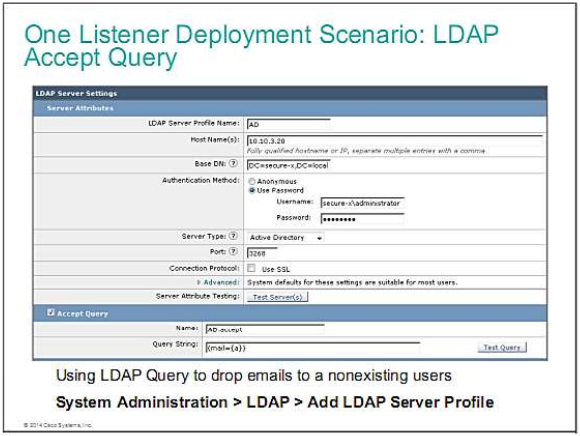

If you store the user information within the LDAP directories in your network infrastructure — for example, in Microsoft Active Directory, — you can configure the Cisco appliance to query your LDAP servers to accept, route, and authenticate messages. You can configure the Cisco appliance to work with one or multiple LDAP servers. On our deployment scenario, HQ Server with IP address 10.10.3.20 is configured as Microsoft Active Directory Server.

Define the LDAP Server:

LDAP Server Profile name: **AD** in our example.

Host Name(s): **10.10.3.20** in our example is the AD Server.

Base DN: **DC=secure-x**, **DC=Local** in our example. A base DN is the starting point for directory server searches.

Authentication Method: **Use Password** in our example, with the **secure-x\administrator** username and the password.

Server Type: Can be **Active Directory**, **OpenLDAP**, **Unknown or Other**. In our example, the directory server type is **Active Directory**.

Port: **3268**.

Connection Protocol: SSL is not used in our example.

Define the the LDAP Query - The Accept Query is defined this example:

Name: **AD.accept** in our example.

Query String: **(mail={a})** is the string used for LDAP Accept Query

- {a} = username@domainname

Test the LDAP Query (not shown in the figure):

Click **Test Query**

Enter the Test Parameters: Envelope Recipient Address and the optional Envelope Sender Address

Click **Run Test**

Perform similar step to enable the other types of LDAP queries such as LDAP group query. The LDAP Active Directory group query string used to determine if a user is a member of the AD group is: **(&(memberOf={g})(proxyAddresses=smtp:{a}))** (not shown in the figure).

Group query allows you to query the LDAP for group information so you can create group based incoming or outgoing mail policies. In the incoming or outgoing mail policies, you can match on **Recipient: AD.group: IT** to match the policy if the recipient belongs to the IT AD group for example.

Once the LDAP Server and all the required LDAP queries have been configured, the LDAP queries also need to be enabled in the listener, which is shown in the next slide.

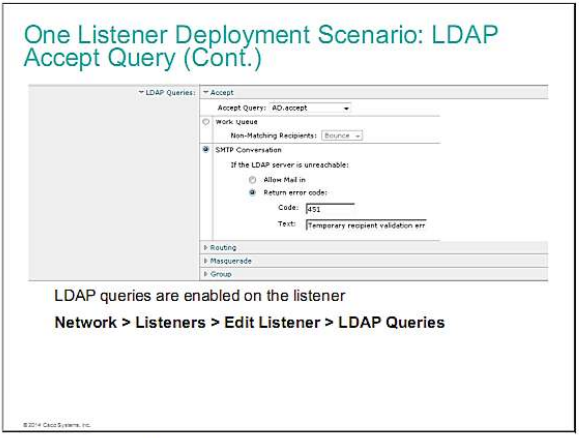

LDAP queries are enabled on the listener.

From the **Network > Listeners** page, select the **listener** to enable the LDAP query on.

Under the **Listener Settings**, LDAP Queries section, enable the different types of LDAP query required per your requirements.

For the **LDAP Accept Query**, using the **Accept Query** drop down box under the **LDAP Queries**: Accept section, select the configured **LDAP Accept Query** to use. In this example, the **AD.accept** LDAP Accept Query configured in the last slide is selected. **AD** is name of the LDAP Server Profile and **.accept** refers to the LDAP Accept Query configured within the LDAP Server Profile.

For the LDAP Accept Query, you can specify whether the LDAP Accept occurs during the work queue processing or during the SMTP conversation.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Using LDAP Accept during the work queue processing means the message must be accepted by the Cisco Email Security Appliance, whether the recipient is valid or not. At this point, if the recipient is invalid, you can specify the behavior for non-matching recipients to bounce or drop. Bounce means composing a message back to the sender indicating the message is undeliverable. Drop means silently throws the message away.

Using the LDAP Accept during the SMTP conversation checks if the recipient is valid or not immediately after the RAT check.

In case the LDAP server is not available, using LDAP Accept occurs during the work queue allows the Cisco ESA to queue the messages while it waits for the LDAP server to become available. Using LDAP Accept during during the SMTP conversation is time sensitive and the Cisco ESA cannot wait for the LDAP server to become available but you can set the Cisco ESA to **Allow Mail In** if the LDAP server is not available under the SMTP conversation setting.

It is recommended to use LDAP Accept during the SMTP conversation in most cases.

In this example, the **Accept Query** during the work queue processing is used to drop emails to a nonexisting users. For example, if email message will be sent to nonexisting email address **xyz@secure-x.public**, such e-mail will be dropped. On the other hand, e-mail message sent to **it1@secure-x.public** will be accepted as this is user's **it1** e-mail address stored in Active Directory Server.

Perform similar step to enable the other types of LDAP queries such as LDAP group query in the listener (not shown in the figure).

Note You may wish to bypass LDAP acceptance queries for special recipients, such as administrators. You can configure this setting from the RAT under **Mail Policies > Recipient Access Table (RAT) > Edit Recipient Access Table** and check the **Bypass LDAP Accept Queries for this Recipient** checkbox (not shown in the figure).

One Listener Deployment Scenario: HAT

This topic describes the Cisco Email Security Appliance HAT.

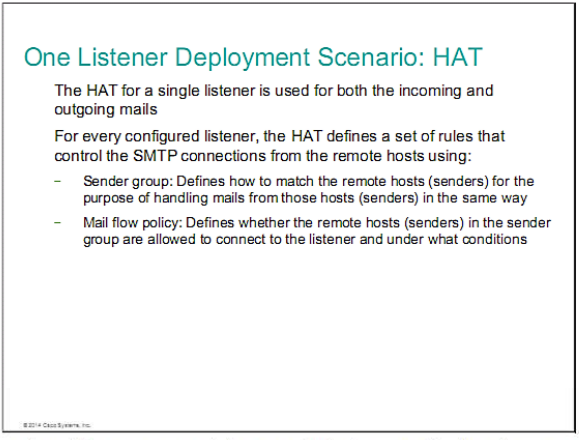

For every configured listener, you must define a set of rules that control the incoming connections from the remote hosts (senders). For example, you can define remote hosts (senders) and whether or not they can connect to the listener. The Cisco ESA allows you to define which hosts (senders) are allowed to connect to the listener using the HAT. The HAT maintains a set of rules that control incoming connections from remote hosts (senders) for a listener. Every configured listener has its own HAT. You configure HATs for both public and private listeners.

The HAT supports the ability to group the remote host definitions into **sender groups**. Each sender group will reference a rule called **mail flow policy**.

In HAT, you apply **mail flow policy** to the **sender group:**

The **sender group** is the way to group the remote hosts that is attempting to connect to the listener (by the sender's SenderBase reputation score, IP address or host name)

You define the way in which the remote hosts (senders) attempt to connect to a listener. The Cisco ESA classify the remote hosts (senders) into sender groups. A sender group is a list of remote hosts (senders) defined for the purpose of handling email from those remote hosts (senders) in the same way.

When a listener receives a SMTP connection, it compares the source IP address against the configured sender groups. It evaluates the sender groups in the order listed in the HAT. When it finds a match, it applies the configured **mail flow policy** to the SMTP connection.

The **mail flow policy** defines whether the remote host specified in the **sender group** can or cannot connect to the listener and under what conditions, like rate limiting the connections.

Each **sender group** in the HAT has a reference to a **mail flow policy**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

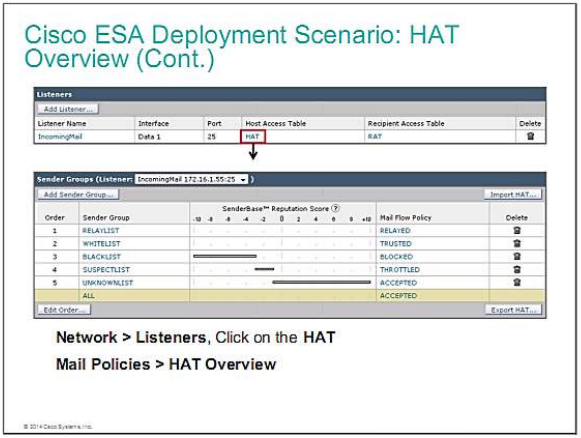

You can define whether the defined remote hosts (senders) in the sender group are allowed to connect to the listener and under what conditions. You define access rules using mail flow policies. For example, you can define that a particular sender group is allowed to connect to the listener, but only allow a maximum number of messages per connection.

The figure shows the HAT configurations of the "**IncomingMail**" listener based on our one listener scenario.

In the **Sender Groups (Listener:)** drop down box, select the listener which this HAT applies to.

When you create a listener, the Cisco ESA creates predefined sender groups and mail flow polices for the listener. You can edit the predefined sender groups and mail flow policies, and also create new sender groups and mail flow policies.

The order that rules appear in the HAT is important. The HAT is read from top to bottom for each host that attempts to connect to the listener. If a rule matches a connecting host, the action is taken for that connection immediately.

Note For all public listeners you create, by default, the HAT is set to *accept* email from *all* hosts. For all private listeners you create, by default, the HAT is set up to *relay* email from the host(s) you specify, and reject *all* other hosts.

You can also export all information stored in a HAT to a file, and you can import the HAT information stored in a file into the appliance for a listener.

There are two ways how to edit HAT settings in Cisco ESA GUI:

Network > Listeners, Click on the **HAT**

Mail Policies > HAT Overview

One Listener Deployment Scenario: HAT > Sender Group

This topic describes the Cisco Email Security Appliance Sender Group HAT component.

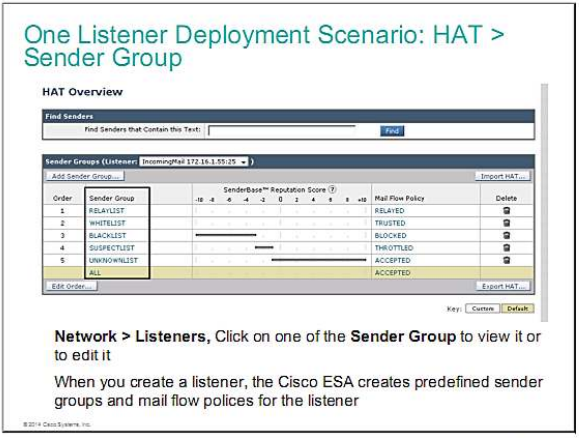

A sender group is simply a list of senders gathered together for the purposes of handling email from those senders in the same way (that is, applying a mail flow policy to a group of senders).

From the **Mail Policies > HAT Overview** page, you can configure how the listener handles messages from senders. Do this by creating, editing, and deleting the sender groups and mail flow policies in the HAT.

When an SMTP server attempts an SMTP connection with the Cisco ESA, the listener evaluates the sender groups in order and assigns the connection to a sender group when it matches any criterion in the sender group, such as the sender's SenderBase reputation score, IP address or host name.

Cisco recommends maintaining the default order of: **RELAYLIST** followed by **WHITELIST**, **BLACKLIST**, **SUSPECTLIST,** then the **UNKNOWNLIST** sender group. The **ALL** sender group at the end of the list is the default sender group.

You can define multiple remote hosts in a sender group by IP address or hostname such as 10.1.1.0/24 or 10.2.3.4 or example.com or .example.com and etc.

You can also define remote hosts by their SBRS.

You define whether the defined remote hosts in the sender group are allowed to connect to the listener and under what conditions using the mail flow policies. For example, you can define that a particular sender group is allowed to connect to the listener, but only allow a maximum number of messages per connection.

When you create a listener, the Cisco ESA creates predefined sender groups and mail flow polices for the listener.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Single listener example as shown in the figure - Single HAT with the predefined sender groups with their associated default mail flow policies:

RELAYLIST: Outgoing mail will be relayed if the mail server IP address is specified on the **RELAYLIST**. It will not be checked against the RAT. The **RELAYED** mail flow policy used by this sender group is typically configured so that the content from those senders is not scanned by the anti-spam scanning engine but is checked by anti-virus and only the internal mail server(s) are listed in the HAT as a relay. The **RELAYED** mail flow policy rate limits for the senders in this group are typically set to use the default which has a high limit.

WHITELIST: Add senders you trust to the **WHITELIST** sender group. The **TRUSTED** mail flow policy used by this sender group is typically configured so that the content from those senders is not scanned by the anti-spam scanning engine but is checked by anti-virus. In our scenario, we don't have any senders defined in the **WHITELIST** sender group. The **TRUSTED** mail flow policy rate limits for the senders in this group are typically set to use the default which has a high limit.

BLACKLIST: Senders in the **BLACKLIST** sender group are rejected (by the default parameters set in the **BLOCKED** mail flow policy). Adding senders to this group rejects connections from those hosts by returning a 5XX SMTP response in the **SMTP HELO** command.

SUSPECTLIST:This sender group uses the **THROTTLED** mail flow policy that throttles, or slows, the rate of incoming mail. For example, it may use a max messages per connection value of 1. If the senders are suspicious, you can add them to this sender group. The **THROTTLED** mail flow policy typically also has both anti-spam and anti-virus turned on.

UNKNOWLIST: This sender group are useful if you are undecided about the mail flow policy you should use for a given sender. The **ACCEPTED** mail flow policy is used for this group and it typically dictates that mail is accepted for the senders in this group, but anti-spam and anti-virus should be used. The **ACCEPTED** mail flow policy rate limits for the senders in this group are typically set to use the default which has a high limit.

ALL: Default sender group the that applies to all other senders. The default sender group also uses the **ACCEPTED** mail flow policy like the **UNKNOWLIST** sender group.

If the Cisco ESA is setup to use two listeners, one public listener and one private listener (not shown in the figure), then:

The public listener HAT should contain the **WHITELIST**, **BLACKLIST**, **SUSPECTLIST**, **UNKNOWNLIST** and the **ALL** sender groups.

Each public listener has its own HAT and RAT

The private listener HAT should contain the **RELAYLIST** and the **ALL** sender groups. For the **ALL** sender group in the private listener, the mail flow policy will be **BLOCKED** to deny any other mail servers from relaying mails to the Cisco ESA.

Each private listener has its own HAT but no RAT.

Two listeners example - HAT for Incoming Mail (Public Listener):

Sender Group Name	Mail Flow Policy
WHITELIST	TRUSTED
BLACKLIST	BLOCKED
SUSPECTLIST	THROTTLED
UNKNOWNLIST	ACCEPTED
ALL	ACCEPTED

https://t.me/learningnets

Two listeners example - HAT for Outgoing Mail (Private Listener):

Sender Group Name	Mail Flow Policy
RELAYLIST	RELAYED
ALL	BLOCKED

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: HAT > Sender Group SBRS

This topic describes the Cisco Email Security Appliance Sender Group SBRS HAT component

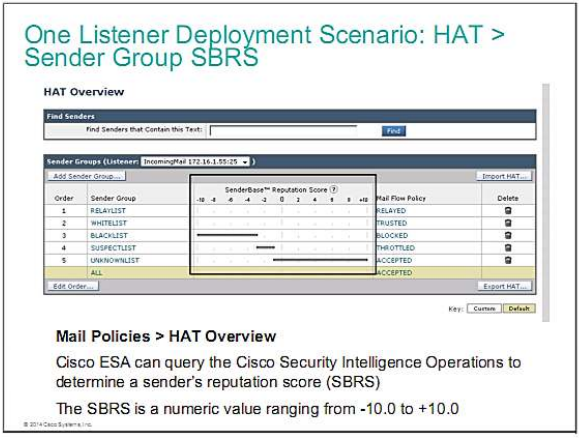

The Cisco ESA can query the Cisco SIO to determine a sender's reputation score (SBRS). The SBRS is a numeric value assigned to an IP address, domain, or organization based on information from the Cisco SIO. The scale of the SBRS ranges from -10.0 to +10.0.

A SBRS of -10.0 represents the source is most likely to be a source of spam. A SBRS of +10.0 represents the source to be a trustworthy sender.

Using the SBRS, you configure the Cisco ESA to apply mail flow policies to the senders based on their trustworthiness.

The figure and the below list shows using a moderate approach which is the Cisco ESA default settings in defining the sender groups and mail flow polices. In this case, the known good senders (**WHITELIST**) is not using the SBRS to match the senders, the bad senders (**BLACKLIST**) are matched if their SRBS score is between -3 to -10, the suspected senders (**SUSPECTLIST**) are matched if their SRBS score is between -1 to -3 and the undecided (**UNKNOWNLIST**) senders are matched if their SRBS score is between -1 to +10 or None.

When defining the SBRS range, you can enable **Include A SBRS Scores of "None"**. For example, the very new domains may not have SenderBase Reputation Scores yet.

The **RELAYLIST** sender group also does not use the SenderBase Reputation Score to match the sender by default.

https://t.me/learningnets

One Listener Deployment Scenario: HAT > BLACKLIST Sender Group

This topic describes the Cisco Email Security Appliance BLACKLIST Sender Group.

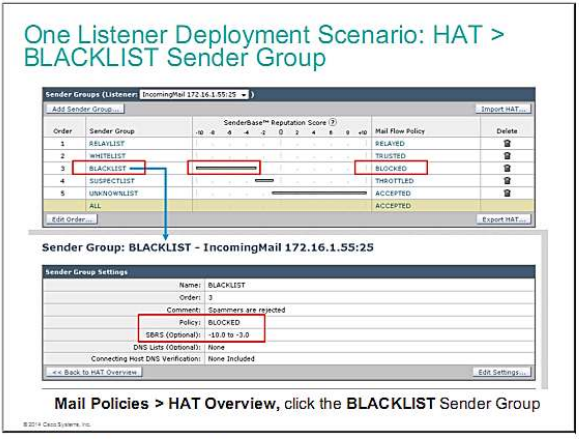

The figure shows the **BLACKLIST** predefined sender group.

Within the sender group, define the name of the sender group, the order in which the sender group will be matched in the HAT, optional comments about the sender group, the mail flow policy, the optional SBRS, the optional DNS lists and the connecting host DNS verification.

In the example, the **BLACKLIST** predefined sender group is third in the order within the HAT and is using the **BLOCKED** mail flow policy.

Any sender with a SBRS ranging from -10 to -3 will be matched by this sender group.

The optional **DNS Lists** mechanism allows you to configure a sender within a group that queries a DNS List (a blacklist) using a third party service so that you can adjust your mail flow policies accordingly. For example, you could reject connections or limit the behavior of the connecting coming from any blacklisted domain.

Senders can be unverified for different reasons. For example, the DNS server could be "down" or not responding, or the domain may not exist. Host DNS verification settings for sender groups allow you to classify unverified senders prior to the SMTP conversation and include different types of unverified senders in your various sender groups. The Cisco appliance attempts to verify the sending domain of the connecting host via DNS for incoming mail. This verification is performed prior to the SMTP conversation. The system acquires and verifies the validity of the remote host's IP address (that is, the domain) by performing a double DNS lookup. A double DNS lookup is defined as a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The appliance then checks that the results of the A lookup match the results of the PTR lookup. If the PTR or A lookups fail, or the results do not match, the system uses only the IP address to match entries in the HAT and the sender is considered as not verified.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Unverified senders are classified into the following categories:

Connecting host PTR record does not exist in the DNS.

Connecting host PTR record lookup fails due to temporary DNS failure.

Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A).

Using the sender group "**Connecting Host DNS Verification**" setting, you can specify a behavior for the unverified senders. For example, senders for which reverse DNS lookups fail can be configured to match the **SUSPECTLIST** sender group and will receive the default action from the **THROTTLED** mail flow policy (not shown in the figure).

One Listener Deployment Scenario: HAT > RELAYLIST Sender Group

This topic describes the Cisco Email Security Appliance RELAYLIST Sender Group.

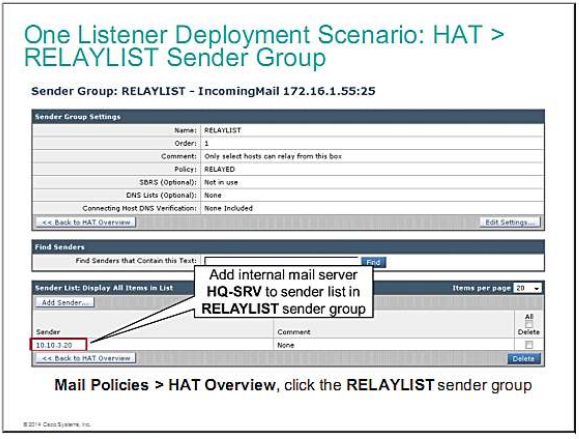

The figure shows the **RELAYLIST** predefined sender group:

- The **RELAYLIST** Sender Group defines the internal mail servers allowed to relay the outgoing emails through the Cisco ESA
- The **RELAYLIST** predefined sender group is first in the order within the HAT and is using the **RELAYED** mail flow policy.
- The **RELAYLIST** sender group is not using the SBRS to match the senders.

Add senders you know that should be allowed to relay the outgoing emails through the Cisco ESA to the **RELAYLIST** sender group. You typically would add the internal mail servers to the **RELAYLIST** sender group.

In this deployment scenario, the **RELAYLIST** sender group will be matched by the IP address **10.10.3.20,** which is the internal email server (**HQ-SRV**) IP address. This way the listener will relay the outgoing mails only from the 10.10.3.20 internal mail server to the Internet.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: HAT > Add Sender Group

This topic describes how to add a Sender Group.

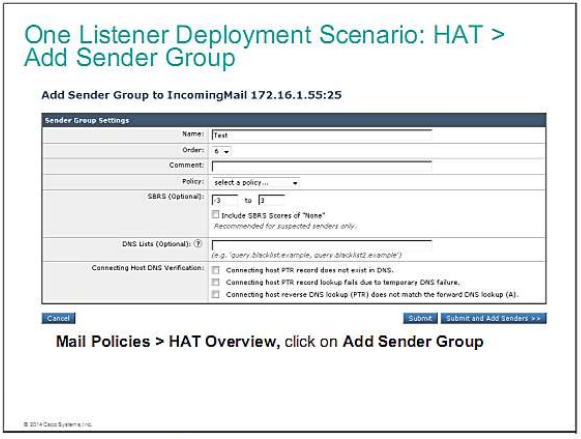

To add a new sender group, follow these steps:

1. From the **Mail Policies > HAT Overview** page, click **Add Sender Group**.
2. Type the name of the sender group, select the order in which it will be placed in the list of sender groups, and a comment (optional) in the fields provided.
3. Choose a mail flow policy from the drop-down list.
4. Select a SBRS range and DNS list (optional). You can also mark the check box to include senders with a SBRS of "None".
5. Configure any Host DNS Verification settings.
6. Click **Submit and Add Senders** to create the sender group and to add senders to it using IP addresses or host names.

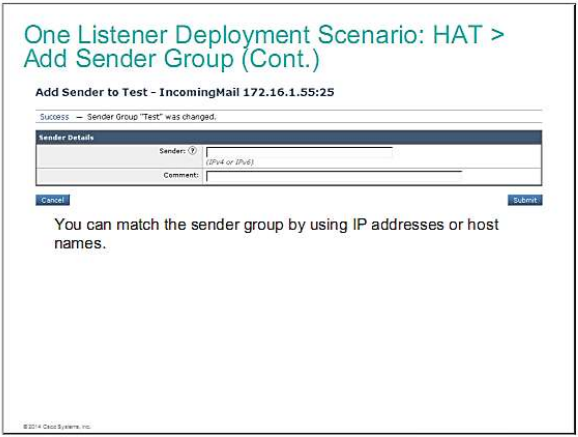

This is the page where you can match the sender group by using IP addresses or host names.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: HAT > Mail Flow Policy

This topic describes the Cisco Email Security Appliance Mail Flow Policy

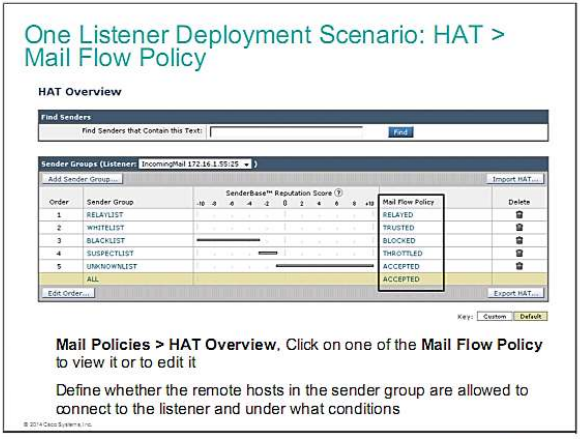

A mail flow policy is referenced in each sender group to define whether the remote hosts are allowed to connect to the listener and under what conditions. It is used to control or rate limit the flow of email messages from a sender to the listener.

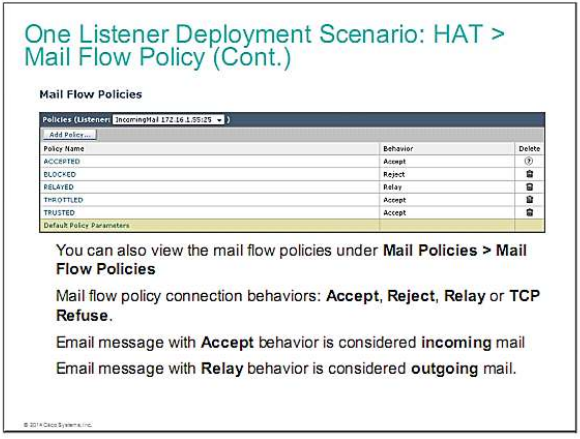

Every message received by the Cisco ESA is classified as incoming or outgoing mail. The determination is made using the connection behavior configuration in the mail flow policy that was applied in the HAT. Every message that is **accepted** is considered incoming mail. Every message that is **relayed** is considered outgoing mail.

Each mail flow policy can have one of the following connection behavior (action):

Accept: The connection will be accepted and the SMTP conversation will start. The sender is limited to the recipients in the domains specified in the RAT and maybe subject to further checks on the recipient addresses.

Reject: The TCP connection is accepted, but the Cisco ESA sends the sender an SMTP 554 banner to indicate that they are not welcome. That is then immediately followed by a TCP FIN to close the connection. The banner code and text is customizable to provide the sender some information about why the connection has been rejected.

TCP Refuse: The TCP connection is closed. The Cisco ESA issues a FIN to the sender to indicate the connection is over before it even starts. In this case, most senders will treat the Cisco ESA as down and may reattempt he connection for hours or days. It is not generally recommended to use this connection behavior for a typical mail flow policy.

Relay: The connection will be accepted and the SMTP conversation will start. The sender is **not** limited to the recipients in the domains specified in the RAT. The RAT is not checked at all. This ability to send to any email address is a high level of privilege to grant any sender, so use it carefully by limiting it to very specific sending hosts which are typically all the internal email servers.

The Cisco ESA uses two different sets of mail policies for message content security:

Incoming mail policies for connections that match a HAT policy with the **Accept** connection behavior in any listener.

Outgoing mail policies for connections that match a HAT policy with the **Relay** connection behavior in any listener.

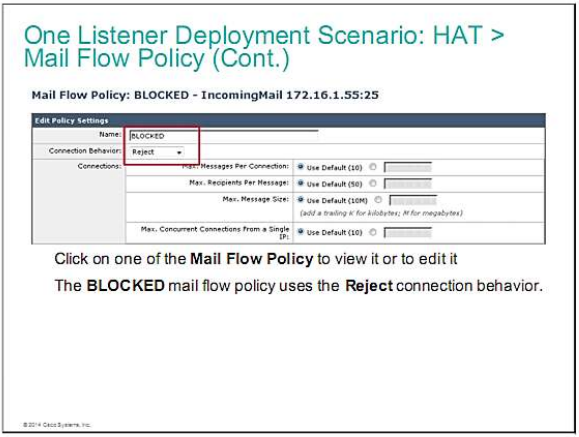

The figure shows mail flow policy example, the **BLOCKED** mail flow policy. To edit mail flow policies, click on the policy under **Mail Policies > Mail Flow Policies** or **Mail Policies > HAT Overview**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The **BLOCKED** mail flow policy uses the **Reject** connection behavior.

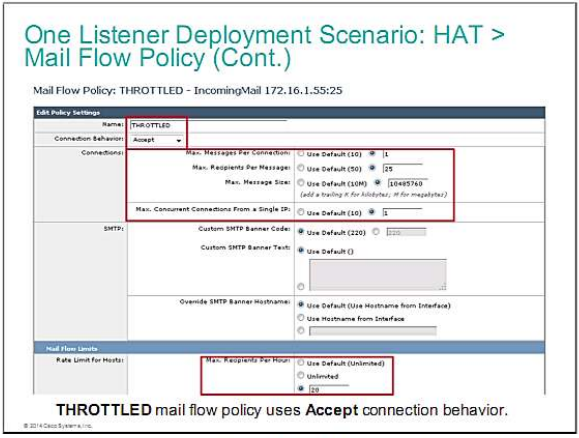

THROTTLED mail flow policy uses **Accept** connection behavior.

The figure shows mail flow policy example, the **THROTTLED** mail flow policy. To edit mail flow policies, click on the policy under **Mail Policies > Mail Flow Policies** or **Mail Policies > HAT Overview**.

The **THROTTLED** mail flow policy uses the **Accept** connection behavior.

If a mail flow policy has a connection behavior of **Accept**, you can further specify exactly how much mail is accepted. This rate limiting capability is important for managing the email injection rates from the senders so that no sender(s) can overwhelm the Cisco ESA. The typical **THROTTLED** mail flow policy make use of these limits to control the possible spam senders.

The figure shows some of the typical rate limiting settings in the **THROTTLED** mail flow policy including the maximum number of messages per connection, the maximum number of recipients per message, the maximum message size, the maximum number of concurrent connections from a single IP, and other mail flow limits.

One Listener Deployment Scenario: HAT > Mail Flow Policy > Anti-Spam and Anti-Virus

This topic describes Cisco Email Security Appliance Mail Flow Policy Anti-Spam and Anti-Virus settings.

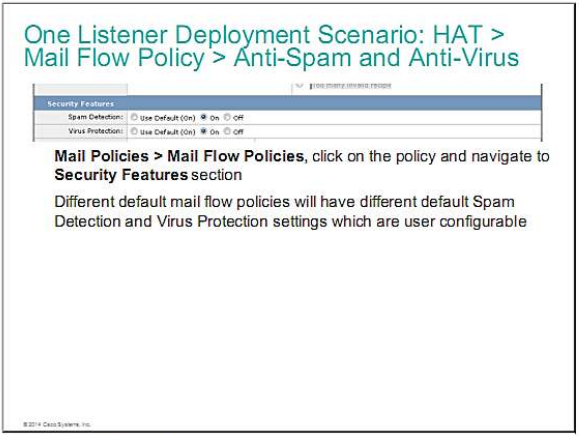

The figure shows part of the mail flow policy security features settings, where the Spam Detection and Virus Protection settings can be enabled or disabled. The different default mail flow policies will have different default Spam Detection and Virus Protection settings which are user configurable.

Note In the Mail Flow Policy, Anti-Spam and Anti-Virus can be turned enabled or disabled in general for every sender that belongs to that sender group. In the Incoming and Outgoing Mail Policies, Anti-Spam and Anti-Virus can also be enabled or disabled for the particular users using user based policies.

Note Refer to the latest Cisco ESA user guide for more information on all the other mail flow policy parameters not discussed here.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

One Listener Deployment Scenario: HAT > Mail Flow Policies Summary

This topic describes the Cisco Email Security Appliance Mail Flow Policies connection behavior, throttling and anti-spam and anti-vrius typical settings.

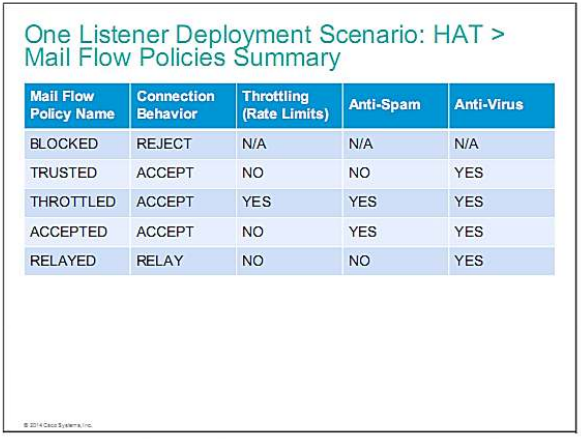

The table shows the system defined mail flow policies that are used in the HAT for a single listener.

The table summarize the connection behavior for each mail flow policy and the typical throttling (rate limits), anti-spam and anti-vrius settings.

One Listener Deployment Scenario: RAT

This topic describes the Cisco Email Security Appliance RAT

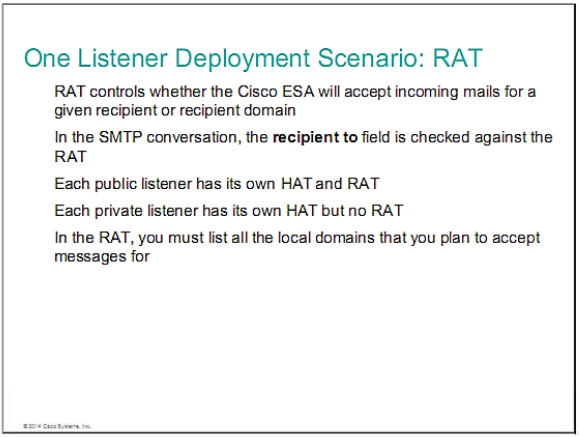

The RAT controls whether the Cisco ESA will accept or reject the incoming mails for a given recipient or recipient domain.

In the SMTP conversation, the recipient to, or rcpt to, field is checked against the RAT. Messages to recipients matched by the RAT will be accepted, and all others will be dropped.

Each public listener has its own HAT and RAT. Each private listener has its own HAT but no RAT.

In the RAT, you must list all the local domains that you plan to accept messages for. The entries in the RAT can be in different forms, and can include usernames or just the host or domain portion after the @ symbol. Each RAT entry has an action of **Accept** or **Reject**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

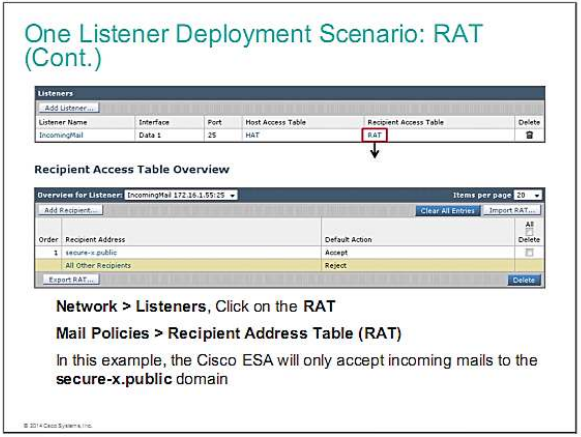

When you create a public listener, you define all local domains that the appliance will accept messages for using the RAT. Many enterprise gateways are configured to receive messages for several local domains. The Recipient Access Table defines which recipients will be accepted by the public listener. The table specifies the address (which may be a partial address, username, domain, or hostname) and whether to accept or reject it.

The RAT allows you to define a recipient or group of recipients. Recipients can be defined by full email address, domain, partial domain, or username:

division.example.com	Fully-qualified domain name.
.partialhost	Everything within the "partialhost" domain.
user@domain	Complete email address.
user@	Anything with the given username.
user@[IP_address]	Username at a specific IP address. Note that the IP address must be between the "[]" characters.

In the figure, the **secure-x.public** recipient address in the RAT has the **Accept** action, therefore, the Cisco ESA will accept incoming mails to the **secure-x.public** domain.

The default entry, **All Other Recipients** has the **Reject** action, therefore, the Cisco ESA will reject incoming mails to all other domains except the **secure-x.public** domain.

You can also export all information stored in a RAT to a file, and you can import the RAT information stored in a file into the appliance for a listener.

One Listener Deployment Scenario: SMTP Routes

This topic describes the Cisco Email Security Appliance SMTP Routes.

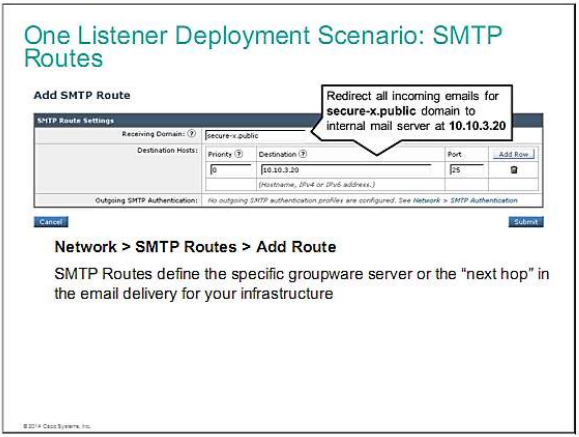

The SMTP Routes table allows you to redirect all email for each domain you enter to a specific delivery host. In typical installations, the SMTP Routes table defines the specific groupware server or the next hop in the email delivery for your infrastructure.

Use the **Network > SMTP Routes > Add Route** page to redirect all email for a particular domain to a specific MX host, without rewriting the envelope recipient. When creating a SMTP route, the receiving domain can be specified using the hostname, domain, or IP address and the destination for each receiving domain can be specified using the hostname or IP address.

A SMTP route is configured on the Cisco ESA for routing the mails to the internal mail server. In the figure, the SMTP Routes table contains one entry. This entry specify the groupware server for the **secure-x.public** domain is **10.10.3.20**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

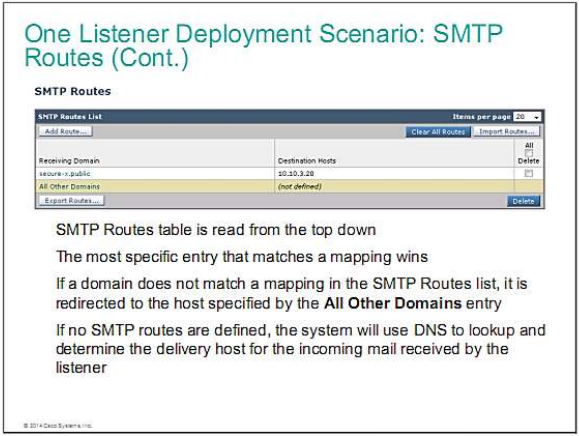

The SMTP Routes table is read from the top down for every email delivery. The most specific entry that matches a mapping wins. The Cisco ESA allows up to 40,000 SMTP Route mappings to be configured.

If a domain does not match a mapping in the SMTP Routes list, it default to being redirected to the host specified by the **All Other Domains** entry.

When creating a SMTP route, the receiving domain can be specified using the hostname, domain, or IP address and the destination for each receiving domain can be specified using the hostname or IP address.

Note	If no SMTP routes are defined, the system will use DNS to lookup and determine the delivery host for the incoming mail received by the listener.

https://t.me/learningnets

One Listener Deployment Scenario: Email Relaying on Internal Mail Server

This topic describes how to configure Email relaying on internal mail server and hence redirect emails from internal server to Cisco Email Security Appliance.

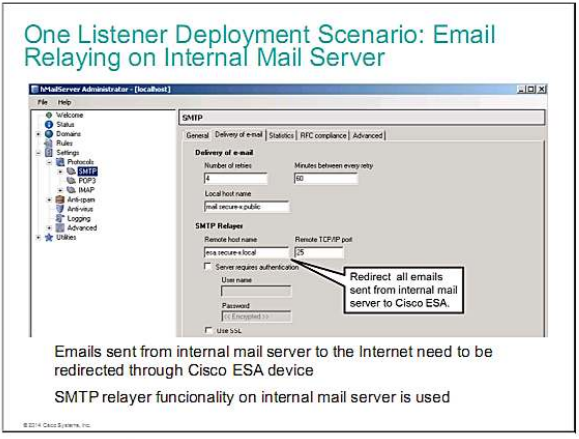

When emails are being sent from internal mail server to the Internet, these emails need to be redirected through Cisco ESA device. The internal mail server on the HQ server (10.10.3.20) in our deployment scenario needs to be configured with a mail relay which points to the Cisco ESA. The internal mail server uses the mail relay (Cisco ESA) to send the outgoing mails from the **secure-x** domain.

In our deployment scenario internal mail server used is the **hMailServer**. When using the hMailServer SMTP server, navigate to **Settings > Protocols > SMTP > Delivery of e-mail** screen and type in the IP address or domain name and the TCP port of the Cisco ESA into **SMTP Relayer** field.

Note Refer to the specific email server configuration guide on how to configure SMTP relay when using any other email server software.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

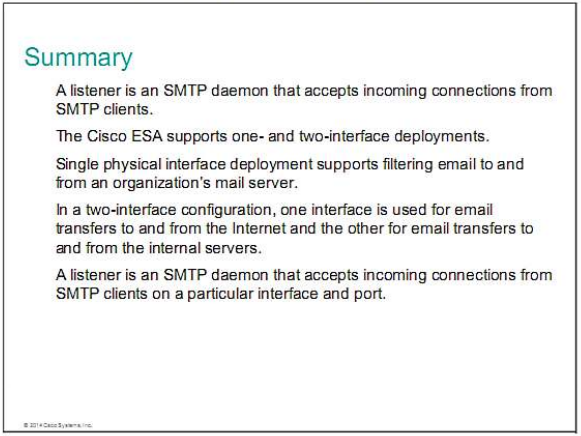

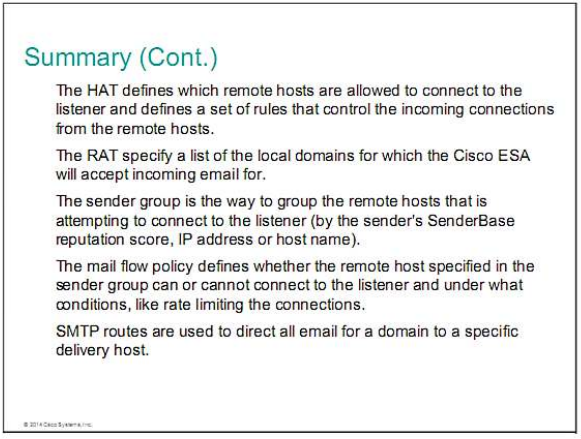

References

For additional information, refer to these references:

Email Security Using Cisco ESA Deployment Guide

Cisco IronPort AsyncOS for Email Configuration Guide

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 3

Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies

Overview

This lesson explain how to configure the different features within the incoming and outgoing mail policies (anti-spam, anti-virus, content filters, outbreak filters, data loss prevention), describe the reporting functions and the tools used for troubleshooting the mail processing flows.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the Cisco ESA Incoming and Outgoing Mail Policies Configurations Overview
- Describe how the Cisco Email Security Appliance matches mail policies.
- Provide an overview of the Anti-Spam filtering
- Describe the Cisco ESA Incoming and Outgoing Mail Policy Anti-Spam Configuration
- Describe the Cisco ESA Spam Quarantine.
- Describe the Cisco ESA Policy, Virus, Outbreak Quarantines Configuration
- Provide an overview of the Anti-Virus filtering
- Describe the Cisco ESA Incoming and Outgoing Mail Policy Anti-Virus Configuration
- Provide an overview of the content filters
- Describe the Cisco ESA Incoming and Outgoing Mail Policy Content Filters Configuration
- Provides an overview of the outbreak filters

https://t.me/learningnets

Describe the Cisco ESA Outbreak Filters Configuration

Provide an overview of the DLP outgoing mail policy feature

Describe how to enable the DLP feature

Provide an overview of the Cisco ESA reporting capabilities

Describe message tracking tool.

Describe the trace tool

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ESA Incoming and Outgoing Mail Policies Overview

This topic describes the Cisco Email Security Appliance Incoming and Outgoing Mail Policies Configurations Overview

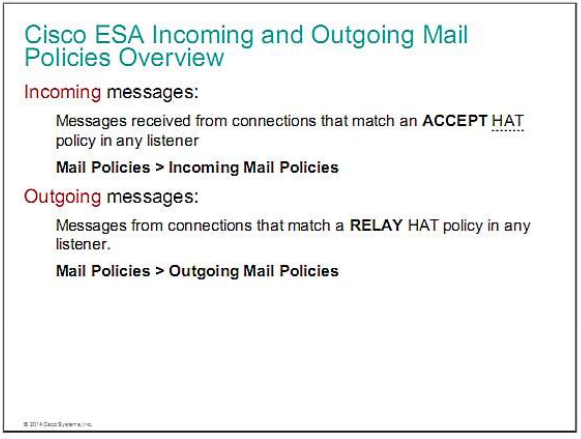

The Cisco ESA uses two different sets of mail policies for message content security:

- Incoming mail policies for messages are messages received from connections that match an ACCEPT HAT policy in any listener.
- Outgoing mail policies for messages are messages from connections that match a RELAY HAT policy in any listener.

Having separate sets of policies allows you to define different security rules for messages sent to your users and messages sent from your users. You manage these policies using the **Mail Policies > Incoming Mail Policies** or **Outgoing Mail Policies** pages in the GUI, or the **policyconfig** command in the CLI. This course will not cover any CLI configurations.

This topic describes the Cisco ESA Incoming and Outgoing Mail Policies Configurations Overview

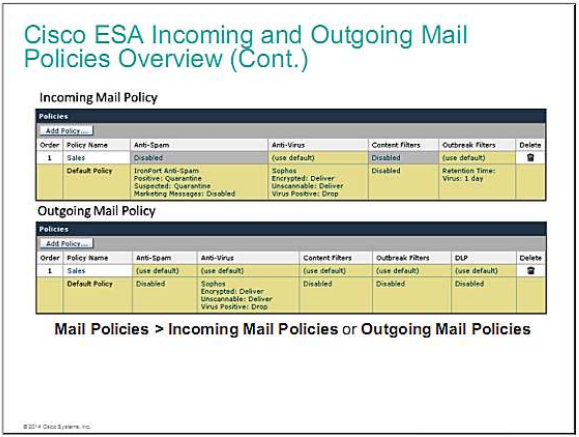

As messages are received by the Cisco ESA, the Cisco ESA will attempt to match each message recipient and sender to a mail policy in the Incoming or Outgoing Mail Policies table, depending on whether it is an incoming or outgoing message. Matches are based either on the recipient's address or the sender's address. Each user (sender or recipient) is evaluated for each mail policy defined the appropriate mail policy table in a top-down fashion. For each user, the first matching policy wins. If a user does not match any specific policy, user will automatically match the Default Policy located at the bottom of the table. The Default Policy are applied to messages that are not covered by any other mail policy. If no other policies are configured, the Default Policy applies to all messages.

Through the **Mail Policies** menu in the GUI, you create and manage incoming or outgoing mail policies. Mail policies are defined as a specific set of users, that map to specific settings for the following features:

- Anti-Spam Scanning
- Anti-Virus Scanning
- Content Filters
- Virus Outbreak Filters
- Data Loss Prevention

Note DLP scanning can only be performed on outgoing messages.

Users can be defined by email address, email domains, or LDAP group queries.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

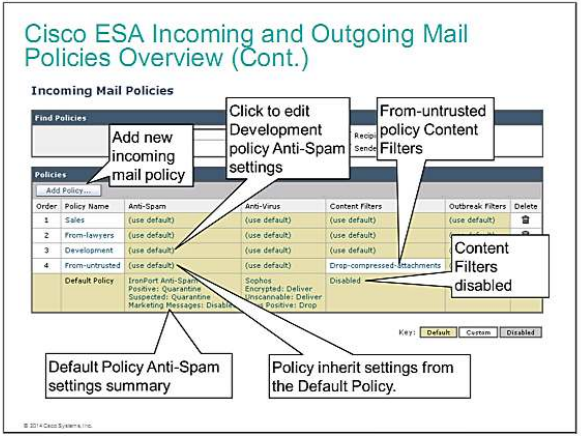

Newly created policies will inherit the Anti-Spam, Anti-Virus, and other settings of the Default Policy as indicated by **(use default)**. As in the HAT operation, starting at the top of the policy list, the first match is applied. The Default Policy is always the last policy in the policy list. No policies, except the matching policy will affect the message processing. As opposed to letting all mail flow through using the Default Policy, users based mail policies can be implemented to allow for finer degrees of control. To create a new incoming mail policy, click on **Add Policy**.

To control the Antispam, Antivirus, and other settings of each mail policy, click on **(use default)** to create individual policy based Anti-Spam, Anti-Virus, and other settings.

For example, as shown in the figure, the **From-untrusted** policy is using the content filters named **Drop-compressed-attachments** and the content filters are disabled for the Default Policy.

The **Sales**, **From-lawyers** and **Development** policies all use the Default Policy Anti-Spam, Anti-Virus, Content Filters and Outbreak Filters settings.

Cisco ESA Mail Policies Matching

This topic describes how the Cisco Email Security Appliance matches mail policies.

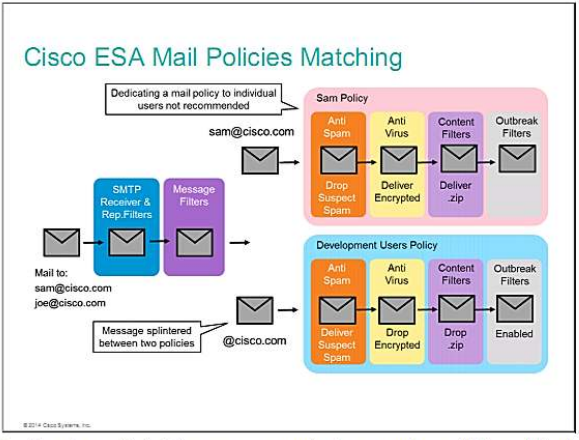

Mail policies allow the email administrator to create rules for managing email flow within the organization. The policies are user-based and are specified to match Cisco filtering tools such as Antispam, Antivirus, Outbreak Filters, and Content Filters.

The mail policies consist of incoming and outgoing policies. Different policies are required depending on whether it is incoming or outgoing mail. The incoming policy table is used for incoming messages from sending hosts that are defined in the **"HAT policy Accept behavior."** The outgoing policy table is used for outgoing (relayed) messages from sending hosts that are defined in the HAT policy **Relay** behavior.

As messages are received, the Cisco ESA attempts to match each message recipient and sender to a mail policy in the Incoming or Outgoing Mail Policies table, depending on whether it is an incoming or outgoing message. Matches are based on either the recipient's address or the sender's address. Each message is evaluated for each policy in the appropriate table (incoming or outgoing) in a top-down fashion. For each message, the first matching policy wins. If a message does not match any specific policy, it will match the default policy of the table.

This figure shows two different policies (based upon Rcpt To) for incoming mails. The top mail policy (Sam Policy) is dedicated to a specific user, while the bottom one (Development Users Policy) is for a user group. An email addressed to two users **sam@cisco.com** and **joe@cisco.com** will be splintered. The email to **sam@cisco.com** will be processed by the top policy (Sam Policy), the email to **joe@cisco.com** will be processed by the bottom policy (Development Users Policy). The mail policies are processed sequentially until a match is found. This match process is performed for each email splinter separately.

Dedicating policies to individual users is not recommended but is shown here to illustrate the message splintering, in which a single mail is processed by two policies. As a general rule, you should group as many users as possible into one policy to avoid splintering and avoid troubleshooting messages that mistakenly dropped into a different policy.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

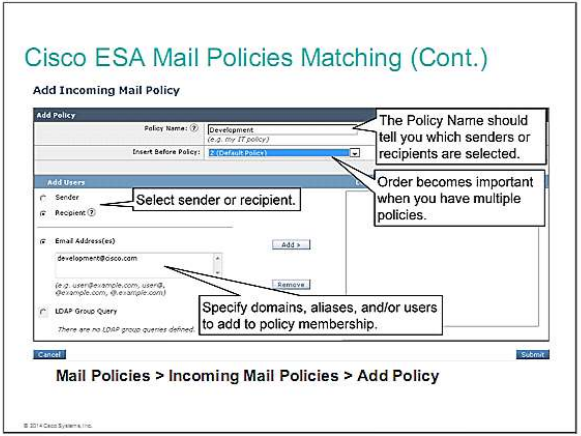

To add a new policy, navigate to **Mail Policies > Incoming Mail Policies > Add Policy**:

Define a unique name for the policy and adjust the order of the policy (if necessary)

Define users for the policy. Users can be sender or recipient and are defined by email address, email domain, or LDAP group by using LDAP group queries

This example illustrates how to configure a mail policy called "Development." This policy will match on incoming mail whose recipient address is equal to the **development@cisco.com** alias.

It is recommended that the policy name have a descriptive phrase with no spaces. The sequence number determines the order. The order is important when there are two or more policies, including the default policy.

Cisco ESA Mail Policies Matching (Cont.)

Recipient address matches the **Envelope Recipient** address

Sender address matches the: **Envelope Sender** (RFC821 MAIL FROM address), Address found in the RFC822 **From: header,** Address found in the RFC822 **Reply-To: header**

- Sender-based policy example: Incoming mails from the partners are exempted from Anti-Spam scan

Sender/Recipient syntax	Meaning
@cisco.com	All users within domain name
@ccnp-security.cisco.com	All users within the subdomain
user1@cisco.com	Complete email address
LDAP Group: Sales	Address checked for LDAP Group membership
user1@	Any recipient with the given username
user1@[1.2.3.4]	Username at a domain literal address (square brackets required)

© 2014 Cisco Systems, Inc.

You can create multiple policies that satisfy the disparate security needs of the different user groups within your organization. The Cisco ESA Email uses the rules defined in these policies to scan each message and, if necessary, perform an action to protect your user.

Users for a given policy can be defined in the following ways:

Full email address: user@example.com

Partial email address: user@

All users in a domain: @example.com

All users in a partial domain: @.example.com

LDAP group using LDAP group query (example shown in the next slide)

https://t.me/learningnets

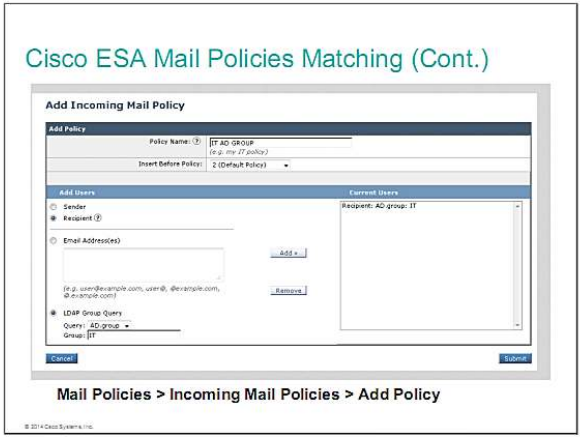

Group query allows you to query LDAP for group information so you can create group based incoming or outgoing mail policies.

In this example, an incoming mail policy named **IT AD GROUP** is matched using **Recipent: AD.group: IT** to match the policy if the recipient belongs to the IT AD group.

When creating the mail policy, under the LDAP Group Query section, in the Query drop down box, select the configured LDAP Server Profile LDAP Group Query. In this example, it is **AD.group** where **AD** is name of the LDAP Server Profile and **.group** refers to the LDAP Group Query configured within the **AD** LDAP Server Profile.

In the Group: field, enter the name of the LDAP group then click the **Add** box to add that to the Current Users. In this example, **IT** was entered then added resulting in the **Recipent:AD.group: IT** setting under the Current Users.

Anti-Spam Overview

This topic provides an overview of the Anti-Spam filtering.

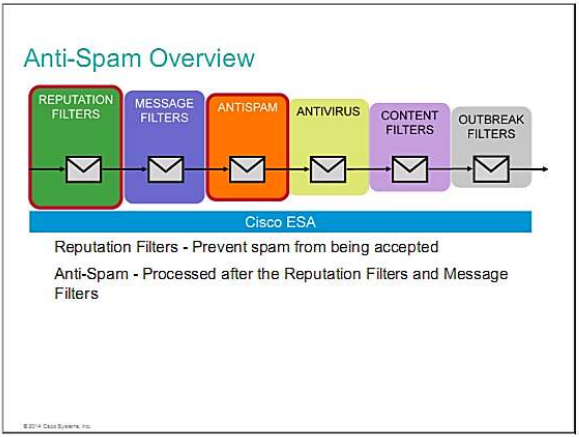

This figure shows the Cisco Email Security Appliance message filtering process flow. This figure lists the processes horizontally. To defend against spam, there are two separate processes: Reputation Filters, which appear at the beginning using the HAT and have very low overhead, and then Anti-Spam, which is further down the process. Anti-Spam has more overhead because it looks at message content and happens after the message has been accepted.

Reputation-based filtering relies on the likelihood that if a server is a known spam sender, it is more likely that email coming from that server is spam compared to a host that does not have a reputation for distributing spam. Reputation filters provide the first layer of defense by looking at the source IP address of the email server and comparing it to the reputation data downloaded from Cisco SenderBase. Cisco SenderBase is the world's largest repository for security data, including sources of spam, botnets, and other malicious hosts. When hosts on the Internet engage in malicious activity, SenderBase lowers the reputation score of that host.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anti-Spam Configuration

This topic describes the Cisco Email Security Appliance Incoming and Outgoing Mail Policy Anti-Spam Configuration

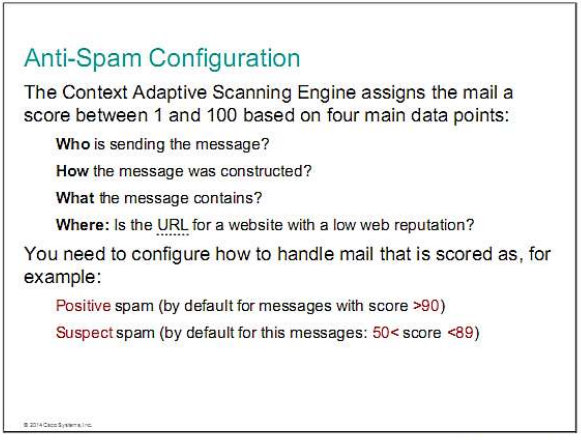

CASE combines content analysis, email reputation, and web reputation to deliver the broadest set of threat prevention factors. Every five minutes, the CASE downloads a new rule set. The engine is designed to take a close look at four main data points:

1. **Who** is sending the message (looks at the reputation score).
2. **How** the message was constructed. Spammers try to make messages look like they were sent via Outlook, for example, when really they may have used an unbranded software application.
3. **What** the message contains, for example, random dots inside an image, or the same color scheme as found in several spam traps.
4. **Where:** Is the URL for a website with a low web reputation? The SenderBase database has both web and email domain reputation scores.

All this information is taken together to assign the mail a score between 1 and 100. For example, suspected spam will be mail that receives a score between 50 and 89. Positive spam is email that has score above 90. These are the defaults. These spam mail is called "gray mail." Because the message was not blocked in the HAT, the Anti-Spam engine will assign the spam score. Use the mail policy Anti-Spam settings to configure how to handle mail that is scored as positive spam and suspect spam.

https://t.me/learningnets

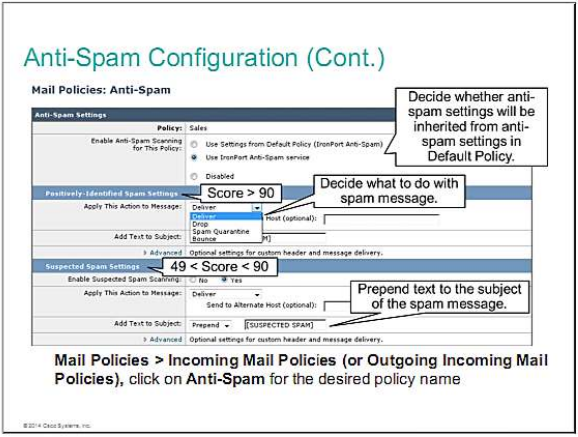

The Anti-Spam processes scan email for incoming and outgoing mail based on the mail policies that you configure. Anti-Spam is enabled by default in the default incoming mail policy and disabled by default in the default outgoing mail policy.

The Anti-Spam scanning engines assign a score to each message. The higher the score, the greater the likelihood that the message is spam. Based on the score, each message is categorized as one of the following:

Not spam

Unwanted marketing email from a legitimate source

Suspected spam

Positively-identified spam

Actions taken on messages positively identified as spam, suspected to be spam, or identified as unwanted marketing messages are not mutually exclusive; you can combine some or all of them differently within different incoming or outgoing policies for different processing needs for groups of users. For each mail policy, you specify the settings that determine which messages are considered spam and what action to take on those messages. For example, you may want to drop messages positively identified as spam, but quarantine suspected spam messages. You also specify which engine will scan messages that the policy applies to.

You can assign different users to different mail policies and define different scanning engines, spam-definition thresholds, and spam-handling actions for each policy.

To configure the Anti-Spam setting for an incoming or outgoing mail police, from the **Mail Policies > Incoming Mail Policies (or Outgoing Incoming Mail Policies)** page, click on **Anti-Spam** for the desired policy name.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

In the **Enable Anti-Spam Scanning for This Policy** section, select the Anti-Spam solution you want to use for the policy. The options you see depend on the Anti-Spam scanning solution(s) that you have enabled:

If you select **Use Settings from Default Policy**, all other options on the page are disabled and anti-spam settings defined in **Default Policy** will get applied, which will quarantine both positive spam and suspect spam, with the marketing email scanning disabled.

Select **Use IronPort Anti-Spam Service** to change default anti-spam settings

Click **Disable** to disable the Anti-Spam service for this policy.

The configurable Anti-Spam settings for positively identified spam, suspected spam, and marketing messages includes:

Choose an action to take on positively identified spam, suspected spam, or unwanted marketing messages:

- **Deliver**: Deliver the message. The message delivery continues with the other checks such as Anti-Virus.
- **Drop**: Drop the message. The message delivery stop.
- **Spam Quarantine**: Store the messages so the recipient can examine them. The decision to quarantine messages or not is usually a business policy decision, not a technical decision. The primary reason for using quarantine is to protect against false positives (messages mistakenly marked as spam but are legitimate). Because the Anti-Spam engine could potentially drop a legitimate message, the choice is made to save the messages so the recipient can sort it out for themselves.
- **Bounce**: A bounce message is sent back to the sender that usually contains information about the original message and why the delivery failed.

Enable Suspected Spam Scanning.

Enable Marketing Email Scanning.

In the figure, the Anti-Spam setting for an **Incoming Mail Policy** named **Sales** is configured with option to **Use IronPort Anti-Spam service**. Both positive spam and suspected spam messages will be delivered with added text to the subject of the message.

Note When IronPort Anti-Spam is enabled during system setup, it is enabled for the default incoming mail policy with the default values for the global settings. If you have not enabled IronPort Anti-Spam in the system setup wizard, select **Security Services > IronPort Anti-Spam**, click **Enable**, scroll to the bottom of the license agreement page and click **Accept** to accept the agreement, click **Edit Global Settings**, select the check box for **Enable IronPort Anti-Spam Scanning** to enable the Anti-Spam feature globally (not shown in the figure).

The Cisco ESA offers the following two Anti-Spam solutions:

IronPort Anti-Spam Filtering

- IronPort Anti-Spam addresses a full range of known threats including spam, phishing and zombie attacks, as well as hard-to-detect low volume, short-lived email threats such as "419" scams. In addition, IronPort Anti-Spam identifies new and evolving blended threats such as spam attacks distributing malicious content through a download URL or an executable.

Cisco IronPort Intelligent Multi-Scan Filtering

- Cisco IronPort Intelligent Multi-Scan incorporates multiple anti-spam scanning engines, including Cisco IronPort Anti-Spam, to provide a multi-layer anti-spam solution.

- When processed by Cisco IronPort Intelligent Multi-Scan, a message is first scanned by third-party anti-spam engines. Cisco IronPort Intelligent Multi-Scan then passes the message and the verdicts of the third-party engines to Cisco IronPort Anti-Spam, which assumes responsibility for the final verdict. After Cisco IronPort Anti-Spam performs its scan, it returns a combined multi-scan score to AsyncOS. Combining the benefits of the third-party scanning engines and Cisco IronPort Anti-Spam results in more caught spam while maintaining Cisco IronPort Anti-Spam's low false positive rate.

You can license and enable both these solutions on the Cisco ESA, but you can only use one in a particular policy. You can specify a different Anti-Spam solution for different groups of users. In the example in the figure only IronPort Anti-Spam Filtering is licensed and enabled.

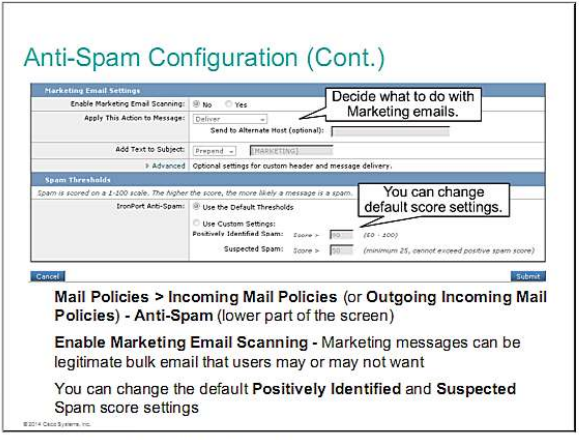

This figure shows the lower part of the Anti-Spam page.

Both Anti-Spam scanning engines can distinguish between spam and unwanted marketing messages from a legitimate source. Marketing messages are broadcast emails that are typically not received until a user signs up for them. Users will register themselves by answering the question, "Would you like to receive advertisements from us or our affiliate partners?" After answering "**Yes**," users receive email that they did not anticipate, which they interpret as spam. Even though marketing messages are not considered spam, your organization or end-users may not want to receive them. Like spam, you have the option to deliver, drop, quarantine, or bounce unwanted marketing messages. You also have the option to tag unwanted marketing messages by adding text to the message's subject to identify it as marketing. Marketing messages are legitimate bulk email that users do not want. These broadcast emails are typically not received until a user signs up for them.

What's the difference between spam and marketing? Marketing makes no attempt to disguise the origin of the communication. Marketing is often cross reported (some users report them as spam and some as not spam). Marketing is not something we can block outright for all users.

In the figure, marketing email scanning is not enabled.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Spam Thresholds - Use the default thresholds or enter a threshold value for positively identified spam and a value for suspected spam. When evaluating messages for spam, both anti-spam scanning solutions apply thousands of rules in order to arrive at an overall spam score for the message. The score is then compared to the thresholds specified in the applicable mail policy to determine whether the message is considered spam. For highest accuracy, the threshold for positive identification as spam is quite high by default: Messages scoring between 90 and 100 are considered to be positively identified as spam. The default threshold for suspected spam is 50. Messages above the suspected threshold but below the positive-identification threshold will be considered to be suspected spam. Messages with scores below the suspected spam threshold will be considered legitimate.

Note	For other advanced settings not mentioned, please refer to the Cisco ESA user guide for more information.

https://t.me/learningnets

Spam Quarantine Configuration

This topic describes the Cisco Email Security Appliance Spam Quarantine.

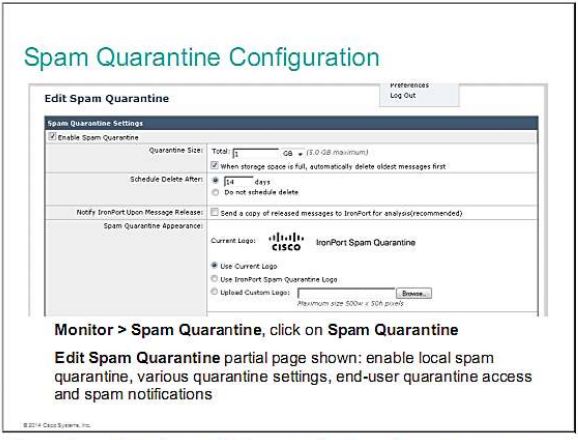

The Cisco ESA can be configured to send both spam and suspected spam to a spam quarantine. Spam quarantines are used specifically to hold spam and suspect spam messages for end users instead of deleting them immediately. A quarantine holds these messages safely for a period of time, to allow the users to review them, or to await an update that will better evaluate the safety of the message. You can also configure the system to send a notification email to users, informing them of quarantined spam and suspected spam messages. This notification contains a summary of the messages currently in the spam quarantine for that user. The user may view the messages and decide whether to have them delivered to their inbox or delete them. Users can also search through their quarantined messages. Users can access the quarantine via the notification or directly via a web browser.

The spam quarantine can be configured to be self-maintaining, meaning that mail is periodically deleted from the spam quarantine automatically in order to keep from consuming all of the quarantine space.

Each Cisco ESA can have a local spam quarantine enabled if the Anti-Spam feature has been enabled. Each Cisco ESA can also refer to an external spam quarantine, configured on another content security appliance (typically a Cisco Content Security Management appliance). However, when both the local and an external spam quarantine are enabled the local spam quarantine is used.

If you use an external spam quarantine using the **Security Services > Centralized Services > Spam Quarantine** page (not shown in the figure), the quarantine settings are configured on that external appliance.

To enable the **Local Spam Quarantine**, from the **Monitor > Spam Quaratine** page, click on the **Quaratine Name** to bring up the **Edit Spam Quarantine** page (partial page shown in the figure), click on the **Enable Spam Quarantine** check box.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

From the **Edit Spam Quarantine** page, configure the various spam quarantine settings such as the quarantine size, delete schedule, appearance and etc.., enable end-user quarantine access to allow end users to access and manage their messages in the spam quarantine directly via a web browser, and enable spam notifications.

Spam notifications are email messages sent to end users when they have messages in the spam quarantine. Notifications contain a listing of quarantined spam or suspected spam for the user. Notifications also include a link for users to use to view their quarantined messages. Once enabled, notifications are sent according to the schedule set.

Spam notifications is one of the method for end-users to access the spam quarantine. Users access the quarantine through the email notification they receive (if notifications are enabled for the quarantine). Clicking on any message subject logs the user into the web UI for the quarantine for the email address to which that notification was sent.

Enable spam notifications by checking the **Enable Spam Notifications** checkbox:

Enter a **From: address** for the notifications. Users may want to add this address to any "whitelist" supported by their email client.

Enter a subject for the notification.

Enter a customized title for the notification.

Customize the message body.

Select a message format (HTML, Text, or HTML/Text).

Specify a bounce address (bounced notifications will be sent to this address).

Optionally, you can consolidate messages sent to the same LDAP user at different addresses.

Set the notification schedule. You can configure the notifications to be sent once a month, once a week, or one or more times a day (with or without weekends).

You can enable end users to create safelists and blocklists to better control which emails are treated as spam. Safelists allow a user to ensure that certain users or domains are never treated as spam, while blocklists ensure that certain users or domains are always treated as spam. The safelists and blocklists settings are configured from the **Monitor > Spam Quarantine** page (not shown in the figure), and you must enable the **Spam Quarantine** to use this feature. When you enable the **safelist/blocklist** feature, each end user can maintain a safelist and blocklist for his or her email account. Blocklists may not effective in preventing spam, because many spammers use a technique called domainkiting, whereby spammers send from different email addresses, domains,and IPs. As a result, blocklist entries will likely be ineffective in preventing new spam messages because the source is continually changing.

Safelists and blocklists prevent email from being treated as spam or ensure that email is treated as spam. However, a safelist or blocklist setting do not prevent the Cisco ESA from scanning an email for viruses or determining if the email meets the criteria for a content-related mail policy. If an email is part of a safelist, it may not be delivered to the end user depending on other scanning settings.

Note Please refer to the Cisco ESA for more information.

https://t.me/learningnets

Policy, Virus, Outbreak Quarantines Configuration

This topic describes the Cisco Email Security Appliance Policy, Virus, Outbreak Quarantines Configuration.

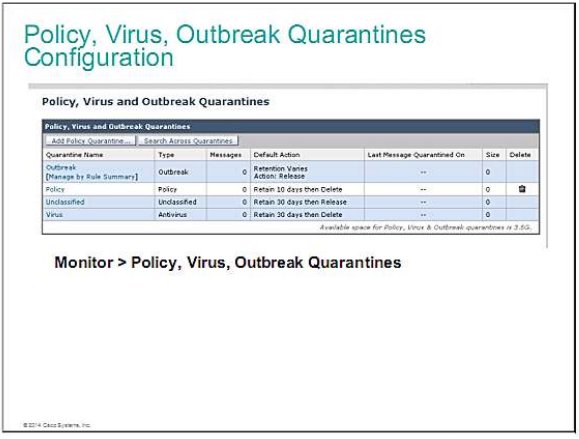

The **Monitor > Policy, Virus, Outbreak Quarantines** page is used to configure the other types of quarantines:

Policy: Holds messages caught by message filters, content filters, and DLP message actions. Let Human Resources personnel or the Legal department review messages that may contain offensive, confidential, or otherwise disallowed information.

Virus: Store messages that are marked as infected, encrypted, or not scannable by the anti-virus scanning engine to prevent the spread of viruses to your users.

Outbreak: Hold messages that are flagged by the Outbreak Filters as possibly being part of a viral outbreak or small-scale malware attack until an anti-virus or anti-spam update is released.

Unclassified: Holds messages only if a quarantine that is specified in a message filter, content filter, or DLP message action has been deleted.

These quarantine is enabled by default using the quarantine names as shown in the figure.

Click on the quarantine name to edit the quarantine settings such as the retention period and etc..

Note	Please refer to the Cisco ESA User Guide on more information.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anti-Virus Overview

This topic provides an overview of the Anti-Virus filtering.

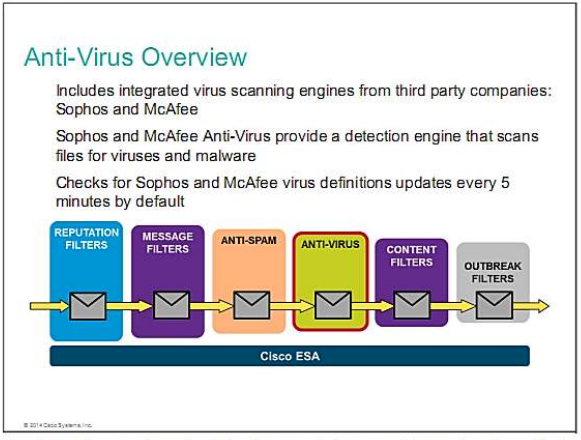

The Cisco Email Security Appliance includes integrated virus scanning engines from third party companies Sophos and McAfee. You can obtain license keys for the Cisco ESA to scan messages for viruses using one or both of these virus scanning engines, and then configure the Cisco ESA to scan for viruses using either anti-virus scanning engine. Both the Sophos and McAfee Anti-Virus scanners are able to detect and remove not only viruses, but also all types of malicious software. Scanning messages with multiple scanning engines provides defense in depth by combining the benefits of both Sophos and McAfee Anti-Virus scanning engines. Each engine has leading anti-virus capture rates. Because each engine relies on a separate base of technology for detecting viruses, the multiscan approach can be more effective. The downside is that multiple scanning engines can lead to reduced system throughput.

The McAfee and Sophos engines contain the program logic necessary to scan files at particular points, process and pattern-match virus definitions with data they find in your files, decrypt and run virus code in an emulated environment, apply heuristic techniques to recognize new viruses, and remove infectious code from legitimate files.

The default operation is to drop the message that contains the infected attachment. However, other settings include removing the infected attachment or repairing it. Both Sophos and McAfee have the ability to disinfect (repair) files. There can be limitations when it comes to disinfecting, because it is not always possible to return a file to its original state. Some viruses overwrite part of the executable program, which cannot be reinstated. In this instance, you define how to handle messages with attachments that cannot be repaired.

Sophos and McAfee frequently update their virus definitions with newly-identified viruses. These updates must be passed to your appliance. By default, the Cisco ESA is configured to check for updates every 5 minutes. For the Sophos and McAfee anti-virus engines, the server updates from a dynamic website. The system does not timeout on updates as long as the update is actively downloading to the appliance. If the update download pauses for too long, then the download times out. The maximum amount of time that the system waits for an update to complete before timing out is a dynamic value, that is defined as 1 minute less than the anti-virus update interval (defined using the **Security Services > Service Updates** page).

Use the **Security Services > Sophos or McAfee** page to verify the Cisco ESA has the latest anti-virus engine and identity files installed, and to confirm when the last update was performed. Manual updates can also be triggered from this page.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anti-Virus Configuration

This topic describes the Cisco Email Security Appliance Incoming and Outgoing Mail Policy Anti-Virus Configuration

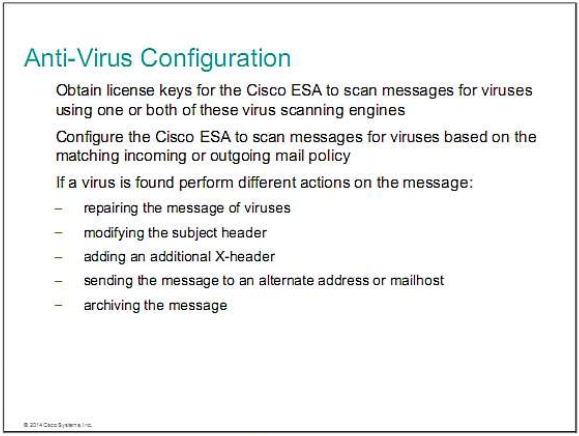

Your Cisco appliance ships with a 30-day evaluation key for each available anti-virus scanning engine. You enable the evaluation key by accessing the license agreement in the System Setup Wizard or **Security Services > Sophos/McAfee Anti-Virus** pages in the GUI. Once you have accepted the agreement, the Anti-Virus scanning engine will be enabled, by default, for the default incoming and outgoing mail policies. You need to obtain license keys for the Cisco ESA to scan messages for viruses using one or both of these virus scanning engines

If a virus is found Cisco ESA can perform different actions on the message:

repairing the message of viruses

modifying the subject header

adding an additional X-header

sending the message to an alternate address or mailhost

archiving the message

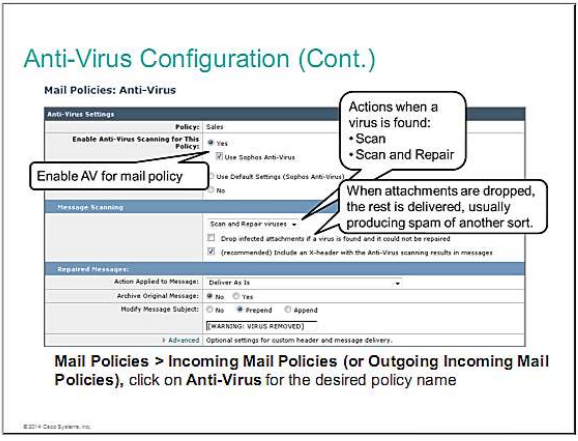

The Anti-Virus processes scan email for incoming and outgoing mail based on the mail policies that you configure. Anti-Virus is enabled by default in the default incoming mail policy and outgoing mail policy.

To configure the Anti-Virus setting for an incoming or outgoing mail police, from the **Mail Policies > Incoming Mail Policies** (or **Outgoing Incoming Mail Policies**) page, click on **Anti-Spam** for the desired policy name:

In the **Enable Anti-Virus Scanning for This Policy** section, select the Anti-Virus solution you want to use for the policy or select **Use Default Settings**. You can click **No** to disable the Anti-Virus service in this policy. In the figure, Sophos is selected, only the Sophos Anti-Virus feature key has been enabled in this example.

In the **Message Scanning** section, select:

- **Scan for Viruses only**: Messages processed by the system are scanned for viruses. Repairs are not attempted for infected attachments. You can choose whether to drop attachments and deliver mail for messages that contain viruses or could not be repaired.
- **Scan and Repair viruses**: Messages processed by the system are scanned for viruses. If a virus is found in an attachment, the system will attempt to "repair" the attachment.
- **Dropping Attachments**: You can choose to drop infected attachments. When infected attachments to messages have been scanned and dropped by the Anti-Virus scanning engine, the attachment is replaced with a new attachment called "Removed Attachment." The attachment type is text/plain and contains the following:

  This attachment contained a virus and was stripped
  Filename: filename
  Content-Type: application/filetype

In the **Repaired Messages** section, once a message has been repaired, you may:

- Deliver As Is
- Deliver as Attachment (text/plain) to New Message.
- Deliver as Attachment (message/rfc822) to New Message.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

- Archive Original Message
- Modify Message Subject
- And other Advanced settings such as to Add Custom Header to Message
- Please refer to the Cisco ESA user guide for more information.

Note You can globally enabled the virus scanning engine when you ran the **System Setup Wizard**. If you have not enabled the virus scanning engine using the system setup wizard, select the **Security Services > McAfee** or **Security Services > Sophos** page, click **Enable**. Depending on the feature keys, you can enable Sophos, McAfee, or both. After reading the license agreement, scroll to the bottom of the page and click **Accept** to accept the agreement. Click **Edit Global Settings** to choose a maximum virus scanning timeout value for the Cisco ESA to stop performing anti-virus scanning on a message. The default value is 60 seconds. Files that are too large to be completed before the timeout are categorized as unscannable. This situation may occur on the lower end Cisco ES appliances that cannot complete the scan in time (during peak loads).

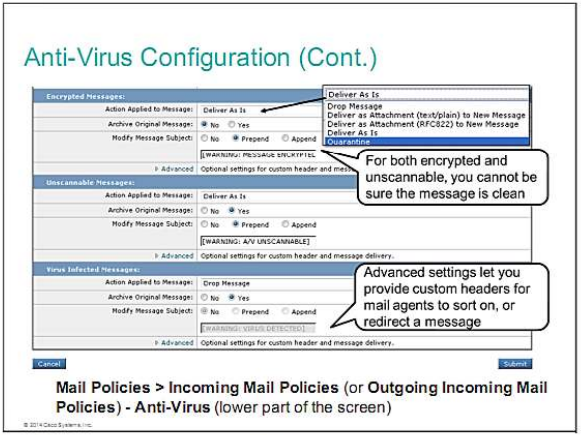

Additional Anti-Virus configuration options include:

Encrypted Messages: This outcome is for a message that is encrypted therefore unscannable. You can drop the attachment, quarantine the message, forward it to the recipient with a warning, and others.

Unscannable Messages: The message cannot be scanned, probably due to file size. The available actions are similar to the encrypted messages.

Virus Infected Messages: The attachment is determined to be virus infected, but could not be repaired or removed from the message. The default action is to drop the message, but you can it deliver as is, or as an attachment to an RFC 822-compliant message.

Content Filters Overview

This topic provides an overview of the content filters.

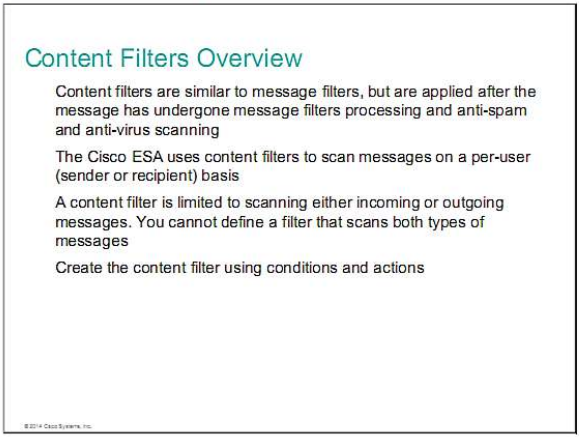

Sometimes the Cisco Email Security Appliance receives a message that should be given special treatment due to its content, whether it's because the content warrants quarantining for later examination, because corporate policy requires certain messages to be encrypted before delivery, or any number of reasons. These are cases when the message cannot be handled by the other content security features, such as anti-virus scanning or DLP. The appliance uses content filters to scan for such content and then take appropriate action on the message.

Content filters are similar to message filters, but are applied after the message has undergone message filters processing and anti-spam and anti-virus scanning. The Cisco ESA uses content filters to scan messages on a per-user (sender or recipient) basis. Content filters are similar to message filters, except that they are applied later in the email pipeline — after a message has been "splintered" into a number of separate messages for each matching mail policy.

A content filter is limited to scanning either incoming or outgoing messages. You cannot define a filter that scans both types of messages. The Cisco ESA has a separate "master list" of content filters for each type of message. The master list also determines in which order the appliance runs the content filters. However, each individual mail policy determines which particular filters will be executed when a message matches the policy.

Create the content filter using conditions and actions:

- Conditions that trigger when the Cisco ESA uses a content filter to scan a message (optional).
- Actions that the appliance takes on a message (required).
- Action variables that the Cisco ESA can add to a message when modifying it (optional) .

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

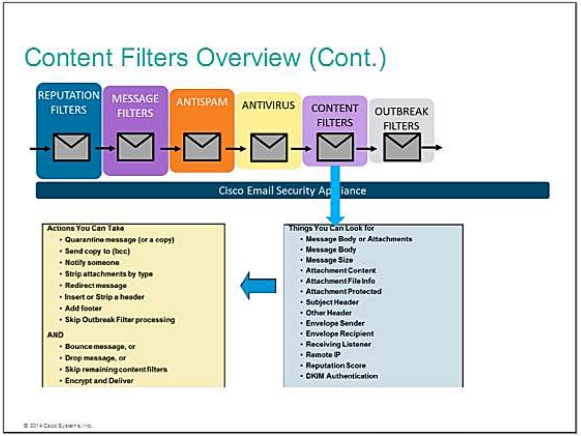

Content filters can be configured in the GUI, which includes a Rule Builder page that allows you to easily create the conditions and actions that constitute a content filter. Once the content filters are configured, the incoming or outgoing mail policies can reference the content filters that are configured and the order in which they will be applied for any given policy.

In the content filter conditions, when you add content filter rules that search for patterns in the message body or attachments, you can specify the minimum threshold for the number of times the pattern must be found. When the Cisco ESA scans the message, it totals the score for the number of matches it finds in the message and attachments. If the minimum threshold is not met, the regular expression does not evaluate to true. You can specify this threshold for text, smart identifiers, or content dictionary terms.

At least one action must be defined for each content filter (whereas the presence of a condition is optional). The actions are performed in order, so consider the order of actions when defining multiple actions for a content filter.

Content Filters Configuration

This topic describes the Cisco Email Security Appliance Incoming and Outgoing Mail Policy Content Filters Configuration

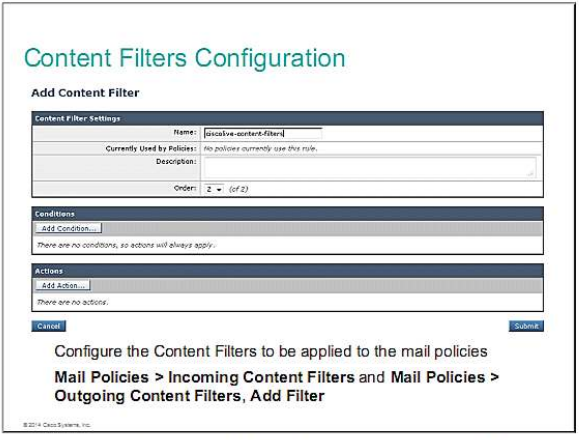

To add an incoming content filter, from the **Mail Policies > Incoming Content Filters** page, click **Add Filter**.

To add an outgoing content filter, from the **Mail Policies > Outgoing Content Filters** page, click **Add Filter**.

In the Add Content Filter screen:

Enter the name and description of the content filter. In the figure, an incoming content filter named **ciscolive-content-filters** is being added.

When editing a content filter that is in used by a mail policy, this page will also show which mail policy currently uses this content filter.

Click **Add Condition** to configure the condition(s) for this content filter. This will be shown in the next slide.

Click **Add Action** to configure the action(s) for this content filter. This will be shown in the next slide.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Content Filters Configuration (Cont.)

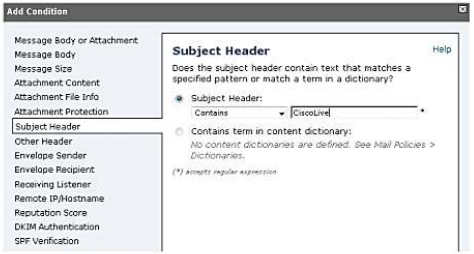

Content Filters > Add Condition

In the figure, the condition is if the **Subject Header** contains the text **CiscoLive**

Click **Add Condition** to configure the condition(s) that trigger when the Cisco ESA uses this content filter to scan a message.

In the figure, the condition is if the **Subject Header** contains the text **CiscoLive**.

Note Please refer to the Cisco ESA User Guide for more information on the other available conditions.

Content Filters Configuration (Cont.)

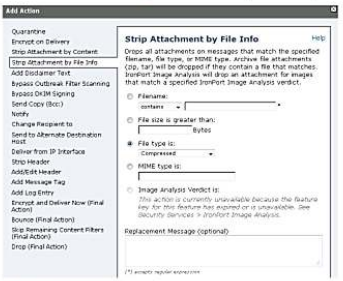

Content Filters > Add Action

In the figure, the action is to strip any compressed attachment

Click **Add Action** to configure the action(s) for this content filter.

In the figure, the action is to strip any compressed attachment.

https://t.me/learningnets

Note Please refer to the Cisco ESA User Guide for more information on the other available actions.

To apply a content filter to the incoming or outgoing mail policy, from the **Mail Policies > Incoming Mail Policies** (or **Outgoing Incoming Mail Policies**) page, click on Content Filters for the desired policy name.

In the **Content Filtering for Policy** drop down box, you can select one of the following option:

Enable Content Filters (Inherit default mail policy settings)

Enable Content Filters (Customized settings)

Disable Content Filters

In the figure, using the the Enable Content Filters (Customized settings), the **ciscolive-content-filters** content filter is enabled in the incoming mail policy named **Sales**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Outbreak Filters Overview

This topic provides an overview of the outbreak filters.

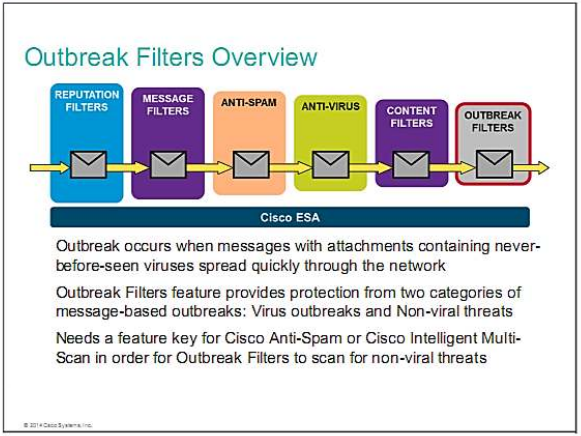

Outbreak Filters protects your network from large-scale virus outbreaks and smaller, non-viral attacks, such as phishing scams and malware distribution, as they occur. Unlike most anti-malware security software, which cannot detect new outbreaks until data is collected and a software update is published, Cisco gathers data on outbreaks as they spread and sends updated information to the Cisco Email Security Appliance in real-time to prevent these messages from reaching your users.

Cisco uses global traffic patterns to develop rules that determine if an incoming message is safe or part of an outbreak. Messages that may be part of an outbreak are quarantined until they're determined to be safe based on updated outbreak information from Cisco or new anti-virus definitions are published by Sophos and McAfee.

Messages used in small-scale, non-viral attacks use a legitimate-looking design, the recipient's information, and custom URLs that point to phishing and malware websites that have been online only for a short period of time and are unknown to web security services. Outbreak Filters analyzes a message's content and searches for URL links to detect this type of non-viral attack. Outbreak Filters can rewrite URLs to redirect traffic to potentially harmful websites through a web security proxy, which either warns users that the website they are attempting to access may be malicious or blocks the website completely.

The Outbreak Filters feature uses three tactics to protect your users from outbreaks:

Delay: Outbreak Filters quarantines messages that may be part of a virus outbreak or non-viral attack . While quarantined, the appliances receives updated outbreak information and rescans the message to confirm whether it's part of an attack.

Redirect: Outbreak Filters rewrites the URLs in non-viral attack messages to redirect the recipient through the Cisco web security proxy if they attempt to access any of the linked websites. The proxy displays a splash screen that warns the user that the website may contain malware, if the website is still operational, or displays an error message if the website has been taken offline.

Modify: In addition to rewriting URLs in non-viral threat messages, Outbreak Filters can modify a message's subject and add a disclaimer above the message body to warn users about the message's content.

The Outbreak Filters feature provides protection from two categories of message-based outbreaks: virus outbreaks, which are messages with never-before-seen viruses in their attachments, and non-viral threats, which includes phishing attempts, scams, and malware distribution through links to an external website. By default, the Outbreak Filters feature scans your incoming and outgoing messages for possible viruses during an outbreak. You can enable scanning for non-viral threats in addition to virus outbreaks if you enable anti-spam scanning on the appliance. The Cisco ESA needs a feature key for Cisco Anti-Spam or Cisco Intelligent Multi-Scan in order for Outbreak Filters to scan for non-viral threats.

Outbreak Filters are powered by Cisco's unique CASE. CASE leverages over 100,000 adaptive message attributes tuned automatically and on a regular basis, based on real-time analysis of messaging threats. For virus outbreaks, CASE analyzes the message content, context and structure to accurately determine likely Adaptive Rule triggers. CASE combines Adaptive Rules and the real-time Outbreak Rules published by Cisco Security Intelligence Operations to evaluate every message and assign a unique threat level. To detect non-viral threats, CASE scans messages for URLs and uses Outbreak Rules from SIO to evaluate a message's threat level if one or more URLs are found. Based on the message's threat level, CASE recommends a period of time to quarantine the message to prevent an outbreak. CASE also determines the rescan intervals so it can reevaluate the message based on updated Outbreak Rules from SIO. The higher the threat level, the more often it rescans the message while it is quarantined. CASE also rescans messages when they're released from the quarantine. A message can be quarantined again if CASE determines that it is spam or contains a virus upon rescan.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Outbreak Filters Configuration

This topic describes the Cisco Email Security Appliance OutBreak Filters Configuration

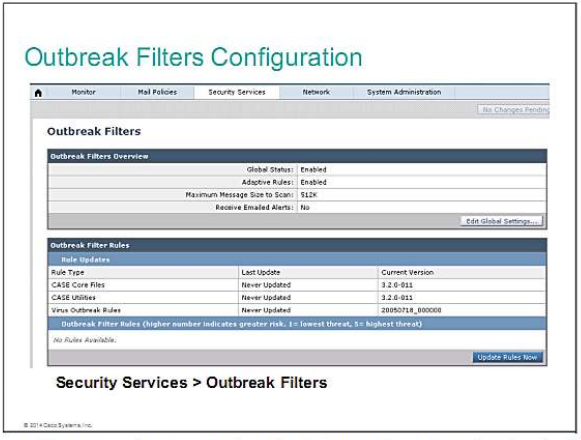

The **Outbreak Filters** page shows two sections: the **Outbreak Filters Overview** showing the outbreak filters global settings and a listing of current **Outbreak Filter Rules**.

Click **Edit Global Settings** to:

Enable the outbreak filters globally on the Cisco ESA

- To enable the Outbreak Filters feature globally, check the box next to **Enable Outbreak Filters** on the **Outbreak Filters Global Settings** page, and click **Submit**. You must have also agree to the Outbreak Filters license agreement.

Enable the adaptive rules scanning

- Adaptive Scanning uses a set of factors (file size, etc.) to determine the likelihood of a message being part of an outbreak when no virus signature or spam criteria relating to the message's content is available.

Set the maximum message size to scan in bytes.

- 512k bytes is the default size.

Enable receive emailed alerts.

- Enabling emailed alerts for Outbreak Filters merely enables the alerting engine to send alerts regarding Outbreak Filters. Specifying which alerts are sent and to which email addresses is configured via the Alerts page in the System Administration tab (not shown in the figure).

The Outbreak Filter Rules section lists the time, date, and version of the latest update for various components (the rules engine as well as the rules themselves), as well as a listing of the current Outbreak Filter rules with threat level.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Outbreak Rules are published by the Cisco Security Intelligence Operations and the Cisco ESA checks for and downloads new outbreak rules every 5 minutes by default. The update interval can be changed using the **Security Services > Service Update** page (not shown in the figure).

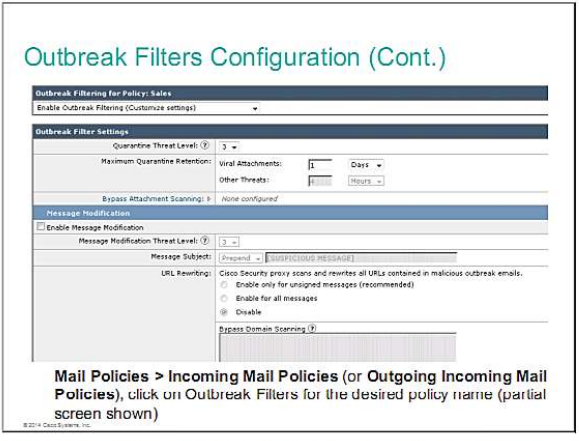

Mail Policies > Incoming Mail Policies (or **Outgoing Incoming Mail Policies**), click on Outbreak Filters for the desired policy name (partial screen shown)

To apply Outbreak Filters to the incoming or outgoing mail policy, from the **Mail Policies > Incoming Mail Policies** (or **Outgoing Incoming Mail Policies**) page, click on **Outbreak Filters** for the desired policy name.

In the **Outbreak Filtering for Policy** drop down box, you can select one of the following option:

Enable Outbreak Filters (Inherit default mail policy settings).

Enable Outbreak Filters (Customized settings).

Disable Outbreak Filters.

The following Outbreak Filter settings can be configured for a mail policy:

Quarantine threat level:

- A smaller number means that you will be quarantining more messages, while a larger number results in fewer messages quarantined. Cisco recommends the default value of 3.

Maximum quarantine retention:

- Specify the maximum amount of time in either hours or days that messages stay in the Outbreak Quarantine. You can specify different retention times for messages that may contain viral attachments and messages that may contain other threats, like phishing or malware links. You cannot quarantine non-viral threats unless you enable Message Modification for the policy.

File extension types for bypassing:

- You can modify a policy to bypass specific file types. To bypass a file extension, click **Bypass Attachment Scanning**, select or type in a file extension, and click **Add Extension**.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Message Modification: Enable Message Modification if you want to scan messages for non-viral threats, such as phishing attempts or links to malware websites. Based on the message's threat level, the Cisco ESA can modify the message to rewrite all of the URLs to redirect the recipient through the Cisco web security proxy if they attempt to open the website from the message. The Cisco ESA can also add a disclaimer to the message to alert the user that the message's content is suspicious or malicious.

- **Message Modification Threat Level**: Select a **Message Modification Threat Level** threshold from the list. This setting determines whether to modify a message based on the threat level returned by CASE. A smaller number means that you will be modifying more messages, while a larger number results in fewer messages being modified. Cisco recommends the default value of 3.
- **Message Subject**: You can alter the text of the Subject header on non-viral threat messages containing modified links by prepending or appending certain text strings to notify users that the message has been modified for their protection.
- **URL rewriting** and **Bypassing Domains**: If the message's threat level exceeds the message modification threshold, the Outbreak Filters feature rewrites all URLs in the message to redirect the user to the Cisco web security proxy's splash page if they click on any of them. The Cisco ESA rewrites all of the URLs inside a message except for the ones pointing to bypassed domains.
- Threat disclaimer: The Cisco ESA can append a disclaimer message above the heading of a suspicious message to warn the user of its content. This disclaimer can be in HTML or plain text, depending on the type of message. Select the disclaimer text you want to use from the Threat Disclaimer list or click the **Mail Policies > Text Resources** link to create a new disclaimer using the Disclaimer Template.

Note For more information, please refer to the Cisco ESA User Guide.

https://t.me/learningnets

Data Loss Prevention Overview

This topic provides an overview of the DLP outgoing mail policy feature

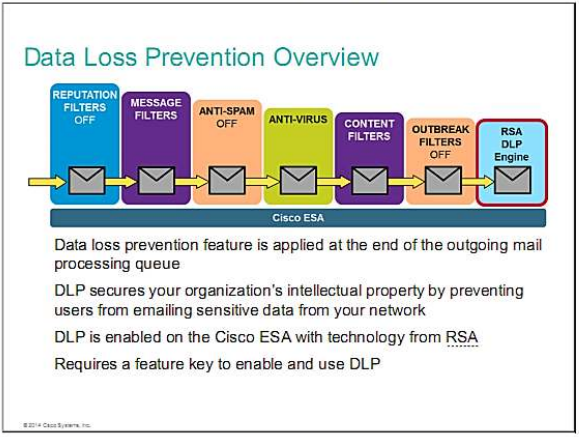

Data loss prevention feature is applied at the end of the outgoing mail processing queue. Altogether, four features are active in the outgoing mail processing: message filters, anti-virus, content filters, and DLP.

Note You can apply DLP to outgoing mail policies only. Data Loss Prevention cannot be applied to the incoming emails.

There are several areas where the law enforces that sensitive data be protected. The laws vary by country and locality. As an example from the U.S., the HIPAA, Public Law 104-191, included Administrative Simplification provisions that required the Department of HHS to adopt national standards for electronic healthcare transactions and code sets, unique health identifiers, and security. At the same time, Congress recognized that advances in electronic technology could erode the privacy of health information. Consequently, Congress incorporated into HIPAA provisions that mandated the adoption of federal privacy protections for individually identifiable health information.

The Cisco Email Security Appliance DLP feature secures your organization's proprietary information and intellectual property and enforces compliance with government regulations by preventing users from maliciously or unintentionally emailing sensitive data from your network. You define the types of data that your employees are not allowed to email by creating DLP policies that are used to scan outgoing messages for any data that may violate laws or corporate policies.

The Cisco ESA uses two engines for in-depth scanning and reducing false positives: Content matching classifier and RSA DLP policy engine and requires a feature key to enable and use DLP.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

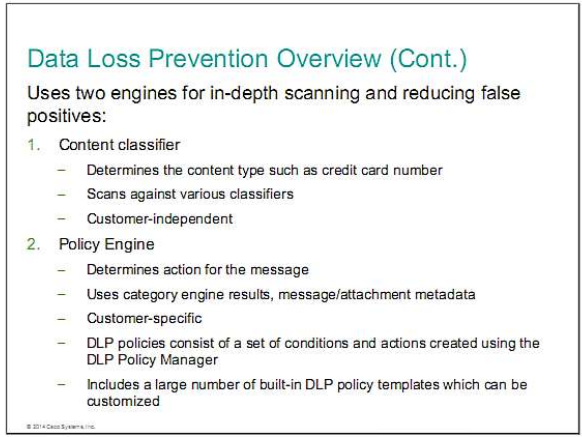

Cisco ESA uses two engines for in-depth scanning and reducing false positives:

1. **Content matching classifier:** The content matching classifier scans an outgoing message and its attachments for sensitive information, such as credit card data or other personal information. A classifier contains a number of detection rules along with context rules that impose additional requirements. As an example, consider the Credit Card Number classifier developed by RSA. This classifier not only requires that the message contains a text string that matches a credit card number pattern, but that it also contains supporting information such as an expiration date, credit card company (Visa, AMEX, and so on), or name and address. Requiring this additional information results in more accurate verdicts of a message content, leading to fewer false positives. A DLP violation occurs when a classifier detects sensitive information in a message that violates your organization's DLP rules.
2. **DLP policy engine:** At the highest level is a DLP policy, which consists of a set of conditions and a set of actions. The conditions include classifiers for the content of a message and tests for message metadata, such as sender, recipient, or attachment file type. The actions specify both the overall action to take on messages (deliver, drop, or quarantine) and secondary actions such as encrypting the message, copying it, altering its header, and sending notifications.

When someone in your organization sends a message to a recipient outside your organization, the Cisco ESA determines which outgoing mail policy applies to the sender or recipient of that message, based on rules that you defined. The Cisco ESA evaluates the content of the message using the DLP policies that are specified in that outgoing mail policy. Specifically, the Cisco ESA scans the message content (including headers and attachments) for text that matches words, phrases, predefined patterns such as social security numbers, or a regular expression that you identified as sensitive content in an applicable DLP policy. The Cisco ESA also evaluates the context of disallowed content in order to minimize false positive matches. For example, a number matching a credit card number pattern is only a violation if it is accompanied by an expiration date, credit card company name (Visa, AMEX, etc.), or a person's name and address.

DLP policies are created using the DLP Policy Manager. The Cisco ESA includes a large number of built-in DLP policy templates which can be customized. The Cisco ESA also has the ability for creating new custom DLP policies in the rare case if none of the built-in DLP policy template meet the specific requirements. If message content matches more than one DLP policy, the first matching DLP policy in the list applies, based on the order that you specified.

When potentially sensitive content appears in a message, the appliance assigns a risk factor score (severity setting) between 0 - 100 to the potential violation. This score indicates the likelihood that the message contains a DLP violation. Based on the DLP policy, DLP violations are classified as Low, Medium, High or Critical severities, and each classification can have a different action assigned.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Data Loss Prevention Configuration

This topic describes how to enable the DLP feature.

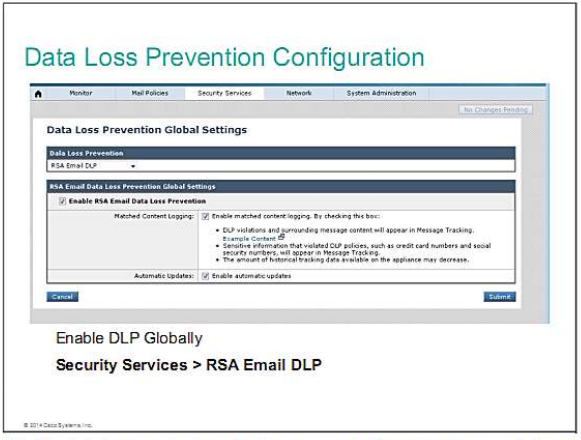

DLP is disabled by default. From the **Security Services > RSA Email DLP** page, click **Enable** to enable DLP globally.

The figure shows the Data Loss Prevention Global Settings page for the **RSA Email DLP** once DLP has been enabled globally.

Matched Content Logging: Enabling it allows all administrators to perform message tracking down into the violation. Companies should consider enabling matched content logging carefully because message tracking is either on or off. Any administrator with message tracking privileges can search DLP violations.

Automatic Updates: Enabling it allows the Cisco Email Security Appliance to get automatic DLP Engine and Content Matching Classifiers updates. Automatic updates occur every 5 minutes by default. The update interval is configurable using the **Security Services > Service Updates** page (not shown in the figure).

In the DLP drop down box, select the **RSA Enterprise Manager** option (not shown in the figure) if using the centralized DLP management solution which allows the DLP policies to be configured in RSA Enterprise Manager then pushed to all the Cisco ESAs on the network, for consistent DLP policies across the organization.

Note For more information, please refer to the Cisco ESA User Guide.

https://t.me/learningnets

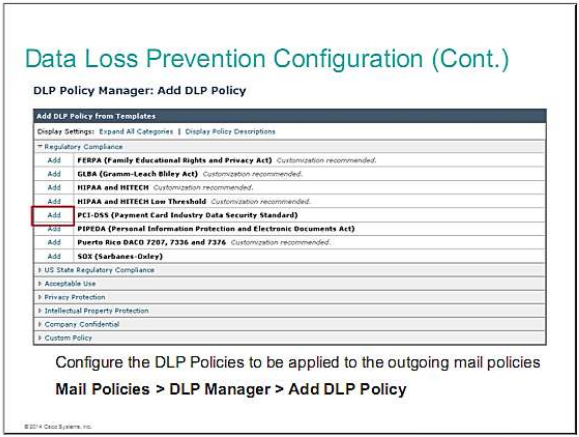

After DLP has been enabled globally, configure the DLP policies and actions using the DLP **Policy Manager**, and then enable them on the outgoing mail policies.

Use the **Mail Policies > DLP Manager > Add DLP Policy** page to add a custom DLP policy or to add a DLP policy using one of the built-in DLP policy template.

DLP policy templates are organized into the following categories:

Regulatory Compliance: Identifies messages and attachments that contain personally identifiable information, credit information, or other protected or nonpublic information.

Acceptable Use: Identifies messages sent to competitors or restricted recipients that contain sensitive information about an organization.

Privacy Protection: Identifies messages and attachments that contain identification numbers for financial accounts, tax records, or national IDs.

Intellectual Property Protection: Identifies popular publishing and design document file types that may contain intellectual property that an organization would want to protect.

Company Confidential: Identifies documents and messages that contain information about upcoming mergers and acquisitions specific to an organization.

Custom Policy: Cisco Email Security Appliance also provides the option to create your own policy from scratch using classifiers developed by RSA or your organization. This option is considered advanced and should be used only in the rare cases when the predefined policy templates do not meet the unique requirements of your network environment.

In the figure, only the top part of the DLP Policy page is shown where a DLP policy is added using the **Regulatory Compliance > PCI-DSS** built-in DLP policy template. Click **Add** next to the option as shown in the figure.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Note PCI-DSS Identifies documents and transmissions that contain credit cardholder data regulated by the PCI DSS 2.0. The Standard is applicable worldwide. The PCI DSS 2.0 defines requirements for protection of commonly used elements of cardholder and sensitive authentication data. Developed and maintained by the PCI Security Standards Council, the Standard is intended to facilitate adoption of consistent data security measures on a global scale. This policy will detect credit card track data and credit cards. It can be customized to detect credit card numbers from any of the following issuers: American Express, Diner's Club, Discover Card, JCB, MasterCard, Visa and China UnionPay.

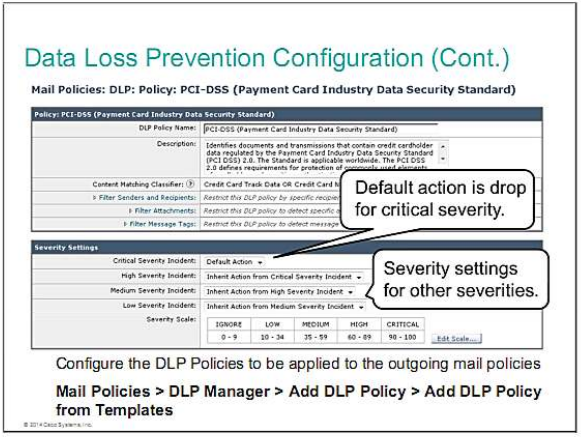

After you click **Add** in the **Mail Policies > DLP Manager > Add DLP Policy** page, the **Mail Policies: DLP: Policy: PCI-DSS (Payment Card Industry Data Security Standard)** screen will apear.

In the first section of the template, you can limit the policy to specific senders or recipients. You can also choose to match on specific attachments. These filters are a good way to test a DLP template in an initial deployment.

In the **Severity Settings** section, choose what to do with messages containing DLP violations. For example, in the figure, the action for the critical severity incidents is using the **Default Action**. The Default Action is **drop** by default. Actions for the other severities are using the action inherited from the critical severity. DLP message actions are configured from the **Mail Policies > DLP Message Actions** page as shown in the next slide.

Note For more information, please refer to the Cisco ESA User Guide.

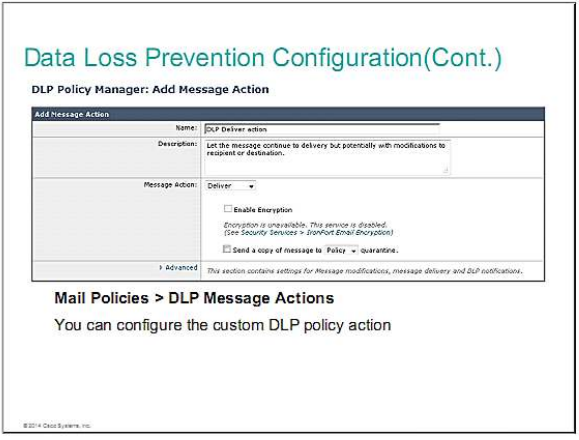

You can configure the custom DLP policy action from the **Mail Policies > DLP Message Actions** page by clicking on **Add Message Action**.

The message actions that DLP support are:

Deliver: Let the message continue to delivery but potentially with modifications to recipient or destination.

Deliver with encryption: Let the message continue to delivery but encrypt it using CRES to secure the content (if encryption service has been enabled).

Quarantine: Hold the message in the Cisco ESA Policy Quarantine.

Drop: Discard the message. No further processing is done.

In the figure, a DLP message action named **DLP Deliver action** is configured using the **Deliver** message action.

After you configure custom DLP message action, you can use this action option In the **Severity Settings** section of **Mail Policies > DLP Manager > Add DLP Policy > Add DLP Policy from Templates** screen.

The **Advanced** section (not shown in the figure) is used to configure message modifications, message delivery and DLP notifications settings.

The Cisco ESA has a default DLP policy action named **Default Action** (not shown in the figure) which has **Drop** message action assigned.

For more information, please refer to the Cisco ESA User Guide.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

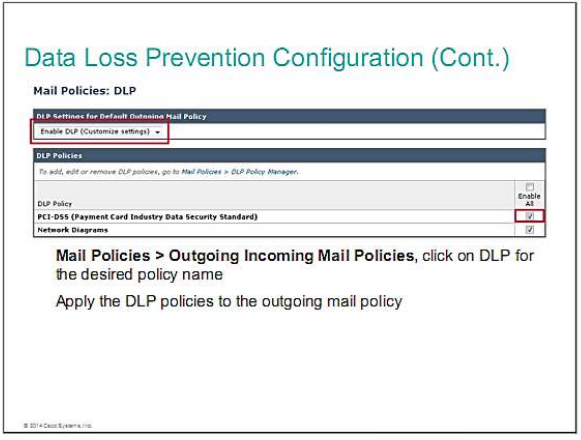

DLP policies are enabled in an outgoing mail policy only.

From the **Mail Policies > Outgoing Incoming Mail Policies** page, click on DLP for the desired policy name.

In the **DLP Setting** drop down box, you can select one of the following option:

Enable DLP (Inherit default mail policy settings)

Enable DLP (Customized settings)

Disable DLP

In the figure, the **PCI-DSS (Payment Card Industry Data Security Standard)** and the **Network Diagrams** DLP policies are applied to the outgoing mail policy.

Reporting Overview

This topic provides an overview of the Cisco Email Security Appliance reporting capabilities.

You can find many Reports options under **Monitor** in the GUI to monitor, sort, analyze, and classify the mail flows of the Cisco ESA and to differentiate between high-volume senders of legitimate mail and potential spammers (senders of high-volume, unsolicited commercial email) or virus senders. These reports can help you troubleshoot the Cisco ESA mail flow processing by including information such as the SBRS, sender groups matched, DLP incidents, content and outbreak filters, detected virus types and etc...

Reporting in the Cisco ESA involves three basic actions:

Scheduled Reports to be run on a daily, weekly, or monthly basis.

Generate a report immediately ("on-demand" report).

View archived versions of previously run reports (both scheduled and on-demand).

Configure scheduled and on-demand reports via the **Monitor > Scheduled Reports** page. View archived reports via the **Monitor > Archived Reports** page.

The Cisco ESA will retain the most recent reports it generates, up to 1000 total versions for all reports. You can define as many recipients for reports as you want, including zero recipients. If you do not specify an email recipient, the system will still archive the reports. If you need to send the reports to a large number of addresses, however, it may be easier to create a mailing list rather than listing the recipients individually. By default, the Cisco ESA archives the twelve most recent reports of each scheduled report. Reports are stored in the **/saved_reports directory** of the Cisco ESA.

The figure shows the different reports type available in the Cisco ESA.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

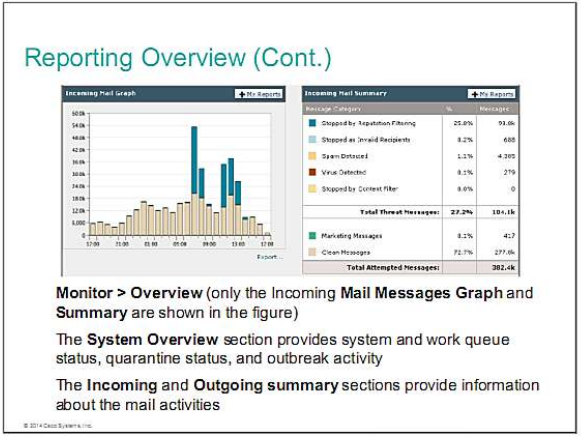

The **Monitor > Overview** page provides a synopsis of the message activity of your appliance, including an overview of your quarantines and Outbreak Filters status (in the System Overview section of the page). The Overview page also includes graphs and detailed message counts for incoming and outgoing messages. You can use this page to monitor the flow of all mail into and out of your gateway. The incoming and outgoing mail Summary Details show the number and percentage of messages categorized as clean, stopped by reputation filtering (SBRS), stopped as invalid recipient, spam detected, virus detected, stopped by content filter, and those considered clean.

The **Overview** page is divided into two sections: **System Overview** and **Incoming** and **Outgoing Mail graphs and summary.** In the figure, only the incoming messages graph and summary portion of the **System Overview** screen are shown. The **System Overview** section of the Overview page serves as a system dashboard, providing details about the appliance including system and work queue status, quarantine status, and outbreak activity.

The **System Overview** page highlights how the Cisco ESA is integrated with the SenderBase Reputation Service for incoming mail (messages stopped by reputation filtering), for example:

View a mail trend graph of all mail flows.

View a graph showing the number of attempted messages, messages stopped by reputation filtering (SBRS), messages with invalid recipients, messages marked as spam, messages marked as virus positive, and clean messages, over time.

View the summary of the system status and local quarantines.

See current virus and non-virus outbreak information.

The **Incoming** and **Outgoing** summary sections provide access to real-time activity of all mail activity on your system and is comprised of the **Incoming and Outgoing Mail Graphs** and **Mail Summaries**. You can select the time frame on which to report via the **Time Range** menu. The time range you select is used throughout all of the **Email Security Monitor** pages.

The mail trend graph on the left side of the figure shows the breakdown of incoming mail in real-time. While the mail trend graph displays a visual representation of the mail flow, the summary table on the right side of the figure provides a numeric breakdown of the information. The summary table includes the percentage and actual number of each type of message, including the total number of attempted, threat, and clean messages.

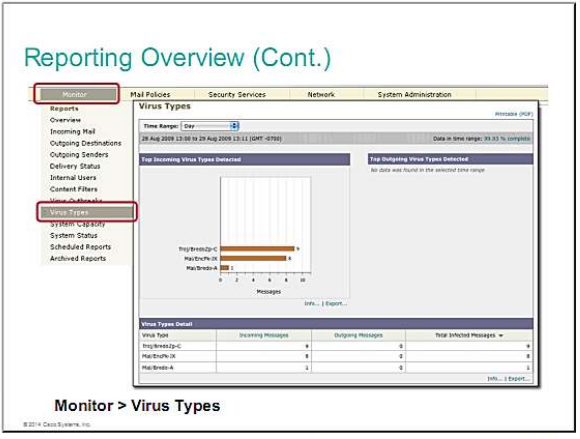

Another example for reporting on Cisco ESA is **Virus Types** report page found under **Monitor > Virus Types** in GUI.

The **Virus Types** page provides an overview of the viruses entering and leaving your network. It displays the viruses that are detected by the virus scanning engines. You can use this report to take a specific action against a particular virus. For example, if you are receiving a high volume of a viruses known to be embedded in PDF files, then you could create a filter action to quarantine messages with PDF attachments.

The **Virus Types** Details listing displays virus specifics, including the infected incoming and outgoing messages and the total infected messages, the name of the virus, and the number of incoming messages that are infected with this virus.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Message Tracking

This topic describes message tracking tool.

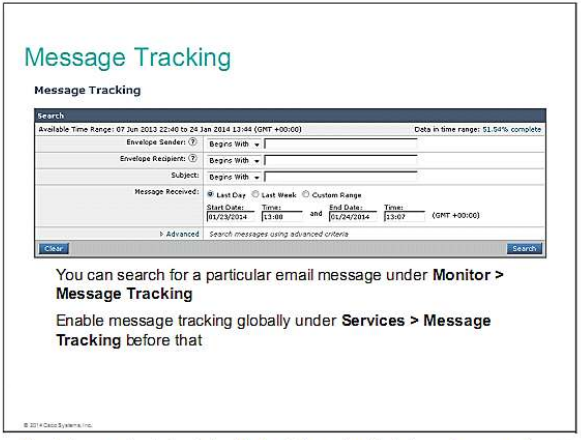

Message tracking helps resolve help desk calls by giving a detailed view of message flow. For example, if a message was not delivered as expected, you can determine if it was found to contain a virus or placed in a spam quarantine — or if it is located somewhere else in the mail stream. You can search for a particular email message or a group of messages that match criteria that you specify.

Note	You cannot use message tracking to read the content of messages.

First you need to enable message tracking globally under **Services > Message Tracking**. After message tracking is enabled, you could search for messages under **Monitor > Message Tracking.**

Enter search criteria such as **Envelope Sender**, **Envelope Recipient**, **Subject** or time range when the message was received. Click **Search** to submit the query. The results are displayed at the bottom of the page.

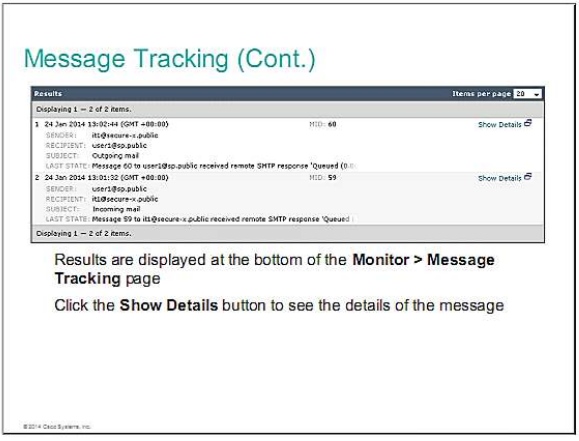

After you click **Search**, are the results displayed at the bottom of the page.

You will see the time that Cisco Email Security Appliance received the message, the sender and recipient of the message and the subject of the message. Two messages are shown in the figure - outgoing one from the **it1@secure-x.public** to the **user1@sp.public** e-mail address, and incoming one from the **user1@sp.public** to the **it1@secure-x.public** e-mail address

You can click the **Show Details** button to see the details of the message.

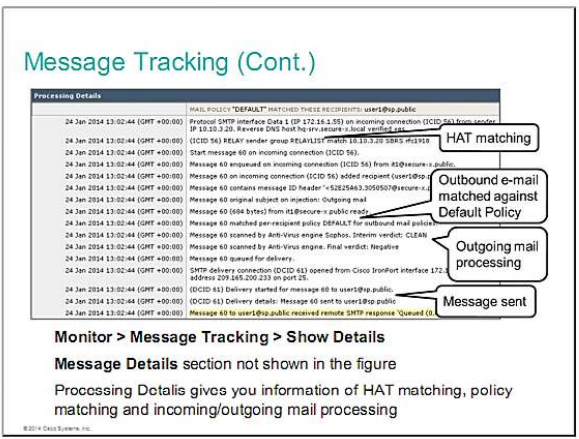

After you click **Show Details**, the details of the message will appear on the screen.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The upper part of the screen (not shown in the figure) shows **Message Details** (**Envelope and Header Summary** and **Sending Host Summary**).

The lower part of the screen shows **Processing Detalis**, as seen in the figure. You can see that mail message was matched against **RELAYLIST** sender group in the **HAT**, which has mail flow policy **RELAYED** applied to it. The connection behavior of the **RELAYED** mail flow policy is **Relay**, this is why the email is considered to be outgoing mail and is matched against default outgoing policy. The email goes through outgoing mail processing before being sent to the recipient.

https://t.me/learningnets

Trace

This topic describes the trace tool.

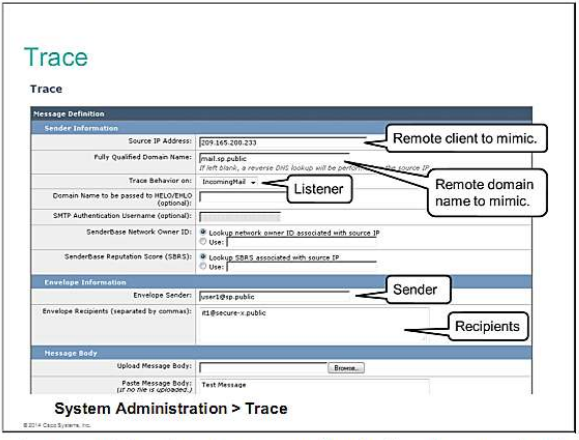

You can use **System Administration > Trace** page to debug the flow of messages through the system by emulating sending a test message. The Trace page emulates a message as being accepted by a listener and prints a summary of features that would have been triggered or affected by the current configuration of the system (including uncommitted changes). The test message is not actually sent. The **Trace** page can be a powerful troubleshooting or debugging tool, especially if you have combined many of the advanced features available on the Cisco appliance

In the figure, an email flow is emulated from the **user1@sp.public** e-mail address to the **it1@secure-x.public** e-mail address with the message body of "**Test Message**". Click the **Start Trace** button when you fill in the required fields.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

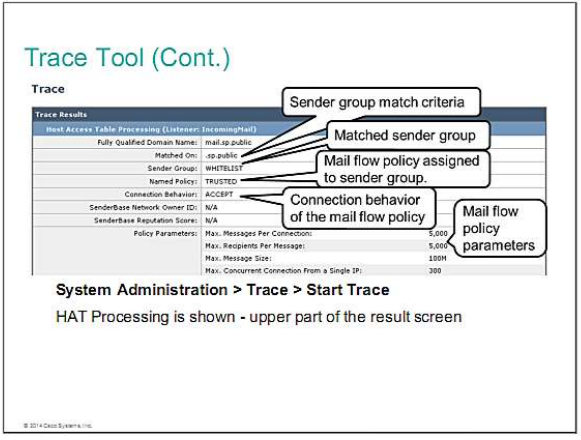

After you click **Start Trace**, the trace results are shown. The upper part of the Trace Results screen represents HAT Processing.

The figure shows that message was matched against **WHITELIST** sender group, which is matching any subdomain of **sp.public**. Mail flow policy assigned to the **WHITELIST** sender group is **TRUSTED**. Mail flow policy **TRUSTED** is defined with connection behavior **ACCEPT**. The figure also shows policy parameters for the macthed mail flow policy, which is **TRUSTED** in the example.

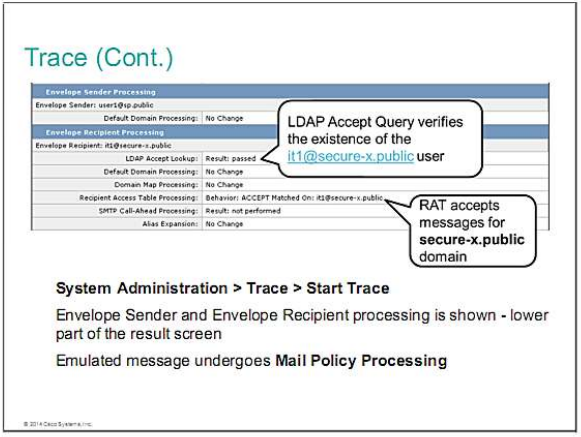

After you click **Start Trace**, the trace results are shown. The lower of the Trace Results screen shows **Envelope Sender** and **Envelope Recipient** processing information (shown in the figure) and many more - **Mail Policy Processing**, **Anti-Spam Processing**, **Anti-virus processing results** and etc... (not shown in the figure).

For example, in the figure, you can see that the **it1@secure-x.public** recipient passed the LDAP Accept Query lookup, meaning the **it1** user is a valid user in the directory and the recipient **secure-x.public** domain matches an entry in the RAT with the accept action.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

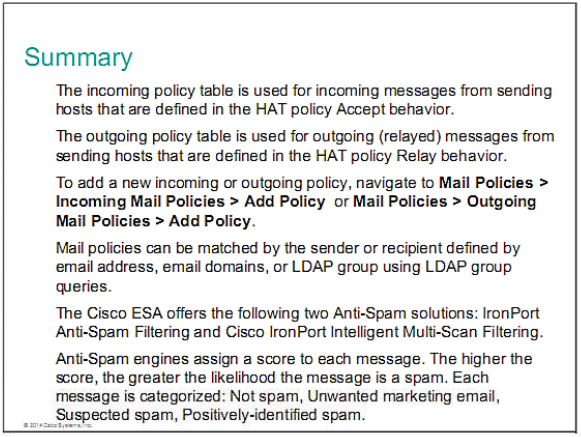

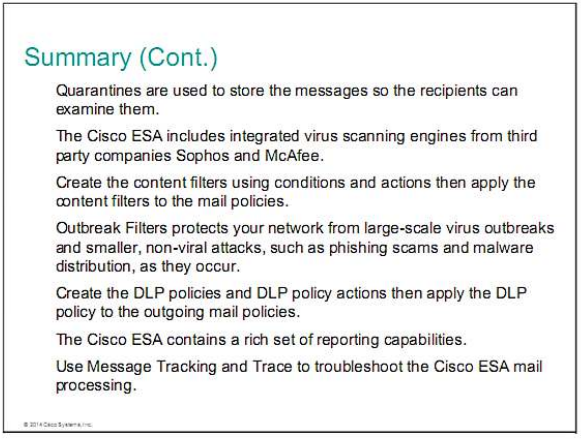

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Summary

Overview

This topic summarizes the key points that were discussed in this module.

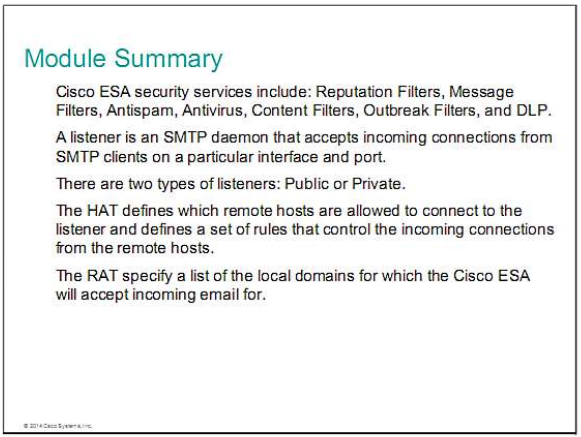

https://t.me/learningnets

Module Summary (Cont.)

- The sender group is the way to group the remote hosts that is attempting to connect to the listener (by the sender's SenderBase reputation score, IP address or host name).
- The mail flow policy defines whether the remote host specified in the sender group can or cannot connect to the listener and under what conditions, like rate limiting the connections.
- Reputation based filtering uses SBRS. SBRS has a range of -10 to +10.
- The Cisco ESA gets Anti-Virus and etc.. updates from the Cisco SIO every 5 minutes by default.
- The incoming policy table is used for incoming messages from sending hosts that are defined in the HAT policy Accept behavior.
- The outgoing policy table is used for outgoing (relayed) messages from sending hosts that are defined in the HAT policy Relay behavior.

© 2014 Cisco Systems, Inc.

Module Summary (Cont.)

- To add a new incoming or outgoing policy, navigate to **Mail Policies > Incoming Mail Policies > Add Policy** or **Mail Policies > Outgoing Mail Policies > Add Policy**.
- Mail policies can be matched by the sender or recipient defined by email address, email domains, or LDAP group using LDAP group queries.
- The Cisco ESA contains a rich set of reporting and troubleshooting capabilities.

© 2014 Cisco Systems, Inc.

References

For additional information refer to this resource:

- Latest Cisco AsyncOS for Email User Guide

https://t.me/learningnets

Module Self-Check

Questions

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

1. Which DNS record is used to locate the IP address of the receiving SMTP server? (Source: Cisco Email Security Appliance Overview)

   A. MX record
   B. PTR record
   C. A record
   D. AA record

2. Which Cisco ESA service is a cloud based service? (Source: Cisco Email Security Appliance Overview)

   A. Message Filters
   B. Outbreak Filters
   C. RSA Email DLP
   D. CRES

3. The HAT and the RAT is referenced within which Cisco ESA configuration component? (Source: Describing the Cisco Email Security Appliance Basic Setup Components)

   A. Listener
   B. SMTP Routes Table
   C. Mail Flow Policies
   D. Incoming and Outgoing Mail Policies

https://t.me/learningnets

4. Which Cisco ESA configuration component references the SBRS? (Source: Describing the Cisco Email Security Appliance Basic Setup Components)

   A. Sender Groups
   B. RAT
   C. DLP Policies
   D. Mail Flow Policies

5. Which Cisco ESA feature uses features from RSA? (Source: Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies)

   A. Anti-Spam
   B. Anti-Virus
   C. DLP
   D. Content Filters

6. Which Cisco ESA feature uses features from McAfee or Sophos? (Source: Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies)

   A. Anti-Spam
   B. Anti-Virus
   C. DLP
   D. Content Filters

7. Which Cisco ESA feature is the last process in the outgoing mail flow? (Source: Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies)

   A. Anti-Spam
   B. Anti-Virus
   C. DLP
   D. Content Filters

8. Which type of Cisco ESA policy configurations can be done using built-in policy templates? (Source: Configuring Cisco Email Security Appliance Basic Incoming and Outgoing Mail Policies)

   A. Message Filters
   B. Content Filters
   C. Outbreak Filters
   D. DLP
   E. LDAP Queries

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Answer Key

1. C
2. D
3. A
4. A
5. C
6. B
7. C
8. D

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module 5

Cisco Intrusion Prevention Systems

Introduction

Protecting your business assets is becoming more complicated as applications become more dynamic and as more users access enterprise and cloud-based resources from mobile devices. Signature-only detection products lead to a one-dimensional response. Cisco traditional and next generation Intrusion Prevention System (IPS) solutions can identify and mitigate attacks with market-leading, context-aware threat prevention that augments your firewall and VPN deployments. Cisco IPS devices intelligently determine the best response by looking at multiple factors, including the network reputation of an incoming flow's source, the target's value to the organization, the target's operating system, and the user identity associated with the flow.

The Cisco SenderBase correlates real-time data from more than 1.5 million devices around the world to create network reputation scores that enable Cisco IPS devices to block threats from known malicious hosts before they pass through the IPS inspection process.

Upon completing this module, you will be able to:

- Describe the basic definitions and approaches to traditional intrusion prevention/detection systems and next generation IPS
- Configure different Cisco IPS sensor interface modes
- Configure the Cisco IPS sensor built-in signatures
- Describe some methodologies for tuning a Cisco IPS sensor to properly manage false positive and negative events
- Describes the methods and configuration procedures to create custom signatures on a Cisco IPS sensor
- Enable the anomaly detection functionality on the Cisco IPS sensor
- Enable the reputation-based features on the Cisco IPS sensor

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 1

Describing Cisco IPS Solutions

Protection of company resources is becoming increasingly important in modern IT infrastructures. Not only must company resources be defended passively, but the protection also must be constantly monitored and enhanced with systems that actively inspect the data that is passed over allowed and open connections. Intrusion prevention and intrusion detection systems are two technologies that can provide such risk reduction. This lesson describes the basic definitions and approaches to traditional intrusion prevention/detection systems and provide a high level overview of the NGIPS.

Upon completing this lesson, you will be able to meet these objectives:

- Identifies how to place the Cisco IPS into a typical modularized network design
- Compare intrusion detection systems and intrusion prevention systems
- Explain the terminology that is used in intrusion prevention and detection
- Describe network intrusion prevention approaches
- Describe the Cisco Sourcefire NGIPS Benefits
- Describe the NGIPS feature with Cisco Next Generation Firewall Services
- Describe the Cisco ASA NGFW Services NGIPS Threat Profile
- Describe how to apply an IPS Threat Profile within a NGFW Services Access Policy
- Describe the Cisco NGFW Services NGIPS Global Settings
- Describe an Cisco NGFW Services NGIPS triggered Event

https://t.me/learningnets

Cisco Modular Network Architecture and Cisco IPS

This topic identifies how to place the Cisco Intrusion Prevention System into a typical modularized network design.

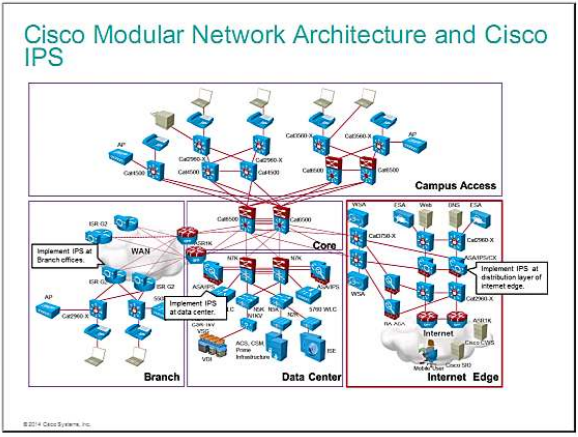

The figure illustrates where in the Cisco modular network architecture the Cisco IPS Appliance or Module should be deployed. Cisco IPS may be implemented in the distribution part of the Internet edge infrastructure. Logically, the IPSs are placed between the firewall and core routers, protecting the enterprise from threats originating from campus and remote users.

The Internet edge is a public-facing network infrastructure and is particularly exposed to large array of external threats.

IPS provide filtering of known network worms and viruses, DoS traffic, and directed hacking attacks. This functionality is highly beneficial in a campus environment by quickly identifying attacks and providing forensic information so that attacks can be cleaned up before substantial damage is done to network assets. IPS is designed to monitor and permit all traffic that is not malicious and can be deployed with a single campus-wide protection policy allowing for a pervasive campus-wide IPS deployment.

To get the most benefit, IPS needs to cover a majority of the network segments in a campus. Since the distribution switches provide the aggregation point for all of the access switches and provides connectivity and policy services for traffic flows between the access layer and the rest of the network, it is recommended that IPS devices be deployed at the distribution layer in the campus network.

IPS are also commonly deployed at the branch offices typically using an integrated IPS module into the branch router of firewall.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IPS deployment is driven by three key design considerations:

Deployment Model: Cisco IPS appliances and modules can be deployed in inline or promiscuous mode, typically referred to as IPS or IDS modes. When deployed in the inline mode, the Cisco IPS is placed in the traffic path. This allows the sensor to stop attacks by dropping malicious traffic before it reaches the intended target, thus providing a strong protective service. IPS inline mode enables automatic threat detection and mitigation capabilities that offer some clear advantages in terms of timely threat mitigation and degree of protection. In addition, signature tuning enables the automated response actions to be tuned according to customer policy. Since IPS is in the data path, however, it is critical to ensure that the deployment is well designed, architected and tuned to ensure that it does not have a negative impact on network latency, convergence, and service availability.Cisco IPS can also be deployed in promiscuous mode. In this mode, the IPS is not in the data path, but rather performs passive monitoring, with traffic being passed to it through a monitoring port using traffic mirroring techniques such as SPAN, RSPAN, or VACL capture. The Cisco IPS sensor analyzes a copy of the monitored traffic rather than the actual forwarded packet. Upon detection of anomalous behavior, management systems are informed of an event and operational staff can subsequently decide what action, if any, to take in response to an incident. The time between threat detection and mitigation may thus be extended and requires a manual response.

Scalability and Availability: For scalability, Cisco offers a range of different IPS platforms that can be deployed according to a particular customer's needs. For networks with a high level of activity and traffic, IPS appliances can be used. For networks with a lower level of activity and traffic or where integrated security is preferred, Cisco IPS modules are viable options. For increased scalability and high availability, multiple IPS sensors can be bundled together using a load-balancing mechanism. IPS load-balancing can be accomplished using the ECLB feature within a switch to perform intelligent load-balancing across the IPS devices. In addition, multiple IPS sensors may also be deployed by using an external load-balancing appliance.

Traffic Symmetry: For maximum visibility, the IPS sensors must be able to see traffic in both directions. For this reason, it is important to ensure the symmetry of the traffic as it traverses or reaches the IPS sensor. Symmetrical traffic flows offer a number of important benefits, including enhanced threat detection, reduced vulnerability to IPS evasion techniques, and improved operations through reduced false-positives and false-negatives. Consequently, this is a key design element.

https://t.me/learningnets

Cisco Modular Network Architecture and Cisco IPS (Cont).

Cisco IPS Solutions Overview

Cisco IPS Appliance	Cisco IPS Software Module	Cisco IPS Hardware Module	Cisco IOS IPS	Cisco NGIPS
Cisco IPS 4300 Series	Cisco ASA 5500-X IPS software module	Cisco ASA 5585-X IPS-SSP	Cisco IOS Software	Cisco ASA NGFW services with NGIPS
Cisco IPS 4500 Series		Cisco ISR G2 NME-IPS		SourceFire NGIPS

© 2014 Cisco Systems, Inc.

The Cisco IPS solutions scale to serve a wide range of deployment scenarios, from small offices and branch locations to enterprise data center architectures.

The figure introduces the Cisco IPS 4300 and 4500 series appliances, Cisco Adaptive Security Appliance 5500-X IPS software module, ASA 5585-X IPS-SSP, Cisco Integrated Services Router G2 NME-IPS, IOS IPS and the latest Cisco Next Generation IPS: ASA NGFW services with NGIPS and SourceFire NGIPS:

The Cisco IPS 4300 series appliances is a one rack-unit form factor designed for edge, campus, and branch environment with performance ranges from 350 Mbps to 2 Gbps.

The Cisco IPS 4500 series appliances provides data center specific protection for servers, databases, enterprise and custom software with performance ranges from 5 Gbps to 20 Gbps.

The Cisco ASA 5500-X midrange models run the Cisco ASA IPS module as a software module with performance ranges from 250 Mbps to 1.3 Gbps.

The Cisco ASA 5585-X runs the Cisco ASA IPS module as a hardware (IPS-SSP) module with performance ranges from 2 Gbps to 10 Gbps.

The Cisco ISR G2 NME-IPS integrates IPS in a full-service branch-office solution. By integrating IPS and branch-office routing, Cisco integrated services routers can secure remote branch-office networks from threats originating from the Internet and reduce the WAN link overload from infected hosts at the branch office. The NME-IPS can monitor up to 45 and 75 Mbps of traffic.

The Cisco IOS IPS is an inline, deep-packet inspection feature that effectively mitigates a wide range of network attacks. A component of the Cisco IOS Integrated Threat Control framework and complemented by Cisco IOS Flexible Packet Matching feature, Cisco IOS IPS provides your network with the intelligence to accurately identify, classify, and stop or block malicious traffic in real time.

Cisco ASA NGFW services (NGFW version 9.2.1 with ASA version 9.1.3 and later) with IPS provide context-driven threat detection and mitigation. The simplified operation puts focus on threat prevention rather than on detection parameters. Inputs from the Cisco Application Visibility and Control and WSE security services optimize the Cisco IPS's operation and efficacy to provide holistic threat prevention.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Sourcefire Next-Generation Network Security solutions set a new standard for advanced threat protection, integrating real-time contextual awareness, intelligent security automation, and unprecedented performance with industry-leading network intrusion prevention.

Cisco IPS solutions deliver:

A wide and deep body of protection, with more than 5500 signatures.

Patented anti-evasion technology to defend and monitor against worms, viruses, Trojans, reconnaissance attacks, spyware, botnets, undesired applications, and malware.

Protocol and behavior analysis for final threat determination.

Cisco Global Correlation to help identify the source of DoS, DDoS, SYN flood, and encrypted attacks so you can block them.

Specific protections for Unified Communications, WLAN, routing, and switching to help guard the network infrastructure.

Management solutions for smaller deployments, as well as enterprise-class coverage. Cisco IPS Manager Express is an all-in-one IPS management and reporting application for up to 10 devices. Cisco Security Manager is an enterprise-class security management application with thousands of real-world deployments.

Subscriptions for IPS signature updates require a license that can be included in the Cisco SMARTnet services contract for an additional cost.

https://t.me/learningnets

Intrusion Detection Systems vs. Intrusion Prevention Systems

This topic compares intrusion detection systems and intrusion prevention systems.

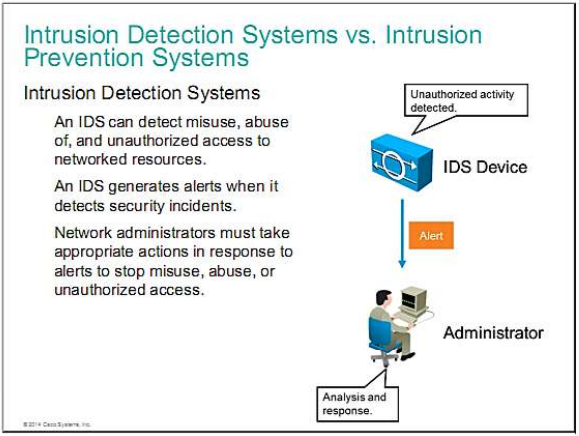

An IDS is a security control (countermeasure) that can detect misuse, abuse of, and unauthorized access to networked resources. The Cisco intrusion protection product portfolio consists of various devices that are called sensors, all of which can monitor traffic from or through a particular network segment, analyze it, detect malicious activity, and perform a response action if the traffic is deemed malicious.

An IDS is usually a dedicated device that monitors network traffic and detects anomalies based on certain criteria. These criteria can be a database of signatures, a statistical knowledge of what represents normal network traffic, or an administrator-specified security policy.

The following are the most commonly detected attacks by a network IDS:

Network sweeps and scans, which can indicate network reconnaissance

Common network anomalies on most OSI layers, which include the following:

- Malformed ARP requests or replies
- Invalid IP datagrams
- Invalid TCP packets (for example, a source or destination port is 0)
- Malformed application layer protocol units (for example, an HTTP request that does not begin with GET, POST, HEAD, or another valid HTTP command)

Flooding DoS attacks, which can come in the form of a large number of ICMP packets or TCP synchronization (TCP SYN) packets. These attacks can affect the resources of a system and severely degrade performance. The attack could force the system to consume all its memory, as in the case with a TCP SYN flood, because the system must reserve a certain amount of memory for each connection that is set up.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Application layer attacks, such as buffer overflows, directory traversal attacks, or various forms of command injection

An IDS generates alerts when it detects offending traffic. The alerts are stored locally on the IDS and may also be retrieved by a management system. An IDS does not stop malicious traffic by itself. Network security administrators must examine the alerts that are generated by IDS and react accordingly.

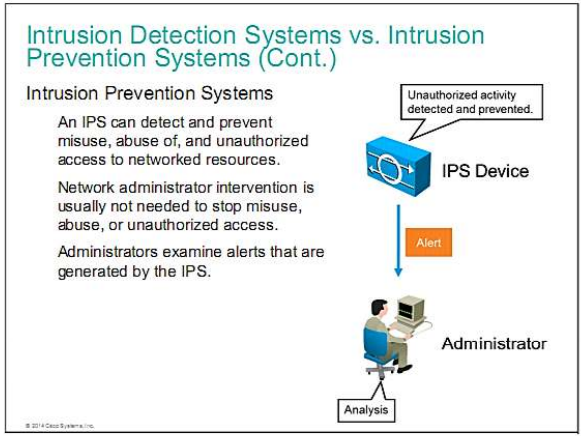

An IPS is a security control (countermeasure) that can detect and prevent misuse, abuse of, and unauthorized access to networked resources. All Cisco sensors can perform inline intrusion prevention.

The Cisco Intrusion Prevention System is a network-based solution designed to accurately identify, classify, and stop malicious traffic—including worms, spyware and adware, network viruses, and application abuse—before it affects business activities.

In contrast to an IDS, an IPS can stop offending traffic. Network security administrator intervention is usually not needed to stop misuse, abuse, or unauthorized access, because IPS sensors take immediate, real-time actions to protect network resources. Network administrators may still need to examine and analyze alerts that are generated by IPS.

Intrusion Prevention Systems Terminologies

This topic explains the terminology that is used in intrusion prevention and detection.

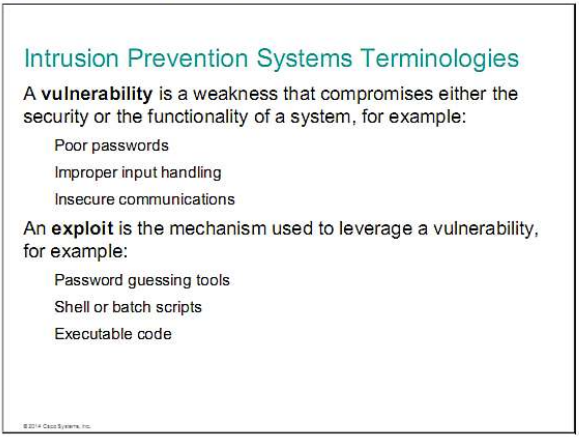

Vulnerability is a weakness that compromises either the security or the functionality of a system, for example:

Poor passwords: Passwords are the first line of defense. Weak or easily guessed passwords are considered vulnerabilities.

Improper input handling: Software that does not properly manage all possible input can have unexpected results. Often, this leads to either a DoS or access to restricted system resources.

Insecure communications: Data that is transferred in plaintext is susceptible to interception. System passwords, employee records, and confidential company documents are examples of data that is vulnerable to interception.

An exploit is the mechanism that is used to leverage a vulnerability to compromise the security or functionality of a system, for example:

Password guessing tools: These tools attempt to "crack" passwords by using knowledge of the algorithm that is used to generate the actual password or by attempting to access a system using permutations and combinations of different character sets. Two popular password cracking tools are L0phtCrack and John the Ripper.

Shell or batch scripts: These scripts are created to automate attacks or perform simple procedures that are known to expose vulnerability.

Executable code: Exploits that are written as executable code require programming knowledge and access to software tools such as a compiler. Consequently, executable code exploits are considered to be more advanced forms of exploitation.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The concepts of threat, vulnerability, risk, and asset value are easy to relate using a simple risk equation, in which risk equals threat multiplied by vulnerability multiplied by asset value, and in which vulnerability and threat have the following further dependencies:

- Vulnerability of an asset depends on the presence and quality of security controls that are deployed to guard an asset.
- Threat depends on the capability of the attacker and on frequency of attacks.

Such an equation is useful when prioritizing risks to business because security engineers are expected to first address the most pressing (highest) risks. Based on this equation, an engineer can intuitively see that high risk results from a combination of high threat (high attacker capacity and attack frequency), high vulnerability (absence or ineffectiveness of security controls), and high asset value.

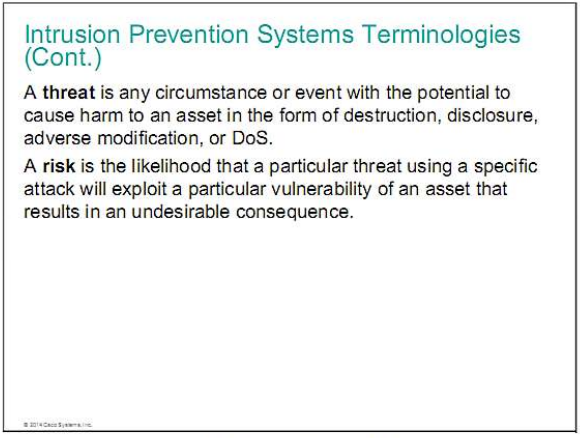

A threat is any circumstance or event with the potential to cause harm to an information system in the form of destruction, disclosure, adverse modification of data, or DoS. Examples of threats are application layer network attacks against exposed application servers, malicious software (malware) targeting workstations, and physical destruction of a server.

A risk is the likelihood that a particular threat using a specific attack will exploit a particular vulnerability of an asset that results in an undesirable consequence.

Intrusion Prevention Systems Terminologies (Cont.)

Accuracy

False positive: A security control acted when malicious activity did not take place.

False negative: A security control did not act when malicious activity took place.

True positive: A security control acted when malicious activity took place.

True negative: A security control did not act because there was no malicious activity.

© 2014 Cisco Systems, Inc.

Security controls such as IPS or IDS deployed in real environments may produce erroneous decisions, either because of their misconfiguration or because of the environment, in which legitimate activity may resemble malicious activity, and vice versa.

All decisions of security controls can be classified as one of the following:

True positives: The security control, such as an IPS or IDS sensor, acted as a consequence of malicious activity. This classification represents normal and optimal operation.

False positives: The security control acted as a consequence of nonmalicious activity. This classification represents an error, generally caused by proactive controls that are too tight (they do not permit all legitimate traffic) or by reactive controls that are too relaxed (with too broad descriptions of the attack).

True negatives: The security control has not acted because there was no malicious activity. This classification represents normal and optimal operation.

False negatives: The security control has not acted, even though there was malicious activity. This classification represents an error, generally caused by proactive controls that are too relaxed (they permit more than just minimal legitimate traffic) or by reactive controls that are too specific (with too specific descriptions of the attack).

These decisions demonstrate that addressing false-positive issues or false-negative issues appears to be a balancing task. As you tune the system to be less restrictive to remove false positives, you will increase the likelihood of false negatives, and vice versa. Therefore, security control systems must often be expertly tuned to achieve an acceptable balance between these decisions.

Preventive controls (for example, IPS sensors) are generally tuned to be less sensitive in order not to block legitimate traffic. Detective controls (for example, IDS sensors) are tuned to be more sensitive at a cost of false positives.You often can combine a relaxed preventive control with a sensitive detective control to gain insight into the false negatives of the preventive control and enable incident response, if the preventive control is bypassed.

https://t.me/learningnets

Traditional Network Intrusion Prevention Systems

This topic describes network intrusion prevention solutions.

Traditional Network Intrusion Prevention Systems

Network IPS sensors are devices tuned for intrusion prevention analysis.

- The operating system is hardened.
- The hardware is dedicated to traffic analysis.

Network IPS Sensor

© 2014 Cisco Systems, Inc.

A network IPS is a security control that analyzes network traffic and attempts to block malicious network traffic to prevent security incidents. A network IPS analyzes network traffic from multiple aspects.

Traditional Network Intrusion Prevention Systems (Cont.)

Signature-Based IPS

Analyze network traffic and act if a known malicious event is detected:

- A database of known malicious traffic is required.
- The database must be continuously updated.

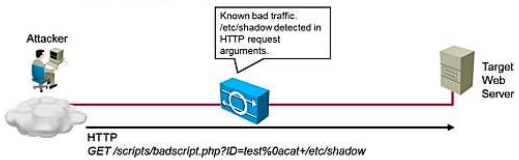

© 2014 Cisco Systems, Inc.

A signature-based network IPS analyzes network traffic and compares the data in the flow against a database of known attack signatures. A signature-based IPS looks at the packet headers or data payloads when it analyzes network traffic. For example, a signature might be a sequence or a string of bytes in a certain context. Here are some examples of types of attack:

Attacks against a web server may be in the form of specially crafted URLs. Therefore, the IPS looks for a malicious data pattern at the start of the HTTP data flow, which begins with an HTTP request from the client.

An attack against an SMTP server could be in the form of a buffer overflow in the "mail from" command of the SMTP session. The IPS looks for an attack signature in the SMTP session that starts with the "mail from" command and includes a malicious data pattern before the end of the line.

An attack on the mail client can be in the form of a buffer overflow in the MIME header of the message itself. The IPS looks for the sequence of bytes that identifies the start of a new MIME part in the message and a sequence of bytes that comprise a buffer overflow following it.

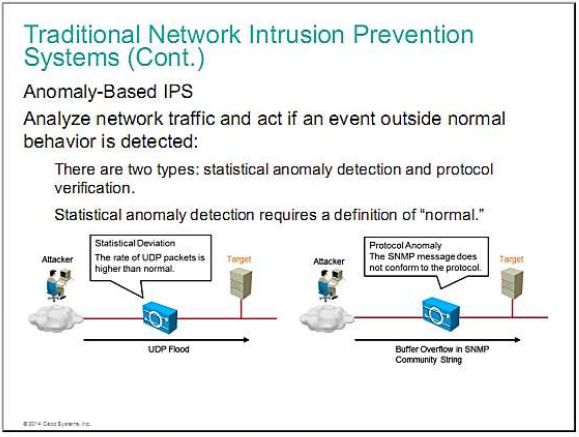

Anomaly-based network IPS devices observe network traffic and act if a network event outside normal network behavior is detected. There are two types of anomaly-based network IPSs:

Statistical anomaly detection (also often called network behavior analysis) observes network traffic over time and builds a statistical profile of normal traffic behavior based on communication patterns, traffic rate, mixture of protocols, traffic volume, and so on. After a normal profile has been established, statistical anomaly detection systems detect or prevent activity that violates the normal profile.

Protocol verification observes network traffic and compares network, transport, and application layer protocols that are used inside network traffic to protocol standards. If a deviation from standards-based protocol behavior is detected (for example, a malformed IP packet or a malformed SNMP message), the system can take appropriate action.

https://t.me/learningnets

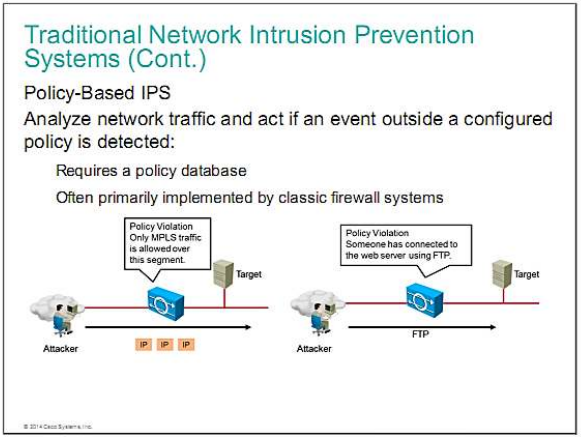

A policy-based IPS analyzes network traffic and acts if it detects a network event outside a configured traffic policy. The traffic policy typically describes a small set of allowed communications over a network segment. Therefore a policy-based IPS is very much like a class firewall and is often implemented using a firewall product.

Sourcefire NGIPS Benefits Overview

This topic describes the Cisco Sourcefire NGIPS benefits.

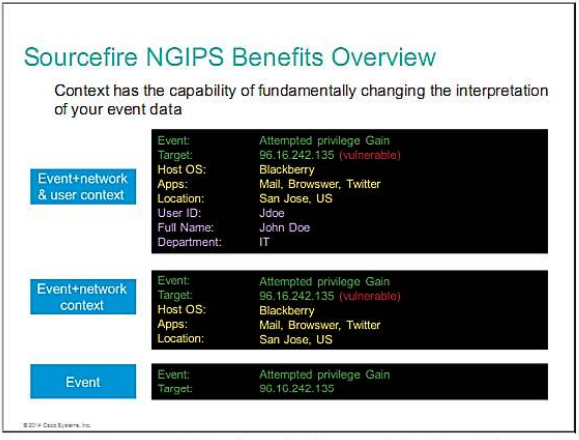

Today's information systems complexity, advanced malware, and relentless attackers are driving the evolution of security into a real-time threat detection game. To stay on top of security, you need to fully know what's taking place on your network and hosts and you need to better understand the context of the events. Today's IT organizations need a dynamic approach to defending the network—one that uses awareness and automation to provide visibility and context while constantly adapting to new threats, new vulnerabilities, and everyday network changes. Security awareness and controls built into NGIPS provide this deeper insight into and management of what's actually happening in and around your network.

NGIPS offers a logical and essential progression of capabilities needed to protect networks from emerging threats. Pioneered by Sourcefire (now part of Cisco), and endorsed by Gartner, the NGIPS builds on typical IPS solutions by providing contextual awareness—about network activity, systems and applications, people, and more—to promptly assess threats, ensure a consistent and appropriate response, and reduce an organization's security expenditures.

These benefits of next generation IPSs not only help manage risks but they also allow information security team members to focus their response efforts and work more efficiently when incidents do occur.

As shown in the figure, imagine you manage a network and you see the events generated without context. Then you see the extra information network context brings. Then you see the extra information user context brings. What would you rather have if you were managing the network? No context, network context or network + user context? What's interesting about this scenario is that in all three cases, the event data does not change. Yet how we respond to the event data fundamentally changes, based on the contextual data available to us at the time.

Cisco Sourcefire's NGIPS with advance malware prevention and content-awareness benefits:

Context: Generate a rich set of context in which to embed into a security event. Such events have immediate meaning and relevance that is simply not clear in non-contextual security controls.

Speed: Sourcefire NGIPS effortlessly scales from 50Mbps to 40Gbps.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Accuracy: Sourcefire IPS detection rates (as measured by NSS Labs) are second to none. Sourcefire NGIPS have the best detection rate in the industry. NSS Labs is a trusted authority in independent product testing, especially when it comes to research and analysis on Intrusion Prevention Systems.

Flexibility: Sourcefire systems are customizable in many ways. In rule sets (Snort). In user interface (workflows). In meta-rules (Sourcefire Management Center correlation rules). In management interface (Sourcefire Event Viewer - eStreamer, Java Database Connectivity - JDBC, Open Platform for Security Application Programming Interface - OPSEC API, Syslog, SMTP, SMNP).

Value: Sourcefire systems' extensive automation capabilities allow you to realize real reductions in analyst time to identify and address intrusions. In doing so, Sourcefire can materially reduce ongoing operational cost. Sourcefire system offer the lowest cost per protected megabyte and also the lowest administrative cost.

https://t.me/learningnets

ASA NGFW Services with NGIPS Overview

This topic describes the NGIPS feature with Cisco Next Generation Firewall Services.

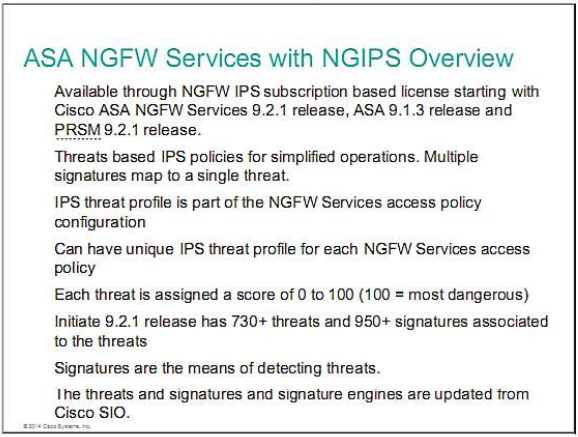

The NGFW NGIPS feature is available through NGFW IPS subscription based license starting with Cisco Adaptive Security Appliance NGFW Services 9.2.1 release, ASA 9.1.3 release and PRSM 9.2.1 release.

NGIPS filtering analyzes network traffic in real time, comparing the traffic contents against known threats.

The Cisco Security Intelligence Operations develops signatures that identify threats. Multiple signatures can map to a single threat. New threats and signature sets are updated by default every hour unless you disable updates.

With NGIPS, threats are the focus not signatures. NGIPS filtering analyzes network traffic in real time, comparing the traffic contents against known threats. If a connection matches a threat, you can drop the connection to block the threat. You can also choose to monitor but allow, or completely ignore, threats that you decide are benign.

NGIPS using threats based IPS policies eliminate the needs for complex individual IPS signatures tunings.

Each threat is assigned a score with a score of 100 being the most dangerous and 0 being least dangerous.

In the initiate 9.2.1 release, it has 730+ threats and 950+ signatures associated to the threats. Signatures are used as the mean to detect the threats. The threats and signatures and signature engines are updated from the Cisco SIO. By default, the update server is checked every 5 minutes throughout the day seven days a week.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

ASA NGFW Services NGIPS Threat Profile

This topic describes the Cisco Adaptive Security Appliance NGFW Services NGIPS Threat Profile.

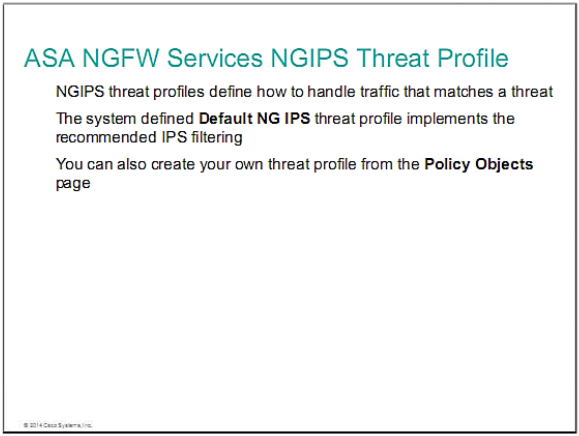

You can configure device-level and per-access policy settings for NGIPS filtering. NGIPS threat profile define how to handle traffic that matches a threat.

The system defined **Default NG IPS profile** implements the recommended IPS filtering, you can also create your own threat profile or select any other already defined threat profile. You can create NGIPS threat profiles from the **Policy Objects** page, or you can create them while editing the access policies.

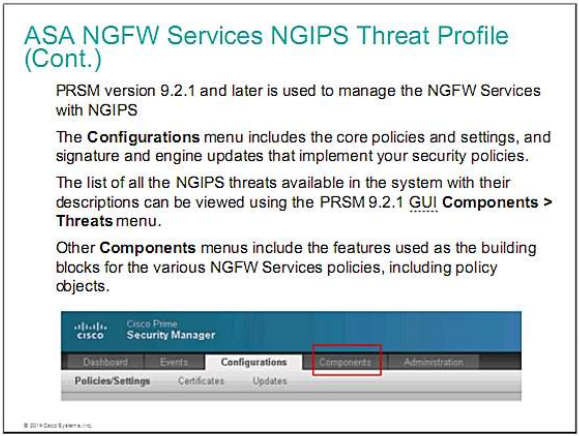

PRSM version 9.2.1 and later is used to manage the NGFW Services with NGIPS. The PRSM 9.2.1 menu structure shown in the figure is different than the PRSM pre-9.2.1 releases. But the basic access, identity and decryption polices configurations are very similar. With the PRSM 9.2.1 and later releases, more ASA features such as ASA ACLs and NAT can also be configured using the PRSM.

The **Configurations** menu includes the core policies and settings, and signature and engine updates that implement your security policies. The Configurations menu options are shown in the figure.

The list of all the NGIPS threats available in the system with their descriptions can be viewed using the PRSM 9.2.1 GUI **Components > Threats** menu (not shown in the figure).

Other **Components** menus include the features used as the building blocks for the various NGFW Services policies, including policy objects and directory realms (not shown in the figure).

PRSM 9.2.1 GUI operations is beyond the scope of this course. Please refer to the latest PRSM User Guide for more information on all the other menu options.

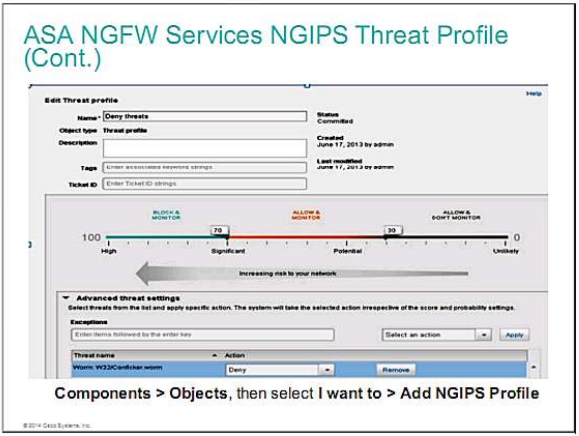

NGIPS Profile is used to apply Next Generation IPS filtering within an access policy which can drop otherwise allowed traffic.

The NGIPS threat profile threat zone sliders are used to define three separate threat zones. The sliders are based on the threat score, with 100 being the most dangerous (on the left) and 0 being least dangerous (on the right). Move the sliders to define the following zones:

Block and Monitor (Deny): Traffic that includes threats in this zone is dropped, and you will see events for the connections in Event Viewer.

Allow and Monitor (Alert): Traffic that includes threats in this zone is allowed, but events are generated for the connections. You can use the dashboards and Event Viewer to monitor these connections.

Allow and Do Not Monitor (Ignore): Traffic that includes threats in this zone is allowed and no events are generated.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Use the **Advanced threat settings** section to define exceptions to the zones defined in the slider. For example, you can identify a threat that should always be blocked regardless of its threat score. Or, you can identify a threat you believe is benign and that you want to ignore.

To add a threat, select it from the Exceptions list, select the **Action** you want for the threat, and click **Apply**. The exception is added to the table. To edit an existing exception, find it in the table and select a different action for that entry. To delete an existing exception, find it in the table and click **Remove** for that entry.

https://t.me/learningnets

Apply NGIPS Threat Profile in NGFW Services Access Policy

This topic describes how to apply an IPS Threat Profile within a NGFW Services Access Policy.

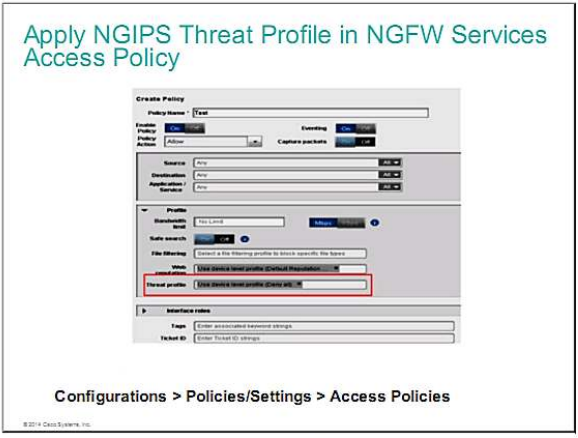

Use the NGIPS threat profile objects to define the IPS filtering policy using these two methods:

Intrusion Prevention global settings: Select a single threat profile that defines your device level default threat profile.

Access policy: Select a profile within an access policy, or specify that the device level threat profile should be used. This defines the NGIPS filtering that you want to apply to traffic that matches the access policy. If an access policy does not specify an NGIPS threat profile, no IPS filtering is applied to the traffic matched by the access policy.

- Different NGIPS threat profiles can be applied to different access policies.
- For each access policy that should be applied a NGIPS threat profile, edit the access policy and select the appropriate NGIPS threat profile.
- Access policies will have the device level threat profile automatically applied by default.

In the figure, the device level IPS threat profile (Deny all) is applied to the Test access policy.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

NGFW Services NGIPS Global Settings

This topic describes the Cisco Next Generation Firewall Services NGIPS Global Settings

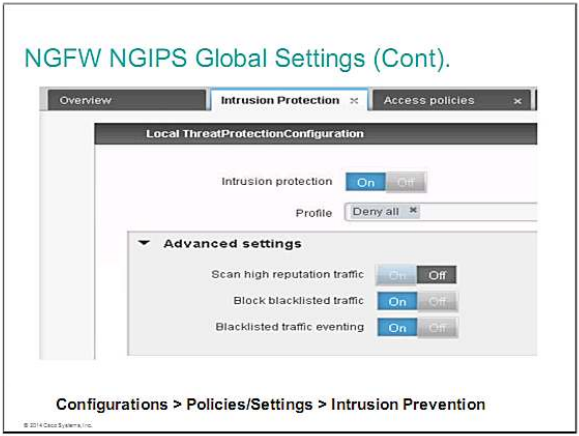

Select **Configurations > Policies/Settings**, open the **Intrusion Prevention** tab to turn NGIPS filtering on or off globally, select a device-level profile, and configure the advanced settings.

The device level threat profile is applied to all traffic matching the access policies configured to use the device level threat profile.

In the example shown in the figure, NGIPS filtering is turned on globally and the device level threat profile is named **Deny all**.

Advanced settings:

Scan High Reputation Traffic: On/Off—Whether to perform Next Generation IPS filtering for traffic flows to high reputation sites, which are defined as low risk, that is, worthy of being bypassed. Next Generation IPS can identify certain types of traffic as being highly unlikely to contain threats, and web reputation can be part of this identification. Be aware that reputation scores are updated regularly with WBRS downloads, and you need to have an active Web Security Essentials license for these downloads to occur; otherwise, the reputation scores used in this evaluation will become dated.

Block Blacklisted Traffic: On/Off—Whether to block the sites on the blacklist downloaded periodically through the engine and signature updates. These sites are consider so threatening that all traffic from them or to them is dangerous. The actual sites on the list change as updates occur. If you enable this option, blacklisted traffic flows are blocked before any other access policies are considered.

Blacklisted Traffic Eventing: On/Off—Whether to generate traffic events when a traffic flow is dropped because it involves a blacklisted site. If you turn this off, blacklisted flows are dropped silently. If you enable eventing, dropped traffic generates Flow Deny events where the policy name is Blacklist (Source) or Blacklist (Destination), depending on whether the blacklist site originated the traffic or was the destination.

https://t.me/learningnets

Viewing NGIPS Threat Information in PRSM Event Viewer

This topic describes an Cisco Next Generation Firewall Services NGIPS triggered event.

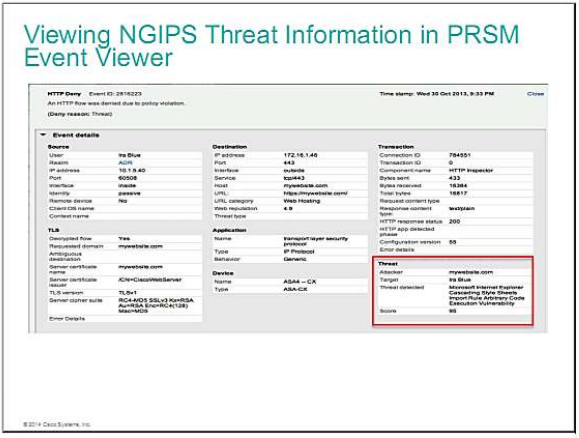

The figure shows a Deny event because a threat with a score of 95 was detected as indicated by the information provided in the **Threat** section of the PRSM event viewer.

In this example, the NGFW services has been configured to perform TLS decryption and the user (the target) was Ira Blue and Ira was accessing the **https://mywebsite.com** URL (the attacker) which triggered the Microsoft Internet Explorer Cascading Style Sheets Import Rule Arbitrary Code Execution Vulnerability threat.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

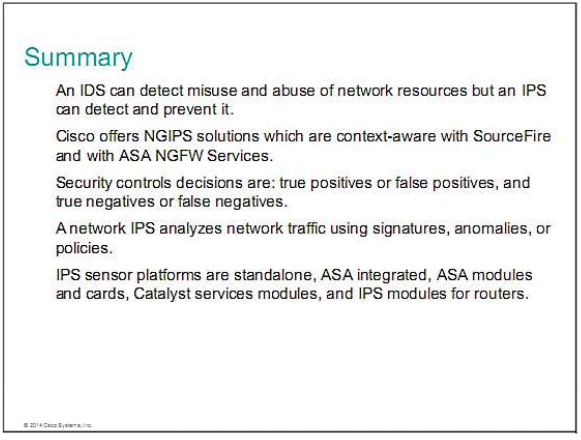

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 2

Integrating Cisco IPS Sensor into a Network

This lesson describes the available Cisco Intrusion Prevention System sensor network integration modes. The lesson also describes how to configure Cisco IPS sensor interface modes using the Cisco IPS Device Manager. IPS sensors can be integrated into the network in several different ways, each of which requires special consideration.

Upon completing this lesson, you will be able to meet these objectives:

- Describe Cisco IPS sensor deployment modes
- Describe the deployment of Cisco IPS sensors in promiscuous mode
- Describe how to deploy sensors in inline interface pair mode
- Describe how to deploy sensors in inline VLAN pair mode
- Describe how to deploy sensors in inline VLAN group mode
- Describe the Cisco IPS Management options
- Describe how to configure interfaces on an IPS sensor
- Describe how to configure promiscuous interfaces
- Describe how to configure inline interface pairs
- Describe how to configure inline VLAN pairs
- Describe how to configure inline VLAN groups
- Describes how to configure traffic redirection for IPS modules

https://t.me/learningnets

Cisco IPS Sensor Deployment Modes Overview

This topic describes Cisco Intrusion Prevention System sensor deployment modes.

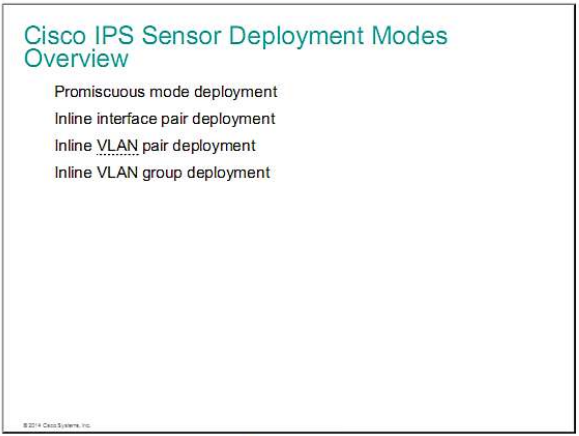

A Cisco IPS sensor can be configured to provide detection or prevention capability to help defend against network attacks. Such sensors can be integrated into the network in several different ways, each of which requires special consideration.

The following are the two most general integration options:

- Promiscuous mode, where the sensor receives a copy of network traffic to analyze while the original traffic flows through the network, bypassing the sensor. In promiscuous mode, the sensor can both detect and prevent attacks, but its preventive options are limited compared to inline mode.
- One of several inline modes where traffic flows through the Cisco IPS sensor itself. In inline mode, the component has many and more reliable aggressive response actions available to it in IPS mode.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Deploy Sensors in Promiscuous Mode

This topic describes the deployment of Cisco Intrusion Prevention System sensors in promiscuous mode.

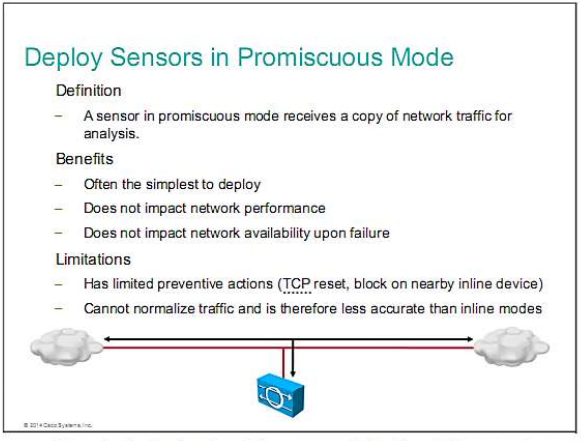

Promiscuous mode is typically the simplest deployment mode for Cisco IPS sensors. In promiscuous mode, packets do not flow through the Cisco IPS sensor. The Cisco IPS sensor receives a copy of network traffic and analyzes copied packets to determine whether they contain signs of suspicious or malicious activity. Although promiscuous mode is simple to deploy, it typically does not impact network performance because the sensor sees only a copy of network traffic—original packets are forwarded unmodified to their original destinations. Moreover, a sensor failure has no impact on network availability because the sensor is out of the packet path.

The main limitation of promiscuous mode is its lack of real-time preventive capability. Although the sensor can disconnect some network sessions (using, for example, generated TCP reset packets) or to dynamically modify access rules on a nearby inline device, such as a router or a firewall, it lacks the capability to drop the original suspicious or malicious packets that are forwarded to their destinations. Additionally, promiscuous mode sensors cannot modify copied traffic in order to normalize it (that is, remove ambiguities or correct suspicious or malicious properties). They therefore are less accurate compared to inline methods, which can use normalization in order to remove ambiguous traffic properties and therefore enhance analysis.

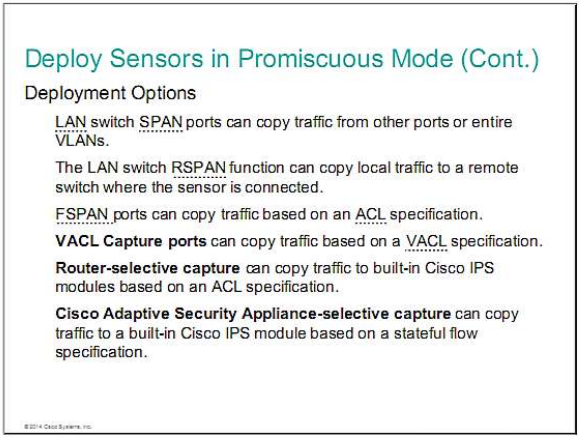

The simplest option for promiscuous mode deployment is to use a hub or a network tap to copy network traffic to a sensing interface of the Cisco IPS sensor. If a hub on a transit or edge segment is used, the Cisco IPS sensor is by design receiving all traffic that is forwarded over the hub. If network taps are used, which are dedicated devices for capturing traffic, you can monitor transit segments in the same manner.

The second promiscuous mode deployment option is to use the SPAN function on LAN switches, which has the advantage of being already present in the infrastructure in most networks. The SPAN function on switches allows the user to flexibly select a subset of traffic (typically per-physical-port or per-VLAN) that is to be copied to a dedicated SPAN port, where the sensing interface of the Cisco IPS sensor is connected.

RSPAN allows LAN switches to forward copied (mirrored) traffic to a VLAN instead of to a destination SPAN port and to use this VLAN (the RSPAN VLAN) to transport copied traffic to another switch, which copies it to its local SPAN destination port. You can use RSPAN to centralize promiscuous mode analysis services and to send captured traffic from remote LAN segments to a central Cisco IPS sensor farm. This method is typically used to analyze low-volume traffic flows using a low number of central sensors. Therefore, RSPAN has the benefit of possibly reducing the number of required sensors.

To increase the granularity of capture on LAN switches, the FSPAN feature allows the creation of ACLs that can granularly select traffic that is copied to the SPAN destination port. FSPAN uses classic stateless ACLs, which can select captured traffic based on various OSI Layer 3 and Layer 4 criteria (IP addresses, protocols, and ports).

Cisco switches use the VACL Capture feature to capture a subset of traffic in a VLAN. This feature is functionally equivalent to the FSPAN feature, and allows the administrator to specify a stateless ACL to limit the scope of capture. By defining a VLAN access map in the Cisco switch configuration, the administrator specifies which intra-VLAN traffic should be copied to the VACL Capture destination port on which the Cisco IPS sensor appliance is connected.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Similar to LAN switches, Cisco routers support built-in Cisco IPS modules, which support promiscuous mode operation. The configuration and functionality of a promiscuous mode deployment on the router is similar to the configuration of FSPAN on switches: The administrator configures a traffic-selecting ACL, and the router then copies traffic that is selected by this ACL from specific interfaces to the built-in Cisco IPS module. This method allows for flexible traffic selection, is available in the infrastructure, and does not affect network performance and availability. The administrator must be careful not to select too much traffic that would oversubscribe the local Cisco IPS module.

On the Cisco ASA IPS modules provide high-performance IPS analysis that is independent of the main Cisco ASA processing. The Cisco ASA can then selectively copy transit traffic to the Cisco IPS module using its MPF CLI. The main difference between the Cisco ASA adaptive security appliance and FSPAN/VACL Capture or Cisco router IPS module promiscuous mode integration is the fact that traffic selection in a Cisco ASA is stateful. The ACL or other criterion specifying traffic to be copied to the Cisco IPS module needs to match only the first packet of a flow, and all subsequent packets of that bidirectional flow are automatically copied to the Cisco IPS module. This practice greatly simplifies capture ACLs and ensures that traffic is always captured bidirectionally to avoid false negatives that are caused by asymmetric capture.

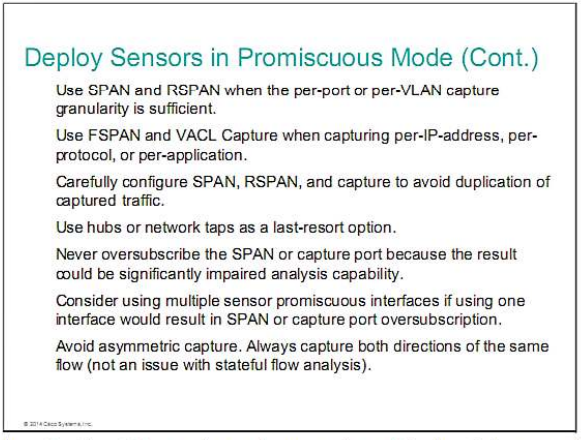

When implementing Cisco IPS sensors in promiscuous mode, consider these deployment guidelines:

Use the SPAN or RSPAN deployment mode when the per-port or per-VLAN capture granularity is sufficient for your purposes and does not oversubscribe the SPAN destination port or the Cisco IPS sensor.

Use FSPAN or VACL Capture when a more granular capture is required, such as per-IP-address, per-protocol, or per-application (per-port) capture. Note that with stateless ACLs governing capture filters, you cannot properly capture traffic of dynamic applications that negotiate their ports during operation unless you capture all traffic to or from an IP address or network.

When using selective capture, carefully specify traffic to be captured to avoid sending duplicate packets to the Cisco IPS sensor.

Use hubs and network taps as an option of last resort. Always prefer capturing traffic using LAN switches, routers, or inside security appliances.

Carefully set the dimensions of the network ports where the Cisco IPS sensor is connected. Always provide enough port bandwidth so that captured traffic is not dropped at the port level, which would cause adverse effects (false negatives) in IPS traffic analysis. If high-speed Cisco IPS sensor interfaces or switch ports are not available, consider using multiple ports on the sensor and switch to capture traffic.

When using selective capture, carefully specify traffic to be captured to avoid capturing network flows partially. Always capture both directions of a particular flow to ensure proper IPS analysis.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Deploy Sensors in Inline Interface Pair Mode

This topic describes how to deploy sensors in inline interface pair mode.

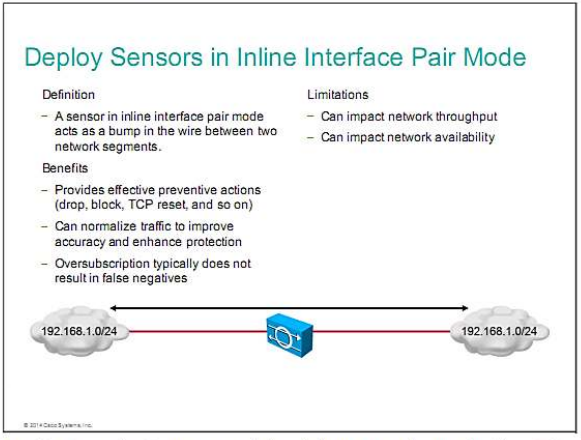

All Cisco Intrusion Prevention System sensor inline deployment modes turn the Cisco IPS sensor into a traffic-forwarding bridging device. The simplest of these modes is the inline interface pair deployment mode, where a sensor uses a pair of network interfaces to interconnect physical or logical (VLAN) networks and acts as a bump in the wire between two network segments. The Cisco IPS sensor inspects all IP (IPv4 and IPv6) traffic that is forwarded over the interface pair and passes non-IP traffic through uninspected. When physical networks are interconnected, such a deployment mode preserves physical separation between networks and cannot be bypassed by LAN switch misconfiguration.

The major benefits of this mode are the following:

Effective preventive responses available with the sensor in traffic-forwarding mode

The dropping of suspicious or malicious packets, blocking a traffic source or destination

The ability to modify traffic in real time (normalization) to remove malicious traffic characteristics and ambiguities

Oversubscription of the sensor does not necessarily lead to false negatives, but can lead to network performance issues

There are two possible limitations:

The impact to network performance if the performance of the sensor does not contain the correct dimensions.

The impact to network reliability upon sensor failure, if path or device redundancy is not provided.

Deploy Sensors in Inline Interface Pair Mode (Cont.)

Deployment Options

- A sensor connecting two physical infrastructures
- A sensor connecting two VLANs

© 2014 Cisco Systems, Inc.

There are two examples of inline interface pair deployment:

A Cisco IPS sensor that is connected to two physically distinct switches using two paired physical interfaces, and the sensor forwarding traffic between two physical parts of the infrastructure. The Cisco IPS sensor forwards all LAN frames between the two switches, and the switch ports connecting to the Cisco IPS sensor can be either in access or trunk mode.

A Cisco IPS sensor that is connected to a single switch using two paired physical interfaces, and the sensor forwarding traffic between logical domains on the same switch. If both interfaces connect to the same switch, they must be configured as access VLANs and must be on different VLANs.

Deploy Sensors in Inline Interface Pair Mode (Cont.)

- Use inline interface pairs whenever you want to inspect all traffic between two network domains.
- Consider a redundant design to avoid single points of failure.
- Consider hardware bypass if Cisco IPS fail-open behavior is acceptable.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

When implementing Cisco IPS sensors in inline interface pair mode, consider these deployment guidelines:

Use the inline interface pair deployment mode when you need to inspect all traffic between two network domains, especially when you need to maintain their physical separation. This mode of deployment is the simplest that can support both access and trunked connections.

Consider using redundant sensors and redundant paths in the network to reroute around possible sensor failures.

If you are unable to provide redundant devices or paths, consider using hardware bypass if the resulting Cisco IPS fail-open behavior is acceptable according to the local security policy.

https://t.me/learningnets

Deploy Sensors in Inline VLAN Pair Mode

This topic describes how to deploy sensors in inline VLAN pair mode.

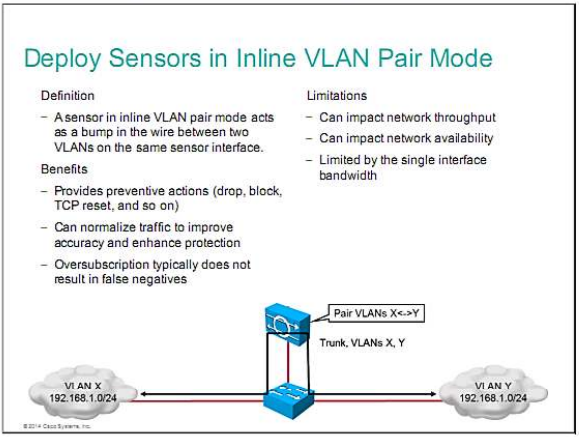

The inline VLAN pair mode is an inline deployment mode where a sensing interface of the Cisco Intrusion Prevention System sensor acts as an 802.1Q trunk port, and the sensor acts as a bump in the wire performing VLAN translation between pairs of VLANs on this trunk interface. Multiple VLAN pairs can be configured on a trunk interface, and each VLAN pair is treated similarly to an inline interface pair. The sensor inspects the traffic that it receives on each VLAN in each pair, and can either forward the packets on the other VLAN in the pair or drop the packet if an intrusion attempt is detected. The sensor replaces the VLAN ID field in the 802.1Q header of each received packet with the identity of the egress VLAN on which the sensor forwards the packet. The sensor drops all packets that are received on any VLANs that are not assigned to inline VLAN pairs.

As with other inline modes, this mode provides preventive actions, can normalize traffic, and can avoid false negatives when oversubscribed. However, this deployment mode can also impact network throughput and reliability, and is vulnerable to bypass if the logical separation (in the adjacent switch) is compromised. Also, because all packets have to cross the trunk interface twice (in and out of the sensor over the same physical interface), traffic can use only up to half the bandwidth made available by the trunk interface.

In the figure, VLANs X and Y use the same Layer 3 IP network addressing (192.168.1.0/24). On the IPS sensor, an inline VLAN pair is configured to connect both VLANs. The IPS inspects traffic between both VLANs.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

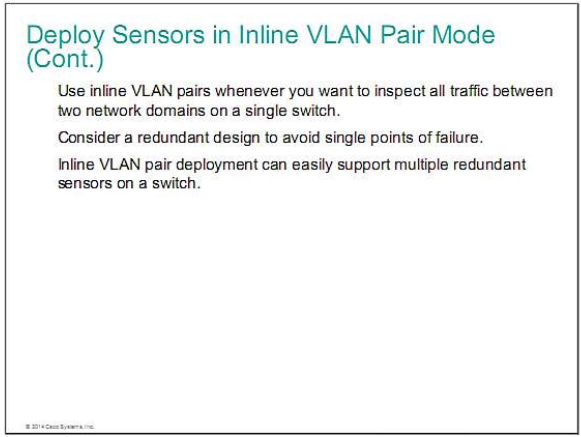

When implementing Cisco IPS sensors in inline VLAN pair mode, consider these deployment guidelines:

- Use the inline VLAN pair deployment mode when you need to inspect all traffic between two network domains using a single LAN switch.
- Consider using redundant sensors and redundant paths in the network to reroute around possible sensor failures. This mode of deployment is especially recommended for redundant and load-sharing designs.
- Inline VLAN pair deployment can easily support multiple redundant sensors on a switch.

Deploy Sensors in Inline VLAN Group Mode

This topic describes how to deploy sensors in inline VLAN group mode.

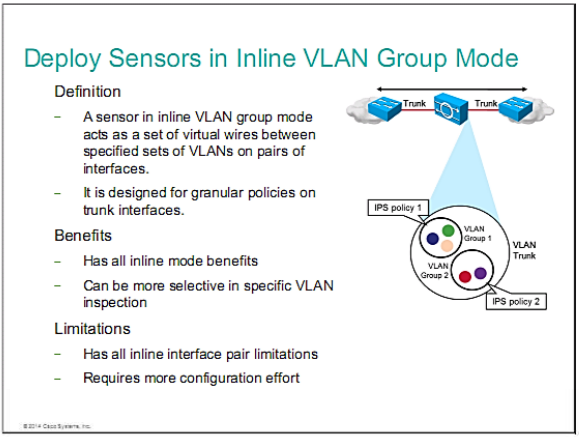

The inline VLAN group mode is an inline deployment mode where a pair of sensing interfaces can be selectively divided (virtualized) into multiple logical “wires” that can be analyzed separately. With the inline VLAN group mode, you can divide each physical interface into VLAN group subinterfaces, each of which consists of a group of VLANs on that interface. Because the Cisco Intrusion Prevention System analysis engine supports multiple virtual sensors, each sensor can monitor traffic on one or more of these VLAN groups (subinterfaces). This practice lets you apply multiple policies to the same sensor. The main advantage is that you can use a sensor with only a few interfaces as if it had many interfaces, and apply granular policies on switched network trunk interfaces.

You can deploy the inline VLAN group mode using Cisco IPS sensors. All the Cisco Catalyst switches feature the necessary trunking capabilities for inline VAN group mode.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

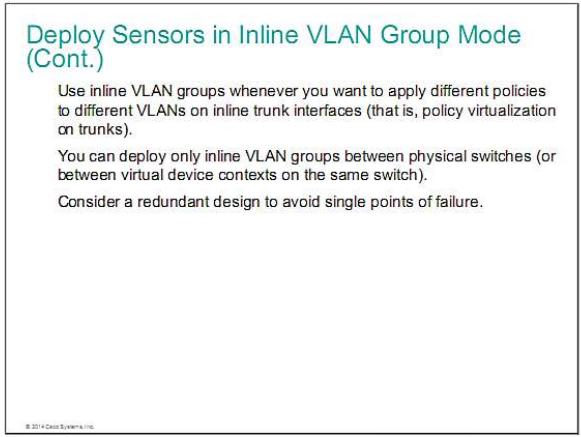

When implementing Cisco IPS sensors in inline VLAN group mode, consider these deployment guidelines:

- Use the Cisco IPS inline VLAN group mode when deploying policy virtualization (that is, different analysis and response policies for different VLANs) on a trunk port.
- Note that you need to deploy the sensor using a physical interface pair that is connected to trunk ports of neighboring switches. You cannot deploy inline VLAN group mode by connecting a Cisco IPS sensor to a single switch using multiple interfaces unless the switch supports virtual device contexts and can therefore support multiple sets of independent VLAN databases.
- As with other inline methods, consider using redundant sensors and redundant paths in the network to reroute around possible sensor failures.

https://t.me/learningnets

Cisco IPS Management Options Overview

This topic describes the Cisco Intrusion Prevention System Management options.

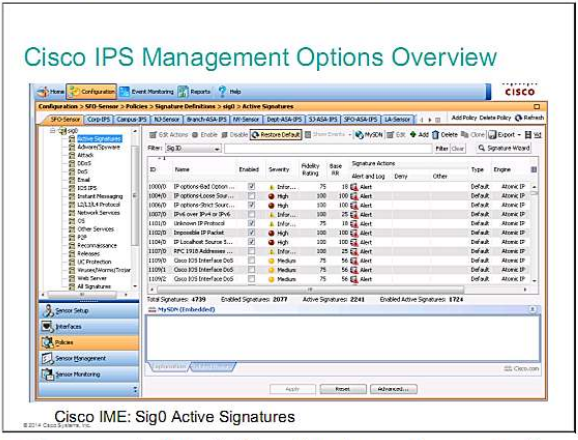

On the Cisco IPS sensor, use the CLI to initially configure the sensor host name, IP address, network ACL to specify networks that are allowed to access the sensor, DNS, NTP, SensorBase Network Participation and etc...

It is recommended to use the IPS GUI management tools, such as the Cisco IPS Device Manager or Cisco IPS Manager Express to configure the Cisco IPS sensor security policy and administer the sensor. You should use the CLI primarily to initialize the sensor and to perform advanced troubleshooting.

This lesson does not cover performing the initial setup using the CLI, please refer to the latest IPS configuration guide for more information.

IPS GUI Management Tools Includes:

Cisco IDM: The Cisco IDM is a web-based Java Web Start application that enables you to configure, administer, and monitor your Cisco IPS sensor. The web server for Cisco IDM application listens on the management interface of the sensor. The Cisco IDM provides element management functionality for a single Cisco IPS sensor. The Cisco IDM GUI was designed to simplify sensor configuration, management, and monitoring tasks. For example, you can use Cisco IDM to easily sort and view all signatures currently stored on the sensor. You can sort by attack type, protocol, service, operating system, action to be performed, engine, signature ID, or signature name. The Cisco IDM also has a Custom Signature Wizard to assist you in creating new signatures. The wizard guides you through the parameters that you must select to configure a custom signature, including selection of the appropriate signature engine. To provide transmission security and authentication, the web server for Cisco IDM secures its HTTP sessions using TLS. Cisco IDM is enabled by default to use TLS. When you enter a URL into your web browser that starts with https://sensor_ip_address, the web browser responds by using the TLS protocol to negotiate an encrypted session with the sensor. Although you can disable the use of TLS, it is highly recommended that you use it because it provides security for communications between the sensor and external systems.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco IME: The Cisco IME is a network management application that provides system health, events, and collaboration monitoring in addition to reporting and configuration for up to ten sensors. IME monitors sensor health using customizable dashboards and provides security alerts through RSS feed integration from the Cisco Security Intelligence Operations. It monitors Global Correlation data, which you can view in events and reports. It monitors events and lets you sort the views. IME also supports tools such, as ping, trace route, DNS lookup, and whois lookup for selected events. It contains a flexible reporting network. It embeds the IDM configuration component to allow for a seamless integration between the monitoring and configuration of IPS devices. Within IME you can set up your sensors, configure policies, monitor IPS events, and generate reports. IME works in single application mode and the entire application is installed on one system and you manage everything from that system.

Cisco Security Manager: The Cisco Security Manager is a comprehensive management solution that enables advanced management and rapid troubleshooting of multiple security devices. Cisco Security Manager provides scalable, centralized management from which administrators can efficiently manage a wide range of Cisco security devices, gain visibility across the network deployment, and securely share information with other essential network services such as compliance systems and advanced security analysis systems. Designed to maximize operational efficiency, Cisco Security Manager also includes a powerful suite of automated capabilities, such as health and performance monitoring, software image management, auto-conflict detection, and integration with ticketing systems. Cisco Security Manager supports a wide range of Cisco security devices, including Cisco ASA 5500 and ASA 5500-X Series ASA Cisco IPS 4200, 4300 and 4500 Series Sensors, Cisco Secure Routers, and the Cisco AnyConnect Secure Mobility Client.

Cisco ASDM: The Cisco Adaptive Security Device Manager **Home Page > Intrusion Prevention tab** can also be used to access the ASA IPS module GUI which is very similar to the Cisco IDM.

– With the IPS 7.1(4)E4 software version, there is a problem specific with Java 7 where the connection to the ASA 5500-X IPS software module will fail when trying to connect to the IPS from within ASDM. This problem has been fixed in the IPS 7.1(7)E4 software version.

The figure shows the Sig0 Active Signatures page of the SFO-Sensor using the Cisco IME. The actual sensor configurations using the Cisco IME or using the Cisco IDM are similar. If you know how to configure a sensor using IDM then you should not have any problems configuring the sensor using IME.

https://t.me/learningnets

Configure Interfaces on IPS Sensor

This topic describes how to configure interfaces on an IPS sensor.

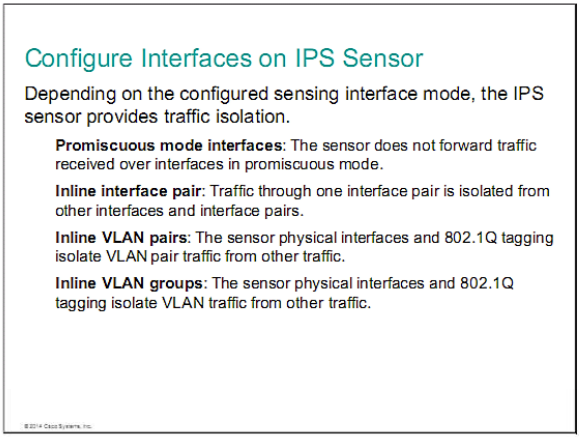

Depending on configured interface mode, the Cisco Intrusion Prevention System sensor provides traffic isolation between different interfaces or interface pairs:

Promiscuous mode: In promiscuous mode, the sensor analyzes a copy of the monitored traffic rather than the actual forwarded packet. The sensor does not forward traffic that is received over interfaces in promiscuous mode.

Inline interface pair mode: Operating in inline interface pair mode puts the Cisco IPS sensor directly into the traffic flow. The sensor receives traffic on one interface and forwards it through the other interface in the interface pair. Traffic through one interface pair is isolated from other interfaces.

Inline VLAN pair mode: You can associate VLANs in pairs on a physical interface that acts as an 802.1Q trunk port, which is known as inline VLAN pair mode. Packets that are received on one of the paired VLANs are analyzed and then forwarded to the other VLAN in the pair. The sensor performs VLAN bridging between pairs of VLANs on the trunk. The physical interfaces of the sensor and 802.1Q tagging isolate VLAN pair traffic from other traffic.

Inline VLAN group mode: You can divide each physical interface or inline interface into VLAN group subinterfaces, each of which consists of a group of VLANs on that interface. The analysis engine supports multiple virtual sensors, each of which can monitor one or more of these interfaces. Physical interfaces of the sensor and 802.1Q tagging isolate VLAN traffic from other traffic.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Promiscuous Interfaces

This topic describes how to configure promiscuous interfaces.

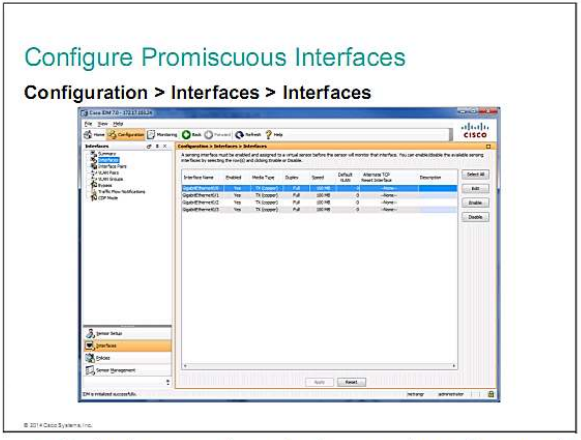

To configure a sensing interface as a promiscuous interface, you need not perform any action in the Interfaces pane. Any interfaces not chosen as members of inline pairs or inline VLANs are automatically considered as promiscuous and can later be assigned to the analysis engine (virtual sensor).

Configure Inline Interface Pairs

This topic describes how to configure inline interface pairs.

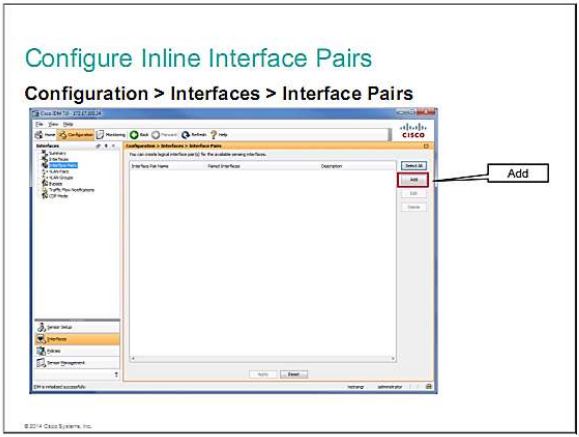

Before you can use your interfaces in inline interface pair operational mode, you must configure an interface pair inside the Cisco Intrusion Prevention System Device Manager. Configure a new inline interface pair by completing these steps:

1. Go to **Configuration > Interfaces > Interface Pairs**. The Interface Pairs panel is displayed.
2. Click **Add** to open the **Add Interface Pair** window and add an interface pair.

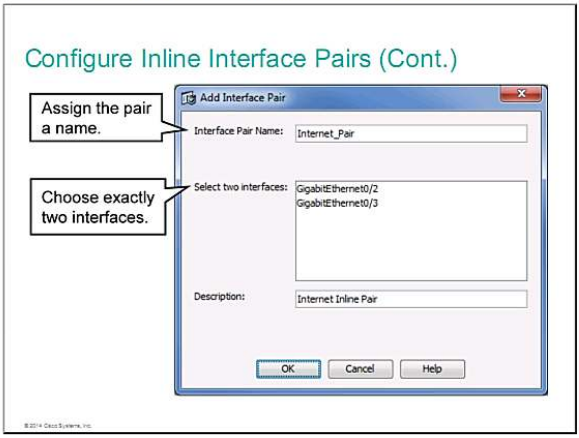

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

3. Enter a name in the **Interface Pair Name** field.
4. Choose the first interface from the **Select Two Interfaces list**, and then hold down the **Shift** key while choosing the second interface of the pair.
5. Optionally, add a description of the interface pair in the **Description** field.
6. Click **OK**.

https://t.me/learningnets

Configure Inline VLAN Pairs

This topic describes how to configure inline VLAN pairs.

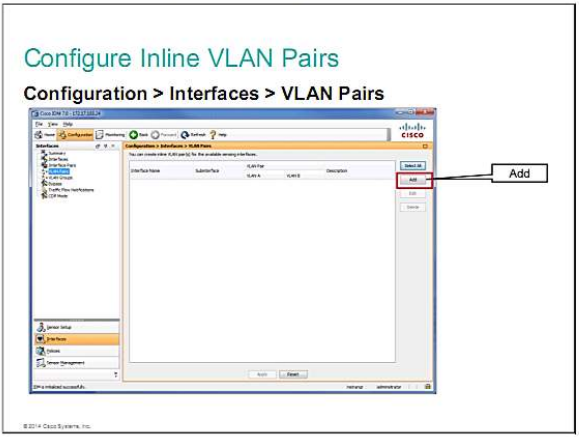

Before you can use your interfaces in inline VLAN pair operational mode, you must configure an interface VLAN pair inside the Cisco Intrusion Prevention System Device Manager. Follow these steps to configure inline VLAN pairs:

1. Go to **Configuration > Interfaces > VLAN Pairs**.
2. Click **Add** to open the **Add Inline VLAN Pair** window and add an inline VLAN pair.

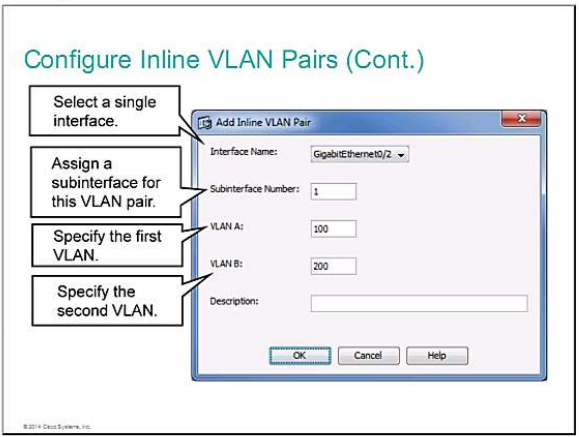

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

3. Choose an interface from the **Interface Name** drop-down list.
4. Enter a subinterface number (1 to 255) for the inline VLAN pair in the Subinterface Number field.
5. Specify the first VLAN (1 to 4095) for this inline VLAN pair in the **VLAN A** field.
6. Specify the other VLAN (1 to 4095) for this inline VLAN pair in the **VLAN B** field.
7. Optionally, add a description of the inline VLAN pair in the **Description** field.
8. Click **OK**.

https://t.me/learningnets

Configure Inline VLAN Groups

This topic describes how to configure inline VLAN groups.

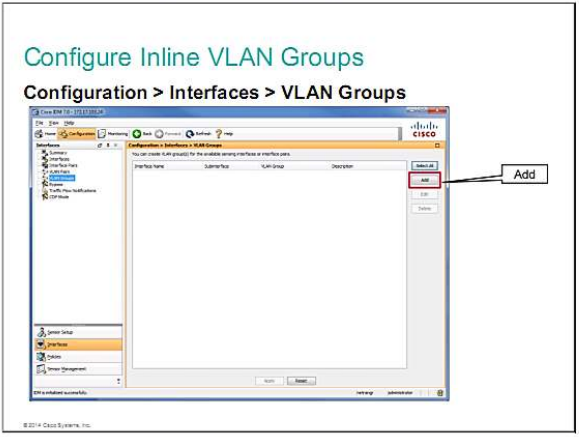

Before you can use your interfaces in inline VLAN group operational mode, you must configure an inline interface pair inside the Cisco Intrusion Prevention System Device Manager and additional interface VLAN groups. Follow these steps to configure VLAN groups:

1. Go to **Configuration > Interfaces > VLAN Groups**.
2. Click **Add** to open the **Add VLAN Group** window and add a VLAN group.

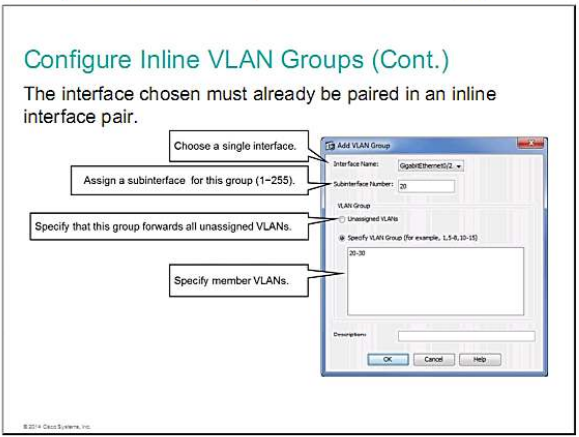

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

3. Choose an interface from the **Interface Name** drop-down list.
4. Enter a subinterface number (1 to 255) for the VLAN group in the **Subinterface Number** field.
5. Under the **VLAN Group** section, specify the VLAN group for this interface by choosing one of the radio buttons:
   - **Unassigned VLANs**: This button lets you assign all of the VLANs that are not already specifically assigned to a subinterface.
   - **Specify VLAN Group (for example, 1, 5–8, 10–15)**: This button lets you specify the VLANs that you want to assign to this subinterface. You can assign more than one VLAN (up to 4096) in this pattern: 1, 5–8, 10–15. This option lets you set up different policies that are based on the VID. For example, you can make VLANs 1 to 10 go to one virtual sensor (VS0) and VLANs 20 to 30 go to another virtual sensor (VS1).
6. Optionally, you can add a description of the VLAN group in the **Description** field.
7. Click **OK**.

The new VLAN group appears in the list in the **VLAN Groups** pane. You must assign this VLAN group to a virtual sensor.

https://t.me/learningnets

Traffic Redirection to IPS Module Overview

This topic describes how to configure traffic redirection to the Cisco Intrusion Prevention System module.

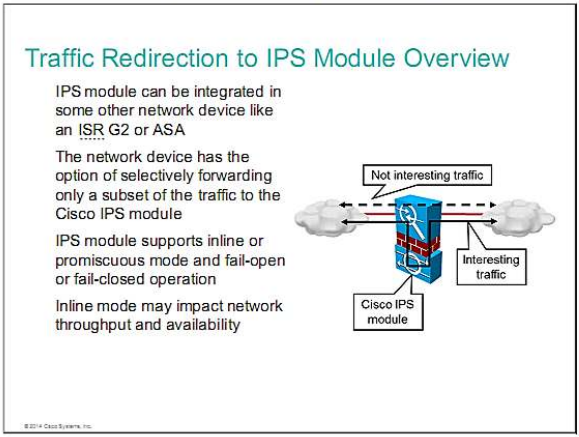

When a Cisco IPS module is integrated in some other network device like Cisco Integrated Services Routers G2 or Cisco Adaptive Security Appliance, the network device has the option of selectively forwarding only a subset of the traffic to the Cisco IPS module. The Cisco IPS module can analyze and forward this traffic in the inline mode if it does not trigger any configured aggressive response options. The IPS administrator can granularly determine which traffic requires inline treatment and can easily engineer the network for an optimal balance of security and performance.

The fail-close option means that if the IPS module fails, then the ISR G2 or Cisco ASA does not let the traffic pass. The fail-open option means if the IPS module fails, the ISR G2 or Cisco ASA lets the traffic pass, but it will not be inspected by the IPS module.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

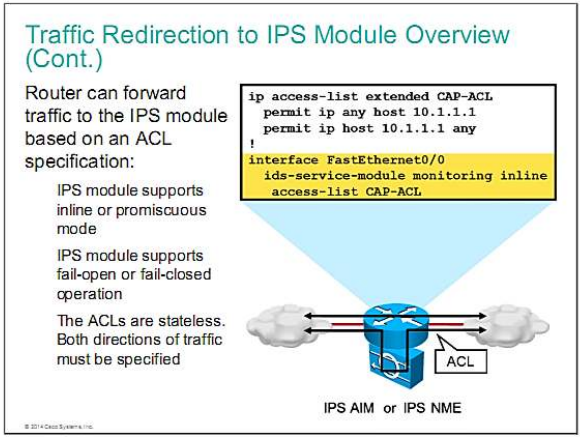

With the Cisco ISR G2 routers, the Cisco IOS Software can selectively forward packets that are selected by a stateless ACL to the Cisco IPS module. Cisco IPS module supports inline or promiscuous mode and fail-open and fail-closed operation. To ensure optimal analysis, the ACL that is used when traffic is selected for analysis should consider both directions of traffic flow. The configuration excerpt in the figure shows an example of the Cisco IOS Software configuration that is needed to redirect a subset of traffic that is selected using an ACL to the Cisco IPS module using the inline mode. Other traffic (traffic that is denied by the ACL) is forwarded without inspection.

The **service-module** CLI command is used to select the fail-close or fail-open operations. The default is fail-open.

```
router(config)# interface ids-sensor 1/0
router(config-if)# service-module {fail-close | fail-open}
```

The Cisco ISR G2 router supports the Cisco IPS NME-IPS module.

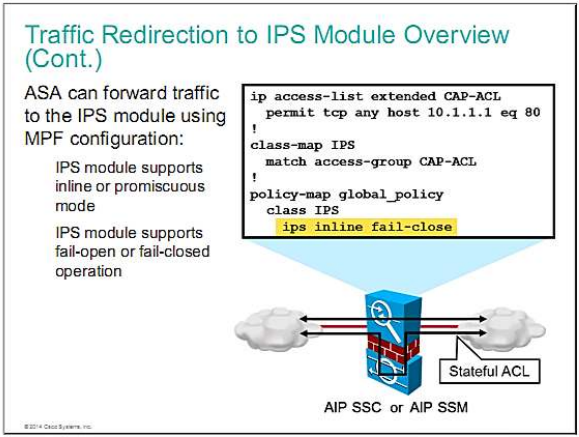

On the Cisco ASA, the MPF configuration is used to forward specific traffic to the Cisco IPS module. The Cisco ASA IPS module supports inline or promiscuous mode and fail-open or fail-closed operation.

The configuration excerpt in the figure shows an example of the Cisco ASA MPF configuration that is needed to statefully redirect a subset of traffic to the Cisco IPS module using inline mode with fail-closed operation. The traffic to be forwarded to the IPS module is selected using a class map matching an ACL. Other traffic not matched by the class map is forwarded through the ASA without IPS inspection.

Within the class map, the ACL or other matching criterion specifying traffic to be forwarded to the Cisco IPS module needs to match only the first packet of a flow, and all subsequent packets of that bidirectional flow are automatically copied to the Cisco IPS module. You can also use other matching criteria besides ACL for selecting traffic in the Cisco ASA class map; for example, IP precedence, DSCP, and TCP or UDP port.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

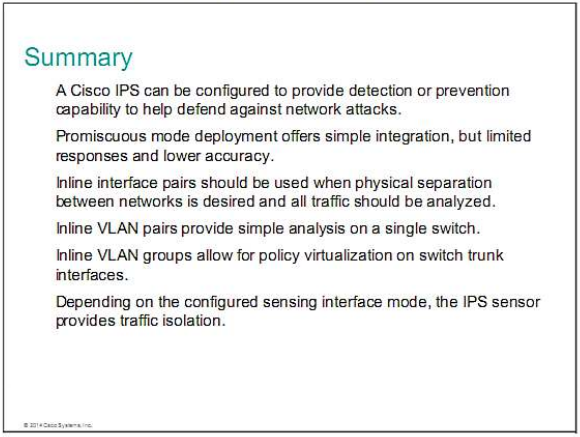

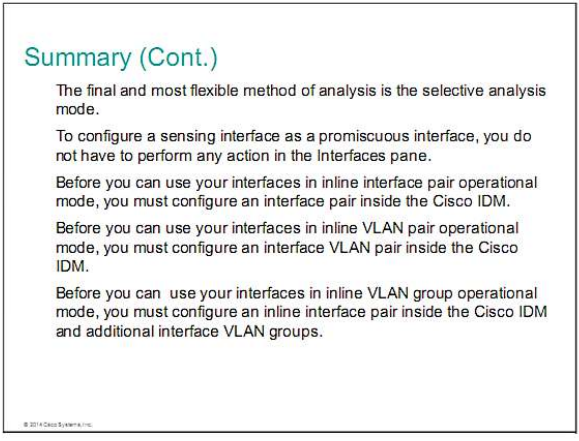

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 3

Configuring Basic Cisco IPS Operations

This lesson introduces the configuration of the built-in signatures in the Cisco Intrusion Prevention System sensor products. You will be able to find individual signatures and classes of signatures, and to perform basic signature-related configuration actions. You will also learn how to configure the actions that you want the sensor to take, and to configure the two configuration mechanisms that allow you to scalably change responses for a large number of signatures.

Upon completing this lesson, you will be able to meet these objectives:

- Describe different types, features, and actions of Cisco IPS signatures
- Describe the signature properties
- Describe the Cisco IPS actions
- Describe how to configure virtual sensor
- Describe the threat rating
- Describe how to configure basic signature properties
- Explain the risk rating
- Describe the inputs for the risk rating calculation
- Describe the threat rating
- Describe event action overrides
- Describe how to configure and verify event action overrides
- Describe event action filters
- Describe how to configure and verify event action filters
- Examine the IPS Events using IDM

https://t.me/learningnets

IPS Signatures

This topic describes different types, features, and actions of Cisco Intrusion Prevention System signatures.

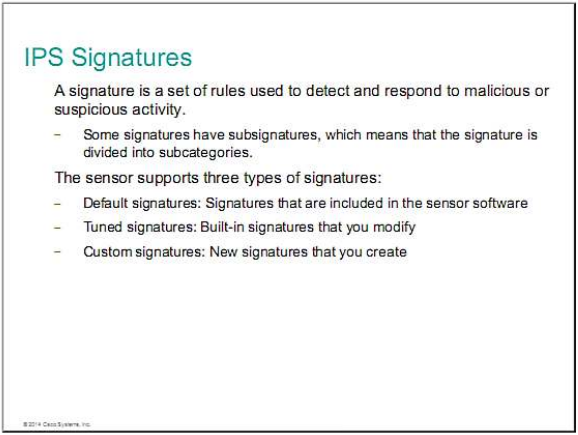

A signature is a set of rules that your sensor uses to detect typical malicious or suspicious network activity such as DoS attacks. Because sensors scan network packets, they use signatures to detect attacks and respond with actions that you define. A signature must be enabled to monitor network traffic. In Cisco IPS, the most critical signatures are enabled in order by default.

Some signatures have subsignatures, which means that the signature is divided into subcategories that typically indicate different variants of the same attack. When you configure a subsignature, changes made to the parameters of one subsignature apply only to that subsignature. For example, if you edit signature 3050 subsignature 1, the change applies only to subsignature 1 and not to signature 3050 subsignature 2, signature 3050 subsignature 3, and so on.

The Cisco IPS sensor distinguishes between three general types of signatures:

Default signatures, which are known attack signatures that are included in the sensor software

Tuned signatures, which are built-in signatures that you modify by adjusting their parameters

Custom signatures, which you create based on criteria that you define

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

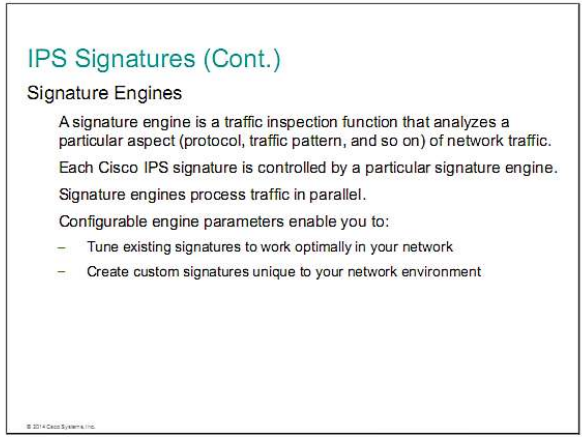

A signature engine is a component of the analysis engine of the sensor that inspects a particular aspect of network traffic and supports a category of signatures. Each Cisco IPS signature is created and controlled by a signature engine that is specifically designed for the type of traffic being monitored. For example, the STRING.TCP engine examines TCP connections searching for string patterns. It controls such signatures as the following:

Signature 3118, rwhoisd format string, which triggers upon detecting an **soa** command that is sent to an rwhois server with a large argument

Signature 3138, Bagle.C virus email attachment, which fires when a pattern matching the C variant of the Bagle virus in an email attachment is detected

Cisco IPS signature engines process traffic in parallel.

An engine is composed of a parser and an inspector. Each engine has a set of legal parameters that have allowable ranges or sets of values. These configurable engine parameters enable you to tune signatures to work optimally in your network and to create new signatures unique to your network environment.

Note Cisco IPS also provides customers with unrivaled industrial control protection. Industrial control protection is a highly specialized branch of network security, and the ramifications of failure can be critical. Cisco IPS provides a set of licensable signatures (over a 1 or 3 year term on a per-device basis) written by experienced industrial control experts to safely inspect and protect these critical network communications. Industrial control systems or ICSs, is the term used to identify several types of control systems, including SCADA systems, PCSs, and other smaller control system types, such as PLCs, used in critical infrastructures and the industrial sector. Cisco IPS industrial control protection signatures require IPS software version 7.0 and later and E4 support.

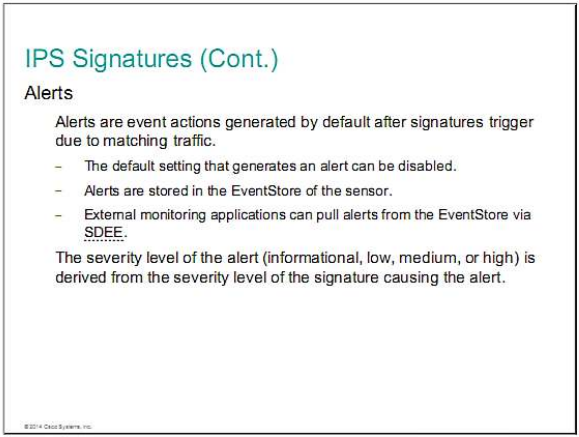

By default, the sensor generates an alert when an enabled signature is triggered. Generating an alert, however, is a configurable signature action that can be disabled. Alerts are stored in the sensor EventStore, which is a fixed-size, indexed store. The Cisco IPS Device Manager or Cisco IPS Manager Express can pull alerts from the sensor via the SDEE. This capability allows a host or hosts to collect alerts on an as-needed basis.

SDEE specifies two types of event requests for external monitoring applications such as Cisco IDM interfacing with the sensor:

Query: Retrieve events that are in the EventStore at the time the query request is issued

Subscription: Establish a live feed

Note Multiple hosts can perform queries and subscribe to the live event feed simultaneously.

Every alert has a severity level that is derived from the severity level of the signature causing the alert. Therefore, an alert, and a signature, can have one of the following security levels:

Informational

Low

Medium

High

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IPS Signature Properties

This topic summarizes the Cisco Intrusion Prevention System signature properties.

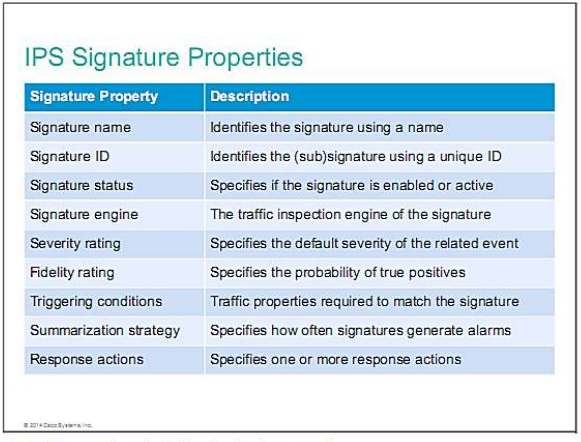

Each Cisco IPS signature has the following basic properties:

The signature name is a user-friendly name of a signature and usually provides a short description of the attack that the signature is designed to detect or prevent.

The signature ID is an integer that uniquely identifies the signature. The Cisco IPS assigns custom signature IDs above 60000. If a signature uses multiple subsignatures, each subsignature is assigned a subsignature ID that is unique in the context of the parent signature.

The signature status describes the signature as enabled or disabled (that is, inspecting traffic or not), and as active or retired (present in the active signature database or not).

The signature engine is the engine on which the signature is based. Signature engines define the Cisco IPS inspection routines that this signature uses to match traffic.

The severity rating of this signature determines the default severity that is assigned to events that cause this signature to trigger. This severity can be adjusted later to reflect other environmental conditions.

The fidelity rating of the signature describes the accuracy with which this signature matches attacks.

The triggering conditions are perhaps the most important part of the signature, and describe network traffic properties that need to be present for this signature to trigger.

The summarization strategy describes how often the signature generates alarms.

The response actions describe the actions that are taken by the sensor when this signature triggers.

IPS Actions

This topic identifies the Cisco Intrusion Prevention System actions that are available in the inline and promiscuous modes.

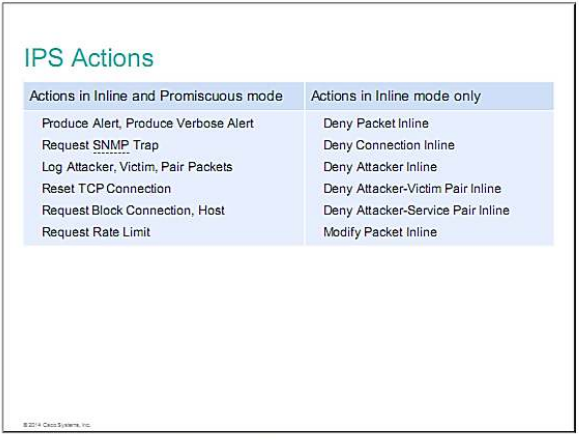

Cisco IPS sensors in inline mode will support all actions. Sensors in promiscuous mode cannot support deny actions because traffic is not flowing through them.

You can try to use the Reset TCP Connection action in promiscuous mode to attempt to block TCP-based attacks in real time. Note, however, that this action is not always reliable, especially with high-packet-rate flows. Instead, you can use the Request Block Host and Request Block Connection actions to prevent some attacks in promiscuous mode.

The Request Block Host, Request Block Connection, and Request Rate Limit actions are always performed on remote (external) network devices, and not on the sensor itself. In inline mode, the Deny Attacker Inline and Deny Connection Inline actions are equivalent to the Request Block Host and Request Block Connection, except that they are performed on the sensor.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure a Virtual Sensor

This topic describes how to configure a virtual sensor.

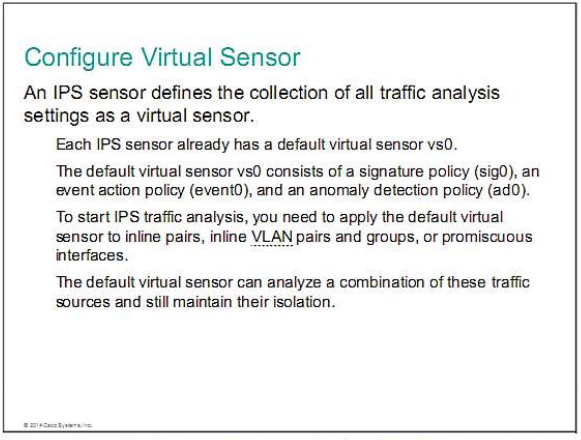

The Cisco Intrusion Prevention System sensor can receive data input from one or many monitored data streams: inline interface pairs, inline VLAN pairs, inline VLAN groups, or promiscuous interfaces. By default, a single sensor policy (configuration) is applied to all monitored data streams. This configuration is called a virtual sensor and includes a signature policy, an event action policy, and an anomaly detection policy that is applied to traffic as a whole.

By default, each Cisco IPS sensor includes a single virtual sensor named vs0, which in turn includes a signature policy named sig0, an event action policy named event0, and an anomaly detection policy named ad0. This virtual sensor cannot be deleted.

To start analyzing traffic from traffic sources, you need to associate these traffic sources with the vs0 virtual sensor. You can either apply a single source (such as an inline interface pair) to the vs0 virtual sensor or you can apply a combination of available sources. Note that the sensor will always maintain traffic isolation between sources; that is, if two inline interface pairs are assigned to the sensor, packets from one interface pair cannot "leak" to the other pair.

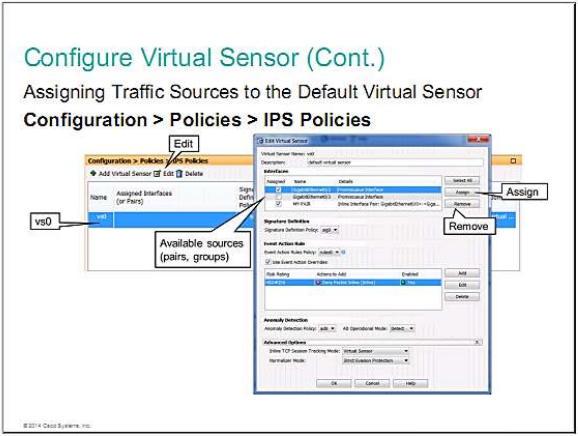

Perform these steps to assign traffic sources to the default vs0 virtual sensor:

1. Navigate to **Configuration > Policies > IPS Policies**. The upper half of the screen displays the list of virtual sensors. By default, only the **vs0** virtual sensor is defined.
2. Click the **vs0** virtual sensor to choose it.
3. To edit the default virtual sensor, click **Edit**. The **Edit Virtual Sensor** window opens.
4. The available interfaces or interface pairs that you can assign to the virtual sensor are displayed. Choose an interface or interface pair from the **Interfaces** list.
5. Click **Assign**. If you want to remove an interface or interface pair from this list, click **Remove**. In this example, the GigabitEthernet0/2 promiscuous interface and the MY-PAIR inline interface pair are assigned to the virtual sensor.
6. (Optionally) Enter a description for the default virtual sensor in the **Description** field.
7. Click **OK**. The **Edit Virtual Sensor** window closes and the **Virtual Sensor** panel displays the interfaces or interface pairs that you have added to the virtual sensor.
8. Click **Apply** to apply your changes.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

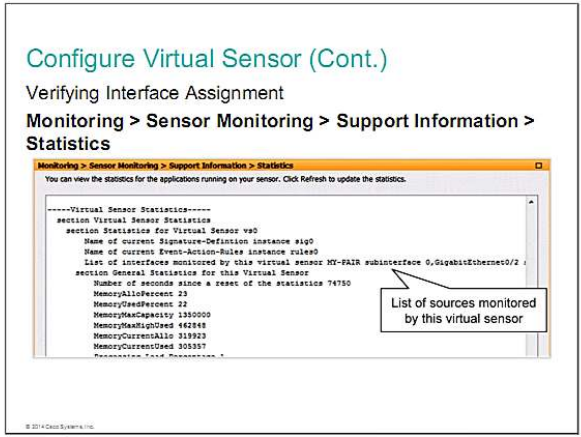

To verify that the interfaces are correctly assigned to the default sensor, perform these steps:

1. Navigate to **Monitoring > Sensor Monitoring > Support Information > Statistics**.
2. Scroll down to the **Virtual Sensor Statistics** section.
3. Under the **Virtual Sensor Statistics** section and **Statistics for Virtual Sensor vs0** section look for the line starting with "List of interfaces monitored by this virtual sensor." This line lists the interfaces that are assigned to the virtual sensor.

IPS Threat Profiles Overview

This topic describes the Cisco Intrusion Prevention System Threat Profiles

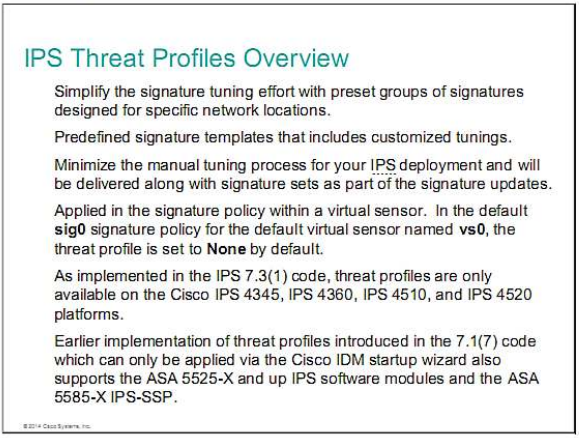

IPS threat profiles provide time savings in tuning and managing the IPS solution. Threat profiles offer guidance of Cisco best practices for different deployments so you can deploy and tune Cisco IPS sensors faster. Using an intuitive GUI interface, these threat profiles simplify the signature tuning effort with preset groups of signatures designed for specific network locations. You can add protection to the data center, Internet edge, Web applications, and others quickly and easily. Threat profiles minimize the manual tuning process for your IPS deployment and will be delivered along with signature sets as part of the signature updates.

An IPS threat profile is a predefined signature template that includes customized tunings. These tunings adjust the signature coverage and response actions to enable the sensor to make better choices in various deployment and threat scenarios.

IPS threat profiles are applied in the signature policy within a virtual sensor. In the default **sig0** signature policy for the default virtual sensor named vs0, the threat profile is set to **None** by default.

The following threat profiles are supported with IPS version 7.1.3:

SCADA - In addition to signatures in the default set, the SCADA signature template includes specialized signatures for general SCADA protocol detections and specific identifiers that address tools and environments common to most device controlled environments. Use this template if the Cisco IPS is primarily used for protecting ICS.

Edge - In addition to signatures in the default set, the Internet Edge signature template includes additional signatures that provide broader protection for desktop operating systems, web browsers, web technologies, and common desktop applications. Use this template if the Cisco IPS is primarily used for securing an Internet connection.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Web Applications - In addition to signatures in the default set, the Web Applications signature template includes additional signatures that provide broader protection for web servers, web development tools and frameworks, content management systems, load balancers, and databases. Use this template if the Cisco IPS is primarily used for protecting web server farms.

Data Center - In addition to signatures in the default set, the Data Center signature template includes additional signatures that provide broader protection for server operating systems, web servers, application servers, databases, content management systems, massaging servers, and virtualization systems. Use this template if the Cisco IPS is primarily used for protecting data centers.

Threat profiles as implemented in the 7.3(1) code are only available on the Cisco IPS 4345, IPS 4360, IPS 4510, and IPS 4520 platforms. However, earlier implementation of threat profiles introduced in the 7.1(7) code which can only be applied via the IDM startup wizard which are less dynamic in nature also supports the Cisco Adaptive Security Appliance 5525-X and up IPS software modules (ASA 5500-X platforms with hardware regex acceleration) and the ASA 5585-X IPS-SSP.

Note	Please refer to the latest IPS configuration guide to see the latest list of platforms supporting threat profiles and more information on using threat profiles.

https://t.me/learningnets

Configure Basic Signature Properties

This topic describes how to configure basic signature properties.

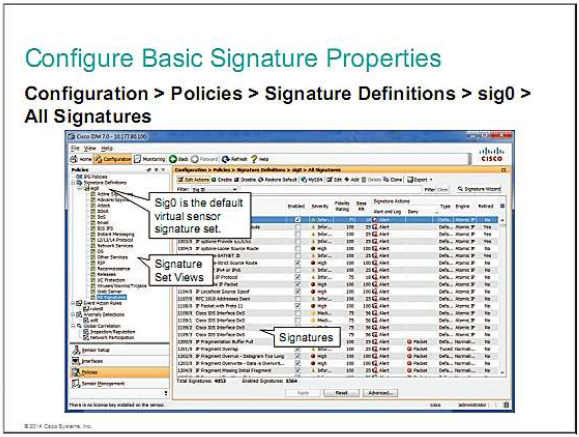

You can access signatures of interest in various ways from the Cisco Intrusion Prevention System Device Manager. To begin, click **Configuration**, choose **Policies > Signature Definitions > sig0**, and then click **All Signatures** to access the **Signature Configuration** panel. By default, the **Signature Configuration** panel displays signatures that are listed by signature ID number. The **All Signatures** database view displays all signatures available in the sensor signature set.

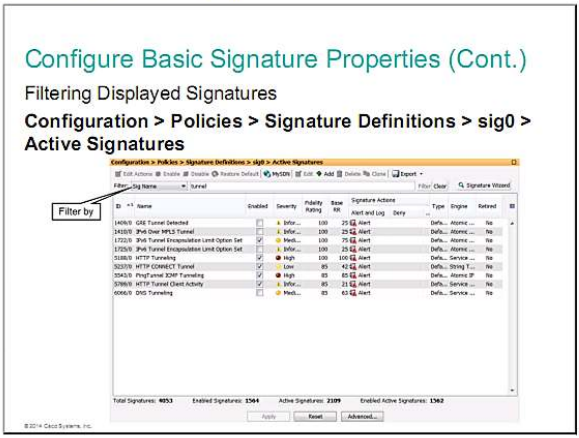

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

You can use the **Filter** drop-down list to display signatures in different ways, such as the types of attack they detect or the services that they inspect. When you change your selection in the **Filter** drop-down list, the **Select Criteria** menu changes to correspond to your selection.

Configure Basic Signature Properties (Cont.)

Enabling and Disabling Signatures

Configuration > Policies > Signature Definitions > sig0 > Active Signatures

Enabling or disabling a signature makes a signature inspect or ignore traffic. Note that the signature must be in the active signature set (therefore, it must not be retired) in order to be enabled.

Configure Basic Signature Properties (Cont.)

Retiring and Activating Signatures

Configuration > Policies > Signature Definitions > sig0 > All Signatures > Edit

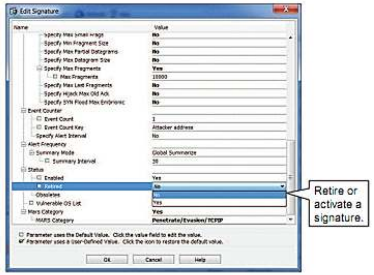

Retiring a signature removes a signature from the set of currently active signatures that are loaded into sensor memory. Retiring signatures conserves sensor memory and enhances sensor performance.

You can later activate retired signatures. When you activate a signature, the signature is added to the set of currently active signatures. Doing so requires the sensor to rebuild its configurations.

Note	Activating and retiring signatures can take 30 minutes or longer.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Risk Rating

This topic describes the risk rating.

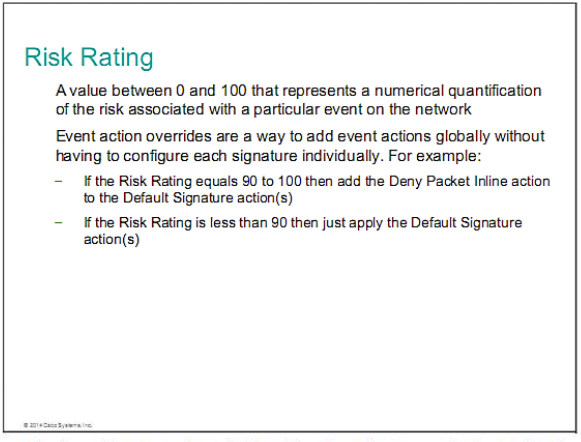

Depending on the size of the enterprise and where Cisco Intrusion Prevention System is placed, your IPS may generate thousands to millions of alerts every day. Analyzing all these alerts can be a daunting task. Cisco IPS risk rating and threat rating are two powerful features that greatly simplify IPS policy management and day-to-day operations. Risk rating is a quantitative measure of your network's threat level before IPS mitigation.

The risk rating provides the administrator with an indication of the relative severity of the event in a particular network environment. The risk rating is an integer value, with values between 0 and 100, that the sensor dynamically calculates after each triggering of a signature. The higher the value, the greater the security risk of the trigger event for the associated alert. Risk rating can help enhance your productivity as it intelligently assesses the level of risk of each event and helps you focus on high-risk events.

The risk rating is a composite value that depends on the proper configuration of some sensor properties, and collection of environmental variables. Therefore, it is imperative to properly configure some sensor aspects to obtain meaningful risk rating values.

Risk rating allows you to easily build powerful policies in the IPS sensor such as configuring an event action override rule to add the Deny Packet Inline action to all events with a risk rating between 90 to 100.

Event action override configurations will be covered in a later topic in this lesson.

Risk Rating Calculation

This topic summarizes the inputs that are used for the risk rating calculation.

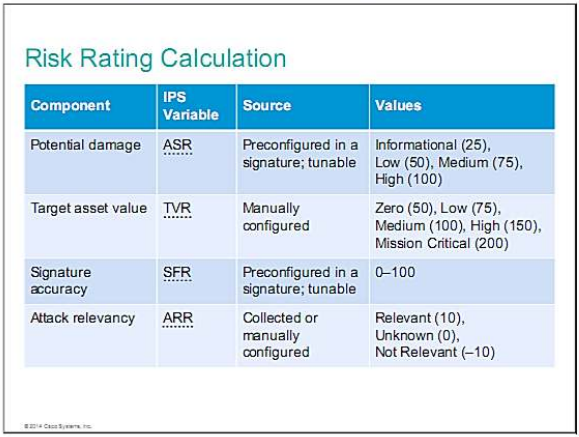

Cisco Intrusion Prevention System uses all of the inputs as shown in the figure in its risk rating calculation. Some of these inputs must be manually configured, some of them are preconfigured and can be tuned, and some are collected automatically from the environment.

Attack Severity

Cisco IPS assigns an ASR to each signature. The ASR is determined by the severity level that you configured for the signature. The severity level can be informational, low, medium, or high.

Each of these severities has an associated numeric value that the risk rating formula uses for the ASR value.

- Informational (25)
- Low (50)
- Medium (75)
- High (100)

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Target Value

With Cisco IPS, you can assign a TVR to each host or network in your environment. The TVR is a user-configurable value that identifies the importance of a network asset, through its IP address. You can develop a security policy that is more stringent for valuable corporate resources and looser for less important resources. For example, you could assign a TVR to the company web server that is higher than the TVR that you assign to a desktop node. In this example, attacks against the company web server have a higher risk rating than attacks against the desktop node.

The following are the current numeric values for the configured targets:

Zero (50)

Low (75)

Medium (100)

High (150)

Mission Critical (200)

By default, the Cisco IPS considers all assets (IP addresses) as having the medium TVR.

Signature Fidelity

Cisco IPS assigns a SFR to each signature. SFRs are tunable by IPS administrators on a per-signature basis, but typically need not be tuned because they are preset to sensible values. The SFR is an indication of the confidence that the signature developer has in the signature accuracy. It is not an indication of the seriousness of the potential attack. SFR is a weight that is associated with how well this signature might perform in the absence of specific knowledge of the target. The SFR has a value between 0 and 100.

Attack Relevancy

Cisco IPS computes the ARR for each signature-triggering event. Each Cisco IPS signature includes a list of vulnerable platforms, which the Cisco IPS compares to the target platform for each detected attack. The target platform is determined either through manual configuration, or through the passive operating system fingerprinting feature of the Cisco IPS sensor.

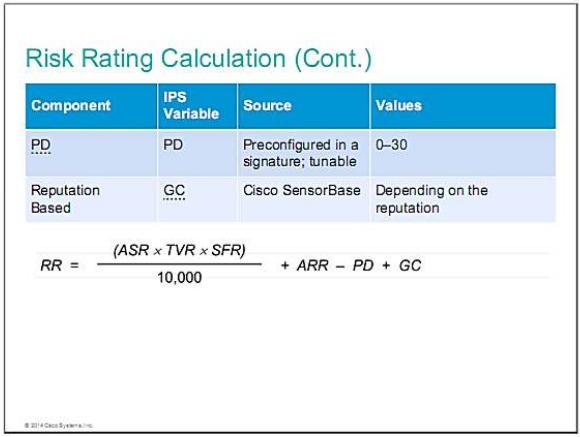

Promiscuous Delta

The risk rating of an IPS deployed in promiscuous mode is reduced by the PD. This is because promiscuous sensing is less accurate than inline sensing. The promiscuous delta can be configured on a per-signature basis, with a value range of 0 to 30.

Global Correlation

The Global Correlation feature uses network reputation scores in two different but complementary ways. First, the reputation of the source of a new flow is tested and the flow is denied without further processing if the reputation is bad. Second, the flow is passed through the IPS inspection engines. These engines determine the threat potential of the flow based on the sensor's policy configuration, and assign a risk rating to the flow. The risk rating is then increased to take into account the reputation of the flow's source. The actual formula that is used to calculate risk rating increase is complex, relies on statistical information, and may change with time. If the resultant risk rating is above a threshold, the flow is denied (or an alternate action is taken, depending on the policy configuration).

The figure shows the formula for calculating the risk rating.

One additional risk rating calculation component that is not shown is the WLR. WLR is a value generated based on data found in the Cisco Security Agent watch list. The Cisco Security Agent watch list contains IP addresses of devices involved in network scans or possibly contaminated by viruses or worms. If an attacker is found on the watch list, the watch list rating for that attacker is added to the risk rating. The Cisco Security Agent product has been end of life.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Threat Rating

This topic describes the threat rating.

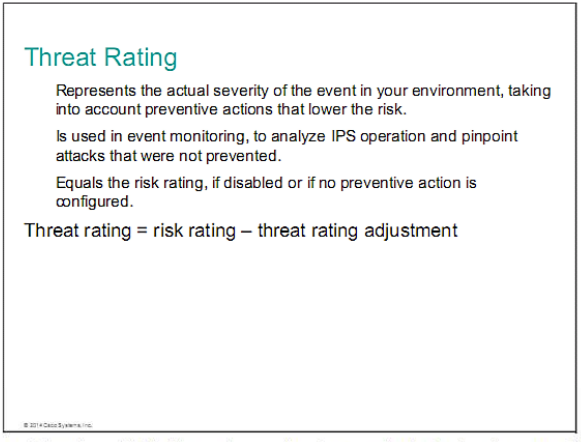

In addition to risk rating, which indicates the severity of an event in the local environment, Cisco Intrusion Prevention System supports the concept of a threat rating. The threat rating indicates the severity of an event in the local environment by also taking into account any preventive actions that the sensor took to reduce the risk for this event.

You should use the threat rating in your event monitoring practices to determine the level of residual risk after the sensor has performed its preventive functions, if any. If you do not configure the sensor with any preventive functions, the threat rating always equals the risk rating.

The threat rating uses threat rating adjustments to lower the risk rating. Every preventive (aggressive) action that is assigned to a signature uses a specific threat rating adjustment that is hard-coded in the Cisco IPS software. If you enable multiple preventive actions for a signature, the risk rating is reduced by the largest threat rating adjustment of the configured actions.

https://t.me/learningnets

Event Action Overrides Overview

This topic describes event action overrides.

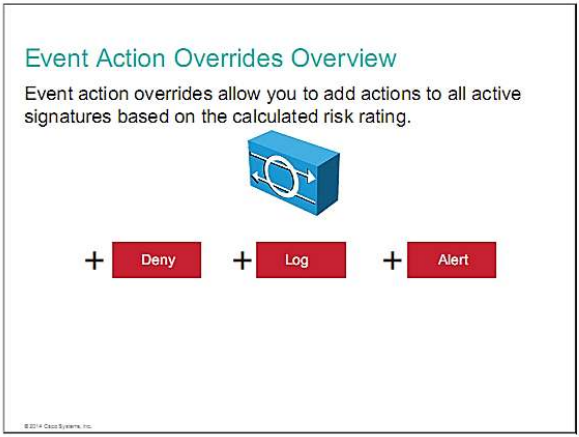

Event action overrides are a configuration tool that allows to you add one or more actions to all active signatures based on the calculated risk rating of each event.

Event action overrides are a way to add event actions globally without the need to configure each signature individually. Each event action has an associated risk rating range. If a signature event occurs and the risk rating for that event falls within the range for an event action, that action is added to the event. For example, if you want any event with a risk rating of 85 or more to generate an SNMP trap, you can set the risk rating range for Request SNMP Trap to 85–100. If you do not want to use action overrides, you can disable the entire event action override component.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

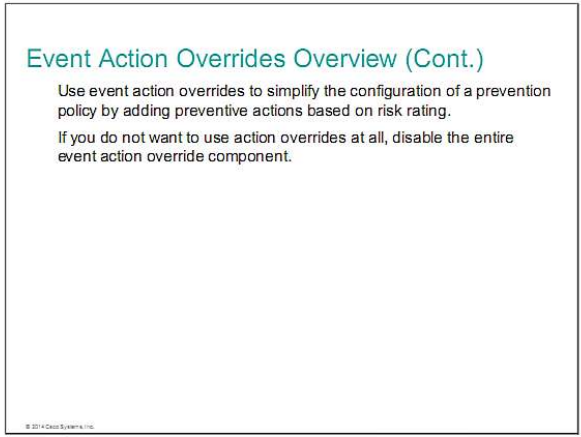

Consider these guidelines when you implement event action overrides on your Cisco Intrusion Prevention System sensor:

- Use event action overrides to simplify the creation of either detective or preventive actions for all signatures based on their risk rating. With this tool you need not configure each signature individually.
- If you do not want to use action overrides, you can disable the entire event action override component.

https://t.me/learningnets

Configure and Verify Event Action Overrides

This topic describes how to configure event action overrides.

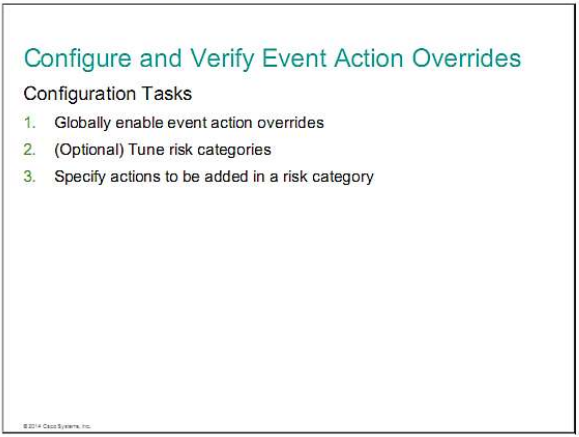

To configure event action overrides on a sensor, perform the following configuration tasks:

1. Globally enable event action overrides. Event action overrides are enabled by default, therefore this task is optional.
2. Optionally, tune risk categories to either change existing risk categories or to introduce your own risk categories that are based on risk rating ranges.
3. Create event action override rules to add actions to events based on a risk category.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

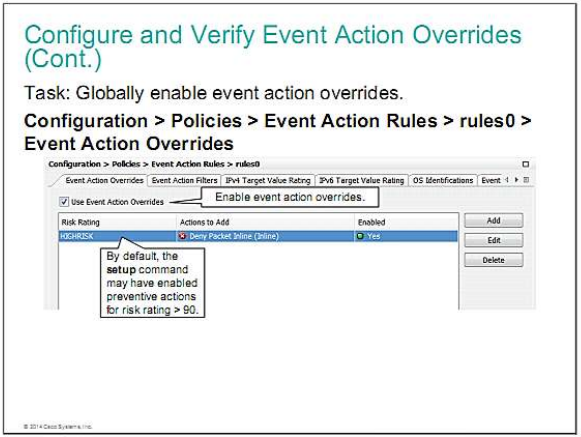

In this task you can globally enable event action overrides if they were previously disabled by your configuration. Perform the following steps to enable event action overrides:

1. Click **Configuration** and choose **Policies > Event Action Rules**.
2. Click **rules0**.
3. From the **rules0** panel, click the **Event Action Overrides** tab. The **Event Action Overrides** panel is displayed.
4. Check the **Use Event Action Overrides** check box.

By default, the **setup** command may have enabled preventive actions for events with a risk rating greater than 90. In such a case, the corresponding event action override is displayed.

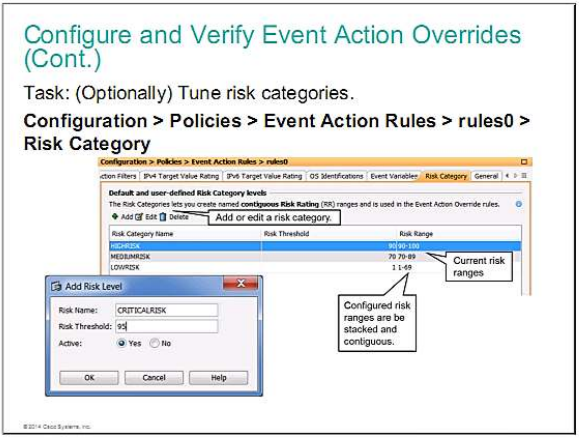

Risk categories link a category name to a numeric range defining the risk rating. You can use predefined risk categories (HIGHRISK, MEDIUMRISK, and LOWRISK) or you can define your own risk categories with custom ranges. You specify only the low threshold for the category to ensure that the ranges are contiguous. The upper threshold for a category is either the next higher configured category or 100 if there is no category with a higher threshold.

To add, edit, or delete a risk category, perform these steps:

1. Click **Configuration**, choose **Policies > Event Action Rules > rules0 > Risk Category**, and click **Add**.
2. In the **Risk Name** field, enter a name for this risk category.
3. In the **Risk Threshold** field, enter a numerical value for the risk threshold (minimum 0, maximum 100). This number represents the lower boundary of risk. The range appears in the **Risk Range** field and in the red, yellow, and green threshold fields.
4. To make this risk category active, click the **Yes** radio button.
5. Click **OK**. The new risk category appears in the list on the **Risk Category** tab.
6. To edit an existing risk category, choose it in the list and click **Edit**. Make any changes needed.
7. Click **OK**. The edited risk category appears in the list on the **Risk Category** tab.
8. To delete a risk category, choose it in the list and click **Delete**. The risk category no longer appears in the list on the **Risk Category** tab.
9. Click **Apply** to apply your changes and save the revised configuration.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure and Verify Event Action Overrides (Cont.)

Task: Specify actions to be added in a risk category.

Configuration > Policies > Event Action Rules > rules0 > Event Action Overrides

To add new event action overrides and therefore specify actions to be added for a specific risk category, perform the following steps:

1. Click **Configuration** and choose **Policies > Event Action Rules**.
2. Click **rules0**.
3. From the rules0 panel, click the **Event Action Overrides** tab.
4. Check the **Use Event Action Overrides** check box.
5. Click **Add** to create a new event action override. The **Add Event Action Override** window opens. You can create multiple event action overrides.
6. From the **Risk Rating** drop-down list, assign a risk rating range to this network asset.
7. From the **Available Actions to Add** list, check the event actions that this event action override will correspond to.
8. Check the **Enabled** check boxes for the actions you want to enable in the override.
9. Click **OK**. The new event action override now appears in the list on the **Event Action Overrides** tab.
10. Check the **Use Event Action Overrides** check box.
11. To edit an existing event action override, choose it in the list and click **Edit**. Make any changes needed.
12. Click **OK**. The edited event action override now appears in the list on the **Event Action Overrides** tab.
13. Check the **Use Event Action Overrides** check box.
14. To delete an event action override, choose it in the list and click **Delete**. The event action override no longer appears in the list in the **Event Action Overrides** tab. You cannot delete the event action override for **Deny Packet Inline** because it is protected. If you do not want to use that override, disable it.
15. To enable or disable an event action override, choose it in the list and click **Edit**.

https://t.me/learningnets

16. To disable an event action override, clear the **Enabled** check boxes for any event actions that you have assigned to that event action override. To enable an event action override, check the **Enabled** check box for any event actions that you have assigned to that event action override.
17. Click **Apply** to apply your changes and save the revised configuration.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Event Action Filters Overview

This topic describes event action filters.

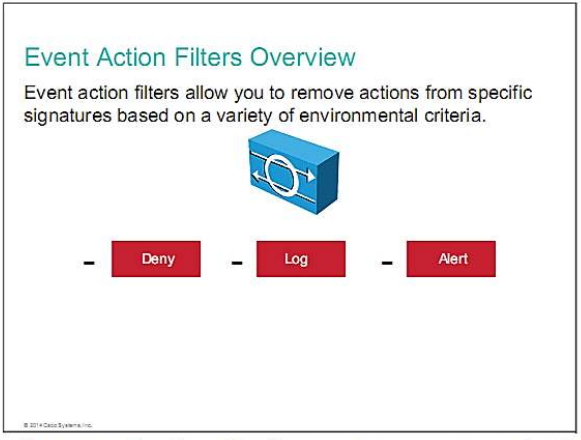

Event action filters are a configuration tool that allows to you remove one or more actions from all active signatures based on various criteria.

You will typically use event action filters to tune your Cisco Intrusion Prevention System sensor when you manage false positive events. When you tune, you will typically remove one, several, or all actions from a particular signature in a particular context (for example, when triggered by a specific source IP address).

Configure and Verify Event Action Filters

This topic describes how to configure event action filters.

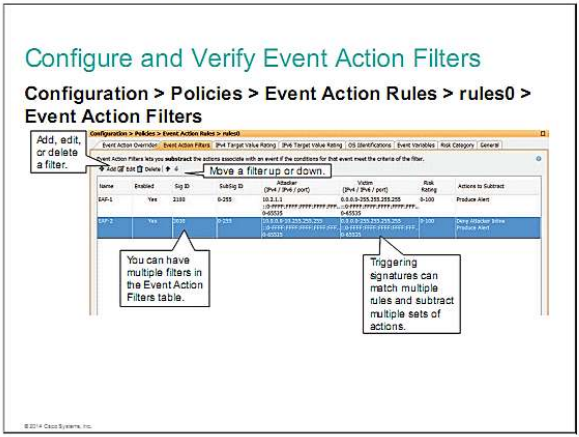

Event action filters are processed as an ordered list. You can move filters up or down in the list. Filters let the sensor perform certain actions in response to the event without requiring the sensor to perform all actions or remove the entire event. Filters work by removing actions from an event. A filter that removes all actions from an event effectively consumes the event.

Use the **Event Action Filters** panel to add and manage event action filters. Choose an event action filter and then click the following buttons to perform the corresponding tasks:

Move Up: This button moves the chosen event action filter up one row in the list. This action results in a change in the processing order of the filters.

Move Down: This button moves the chosen filter down one row in the list. This action results in a change in the processing order of the filters.

Edit: This button opens the **Edit Event Action Filter** window. This window enables you to modify the filter values.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

ConfigConfigure and Verify Event Action Filters (Cont.)

Configuring an Event Action Filter

Configuration > Policies > Event Action Rules > rules0 > Event Action Filters

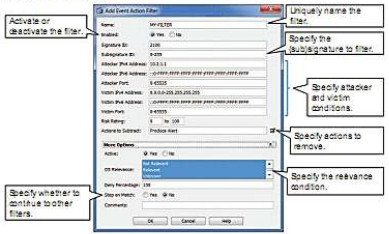

Perform the following steps to create a new event action filter:

Click **Configuration** and choose **Policies > Event Action Rules**.

Click **rules0**.

From the **rules0** panel, click the **Event Action Filters** tab.

Click **Add**. The **Add Event Action Filter** window opens.

The following fields are available to define the Event Action Filters:

Name—Lets you name the filter you are adding.

Enabled—Indicates whether or not this filter is enabled.

Sig ID—Identifies the unique numerical value assigned to this signature. This value lets the sensor identify a particular signature. You can also enter a range of signatures.

SubSig ID—Identifies the unique numerical value assigned to this subsignature. The subSig ID identifies a more granular version of a broad signature. You can also enter a range of subSig IDs.

Attacker (address/port)—Identifies the IP address and/or port of the host that sent the offending packet. You can also enter a range of addresses.

Victim (address/port)—Identifies the IP address and/or port used by the attacker host. This is the port from where the offending packet originated. You can also enter a range of ports.

Risk Rating—Indicates the risk rating range between 0 and 100 that should be used to trigger this event action filter. If an event occurs with a risk rating that falls within the minimum-maximum range you configure here, the event is processed against the rules of this event filter.

Actions to Subtract—Indicates the actions that should be removed from the event, should the conditions of the event meet the criteria of the event action filter.

https://t.me/learningnets

Examine IPS Events using IDM

This topic describes how to examine the Cisco Intrusion Prevention System Events using Cisco IPS Device Manager.

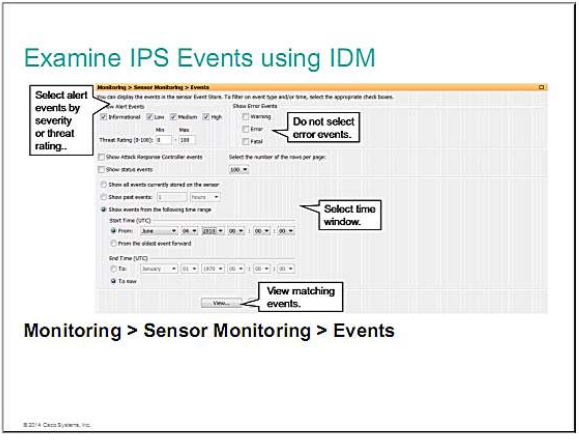

The Cisco IDM Events pane lets you filter and view the IPS sensor event data. You can filter events based on type, time, or both.

1. Go to **Monitoring > Sensor Monitoring** and choose **Events.**
2. Complete the filter choices in the Events pane. You can choose to view events based on criteria such as event severity, types of events, number of events per page, and time consideration.
3. The **Show Error Events** check boxes let you configure the type of errors (Warning, Error, Fatal) you want to view. Do not select any error events if you only want to see the IPS alerts.
4. Click the **View** button.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

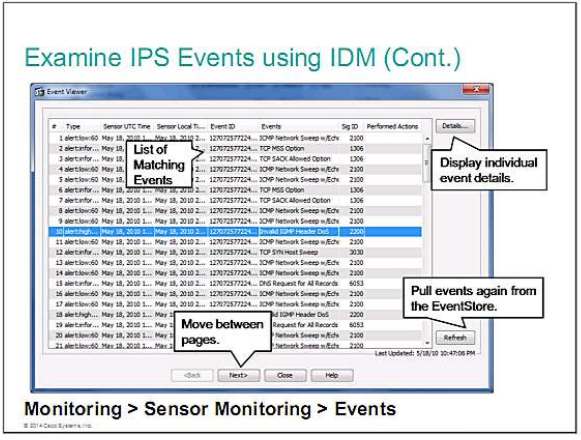

After event property selection, an **Event Viewer** window will open displaying all events in the selected time frame, with the selected event properties. You can browse through the events, and view the details of a specific event.

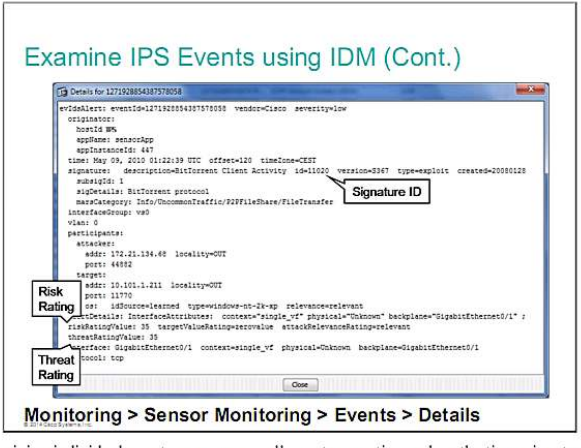

When examining individual events, you can see all event properties such as the time, signature description and signature ID and etc.. At the bottom, you can observe the event risk rating and threat rating. In this example, because the signature associated with this event did not include any preventive actions, the threat rating of the event is the same as its risk rating.

Summary

This topic summarizes the key points that were discussed in this lesson.

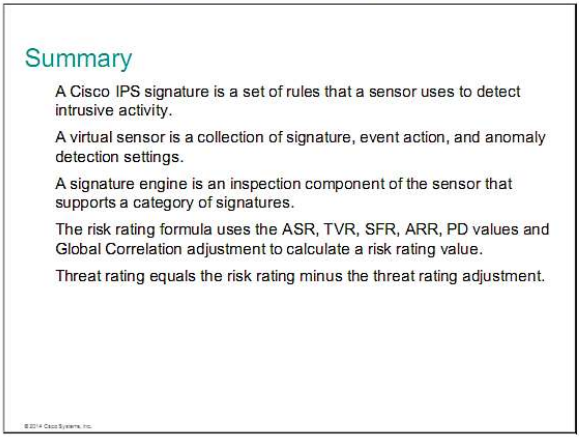

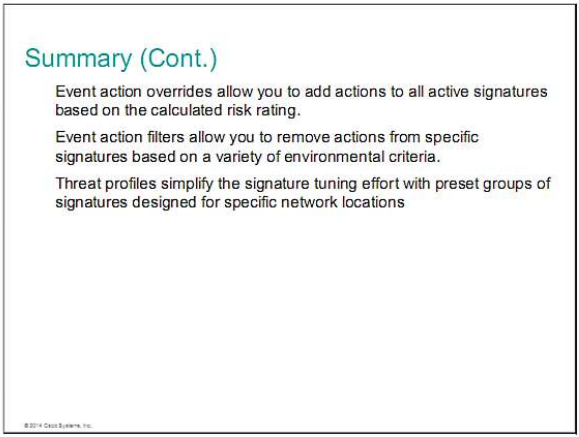

https://t.me/learningnets

Lesson 4

Tuning Cisco IPS Signatures

Overview

Network intrusion prevention technologies are much more effective when they are customized for the environment in which they operate. This lesson examines some methodologies for tuning a Cisco Intrusion Prevention System sensor to properly manage false positive and negative events.

Upon completing this lesson, you will be able to meett hese objectives:

- Describe false negative events
- Describe false positive events
- Describe false positives
- Describe Cisco IPS tuning approaches
- Tune Cisco IPS to reduce false positives
- Reduce false positives by narrowing the search context
- Reduce false positives by narrowing the header values
- Reduce false positives by limiting the number of matched patterns
- Describe an example of reducing false positives by limiting the number of matched patterns
- Reduce false positives by decreasing the attention span
- Reduce false positives by increasing the number of events
- Tune Cisco IPS to reduce false negatives
- Reduce false negatives using IP reassembly
- Reduce false negatives using TCP reassembly
- Reduce false negatives using deobfuscation

https://t.me/learningnets

False Negatives

This topic describes false negative events.

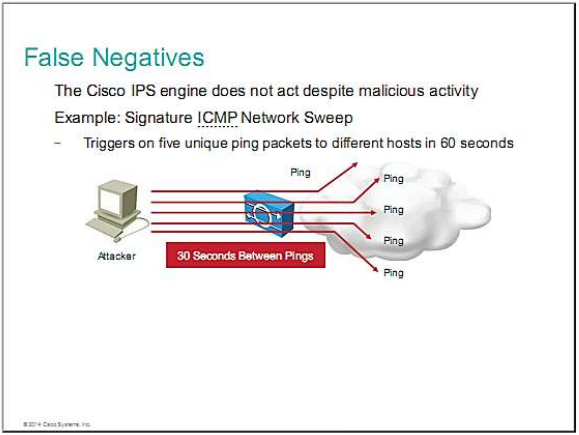

False negatives are events where the Cisco Intrusion Prevention System sensor does not react, even though there is malicious activity. This represents an error, generally caused by signatures that are too specific (that is, with descriptions of the attack that are too narrow), evasion, or the lack of any signatures for a particular attack.

The figure shows a common example of a false negative. The ICMP Network Sweep signature classifies a series of ICMP echo (ping) requests to at least five unique destination hosts within 60 seconds as an ICMP reconnaissance scan of the network. If an attacker scans the network more slowly, by generating fewer than five unique scan packets in 60 seconds, the sensor will not be able to correlate these packets into a reconnaissance event. You could tune these thresholds and time intervals, with the possible consequence that legitimate pings will be classified as a reconnaissance scan (that is, a false positive).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

False Positives

This topic describes false positive events.

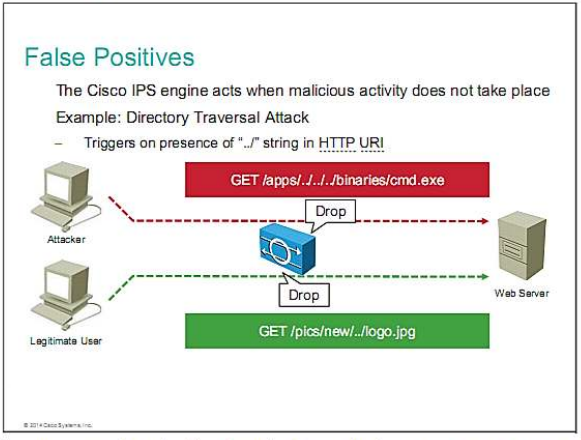

False positives are events where the Cisco Intrusion Prevention System sensor reacts as a consequence of nonmalicious activity. Such a response represents an error that is generally caused by signatures that are too relaxed (with attack descriptions that are too broad) and match attacks as well as legitimate traffic.

The figure shows common examples of a true positive and a false positive that are caused by the same sensor signature. The signature classifies all HTTP requests that contain the "../" string in the HTTP URI as directory traversal attacks. In directory traversal attacks, the attacker attempts to fool the web server into accessing files outside its normal working directories by referencing files using the parent directory (..) once or multiple times inside the requested path to a file.

In the upper part of the figure, the attacker is attempting to maliciously execute the Microsoft Windows cmd.exe command shell outside the applications directory of the normal web server. This event is a true positive, and the Cisco IPS or IDS sensor should react to it.

In the lower part of the figure, a user is accessing a web page that contains links to JPEG pictures. These links contain the "../" string because the web page designers have created relative links inside HTML content. Because the IPS sensor only looks for the "../" string inside an HTTP URI, it will classify these benign requests as attacks, possibly impacting the user session by dropping such legitimate connections.

You could change this signature to be more specific, and match on more than the "../" string, with the possible consequence of attacks not matching the new signature (that is, a false negative).

False Positive Examples

This topic describes examples of common false positives.

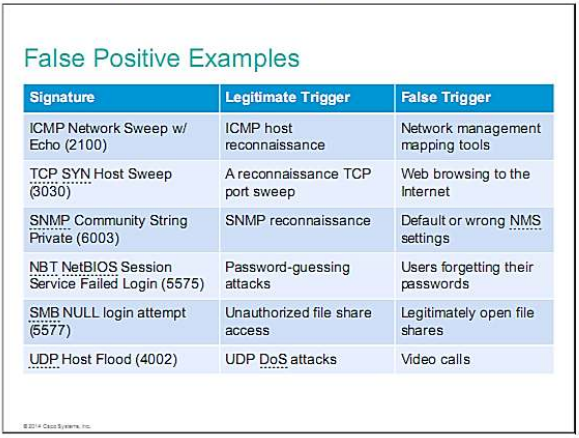

This table shows some common examples of false positives in enterprise networks. The listed signatures will likely require tuning in many networks, because their descriptions of malicious traffic will match legitimate enterprise applications and protocols.

https://t.me/learningnets

Cisco IPS Tuning Approaches

This topic describes the tuning of false positives and false negatives.

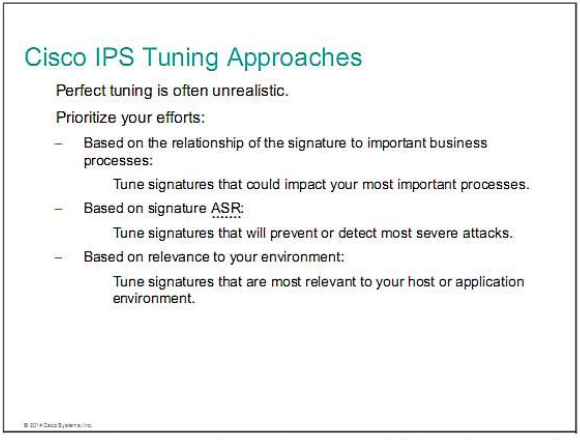

Often, the tuning process for each sensor is a complex task, especially in environments with many hosts and applications sending traffic through a sensor. In these cases, you will often not be able to tune a sensor perfectly because such a task would require too much of your time.

Consider prioritizing your tuning efforts based on one or more of the following criteria:

Tune signatures that are most relevant to your critical business processes. For example, tune false positives and negatives for signatures that are triggered by traffic to and from your most critical assets. This action ensures that you can rely on these tuned signatures and immediately use them in Cisco Intrusion Prevention System mode for your critical assets. Other signatures that have not been tuned yet may continue to exhibit false positives and negatives, and will need to be tuned later. Most often, tuning may be the most effective approach, and if carried out in multiple stages that are based on decreasing asset importance, can enable you to eventually fully tune each sensor.

Tune signatures that address the most serious attacks based on the signature ASR. This will ensure that your sensors do not misinform you (with false positives) about the most critical threats, and ensure that your sensors do not miss (with false negatives) the most devastating attacks against any of your assets. Therefore, you can enable preventive actions for these tuned signatures, and tune other signatures (that remain in detection mode) later. Because there are many signatures with high ASR values, this approach may be much more time-consuming compared to the previous one.

Tune signatures that are most relevant to your environment, regardless of severity or asset value. This will ensure that signatures that are relevant to your environment are very accurate and are ready to operate in IPS mode. The efficiency of this approach depends on the diversity and complexity of your environment, where diverse environments will require you to tune many signatures initially.

https://t.me/learningnets

Tune Cisco IPS to Reduce False Positives

This topic describes the possible approaches to reducing false positives.

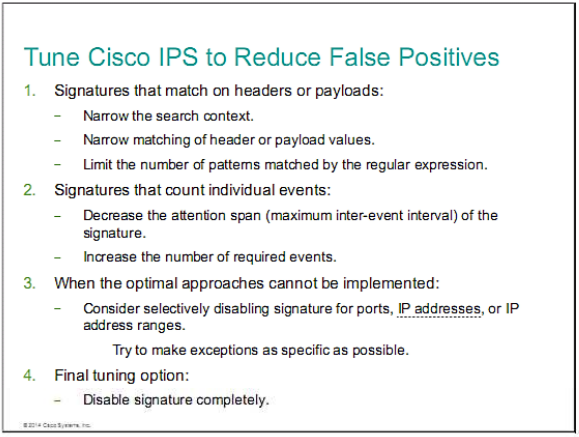

The fundamental reason for false positive events is signatures that are very general and that match legitimate traffic in addition to network attacks. Because management of false positives requires you to first separate false positives from true positives in your environment, and then change the manner in which these signatures operate, it cannot be automated.

There are three basic approaches to management of false positives:

1. The best but most difficult to deploy option is to change problematic signature-triggering conditions to make signatures slightly more specific and therefore avoid triggering on legitimate traffic. This option requires you to be proficient in signature design, and to thoroughly understand network attacks and the way they use network protocols. Depending on the signature engine and definition, you will adopt different approaches to limit its detection scope.
2. The second option is to selectively disable problematic signatures so that they do not trigger for specific source or destination hosts in network traffic. This approach is generally less optimal, because it will cause false negatives for specific attacks, because the sensor will not be able to use a particular signature for specific flows.
3. The third and least desirable option is to completely disable the problematic signature. While this approach is very simple to implement, it causes false negatives for all hosts (for the problematic signature), and should only be used if the attacks to which the signature applies do not present a major risk in your environment.

Note If you employ the first or second approach, it is very important to change the policy of the sensor in the narrowest sense possible. In other words, change the signature only for the hosts that cause false positives, or disable the signature only for the hosts that cause false positives.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reduce False Positives: Narrow Search Context

This topic explains how to reduce false positives by narrowing the search context.

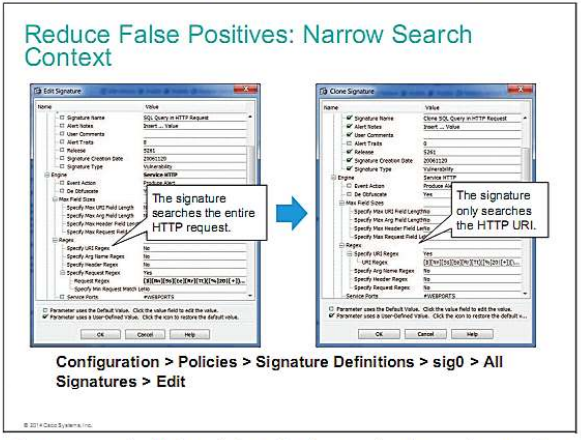

This figure shows an example of false positive tuning by narrowing the search context for a SERVICE.HTTP-based signature that looks for SQL injection patterns in the entire HTTP request (Cisco Intrusion Prevention System signature ID 5474, Query in HTTP Request).

If this signature fires incorrectly for some specific hosts or host pairs, you can clone the original signature to produce a copy of the original signature. Then, edit the cloned signature and copy the SQL injection-detecting regular expression from the Request Regex container to the URI Regex container, which only matches patterns inside the HTTP URI (a subset of the entire HTTP request).

You should then use event action filters to disable the original signature for the hosts that cause false positives with it, and only enable the cloned signature for the hosts that cause false positives with the original signature.

As a consequence, the new (cloned) signature will not be able to detect SQL injection attempts in the entire HTTP request (false negative) for problematic hosts, due to its more specific matching nature.

https://t.me/learningnets

Reduce False Positives: Narrow Header Values

This topic explains how to reduce false positives by narrowing the header values.

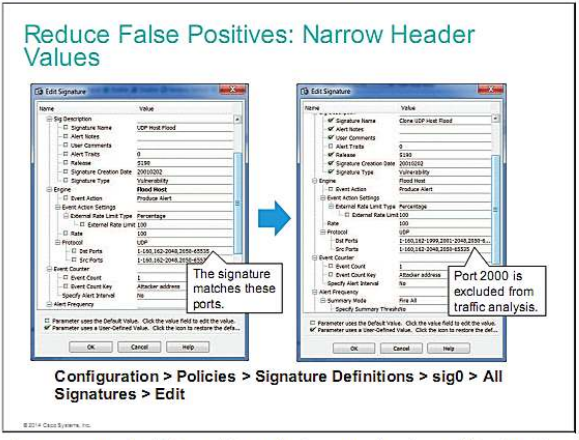

This figure shows an example of false positive tuning by narrowing the specific values from the original value set that is inside a signature. The example signature detects UDP host floods on various UDP ports (Cisco Intrusion Prevention System signature ID 4002, UDP Host Flood).

If this signature fires incorrectly for some specific hosts or host pairs because legitimate UDP traffic on some port (port 2000 in this example) exceeds the threshold rate, you can first clone the original signature to produce a copy of the original signature. Edit the cloned signature and change the Dst Ports value to exclude port 2000 from the port set.

You should then use event action filters to disable the original signature for the hosts that cause false positives with it, and only enable the cloned signature for the hosts that cause false positives with the original signature. As a consequence, the new (cloned) signature will not be able to detect UDP host floods on destination port 2000 for problematic hosts.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reduce False Positives: Limit Number of Matched Patterns

This topic explains how to reduce false positives by limiting the number of matched patterns.

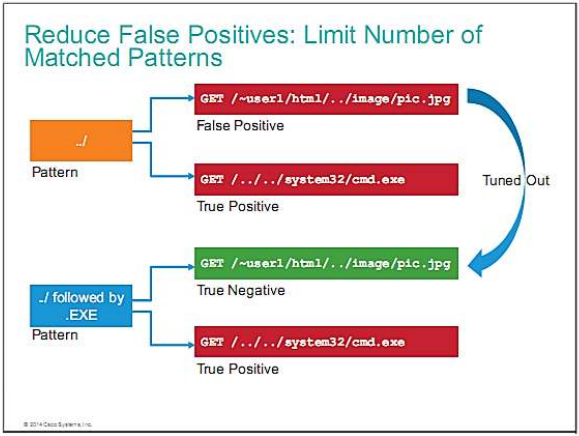

As an example of redesigning regular expressions, the directory traversal signature that matches on the "../" string is a good candidate for tuning. The original regular expression pattern will match both attacks and legitimate traffic. This example changes the regular expression pattern to be more specific, and only trigger if it detects the "../" string followed by access to a sensitive file, in this case any executable file that is identified by the .exe extension suffix. You could also include other sensitive file extensions or paths (/etc, SYSTEM32, WINDOWS, and so on) in the new signature that would apply only to problematic hosts causing false positives.

Reduce False Positives: Limit Number of Matched Patterns Example

This topic describes how to reduce false positives by limiting the number of matched patterns.

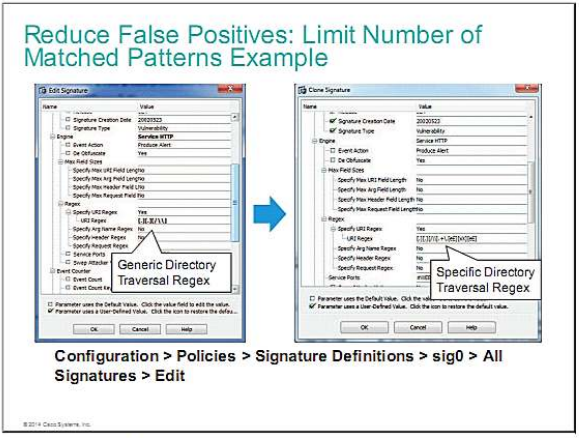

This figure shows an example of false positive tuning by changing the regular expression to match a smaller number of patterns. Take the original directory traversal signature (Cisco Intrusion Prevention System signature ID 5256, Dot Dot Slash in URI), and clone it to produce a copy of the original signature. Edit the cloned signature and change the URI Regex specification from **|.||.||/\\|** to **|.||.||/\\|.+\.|eE||xX||eE|**, which will match traffic where the dot dot slash pattern in the URI is followed by an .exe extension.

Next, use event action filters to disable the original signature for the hosts that cause false positives with it, and only enable the cloned signature for the hosts that cause false positives with the original signature. As a consequence, the new (cloned) signature will not be able to detect some directory traversal attacks that do not include references to .exe files, between hosts that had false positive issues with the original signature.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reduce False Positives: Decrease Attention Span

This topic explains how to reduce false positives by decreasing the attention span.

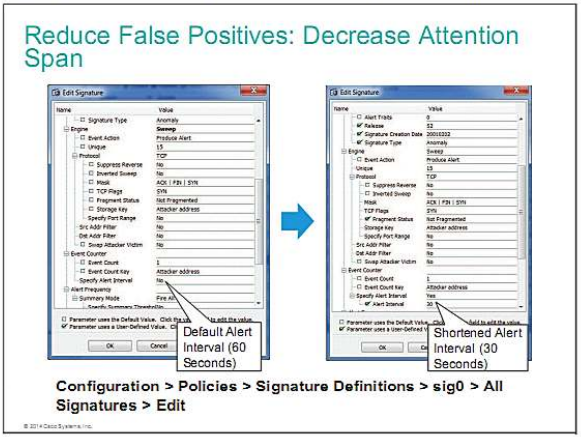

This figure shows an example of false positive tuning that involves lowering the attention span of a signature by lowering the time interval in which a set of events must be correlated. The example signature (Cisco Intrusion Prevention System signature ID 3030, TCP SYN Host Sweep) detects reconnaissance attacks that use TCP SYN packets to discover a specific service on multiple hosts in the network.

If this signature fires incorrectly for some specific hosts because there are multiple legitimate SYN packets to different destination hosts, you can clone the original signature and produce a copy of the original signature. Edit the cloned signature. Change the Specify Alert Interval value to Yes, and specify an Alert Interval of 30 seconds, which is shorter than the default of 60 seconds.

You should then use event action filters to disable the original signature for the hosts that cause false positives with it, and only enable the cloned signature for the hosts that cause false positives with the original signature. As a consequence, the new (cloned) signature will not be able to detect some reconnaissance attempts that do not fit in the new time window, but only from hosts that had false positive issues with the original signature.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Reduce False Positives: Increase Number of Events

This topic explains how to reduce false positives by increasing the number of expected events.

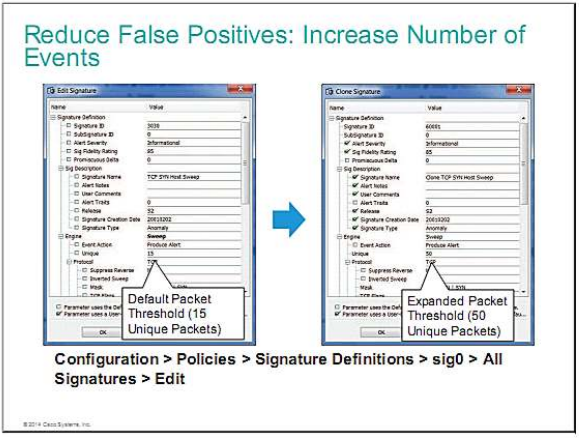

As an alternative or additional customization of a **counting** signature, you can also increase the number of packets that are required to trigger a **composite** signature. This figure shows the same cloned TCP SYN Host Sweep signature and its Unique parameter increased to require 50 instead of 15 unique packets to declare a reconnaissance attempt.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Tune Cisco IPS to Reduce False Negatives

This topic describes how to tune Cisco Intrusion Prevention System to reduce false negatives.

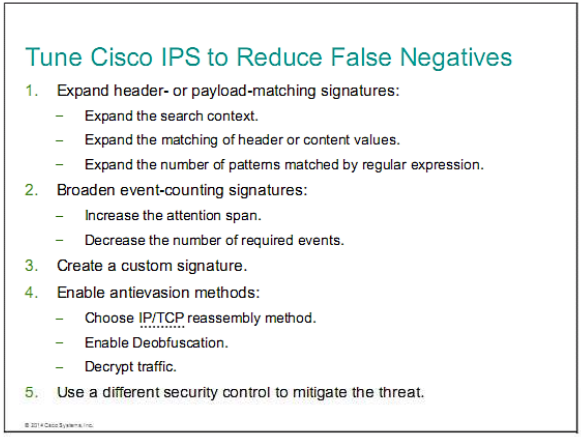

False negative events occur when a sensor does not react to malicious activity. By definition, they are not detected or prevented by the sensor, and must be discovered through secondary means.

The primary reasons for false negative events are signatures that are too specific, and consequently do not match variations of a specific attack, and the lack of signatures to address specific attacks and attacker evasion techniques.

There are several approaches to management of false negatives:

1. One of the best options, and one of the most difficult to deploy, is to change the triggering conditions of problematic signatures to make signatures slightly more generic and therefore allow them to trigger on additional malicious traffic patterns. As with false positives and signature tuning, this option requires you to be proficient in signature design, and to thoroughly understand network attacks and the way they use network protocols. In this approach you can expand the search context by examining more data than the original signature, expand the number of specific values that you are matching inside parsed protocol fields, or if using regular expressions, you can also change the regular expression itself to match a larger number of patterns.
2. The second option is to create a custom signature to address an attack that was previously undetected by the sensor. Again, this approach often requires expert knowledge.

3. The third option is to enable additional antievasion methods, if all antievasion options have not yet been exhausted, and if false negatives are caused by evasion techniques. The sensor uses IP and TCP traffic reassembly methods to ensure that application-layer inspection routines can properly parse byte streams that are exchanged between network applications. Misconfiguring, or creating exceptions in the sensor low-layer IP and TCP traffic reassembly routines, can therefore have a serious impact on application-layer signatures and can cause serious false negatives. Therefore, you should be extremely careful when tuning sensor low-level parameters. If you are forced to do so, consider placing the sensor close to another network device that performs protocol verification and normalization, such as an application-layer gateway (proxy) or SPF device with AIC, such as the Cisco Adaptive Security Appliance.
4. The fourth option is to deploy a different security control in addition to the Cisco IPS sensor, with the new control able to address the threat in a more reliable way.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reduce False Negatives: IP Reassembly

This topic explains how to reduce false negatives by tuning the IP reassembly function.

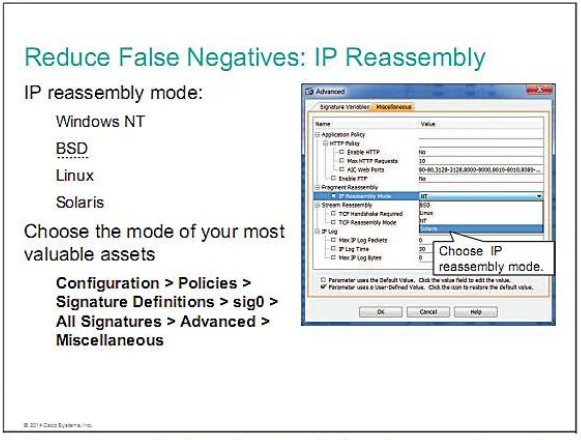

When a (virtual) sensor is operating in promiscuous mode, it receives a copy of network traffic and must be aware of the manner in which local hosts reassemble IP fragments. In promiscuous mode, the Cisco Intrusion Prevention System sensor allows you to configure its reassembly mode to match the local environment. The IP reassembly mode is a global setting for an entire virtual sensor, and can be configured in NT (Microsoft Windows stacks), BSD, Linux, and Solaris IP reassembly modes. You should choose the IP reassembly mode that covers the most critical assets in your network, or consider migrating the sensor to inline mode. Inline mode is especially recommended in heterogeneous environments, where the sensor inspects original network traffic and is able to change packets in order to fix IP reassembly ambiguities.

Reduce False Negatives: TCP Reassembly

This topic explains how to reduce false negatives by tuning the TCP reassembly function.

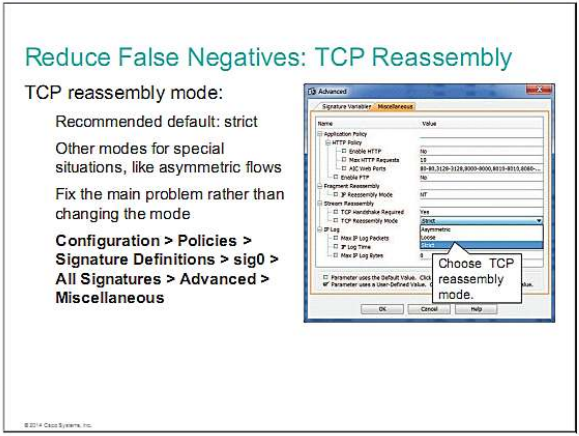

By default, the promiscuous or inline-mode sensor uses a strict TCP reassembly mode. In this mode, the sensor tracks TCP sessions only when it sees that the sessions are established normally, using a three-way handshake, and sees all packets of each TCP session. If possible, you should always run all virtual sensors in the strict mode.

The first alternative is to run a virtual sensor in loose TCP reassembly mode, where the sensor will also analyze traffic even if it does not receive all TCP segments of a session. You could, for example, encounter such a situation with a sensor in promiscuous mode that experiences packet drops on the SPAN port, or when the sensor is unable to keep up with traffic rates. Use the loose reassembly mode only as a final option, and prefer instead to fix the original problem, because loose reassembly can allow attackers to perform several evasion attacks and therefore cause false negatives.

The second alternative is to run a virtual sensor in the asymmetric TCP reassembly mode, where the sensor analyzes only one direction of a particular TCP session, typically as a result of asymmetric routing conditions in the network. Because this mode is also susceptible to evasion, it should only be used as a final option, if the original asymmetric routing conditions cannot be eliminated.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reduce False Negatives: Deobfuscation

This topic explains how to reduce false negatives by tuning the deobfuscation function.

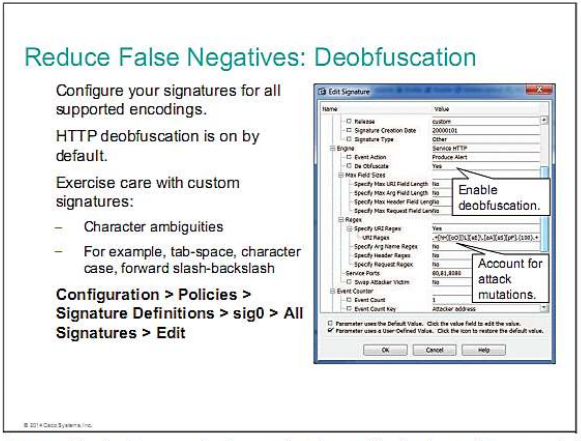

You can counter application-layer evasion by ensuring that application-layer signatures take into account variations of attacks and different formats of data encoding that attackers could use. The SERVICE.HTTP engine includes HTTP deobfuscation functionality (enabled by default for all SERVICE.HTTP signatures), where the sensor converts HTTP data to a canonical ASCII form before matching it against its signatures, which prevents format-based attacks. When creating custom signatures, ensure that you account for other evasion possibilities in your regular expressions, such as case sensitivity, using tabs instead of spaces, and using forward slashes instead of backslashes.

Summary

This topic summarizes the key points that were discussed in this lesson.

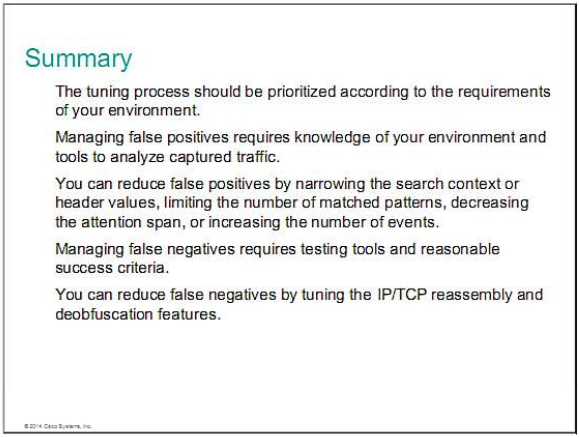

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 5

Configuring Custom Cisco IPS Signatures

Overview

When you need to address a threat that is uniquely specific to your environment, or otherwise do not have an appropriate signature in the default signature set to address a particular threat, you can create custom signatures on the Cisco Intrusion Prevention System sensor. This lesson describes the methods and configuration procedures that allow you to create custom signatures on a Cisco IPS sensor.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the need for custom signatures
- Identify the available signature engines
- Describe the custom signature configuration procedure
- Describe the custom signature matching strategies
- Configure a custom signature using three available methods
- Describe the configuration options in the Custom Signature wizard
- Describe the procedure for using the Custom Signature wizard with and without specifying the signature engine
- Verify custom signatures

https://t.me/learningnets

Custom Signatures Overview

This topic describes the need for custom signatures and matching strategies.

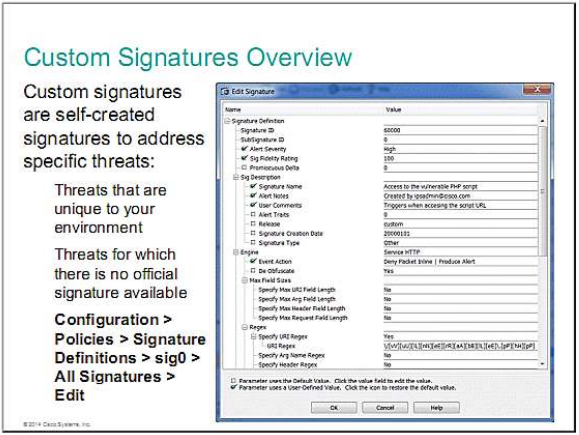

While Cisco Intrusion Prevention System sensors include a rich signature set, there are scenarios and environments where you may need to customize traffic analysis by creating custom signatures. This situation occurs most often in these cases:

- There are threats that are unique to your environment, and the signature set that Cisco provides cannot address them because the threats are globally unknown.
- There are threats that the default signature database does not address. For example, you may be using devices or operating systems with a relatively low market share, which do not warrant a prioritized creation of signatures for these platforms.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Signature Engines

This topic reviews the signature engines that you can use to define custom signatures.

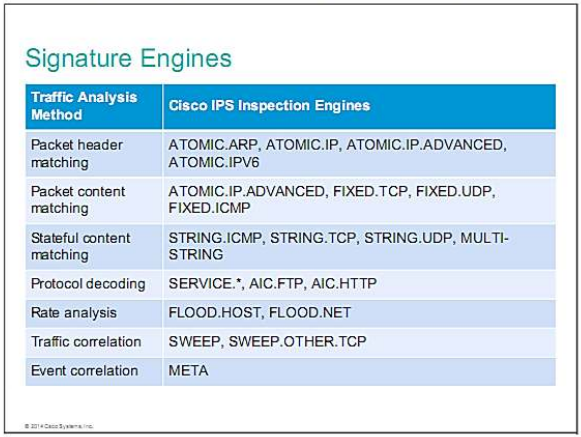

The table shows the Cisco Intrusion Prevention System software signature engines that you can use to create custom signatures.

Custom Signature Creation Procedure

This topic describes the procedure for creating custom signatures.

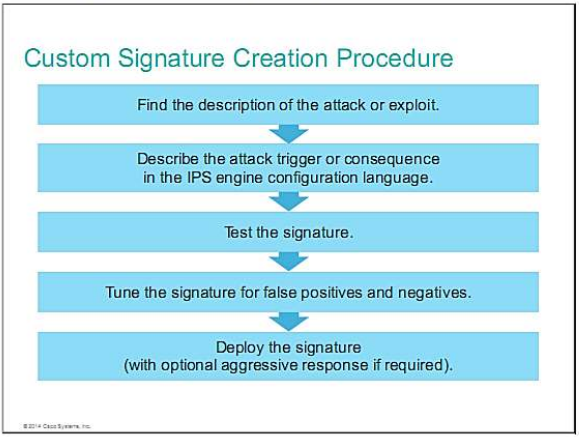

When creating custom signatures, follow this creation and deployment procedure:

1. Find the description of the attack or exploit, and a sample or description of the network traffic it generates.
2. Describe the network traffic unique to this attack (or network traffic that represents a consequence of this attack) using the configuration language of the Cisco Intrusion Prevention System sensor (that is, the signature engines).
3. Test this signature in a test environment to verify that it triggers with simulated attack traffic.
4. Test this signature in a production environment without aggressive actions, to determine its false positives and possible false negatives (with evasion testing). Tune the signature to achieve an acceptable balance between false positives and false negatives.
5. Deploy the signature in the production signature set, with an optional preventive (aggressive) response if required.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Custom Signature Matching Strategies

This topic lists possible matching strategies when defining custom signatures.

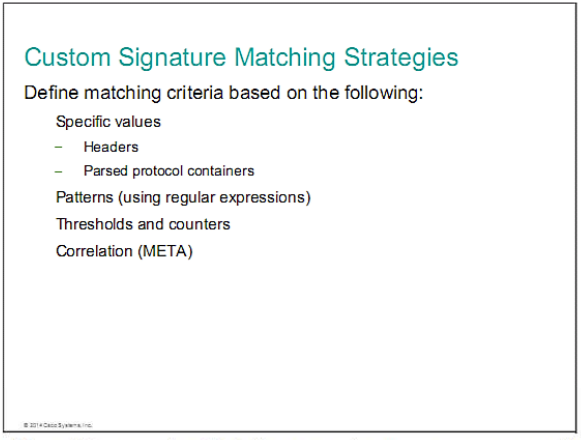

When describing malicious network traffic inside a custom signature, you can create matching criteria that are based on the following:

Specific values that are found in protocol units (packet headers, or parsed application-layer containers of application-layer protocols).

Data patterns, found anywhere in a packet or session, or inside parsed application-layer containers. Typically, you will use the regular expression language to describe such data patterns.

Thresholds and counters that are based on some form of stateful counting that takes into account events inside traffic that was seen in the past.

Correlation of multiple individual (custom or built-in) signatures to detect high-level activity or complex attacks.

Configure Custom Signatures

This topic describes the available custom signature configuration methods.

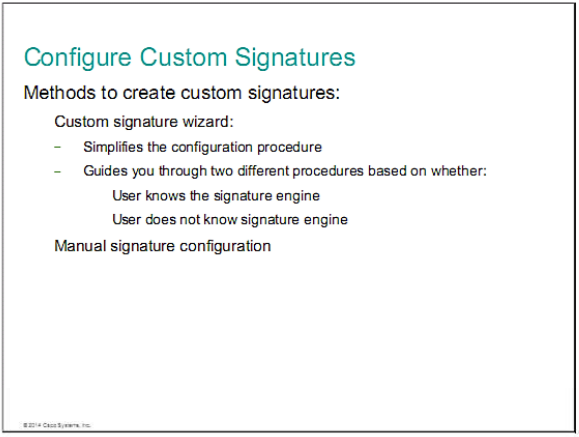

The simplest method of creating a custom signature on a Cisco Intrusion Prevention System sensor is to use the Custom Signature wizard. The Custom Signature wizard will guide you through the signature-creating process through a series of simple steps. The procedure differs depending on whether the user knows the signature engine.

Alternatively, you can create custom signatures using a manual configuration.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Custom Signature Wizard

This topic illustrates the operation of the Custom Signature wizard.

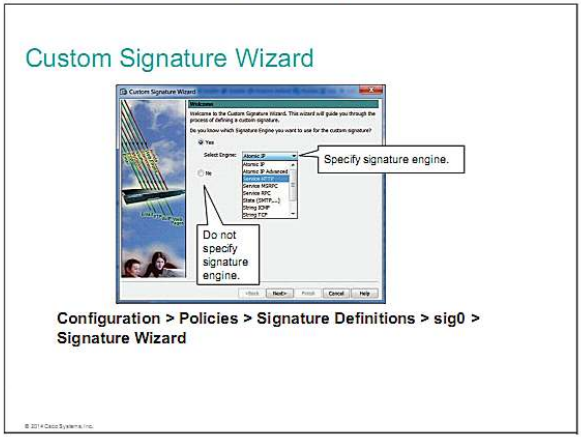

The graphic shows the prompt in the Custom Signature wizard. The following steps will depend on whether you have specified the signature engine.

Custom Signature Wizard Options

This topic describes the procedure for using the Custom Signature izard with and without specifying the signature engine.

The figure illustrates the selections that are offered by the Custom Signature wizard using both configuration methods.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Custom Signatures

This topic describes how to verify custom signatures.

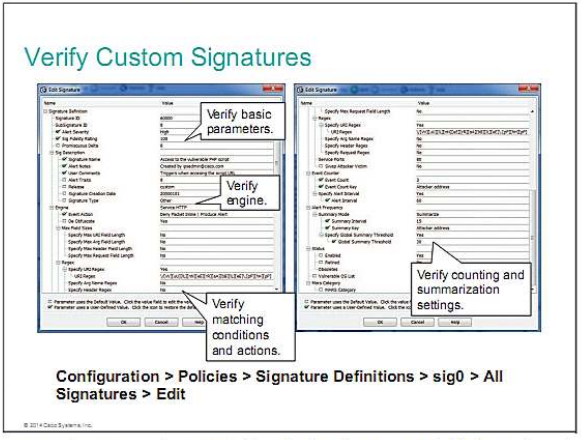

To verify a custom signature, navigate to the list of active signatures and edit the configured custom signature to verify all signature parameters.

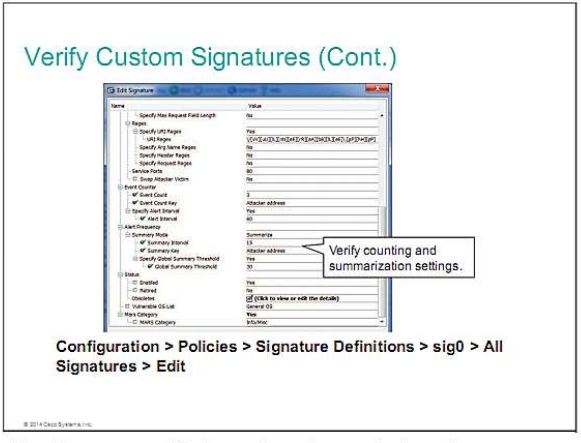

This figure shows how you can verify the counting and summarization settings.

Summary

This topic summarizes the key points that were discussed in this lesson.

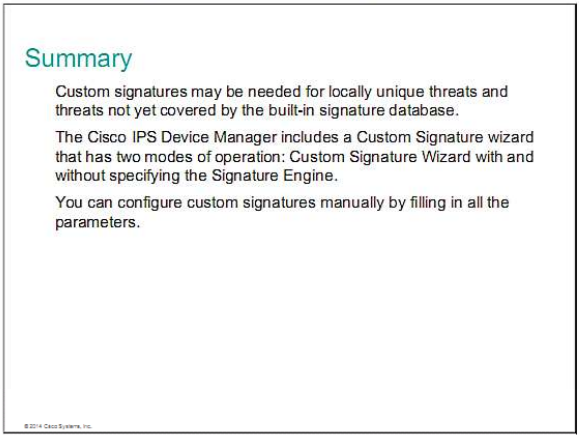

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 6

Configuring Cisco IPS Anomaly Detection

Overview

The anomaly detection component of the Cisco Intrusion Prevention System Sensor Software detects known and yet-unknown network threats and can take appropriate preventive actions to prevent their spread in the network. Anomaly detection enables the sensor to be less dependent on signature updates by letting the Cisco IPS sensor learn normal activity, send alerts, and take dynamic response actions for behavior that deviates from what it has learned as normal behavior. In this lesson, you will learn to deploy and troubleshoot the anomaly detection functionality of the Cisco IPS sensor.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the principles behind anomaly detection
- Describe the common worm scanning methods
- Describe scanners and histograms
- Describe the anomaly detection zones
- Describe the anomaly detection learning and knowledge base
- Describe anomaly detection and actions
- Describe an anomaly detection scenario
- Describe the anomaly detection configuration procedure
- Describe how to verify knowledge base creation
- Describe how to verify the operational mode of anomaly detection
- Describe how to verify anomaly detection statistics

https://t.me/learningnets

Anomaly Detection

This topic describes the principles behind anomaly detection and its functional components.

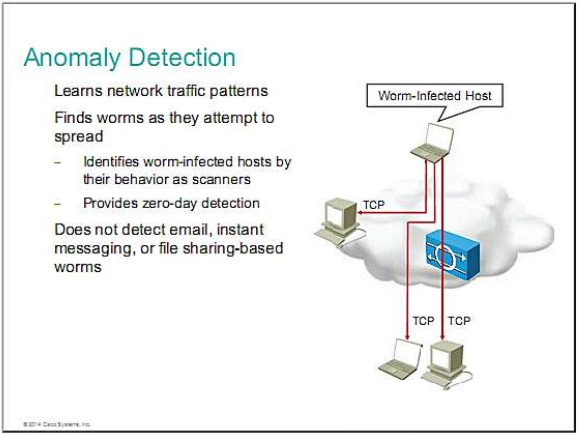

One aspect of the anomaly-based approach to intrusion detection and prevention involves the learning of statistical profiles of normal network traffic, and then detecting any deviation from such normal traffic behavior.

The Cisco Intrusion Prevention System software anomaly detection feature is a statistical anomaly-based intrusion detection and prevention technique that is designed to detect and prevent scanning network worms. Scanning network worms are malware that autonomously spread through the network by aggressively scanning the network to find appropriate targets, and then compromising those targets through network attacks, typically against exposed network services.

The Cisco IPS software anomaly detection feature dynamically learns network patterns, and after learning, is able to accurately identify known and yet-unknown worms as they attempt to spread through the network. The feature identifies individual infected hosts, and can alert you when worm activity goes beyond individual infections.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Worm Scanning Methods

This topic describes the common worm scanning methods.

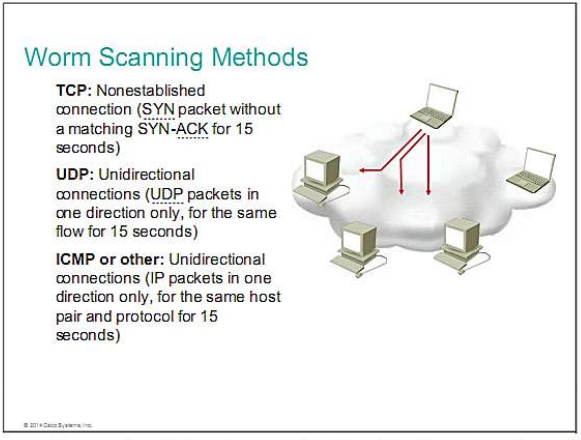

A scanning worm locates vulnerable hosts by generating a list of IP addresses to probe, and then contacts the hosts. Code Red worm, Sasser worm, Blaster worm, and the Slammer worm are examples of worms that spread in this manner.

Anomaly detection identifies worm-infected hosts by their behavior as scanners. A scanner is an IP address that generates scan events (connections from the same address to multiple addresses) on the same TCP or UDP destination port for too many destination IP addresses.

The anomaly detection algorithm considers the following events as an anomaly detection scan event:

For TCP connections, a scan event consists of nonestablished TCP connections, that is, a TCP SYN packet that does not receive its ACK response within 15 seconds. A worm-infected host that scans using the TCP protocol generates many nonestablished connections on the same destination port for many IP addresses because it blindly scans the network and attempts to contact nonexistent hosts.

For UDP flows, a scan event consists of unidirectional UDP flows, that is, a UDP flow where an initial UDP packet to a service goes unanswered for 15 seconds. A worm-infected host that scans using the UDP protocol generates many unidirectional flows on the same destination port for many IP addresses because it blindly scans the network and attempts to contact nonexistent hosts.

For other protocols, including ICMP messages, a scan event consists of unidirectional flows of a particular protocol, similar to how a UDP scanner is categorized.

Scanners and Histograms

This topic describes scanners and histograms.

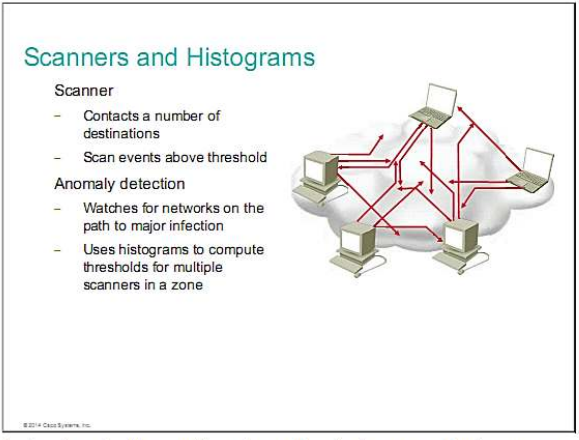

The anomaly detection algorithm considers a host to be a single scanner when it contacts a number of destinations inside a scan event that is greater than the scanner threshold. The scanner threshold is a value that is dynamically discovered or statically set in the anomaly detection database.

Cisco Intrusion Prevention System Sensor Software anomaly detection can also detect and react to situations in which multiple hosts become infected, and the network starts on its path to major infection. This is identified by multiple scanners that ramp up the overall scanning rate.

A histogram is a chart that represents a frequency distribution, and is often represented as a table that reflects this frequency distribution. Anomaly detection uses histograms to build a statistical profile of groups of hosts that perform scanner-like activities. The values in these histograms are dynamically learned or configured manually.

Anomaly detection builds separate histograms for each configured service in a zone. If you do not configure a service to be tracked, anomaly detection will use a default zone histogram for all other services.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

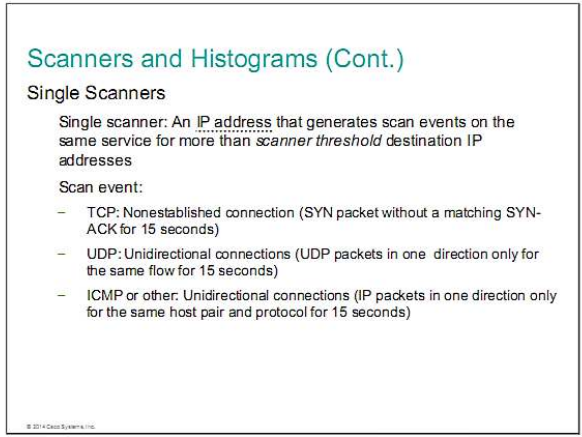

A scanner is an IP address that generates scan events (connections from the same address to multiple addresses) on the same TCP or UDP destination port for too many destination IP addresses. The anomaly detection algorithm considers the following events as an anomaly detection scan event:

For TCP connections, a scan event consists of nonestablished TCP connections, that is, a TCP SYN packet that has not received its ACK response within 15 seconds. A worm-infected host that scans using the TCP protocol generates many nonestablished connections on the same destination port for many IP addresses, because it blindly scans the network and attempts to contact nonexistent hosts.

For UDP flows, a scan event consists of unidirectional UDP flows, that is, a UDP flow where an initial UDP packet to a service goes unanswered for 15 seconds. A worm-infected host that scans using the UDP protocol generates many unidirectional flows on the same destination port for many IP addresses, because it blindly scans the network and attempts to contact nonexistent hosts.

For other protocols, including ICMP messages, a scan event consists of unidirectional flows of a particular protocol, similarly to how a UDP scanner is categorized.

The anomaly detection algorithm considers a host to be a **single scanner** when it contacts a number of destinations inside a scan event that is greater than the **scanner threshold**. The scanner threshold is a value that is dynamically discovered or statically set in the anomaly detection database.

Scanners and Histograms (Cont.)

Histograms

A histogram is a table that represents the distribution of the source IPs according to their expected scanning behavior.

- Histograms are learned or configured manually.
- Destination IP address row is the same for all histograms.
- Source IP address row can be learned or configured.

You can enable a separate histogram and scanner threshold for specific services, or use the default one for all other services.

Number of source IP addresses	A	B	C
Number of destination IP addresses	5	20	100

©2014 Cisco Systems, Inc.

Anomaly detection histograms track source and destination IP addresses. The destination address row in an anomaly detection histogram is always the same, and represents a threshold of dynamic addresses that sources will scan. The source address row in anomaly detection histograms is typically dynamically learned, and shows the number of sources that would normally scan a particular number of destinations that are shown in the corresponding destination field.

In the table, the first column represents a certain number of sources, scanning five different destinations, where A represents the learned normal number of sources in a particular network. If the sensor observed more than A sources scanning more than five different destinations, it would consider this an anomaly. The second column represents a certain number of sources, scanning 20 different targets, where B represents the learned number of sources, scanning up to 20 different targets. The last column is a certain number of sources, scanning 100 different targets, where C represents the learned number of sources, scanning up to 100 different targets. Collectively, the table represents a frequency distribution for network scanning.

https://t.me/learningnets

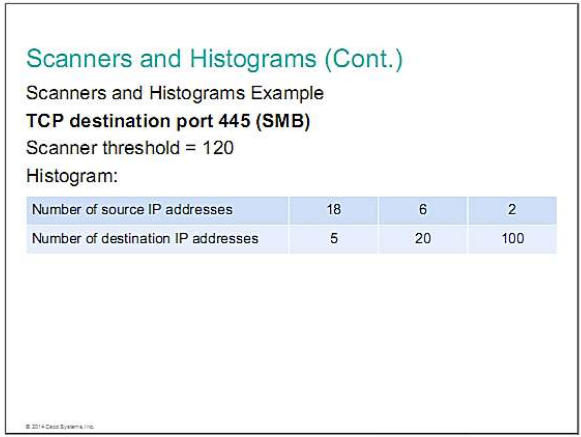

In the example, the concepts of scanner threshold and the histogram are combined. This example shows the learned scanner threshold and the histogram for the TCP protocol, destination port 445 (the Server Message Block protocol that is used in Microsoft file sharing and remote procedure calls).

From a single source you do not expect to see nonestablished connections to more than 120 different destination IP addresses in this zone.

You do not expect to see more than 18 sources generate nonestablished connections to five or more different destinations in this zone in 60 seconds.

You do not expect to see more than 6 sources generate nonestablished connections to 20 or more different destinations in this zone in 60 seconds.

You do not expect to see more than 2 sources generate nonestablished connections to 100 or more different destinations in this zone in 60 seconds.

Given a scanner threshold of 120, the example says that not more than 120 incomplete connections to different destinations are expected to be seen from a single IP address. If that occurs, the IP address is designated as a single scanner.

The histogram defines the rest of the expectations; however, the histogram does not consider individual hosts, but more general scanning behaviors. This histogram does not expect to see more than 18 different scanners that each scan up to five destination addresses. It also does not expect to more than six different scanners that each scan up to 20 destination addresses. Finally, this histogram example does not expect to see more than two different scanners that generate incomplete connections to 100 or more different destinations. If these thresholds are violated, anomaly detection would recognize this behavior as multiple scanners being present in the network.

All nonestablished connections are counted in 60-second time intervals.

Anomaly Detection Zones

This topic describes anomaly detection zones.

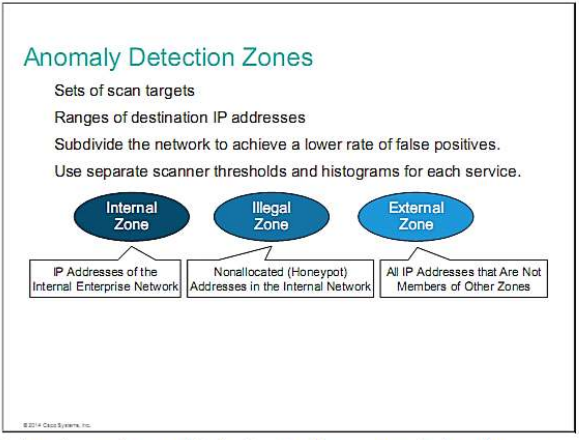

An anomaly detection zone is a set of destination IP addresses. Anomaly detection uses zones to achieve a lower false positive rate by subdividing the network into zones. Anomaly detection will separately track single scanners and histograms for each service in each zone. Because each zone will typically exhibit different behavior, the thresholds in each zone will be different and adapted to the host behavior of that particular zone.

There are three types of zones, each with its own thresholds: internal, illegal, and external zones.

It is recommended that you configure the internal zone with the IP address range of your internal network. If you configure the internal zone in this way, the internal zone tracks all connections that terminate in your IP internal address range and the external zone tracks all connections that terminate externally, for example in the Internet.

You can configure the illegal zone with IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses, or part of your internal IP address range that is unoccupied. An illegal zone can be very helpful for accurate detection, because no legal traffic is expected to reach this zone. This configuration allows very low thresholds, which in turn can lead to very quick worm virus detection.

The external zone is the default zone with the default range of 0.0.0.0–255.255.255.255. By default, the internal and illegal zones contain no IP addresses. Packets that do not match the set of IP addresses in the internal or illegal zone are managed by the external zone.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anomaly Detection Learning and Knowledge Base

This topic describes the anomaly detection learning and knowledge base.

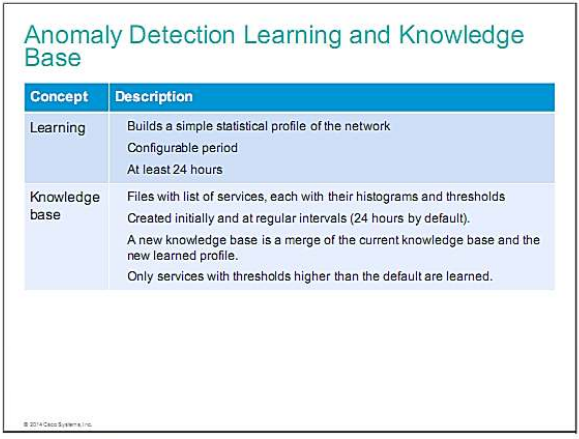

Anomaly detection initially conducts a learning process when the most normal state of the network is reflected. Anomaly detection then derives a set of scanner and histogram thresholds that best fit the normal network. This learning is done in two phases:

- During the initial anomaly detection learning mode, where the initial network profile is built.
- During anomaly detection detection mode, where the sensor still builds statistical network profiles, as long as it does not consider the network under worm attack.

In the initial setup, the sensor is in learning mode. It is assumed that during this phase no attack is being carried out. Anomaly detection creates an initial baseline of the network traffic. This initial baseline is known as a knowledge base. The default amount of time for anomaly detection to be in the learning mode is 24 hours, but depending on your network complexity, you may want to change this default. After the learning mode time has expired, terminate this phase by configuring anomaly detection to operate in detection mode.

For ongoing operation, the sensor is in learning plus detection mode. The sensor is in this state 24 hours, 7 days a week. Once the sensor creates a knowledge base, anomaly detection detects attacks based on the knowledge base. The sensor looks at the network traffic flows that violate thresholds in the knowledge base and sends alerts. Because anomaly detection looks for anomalies, it also records gradual changes to the knowledge base that do not violate the thresholds, and in doing so, creates a new knowledge base. The new knowledge base is periodically saved and takes the place of the old one so that an up-to-date knowledge base is maintained. You can also configure anomaly detection to avoid automatic updates to the knowledge base. Instead, you can approve and load the new knowledge bases manually.

By default, anomaly detection functions even if you do not follow the two phases and manually change the operational mode from learning to detection.

Note	Anomaly detection does not detect attacks when working with an empty knowledge base. After the default of 24 hours, the default operational mode is changed to detection. The knowledge base is saved and loaded, and anomaly detection now also detects attacks.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anomaly Detection and Actions

This topic describes anomaly detection and actions.

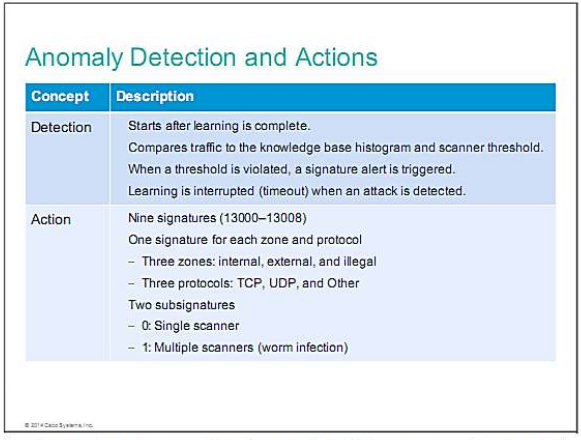

Anomaly detection monitors the network and constantly looks for worms and scanners. When the scanner threshold is crossed, a single signature alert is triggered. When a histogram threshold is crossed, the scanner is assumed to be a worm and the multiple scanner signature is triggered. During the time when the sensor believes that there is a serious worm attack (multiple scanners), learning is suspended so that the anomalous traffic is not included as part of normal traffic. In this period, anomaly detection also lowers the scanner threshold for the offending service in order to quickly trigger single scanner signatures for all infected hosts.

Because learning is suspended, the learned baseline of normal traffic should not be affected. Once the worm attack is over, learning resumes. The time period for resuming learning is configurable.

The Cisco Intrusion Prevention System Sensor Software anomaly detection functionality uses the TRAFFIC anomaly signature engine to activate its rules, which it also calls signatures (although this is misleading). The TRAFFIC anomaly signature engine contains nine anomaly detection signatures covering the three zones, and three protocols: TCP, UDP, and Other. Each signature has two subsignatures, one for the single scanner and the other for multiple scanners (indicating a more serious, networkwide infection). When anomaly detection discovers an anomaly, it triggers an alert for these signatures. All anomaly detection signatures are enabled by default and the alert severity for each one is set to High. No preventive actions are configured inside these signatures by default.

Because the multiple scanners signature has no specific attacker IP address field, you should only assign preventive actions to the single scanner signature. Anomaly detection automatically lowers scanner thresholds to help identify individual infected hosts when a wider infection is detected.

Anomaly Detection Scenario

This topic describes an anomaly detection scenario.

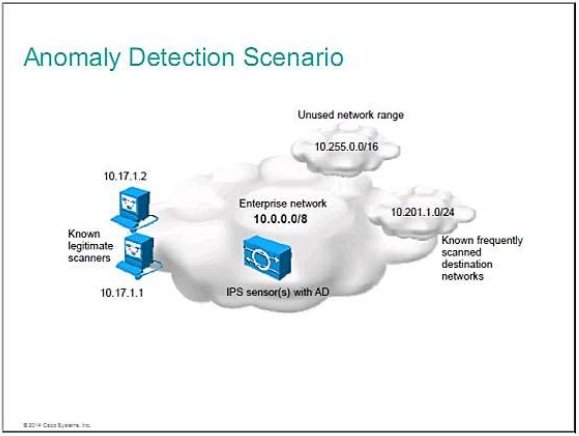

This figure shows a typical configuration scenario. The Cisco Intrusion Prevention System sensor will operate in anomaly detection mode (in addition to its normal, signature-based mode) in an enterprise network. The enterprise network is addressed using the 10.0.0.0/8 address range. There are two legitimate scanners (network management workstations) addresses with the 10.17.1.1 and 10.17.1.2 addresses. The internal network 10.201.1.0/24 is a frequent destination of legitimate internal scans. The 10.255.0.0/16 internal address range is unused, and can be used as the range for the anomaly detection illegal zone.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anomaly Detection Configuration Procedure

This topic describes the anomaly detection configuration procedure.

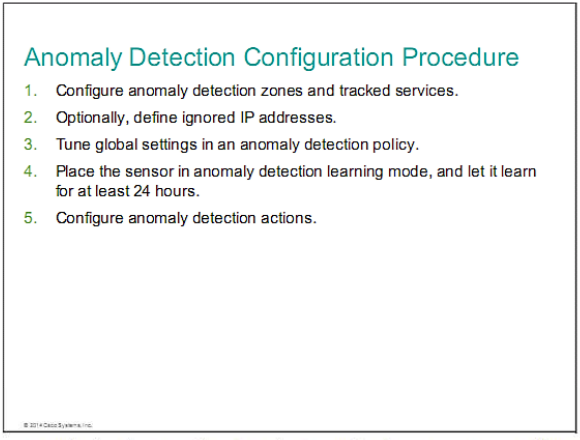

To configure anomaly detection on a Cisco Intrusion Prevention System sensor, you will need to perform the following configuration tasks:

1. Configure anomaly detection zones and specific tracked services in an anomaly detection policy. Optionally, you can also configure custom thresholds for the default group of services.
2. Optionally specify ignored IP addresses. The ignored IP addresses are exempted from scanner detection. These addresses will belong to legitimate scanners in your network.
3. Tune global anomaly detection settings in an anomaly detection policy. These values will further adapt anomaly detection to your environment.
4. Place the sensor into the learning mode, and leave the configured anomaly detection policy in learning mode for at least 24 hours to establish the initial baseline knowledge base. After the learning period, the Cisco IPS sensor will automatically switch its anomaly detection policy from learning to detection mode.
5. Configure anomaly detection actions to optionally take preventive action when infected hosts are detected by anomaly detection.

https://t.me/learningnets

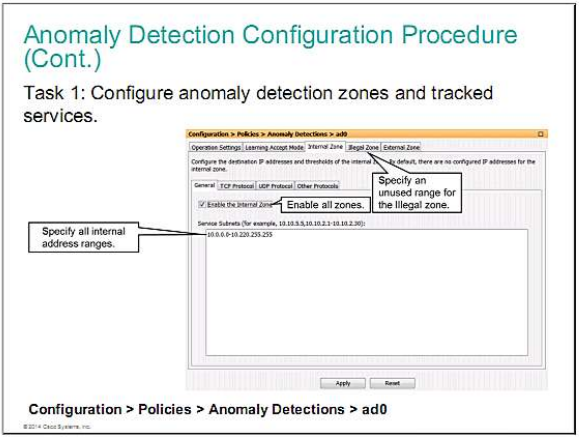

In Task 1, you need to configure all three anomaly detection zones (Internal, External, and Illegal) with appropriate parameters. You should enable all zones, and specify the destination addresses for the Internal and Illegal zones first. Perform the following steps:

Follow these steps to configure the learning accept mode for anomaly detection:

1. Click **Configuration**, choose **Policies > Anomaly Detections > ad0**, and click the **Internal Zone** tab.
2. Verify that the **Enable the Internal Zone** check box is checked.
3. In the Service Subnets section, specify all address ranges that comprise your internal network. Separate multiple ranges with a comma.
4. Switch to the **Illegal Zone** tab.
5. Verify that the **Enable the Illegal Zone** check box is checked.
6. In the Service Subnets section, specify address ranges that are unused in your network but would get routed across this sensor.
7. Click **Apply** to save your configuration.

In the example, 10.0.0.0–10.200.255.255 address range is used for internal network zone.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Anomaly Detection Configuration Procedure (Cont.)

Task 1: Configure anomaly detection zones and tracked services (cont.).

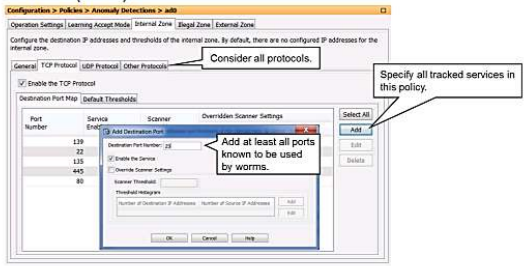

Configuration > Policies > Anomaly Detections > ad0

Next, you will need to specify which services are explicitly tracked by the anomaly detection algorithm. This step is optional, but recommended as anomaly detection will build separate profiles for each tracked service, making anomaly detection more accurate. You should add services (destination ports) that are known to be used by worms to all zones.

Perform the following task in all three zones, for all three protocols (TCP, UDP, and Other protocols):

8. Click **Configuration**, choose **Policies > Anomaly Detections > ad0**, and click the appropriate zone tab.
9. Select the appropriate protocol tab in the zone tab, and verify that the protocol is enabled.
10. In the protocol tab, click the **Add** button and add all ports that you consider to be a known service that worms would use.
11. When finished with all zones and protocols, click **OK**.

In the destination port settings, you can optionally also override default histograms and scanner thresholds. Only use these settings if you are seeing false positives or false negatives with your anomaly detection policy.

In the figure, as an example for the internal zone, the TCP protocol is enabled and explicitly tracking for port 25 is added using the default histograms and scanner thresholds.

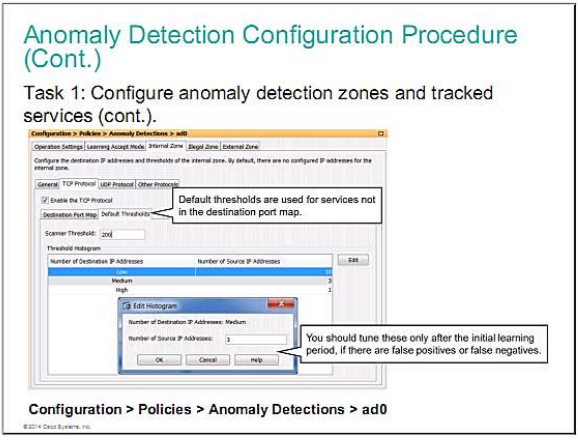

Next, you can optionally alter the built-in default thresholds. These thresholds are used for all services that you have not explicitly specified in the destination port maps for the three protocols. It is recommended that you use these default thresholds, and only change them if you are seeing false positives or false negatives with your anomaly detection policy.

The default thresholds are as follows:

Scanner threshold: 200 scanners

Histogram thresholds:

- **Low:** Ten source IP addresses where there are 5 destination IP addresses
- **Medium:** Three source IP addresses where there are 20 destination IP addresses
- **High:** One source IP address where there are 100 destination IP addresses

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

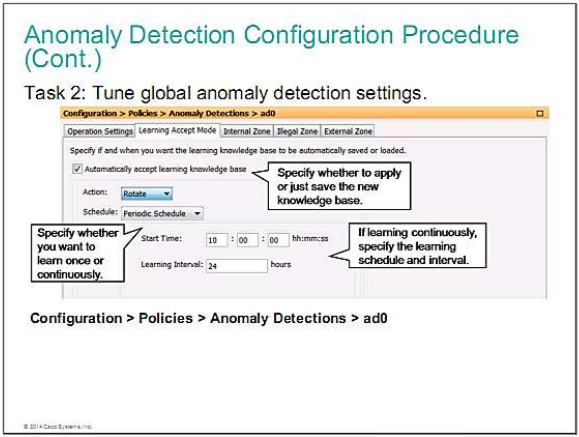

In Task 2 of anomaly detection configuration, you will configure global anomaly detection settings that will influence anomaly detection engine operation.

Use the Learning Accept Mode tab to configure whether the sensor will automatically create a new knowledge base every so many hours. You can configure whether the knowledge base will be created and loaded (Rotate) or saved (Save Only). You can schedule how often and when the knowledge base will be loaded or saved. The new updated knowledge base is saved as KB_*current-date*.

Note You must be an administrator or operator to configure the learning accept mode.

Follow these steps to configure the learning accept mode for anomaly detection:

1. Log into the Cisco IPS Device Manager using an account with administrator or operator privileges.
2. Click **Configuration**, choose **Policies > Anomaly Detections > ad0**, and click the **Learning Accept Mode** tab.
3. To have anomaly detection automatically update the knowledge base, check the **Automatically Accept Learning Knowledge Base** check box.
4. From the Action drop-down list, choose one of the following action types:

   Rotate: With this action option, a new knowledge base is created and loaded. This option is the default.

   Save Only: With this action option, a new knowledge base is created but not loaded. You can view it to decide if you want to load it.
5. From the Schedule drop-down list, choose one of the following schedule types:

   Calendar Schedule

   Periodic Schedule
6. To configure the periodic schedule (the default):

In the Start Time fields, enter the start time in hours, minutes, and seconds using the 24-hour time format.

In the Learning Interval field, enter the interval of the subsequent knowledge base snapshots.

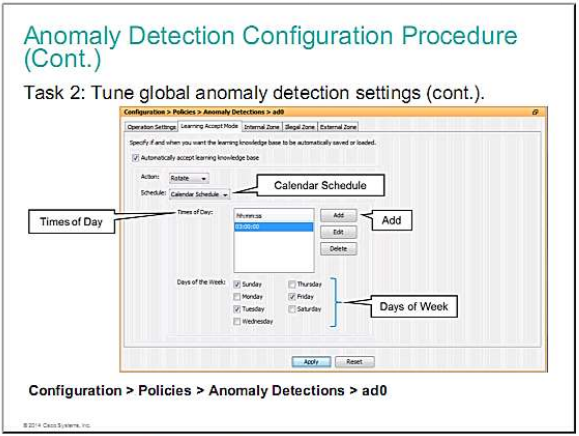

7. To configure the calendar schedule, follow these steps:

   Click **Add** to add the start time. The Add Start Time dialog box appears.

   Enter the start time in hours, minutes, and seconds using the 24-hour time format.

   Click **OK**.

   In the Days of the Week field, check the check boxes of the days that you want the anomaly detection module to capture knowledge base snapshots.

8. Click **Apply** to apply your changes and save the revised configuration.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

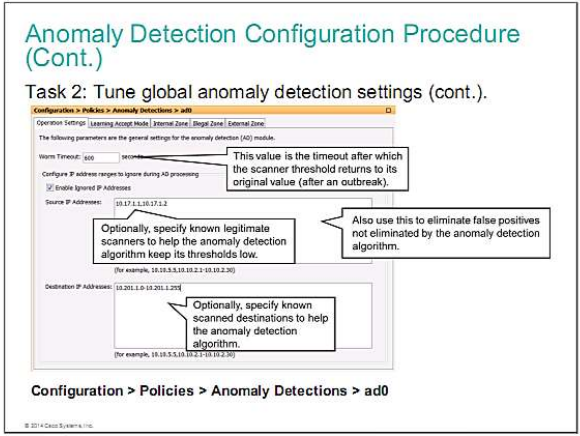

Next, you will configure some anomaly detection engine parameters that will help it operate better in your environment.

From the Operation Settings tab, you can set the worm detection timeout. If anomaly detection detects a worm outbreak, it lowers thresholds in order to "catch" individual scanners (infected hosts). After this timeout, the scanner threshold returns to the configured value, and anomaly detection stops indicating infected hosts until the next outbreak. You may consider increasing this timeout if the default timeout does not allow you to pinpoint all infected hosts in your network.

You can also configure source and destination IP addresses that you want the sensor to ignore when anomaly detection is gathering information for a knowledge base. Anomaly detection does not track these source and destination IP addresses, and the knowledge base thresholds are not affected by these IP addresses. Refer to the guidelines at the beginning of this topic for guidance on these exceptions.

Follow these steps to configure anomaly detection operation settings:

1. Log in to the Cisco IDM using an account with administrator or operator privileges.
2. Click **Configuration** and choose **Policies > Anomaly Detections > ad0** and click the **Operation Settings** tab.
3. In the Worm Timeout field, enter the number of seconds that you want to wait for worm detection to time out. The range is 120 to 10,000,000 seconds. The default is 600 seconds.
4. To enable the list of ignored IP addresses, check the **Enable Ignored IP Addresses** check box.
5. In the Source IP Addresses field, enter the addresses or range of source IP addresses that you want anomaly detection to ignore. In our example as shown in the figure, there are two legitimate scanners (network management workstations) addresses with the 10.17.1.1 and 10.17.1.2 addresses.
6. In the Destination IP Addresses field, enter the addresses or range of destination IP addresses that you want anomaly detection to ignore. In our example as shown in the figure, the internal network 10.201.1.0/24 is a frequent destination of legitimate internal scans.
7. Click **Apply** to apply your changes and save the revised configuration.

https://t.me/learningnets

Anomaly Detection Configuration Procedure (Cont.)

Task 3: Place the sensor in anomaly detection learning mode.

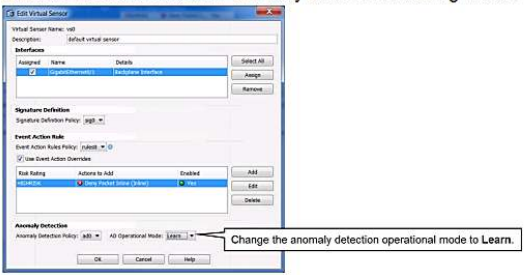

Configuration > Policies > IPS Policies > Edit Virtual Sensor

In Task 3, you will start learning the initial knowledge base by putting the sensor into anomaly detection learning mode.

Follow these steps to set the proper operational mode:

1. Click **Configuration** and choose **Policies > IPS Policies**.
2. To edit a virtual sensor, choose the virtual sensor and click **Edit**.
3. Choose **Learn** from the AD Operational Mode drop-down list.

Anomaly Detection Configuration Procedure (Cont.)

Task 4: Configure anomaly detection actions.

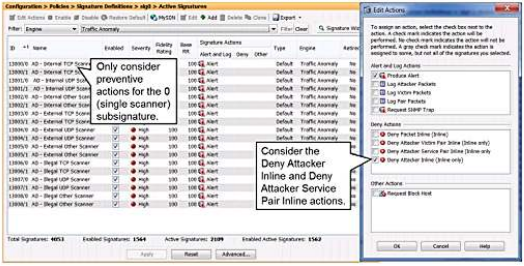

Configuration > Policies > Signature Definitions > All Signatures

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Finally, in Task 4, you can enable preventive actions for when anomaly detection triggers its anomaly signatures. By default, all anomaly detection signatures are enabled and produce alerts.

Note that you should enable preventive actions only after you verify that anomaly detection signatures do not cause false positives, by letting them operate in your network in detection mode for a while.

To enable preventive actions for the anomaly detection signatures, select the relevant anomaly detection signatures from the signatures list. It is easiest to find them if you filter the signature list by engine, and select the TRAFFIC Anomaly Engine. All of these anomaly detection signatures are enabled by default and the alert severity for each one is set to High.

You can consider enabling the Deny Attacker Inline and Deny Attacker-Service Pair Inline actions to stop infected hosts from sending traffic across the sensor. Note that you should only enable these actions for the 0 subsignature, because the 1 subsignature will not identify individual hosts.

https://t.me/learningnets

Verify Knowledge Base Creation

This topic describes the verification of knowledge base creation.

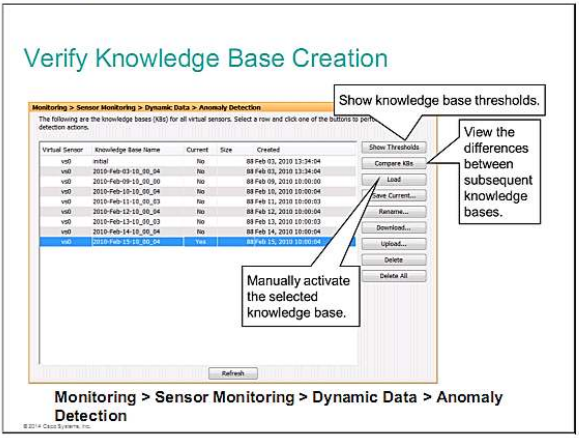

The Anomaly Detection monitoring pane displays the knowledge bases for all virtual sensors. Choose the **Monitoring > Sensor Monitoring > Dynamic Data > Anomaly Detection** pane to verify anomaly detection operation. If you see knowledge bases generating periodically, the scheduled generation is working properly.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Anomaly Detection Operational Mode

This topic describes how to verify the operational mode of anomaly detection.

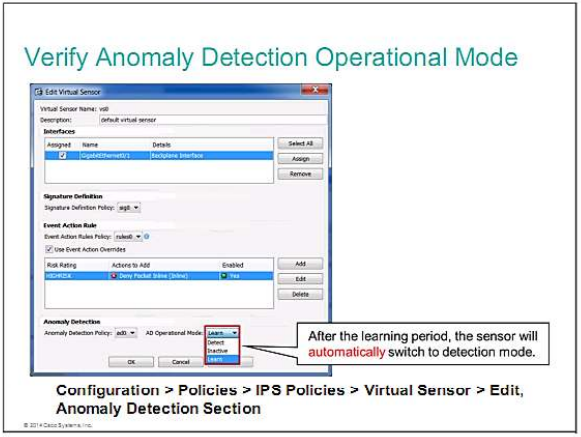

In the **Configuration > Policies > IPS Policies > Virtual Sensor > Edit** page, under the **Anomaly Detection section, AD Operational Mode**, the sensor should automatically switch from the **Learn** mode to **Detect** mode after the configured learning period.

Verify Anomaly Detection Statistics

This topic explains how to verify statistical information on anomaly detection.

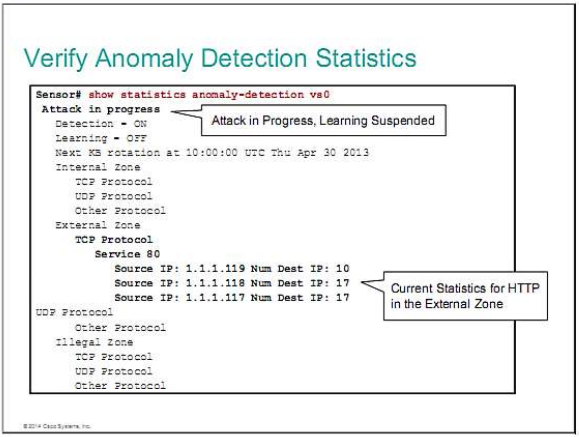

The **show statistics anomaly-detection** CLI command is used to display anomaly statistics. It also reveals whether an anomaly has been detected and the source of the worm infestation. In the example, an attack has been detected, which has in turn suspended the learning process so that malicious activity is not learned as normal. A list of source IP addresses is also displayed.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

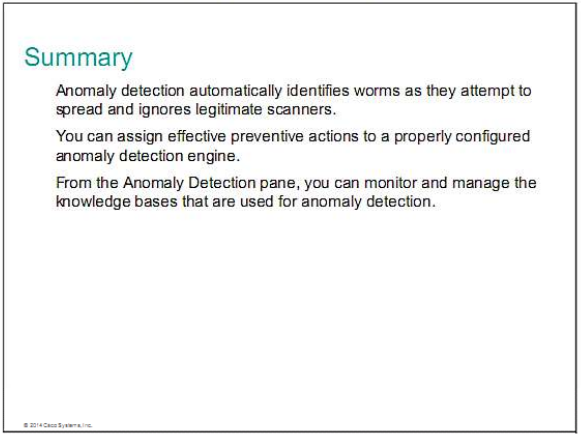

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 7

Configuring Cisco IPS Reputation-Based Features

Overview

When integrating the Cisco Intrusion Prevention System sensor with your network and system environment, you can enable several Cisco IPS sensor features that can increase the quality of the alarms and responses that are produced. This lesson provides configuration guidance for these reputation-based features.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the two options for using reputation data from the Cisco SensorBase
- Describe the traffic processing flow in the IPS sensor Global Correlation and Reputation Filter active
- Describe the reputation filters operations
- Describe global correlation operations
- Describe the process for receiving updates from Cisco SensorBase
- Describe how the IPS sensors send information to Cisco SensorBase using network participation
- Describe the global correlation inspection and reputation-based filtering configurations
- Verify Global Correlation and Reputation Filter operations

https://t.me/learningnets

Global Correlation and Reputation-Based Filtering Overview

This topic describes the two ways of using reputation data from the Cisco SensorBase.

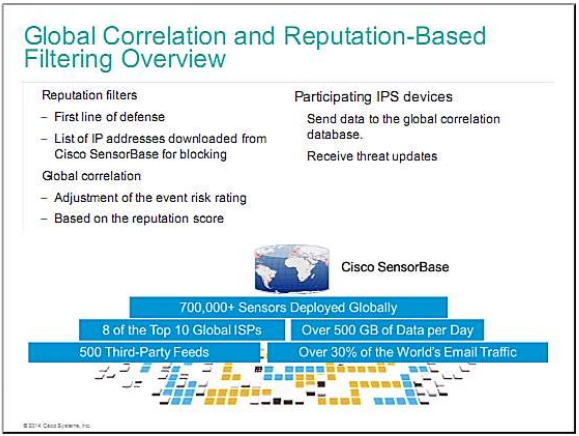

With Cisco Global Correlation, Cisco Intrusion Prevention System sensors receive threat updates from the Cisco SensorBase network, which contains detailed information about known threats on the Internet, including serial attackers, botnet harvesters, malware outbreaks, and dark nets. Cisco IPS uses this information to filter out the worst attackers before they have a chance to attack critical assets. It then incorporates the global threat data into its system to detect and prevent malicious activity even earlier.

Participating Cisco IPS devices receive and absorb global correlation updates from the SensorBase network. The participating Cisco IPS devices send data back to the global correlation database, which results in a feedback loop that keeps the updates current and global. Using network participation, Cisco collects nearly real-time data from sensors around the world.

There are two mechanisms that Cisco IPS sensors use to correlate network traffic with SensorBase data:

Reputation filters are a traffic-filtering mechanism that denies hosts that are on the SensorBase blacklist from forwarding traffic through the sensor. Using reputation filters, you can establish a first line of defense that drops traffic regardless of malicious activity, and automatically denies known hostile hosts from exchanging any data with resources that are protected by the sensor.

Global correlation is a mechanism that influences the risk rating and increases the risk rating for events that are caused by hosts with a negative reputation score. The reputation score of a host depends on its past behavior, and provides a method for increasing the reliability of prevention by increasing the probability of aggressive response for hosts with a known history of malicious behavior. The actual formula that is used to calculate risk rating increase is complex, relies on statistical information, and may change with time. For example, an event with a calculated risk rating of 85 has triggered, and the sensor has determined (using a query to the SensorBase network) that the attacker has a reputation of -5. the sensor will increase the risk rating of this event to 99, increasing the probability of prevention when using event action overrides.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Note	Originally, SensorBase was called SenderBase. The Cisco IPS documentations use the term SensorBase since reputation filtering was introduced into the Cisco IPS product. When the SenderBase framework grew to include data collected from hundreds of thousands of Cisco IPS sensors deployed around the world, the SenderBase name was changed to SensorBase. Both terms are used in the various Cisco documentations, depending on the source.

https://t.me/learningnets

IPS Processing Flow with Global Correlation and Reputation Filter

This topic describes the traffic processing flow in the Cisco Intrusion Prevention System sensor Global Correlation and Reputation Filter active.

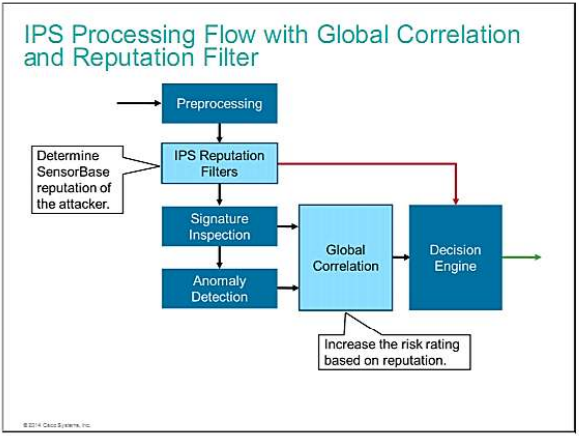

This figure shows the traffic processing flow inside a Cisco Intrusion Prevention System sensor when global correlation and reputation-based filtering are active. The Cisco IPS reputation filters are applied to traffic flows early in the packet processing pipeline, and determine the reputation score for all sources (IP addresses) of traffic flowing through the sensor.

In this early stage, the Cisco IPS reputation filters can already drop traffic from hosts that are on the Cisco SensorBase blacklist, that is, hosts that are confirmed sources of current significant malicious activity. The reputation filters can also create a state inside the sensor that makes the sensor aware of the negative reputation of the source, without performing any immediate aggressive action.

When a Cisco IPS signature is activated, its risk rating is automatically adjusted based on the reputation of the attacker IP address. This adjustment may increase the probability that malicious activity will be prevented, when using event action overrides to assign aggressive actions based on a risk rating threshold.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Reputation Filter Operations

This topic describes the reputation filters operations.

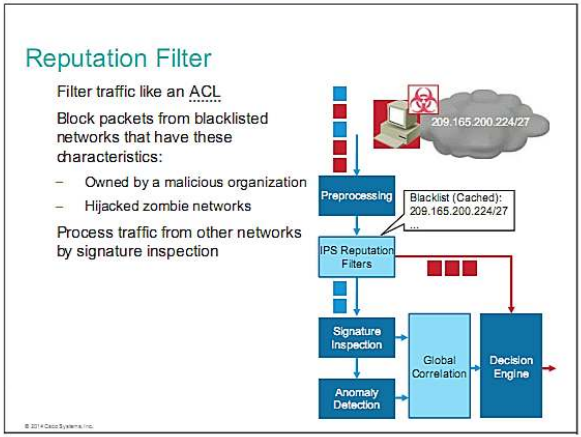

Reputation filters are a traffic-filtering mechanism that denies hosts that are on the Cisco SensorBase blacklist from forwarding traffic through the sensor. Some networks on the Internet are owned by malicious organizations or are hijacked "zombie" networks. It can be assumed that the traffic from these networks is not legitimate traffic and that it can be safely dropped by Cisco Intrusion Prevention System sensor reputation filters. Reputation filters block traffic like an ACL. In this way, no signature inspection or anomaly detection is needed for this traffic.

If an attack enters a Cisco IPS sensor from a network without a reputation, the traffic is processed by signature inspection engines. If the sensor is configured for network participation, it will send some data about the attacker and actions that are taken by the sensor back to Cisco SensorBase.

In the example, an attacker is trying to establish a session from the network 209.165.200.224/27. This traffic reaches a Cisco IPS sensor with reputation filtering enabled. The Cisco IPS sensor received a threat update from the Cisco SensorBase network. Using that update, the sensor refreshes the list of the blacklisted Internet networks and caches it on the sensor. The 209.165.200.224/27 network is also on the blacklist. The traffic from that network is filtered by the reputation filters and forwarded to the decision engine. Consequently, the traffic is dropped without being processed by signature inspection or anomaly detection. However, traffic from networks that are not listed on the blacklist are not blocked by the reputation filters and are processed by signature inspection engines and anomaly detection (if anomaly detection is enabled).

Global Correlation Operations

This topic describes the global correlation operations.

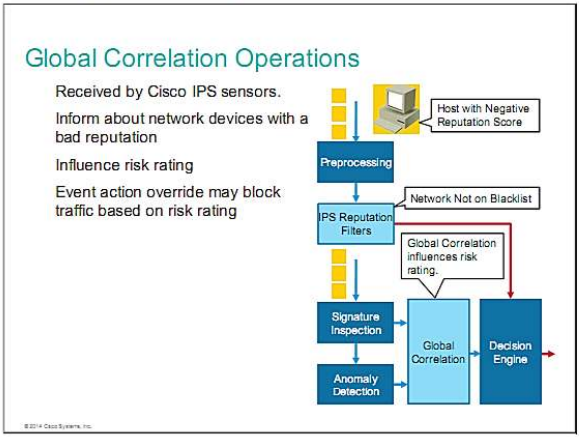

You can configure global correlation so that your Cisco Intrusion Prevention System sensors are aware of network devices with a reputation for malicious activity, and can take action against them. Participating Cisco IPS devices receive and absorb global correlation updates from the Cisco SensorBase. The reputation data that is contained in the global correlation updates is factored into the analysis of network traffic, which increases IPS efficacy, because traffic is denied or allowed based on the reputation of the source IP address.

In the example, traffic from an Internet host with a negative reputation score reaches a Cisco IPS sensor. Its negative reputation score has been determined based on the past network activities of the Internet host and it is kept in the SensorBase. The Cisco IPS sensor has the Global Correlation feature enabled. Its cached global correlation data is regularly updated, so that it already includes the negative reputation score of the Internet host. However, the network of the host is not listed on the reputation filter's blacklist. Therefore, the traffic from the host is processed by signature inspection (and anomaly detection, if enabled). If any Cisco IPS signatures are triggered by inspection, a risk rating is calculated. Based on the negative reputation score for the Internet host, global correlation increases the risk rating. When the risk rating is increased, so is the probability that potential malicious activity is prevented. Based on risk rating thresholds, event action overrides may deploy aggressive actions to block the offending traffic.

After the traffic from the Internet host is processed, the new information on this host is sent back to Cisco Security Intelligence Operations (if network participation is enabled), so that the reputation for that particular IP address can potentially be updated. For example, if there is a lot of offending traffic from the same host and that information is reported to Cisco SIO, then the reputation score for that network address could be decreased (for example, from –6 to –7).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

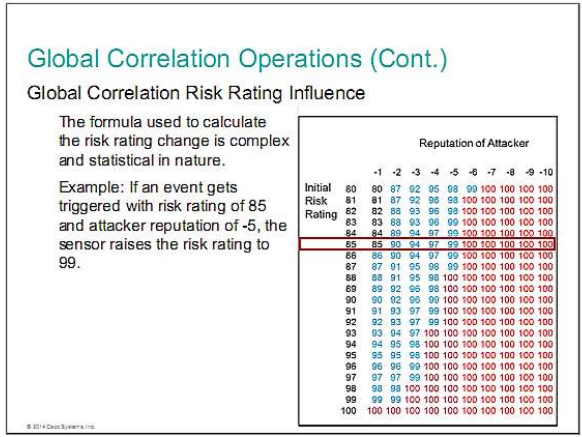

This table shows the influence of attacker reputation to event risk rating. The actual formula that is used to calculate risk rating increase is complex, relies on statistical information, and may change with time. In this example, an event with a calculated risk rating of 85 has triggered, and the sensor has determined (using a query to the SensorBase network) that the attacker has a reputation of -5. Using a table like the one displayed in this example, the sensor will increase the risk rating of this event to 99, increasing the probability of prevention when using (even the default setting of) event action overrides.

Dynamic Updates from Cisco SensorBase

This topic illustrates how the sensors receive updates from Cisco SensorBase.

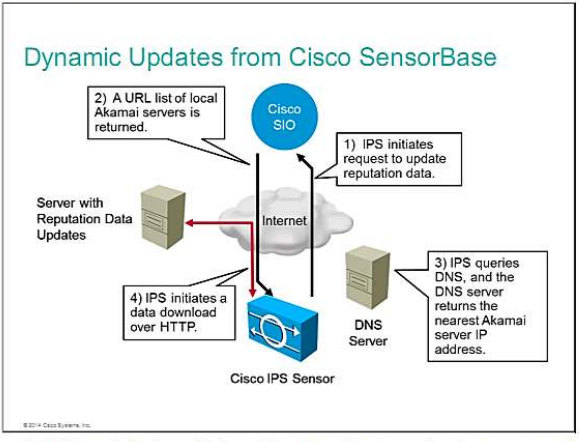

When Cisco Global Correlation is enabled on a Cisco Intrusion Prevention System sensor, the sensor tries to periodically update the reputation data. The procedure is as follows:

1. The sensor initiates a request to update the reputation data. The sensor connects to the following URI: https://update-manifest.ironport.com at Cisco Security Intelligence Operations. DNS server IP addresses must be configured on the Cisco IPS sensor for proper address resolution.
2. A URL list of local Akamai servers is returned to the sensor.
3. The Cisco IPS sensor queries the configured DNS server or servers and the DNS server returns the nearest Akamai server IP address.
4. The IPS sensor initiates a reputation update data download over HTTP. The data is transferred to the IPS sensor in chunks.

Reputation data is received in the form of multiple files that are downloaded as needed during the updates. When the data transfer is finished, the sensor uses the updated reputation data.

The sensor does not try to query the update server at the time when the attack happens. Instead, at the time of the attack, the sensor uses the reputation data that is currently cached on the sensor. The sensor uses the described procedure to update its reputation data at regular intervals.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IPS Sensor Network Participation to Cisco SensorBase

This topic describes how the Cisco Intrusion Prevention System sensors sends information to Cisco SensorBase using network participation.

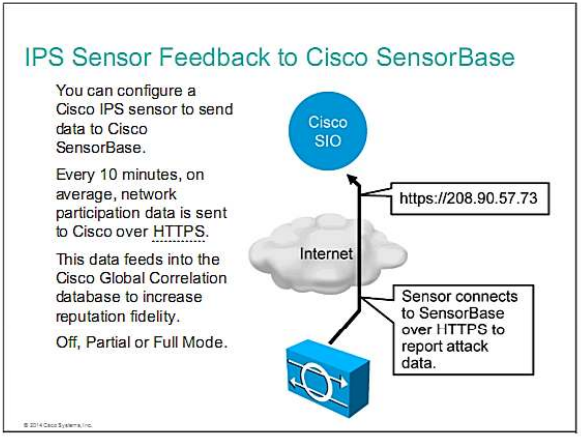

You can configure a Cisco IPS sensor to send data to Cisco SensorBase. Every 10 minutes, on average, network participation data is sent to Cisco over HTTPS. This data feeds into the global correlation database to increase reputation fidelity. Communication between sensors and the SensorBase involves an HTTPS request and response over TCP/IP. The data that is sent includes traffic header information, attacker IP address and port, signature ID, some information about sensor configuration, and so on. No data content or packet content is sent to Cisco SensorBase. This data does not include private addresses.

The scope on the how sensor will contribute to the SensorBase can be off, partial or full mode.

Global Correlation Configurations

This topic describes the global correlation inspection and reputation-based filtering configurations.

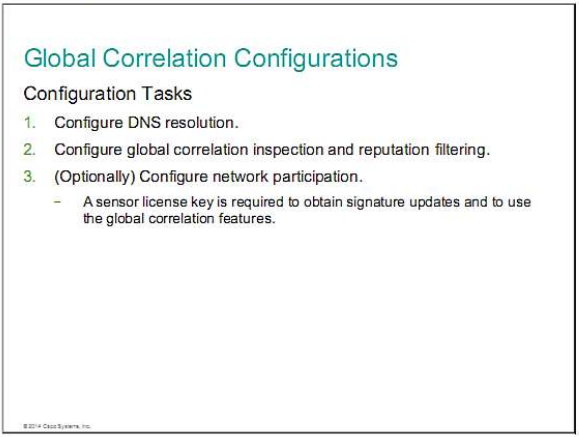

You need to perform the following tasks to enable and configure global correlation and reputation filtering:

1. You must configure DNS resolution on the Cisco Intrusion Prevention System sensor, because the sensor uses DNS to query the SensorBase network.
2. You must enable global correlation and reputation filtering and configure the actions that will be taken when attackers are on the blacklist, or when risk rating should be adjusted based on attacker reputation.
3. Optionally, you can configure the IPS sensor to provide feedback to the SensorBase network, increasing its global awareness.

Note Although the IPS sensor can function without a valid IPS license, you must have a valid IPS license to obtain signature updates and use the global correlation features. You can view the status of the IPS sensor license in the Cisco IPS Device Manager Home page Dashboard using the Licensing gadget or from the IDM Licensing pane (**Configuration > Sensor Management > Licensing**).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

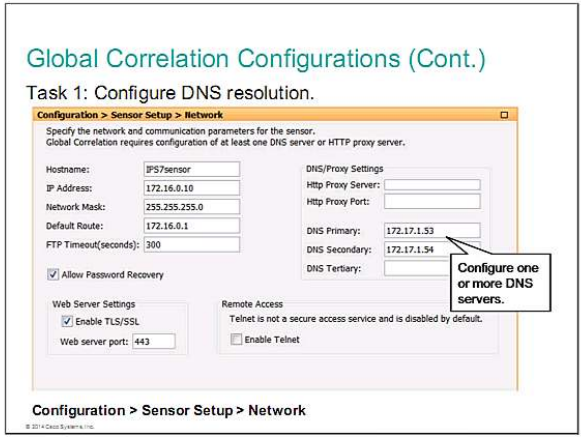

To configure one or more DNS servers, use the Network panel, which you can access by choosing **Configuration > Sensor Setup > Network**. The following relevant fields and check boxes are available inside the DNS/Proxy Settings section of the Network panel:

DNS Primary: Lets you enter the primary DNS server IP address.

DNS Secondary: Lets you enter the secondary DNS server IP address.

DNS Tertiary: Lets you enter tertiary DNS server IP address.

HTTP Proxy Server: Lets you enter a proxy server IP address. You may need proxy servers to download global correlation updates if your network uses proxy.

HTTP Proxy Port: Lets you enter the port number for the proxy server.

You can also use an HTTP proxy to access the SensorBase network, if DNS is not available to the sensor.

If you are using a DNS server, you must configure at least one DNS server and it must be reachable for global correlation updates to be successful. You can configure other DNS servers as backup servers. DNS queries are sent to the first server in the list. If it is unreachable, DNS queries are sent to the next configured DNS server.

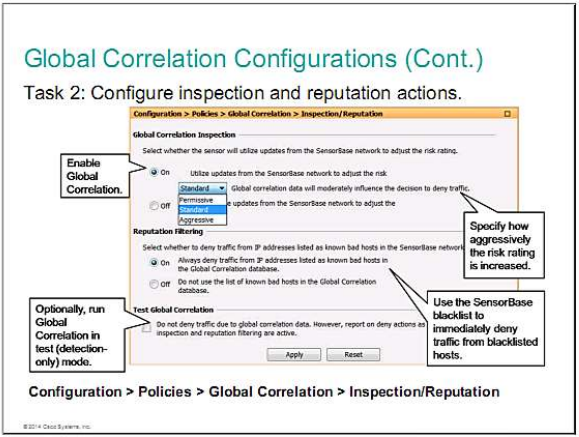

In Task 2 of this configuration sequence, you will enable global correlation or reputation filtering, and configure the actions resulting from reputation-based analysis on this IPS sensor.

To complete this task, perform the following steps using Cisco IDM:

1. Navigate to **Configuration > Policies > Global Correlation > Inspection/Reputation**.
2. In this pane, the Global Correlation Inspection section lets you turn global correlation to **On** or **Off**. Enable global correlation by selecting the **On** radio button. The default is **On**.
3. If using Global Correlation, choose how aggressively the sensor should adjust its risk rating values when correlating them to attacker reputation. You have the following options in the Global Correlation Inspection section:
   - **Permissive**: When using this setting, the sensor will mildly influence the risk rating value using its statistical calculation formula.
   - **Standard**: When using this setting, the sensor will moderately influence the risk rating value using its statistical calculation formula. This is the default.
   - **Aggressive**: When using this setting, the sensor will significantly influence the risk rating value using its statistical calculation formula.

Note	It is recommended that you use the Global Correlation feature in its nonaggressive (test) mode for a while to determine the most suitable option for your environment.

4. In this pane, in the Reputation Filtering section, you can set reputation filtering to **On** or **Off**. Setting it to **On** will deny access to malicious hosts that are listed on the SensorBase blacklist. Reputation filtering is turned **On** by default.
5. Optionally, you can check the check box in the Test Global Correlation section to enable the nonaggressive test mode for both the Global Correlation and the Reputation Filtering features. This test mode allows you to try the global correlation features without actually denying any hosts but enables the reporting of deny actions that are affected by global correlation.
6. Click **Apply** to apply your settings and save your configuration to the sensor.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

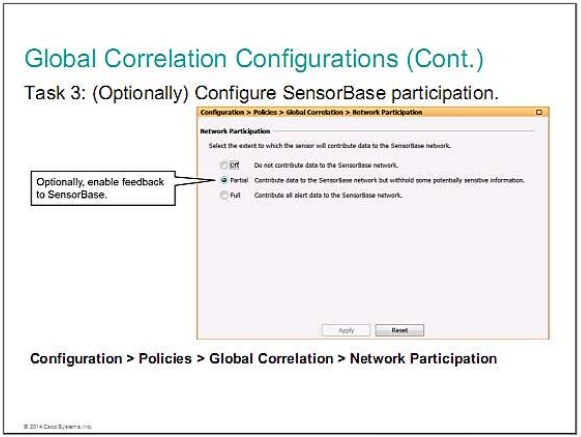

In Task 3, you will optionally configure the sensor to send summary data and feedback to the SensorBase network. If you agree to participate in the SensorBase network, Cisco will collect aggregated statistics about traffic that is sent to your IPS. This includes summary data on the Cisco IPS network traffic properties and how this traffic was managed by the Cisco appliances. Cisco does not collect the data content of traffic or other confidential business or personal information. All data is aggregated and sent by secure HTTPS to the Cisco SensorBase network servers in periodic intervals. All data that is shared with Cisco will be anonymous and treated as strictly confidential.

To complete this task, perform the following steps using Cisco IDM:

1. Navigate to **Configuration > Policies > Global Correlation > Network Participation**.
2. You have the following options in the Network Participation section:

   Off: When this setting is enabled, the sensor will not send any data to the SensorBase network. This is the default.

   Partial: When this setting is enabled, the sensor will send some data to the SensorBase network, but will sanitize data that could be potentially sensitive.

   Full: When this setting is enabled, the sensor will send all alert data to the SensorBase network.
3. Click **Apply** to apply your settings to the sensor. The Network Participation Disclaimer screen will appear.

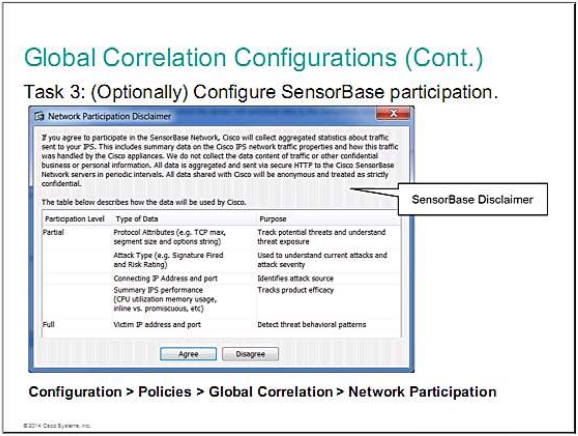

4. This figure shows the contents of the data that is sent to the SensorBase network, and requires you to agree to such participation. If you agree with the configured level of participation, click **Agree**. If you disagree with the configured level of participation, click **Disagree**, and the network participation settings will be reverted to their previous setting.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Global Correlation and Reputation Filter

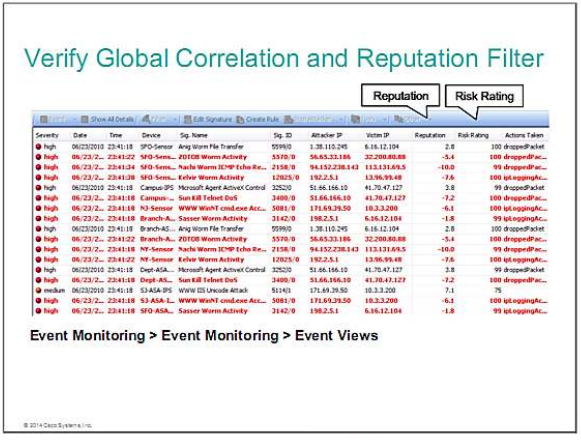

You can observe reputation information for an attacker and risk rating that is calculated for an event. For example, when using Cisco Intercompany Media Engine, navigate to **Event Monitoring > Event Monitoring > Event Views**. In the View Settings, make sure that you select the **Reputation** and **Risk Rating** fields to be displayed.

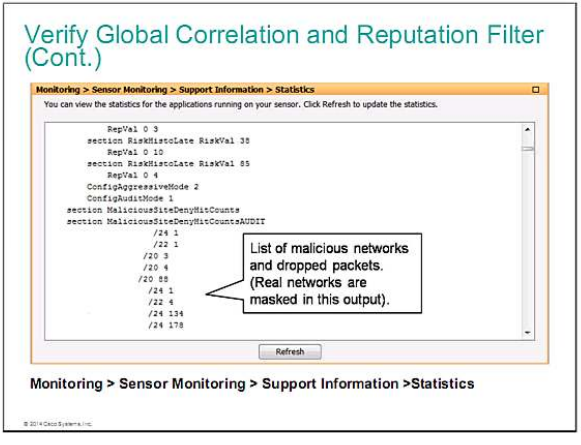

https://t.me/learningnets

You can verify the operation of Reputation Filtering and observe the amount of packets that are dropped due to reputation filters. Inside Cisco Intrusion Prevention System Device Manager, navigate to **Monitoring > Sensor Monitoring > Support Information > Statistics**, and then scroll down to the MaliciousSiteDenyHitCounts section. In this section, you can observe the number of packets that are denied from specific networks. In this output, specific networks have been erased from the figure.

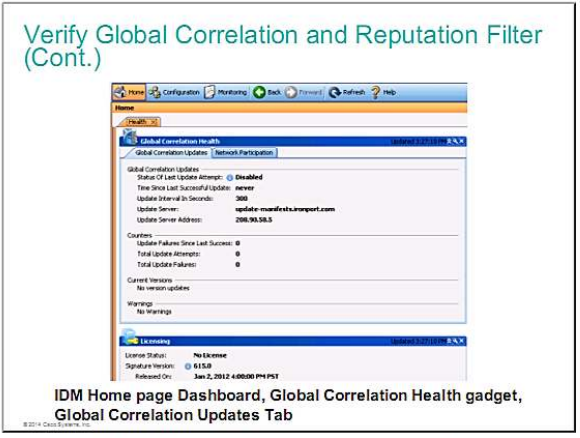

You can view the status of the **Global Correlation Updates** and **Network Participation** in the IDM Home page Dashboard using the **Global Correlation Health** gadget.

Click the **Global Correlation Updates** tab in the **Global Correlation Health** gadget to display the **Global Correlation Updates** status as shown in the figure.

In the example shown in the figure, the IPS sensor does not have a valid IPS license installed, therefore, Global Correlation is disabled.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Global Correlation and Reputation Filter (Cont.)

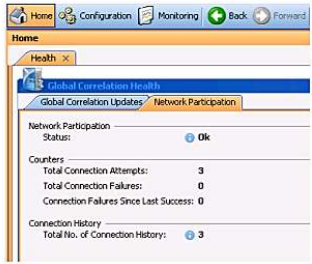

IDM Home page Dashboard, Global Correlation Health gadget, Network Participation Tab

Click the **Network Participation** tab in the **Global Correlation Health** gadget to display the **Network Participation** status as shown in the figure.

In the example shown in the figure, there were three successful connections for sending the IPS sensor network participation data to Cisco.

There is also a **Global Correlation Report** gadget (not shown in figure) that displays the following information:

Packets Denied Due to **Global Correlation—**Displays the percentage of malicious packets identified and whether any have been dropped due to global correlation.

Total **Packets Denied—**Displays the total number of malicious packets that were identified and which ones were dropped because of global correlation criteria.

Summary

This topic summarizes the key points that were discussed in this lesson.

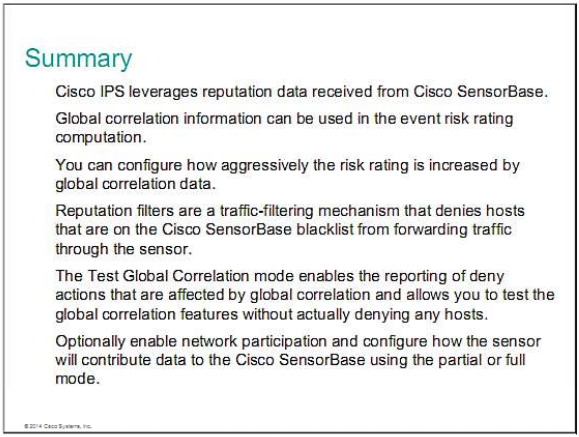

https://t.me/learningnets

Module Summary

Overview

This topic summarizes the key points that were discussed in this module.

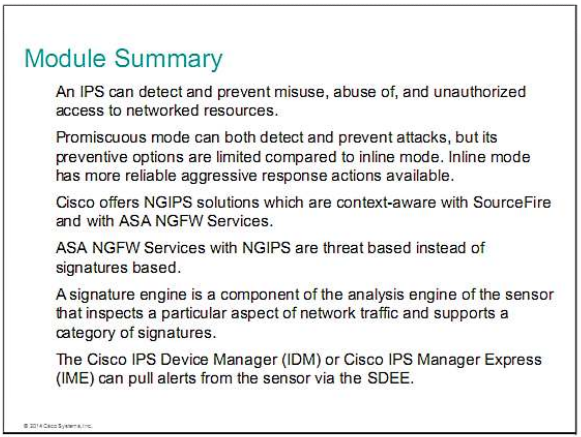

https://t.me/learningnets

Module Summary (Cont.)

By default, each Cisco IPS sensor includes a single virtual sensor named vs0, which in turn includes a signature policy named sig0, an event action policy named event0, and an anomaly detection policy named ad0. This virtual sensor cannot be deleted.

The risk rating can be between 0 and 100 which the sensor dynamically calculates after each triggering of a signature. The higher the value, the greater the security risk of the trigger event.

The threat rating indicates the severity of an event in the local environment by also taking into account any preventive actions that the sensor took to reduce the risk for this event.

Event action overrides are a configuration tool that allows to you add one or more actions to all active signatures based on the calculated risk rating of each event.

Event action filters are a configuration tool that allows to you remove one or more actions from all active signatures based on various criteria.

©2014 Cisco Systems, Inc.

Module Summary (Cont.)

IPS signature threat Profiles provides time savings in tuning and managing the IPS solution. Threat profiles offer guidance of Cisco best practices for different deployments so you can deploy and tune Cisco IPS sensors faster.

The fundamental reason for false positive events is signatures that are very general and that match legitimate traffic in addition to network attacks. Because management of false positives requires you to first separate false positives from true positives in your environment, and then change the manner in which these signatures operate, it cannot be automated.

When you need to address a threat that is uniquely specific to your environment, or otherwise do not have an appropriate signature in the default signature set to address a particular threat, you can create custom signatures on the Cisco IPS sensor.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

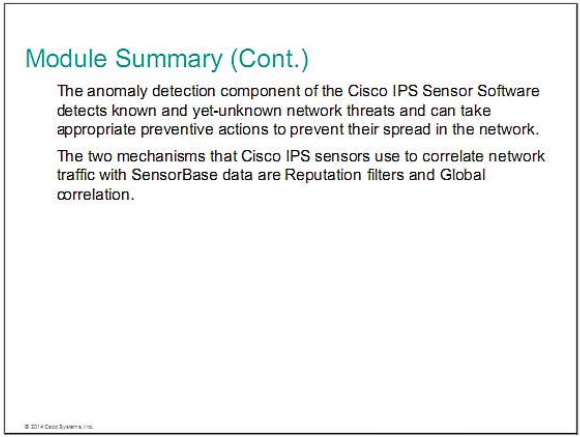

References

For additional information refer to this resource:

Latest Cisco Intrusion Prevention System Sensor Configuration Guide

User Guide for ASA CX and Cisco Prime Security Manager 9.2

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Self-Check

Questions

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

1. Which Cisco IPS appliance deployment mode is where the sensing interface acts as an 802.1Q trunk port, and the sensor acts as a bump in the wire performing VLAN translation between pairs of VLANs on this trunk interface? (Source: Integrating Cisco IPS Sensor into a Network)

   A. Inline Interface pair mode
   B. Inline VLAN group mode
   C. Inline VLAN pair mode
   D. Promiscuous mode

2. How is the Threat Rating calculated? (Source: Configuring Basic Cisco IPS Settings)

   A. Threat rating = risk rating – signature fidelity rating
   B. Threat rating = risk rating – threat rating adjustment
   C. Threat rating = risk rating – attack severity rating
   D. Threat rating = risk rating – attack relevancy rating

3. Which configuration tool that allows to you add one or more actions to all active signatures based on the calculated risk rating of each event? (Source: Configuring Basic Cisco IPS Settings)

   A. Event action overrides
   B. Event action filters
   C. Threat ratings
   D. Anomaly detections
   E. Global correlations

https://t.me/learningnets

4. Which signature engine is best used for packet header or content matching? (Source: Configuring Custom Cisco IPS Signatures)

   A. String
   B. Atomic
   C. Service
   D. Flood
   E. Sweep

5. Which Cisco IPS feature monitors the network and constantly looks for worms and scanners?

   A. Global correlations
   B. Botnet traffic filter
   C. Threat detection
   D. Anomaly detection

6. Which two Cisco IPS features use the reputation data from the Cisco SensorBase? (Choose two.)

   A. Reputation filters
   B. Anomaly detections
   C. Botnet Traffic filter
   D. Global correlation

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Answer Key

1. C
2. B
3. A
4. B
5. D
6. A, D

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Glossary

Term	Definition
ACK	acknowledgment. Notification sent from one network device to another to acknowledge that some event occurred (for example, the receipt of a message).
ACL	access control list. A list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router).
AD	advertised distance.
AIC	Application Inspection and Control.
ARP	Address Resolution Protocol. Internet protocol that is used to map an IP address to a MAC address. Defined in RFC 826.
ARR	Attack Relevancy Rating
ASCII	American Standard Code for Information Interchange. 8-bit code for character representation (7 bits plus parity).
ASR	attack severity rating.
BSD	Berkeley Standard Distribution. A term used to describe any of a variety of UNIX-type operating systems based on the University of California, Berkeley BSD operating system.
CASE	Context Adaptive Scanning Engine
CLI	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.
DDoS	distributed denial of service
DLP	data loss prevention.
DMZ	demilitarized zone.
DN	distinguished name. Global, authoritative name of an entry in the OSI Directory (X.500).
DNS	Domain Name System. System used on the Internet for translating names of network nodes into addresses.
DoS	denial of service. An intentional or unintentional attack on a device that makes the resource unavailable to perform its normal function.
DSCP	differentiated services code point.
ECLB	EtherChannel load balancing.
EHLO	Extended HELLO

https://t.me/learningnets

Term	Definition
ESMTP	Extended Simple Mail Transfer Protocol. Extended version of the Simple Mail Transfer Protocol (SMTP), which includes additional functionality, such as delivery notification and session delivery. ESMTP is described in RFC 1869, *SMTP Service Extensions*.
eStreamer	The Sourcefire Event Streamer is also known as eStreamer
Ethernet	Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mb/s. Ethernet is similar to the IEEE 802.3 series of standards.
FQDN	fully qualified domain name. FQDN is the full name of a system, rather than just its host name. For example, aldebaran is a host name, and aldebaran.interop.com is an FQDN.
FSPAN	Flow-based SPAN
GC	Global Correlation
GET	Group Encrypted Transport.
Gigabit Ethernet	Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.
GLBA	Gramm-Leach-Bliley Act.
GUI	graphical user interface. A user environment that uses pictorial as well as textual representations of the input and the output of applications and the hierarchical or other data structure in which information is stored. Such conventions as buttons, icons, and windows are typical, and many actions are performed using a pointing device (such as a mouse). Microsoft Windows and the Apple Macintosh are prominent examples of platforms using a GUI.
HAT	Host Access Table
HHS	Health and Human Services
HIPAA	Health Insurance Portability and Accountability Act.
HTTP	Hypertext Transfer Protocol. The protocol used by Web browsers and Web servers to transfer files, such as text and graphic files.
HTTPS	Hypertext Transfer Protocol Secure
ICMP	Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792.
ICS	Industrial control system
IDS	intrusion detection system.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
IMAP	Internet Message Access Protocol. Method of accessing email or bulletin board messages kept on a mail server that can be shared. IMAP permits client email applications to access remote message stores as if they were local without actually transferring the message.
IME	Cisco Intercompany Media Engine.
IP address	32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address.
IPS	intrusion prevention system.
IPsec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.
IPv4	IP version 4. Internet Protocol version 4 is the fourth version in the development of the IP and the first version of the protocol to be widely deployed. Along with IPv6, IPv4 is at the core of standards-based internetworking methods of the Internet. IPv4 is still used to route most traffic across the Internet. IPv4 is a connectionless protocol for use on packet-switched link layer networks (for example, Ethernet). It operates on a best-effort delivery model in that it does not guarantee delivery and does not assure proper sequencing or avoidance of duplicate delivery.
IPv6	IP version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation).
ISR	integrated services router. An ISR specifies the elements to guarantee quality of service (QoS) on networks. For example, an ISR can be used to allow video and sound to reach the receiver without interruption. Every application that requires some kind of guarantee has to make an individual reservation.
JDBC	Java Database Connectivity - An interface for direct querying of the SourceFire Defense Center database by third-party reporting and database query tools, such as Crystal Reports.

https://t.me/learningnets

Term	Definition
LAN	local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
LDAP	Lightweight Directory Access Protocol. Protocol that provides access for management and browser applications that provide read/write interactive access to the X.500 Directory.
MAPI	Messaging Application Programming Interface
MIME	Multipurpose Internet Mail Extension. Standard for transmitting nontext data (or data that cannot be represented in plain ASCII code) in Internet mail, such as binary, foreign language text (such as Russian or Chinese), audio, or video data. MIME is defined in RFC 2045.
MPF	Modular Policy Framework.
MTA	Mail Transfer Agent
MUA	Mail User Agents
MX	Mail Exchange
NAT	Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. Also known as Network Address Translator.
NBT	NetBIOS over TCP/IP
NetBIOS	Network Basic Input/Output System. API used by applications on an IBM LAN to request services from lower-level network processes. These services might include session establishment and termination, and information transfer.
NGFW	next generation firewall
NGIPS	next generation intrusion prevention systems
NME-IPS	Network Module Enhanced Intrusion Prevention System
NMS	network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer, such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources.
NSS	NSS Labs is a trusted authority in independent product testing, especially when it comes to research and analysis on Intrusion Prevention Systems

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
NTP	Network Time Protocol. Protocol built on top of TCP that ensures accurate local timekeeping with reference to radio and atomic clocks located on the Internet. This protocol is capable of synchronizing distributed clocks within milliseconds over long time periods.
OOB	out-of-band.
OPSEC	Open Platform for Security
OPSEC API	Open Platform for Security Application Programming Interface. Once Sourcefire NGIPS has identified an out of compliance system, it's necessary for the security team to respond and resolve the issue. Manually responding to the myriad of these issues in the typical network can cause a significant drain on staff. Users can automate many of these activities using the Remediation and OPSEC APIs supported by Sourcefire NGIPS. The APIs are highly flexible and support a range of possible responses.
OSI	Open Systems Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.
PCI	Payment Card Industry
PCI DSS	Payment Card Industry Data Security Standard
PCS	process control system
PD	Promiscuous Delta
PLC	programmable logic controller
POP	Post Office Protocol. Protocol that client email applications use to retrieve mail from a mail server.
POST	power-on self test. Set of hardware diagnostics that runs on a hardware device when this device is powered on.
PRSM	Cisco Prime Security Manager.
PTR	pointer record.
RAT	Recipient Access table
RFC	Request for Comments. Document series that is used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some RFCs are humorous or historical. RFCs are available online from numerous sources.
RSA	Acronym stands for Rivest, Shamir, and Adleman, the inventors of the technique. Public-key cryptographic system that can be used for encryption and authentication.
RSPAN	Remote SPAN
RSS	Really Simple Syndication

https://t.me/learningnets

Term	Definition
SaaS	software as a service.
SBRS	SenderBase Reputation Scores
SCADA	supervisory control and data acquisition
SDEE	Security Device Event Exchange
SFR	signature fidelity rating.
SMB	Server Message Block
SMTP	Simple Mail Transfer Protocol. Internet protocol providing email services.
SMTPD	SMTP daemon
SNMP	Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.
SNORT	Snort is a free and open source network intrusion prevention system created by Martin Roesch in 1998. Snort is now developed by Sourcefire. Martin Roesch founded Sourcefire in 2001 where he was Chief Technology Officer (CTO) and a member of its Board of Directors. He is now vice president and chief architect for Cisco's Security Business Group.
SPAN	Switched Port Analyzer. SPAN is a feature that is available on switches based on Cisco IOS and NX-OS Software that allows traffic received on a port or VLAN to be copied to another port for analysis. It is also referred to as "port mirroring."
SPF	shortest path first. Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms. Sometimes called Dijkstra's algorithm.
SQL	Structured Query Language. International standard language for defining and accessing relational databases.
SSP	Security Services Processor.
SYN	synchronization.
syslog	system logging.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.
TCP/IP	Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite.
TLS	Transport Layer Security. A future IETF protocol to replace SSL.
TVR	target value rating.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
URI	uniform resource identifier. Type of formatted identifier that encapsulates the name of an Internet object, and labels it with an identification of the name space, thus producing a member of the universal set of names in registered name spaces and of addresses referring to registered protocols or name spaces. [RFC 1630]
URL	uniform resource locator. Type of formatted identifier that describes the access method and the location of an information resource object on the Internet. [RFC 1738]
URL	uniform resource locator. Type of formatted identifier that describes the access method and the location of an information resource object on the Internet. [RFC 1738]
VACL	VLAN access control list.
VID	VLAN ID
VLAN	virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
WBRS	Web-Based Reputation Score.
WLAN	wireless LAN. A LAN is a high-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
WLR	watch list rating

https://t.me/learningnets