



Cisco TrustSec Configuration Guide, Cisco IOS XE Cupertino 17.7.x (Catalyst 9300 Switches)

First Published: 2021-12-07

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

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: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. 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. (1721R)

© 2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Cisco TrustSec Overview	1
Restrictions for Cisco TrustSec	1
Information About Cisco TrustSec Architecture	1
Authentication	4
Cisco TrustSec and Authentication	4
Cisco TrustSec Enhancements to EAP-FAST	5
802.1X Role Selection	6
Cisco TrustSec Authentication Summary	6
Device Identities	6
Device Credentials	7
User Credentials	7
Security Group-Based Access Control	7
Security Groups and SGTs	7
Security Group ACL Support	8
SGACL Policies	8
Ingress Tagging and Egress Enforcement	9
Determining the Source Security Group	10
Determining the Destination Security Group	11
SGACL Enforcement on Routed and Switched Traffic	11
SGACL Logging and ACE Statistics	11
VRF-aware SGACL Logging	12
SGACL Monitor Mode	12
Authorization and Policy Acquisition	13
Environment Data Download	14
RADIUS Relay Functionality	14
Link Security	15

Configuring SAP-PMK for Link Security 15

SXP for SGT Propagation Across Legacy Access Networks 16

Layer 3 SGT Transport for Spanning Non-TrustSec Regions 18

VRF-Aware SXP 19

 Layer 2 VRF-Aware SXP and VRF Assignment 19

Feature History for Cisco TrustSec Overview 19

CHAPTER 2

SGACL and Environment Data Download over REST 21

Prerequisites for SGACL and Environment Data Download over REST 21

Restrictions for SGACL and Environment Data Download over REST 21

Information About SGACL and Environment Data Download over REST 22

 SGACL and Environment Data Download over REST Overview 22

 Cisco TrustSec Environment Data 22

 Message Flow Between a Network Device and a Server 23

 Policy Server Selection Criteria 25

 Server and IP Address Selection Process 25

 Server Liveliness Check 26

How to Configure SGACL and Environment Data Download over REST 27

 Configuring the Username and Password 27

 Configuring Certificate Enrollment 28

 Downloading Cisco TrustSec Policies 28

 Downloading Environment Data 30

Verifying the SGACL and Environment Data Download over REST 31

Debugging the SGACL and Environment Data over REST Configuration 32

Configuration Examples for SGACL and Environment Data Download over REST 32

 Example: Configuring the Username and Password 32

 Example: Downloading Cisco TrustSec Policies 33

 Example: Downloading Environment Data 33

Feature History for SGACL and Environment Data Download over REST 33

CHAPTER 3

Configuring Security Group ACL Policies 35

Restrictions for Configuring Security Group ACL Policies 35

Information About Security Group ACL Policies 36

 SGACL Logging 36

How to Configure Security Group ACL Policies	36
SGACL Policy Configuration Process	36
Enabling SGACL Policy Enforcement Globally	37
Enabling SGACL Policy Enforcement Per Interface	37
Enabling SGACL Policy Enforcement on VLANs	38
Configuring SGACL Monitor Mode	39
Manually Configuring SGACL Policies	40
Configuring and Applying IPv4 SGACL Policies	40
Configuring IPv6 SGACL Policies	42
Manually Applying SGACL Policies	43
Displaying SGACL Policies	43
Refreshing the Downloaded SGACL Policies	45
Configuration Examples for Security Group ACL Policies	45
Example: Enabling SGACL Policy Enforcement Globally	45
Example: Enabling SGACL Policy Enforcement Per Interface	45
Example: Enabling SGACL Policy Enforcement on VLANs	46
Example: Configuring SGACL Monitor Mode	46
Example: Manually Configuring SGACL Policies	46
Example: Manually Applying SGACLs	47
Example: Displaying SGACL Policies	47
Feature History for Security Group ACL Policies	47

CHAPTER 4

Cisco TrustSec SGACL High Availability	49
Prerequisites for Cisco TrustSec SGACL High Availability	49
Restrictions for Cisco TrustSec SGACL High Availability	49
Information About Cisco TrustSec SGACL High Availability	49
Verifying Cisco TrustSec SGACL High Availability	50
Feature History for SGACL High Availability	52

CHAPTER 5

Configuring SGT Exchange Protocol	53
Prerequisites for SGT Exchange Protocol	53
Restrictions for SGT Exchange Protocol	54
Information About SGT Exchange Protocol	54
SGT Exchange Protocol Overview	54

Security Group Tagging	54
SGT Assignment	55
How to Configure SGT Exchange Protocol	55
Configuring a Device SGT Manually	55
Configuring an SXP Peer Connection	56
Configuring the Default SXP Password	57
Configuring the Default SXP Source IP Address	58
Changing the SXP Reconciliation Period	58
Changing the SXP Retry Period	59
Creating Syslogs to Capture Changes of IP Address-to-SGT Mapping Learned Through SXP	60
Configuration Examples for SGT Exchange Protocol	60
Example: Enabling Cisco TrustSec SXP and an SXP Peer Connection	60
Example: Configuring the Default SXP Password and Source IP Address	61
Verifying SGT Exchange Protocol Connections	61
Feature History for SGT Exchange Protocol	62

CHAPTER 6

Configuring Security Group Tag Mapping	63
Restrictions for SGT Mapping	63
Information About SGT Mapping	63
Overview of Subnet-to-SGT Mapping	64
Overview of VLAN-to-SGT Mapping	64
Overview of Layer 3 Logical Interface-to-SGT Mapping (L3IF-SGT Mapping)	65
Binding Source Priorities	65
Default Route SGT	65
How to Configure SGT Mapping	66
Configuring a Device SGT Manually	66
Configuring Subnet-to-SGT Mapping	66
Configuring VLAN-to-SGT Mapping	68
Configuring L3IF-to-SGT Mapping	70
Emulating the Hardware Keystore	71
Configuring Default Route SGT	72
Verifying SGT Mapping	72
Verifying Subnet-to-SGT Mapping Configuration	73
Verifying VLAN-to-SGT Mapping	73

Verifying L3IF-to-SGT Mapping	73
Verifying Default Route SGT Configuration	73
Configuration Examples for SGT Mapping	74
Example: Configuring a Device SGT Manually	74
Example: Configuration for Subnet-to-SGT Mapping	74
Example: Configuration for VLAN-to-SGT Mapping for a Single Host Over an Access Link	75
Example: Configuration for L3IF-to-SGT Mapping on an Ingress Port	76
Example: Emulating the Hardware Keystore	77
Example: Configuring Device Route SGT	77
Feature History for Security Group Tag Mapping	77

CHAPTER 7**Cisco TrustSec VRF-Aware SGT 79**

VRF-Aware SXP	79
How to Configure Cisco TrustSec VRF-Aware SGT	79
Configuring VRF-to-Layer-2-VLAN Assignments	80
Configuring VRF-to-SGT Mapping	81
Configuration Examples for Cisco TrustSec VRF-Aware SGT	81
Example: Configuring VRF-to-Layer2-VLAN Assignments	81
Example: Configuring VRF-to-SGT Mapping	81
Feature History for Cisco TrustSec VRF-Aware SGT	82

CHAPTER 8**IP-Prefix and SGT-Based SXP Filtering 83**

Restrictions for IP-Prefix and Security Group Tag (SGT)-Based Security Exchange Protocol (SXP) Filtering	83
Information About IP-Prefix and SGT-Based SXP Filtering	84
How to Configure IP-Prefix and SGT-Based SXP Filtering	84
Configuring an SXP Filter List	85
Configuring an SXP Filter Group	85
Configuring a Global Listener or Speaker Filter Group	86
Enabling SXP Filtering	87
Configuring the Default or Catch-All Rule	87
Configuration Examples for IP-Prefix and SGT-Based SXP Filtering	88
Example: Configuring an SXP Filter List	88
Example: Configuring an SXP Filter Group	88

Example: Enabling SXP Filtering	88
Example: Configuring the Default or Catch-All Rule	89
Verifying IP-Prefix and SGT-Based SXP Filtering	89
Syslog Messages for SXP Filtering	91
Feature History for IP-Prefix and SGT-Based SXP Filtering	92

CHAPTER 9**Flexible NetFlow Export of Cisco TrustSec Fields 93**

Flexible NetFlow Export of Cisco TrustSec Fields	93
Restrictions for Flexible NetFlow Export of Cisco TrustSec Fields	93
Information About Flexible NetFlow Export of Cisco TrustSec Fields	93
Cisco TrustSec Fields in Flexible NetFlow	93
How to Configure Flexible NetFlow Export of Cisco TrustSec Fields	94
Configuring Cisco TrustSec Fields as Key Fields in the Flow Record	94
Configuring SGT Name Export in NetFlow	96
Configuration Examples for Flexible NetFlow Export of Cisco TrustSec Fields	97
Example: Configuring Cisco TrustSec Fields as Key Fields in the Flow Record	97
Example: Configuring SGT Name Export in NetFlow	97
Feature History for Flexible NetFlow Export of Cisco TrustSec Fields	98

CHAPTER 10**Configuring SGT Inline Tagging 99**

Restrictions for SGT Inline Tagging	99
Information About SGT Inline Tagging	99
SGT Inline Tagging on a NAT Enabled Device	100
Configuring SGT Inline Tagging	101
Example: Configuring SGT Static Inline Tagging	103
Feature History for SGT Inline Tagging	103

CHAPTER 11**Configuring Endpoint Admission Control 105**

Information About Endpoint Admission Control	105
Example: 802.1X Authentication Configuration	106
Example: MAC Authentication Bypass Configuration	106
Example: Web Authentication Proxy Configuration	106
Example: Flexible Authentication Sequence and Failover Configuration	107
802.1X Host Modes	107

Pre-Authentication Open Access	107
Example: DHCP Snooping and SGT Assignment	107
Feature History for Endpoint Admission Control	108



CHAPTER 1

Cisco TrustSec Overview

Cisco TrustSec builds secure networks by establishing domains of trusted network devices. Each device in the domain is authenticated by its peers. Communication on the links between devices in the domain is secured with a combination of encryption, message integrity check, and data-path replay protection mechanisms.

- [Restrictions for Cisco TrustSec, on page 1](#)
- [Information About Cisco TrustSec Architecture , on page 1](#)
- [Authentication, on page 4](#)
- [Security Group-Based Access Control, on page 7](#)
- [Authorization and Policy Acquisition, on page 13](#)
- [Environment Data Download, on page 14](#)
- [RADIUS Relay Functionality, on page 14](#)
- [Link Security, on page 15](#)
- [SXP for SGT Propagation Across Legacy Access Networks, on page 16](#)
- [Layer 3 SGT Transport for Spanning Non-TrustSec Regions, on page 18](#)
- [VRF-Aware SXP, on page 19](#)
- [Feature History for Cisco TrustSec Overview, on page 19](#)

Restrictions for Cisco TrustSec

- Protected access credential (PAC) provisioning fails and remains in hung state, when an invalid device ID is specified. Even after clearing the PAC, and configuring the correct device ID and password, PAC still fails.

As a workaround, in the Cisco Identity Services Engine (ISE), uncheck the Suppress Anomalous Clients option in the Administration> System> Settings> Protocols> Radius menu for PAC to work.

- Cisco TrustSec is not supported in FIPS mode.

Information About Cisco TrustSec Architecture

The Cisco TrustSec security architecture builds secure networks by establishing domains of trusted network devices. Each device in the domain is authenticated by its peers. Communication on the links between devices in the domain is secured with a combination of encryption, message integrity check, and data-path replay protection mechanisms. Cisco TrustSec uses the device and user credentials acquired during authentication

for classifying the packets by security groups (SGs) as they enter the network. This packet classification is maintained by tagging packets on ingress to the Cisco TrustSec network so that they can be properly identified for the purpose of applying security and other policy criteria along the data path. The tag, called the security group tag (SGT), allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic.




Note Cisco TrustSec IEEE 802.1X links are not supported on platforms supported in the Cisco IOS XE Denali, Cisco IOS XE Everest, and Cisco IOS XE Fuji releases, and hence only the Authenticator is supported; the Supplicant is not supported.

The Cisco TrustSec architecture incorporates three key components:

- **Authenticated networking infrastructure**—After the first device (called the seed device) authenticates with the authentication server to begin the Cisco TrustSec domain, each new device added to the domain is authenticated by its peer devices already within the domain. The peers act as intermediaries for the domain's authentication server. Each newly-authenticated device is categorized by the authentication server and assigned a security group number based on its identity, role, and security posture.
- **Security group-based access control**—Access policies within the Cisco TrustSec domain are topology-independent, based on the roles (as indicated by security group number) of source and destination devices rather than on network addresses. Individual packets are tagged with the security group number of the source.
- **Secure communication**—With encryption-capable hardware, communication on each link between devices in the domain can be secured with a combination of encryption, message integrity checks, and data-path replay protection mechanisms.

The following figure shows an example of a Cisco TrustSec domain. In this example, several networking devices and an endpoint device are inside the Cisco TrustSec domain. One endpoint device and one networking device are outside the domain because they are not Cisco TrustSec-capable devices or because they have been refused access. The authentication server is considered to be outside of the Cisco TrustSec domain; it is either a Cisco Identities Service Engine (Cisco ISE), or a Cisco Secure Access Control System (Cisco ACS).

Figure 1: Cisco TrustSec Network Domain Example



Each participant in the Cisco TrustSec authentication process acts in one of the following roles:

- **Supplicant**—An unauthenticated device connected to a peer within the Cisco TrustSec domain, and attempting to join the Cisco TrustSec domain.
- **Authentication server**—The server that validates the identity of the supplicant and issues the policies that determine the supplicant's access to services within the Cisco TrustSec domain.
- **Authenticator**—An authenticated device that is already part of the Cisco TrustSec domain and can authenticate new peer supplicants on behalf of the authentication server.

When the link between a supplicant and an authenticator first comes up, the following sequence of events typically occurs:

1. **Authentication (802.1X)**—The supplicant is authenticated by the authentication server, with the authenticator acting as an intermediary. Mutual authentication is performed between the two peers (supplicant and authenticator).
2. **Authorization**—Based on the identity information of the supplicant, the authentication server provides authorization policies, such as security group assignments and ACLs, to each of the linked peers. The authentication server provides the identity of each peer to the other, and each peer then applies the appropriate policy for the link.
3. **Security Association Protocol (SAP) negotiation**—When both sides of a link support encryption, the supplicant and the authenticator negotiate the necessary parameters to establish a security association (SA).



Note SAP is not supported on 100G interfaces. We recommend that you use MACsec Key Agreement protocol (MKA) with extended packet numbering (XPN) on 100G interfaces.

When all three steps are complete, the authenticator changes the state of the link from the unauthorized (blocking) state to the authorized state, and the supplicant becomes a member of the Cisco TrustSec domain.

Cisco TrustSec uses ingress tagging and egress filtering to enforce access control policy in a scalable manner. Packets entering the domain are tagged with a security group tag (SGT) containing the assigned security group number of the source device. This packet classification is maintained along the data path within the Cisco TrustSec domain for the purpose of applying security and other policy criteria. The final Cisco TrustSec device on the data path, either the endpoint or network egress point, enforces an access control policy based on the security group of the Cisco TrustSec source device and the security group of the final Cisco TrustSec device. Unlike traditional access control lists based on network addresses, Cisco TrustSec access control policies are a form of role-based access control lists (RBACLs) called security group access control lists (SGACLs).



Note Ingress refers to packets entering the first Cisco TrustSec-capable device encountered by a packet on its path to the destination and egress refers to packets leaving the last Cisco TrustSec-capable device on the path.


Authentication

Cisco TrustSec and Authentication

Using Network Device Admission Control (NDAC), Cisco TrustSec authenticates a device before allowing it to join the network. NDAC uses 802.1X authentication with Extensible Authentication Protocol Flexible Authentication via Secure Tunnel (EAP-FAST) as the Extensible Authentication Protocol (EAP) method to perform the authentication. EAP-FAST conversations provide for other EAP method exchanges inside the EAP-FAST tunnel using chains. Administrators can use traditional user-authentication methods, such as Microsoft Challenge Handshake Authentication Protocol Version 2 (MSCHAPv2), while still having security provided by the EAP-FAST tunnel. During the EAP-FAST exchange, the authentication server creates and delivers to the supplicant a unique protected access credential (PAC) that contains a shared key and an encrypted token to be used for future secure communications with the authentication server.

The following figure shows the EAP-FAST tunnel and inner methods as used in Cisco TrustSec.

Figure 2: Cisco TrustSec Authentication



Cisco TrustSec Enhancements to EAP-FAST

The implementation of EAP-FAST for Cisco TrustSec has the following enhancements:

- **Authenticate the authenticator**—Securely determines the identity of the authenticator by requiring the authenticator to use its PAC to derive the shared key between itself and the authentication server. This feature also prevents you from configuring RADIUS shared keys on the authentication server for every possible IP address that can be used by the authenticator.
- **Notify each device of the identity of its peer**—By the end of the authentication exchange, the authentication server has identified both the supplicant and the authenticator. The authentication server conveys the identity of the authenticator, and whether the authenticator is Cisco TrustSec-capable, to the supplicant by using additional type-length-value parameters (TLVs) in the protected EAP-FAST termination. The authentication server also conveys the identity of the supplicant, and whether the supplicant is Cisco

TrustSec-capable, to the authenticator by using RADIUS attributes in the Access- Accept message. Because each device knows the identity of its peer, it can send additional RADIUS Access-Requests to the authentication server to acquire the policy to be applied on the link.

802.1X Role Selection

In 802.1X, the authenticator must have IP connectivity with the authentication server because it has to relay the authentication exchange between the supplicant and the authenticator using RADIUS over UDP/IP. When an endpoint device, such as a PC, connects to a network, it is obvious that it should function as a supplicant. However, in the case of a Cisco TrustSec connection between two network devices, the 802.1X role of each network device might not be immediately apparent to the other network device.

Instead of requiring manual configuration of the authenticator and supplicant roles for two adjacent switches, Cisco TrustSec runs a role-selection algorithm to automatically determine which switch functions as the authenticator and which functions as the supplicant. The role-selection algorithm assigns the authenticator role to the switch that has IP reachability to a RADIUS server. Both switches start both the authenticator and supplicant state machines. When a switch detects that its peer has access to a RADIUS server, it terminates its own authenticator state machine and assumes the role of the supplicant. If both switches have access to a RADIUS server, the first switch to receive a response from the RADIUS server becomes the authenticator and the other switch becomes the supplicant.

Cisco TrustSec Authentication Summary

By the end of the Cisco TrustSec authentication process, the authentication server has performed the following actions:

- Verified the identities of the supplicant and the authenticator.
- Authenticated the user if the supplicant is an endpoint device.

At the end of the Cisco TrustSec authentication process, both the authenticator and the supplicant know the following:

- Device ID of the peer
- Cisco TrustSec capability information of the peer
- Key used for the SAP

Device Identities

Cisco TrustSec does not use IP addresses or MAC addresses as device identities. Instead, you assign a name (device ID) to each Cisco TrustSec-capable switch to identify it uniquely in the Cisco TrustSec domain. This device ID is used for the following:

- Looking up the authorization policy
- Looking up passwords in the databases during authentication

Device Credentials

Cisco TrustSec supports password-based credentials. Cisco TrustSec authenticates the supplicants through passwords and uses MSCHAPv2 to provide mutual authentication.

The authentication server uses these credentials to mutually authenticate the supplicant during the EAP-FAST phase 0 (provisioning) exchange where a PAC is provisioned in the supplicant. Cisco TrustSec does not perform the EAP-FAST phase 0 exchange again until the PAC expires, and only performs EAP-FAST phase 1 and phase 2 exchanges for future link bringups. The EAP-FAST phase 1 exchange uses the PAC to mutually authenticate the authentication server and the supplicant. Cisco TrustSec uses the device credentials only during the PAC provisioning (or reprovisioning) steps.

When the supplicant first joins the Cisco TrustSec domain, the authentication server authenticates the supplicant and pushes a shared key and encrypted token to the supplicant with the PAC. The authentication server and the supplicant use this key and token for mutual authentication in all future EAP-FAST phase 0 exchanges.

User Credentials

Cisco TrustSec does not require a specific type of user credential for endpoint devices. You can choose any type of user authentication method that is supported by the authentication server, and use the corresponding credentials. For example, the Cisco Secure Access Control System (ACS) version 5.1 supports MSCHAPv2, generic token card (GTC), or RSA one-time password (OTP).

Security Group-Based Access Control

This section provides information about security group-based access control lists (SGACLs).

Security Groups and SGTs

A security group is a grouping of users, endpoint devices, and resources that share access control policies. Security groups are defined by the administrator in the Cisco ISE or Cisco Secure ACS. As new users and devices are added to the Cisco TrustSec domain, the authentication server assigns these new entities to appropriate security groups. Cisco TrustSec assigns to each security group a unique 16-bit security group number whose scope is global within a Cisco TrustSec domain. The number of security groups in the device is limited to the number of authenticated network entities. You do not have to manually configure security group numbers.

Once a device is authenticated, Cisco TrustSec tags any packet that originates from that device with a security group tag (SGT) that contains the security group number of the device. The packet carries this SGT throughout the network within the Cisco TrustSec header. The SGT is a single label that determines the privileges of the source within the entire enterprise.

Because the SGT contains the security group of the source, the tag can be referred to as the source SGT. The destination device is also assigned to a security group (the destination SG) that can be referred to for simplicity as the destination group tag (DGT), although the actual Cisco TrustSec packet tag does not contain the security group number of the destination device.

Security Group ACL Support

Security group access control lists (SGACLs) is a policy enforcement through which the administrator can control operations performed by an user, based on security group assignments and destination resources. Policy enforcement within the Cisco TrustSec domain is represented by a permissions matrix, with source security group number on one axis and destination security group number on the other axis. Each cell in the matrix contains an ordered list of SGACLs, which specifies permissions that should be applied to packets originating from an IP belonging to a source security group and having a destination IP that belongs to the destination security group.

SGACL provides stateless access control mechanism based on the security association or security group tag value instead of IP addresses and filters. There are three ways to provision an SGACL policy:

- **Static policy provisioning:** The SGACL policies are defined by the user using the command **cts role-based permission**.
- **Dynamic policy provisioning:** Configuration of SGACL policies should be done primarily through the policy management function of the Cisco Secure ACS or the Cisco Identity Services Engine.
- **Change of Authorization (CoA):** The updated policy is downloaded when the SGACL policy is modified on the ISE and CoA is pushed to the Cisco TrustSec device.


The device data plane receives the CoA packets from the policy provider (ISE) and applies the policy to the CoA packets. The packets are then forwarded to the device control plane where the next level of policy enforcement happens for the incoming CoA packets. To view the hardware and software policy counter hit information, run the **show cts role-based counters** command in privileged EXEC mode.

SGACL Policies

Using security group access control lists (SGACLs), you can control the operations that users can perform based on the security group assignments of users and destination resources. Policy enforcement within the Cisco TrustSec domain is represented by a permissions matrix, with source security group numbers on one axis and destination security group numbers on the other axis. Each cell in the body of the matrix can contain an ordered list of SGACLs which specifies the permissions that should be applied to packets originating from the source security group and destined for the destination security group.

The following figure shows an example of a Cisco TrustSec permissions matrix for a simple domain with three defined user roles and one defined destination resource. Three SGACL policies control access to the destination server based on the role of the user.

Figure 3: SGACL Policy Matrix Example



By assigning users and devices within the network to security groups and applying access control between the security groups, Cisco TrustSec achieves role-based topology-independent access control within the network. Because SGACLs define access control policies based on device identities instead of IP addresses as in traditional ACLs, network devices are free to move throughout the network and change IP addresses. As long as the roles and the permissions remain the same, changes to the network topology do not change the security policy. When a user is added to the device, you simply assign the user to an appropriate security group and the user immediately receives the permissions of that group.



Note SGACL policies are applied to traffic that is generated between two host devices, not to traffic that is generated from a device to an end host device.


Using role-based permissions greatly reduces the size of ACLs and simplifies their maintenance. With Cisco TrustSec, the number of access control entries (ACEs) configured is determined by the number of permissions specified, resulting in a much smaller number of ACEs than in a traditional IP network. The use of SGACLs in Cisco TrustSec typically results in a more efficient use of TCAM resources compared with traditional ACLs. A maximum of 5,000 SGACL policies are supported.

Ingress Tagging and Egress Enforcement

Cisco TrustSec access control is implemented using ingress tagging and egress enforcement. At the ingress point to the Cisco TrustSec domain, traffic from the source is tagged with an SGT containing the security group number of the source entity. The SGT is propagated with the traffic across the domain. At the egress point of the Cisco TrustSec domain, an egress device uses the source SGT and the security group number of the destination entity (the destination SG, or DGT) to determine which access policy to apply from the SGACL policy matrix.

The following figure shows how the SGT assignment and the SGACL enforcement operate in a Cisco TrustSec domain.

Figure 4: SGT and SGACL in a Cisco TrustSec Domain



1. The host PC transmits a packet to the web server. Although the PC and the web server are not members of the Cisco TrustSec domain, the data path of the packet includes the Cisco TrustSec domain.
2. The Cisco TrustSec ingress device modifies the packet to add an SGT with security group number 3, the security group number assigned by the authentication server for the host PC.
3. The Cisco TrustSec egress device enforces the SGACL policy that applies to source group 3 and destination group 4, the security group number assigned by the authentication server for the web server.
4. If the SGACL allows the packet to be forwarded, the Cisco TrustSec egress switch modifies the packet to remove the SGT and forwards the packet to the web server.

Determining the Source Security Group

A network device at the ingress of Cisco TrustSec domain must determine the SGT of the packet entering the Cisco TrustSec domain so that it can tag the packet with that SGT when it forwards it into the Cisco TrustSec domain. The egress network device must determine the SGT of the packet in order to apply an SGACL.

The network device can determine the SGT for a packet in one of the following methods:

- Obtain the source SGT during policy acquisition—After the Cisco TrustSec authentication phase, a network device acquires policy information from the authentication server, which indicates whether the peer device is trusted or not. If a peer device is not trusted, then the authentication server can also provide an SGT to apply to all packets coming from the peer device.
- Obtain the source SGT from the packet—If a packet comes from a trusted peer device, the packet carries the SGT. This applies to a network device that is not the first network device in Cisco TrustSec domain for the packet.
- Look up the source SGT based on the source identity—With Identity Port Mapping (IPM), you can manually configure the link with the identity of the connected peer. The network device requests policy information, including SGT and trust state, from the authentication server.
- Look up the source SGT based on the source IP address—In some cases, you can manually configure the policy to decide the SGT of a packet based on its source IP address. The SGT Exchange Protocol (SXP) can also populate the IP-address-to-SGT mapping table.

Determining the Destination Security Group

The egress network device in a Cisco TrustSec domain determines the destination group (DGT) for applying the SGACL. The network device determines the destination security group for the packet using the same methods used for determining the source security group, with the exception of obtaining the group number from a packet tag. The destination security group number is not included in a packet tag.

In some cases, ingress devices or other non-egress devices might have destination group information available. In those cases, SGACLs might be applied in these devices rather than egress devices.

SGACL Enforcement on Routed and Switched Traffic

SGACL enforcement is applied only on IP traffic, but enforcement can be applied to either routed or switched traffic.

For routed traffic, SGACL enforcement is performed by an egress switch, typically a distribution switch or an access switch with a routed port connecting to the destination host. When you enable SGACL enforcement globally, enforcement is automatically enabled on every Layer 3 interface except for SVI interfaces.

For switched traffic, SGACL enforcement is performed on traffic flowing within a single switching domain without any routing function. An example would be SGACL enforcement performed by a data center access switch on server-to-server traffic between two directly connected servers. In this example, the server-to-server traffic would typically be switched. SGACL enforcement can be applied to packets switched within a VLAN or forwarded to an SVI associated with a VLAN, but enforcement must be enabled explicitly for each VLAN.

SGACL Logging and ACE Statistics

When logging is enabled in SGACL, the device logs the following information:

- The source security group tag (SGT) and destination SGT
- The SGACL policy name
- The packet protocol type
- The action performed on the packet

The log option applies to individual ACEs and causes packets that match the ACE to be logged. The first packet logged by the log keyword generates a syslog message. Subsequent log messages are generated and reported at five-minute intervals. If the logging-enabled ACE matches another packet (with characteristics identical to the packet that generated the log message), the number of matched packets is incremented (counters) and then reported.

To enable logging, use the **log** keyword in front of the ACE definition in the SGACL configuration. For example, **permit ip log**.

When SGACL logging is enabled, ICMP Request messages from the device to the client are not logged for IPv4 and IPv6 protocols. However, ICMP Response messages from the client to the device are logged.

The following is a sample log, displaying source and destination SGTs, ACE matches (for a permit or deny action), and the protocol, that is, TCP, UDP, IGMP, and ICMP information:

```
*Jun 2 08:58:06.489: %C4K_IOSINTF-6-SGACLHIT: list deny_udp_src_port_log-30 Denied
udp 24.0.0.23(100) -> 28.0.0.91(100), SGT8 DGT 12
```

In addition to the existing ‘per cell’ SGACL statistics, which can be displayed using the **show cts role-based counters** command, you can also display ACE statistics, by using the **show ip access-list *sgacl_name*** command. No additional configuration is required for this.

The following example shows how you can use the show ip access-list command to display the ACE count:

```
Device# show ip access-control deny_udp_src_port_log-30

Role-based IP access list deny_udp_src_port_log-30 (downloaded)
10 deny udp src eq 100 log (283 matches)
20 permit ip log (50 matches)
```



Note When the incoming traffic matches the cell, but does not match the SGACL of the cell, the traffic is allowed and the counters are incremented in the HW-Permit for the cell.

The following example shows how the SGACL of a cell works:

The SGACL policy is configured from 5 to 18 with “deny icmp echo” and there is incoming traffic from 5 to 18 with TCP header. If the cell matches from 5 to 18 but traffic does not match with icmp, traffic will be allowed and HW-Permit counter of cell 5 to 18 will get incremented.

```
Device# show cts role-based permissions from 5 to 18

IPv4 Role-based permissions from group 5:sgt_5_Contractors to group
18:sgt_18_data_user2:sgacl_5_18-01
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE

Device# show ip access-lists sgACL_5_18-01
Role-based IP access list sgACL_5_18-01 (downloaded)
10 deny icmp echo log (1 match)

Device# show cts role-based counters from 5 to 18
Role-based IPv4 counters
From To SW-Denied HW-Denied SW-Permitt HW-Permitt SW-Monitor HW-Monitor
5 18 0 0 0 1673202 0 0
```

VRF-aware SGACL Logging

The SGACL system logs will include VRF information. In addition to the fields that are currently logged, the logging information will include the VRF name. The updated logging information will be as shown below:

```
*Nov 15 02:18:52.187: %RBM-6-SGACLHIT_V6: ingress_interface='GigabitEthernet1/0/15'
sgacl_name='IPV6_TCP_DENY' action='Deny' protocol='tcp' src-vrf='CTS-VRF' src-ip='25::2'
src-port='20'
dest-vrf='CTS-VRF' dest-ip='49::2' dest-port='30' sgt='200' dgt='500'
logging_interval_hits='1'
```

SGACL Monitor Mode

During the pre-deployment phase of Cisco TrustSec, an administrator will use the monitor mode to test the security policies without enforcing them to make sure that the policies function as intended. If the security policies do not function as intended, the monitor mode provides a convenient mechanism for identifying that and provides an opportunity to correct the policy before enabling SGACL enforcement. This enables

administrators to have increased visibility to the outcome of the policy actions before they enforce it, and confirm that the subject policy meets the security requirements (access is denied to resources if users are not authorized).

The monitoring capability is provided at the SGT-DGT pair level. When you enable the SGACL monitoring mode feature, the deny action is implemented as an ACL permit on the line cards. This allows the SGACL counters and logging to display how connections are handled by the SGACL policy. Since all the monitored traffic is permitted, there is no disruption of service due to SGACLs while in the SGACL monitor mode.

Authorization and Policy Acquisition

After device authentication ends, both the supplicant and authenticator obtain the security policy from the authentication server. The two peers then perform link authorization and enforce the link security policy against each other based on their Cisco TrustSec device IDs. The link authentication method can be configured as either 802.1X or manual authentication. If the link security is 802.1X, each peer uses a device ID received from the authentication server. If the link security is manual, you must assign the peer device IDs.

The authentication server returns the following policy attributes:


- Cisco TrustSec trust—Indicates whether the peer device is to be trusted for the purpose of putting the SGT in the packets.
- Peer SGT—Indicates the security group to which the peer belongs. If the peer is not trusted, all packets received from the peer are tagged with this SGT. If the device does not know whether any SGACLs are associated with the peer's SGT, the device may send a follow-up request to the authentication server to download the SGACLs.
- Authorization expiry time—Indicates the number of seconds before the policy expires. A Cisco TrustSec device should refresh its policy and authorization before it times out. The device can cache the authentication and policy data and reuse it after a reboot if the data has not expired.



Note Each Cisco TrustSec device should support some minimal default access policy in case it is not able to contact the authentication server to get an appropriate policy for the peer.

The NDAC and SAP negotiation process is shown in the following figure

Figure 5: NDAC and SAP Negotiation



Environment Data Download

The Cisco TrustSec environment data is a collection of information or policies that assists a device to function as a Cisco TrustSec node. The device acquires the environment data from the authentication server when the device first joins a Cisco TrustSec domain, although you might also manually configure some of the data on a device. For example, you must configure the seed Cisco TrustSec device with the authentication server information, which can later be augmented by the server list that the device acquires from the authentication server.

The device must refresh the Cisco TrustSec environment data before it expires. The device can also cache the environment data and reuse it after a reboot if the data has not expired.

The device uses RADIUS to acquire the following environment data from the authentication server:

- Server lists: List of servers that the client can use for future RADIUS requests (for both authentication and authorization). PAC refresh happens through these servers.
- Device SG: Security group to which the device itself belongs.
- Expiry timeout: Interval that controls how often the Cisco TrustSec device should refresh its environment data.

RADIUS Relay Functionality

The device that plays the role of the Cisco TrustSec authenticator in the 802.1X authentication process has IP connectivity to the authentication server, allowing the device to acquire the policy and authorization from the authentication server by exchanging RADIUS messages over UDP/IP. The supplicant device may not have IP connectivity with the authentication server. In such cases, Cisco TrustSec allows the authenticator to act as a RADIUS relay for the supplicant.

The supplicant sends a special EAPOL message to the authenticator that contains the RADIUS server IP address and UDP port and the complete RADIUS request. The authenticator extracts the RADIUS request from the received EAPOL message and sends it over UDP/IP to the authentication server. When the RADIUS

response returns from the authentication server, the authenticator forwards the message back to the supplicant, encapsulated in an EAPOL frame.

Link Security

When both sides of a link support 802.1AE Media Access Control Security (MACsec), a security association protocol (SAP) negotiation is performed. An EAPOL-Key exchange occurs between the supplicant and the authenticator to negotiate a cipher suite, exchange security parameters, and manage keys. Successful completion of all three tasks results in the establishment of a security association (SA).

Depending on your software version, crypto licensing, and link hardware support, SAP negotiation can use one of the following modes of operation:

- Galois/Counter Mode (GCM)—Specifies authentication and encryption
- GCM authentication (GMAC)—Specifies authentication and no encryption
- No Encapsulation—Specifies no encapsulation (clear text)
- Null—Specifies encapsulation, no authentication and no encryption

All modes except No Encapsulation require Cisco TrustSec-capable hardware.


Configuring SAP-PMK for Link Security

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Device(config)# interface TenGigabitEthernet 1/1/4	Configures an interface and enters interface configuration mode.
Step 4	switchport mode trunk Example: Device(config-if)# switchport mode trunk	Specifies a trunking VLAN Layer 2 interface.
Step 5	cts manual Example: Device(config-if)# cts manual	Enters Cisco TrustSec manual configuration mode.

SXP typically operates between ingress access layer devices at the Cisco TrustSec domain edge and distribution layer devices within the Cisco TrustSec domain. The access layer device performs Cisco TrustSec authentication of external source devices to determine the appropriate SGTs for ingress packets. The access layer device learns the IP addresses of the source devices using IP device tracking and (optionally) DHCP snooping, then uses SXP to pass the IP addresses of the source devices along with their SGTs to the distribution devices. Distribution devices with Cisco TrustSec-capable hardware can use this IP-to-SGT mapping information to tag packets appropriately and to enforce SGACL policies.

Figure 6: SXP Protocol to Propagate SGT Information



You must manually configure an SXP connection between a peer without Cisco TrustSec hardware support and a peer with Cisco TrustSec hardware support. The following tasks are required when configuring the SXP connection:

- If you require SXP data integrity and authentication, you must configure the same SXP password on both peer devices. You can configure the SXP password either explicitly for each peer connection or globally for the device. Although an SXP password is not required, we recommend its use.
- You must configure each peer on the SXP connection as either an SXP speaker or an SXP listener. The speaker device distributes the IP-to-SGT mapping information to the listener device.
- You can specify a source IP address to use for each peer relationship or you can configure a default source IP address for peer connections where you have not configured a specific source IP address. If you do not specify any source IP address, the device will use the interface IP address of the connection to the peer.


SXP allows multiple hops. That is, if the peer of a device lacking Cisco TrustSec hardware support also lacks Cisco TrustSec hardware support, the second peer can have an SXP connection to a third peer, continuing the propagation of the IP-to-SGT mapping information until a hardware-capable peer is reached. A device can be configured as an SXP listener for one SXP connection as an SXP speaker for another SXP connection.

A Cisco TrustSec device maintains connectivity with its SXP peers by using the TCP keepalive mechanism. To establish or restore a peer connection, the device will repeatedly attempt the connection setup using a configurable retry period until the connection is successful or until the connection is removed from the configuration.

Layer 3 SGT Transport for Spanning Non-TrustSec Regions

When a packet leaves the Cisco TrustSec domain for a non-TrustSec destination, the egress Cisco TrustSec device removes the Cisco TrustSec header and SGT before forwarding the packet to the outside network. If, however, the packet is merely traversing a non-TrustSec domain on the path to another Cisco TrustSec domain, as shown in the following figure, the SGT can be preserved by using the Cisco TrustSec Layer 3 SGT Transport feature. In this feature, the egress Cisco TrustSec device encapsulates the packet with an ESP header that includes a copy of the SGT. When the encapsulated packet arrives at the next Cisco TrustSec domain, the ingress Cisco TrustSec device removes the ESP encapsulation and propagates the packet with its SGT.

Figure 7: Spanning a Non-TrustSec domain



To support Cisco TrustSec Layer 3 SGT Transport, any device that will act as a Cisco TrustSec ingress or egress Layer 3 gateway must maintain a traffic policy database that lists eligible subnets in remote Cisco TrustSec domains as well as any excluded subnets within those regions. You can configure this database manually on each device if they cannot be downloaded automatically from the Cisco Secure ACS.

A device can send Layer 3 SGT Transport data from one port and receive Layer 3 SGT Transport data on another port, but both the ingress and egress ports must have Cisco TrustSec-capable hardware.



Note Cisco TrustSec does not encrypt the Layer 3 SGT Transport encapsulated packets. To protect the packets traversing the non-TrustSec domain, you can configure other protection methods, such as IPsec.

VRF-Aware SXP

The SXP implementation of Virtual Routing and Forwarding (VRF) binds an SXP connection with a specific VRF. It is assumed that the network topology is correctly configured for Layer 2 or Layer 3 VPNs, with all VRFs configured before enabling Cisco TrustSec.

SXP VRF support can be summarized as follows:

- Only one SXP connection can be bound to one VRF.
- Different VRFs may have overlapping SXP peer or source IP addresses.
- IP-SGT mappings learned (added or deleted) in one VRF can be updated only in the same VRF domain. The SXP connection cannot update a mapping bound to a different VRF. If no SXP connection exists for a VRF, IP-SGT mappings for that VRF won't be updated by SXP.
- Multiple address families per VRF is supported. Therefore, one SXP connection in a VRF domain can forward both IPV4 and IPV6 IP-SGT mappings.
- SXP has no limitation on the number of connections and number of IP-SGT mappings per VRF.

Layer 2 VRF-Aware SXP and VRF Assignment

VRF to Layer 2 VLANs assignments are specified with the **cts role-based l2-vrf vrf-name vlan-list** global configuration command. A VLAN is considered a Layer 2 VLAN as long as there is no switch virtual interface (SVI) with an IP address configured on the VLAN. The VLAN becomes a Layer 3 VLAN once an IP address is configured on its SVI.

The VRF assignments configured by the **cts role-based l2-vrf** command are active as long as a VLAN remains a Layer 2 VLAN. The IP-SGT bindings learned while a VRF assignment is active are also added to the Forwarding Information Base (FIB) table associated with the VRF and the IP protocol version. If an SVI becomes active for a VLAN, the VRF to VLAN assignment becomes inactive and all the bindings learned on the VLAN are moved to the FIB table associated with the SVI's VRF.

The VRF to VLAN assignment is retained even when the assignment becomes inactive. It is reactivated when the SVI is removed or when the SVI IP address is deconfigured. When reactivated, the IP-SGT bindings are moved back from the FIB table associated with the SVI's VRF to the FIB table associated with the VRF assigned by the **cts role-based l2-vrf** command.

Feature History for Cisco TrustSec Overview

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Cisco TrustSec Overview	Cisco TrustSec builds secure networks by establishing domains of trusted network devices.

Release	Feature	Feature Information
Cisco IOS XE Cupertino 17.7.1	Cisco TrustSec Overview	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 2

SGACL and Environment Data Download over REST

This module describes the downloading of SGACL and environment data over REST APIs.

- [Prerequisites for SGACL and Environment Data Download over REST, on page 21](#)
- [Restrictions for SGACL and Environment Data Download over REST, on page 21](#)
- [Information About SGACL and Environment Data Download over REST, on page 22](#)
- [How to Configure SGACL and Environment Data Download over REST, on page 27](#)
- [Verifying the SGACL and Environment Data Download over REST, on page 31](#)
- [Debugging the SGACL and Environment Data over REST Configuration, on page 32](#)
- [Configuration Examples for SGACL and Environment Data Download over REST, on page 32](#)
- [Feature History for SGACL and Environment Data Download over REST, on page 33](#)

Prerequisites for SGACL and Environment Data Download over REST

- Cisco Identity Services Engine (ISE) Version should be 2.7 and above.
- Cisco TrustSec-enabled devices must use Cisco IOS XE Amsterdam 17.1.1 and later releases.
- The network device configuration on Cisco ISE must be updated to include the configuration to allow REST API calls from a network device IP address (NAS-IP). The device ID and password specified in the Cisco ISE configuration is included as the username and password by the network device that makes REST API calls to Cisco ISE.

Restrictions for SGACL and Environment Data Download over REST

- Cisco TrustSec Change of Authorization (CoA) uses RADIUS as the protocol.
- Only port 9063 is supported as the ERS server port.

- In Cisco IOS XE Amsterdam 17.1.1, third-party Certificate Authority (CA) certificate is not supported, only self-signed certificates are supported.
- Server statistics is not persistent after a refresh of the environment data.
- In Cisco IOS XE Amsterdam 17.1.1, IPv6 server is not supported. In Cisco IOS XE 17.2.1, IPv6 servers are supported.
- In Cisco IOS XE Amsterdam 17.1.1, only one IPv4 address per server is supported.
- Only one Fully Qualified Domain Name (FQDN) per server is supported.

Information About SGACL and Environment Data Download over REST

SGACL and Environment Data Download over REST Overview

In Cisco IOS XE Amsterdam 17.1.1 and later releases, Cisco TrustSec uses the REST-based transport protocol for policy provisioning and environment data download from Cisco Identity Services Engine (ISE). The REST-based protocol is more secure, and provides reliable, and faster Security Group access control list (SGACL) policy and environment data provisioning, than the RADIUS protocol that is used in previous releases.

Both the REST API-based and RADIUS-based download of Cisco TrustSec data is supported. However, only one protocol can be active on a device. In Cisco IOS XE Amsterdam 17.1.1, REST-based protocol is the default. However, you can change the protocol to RADIUS by configuring the **cts authorization list** command.



Note Cisco TrustSec Change of Authorization (CoA) will still use RADIUS as the protocol.

Cisco TrustSec Security Group Access Control List (SGACL) and environment data are synchronized from the active device to the standby device, after the policy is installed. However, REST API connections or sessions are not synchronized during a switchover.

In Cisco IOS XE Amsterdam 17.1.1, only one IPv4 address per server is supported. In Cisco IOS XE Amsterdam 17.2.1 and later releases, 8 IPv4 and 8 IPv6 addresses per server are supported.

In Cisco IOS XE Amsterdam 17.2.1, the Cisco TrustSec device will honor the 429 response code from Cisco ISE. This response code is sent by Cisco ISE, when it is overloaded. Once a 429 response code is received for a particular server, the device marks the server as dead, and switches to the next server in the list (private or public). The next retry attempt is done after 60 seconds.

Cisco TrustSec Environment Data

Environment data comprises of operational data that supplement Cisco TrustSec functions. The environment data request from a device to Cisco ISE consists of the following data:

- Device name: Specifies the name of the device.
- Device capability: Specifies additional data.

The environment data response from Cisco ISE to a device consists of the following data:

- Device security group tag (SGT): Derived from Cisco ISE based on the device name.
- Server list: Displays the list of Cisco TrustSec servers specified in Cisco ISE.
- SG-Name Table: Displays the mapping between SGT and the device name. SGT is displayed in numerals and the device name in text format.
- Refresh time: Indicates the time when the environment data will be refreshed.

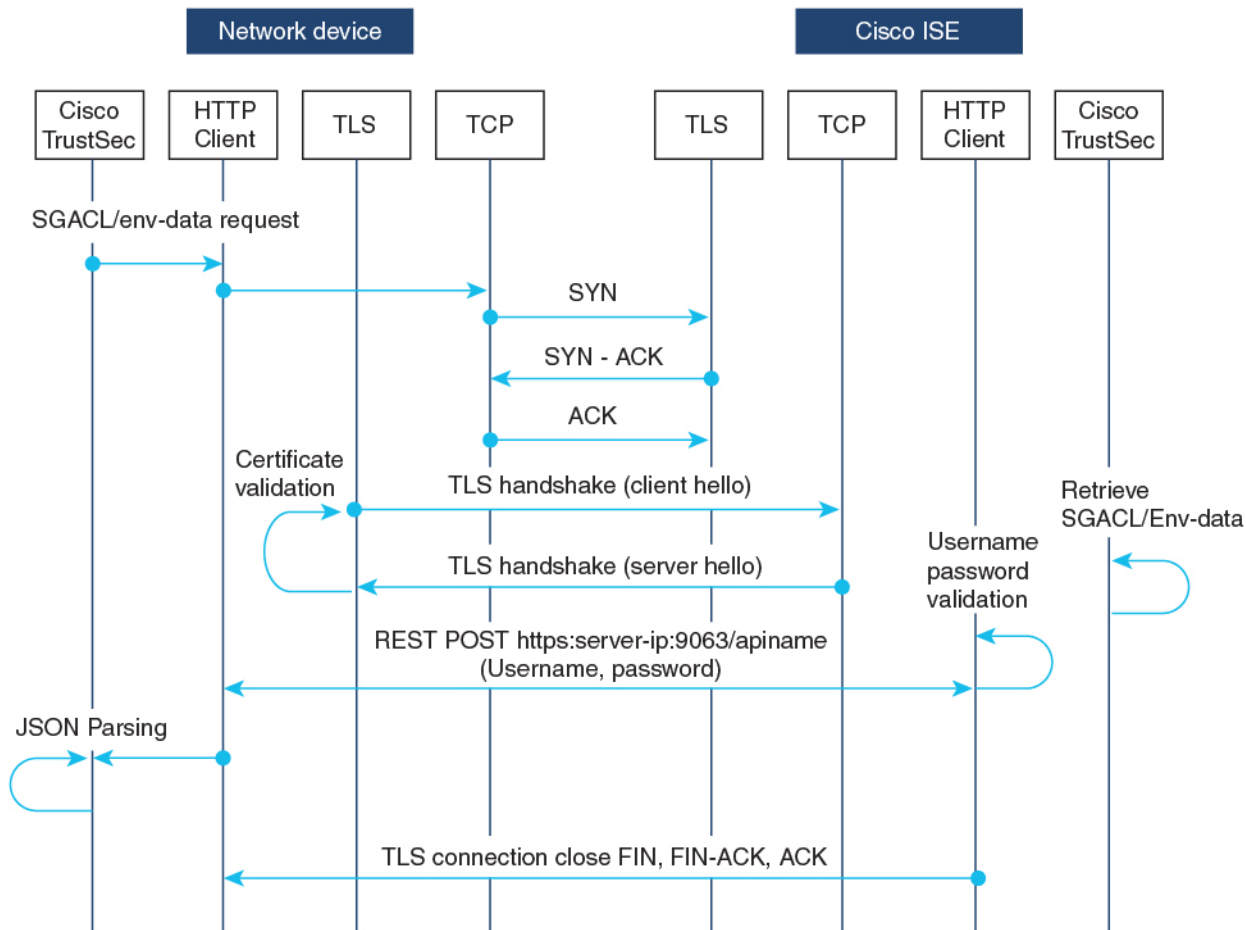


Note As part of Cisco TrustSec environment data refresh, the last received servers are deleted and newly received servers are added to the server list. As a result of the refresh, the server list statistics restarts from zero, and the server status is set to *Inactive* and the IP address state is set to *Reachable*. The device then updates the server statistics and status based on the subsequent policy request and response.

Message Flow Between a Network Device and a Server

The following illustration displays the connection management for REST calls between a network device and server.

Figure 8: Message Flow Between a Network Device and a Server



- Cisco ISE REST API service runs on a secure socket that runs Transport Layer Security (TLS) 1.2 server on port 9063 to service network device requests for SGACL and environment data.
- The device uses a make or break approach to the TLS connection establishment, and there is no persistent TLS connection between the device and Cisco ISE. After the TLS connection is established, the device can use this connection to submit multiple REST API calls to specific resource uniform resource locators (URLs). After all REST requests are processed, the server terminates the connection through a TCP-FIN message. For new REST API calls a new connection must be established with the server.
- The REST API call from the device to Cisco ISE starts with a TCP connection establishment. Cisco ISE must be configured with device IP address to allow ingress connections from the device. TCP connection requests from source IP addresses that are not configured on Cisco ISE are dropped, and an audit log created.
- Username and password: Every RESTAPI call must include the username and password authentication while requesting access to a resource uniform resource identifier (URI). The authentication helps the server to determine if the caller should be given access to the resource or to deny the request.
- A successful TLS connection establishment with Cisco ISE requires its server-certificate signature or PEM to be installed as the trustpoint (by using the **crypto pki trustpoint** command) on the device to

trust the server. Only fingerprint or signature of the server certificate need to be exported and installed on the device under a trustpoint. Import of private-key of the server certificate is not necessary.

- After establishing the TLS connection, the HTTP client on the device initiates a REST call to Cisco ISE on the specified resource.

Policy Server Selection Criteria

Multiple HTTP policy servers are configured on a Cisco TrustSec device. Once a server is selected, the device use this server to interact with Cisco ISE until the server is marked as dead.

There are two types of server selection:

- **In-Order Selection:** This is the default behavior, where servers are picked in the order in which they are configured (from the public server list) or downloaded (from the private server list). Once a server is selected, the device is used till it is marked as dead, and then the next server in the list is selected.

When environment data is successfully downloaded, and a server-list is available, these servers are added to the private server list.

- **Random Server-Selection:** When multiple HTTP policy servers are configured on a device, a single Cisco ISE instance may get overloaded if the device always selects the first configured server. To avoid this situation, each device will randomly select a server. A random number is generated by the device and based on this number a server is selected. For different devices to generate random numbers, the unique board ID and the Cisco TrustSec process ID of the device is used to initialize the random number generator.

Once a server is selected, all future requests go this server until the server is marked as dead. Once a server is in the dead state, the random server selection logic picks up the next alive server. The dead server is not added to the count of active servers when picking the new server. The server numbering starts with zero.

Selected Server = (Generated Random Number) % (Total Number of Active Servers).

To change the server selection logic to random, use the **cts policy-server order random** command.

Server and IP Address Selection Process

The order of server-selection is the private server-list (received as part of server-list download), followed by the public server-list (configured servers). Within these server lists, the order can either be random selection or in-order selection based on whether the **cts policy-server order random** command is enabled or not.

In Cisco IOS XE 17.2.1 and later releases, multiple IP (both IPv4 and IPv6) addresses per server are supported. The order of IP selection is IPv4 addresses, followed by IPv6 addresses, and then FQDN.

This section describes how the server and IP address selection works:

1. When a device boots up for the first time, a server from the public (configured) list is selected.
2. If the **cts environment-data enable** command is configured, the device uses the public server to download the private server-list from Cisco ISE.
3. After successfully receiving the private list, all subsequent requests will use the private list.
4. After the server and IP address are selected, the device connects to Cisco ISE using the server/IP address combination. This server will interact with Cisco ISE until it fails to get a response.

5. If no response is received from the current active server in the private list, the device switches to the next server in the list. If the server is selected for the first time, the IP selection logic searches for the first reachable IP or IPv6 address.
6. After the server and IP address selection, the device is used until it goes down.
7. If none of the servers in the private list are reachable, the device attempts to connect to the servers in the public list. The server switching logic and IP selection are the same for private and public list.
8. The server change happens only when the server list is refreshed.
9. If all servers in both the private and public server list are dead, the device restarts the server/IP address selection logic from the start of the private list.
10. When a specific server/IP address combination fails, the device waits for 60 seconds before it attempts a new combination.

Server Liveliness Check

Whether a server is alive is determined after sending an environment-data or an SGACL request to Cisco ISE. There is no liveliness detection phase after a server is configured or downloaded as part of a server list. The default server status is alive for all types of servers.

When a request is sent to Cisco ISE, and if the server is not reachable or the response is lost, the server is moved to dead state. The server selection logic will pick the same server and the next IP address (if multiple addresses are configured) to send the next set of Cisco ISE requests. The logic will pick the next server in the list, if the device receives the overloaded response (HTTP 429) from Cisco ISE.

A server can be marked as dead because of any of the following reasons:

- The configured IP address is not reachable.
- Incorrect port number.
- The Cisco ISE instance with the IP address is down.
- The interface towards Cisco ISE is down.
- A Transport Layer Security (TLS) handshake failure.
- An HTTP response timeout.
- An incorrectly configured domain name (if a domain name is used).

If a server has both the static IP address and the domain name configured, preference is given to the static IP address. If there is no response to the static IP address, the device tries with the domain name. When no response is received with both the static IP address and the domain-name, the server is marked as dead.

When all servers of the private list are marked as dead, the device uses the public list. If all remaining servers are also marked as dead, then the recovery mechanism starts. The device waits for the next Cisco TrustSec request (for policy refresh, environment data download or refresh, and so on), and marks all the servers as alive to retry the download. If there is no trigger for a new Cisco TrustSec request, the servers remain in the dead state.

How to Configure SGACL and Environment Data Download over REST

Configuring the Username and Password

Configure the username and password in Cisco ISE as the REST API access credentials, before configuring it on the device. See the [Cisco TrustSec HTTP Servers](#) section of the "Cisco TrustSec Policies Configuration" chapter for more information.



Note If you try to configure RADIUS-based configuration by using the `cts authorization-list` command, when the HTTP-based configurations are already enabled, the following error message is displayed on the console:

```
Error: 'cts policy-server or cts environment-data' related configs are enabled.
Disable http-based configs, to enable 'cts authorization'
```

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts policy-server name <i>server-name</i> Example: Device(config)# cts policy-server name ISE-server	Configures a Cisco TrustSec policy server and enters policy-server configuration mode.
Step 4	exit Example: Device(config-policy-server)# exit	Exits policy-server configuration mode and returns to global configuration mode.
Step 5	cts policy-server username <i>username</i> password {0 6 7 <i>password</i>} {<i>password</i>} Example: Device(config)# cts policy-server username admin password 6 password1	Configures an username and password. Note This username and password must be created on Cisco ISE as the REST API access credentials before configuring it on the device.

	Command or Action	Purpose
Step 6	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Configuring Certificate Enrollment

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	crypto pki trustpoint <i>name</i> Example: Device(config)# crypto pki trustpoint mytp	Declares the trustpoint and a given name, and enters ca-trustpoint configuration mode.
Step 4	exit Example: Device(ca-trustpoint)# exit	Exits ca-trustpoint configuration mode and returns to global configuration mode.
Step 5	crypto pki authenticate <i>name</i> Example: Device(config)# crypto pki authenticate mytp	Retrieves the Certificate Authority (CA) certificate and authenticates it. Check the certificate fingerprint if prompted. <p>Note This command is optional if the CA certificate is already loaded into the configuration.</p>
Step 6	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Downloading Cisco TrustSec Policies

The **cts role-based enforcement** must already be configured to download Cisco TrustSec Policies.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts policy-server name <i>server-name</i> Example: Device(config)# cts policy-server name ISE-server	Configures a Cisco TrustSec policy server and enters policy-server configuration mode.
Step 4	address domain-name <i>name</i> Example: Device(config-policy-server)# address domain-name domain1	Configures the domain name address of the policy server.
Step 5	address {ipv4 ipv6} <i>policy-server-address</i> Example: Device(config-policy-server)# address ipv4 10.1.1.1 Device(config-policy-server)# address ipv6 2001.DB8::1	Configures the IPv4 or IPv6 address of the policy server. <ul style="list-style-type: none">• In Cisco IOS XE Amsterdam 17.1.1, only IPv4 addresses are supported.
Step 6	tls server-trustpoint <i>name</i> Example: Device(config-policy-server)# tls server-trustpoint tls1	Configures the Transport Layer Security trustpoint.
Step 7	timeout <i>seconds</i> Example: Device(config-policy-server)# timeout 15	(Optional) Configures the response timeout in seconds. <ul style="list-style-type: none">• The default is 5 seconds.
Step 8	retransmit <i>number-of-retries</i> Example: Device(config-policy-server)# retransmit 4	(Optional) Configures the maximum number of retries from the server. <ul style="list-style-type: none">• The default is 4.
Step 9	port <i>port-number</i> Example: Device(config-policy-server)# port 9063	(Optional) Configures the policy server port number. Note The ERS server port number must be 9063. You cannot change this port number.

	Command or Action	Purpose
Step 10	content-type json Example: <pre>Device(config-policy-server)# content-type json</pre>	(Optional) Configures the content type to source SGACL and environment data from Cisco ISE. Note By default, JSON is used as the content type, even if this command is not configured.
Step 11	end Example: <pre>Device(config-policy-server)# end</pre>	Exits policy-server configuration mode and returns to privileged EXEC mode.

Downloading Environment Data

The source interface to use for HTTP connections must be specified in the **ip http client source-interface** command.

Procedure

	Command or Action	Purpose
Step 1	enable Example: <pre>Device> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.
Step 3	cts policy-server device-id <i>device-ID</i> Example: <pre>Device(config)# cts policy-server device-id server1</pre>	Configures the policy server device ID to send environment data requests to Cisco ISE. <ul style="list-style-type: none"> • This device-ID must be the one used to add the network access device (NAD) on Cisco ISE.
Step 4	cts environment-data enable Example: <pre>Device(config)# cts environment-data enable</pre>	Enables the downloading of environment data from Cisco ISE. Note The cts environment-data enable command and the cts authorization list command are mutually exclusive. These commands cannot be configured together.
Step 5	end Example:	Exits global configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
	Device(config)# end	

Verifying the SGACL and Environment Data Download over REST

Use the following commands in any order:

- **show cts policy-server details name**

Displays information about the specified policy server.

```
Device# show cts policy-server details name ise_server_1

Server Name      : ise_server_1
Server Status   : Active
  IPv4 Address    : 10.64.69.84
  IPv6 Address    : 2001:DB::2
  Trustpoint      : ISE84
  Port-num        : 9063
  Retransmit count : 3
  Timeout         : 15
  App Content type : JSON
```

- **show cts policy-server statistics active**

Displays statistics information about active policy servers.

When you use the command without the **active** the statistics of all servers are listed.

```
Device# show cts policy-server statistics active

Server Name      : ise_server_1
Server State     : ALIVE
  Number of Request sent      : 7
  Number of Request sent fail : 0
  Number of Response received : 4
  Number of Response rcv fail : 3
  HTTP 200 OK                 : 4
  HTTP 400 BadReq             : 0
  HTTP 401 Unauthorized Req   : 0
  HTTP 403 Req Forbidden      : 0
  HTTP 404 NotFound           : 0
  HTTP 408 ReqTimeout         : 0
  HTTP 415 Unsupported Media   : 0
  HTTP 500 ServerErr          : 0
  HTTP 501 Req NoSupport      : 0
  HTTP 503 Service Unavailable: 0
  TCP or TLS handshake error  : 3
  HTTP Other Error            : 0
```

- **show cts server-list**

Displays the list of servers that are downloaded as part of the environment data. These servers will be part of private server-list.



Note The following output displays the HTTP-based download information:

```
Device# show cts server-list

HTTP Server-list:
  Server Name       : cts_private_server_0
  Server State      : ALIVE
  IPv4 Address      : 10.64.69.151
  IPv6 Address      : 2001:DB8:8086:6502::
  IPv6 Address      : 2001:db8::2
  IPv6 Address      : 2001:db8::402:99
  IPv6 Address      : 2001:DB8:4::802:16
  Domain-name       : ise-267.cisco.com
  Trustpoint        : cts_trustpoint_0

  Server Name       : cts_private_server_1
  Server State      : ALIVE
  IPv4 Address      : 10.10.10.3
  IPv4 Address      : 10.10.10.2
  IPv6 Address      : 2001:DB8::20
  IPv6 Address      : 2001:DB8::21
  Domain-name       : www.ise.cisco.com
  Trustpoint        : cts_trustpoint_1
```

Debugging the SGACL and Environment Data over REST Configuration

Use the following **debug** commands for debugging the configuration.

- **debug cts policy-server http**
Enables HTTP client debugging.
- **debug cts policy-server json**
Enables JSON client debugging.

Configuration Examples for SGACL and Environment Data Download over REST

Example: Configuring the Username and Password

```
Device> enable
Device# configure terminal
Device(config)# cts policy-server name ISE-server
Device(config-policy-server)# exit
Device(config)# cts policy-server username admin 6 password1
```

```
Device(config)# end
```

Example: Downloading Cisco TrustSec Policies

```
Device> enable
Device# configure terminal
Device(config)# cts role-based enforcement
Device(config)# cts policy-server name ISE-server
Device(config-policy-server)# address domain-name domain1
Device(config-policy-server)# address ipv4 10.1.1.1
Device(config-policy-server)# address ipv6 2001:DB8::1
Device(config-policy-server)# tls server-trustpoint tls1
Device(config-policy-server)# timeout 15
Device(config-policy-server)# retransmit 4
Device(config-policy-server)# port 2010
Device(config-policy-server)# end
```

Example: Downloading Environment Data

```
Device> enable
Device# configure terminal
Device(config)# cts policy-server name ISE-server
Device(config-policy-server)# exit
Device(config)# cts policy-server device-id server1
Device(config)# cts env-data enable
Device(config)# end
```

Feature History for SGACL and Environment Data Download over REST

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Amsterdam 17.1.1	SGACL and Environment Data Download over REST	Cisco TrustSec uses the REST-based transport protocol for SGACL policy provisioning and data download from Cisco ISE.
Cisco IOS XE Amsterdam 17.2.1	HTTP SGACL Enforcement with IPv6 Policy Server	IPv6 addresses for policy servers are supported.

Release	Feature	Feature Information
Cisco IOS XE Cupertino 17.7.1	SGACL and Environment Data Download over REST HTTP SGACL Enforcement with IPv6 Policy Server	These features were implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 3

Configuring Security Group ACL Policies

Using security group access control lists (SGACLs), you can control the operations that users can perform based on the security group assignments of users and destination resources. Policy enforcement within the Cisco TrustSec domain is represented by a permissions matrix, with source security group numbers on one axis and destination security group numbers on the other axis. Each cell in the body of the matrix can contain an ordered list of SGACLs, which specifies the permissions that should be applied to packets originating from the source security group and destined for the destination security group.

- [Restrictions for Configuring Security Group ACL Policies, on page 35](#)
- [Information About Security Group ACL Policies, on page 36](#)
- [How to Configure Security Group ACL Policies, on page 36](#)
- [Configuration Examples for Security Group ACL Policies, on page 45](#)
- [Feature History for Security Group ACL Policies, on page 47](#)

Restrictions for Configuring Security Group ACL Policies

- Due to hardware limitations, Cisco TrustSec SGACLs cannot be enforced for punt (CPU bound) traffic in hardware. SGACL enforcement in software is bypassed wfor CPU-bound traffic for switch virtual interface (SVI) and Layer 2 and Layer 3 Location Identifier Separation Protocol (LISP), and loopback interfaces.
- When configuring SGACL policies, if you change the IP version dynamically from **IPv4** or **IPv6** to **Agnostic** (applies to both IPv4 and IPv6) and vice-versa, the corresponding SGACL policies for IPv4 and IPv6 are not downloaded completely through the management VRF interface.
- When configuring SGACL policies, if you change the existing IP version to any other version (**IPv4**, **IPv6**, or **Agnostic**) and vice-versa, Change of Authorization (CoA) from Cisco Identity Services Engine (ISE) cannot be performed using RADIUS. Instead, use SSH and run the **cts refresh policy** command to perform a manual policy refresh.
- When using an allowed SGT model with default action as **deny all**, in some cases, Cisco TrustSec policies are only partially downloaded from the ISE server after a device reload.

To prevent this, define a static policy on the device. Even if the **deny all** option is applied, the static policy permits traffic that allows the device to download policies from the ISE server and overwrite the defined static policies. For device SGT, configure the following commands in global configuration mode:

- **cts role-based permissions from <sgt_num> to unknown**
- **cts role-based permissions from unknown to <sgt_num>**

Information About Security Group ACL Policies

The following sections provide information about configuring SGACL policies.

SGACL Logging

A device can provide logging messages about packets that are permitted or denied by a standard IP access list. That is, any packet that matches an SGACL causes an informational logging message about the packet to be sent to the console. The limit of messages logged to the console is controlled by the **logging console** command that controls the syslog messages. In releases prior to Cisco IOS XE Amsterdam 17.3.1, SGACL logging was done as a CPU-intensive mechanism. From Cisco IOS XE Amsterdam 17.3.1 release, SGACL logging has been enhanced to use NetFlow hardware, which allows much larger logging rates.



Note SGACL logging in hardware is only supported for Role-Based access control list (RBACL).

The first packet that triggers the SGACL creates a flow, and logging is done at the NetFlow timeout of 30 seconds and 1 minute for inactive and active flows respectively. Subsequent packets are collected over 5-minute intervals before they are logged. The logging message includes the access list number, whether the packet was permitted or denied, the source and destination IP addresses of the packet, the interface on which the packet was ingress, and the number of packets from that source permitted or denied in the previous 5-minute interval.



-
- Note**
- Because SGACL logging in the hardware is done using NetFlow, if a NetFlow-based feature is applied to an interface, logging for that interface falls back to the old mechanism. Logging through NetFlow hardware starts again for that interface after the NetFlow-based feature is removed. The rest of the interfaces continue logging through NetFlow hardware.
 - Only 15 NetFlow monitors can be attached to the device at a given time. SGACL logging requires one NetFlow monitor each for IPv4 and IPv6 logging. If NetFlow monitors are not available for logging, SGACL logging is done through the earlier mechanism. Once the required number of NetFlow monitors are available, run the **cts role-based permissions** command to trigger logging through the NetFlow hardware again.
 - If a log access control entry (ACE) has fields other than source port number, destination port number and the protocol in use, logging is done through the earlier mechanism.
-

How to Configure Security Group ACL Policies

The following sections provide information about various SGACL policy configurations.

SGACL Policy Configuration Process

Follow these steps to configure and enable SGACL policies:

1. Configuration of SGACL policies should be done primarily through the Policy Management function of the Cisco Secure Access Control Server (ACS) or the Cisco Identity Services Engine (ISE).

If you are not using AAA on a Cisco Secure ACS or a Cisco ISE to download the SGACL policy configuration, you can manually configure the SGACL mapping and policies.



Note An SGACL policy that is downloaded dynamically from the Cisco Secure ACS or a Cisco ISE will override any conflicting locally-defined policy.

2. To enable SGACL policy enforcement on egress traffic on routed ports, enable SGACL policy enforcement globally as described in the *Enabling SGACL Policy Enforcement Globally* section.
3. To enable SGACL policy enforcement on switched traffic within a VLAN, or on traffic that is forwarded to an SVI that is associated with a VLAN, enable SGACL policy enforcement for specific VLANs, as described in the *Enabling SGACL Policy Enforcement on VLANs* section.

Enabling SGACL Policy Enforcement Globally

You must enable SGACL policy enforcement globally for Cisco TrustSec-enabled routed interfaces.

To enable SGACL policy enforcement on routed interfaces, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts role-based enforcement Example: Device(config)# cts role-based enforcement	Enables Cisco TrustSec SGACL policy enforcement on routed interfaces.
Step 4	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Enabling SGACL Policy Enforcement Per Interface

You must first enable SGACL policy enforcement globally for Cisco TrustSec-enabled routed interfaces. This feature is not supported on port channel interfaces.

To enable SGACL policy enforcement on Layer 3 interfaces, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface type slot/port Example: Device(config)# interface gigabitethernet 6/2	Configures an interface and enters interface configuration mode.
Step 4	cts role-based enforcement Example: Device(config-if)# cts role-based enforcement	Enables Cisco TrustSec SGACL policy enforcement on routed interfaces.
Step 5	end Example: Device(config-if)# end	Exits interface configuration mode and returns to privileged EXEC mode.
Step 6	show cts interface Example: Device# show cts interface	(Optional) Displays Cisco TrustSec states and statistics per interface.

Enabling SGACL Policy Enforcement on VLANs

You must enable SGACL policy enforcement on specific VLANs to apply access control to switched traffic within a VLAN, or to traffic that is forwarded to an SVI associated with a VLAN.

To enable SGACL policy enforcement on a VLAN or a VLAN list, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.

	Command or Action	Purpose
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	cts role-based enforcement vlan-list <i>vlan-list</i> Example: Device(config)# <code>cts role-based enforcement vlan-list 31-35,41</code>	Enables Cisco TrustSec SGACL policy enforcement on the VLAN or VLAN list.
Step 4	end Example: Device(config)# <code>end</code>	Exits global configuration mode and returns to privileged EXEC mode.

Configuring SGACL Monitor Mode

Before configuring SGACL monitor mode, ensure the following:

- Cisco TrustSec is enabled
- Counters are enabled

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# <code>enable</code>	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	cts role-based monitor all Example: Device(config)# <code>cts role-based monitor all</code>	Enables global monitor mode.
Step 4	cts role-based monitor permissions from {<i>sgt_num</i>} to {<i>dgt_num</i>} [<i>ipv4</i> <i>ipv6</i>] Example: Device(config)# <code>cts role-based permissions from 2 to 3 ipv4</code>	Enables monitor mode for IPv4 or IPv6 Role-Based Access Control List (RBACL) (Security Group Tag-Destination Group Tag [SGT-DGT] pair).
Step 5	end Example:	Exits global configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
	Device(config)# end	
Step 6	show cts role-based permissions from <i>{sgt_num} to {dgt_num} [ipv4 ipv6] [details]</i> Example: Device# show cts role-based permissions from 2 to 3 ipv4 details	(Optional) Displays the SGACL policies and details about the monitor mode functionality for each pair. The command output displays if per-cell monitor mode is enabled for the <SGT-DGT> pair.
Step 7	show cts role-based counters [ipv4 ipv6] Example: Device# show cts role-based counters ipv4	(Optional) Displays all the SGACL enforcement statistics for IPv4 and IPv6 events.

Manually Configuring SGACL Policies

A role-based access control list bound to a range of SGTs and DGTs forms an SGACL, a Cisco TrustSec policy enforced on egress traffic. Configuration of SGACL policies are best done through the policy-management functions of Cisco ISE or Cisco Secure ACS. To manually, that is, locally, configure SGACL policies, configure a role-based ACL and bind this role-based ACL to a range of SGTs.



Note An SGACL policy downloaded dynamically from Cisco ISE or Cisco ACS overrides conflicting manually configured policies, if any.

Configuring and Applying IPv4 SGACL Policies



Note When configuring SGACLs and RBACLs, the named access control lists (ACLs) must start with an alphabet.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	ip access-list role-based <i>rbacl-name</i> Example:	Creates an RBACL and enters Role-based ACL configuration mode.

	Command or Action	Purpose
	Device(config)# ip access-list role-based allow_webtraff	
Step 4	<p>{[<i>sequence-number</i>] default permit deny remark}</p> <p>Example:</p> <pre>Device(config-rb-acl)# 10 permit tcp dst eq 80 dst eq 20</pre>	<p>Specifies the access control entries (ACEs) for the RBACL.</p> <p>You can use most of the commands and options allowed in extended named access list configuration mode, with the source and destination fields omitted.</p> <p>The following ACE keywords are not supported:</p> <ul style="list-style-type: none"> • reflect • evaluate • time-range
Step 5	<p>exit</p> <p>Example:</p> <pre>Device(config-rb-acl)# exit</pre>	Exits role-based ACL configuration mode and returns to global configuration mode.
Step 6	<p>cts role-based permissions {default [from {<i>sgt_num</i> unknown} to {<i>dgt_num</i> unknown}] {<i>rbac</i>s ipv4 rbacs}</p> <p>Example:</p> <pre>Device(config)# cts role-based permissions from 55 to 66 allow_webtraff</pre>	<p>Binds SGTs and DGTs to the RBACL. The configuration is analogous to populating the permission matrix configured on Cisco ISE or Cisco Secure ACS.</p> <ul style="list-style-type: none"> • default: Default permissions list. • <i>sgt_num</i>: 0 to 65,519. Source Group Tag. • <i>dgt_num</i>: 0 to 65,519. Destination Group Tag. • unknown: SGACL applies to packets where the security group (source or destination) cannot be determined. • ipv4: Indicates the RBACLs are IPv4. • <i>rbac</i>s: Names of RBACLs.
Step 7	<p>end</p> <p>Example:</p> <pre>Device(config)# end</pre>	Exits global configuration mode and returns to privileged EXEC mode.
Step 8	<p>show cts role-based permissions</p> <p>Example:</p> <pre>Device# show cts role-based permissions</pre>	(Optional) Displays permission to RBACL configurations.

	Command or Action	Purpose
Step 9	show ip access-lists { <i>rbacIs</i> ipv4 rbacIs } Example: Device# show ip access-lists allow_webtraff	(Optional) Displays ACEs of all RBACLs or a specified RBACL.

Configuring IPv6 SGACL Policies

To manually configure IPv6 SGACL policies, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	ipv6 access-list role-based <i>sgacl-name</i> Example: Device(config)# ipv6 access-list role-based sgaclname	Creates a named IPv6 SGACL and enters IPv6 role-based ACL configuration mode.
Step 4	{permit deny } protocol [dest-option dest-option-type {<i>doh-number</i> <i>doh-type</i>}] [dscp <i>cp-value</i>] [flow-label <i>fl-value</i>] [mobility mobility-type {<i>mh-number</i> <i>mh-type</i>}] [routing routing-type <i>routing-number</i>] [fragments] [log log-input] [sequence <i>seqno</i>] Example: Device(config-ipv6rb-acl)# permit 33 dest-option dscp af11	Specifies the access control entries (ACEs) for the RBACL. You can use most of the commands and options allowed in extended named access list configuration mode, with the source and destination fields omitted. The following ACE keywords are not supported: <ul style="list-style-type: none"> • reflect • evaluate • time-range
Step 5	end Example: Device(config-ipv6rb-acl)# end	Exits IPv6 role-based ACL configuration mode and returns to privileged EXEC mode.

Manually Applying SGACL Policies

To manually apply SGACL policies, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts role-based permissions default [ipv4 ipv6] sgacl-name1 [sgacl-name2 [sgacl-name3 ...]]] Example: Device(config)# cts role-based permissions default MYDEFAULTSGACL	Specifies the default SGACL. The default policies are applied when no explicit policy exists between the source and destination security groups.
Step 4	cts role-based permissions from {source-sgt unknown} to {dest-sgt unknown} [ipv4 ipv6] sgacl-name1 [sgacl-name2 [sgacl-name3 ...]]] Example: Device(config)# cts role-based permissions from 3 to 5 SRB3 SRB5	Specifies the SGACLs to be applied for an SGT and a DGT. Values for <i>source-sgt</i> and <i>dest-sgt</i> range from 1 to 65533. By default, SGACLs are considered to be IPv4. <ul style="list-style-type: none"> • from: Specifies the source SGT. • to: Specifies the destination security group. • unknown: SGACL applies to packets where the security group (source or destination) cannot be determined. <p>Note An SGACL policy downloaded dynamically from the ACS will override conflicting manual policies, if any.</p>
Step 5	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Displaying SGACL Policies

After configuring the Cisco TrustSec device credentials and AAA, you can verify the Cisco TrustSec SGACL policies that are downloaded from the authentication server or configured manually. Cisco TrustSec downloads

the SGACL policies when it learns of a new SGT Exchange Protocol (SXP) through authentication and authorization on an interface, from SXP, or from manual IP address to SGT mapping.

By using or omitting keywords, you can display all or part of the permissions matrix:

- If the **from** keyword is omitted, a column from the permissions matrix is displayed.
- If the **to** keyword is omitted, a row from the permissions matrix is displayed.
- If the **from** and **to** keywords are omitted, the entire permissions matrix is displayed.
- If the **from** and **to** keywords are specified, a single cell from the permissions matrix is displayed, and the **details** keyword is available. When **details** is entered, the ACEs of the SGACL of the single cell are displayed.

To display the contents of the SGACL policies' permissions matrix, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	show cts role-based permissions default [ipv4 ipv6 details] Example: Device# show cts role-based permissions default MYDEFAULTSGACL	Displays the list of SGACL, of the default policy.
Step 3	show cts role-based permissions from {source-sgt unknown} to {dest-sgt unknown}] [ipv4 ipv6 details] Example: Device# show cts role-based permissions from 3	Specifies the SGACLs to be applied for an SGT and a DGT. Values for <i>source-sgt</i> and <i>dest-sgt</i> range from 1 to 65533. By default, SGACLs are considered to be IPv4. <ul style="list-style-type: none"> • from: Specifies the source SGT. • to: Specifies the destination security group. • unknown: SGACL applies to packets where the security group (source or destination) cannot be determined. <p>Note An SGACL policy downloaded dynamically from the ACS will override conflicting manual policies, if any.</p>
Step 4	exit Example: Device# exit	Exits privileged EXEC mode.

Refreshing the Downloaded SGACL Policies

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	cts refresh policy {peer [peer-id] sgt [sgt_number] default unknown} Example: Device# cts refresh policy peer my_cisco_ise	Performs an immediate refresh of the SGACL policies from the authentication server. <ul style="list-style-type: none"> • If a <i>peer-id</i> is specified, only the policies related to the specified peer connection are refreshed. To refresh all the peer policies, press Enter without specifying an ID. • If an SGT number is specified, only the policies related to that SGT are refreshed. To refresh all the SGT policies, press Enter without specifying an SGT number. Select default to refresh the default policy. Select unknown to refresh an unknown policy.
Step 3	exit Example: Device# exit	Exits privileged EXEC mode.

Configuration Examples for Security Group ACL Policies

The following sections provide examples of various SGACL policy configurations.

Example: Enabling SGACL Policy Enforcement Globally

The following example shows how to enable SGACL policy enforcement globally:

```
Device> enable
Device# configure terminal
Device(config)# cts role-based enforcement
```

Example: Enabling SGACL Policy Enforcement Per Interface

The following example shows how to enable SGACL policy enforcement per interface:

Example: Enabling SGACL Policy Enforcement on VLANs

```

Device> enable
Device# configure terminal
Device(config)# interface gigabitethernet 1/0/2
Device(config-if)# cts role-based enforcement
Device(config-if)# end

```

Example: Enabling SGACL Policy Enforcement on VLANs

The following example shows how to enable SGACL policy enforcement on VLANs:

```

Device> enable
Device# configure terminal
Device(config)# cts role-based enforcement vlan-list 31-35,41
Device(config)# exit

```

Example: Configuring SGACL Monitor Mode

The following example shows how to configure SGACL monitor mode:

```

Device> enable
Device# configure terminal
Device(config)# cts role-based monitor enable
Device(config)# cts role-based permissions from 2 to 3 ipv4
Device# show cts role-based permissions from 2 to 3 ipv4

IPv4 Role-based permissions from group 2:sgt2 to group 3:sgt3 (monitored):
  denytcpudpicmp-10
  Deny IP-00

Device# show cts role-based permissions from 2 to 3 ipv4 details

IPv4 Role-based permissions from group 2:sgt2 to group 3:sgt3 (monitored):
  denytcpudpicmp-10
  Deny IP-00

Details:
Role-based IP access list denytcpudpicmp-10 (downloaded)
  10 deny tcp
  20 deny udp
  30 deny icmp
Role-based IP access list Permit IP-00 (downloaded)
  10 permit ip

Device# show cts role-based counters ipv4

Role-based IPv4 counters
From      To      SW-Denied  HW-Denied  SW-Permitt  HW_Permitt  SW-Monitor  HW-Monitor
*         *       0          0          8           18962      0           0
2         3       0          0          0           0          0           341057

```

Example: Manually Configuring SGACL Policies

The following example shows how to manually configure SGACL policies:


```

Device> enable
Device# configure terminal
Device(config)# ip access role allow_webtraff
Device(config-rb-acl)# 10 permit tcp dst eq 80
Device(config-rb-acl)# 20 permit tcp dst eq 443
Device(config-rb-acl)# 30 permit icmp
Device(config-rb-acl)# 40 deny ip
Device(config-rb-acl)# exit
Device(config)# cts role-based permissions from 55 to 66 allow_webtraff

Device# show ip access allow_webtraff

Role-based IP access list allow_webtraff
 10 permit tcp dst eq www
 20 permit tcp dst eq 443
 30 permit icmp
 40 deny ip

Device# show cts role-based permissions from 2 to 5

Role-based permissions from group 2 to group 5:
srb2
srb5

```

Example: Manually Applying SGACLs

The following example shows how to manually apply SGACL policies:

```

Device> enable
Device# configure terminal
Device(config)# cts role-based permissions default MYDEFAULTSGACL
Device(config)# cts role-based permissions from 3 to 5 SRB3 SRB5
Device(config)# exit

```

Example: Displaying SGACL Policies

This example shows how to display the content of the SGACL policies permissions matrix for traffic sourced from security group 3:

```

Device> enable
Device# show cts role-based permissions from 3

Role-based permissions from group 3 to group 5:
  SRB3
  SRB5
Role-based permissions from group 3 to group 7:
  SRB4

```

Feature History for Security Group ACL Policies

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Security Group ACL Policies	Using SGACLs, you can control the operations that users can perform based on the security group assignments of users and destination resources.
Cisco IOS XE Amsterdam 17.3.1	Enhanced SGACL Logging	Enhanced ACL logging allows logging to be done at much higher rates than using the NetFlow hardware.
Cisco IOS XE Cupertino 17.7.1	Security Group ACL Policies	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 4

Cisco TrustSec SGACL High Availability

Cisco TrustSec Security Group access control lists (SGACLs) support the high availability functionality in switches that support the Cisco StackWise technology. This technology provides stateful redundancy and allows a switch stack to enforce and process access control entries.

- [Prerequisites for Cisco TrustSec SGACL High Availability, on page 49](#)
- [Restrictions for Cisco TrustSec SGACL High Availability, on page 49](#)
- [Information About Cisco TrustSec SGACL High Availability, on page 49](#)
- [Verifying Cisco TrustSec SGACL High Availability, on page 50](#)
- [Feature History for SGACL High Availability, on page 52](#)

Prerequisites for Cisco TrustSec SGACL High Availability

This document assumes the following:

- An understanding of Cisco TrustSec and the Security Group access control lists (SGACL) configuration.
- Devices are configured to function as a stack.
- All the devices in the stack are running an identical version of Cisco IOS XE software.

Restrictions for Cisco TrustSec SGACL High Availability

- When both active and standby switches fail simultaneously, stateful switchover of SGACL does not occur.

Information About Cisco TrustSec SGACL High Availability

Cisco TrustSec Security Group access control lists (SGACLs) support the high availability functionality in switches that support the Cisco StackWise technology. This technology provides stateful redundancy and allows a switch stack to enforce and process access control entries.

There is no Cisco TrustSec-specific configuration to enable this functionality, which is supported in Cisco IOS XE Denali 16.2.1 and later releases.

High Availability Overview

In a switch stack, the stack manager assigns the switch with the highest priority as the active switch, and the switch with the next highest priority as the standby switch. During an automatic or a CLI-based stateful switchover, the standby switch becomes the active switch and the switch with the next highest priority becomes the standby switch and so on.

Operation data is synchronized from the active switch to the standby switch, during initial system bootup, changes in the operational data (also called Change of Authorization [CoA]), or operational data refresh.

During a stateful switchover, the newly active switch, requests and downloads the operation data. The environment data (ENV-data) and the Role-Based access control lists (RBACLs) are not updated until the refresh time is complete.

The following operation data is downloaded to the active switch:

- Environment Data (ENV-data)—A variable length field that consists of the preferred server list to get the RBACL information at the time of refresh or initialization.
- Protected Access Credential (PAC)—A shared secret that is mutually and uniquely shared between the switch and the authenticator to secure an Extensible Authentication Protocol Flexible Authentication via the Secure Tunneling (EAP-FAST) tunnel.
- Role-Based Policy (RBACL or SGACL)—A variable-length role-based policy list that consists of policy definitions for all the Security Group Tag (SGT) mappings on the switch.



Note Cisco TrustSec credential that consists of the device ID and password details is run as a command on the active switch.

Verifying Cisco TrustSec SGACL High Availability

To verify the Cisco TrustSec SGACL high availability configuration, run the **show cts role-based permissions** command on both the active and standby switches. The output from the command must be the same on both switches.

The following is sample output from the **show cts role-based permissions** command on the active switch:

```
Device# show cts role-based permissions

IPv4 Role-based permissions default (monitored):
    default_sgACL-01
    Deny IP-00
IPv4 Role-based permissions from group 10:SGT_10 to group 15:SGT_15:
    SGACL_3-01
IPv4 Role-based permissions from group 14:SGT_14 to group 15:SGT_15:
    multiple_ace-14
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
```

The following is sample output from the **show cts role-based permissions** command on the standby switch:

```
Device-stby# show cts role-based permissions

IPv4 Role-based permissions default (monitored):
```

```

        default_sgacl-01
        Deny IP-00
IPv4 Role-based permissions from group 10:SGT_10 to group 15:SGT_15:
        SGACL_3-01
IPv4 Role-based permissions from group 14:SGT_14 to group 15:SGT_15:
        multiple_ace-14
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE

```

After a stateful switchover, run the following commands on the active switch to verify the feature:

The following is sample output from the **show cts pacs** command:

```

Device# show cts pacs

AID: A3B6D4D8353F102346786CF220FF151C
PAC-Info:
  PAC-type = Cisco Trustsec
  AID: A3B6D4D8353F102346786CF220FF151C
  I-ID: CTS_ED_21
  A-ID-Info: Identity Services Engine
  Credential Lifetime: 17:22:32 IST Mon Mar 14 2016
PAC-Opaque:
000200B80003000100040010A3B6D4D8353F102346786CF220FF151C0006009C00030100E044B2650D8351FD06
F23623C470511E0000001356DEA96C00093A80538898D40F633C368B053200D4C9D2422A7FEB4837EA9DDBB89D1
E51DA4E7B184E66D3D5F2839C11E5FB386936BB85250C61CA0116FDD9A184C6E96593EEAF5C39BE08140AFBB19
4EE701A0056600CFF5B12C02DD7ECEAA3CCC8170263669C483BD208052A46C31E39199830F794676842ADEECBB
A30FC4A5A0DEDA93
Refresh timer is set for 01:00:05

```

The following is sample output from the **show cts environment-data** command:

```

Device# show cts environment-data

CTS Environment Data
=====
Current state = COMPLETE
Last status = Successful
Local Device SGT:
  SGT tag = 0:Unknown
Server List Info:
Installed list: CTSServerList1-000D, 1 server(s):
  *Server: 10.78.105.47, port 1812, A-ID A3B6D4D8353F102346786CF220FF151C
  Status = ALIVE
  auto-test = FALSE, keywrap-enable = FALSE, idle-time = 60 mins, deadtime = 20 secs
Multicast Group SGT Table:
Security Group Name Table:
0001-45 :
  0-00:Unknown
  2-ba:SGT_2
  3-00:SGT_3
  4-00:SGT_4
  5-00:SGT_5
  6-00:SGT_6
  7-00:SGT_7
  8-00:SGT_8
  9-00:SGT_9
  10-16:SGT_10
!
!
!
Environment Data Lifetime = 3600 secs
Last update time = 14:32:53 IST Mon Mar 14 2016
Env-data expires in 0:00:10:04 (dd:hr:mm:sec)

```

```
Env-data refreshes in 0:00:10:04 (dd:hr:mm:sec)
Cache data applied = NONE
State Machine is running
```

The following is sample output from the **show cts role-based permissions** command after a stateful switchover:

```
Device# show cts role-based permissions

IPv4 Role-based permissions default:
    default_sgACL-01
    Deny IP-00
IPv4 Role-based permissions from group 10:SGT_10 to group 15:SGT_15:
    SGACL_3-01
IPv4 Role-based permissions from group 14:SGT_14 to group 15:SGT_15:
    multiple_ace-14
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
```

Feature History for SGACL High Availability

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	SGACL High Availability	Cisco TrustSec SGACLs support the high availability functionality in switches that support the Cisco StackWise technology. This technology provides stateful redundancy and allows a switch stack to enforce and process access control entries.
Cisco IOS XE Cupertino 17.7.1	SGACL High Availability	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 5

Configuring SGT Exchange Protocol

You can use the SGT Exchange Protocol (SXP) to propagate the Security Group Tags (SGTs) across network devices that do not have hardware support for Cisco TrustSec. This module describes how to configure Cisco TrustSec SXP on switches in your network.

Cisco TrustSec builds secure networks by establishing domains of trusted network devices. Each device in the domain is authenticated by its peers. Communication on the links between devices in the domain is secured with a combination of encryption, message integrity check, and data-path replay protection mechanisms.

The Security Group Tag (SGT) Exchange Protocol (SXP) is one of several protocols that supports CTS and is referred to in this document as Cisco TrustSec-SXP. Cisco TrustSec-SXP is a control protocol for propagating IP-to-SGT binding information across network devices that do not have the capability to tag packets. Cisco TrustSec-SXP passes IP to SGT bindings from authentication points to upstream devices in the network. This process allows security services on switches, routers, or firewalls to learn identity information from access devices.

- [Prerequisites for SGT Exchange Protocol, on page 53](#)
- [Restrictions for SGT Exchange Protocol, on page 54](#)
- [Information About SGT Exchange Protocol, on page 54](#)
- [How to Configure SGT Exchange Protocol, on page 55](#)
- [Configuration Examples for SGT Exchange Protocol, on page 60](#)
- [Verifying SGT Exchange Protocol Connections, on page 61](#)
- [Feature History for SGT Exchange Protocol, on page 62](#)

Prerequisites for SGT Exchange Protocol

The Cisco TrustSec-SGT Over Exchange Protocol (SXP) network needs to be established before implementing SXP. This network has the following prerequisites:

- To use the Cisco TrustSec functionality on your existing router, ensure that you have purchased a Cisco TrustSec security license. If the router is being ordered and needs the Cisco TrustSec functionality, ensure that this license is pre-installed on your router before it is shipped to you
- Cisco TrustSec SXP software must run on all network devices.
- Connectivity should exist between all network devices.
- The Cisco Identity Services Engine 1.0 is required for authentication. The Secure Access Control Server (ACS) Express Appliance server can also be used for authentication, however not all ACS features are supported by Cisco TrustSec. ACS 5.1 operates with a Cisco TrustSec-SXP license

- Configure the **retry open timer** command to a different value on different routers

Restrictions for SGT Exchange Protocol

- Cisco TrustSec Exchange Protocol is not supported on logical interfaces; supported only on physical interfaces.
- In Cisco IOS XE Everest 16.6.4 and later releases, when the Dynamic Host Control Protocol (DHCP) snooping is enabled, Cisco TrustSec enforcement for DHCP packets are bypassed by enforcement polices.

Information About SGT Exchange Protocol

This section provides information about SGT Exchange Protocol.

SGT Exchange Protocol Overview

Cisco TrustSec builds secure networks by establishing domains of trusted network devices. Each device in the domain is authenticated by its peers. Communication on the links between devices in the domain is secured with a combination of encryption, message integrity check, and data-path replay protection mechanisms.

The Security Group Tag (SGT) Exchange Protocol (SXP) is one of several protocols that supports Cisco TrustSec. SXP is a control protocol for propagating IP-to-SGT binding information across network devices that do not have the capability to tag packets. Cisco TrustSec filters packets at the egress interface. During endpoint authentication, a host accessing the Cisco TrustSec domain (the endpoint IP address) is associated with an SGT at the access device through Dynamic Host Control Protocol (DHCP) snooping and IP device tracking. The access device transmits that association or binding through SXP to Cisco TrustSec hardware-capable egress devices. These devices maintain a table of source IP-to-SGT bindings. Packets are filtered on the egress interface by Cisco TrustSec hardware-capable devices by applying security group access control lists (SGACLs). SXP passes IP-to-SGT bindings from authentication points to upstream devices in the network. This process allows security services on switches, routers, or firewalls to learn identity information from access devices.

SGTs can be assigned through any of the following Endpoint Admission Control (EAC) access methods:

- 802.1X port-based authentication
- MAC Authentication Bypass (MAB)
- Web Authentication

SXP uses TCP as the transport protocol, and the TCP port 64999 for connection initiation. SXP uses Message Digest 5 (MD5) for authentication and integrity check. It has two defined roles—speaker (initiator) and listener (receiver).

Security Group Tagging

Security Group Tag is a unique 16 bit tag that is assigned to a unique role. It represents the privilege of the source user, device, or entity and is tagged at the ingress of the Cisco TrustSec domain. SXP uses the device and user credentials acquired during authentication for classifying packets by security groups (SGs) as they

enter a network. This packet classification is maintained by tagging packets on the ingress to the Cisco TrustSec network so that they can be identified for the purpose of applying security and other policy criteria along the data path. The Security Group Tag (SGT) allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic. Static port Identification is used to lookup the SGT value for a particular endpoint connected to a port.

SGT Assignment

The Security Group Tag (SGT) of a packet can be assigned at the port level when the packet comes tagged on a Cisco TrustSec link, or when a single endpoint authenticates on a port. SGT of an incoming packet is determined in the following ways:

- When a packet that is tagged with an SGT comes on a trust port, the tag of the packet is considered as the SGT of the packet.
- When a packet is tagged with an SGT, but comes on an untrusted port, the SGT of the packet is ignored and the peer SGT is configured for the port.
- When a packet does not have an SGT, the peer SGT is configured for a port.

The following methods of assigning SGTs are supported:

- IPM (dot1x, MAB, and Web Authentication)
- VLAN-to-SGT mapping Established when an authentication method provides an SGT for an authenticated entry already has an assigned IP address. A device process monitors endpoint sessions and detects changes or removal of IP-to-SGT binding.
- SXP (SGT Exchange Protocol) Listener

How to Configure SGT Exchange Protocol

This section describes how to configure SGT Exchange Protocol.

Configuring a Device SGT Manually

In normal Cisco TrustSec operation, the authentication server assigns an SGT to the device for packets originating from the device. You can manually configure an SGT to be used if the authentication server is not accessible, but an authentication server-assigned SGT will take precedence over a manually-assigned SGT.

To manually configure an SGT on the device, perform this task:

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Device# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	cts sgt tag Example: Device(config)# cts sgt tag	Configures the SGT for packets sent from the device. The tag argument is in decimal format. The range is 1 to 65533.
Step 3	exit Example: Device(config)# exit	Exits configuration mode.

Configuring an SXP Peer Connection

You must configure the SXP peer connection on both of the devices. One device is the speaker and the other is the listener. When using password protection, make sure to use the same password on both ends.



Note If a default SXP source IP address is not configured and you do not configure an SXP source address in the connection, the Cisco TrustSec software derives the SXP source IP address from existing local IP addresses. The SXP source address might be different for each TCP connection initiated from the device.

To configure an SXP peer connection, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts sxp connection peer peer-ipv4-addr[source src-ipv4-addr] password {default none} mode {local peer} {speaker listener} {vrf vrf-name} Example: Device(config)# cts sxp connection peer 10.10.1.1 password default mode local listener	Configures the SXP address connection. The optional source keyword specifies the IPv4 address of the source device. If no address is specified, the connection will use the default source address, if configured, or the address of the port. The password keyword specifies the password that SXP will use for the connection using the following options:

	Command or Action	Purpose
		<ul style="list-style-type: none"> • default—Use the default SXP password you configured using the cts sxp default password command. • none—Do not use a password. <p>The mode keyword specifies the role of the remote peer device:</p> <ul style="list-style-type: none"> • local—The specified mode refers to the local device. • peer—The specified mode refers to the peer device. • speaker—Default. Specifies that the device is the speaker in the connection. • listener—Specifies that the device is the listener in the connection. <p>The optional vrf keyword specifies the VRF to the peer. The default is the default VRF.</p>
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode
Step 5	show cts sxp connections Example: Device# show cts sxp connections	(Optional) Displays the SXP connection information.

Configuring the Default SXP Password

By default, SXP uses no password when setting up connections.

To configure a default SXP password, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 3	cts sxp default password [0 6 7] <i>password</i> Example: Device(config)# cts sxp default password 0 hello	Configures the SXP default password. You can enter either a clear text password (using the 0 or no option) or an encrypted password (using the 6 or 7 option). The maximum password length is 32 characters.
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode

Configuring the Default SXP Source IP Address

SXP uses the default source IP address for all new TCP connections where a source IP address is not specified. There is no effect on existing TCP connections when you configure the default SXP source IP address.

To configure a default SXP source IP address, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts sxp default source-ip <i>src-ip-addr</i> Example: Device(config)# cts sxp default source-ip 10.0.1.2	Configures the SXP default source IP address.
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode.

Changing the SXP Reconciliation Period

After a peer terminates an SXP connection, an internal hold-down timer starts. If the peer reconnects before the internal hold-down timer expires, the SXP reconciliation period timer starts. While the SXP reconciliation period timer is active, the Cisco TrustSec software retains the SGT mapping entries learned from the previous connection and removes invalid entries. The default value is 120 seconds (2 minutes). Setting the SXP reconciliation period to 0 seconds disables the timer and causes all entries from the previous connection to be removed.

To change the SXP reconciliation period, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts sxp reconciliation period <i>seconds</i> Example: Device(config)# cts sxp reconciliation period 360	Changes the SXP reconciliation timer. The default value is 120 seconds (2 minutes). The range is from 0 to 64000.
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode.

Changing the SXP Retry Period

The SXP retry period determines how often the Cisco TrustSec software retries an SXP connection. When an SXP connection is not successfully set up, the Cisco TrustSec software makes a new attempt to set up the connection after the SXP retry period timer expires. The default value is 120 seconds. Setting the SXP retry period to 0 seconds disables the timer and retries are not attempted.

To change the SXP retry period, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts sxp retry period <i>seconds</i> Example: Device(config)# cts sxp retry period 360	Changes the SXP retry timer. The default value is 120 seconds (2 minutes). The range is from 0 to 64000.

	Command or Action	Purpose
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode.

Creating Syslogs to Capture Changes of IP Address-to-SGT Mapping Learned Through SXP

When the **cts sxp log binding-changes** command is configured in global configuration mode, SXP syslogs (sev 5 syslog) are generated whenever a change to IP address to SGT binding occurs (add, delete, change). These changes are learned and propagated on the SXP connection. The default is **no cts sxp log binding-changes**.

To enable logging of binding changes, perform the following task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts sxp log binding-changes Example: Device(config)# cts sxp log binding-changes	Enables logging for IP to SGT binding changes.
Step 4	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode.

Configuration Examples for SGT Exchange Protocol

The following sections show configuration examples of SGT Exchange Protocol:

Example: Enabling Cisco TrustSec SXP and an SXP Peer Connection

The following example shows how to enable SXP and configure an SXP peer connection between device A, the speaker, and device B, the listener:

```

Device# configure terminal
Device(config)# cts sxp enable
Device(config)# cts sxp default password Cisco123
Device(config)# cts sxp default source-ip 10.10.1.1
Device(config)# cts sxp connection peer 10.20.2.2 password default mode local speaker

```

The following example shows how to configure the SXP peer connection between device B, the listener, and device A, the speaker:

```

Device# configure terminal
Device(config)# cts sxp enable
Device(config)# cts sxp default password Cisco123
Device(config)# cts sxp default source-ip 10.20.2.2
Device(config)# cts sxp connection peer 10.10.1.1 password default mode local listener

```

Example: Configuring the Default SXP Password and Source IP Address

The following example shows how to configure a default SXP password and source IP address:

```

Device# configure terminal
Device(config)# cts sxp default password Cisco123
Device(config)# cts sxp default source-ip 10.20.2.2
Device(config)# end

```

Verifying SGT Exchange Protocol Connections

To view SXP connections, perform this task:

Command	Purpose
show cts sxp connections	Displays detailed information about the SXP status and connections.
show cts sxp connections [brief]	Displays brief information about the SXP status and connections.

The following is sample output from the **show cts sxp connections** command:

```

Device# show cts sxp connections

SXP                               : Enabled
Default Password                   : Set
Default Source IP                  : 10.10.1.1
Connection retry open period: 10 secs
Reconcile period                   : 120 secs
Retry open timer is not running
-----
Peer IP                             : 10.20.2.2
Source IP                           : 10.10.1.1
Conn status                          : On
Conn Version                         : 2
Connection mode                     : SXP Listener
Connection inst#                    : 1

```

```
TCP conn fd          : 1
TCP conn password    : default SXP password
Duration since last state change: 0:00:21:25 (dd:hr:mm:sec)
Total num of SXP Connections = 1
```

The following is sample output from the **show cts sxp connections brief** command:

```
Device# show cts sxp connections brief

SXP                : Enabled
Default Password   : Set
Default Source IP  : Not Set
Connection retry open period: 120 secs
Reconcile period   : 120 secs
Retry open timer is not running
-----
Peer_IP           Source_IP           Conn Status         Duration
-----
10.1.3.1          10.1.3.2           On                  6:00:09:13 (dd:hr:mm:sec)
Total num of SXP Connections = 1
```

Feature History for SGT Exchange Protocol

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	SGT Exchange Protocol	The SXP propagates the SGTs across network devices that do not have hardware support for Cisco TrustSec.
Cisco IOS XE Cupertino 17.7.1	SGT Exchange Protocol	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 6

Configuring Security Group Tag Mapping

Subnet to security group tag (SGT) mapping binds an SGT to all host addresses of a specified subnet. Once this mapping is implemented, Cisco TrustSec imposes the SGT on any incoming packet that has a source IP address which belongs to the specified subnet.

- [Restrictions for SGT Mapping, on page 63](#)
- [Information About SGT Mapping, on page 63](#)
- [How to Configure SGT Mapping, on page 66](#)
- [Verifying SGT Mapping, on page 72](#)
- [Configuration Examples for SGT Mapping, on page 74](#)
- [Feature History for Security Group Tag Mapping, on page 77](#)

Restrictions for SGT Mapping

Restrictions for Subnet-to-SGT Mapping

- An IPv4 subnetwork with a /31 prefix cannot be expanded.
- Subnet host addresses cannot be bound to Security Group Tags (SGT)s when the **network-map bindings** parameter is less than the total number of subnet hosts in the specified subnets, or when bindings is 0.
- IPv6 expansions and propagation only occurs when Security Exchange Protocol (SXP) speaker and listener are running SXPv3, or more recent versions.

Restriction for Default Route SGT Mapping

- Default route configuration is accepted only with the subnet /0. Entering only the host-ip without the subnet /0 displays the following message:

```
Device(config)#cts role-based sgt-map 0.0.0.0 sgt 1000
Default route configuration is not supported for host ip
```

Information About SGT Mapping

This section provides information about SGT mapping.

Overview of Subnet-to-SGT Mapping

Subnet-to-SGT mapping binds an SGT to all host addresses of a specified subnet. Cisco TrustSec imposes the SGT on an incoming packet when the packet's source IP address belongs to the specified subnet. The subnet and SGT are specified in the CLI with the **cts role-based sgt-map** *net_address/prefix sgt sgt_number* global configuration command. A single host may also be mapped with this command.

In IPv4 networks, Security Exchange Protocol (SXP)v3, and more recent versions, can receive and parse subnet *net_address/prefix* strings from SXPv3 peers. Earlier SXP versions convert the subnet prefix into its set of host bindings before exporting them to an SXP listener peer.

For example, the IPv4 subnet 192.0.2.0/24 is expanded as follows (only 3 bits for host addresses):

- Host addresses 198.0.2.1 to 198.0.2.7—tagged and propagated to SXP peer.
- Network and broadcast addresses 198.0.2.0 and 198.0.2.8—not tagged and not propagated.

To limit the number of subnet bindings SXPv3 can export, use the **cts sxp mapping network-map** global configuration command.

Subnet bindings are static, there is no learning of active hosts. They can be used locally for SGT imposition and SGACL enforcement. Packets tagged by subnet-to-SGT mapping can be propagated on Layer 2 or Layer 3 Cisco TrustSec links.

For IPv6 networks, SXPv3 cannot export subnet bindings to SXPv2 or SXPv1 peers.

Overview of VLAN-to-SGT Mapping

The VLAN-to-SGT mapping feature binds an SGT to packets from a specified VLAN. This simplifies the migration from legacy to Cisco TrustSec-capable networks as follows:

- Supports devices that are not Cisco TrustSec-capable but are VLAN-capable, such as, legacy switches, wireless controllers, access points, VPNs, etc.
- Provides backward compatibility for topologies where VLANs and VLAN ACLs segment the network, such as, server segmentation in data centers.

The VLAN-to-SGT binding is configured with the **cts role-based sgt-map vlan-list** global configuration command.

When a VLAN is assigned a gateway that is a switched virtual interface (SVI) on a Cisco TrustSec-capable switch, and IP Device Tracking is enabled on that switch, then Cisco TrustSec can create an IP-to-SGT binding for any active host on that VLAN mapped to the SVI subnet.

IP-SGT bindings for the active VLAN hosts are exported to SXP listeners. The bindings for each mapped VLAN are inserted into the IP-to-SGT table associated with the VRF the VLAN is mapped to by either its SVI or by the **cts role-based l2-vrf** command.

VLAN-to-SGT bindings have the lowest priority of all binding methods and are ignored when bindings from other sources are received, such as from SXP or CLI host configurations. Binding priorities are listing in the Binding Source Priorities section.

Overview of Layer 3 Logical Interface-to-SGT Mapping (L3IF-SGT Mapping)

L3IF-SGT mapping can directly map SGTs to traffic of any of the following Layer 3 interfaces regardless of the underlying physical interface:

- Routed port
- SVI (VLAN interface)
- Layer 3 subinterface of a Layer 2 port
- Tunnel interface

Use the **cts role-based sgt-map interface** global configuration command to specify either a specific SGT number, or a Security Group Name (whose SGT association is dynamically acquired from a Cisco ISE or a Cisco ACS access server).

In cases where Identity Port Mapping (cts interface manual sub mode configuration) and L3IF-SGT require different IP to SGT bindings, IPM takes precedence. All other conflicts among IP to SGT binding are resolved according to the priorities listing in the Binding Source Priorities section.

Binding Source Priorities

Cisco TrustSec resolves conflicts among IP-SGT binding sources with a strict priority scheme. For example, an SGT may be applied to an interface with the **policy {dynamic identity peer-name | static sgt tag}** Cisco Trustsec Manual interface mode command (Identity Port Mapping). The current priority enforcement order, from lowest (1) to highest (7), is as follows:

1. VLAN—Bindings learned from snooped ARP packets on a VLAN that has VLAN-SGT mapping configured.
2. CLI— Address bindings configured using the IP-SGT form of the cts role-based sgt-map global configuration command.
3. Layer 3 Interface—(L3IF) Bindings added due to FIB forwarding entries that have paths through one or more interfaces with consistent L3IF-SGT mapping or Identity Port Mapping on routed ports.
4. SXP—Bindings learned from SXP peers.
5. IP_ARP—Bindings learned when tagged ARP packets are received on a CTS capable link.
6. LOCAL—Bindings of authenticated hosts which are learned via EPM and device tracking. This type of binding also include individual hosts that are learned via ARP snooping on L2 [I]PM configured ports.
7. INTERNAL—Bindings between locally configured IP addresses and the device own SGT.

Default Route SGT

Default Route Security Group Tag (SGT) assigns an SGT number to default routes.

Default Route is that route which does not match a specified route and therefore is the route to the last resort destination. Default routes are used to direct packets addressed to networks not explicitly listed in the routing table.

How to Configure SGT Mapping

This section describes how to configure SGT mapping.

Configuring a Device SGT Manually

In normal Cisco TrustSec operation, the authentication server assigns an SGT to the device for packets originating from the device. You can manually configure an SGT to be used if the authentication server is not accessible, but an authentication server-assigned SGT will take precedence over a manually-assigned SGT.

To manually configure an SGT on the device, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# <code>enable</code>	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	cts sgt tag Example: Device(config)# <code>cts sgt 1234</code>	Enables SXP for Cisco TrustSec.
Step 4	exit Example: Device(config)# <code>exit</code>	Exits global configuration mode and returns to privileged EXEC mode

Configuring Subnet-to-SGT Mapping

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# <code>enable</code>	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 3	cts sxp mapping network-map <i>bindings</i> Example: <pre>Device(config)# cts sxp mapping network-map 10000</pre>	<ul style="list-style-type: none"> Configures the Subnet to SGT Mapping host count constraint. The <i>bindings</i> argument specifies the maximum number of subnet IP hosts that can be bound to SGTs and exported to the SXP listener. <i>bindings</i>—(0 to 65,535) default is 0 (no expansions performed)
Step 4	cts role-based sgt-map <i>ipv4_address/prefix</i> sgt <i>number</i> Example: <pre>Device(config)# cts role-based sgt-map 10.10.10.10/29 sgt 1234</pre>	<p>(IPv4) Specifies a subnet in CIDR notation.</p> <ul style="list-style-type: none"> Use the <i>no</i> form of the command to unconfigure the Subnet to SGT mapping. The number of bindings specified in Step 2 should match or exceed the number of host addresses in the subnet (excluding network and broadcast addresses). The <i>sgt number</i> keyword specifies the Security Group Tag to be bound to every host address in the specified subnet. <i>ipv4_address</i>—Specifies the IPv4 network address in dotted decimal notation. <i>prefix</i>—(0 to 30) Specifies the number of bits in the network address. <i>sgt number</i>—(0–65,535) Specifies the Security Group Tag (SGT) number.
Step 5	cts role-based sgt-map <i>ipv6_address::prefix</i> sgt <i>number</i> Example: <pre>Device(config)# cts role-based sgt-map 2020::/64 sgt 1234</pre>	<p>(IPv6) Specifies a subnet in colon hexadecimal notation. Use the <i>no</i> form of the command to unconfigure the Subnet to SGT mapping.</p> <p>The number of bindings specified in Step 2 should match or exceed the number of host addresses in the subnet (excluding network and broadcast addresses). The <i>sgt number</i> keyword specifies the Security Group Tag to be bound to every host address in the specified subnet.</p> <ul style="list-style-type: none"> <i>ipv6_address</i>—Specifies IPv6 network address in colon hexadecimal notation. <i>prefix</i>—(0 to 128) Specifies the number of bits in the network address. <i>sgt number</i>—(0–65,535) Specifies the Security Group Tag (SGT) number.
Step 6	exit Example:	Exits global configuration mode and returns to privileged EXEC mode..

	Command or Action	Purpose
	Device (config)# exit	

Configuring VLAN-to-SGT Mapping

Task Flow for Configuring VLAN-SGT Mapping on a Cisco TrustSec device.

- Create a VLAN on the device with the same VLAN_ID of the incoming VLAN.
- Create an SVI for the VLAN on the device to be the default gateway for the endpoint clients.
- Configure the device to apply an SGT to the VLAN traffic.
- Enable IP Device tracking on the device.
- Attach a device tracking policy to a VLAN.



Note In a multi-switch network, SISF-based device tracking provides the capability to distribute binding table entries between switches running the feature. This assumes that binding entries are created on the switches where the host appears on an access port, and no entry is created for a host that appears over a trunk port. To achieve this in a multi-switch setup, we recommend that you configure another policy and attach it to the trunk port, as described in the *Configuring a Multi-Switch Network to Stop Creating Binding Entries from a Trunk Port* procedure, in the *Configuring SISF-Based Device Tracking* chapter of the *Security Configuration Guide*.

- Verify that VLAN-to-SGT mapping occurs on the device.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	vlan <i>vlan_id</i> Example: Device (config)# vlan 100	Creates VLAN 100 on the TrustSec-capable gateway device and enters VLAN configuration mode.
Step 4	[no] shutdown Example: Device (config-vlan)# no shutdown	Provisions VLAN 100.

	Command or Action	Purpose
Step 5	exit Example: Device(config-vlan)# exit	Exits VLAN configuration mode and returns to global configuration mode.
Step 6	interface <i>type slot/port</i> Example: Device(config)# interface vlan 100	Specifies the interface type and enters interface configuration mode.
Step 7	ip address <i>slot/port</i> Example: Device(config-if)# ip address 10.1.1.2 255.0.0.0	Configures Switched Virtual Interface (SVI) for VLAN 100.
Step 8	[no] shutdown Example: Device(config-if)# no shutdown	Enables the SVI.
Step 9	exit Example: Device(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 10	cts role-based sgt-map vlan-list <i>vlan_id sgt sgt_number</i> Example: Device(config)# cts role-based sgt-map vlan-list 100 sgt 10	Assigns the specified SGT to the specified VLAN.
Step 11	device-tracking policy <i>policy-name</i> Example: Device(config)# device-tracking policy policy1	Specifies the policy and enters device-tracking policy configuration mode.
Step 12	tracking enable Example: Device(config-device-tracking)# tracking enable	Overrides the default device tracking settings for the policy attribute.
Step 13	exit Example: Device(config-device-tracking)# exit	Exits device-tracking policy configuration mode and returns to global configuration mode.
Step 14	vlan configuration <i>vlan_id</i> Example: Device(config)# vlan configuration 100	Specifies the VLAN to which the device tracking policy will be attached, and enters the VLAN configuration mode.

	Command or Action	Purpose
Step 15	device-tracking attach-policy <i>policy-name</i> Example: Device (config-vlan-config) # device-tracking attach-policy policy1	Attaches a device tracking policy to the specified VLAN.
Step 16	end Example: Device (config-vlan-config) # end	Exits VLAN configuration mode and returns to privileged EXEC mode.
Step 17	show cts role-based sgt-map { <i>ipv4_netaddr</i> <i>ipv4_netaddr/prefix</i> <i>ipv6_netaddr</i> <i>ipv6_netaddr/prefix</i> all [ipv4 ipv6] host { <i>ipv4_addr</i> <i>ipv6_addr</i> } summary [ipv4 ipv6] Example: Device# show cts role-based sgt-map all	(Optional) Displays the VLAN-to-SGT mappings.
Step 18	show device-tracking policy <i>policy-name</i> Example: Device# show device-tracking policy policy1	(Optional) Displays the current policy attributes.

Configuring L3IF-to-SGT Mapping

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts role-based sgt-map interface <i>type slot/port</i> [<i>security-group name</i> <i>sgt number</i>] Example: Device (config) # cts role-based sgt-map interface gigabitEthernet 1/1 sgt 77	An SGT is imposed on ingress traffic to the specified interface. <ul style="list-style-type: none"> • interface <i>type slot/port</i> —Displays list of available interfaces. • security-group <i>name</i>— Security Group name to SGT pairings are configured on the Cisco ISE or Cisco ACS.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • <i>sgt number</i> —(0 to 65,535). Specifies the Security Group Tag (SGT) number.
Step 4	exit Example: Device(config)# exit	Exits configuration mode.
Step 5	show cts role-based sgt-map all Example: Device# cts role-based sgt-map all	Verify that ingressing traffic is tagged with the specified SGT.

Emulating the Hardware Keystore

In cases where a hardware keystore is not present or is unusable, you can configure the switch to use a software emulation of the keystore. To configure the use of a software keystore, perform this task:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device# enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts keystore emulate Example: Device(config)# cts keystore emulate	Configures the switch to use a software emulation of the keystore instead of the hardware keystore.
Step 4	exit Example: Device(config)# exit	Exits configuration mode.
Step 5	show keystore Example: Device# show keystore	Displays the status and contents of the keystore. The stored secrets are not displayed.

Configuring Default Route SGT

Before you begin

Ensure that you have already created a default route on the device using the **ip route 0.0.0.0** command. Otherwise, the default route (which comes with the Default Route SGT) gets an unknown destination and therefore the last resort destination will point to CPU.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts role-based sgt-map 0.0.0.0/0 sgt number Example: Device(config)# cts role-based sgt-map 0.0.0.0/0 sgt 3	Specifies the SGT number for the default route. Valid values are from 0 to 65,519. <p>Note</p> <ul style="list-style-type: none"> • The host_address/subnet can be either IPv4 address (0.0.0.0/0) or IPv6 address (0:0::/0) • The default route configuration is accepted only with the subnet /0. Entering only the host-ip without the subnet /0 displays the following message: <pre>Device(config)#cts role-based sgt-map 0.0.0.0 sgt 1000 Default route configuration is not supported for host ip</pre>
Step 4	exit Example: Device(config)# exit	Exits global configuration mode.

Verifying SGT Mapping

The following sections show how to verify SGT mapping:

Verifying Subnet-to-SGT Mapping Configuration

To display Subnet-to-SGT Mapping configuration information, use one of the following show commands:

Command	Purpose
<code>show cts sxp connections</code>	Displays the SXP speaker and listener connections with their operational status.
<code>show cts sxp sgt-map</code>	Displays the IP to SGT bindings exported to the SXP listeners.
<code>show running-config</code>	Verifies that the subnet-to-SGT configurations commands are in the running configuration file.

Verifying VLAN-to-SGT Mapping

To display VLAN-to-SGT configuration information, use the following show commands:

Table 1:

Command	Purpose
<code>show device-tracking policy</code>	Displays the current policy attributes of the device tracking policy.
<code>show cts role-based sgt-map</code>	Displays IP address-to-SGT bindings.

Verifying L3IF-to-SGT Mapping

To display L3IF-to-SGT configuration information, use the following show command:

Command	Purpose
<code>show cts role-based sgt-map all</code>	Displays all IP address-to-SGT bindings..

Verifying Default Route SGT Configuration

Verify the Default Route SGT configuration:

```
device# show role-based sgt-map all
Active IPv4-SGT Bindings Information

IP Address          SGT      Source
=====
0.0.0.0/0           3        CLI
11.0.0.0/8          11       CLI
11.0.0.10           1110    CLI
11.1.1.1            1111    CLI
21.0.0.2            212     CLI

IP-SGT Active Bindings Summary
=====
```

```
Total number of CLI bindings = 5
Total number of active bindings = 5
```

Configuration Examples for SGT Mapping

The following sections show configuration examples of SGT mapping:

Example: Configuring a Device SGT Manually

```
Device# configure terminal
Device(config)# cts sgt 1234
Device(config)# exit
```

Example: Configuration for Subnet-to-SGT Mapping

The following example shows how to configure IPv4 Subnet-to-SGT Mapping between devices running SXPv3 (Device1 and Device2):

1. Configure SXP speaker/listener peering between devices.

```
Device1# configure terminal
Device1(config)# cts sxp enable
Device1(config)# cts sxp default source-ip 1.1.1.1
Device1(config)# cts sxp default password 1syzygy1
Device1(config)# cts sxp connection peer 2.2.2.2 password default mode local speaker
```

2. Configure Device2 as SXP listener of Device1.

```
Device2(config)# cts sxp enable
Device2(config)# cts sxp default source-ip 2.2.2.2
Device2(config)# cts sxp default password 1syzygy1
Device2(config)# cts sxp connection peer 1.1.1.1 password default mode local listener
```

3. On Device2, verify that the SXP connection is operating:

```
Device2# show cts sxp connections brief | include 1.1.1.1
1.1.1.1          2.2.2.2          On          3:22:23:18
(dd:hr:mm:sec)
```

4. Configure the subnetworks to be expanded on Device1.

```
Device1(config)# cts sxp mapping network-map 10000
Device1(config)# cts role-based sgt-map 10.10.10.0/30 sgt 101
Device1(config)# cts role-based sgt-map 11.11.11.0/29 sgt 11111
Device1(config)# cts role-based sgt-map 192.168.1.0/28 sgt 65000
```

5. On Device2, verify the subnet-to-SGT expansion from Device1. There should be two expansions for the 10.10.10.0/30 subnetwork, six expansions for the 11.11.11.0/29 subnetwork, and 14 expansions for the 192.168.1.0/28 subnetwork.

```
Device2# show cts sxp sgt-map brief | include 101|11111|65000
IPv4,SGT: <10.10.10.1 , 101>
IPv4,SGT: <10.10.10.2 , 101>
IPv4,SGT: <11.11.11.1 , 11111>
IPv4,SGT: <11.11.11.2 , 11111>
IPv4,SGT: <11.11.11.3 , 11111>
IPv4,SGT: <11.11.11.4 , 11111>
```

```

IPv4,SGT: <11.11.11.5 , 11111>
IPv4,SGT: <11.11.11.6 , 11111>
IPv4,SGT: <192.168.1.1 , 65000>
IPv4,SGT: <192.168.1.2 , 65000>
IPv4,SGT: <192.168.1.3 , 65000>
IPv4,SGT: <192.168.1.4 , 65000>
IPv4,SGT: <192.168.1.5 , 65000>
IPv4,SGT: <192.168.1.6 , 65000>
IPv4,SGT: <192.168.1.7 , 65000>
IPv4,SGT: <192.168.1.8 , 65000>
IPv4,SGT: <192.168.1.9 , 65000>
IPv4,SGT: <192.168.1.10 , 65000>
IPv4,SGT: <192.168.1.11 , 65000>
IPv4,SGT: <192.168.1.12 , 65000>
IPv4,SGT: <192.168.1.13 , 65000>
IPv4,SGT: <192.168.1.14 , 65000>

```

6. Verify the expansion count on Device1:

```

Device1# show cts sxp sgt-map
      IP-SGT Mappings expanded:22
      There are no IP-SGT Mappings

```

7. Save the configurations on Device1 and Device2 and exit global configuration mode.

```

Device1(config)# copy running-config startup-config
Device1(config)# exit
Device2(config)# copy running-config startup-config
Device2(config)# exit

```

Example: Configuration for VLAN-to-SGT Mapping for a Single Host Over an Access Link

In the following example, a single host connects to VLAN 100 on an access device. A switched virtual interface on the TrustSec device is the default gateway for the VLAN 100 endpoint (IP Address 10.1.1.1). The TrustSec device imposes Security Group Tag (SGT) 10 on packets from VLAN 100.

1. Create VLAN 100 on an access device.

```

access_device# configure terminal
access_device(config)# vlan 100
access_device(config-vlan)# no shutdown
access_device(config-vlan)# exit
access_device(config)#

```

2. Configure the interface to the TrustSec device as an access link. Configurations for the endpoint access port are omitted in this example.

```

access_device(config)# interface gigabitEthernet 6/3
access_device(config-if)# switchport
access_device(config-if)# switchport mode access
access_device(config-if)# switchport access vlan 100

```

3. Create VLAN 100 on the TrustSec device.

```

TS_device(config)# vlan 100
TS_device(config-vlan)# no shutdown
TS_device(config-vlan)# end
TS_device#

```

4. Create an SVI as the gateway for incoming VLAN 100.

```
TS_device(config)# interface vlan 100
TS_device(config-if)# ip address 10.1.1.2 255.0.0.0
TS_device(config-if)# no shutdown
TS_device(config-if)# end
TS_device(config)#
```

5. Assign Security Group Tag (SGT) 10 to hosts on VLAN 100.

```
TS_device(config)# cts role-based sgt-map vlan 100 sgt 10
```

6. Enable IP Device Tracking on the TrustSec device. Verify that it is operating.

```
TS_device(config)# ip device tracking
TS_device# show ip device tracking all
```

```
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 100
```

```
-----
   IP Address      MAC Address   Vlan   Interface      STATE
-----
Total number interfaces enabled: 1
Vlan100
```

7. (Optional) PING the default gateway from an endpoint (in this example, host IP Address 10.1.1.1). Verify that SGT 10 is being mapped to VLAN 100 hosts.

```
TS_device# show cts role-based sgt-map all
```

```
Active IP-SGT Bindings Information
```

```
IP Address      SGT      Source
=====
10.1.1.1        10       VLAN
```

```
IP-SGT Active Bindings Summary
```

```
=====
Total number of VLAN bindings = 1
Total number of CLI bindings = 0
Total number of active bindings = 1
```

Example: Configuration for L3IF-to-SGT Mapping on an Ingress Port

In the following example a Layer 3 interface of a device linecard is configured to tag all ingressing traffic with SGT 3. Prefixes of attached subnets are already known.

1. Configure the interface.

```
Device# configure terminal
Device(config)# interface gigabitEthernet 6/3 sgt 3
Device(config)# exit
```

2. Verify that the ingressing traffic to the interface is tagged appropriately.

```
Device# show cts role-based sgt-map all
IP Address      SGT      Source
=====
15.1.1.15       4        INTERNAL
17.1.1.0/24     3        L3IF
```

```

21.1.1.2          4          INTERNAL
31.1.1.0/24      3          L3IF
31.1.1.2          4          INTERNAL
43.1.1.0/24      3          L3IF
49.1.1.0/24      3          L3IF
50.1.1.0/24      3          L3IF
50.1.1.2          4          INTERNAL
51.1.1.1          4          INTERNAL
52.1.1.0/24      3          L3IF
81.1.1.1          5          CLI
102.1.1.1         4          INTERNAL
105.1.1.1         3          L3IF
111.1.1.1         4          INTERNAL
IP-SGT Active Bindings Summary
=====
Total number of CLI      bindings = 1
Total number of L3IF    bindings = 7
Total number of INTERNAL bindings = 7
Total number of active  bindings = 15

```

Example: Emulating the Hardware Keystore

This example shows how to configure and verify the use of a software keystore:

```

Device# configure terminal
Device(config)# cts keystore emulate
Device(config)# exit
Device#show keystore
No hardware keystore present, using software emulation.
Keystore contains the following records (S=Simple Secret, P=PAC, R=RSA):
Index   Type   Name
-----
0       S      CTS-password
1       P      ECF05BB8DFAD854E8376DEA4EF6171CF

```

Example: Configuring Device Route SGT

```

Device# configure terminal
Device(config)# cts role-based sgt-map 0.0.0.0/0 sgt 3
Device(config)# exit

```

Feature History for Security Group Tag Mapping

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Security Group Tag Mapping	Subnet to SGT mapping binds an SGT to all host addresses of a specified subnet. Once this mapping is implemented, Cisco TrustSec imposes the SGT on any incoming packet that has a source IP address which belongs to the specified subnet.
Cisco IOS XE Gibraltar 16.11.1	Default Route SGT Classification	Default Route SGT assigns an SGT tag number to those routes that do not match a specified route.
Cisco IOS XE Cupertino 17.7.1	Security Group Tag Mapping Default Route SGT Classification	These features were implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 7

Cisco TrustSec VRF-Aware SGT

The Cisco TrustSec VRF-Aware SGT feature binds a Security Group Tag (SGT) Exchange Protocol (SXP) connection with a specific virtual routing and forwarding (VRF) instance.

- [VRF-Aware SXP, on page 79](#)
- [How to Configure Cisco TrustSec VRF-Aware SGT, on page 79](#)
- [Configuration Examples for Cisco TrustSec VRF-Aware SGT, on page 81](#)
- [Feature History for Cisco TrustSec VRF-Aware SGT, on page 82](#)

VRF-Aware SXP

The SXP implementation of Virtual Routing and Forwarding (VRF) binds an SXP connection with a specific VRF. It is assumed that the network topology is correctly configured for Layer 2 or Layer 3 VPNs, with all VRFs configured before enabling Cisco TrustSec.

SXP VRF support can be summarized as follows:

- Only one SXP connection can be bound to one VRF.
- Different VRFs may have overlapping SXP peer or source IP addresses.
- IP-SGT mappings learned (added or deleted) in one VRF can be updated only in the same VRF domain. The SXP connection cannot update a mapping bound to a different VRF. If no SXP connection exists for a VRF, IP-SGT mappings for that VRF won't be updated by SXP.
- Multiple address families per VRF is supported. Therefore, one SXP connection in a VRF domain can forward both IPV4 and IPV6 IP-SGT mappings.
- SXP has no limitation on the number of connections and number of IP-SGT mappings per VRF.

How to Configure Cisco TrustSec VRF-Aware SGT

This section describes how to configure Cisco TrustSec VRF-Aware SGT.

Configuring VRF-to-Layer-2-VLAN Assignments

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Device(config)# interface vlan 101	Enables an interface and enters interface configuration mode.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding vrf-intf	Associates a VRF instance or a virtual network with an interface or subinterface. Note Do not configure VRFs on the management interface.
Step 5	exit Example: Device(config-if)# end	Exits interface configuration mode and returns to global configuration mode.
Step 6	cts role-based l2-vrf vrf1 vlan-list 20 Example: Device(config)# cts role-based l2-vrf vrf1 vlan-list 20	Selects a VRF instance for Layer 2 VLANs.
Step 7	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Configuring VRF-to-SGT Mapping

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	cts role-based sgt-map vrf vrf-name {ip4_netaddress ipv6_netaddress host {ip4_address ip6_address}}] sgt sgt_number Example: Device(config)# cts role-based sgt-map vrf red 10.0.0.3 sgt 23	Applies the SGT to packets in the specified VRF. The IP-SGT binding is entered into the IP-SGT table associated with the specified VRF and the IP protocol version implied by the type of IP address.
Step 4	end Example: Device(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Configuration Examples for Cisco TrustSec VRF-Aware SGT

The following sections show configuration examples of Cisco TrustSec VRF-Aware SGT:

Example: Configuring VRF-to-Layer2-VLAN Assignments

```
Device> enable
Device# configure terminal
Device(config)# interface vlan 101
Device(config-if)# vrf forwarding vrf-intf
Device(config-if)# exit
Device(config)# cts role-based l2-vrf vrf1 vlan-list 20
Device(config)# end
```

Example: Configuring VRF-to-SGT Mapping

```
Device> enable
Device# configure terminal
```

```
Device(config)# cts role-based sgt-map vrf red 23.1.1.2 sgt 23
Device(config)# end
```

Feature History for Cisco TrustSec VRF-Aware SGT

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Cisco TrustSec VRF-Aware SGT	The Cisco TrustSec VRF-Aware SGT feature binds a SGT SXP connection with a specific VRF instance.
Cisco IOS XE Cupertino 17.7.1	Cisco TrustSec VRF-Aware SGT	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 8

IP-Prefix and SGT-Based SXP Filtering

The Security Group Tag (SGT) Exchange Protocol (SXP) is one of the several protocols that supports Cisco TrustSec. SXP is a control protocol for propagating IP-to-SGT binding information across network devices that do not have the capability to tag packets. SXP passes IP-to-SGT bindings from authentication points to upstream devices in a network. This process allows security services on switches, routers, or firewalls to learn user identity information from access devices.

The IP-Prefix and SGT-Based SXP Filtering feature allows IP-to-SGT bindings to be filtered, when they are exported or imported. This filtering can be done based on the IP prefix, SGT, or a combination of both.

- [Restrictions for IP-Prefix and Security Group Tag \(SGT\)-Based Security Exchange Protocol \(SXP\) Filtering, on page 83](#)
- [Information About IP-Prefix and SGT-Based SXP Filtering, on page 84](#)
- [How to Configure IP-Prefix and SGT-Based SXP Filtering, on page 84](#)
- [Configuration Examples for IP-Prefix and SGT-Based SXP Filtering, on page 88](#)
- [Verifying IP-Prefix and SGT-Based SXP Filtering, on page 89](#)
- [Syslog Messages for SXP Filtering, on page 91](#)
- [Feature History for IP-Prefix and SGT-Based SXP Filtering, on page 92](#)

Restrictions for IP-Prefix and Security Group Tag (SGT)-Based Security Exchange Protocol (SXP) Filtering

- No high availability support for the stateful synchronization of IP-Security Group Tag (SGT) bindings in a Security Exchange Protocol (SXP) database between active and standby devices.
- Filters applied to an existing connection will take effect only on the subsequent bindings that are exported or imported. The filters do not apply to any bindings that have been exported or imported prior to applying the filters.
- Virtual Routing and Forwarding (VRF)-specific filtering is not supported, and a filter specified for a peer IP is applicable across all VRFs on the device.
- SGT values in filter rules will be a list of single SGT numbers. SGT ranges are not supported.

Information About IP-Prefix and SGT-Based SXP Filtering

Overview

The IP-Prefix and SGT-Based SXP Filtering feature allows IP-to-SGT bindings to be filtered, when they are exported or imported. This filtering can be done based on the IP prefix, SGT, or a combination of both.

The Security Group Tag (SGT) Exchange Protocol (SXP) is one of the several protocols that supports Cisco TrustSec. SXP is a control protocol for propagating IP-to-SGT binding information across network devices that do not have the capability to tag packets. SXP passes IP-to-SGT bindings from authentication points to upstream devices in a network. This process allows security services on switches, routers, or firewalls to learn user identity information from access devices.

The IP-to-SGT filtering allow systems to selectively import or export only bindings of interest. In an SXP connection, a filter can be configured on a device that acts either as a speaker or a listener, based on the filtering that happens during the export or import of bindings.

In the case of bidirectional SXP connections, filters are applied in either of the directions, based on whether a speaker or listener filter is configured. If a peer is a part of both the speaker and the listener filter groups, then filtering is applied in both directions.

Filters can be applied either on a peer-to-peer basis or globally (applicable to all SXP connections). In both cases, the filter can be applied on the speaker or the listener.

Filter Rules

A filter that needs to be applied on a device is created with a set of filter rules. Each filter rule specifies the action or actions to be taken for bindings with specific SGT values and/or IP-prefix values. Each binding is matched against the values specified in the filter rules; if a match is found, the corresponding action specified in the filter rule is applied. An action that can be applied on a selected binding is either a permit or a deny action. When a filter is enabled on the speaker or listener during the export or import of IP-SGT bindings, the bindings are filtered based on the filter rules.

If a rule is not specified for a binding in a filter list, the catch-all rule that is configured in the filter-list is executed. In the absence of a catch-all rule, the corresponding binding is implicitly denied.

Types of SXP Filtering

IP-SGT bindings are filtered in one of the following ways:

- SGT-based filtering: Filters IP-SGT bindings in an SXP connection based on the SGT value.
- IP-prefix based filtering: Filters IP-SGT bindings in an SXP connection based on the IP-prefix value.
- SGT and IP-prefix based filtering: Filter IP-SGT bindings in an SXP connection based on the SGT value and IP-prefix value.

A filter rule is applied on each of the IP-SGT binding.

How to Configure IP-Prefix and SGT-Based SXP Filtering

This section describes how to configure IP-prefix and SGT-based SXP filtering.

Configuring an SXP Filter List

In this step, a filter list is created to hold a set of rules. These rules filter the IP-SGT bindings by allowing bindings that are permitted, and blocking bindings that are denied. Each rule can be based on an SGT, IP prefix, or a combination of both the SGT and IP prefix.

If a filter list does not have a rule that matches a specific IP-SGT binding, the binding is implicitly denied unless a default or catch-all rule is defined.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal	Enters global configuration mode.
Step 3	cts sxp filter-list <i>filter-name</i>	Configures a Cisco TrustSec filter list and enters filter-list configuration mode.
Step 4	<i>sequence-number</i> permit ipv4 <i>ip-address/prefix</i> deny sgt <i>sgt-value</i>	Configures a filter list rule.
Step 5	exit	Exits filter-list configuration mode and returns to global configuration mode.
Step 6	cts sxp filter-list <i>filter-name</i>	Configures a Cisco TrustSec filter list and enters filter list configuration mode.
Step 7	[<i>sequence-number</i>] deny sgt <i>sgt-value</i> permit ipv6 <i>ipv6-address/prefix</i>	Configures a filter list rule.
Step 8	exit	Exits filter-list configuration mode and returns to global configuration mode.
Step 9	cts sxp filter-list <i>filter-name</i>	Configures a Cisco TrustSec filter list and enters filter list configuration mode.
Step 10	[<i>sequence-number</i>] permit ipv6 <i>ipv6-address/prefix</i> permit sgt-value permit	Configures a filter list rule.
Step 11	end	Exits filter-list configuration mode and returns to privileged EXEC mode.

Configuring an SXP Filter Group

In this step, a set of peers are combined into a group, and a filter list is applied to the group. A filter-group can either be defined as a speaker group or listener group. To apply the same filter list to all speakers or all listeners, you can create a global speaker filter group or a global listener filter group.



Note Only one filter list can be attached to a filter group.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal	Enters global configuration mode.
Step 3	cts sxp filter-group listener <i>listener-name</i>	Configures an SXP filter-group listener, and enters filter-group configuration mode.
Step 4	filter <i>filter-list-name</i>	Configures a filter list rule.
Step 5	peer <i>ipv4-address</i>	Configures the IP address of a peer.
Step 6	exit	Exits filter-group configuration mode and returns to global configuration mode.
Step 7	cts sxp filter-group speaker <i>speaker-name</i>	Configures a voice VLAN on a multiple VLAN access port.
Step 8	filter <i>filter-list-name</i>	Configures a filter list name.
Step 9	peer <i>ipv4-address</i>	Configures the IP address of a peer.
Step 10	end	Exits filter-group configuration mode and returns to privileged EXEC mode.

Configuring a Global Listener or Speaker Filter Group

When configuring a global listener and global speaker filter group, the filter is applied to across the box for all SXP connections that are in listener or speaker mode.

When adding a filter-list to a filter group the currently configured set of filter lists on the box is displayed as a help string.



Note The **peer** command is not available for the global listener and global speaker filter-group.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.

	Command or Action	Purpose
Step 2	configure terminal	Enters global configuration mode.
Step 3	cts sxp filter-group listener global <i>filter-list-name</i>	Configures a global listener filter group.
Step 4	cts sxp filter-group speaker global <i>filter-list-name</i>	Configures a global speaker filter group.
Step 5	end	Exits global configuration mode and returns to privileged EXEC mode.

Enabling SXP Filtering

After the SXP filter list and filter groups are configured, you must enable filtering.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal	Enters global configuration mode.
Step 3	cts sxp filter enable	Configures a source template for the interface.
Step 4	exit	Exits global configuration mode and returns to privileged EXEC mode.
Step 5	show cts sxp filter-list <i>filter_name</i>	Displays the filter lists configured on the device along with the filter rules in each of the filter list.

Configuring the Default or Catch-All Rule

The default or catch-all rule is applied on IP-SGT bindings for which there was no match with any of the rules in the filter list. If a default rule is not specified, these IP-SGT bindings are denied.

Define the default or catch-all rule in the filter-list configuration mode of the corresponding filter list.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 3	<code>cts sxp filter-list filter-name</code>	Configures a Cisco TrustSec filter list and enters filter-list configuration mode.
Step 4	<code>permit ipv4 ip-address/prefix</code>	Permits access if the conditions are matched.
Step 5	<code>deny ipv6 ipv6-address/prefix</code>	Denies access if the conditions are matched.
Step 6	<code>permit sgt all</code>	Permits bindings corresponding to all SGTs.
Step 7	<code>end</code>	Exits filter-list configuration mode and returns to privileged EXEC mode.

Configuration Examples for IP-Prefix and SGT-Based SXP Filtering

The following sections show configuration examples of IP-prefix and SGT-based SXP filtering.

Example: Configuring an SXP Filter List

```
Device> enable
Device# configure terminal
Device(config)# cts sxp filter-list filter1
Device(config-filter-list)# permit ipv4 10.1.1.0/24 deny sgt 3 4
Device(config-filter-list)# exit
Device(config)# cts sxp filter-list filter2
Device(config-filter-list)# permit sgt all
Device(config-filter-list)# exit
Device(config)# cts sxp filter-list filter3
Device(config-filter-list)# deny ipv6 2001:db8::1/64 permit sgt 67
Device(config-filter-list)# end
```

Example: Configuring an SXP Filter Group

```
Device> enable
Device# configure terminal
Device(config)# cts sxp filter-group listener group1
Device(config-filter-group)# filter filter1
Device(config-filter-group)# peer 172.16.0.1 192.168.0.1
Device(config-filter-group)# exit
Device(config)# cts sxp filter-group listener global group2
Device(config)# end
```

Example: Enabling SXP Filtering

```
Device> enable
Device# configure terminal
```

```
Device(config)# cts sxp filter-enable
Device(config)# end
```

Example: Configuring the Default or Catch-All Rule

The following example shows how to create a default prefix rule that permits bindings corresponding to all IPv4 and IPv6 addresses:

```
Device(config)# cts sxp filter-list filter1
Device(config-filter-list)# permit ipv4 10.0.0.0/0
Device(config-filter-list)# deny ipv6 2001:db8::1/0
```

The following example shows how to create a default SGT rule that permits bindings corresponding to all SGTs:

```
Device(config)# cts sxp filter-list filter_1
Device(config-filter-list)# permit sgt all
```

Verifying IP-Prefix and SGT-Based SXP Filtering

To verify the configuration, use the following commands:

The **debug cts sxp filter events** command is used to log events related to the creation, removal, and update of filter-lists and filter-groups. This command is also used to capture events related to the matching actions in a filtering process.

```
Device# debug cts sxp filter events
```

The following sample output from the **show cts sxp filter-group speaker** command displays SXP speaker filter groups:

```
Device# show cts sxp filter-group speaker group1
  Filter-group: group1
  Filter-name: filter1
  Peer-list: 172.16.0.1 192.168.0.1
```

The following sample output from the **show cts sxp filter-group listener** command displays SXP speaker listener groups:

```
Device# show cts sxp filter-group listener

Global Listener Filter: Not configured
  Filter-group: group1
  Filter-name: filter1
  Peer-list: 172.16.0.1 192.168.0.1
  Filter-group: group2
  Filter-name: filter1
  Peer-list: 192.0.2.1, 198.51.100.1, 203.0.113.1
```

The following sample output from the **show cts sxp filter-group speaker detailed** command displays detailed information about SXP speaker filter groups:

```
Device# show cts sxp filter-group speaker group1 detailed

Filter-group: group1
Filter-name: filter1
Filter-rules:
  10 deny sgt 30
  20 deny prefix 10.1.0.0/16
  30 permit sgt 60-100
Peer-list: 172.16.0.1 192.168.0.1
```

The following sample output from the **show cts sxp filter-group** command displays information about all configured filter groups:

```
Device# show cts sxp filter-group

Global Listener Filter: Not configured

Global Speaker Filter: Not configured

Listener Group:
  Filter-group: group1
  Filter-name: filter1
  Peer-list: 172.16.0.1 192.168.0.1
  Filter-group: group2
  Filter-name: filter1
  Peer-list: 192.0.2.1, 198.51.100.1, 203.0.113.1

Speaker Group:
  Filter-group: group3
  Filter-name: filter1
  Peer-list: 172.16.0.1 192.168.0.13
  Filter-group: group2
  Filter-name: filter1
  Peer-list: 192.0.2.1, 198.51.100.1, 203.0.113.1
```

The following sample output from the **show sxp filter-group detailed** command displays detailed information about all configured SXP filter groups:

```
Device# show cts sxp filter-group detailed

Global Listener Filter: Configured
  Filter-name: global1
  Filter-rules:
    10 deny 192.168.0.13/32
    20 deny sgt 100-200

Global Speaker Filter: Configured
  Filter-name: global2
  Filter-rules:
    10 deny 192.168.0.13/32
    20 deny sgt 100-200

Listener Group:
  Filter-group: group1
  Filter-name: filter1
  Filter-rules:
    10 deny sgt 30
```

```

    20 deny prefix 172.16.0.0/16
    30 permit sgt 60-100
Peer-list: 172.16.0.1, 192.168.0.13

Filter-group: group2
Filter-name: filter1
Filter-rules:
    10 deny sgt 30
    20 deny prefix 172.16.0.0/16
    30 permit sgt 60-100
Peer-list: 192.0.2.1, 198.51.100.1, 203.0.113.1

Speaker Group
Filter-group: group3
Filter-name: filter1
Filter-rules:
    10 deny sgt 30
    20 deny prefix 172.16.0.0/16
    30 permit sgt 60-100
Peer-list: 10.10.10.1, 172.16.0.1, 192.168.0.13

Filter-group: group2
Filter-name: filter1
Filter-rules:
    10 deny sgt 30
    20 deny prefix 172.16.0.0/16
    30 permit sgt 60-100
Peer-list: 192.0.2.1, 198.51.100.1, 203.0.113.1

```

Syslog Messages for SXP Filtering

Syslog messages for SXP filtering are generated to indicate the various events related to filtering.

Syslog Messages for Filter Rules

The maximum number of rules that can be configured in a single filter is 128. The following message is generated every time the number of filter rules that is configured in a single filter increases by 20% of the limit:

```
CTS SXP filter rules exceed %[ ] threshold. Reached count of [count] out of [max] in filter [filter-name].
```

The following message is generated when the number of rules configured in a single filter reaches 95% of the maximum number of rules allowed for a filter list:

```
CTS SXP filter rules exceed [ ] threshold. Reached count of [count] out of [max] in filter [filter-name].
```

The following message is generated when the number of rules configured in a single filter reaches the maximum number of allowed rules, and no more rules can be added.

```
Reached maximum filter rules. Could not add new rule in filter [filter-name]
```

Syslog Messages for Filter Lists

The maximum number of filter lists that can be configured is 256. The following message is generated every time the number of filter lists that is configured increases by 20% of this limit:

```
CTS SXP filter rules exceed %[ ] threshold. Reached count of [count] out of [max] in filter
[filter-name].
```

The following message is generated when the number of filter lists that is configured reaches 95% of the maximum number of allowed filter lists:

```
CTS SXP filter rules exceed %[ ] threshold. Reached count of [count] out of [max]
```

The following message is generated when the number of filter lists that is configured reaches the maximum number of allowed filter lists, and no more filter lists can be added:

```
Reached maximum filter count. Could not add new filter
```

Feature History for IP-Prefix and SGT-Based SXP Filtering

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	IP-Prefix and SGT-Based SXP Filtering	The IP-Prefix and SGT-Based SXP Filtering feature provides a filtering mechanism to solve the high IP-SGT bindings scale issue.
Cisco IOS XE Cupertino 17.7.1	IP-Prefix and SGT-Based SXP Filtering	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 9

Flexible NetFlow Export of Cisco TrustSec Fields

- [Flexible NetFlow Export of Cisco TrustSec Fields, on page 93](#)

Flexible NetFlow Export of Cisco TrustSec Fields

The Flexible NetFlow Export of Cisco TrustSec Fields feature supports the Cisco TrustSec fields in the Flexible NetFlow (FNF) flow record and helps to monitor, troubleshoot, and identify nonstandard behavior for Cisco TrustSec deployments.

This module describes the interaction between Cisco TrustSec and FNF and how to configure and export Cisco TrustSec fields in the NetFlow Version 9 flow records.

Restrictions for Flexible NetFlow Export of Cisco TrustSec Fields

- The security group tag (SGT) value that is exported in FNF records is zero in the following scenarios:
 - The corresponding packet is received with an SGT value of zero from a trusted interface.
 - The corresponding packet is received without an SGT.
 - The SGT is not found during the IP-SGT lookup. (The SGT is not found in the same packet because the packet is received without an SGT.)
 - When a flow record has SGT and Destination Group Tag (DGT) fields (or only either of the two), and if both these values are not applicable, a flow will still be created with zero values for SGT and DGT. The flow records are expected to include source and destination IP addresses, along with SGT and DGT fields.

Information About Flexible NetFlow Export of Cisco TrustSec Fields

Cisco TrustSec Fields in Flexible NetFlow

The Cisco TrustSec fields, source SGT and destination sSGT, in FNF flow records help administrators correlate the flow with identity information. It enables network engineers to gain a detailed understanding how customers use the network and application resources. This information can then be used to efficiently plan and allocate access and application resources, and to detect and resolve potential security and policy violations.

Cisco TrustSec fields are supported for ingress and egress FNF and for unicast and multicast traffic.

The following table lists NetFlow Version 9 enterprise-specific field types for Cisco TrustSec, which are used in FNF templates for the Cisco TrustSec source and destination SGTs.

Flow Field Type	Description
CTS_SRC_GROUP_TAG	Cisco TrustSec sourceSGT
CTS_DST_GROUP_TAG	Cisco TrustSec destination SGT

Cisco TrustSec fields are configured in addition to the existing match fields under the FNF flow record. The following configurations are used to add Cisco TrustSec flow objects to the FNF flow record as key or nonkey fields and to configure source and destination SGTs for a packet.

The **match flow cts {source | destination} group-tag** command is configured under the corresponding flow record to specify Cisco TrustSec fields as key fields. The key fields differentiate flows, with each flow having a unique set of values. A flow record requires at least one key field, before it can be used in a flow monitor. You can configure the **match** command to a source SGT, destination SGT or both, at the same time.

The flow record is then configured under the flow monitor, and the flow monitor is applied to an interface. To export the FNF data, a flow exporter needs to be configured and then added under the flow monitor.

How to Configure Flexible NetFlow Export of Cisco TrustSec Fields

The following sections provide information about the various tasks that comprise FNF export of Cisco TrustSec fields.

Configuring Cisco TrustSec Fields as Key Fields in the Flow Record

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	flow record record-name Example: Device(config)# flow record cts-record-ipv4	Creates a new FNF flow record, or modifies an existing FNF flow record, and enters Flexible NetFlow flow record configuration mode. • This command also allows you to modify an existing flow record.

	Command or Action	Purpose
Step 4	match ipv4 protocol Example: <pre>Device(config-flow-record)# match ipv4 protocol</pre>	(Optional) Configures the IPv4 protocol as a key field for a flow record.
Step 5	match ipv4 source address Example: <pre>Device(config-flow-record)# match ipv4 source address</pre>	(Optional) Configures the IPv4 source address as a key field for a flow record.
Step 6	match ipv4 destination address Example: <pre>Device(config-flow-record)# match ipv4 destination address</pre>	(Optional) Configures the IPv4 destination address as a key field for a flow record.
Step 7	match transport source-port Example: <pre>Device(config-flow-record)# match transport source-port</pre>	(Optional) Configures the transport source port as a key field for a flow record.
Step 8	match transport destination-port Example: <pre>Device(config-flow-record)# match transport destination-port</pre>	(Optional) Configures the transport destination port as a key field for a flow record.
Step 9	match flow direction Example: <pre>Device(config-flow-record)# match flow direction</pre>	(Optional) Configures the direction in which the flow is monitored as a key field.
Step 10	match flow cts {source destination} group-tag Example: <pre>Device(config-flow-record)# match flow cts source group-tag Device(config-flow-record)# match flow cts destination group-tag</pre>	Configures the Cisco TrustSec source group tag or destination group tag as a key field for the record in the FNF flow record. <ul style="list-style-type: none"> • Ingress: <ul style="list-style-type: none"> • In an incoming packet, if a header is present, SGT reflects the same value as the header. If no value is present, it will show zero. • The DGT value does not depend on the ingress port SGACL configuration.

	Command or Action	Purpose
		<ul style="list-style-type: none"> Egress: <ul style="list-style-type: none"> If either the propagate-sgt command, or Cisco TrustSec is disabled on the egress interface, SGT will be zero. In an outgoing packet, if the SGACL configuration that corresponds to the SGT or DGT exists, DGT will be a numeral other than zero. If SGACL is disabled on the egress port or VLAN, or if global SGACL enforcement is disabled, DGT will be zero.
Step 11	end Example: <pre>Device(config-flow-record)# end</pre>	Exits Flexible NetFlow flow record configuration mode and returns to privileged EXEC mode.

Configuring SGT Name Export in NetFlow

Each flow exporter supports only one destination. If you want to export the data to multiple destinations, you must configure multiple flow exporters and assign them to the flow monitor.

Procedure

	Command or Action	Purpose
Step 1	enable Example: <pre>Device> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.
Step 3	flow exporter <i>exporter-name</i> Example: <pre>Device(config)# flow exporter EXPORTER-1</pre>	Creates a flow exporter or modifies an existing flow exporter, and enters Flexible NetFlow flow exporter configuration mode.
Step 4	destination { <i>ip-address</i> <i>hostname</i> } [vrf <i>vrf-name</i>] Example:	Specifies the IP address or hostname of the destination system for the exporter.

	Command or Action	Purpose
	Device(config-flow-exporter)# destination 172.16.10.2	
Step 5	option cts-sgt-table [timeout seconds] Example: Device(config-flow-exporter)# option cts-sgt-table timeout 1200	Selects the SGT ID-to-name table option for the exporter. <ul style="list-style-type: none"> • This option allows FNF to export Cisco TrustSec environmental data tables that map SGTs to Security Group Names.
Step 6	end Example: Device(config-flow-exporter)# end	Exits Flexible NetFlow flow exporter configuration mode and returns to privileged EXEC mode.

Configuration Examples for Flexible NetFlow Export of Cisco TrustSec Fields

The following sections provide examples relating to the configuration of FNF export of Cisco TrustSec fields.

Example: Configuring Cisco TrustSec Fields as Key Fields in the Flow Record

The following example shows how to configure the Cisco TrustSec flow objects as key fields in an IPv4 Flexible NetFlow flow record:

```
Device> enable
Device# configure terminal
Device(config)# flow record cts-record-ipv4
Device(config-flow-record)# match ipv4 protocol
Device(config-flow-record)# match ipv4 source address
Device(config-flow-record)# match ipv4 destination address
Device(config-flow-record)# match transport source-port
Device(config-flow-record)# match transport destination-port
Device(config-flow-record)# match flow direction
Device(config-flow-record)# match flow cts source group-tag
Device(config-flow-record)# match flow cts destination group-tag
Device(config-flow-record)# end
```

Example: Configuring SGT Name Export in NetFlow

The following example shows how to configure SGT name export in NetFlow

```
Device> enable
Device# configure terminal
Device(config)# flow exporter EXPORTER-1
Device(config-flow-exporter)# destination 172.16.10.2
Device(config-flow-exporter)# option cts-sgt-table timeout 1200
Device(config-flow-exporter)# end
```

Feature History for Flexible NetFlow Export of Cisco TrustSec Fields

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Flexible NetFlow Export of Cisco TrustSec Fields	The Flexible NetFlow Export of Cisco TrustSec Fields feature supports the Cisco TrustSec fields in the FNF flow record and helps to monitor, troubleshoot, and identify nonstandard behavior for Cisco TrustSec deployments.
Cisco IOS XE Cupertino 17.7.1	Flexible NetFlow Export of Cisco TrustSec Fields	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 10

Configuring SGT Inline Tagging

- [Restrictions for SGT Inline Tagging, on page 99](#)
- [Information About SGT Inline Tagging, on page 99](#)
- [SGT Inline Tagging on a NAT Enabled Device, on page 100](#)
- [Configuring SGT Inline Tagging, on page 101](#)
- [Example: Configuring SGT Static Inline Tagging, on page 103](#)
- [Feature History for SGT Inline Tagging, on page 103](#)

Restrictions for SGT Inline Tagging

- Cisco TrustSec manual configurations and 802.1x configurations can coexist only if Security Association Protocol is not configured.

Information About SGT Inline Tagging

Each security group in a Cisco TrustSec domain is assigned a unique 16 bit tag called the Security Group Tag (SGT). The SGT is a single label indicating the privileges of the source within the entire network. It is in turn propagated between network hops allowing any intermediary devices (switches, routers) to enforce policies based on the identity tag.

Cisco TrustSec-capable devices have built-in hardware capabilities that can send and receive packets with SGT embedded in the MAC (L2) layer. This feature is called Layer 2 (L2)-SGT Imposition. It allows ethernet interfaces on the device to be enabled for L2-SGT imposition so that the device can insert an SGT in the packet to be carried to its next hop ethernet neighbor. SGT-over-Ethernet is a method of hop-by-hop propagation of SGT embedded in clear-text (unencrypted) ethernet packets. The inline identity propagation is scalable, provides near line-rate performance and avoids control plane overhead.

The Cisco TrustSec with SGT Exchange Protocol V4 (SXPv4) feature supports Cisco TrustSec metadata-based L2-SGT. When a packet enters a Cisco TrustSec-enabled interface, the IP-SGT mapping database (with dynamic entries built by SXP and/or static entries built by configuration commands) is analyzed to learn the SGT corresponding to the source IP address of the packet, which is then inserted into the packet and carried throughout the network within the Cisco TrustSec header.

As the tag represents the group of the source, the tag is also referred to as the Source Group Tag (SGT). At the egress edge of the network, the group assigned to the packet's destination becomes known. At this point, access control can be applied. With Cisco TrustSec, access control policies are defined between the security

groups and are referred to as Security Group Access Control Lists (SGACL). From the view of any given packet, SGACL is simply being sourced from a security group and destined for another security group.

The SGT tag received in a packet from a trusted interface is propagated to the network, and is also used for Identity firewall classification. When IPsec support is added, the received SGT tag is shared with IPsec for SGT tagging.

A network device at the ingress of Cisco TrustSec cloud needs to determine the SGT of the packet entering the Cisco TrustSec cloud so that it can tag the packet with that SGT when it forwards it into the Cisco TrustSec cloud. The SGT of a packet can be determined with these methods:

- SGT field on Cisco TrustSec header: If a packet is coming from a trusted peer device, it is assumed that the Cisco TrustSec header carries the correct SGT field. This situation applies to a network that is not the first network device in the Cisco TrustSec cloud for the packet.
- SGT lookup based on source IP address: In some cases, the administrator may manually configure a policy to decide the SGT of a packet based upon the source IP address. An IP address to SGT table can also be populated by the SXP protocol.

L2 Inline Tagging is supported for IPv6 multicast traffic with unicast source IPv6 addresses.

SGT Inline Tagging on a NAT Enabled Device

The following scenarios explain how SGT is determined for a packet that flows from a primary device, which has Network Address Translation (NAT) enabled on both ingress and egress ports, to a secondary device:



Note All ports that are used for the flow must have **CTS manual** and trusted configured on both devices.

- If inline tagging is enabled between both devices and SGT tag is not changed with CLI:

In this case, on the primary device Cisco TrustSec is enforced on the SGT tag corresponding to the packet's source IP. The same SGT tag is tagged to the NAT IP. On the secondary device, Cisco TrustSec is enforced on the SGT tag corresponding to the packet's source IP also.

For example, a packet is received on the primary device with a source IP 192.0.2.5 and SGT tag 133. Cisco TrustSec is enforced for the SGT tag 133 on the primary device. After NAT translation the packet's IP changes to 198.51.100.10 and tagged to the SGT tag 133. On the secondary device, the packet is received with IP address 198.51.100.10 and SGT tag 133. Cisco TrustSec is enforced with SGT tag 133 on the secondary device.

- If inline tagging is enabled between both devices and SGT tag is changed with CLI:

In this case, on the primary device Cisco TrustSec is enforced on the SGT tag corresponding to the packet's source IP. The SGT tag is changed by CLI but the SGT tag corresponding to the packet's source IP is tagged to the packet's NAT IP. On the secondary device, Cisco TrustSec is enforced on the SGT tag corresponding to the packet's source IP also.

For example, a packet is received on the primary device with a source IP 192.0.2.5 and SGT tag 133. Cisco TrustSec is enforced for the SGT tag 133 on the primary device. The SGT tag is changed to 200 with CLI. After NAT translation the packet's IP changes to 198.51.100.10 but tagged to the SGT tag 133. On the secondary device, the packet is received with IP address 198.51.100.10 and SGT tag 133. Cisco TrustSec is enforced on the SGT tag 133 on the secondary device.

- If inline tagging is disabled (SGT is populated through SXP protocol on the secondary device) and SGT tag is changed with CLI:

In this case, on the primary device Cisco TrustSec is enforced on the SGT tag corresponding to the packet's source IP. The SGT to Post Nat IP is defined through CLI and is learnt on the primary device. On the secondary device, Cisco TrustSec is enforced on the SGT tag corresponding to the NAT IP, if there is no direct Cisco TrustSec link between primary and secondary device and IP to SGT bindings are learnt through SXP in secondary device.

For example, a packet is received on the primary device with a source IP 192.0.2.5 and SGT tag 133. After NAT translation the source IP changes to 198.51.100.10, for which the SGT is defined through CLI as 200. Cisco TrustSec is enforced for the SGT tag 133 on the primary device. On the secondary device, IP to SGT binding is received through SXP and Cisco TrustSec is enforced on the SGT tag 200 on the secondary device.

Configuring SGT Inline Tagging

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface {gigabitethernet port vlan number} Example: Device(config)# interface gigabitethernet 1/0/1	Configures the interface on which Cisco TrustSec SGT authorization and forwarding is enabled, and enters interface configuration mode.
Step 4	switchport mode access Example: Device(config-if)# switchport mode access	Sets the trunking mode to access mode.

	Command or Action	Purpose
Step 5	cts manual Example: <pre>Device(config-if)# cts manual</pre>	Enables Cisco TrustSec SGT authorization and forwarding on the interface, and enters Cisco TrustSec manual interface configuration mode.
Step 6	propagate sgt Example: <pre>Device(config-if-cts-manual)# propagate sgt</pre>	Enables Cisco TrustSec SGT propagation on an interface. Note Use this command in situations where the peer device is not capable of receiving SGT over Ethernet packets (that is, when a peer device does not support Cisco Ethertype CMD 0x8909 frame format).
Step 7	policy static sgt tag [trusted] Example: <pre>Device(config-if-cts-manual)# policy static sgt 77 trusted</pre>	Configures a static SGT ingress policy on the interface and defines the trustworthiness of an SGT received on the interface. Note The trusted keyword indicates that the interface is trustworthy for Cisco TrustSec. The SGT value received in the Ethernet packet on this interface is trusted and will be used by the device for any SG-aware policy enforcement or for the purpose of egress-tagging.
Step 8	exit Example: <pre>Device(config-if-cts-manual)# exit</pre>	Exits Cisco TrustSec manual interface configuration mode and enters interface configuration mode.
Step 9	dot1x pae authenticator Example: <pre>Device(config-if)# dot1x pae authenticator</pre>	Enables 802.1x authentication on the port.
Step 10	dot1x authenticator eap profile name Example: <pre>Device(config-if)# dot1x authenticator eap profile md5</pre>	Specifies the Extensible Authentication Protocol (EAP) profile to use during 802.1x authentication.

	Command or Action	Purpose
Step 11	end Example: Device(config-if)# end	Exits interface configuration mode and enters privileged EXEC mode.

Example: Configuring SGT Static Inline Tagging

This example shows how to enable an interface on the device for L2-SGT tagging or imposition and defines whether the interface is trusted for Cisco TrustSec

```
Device# configure terminal
Device(config)# interface gigabitethernet 1/0/1
Device(config-if)# cts manual
Device(config-if-cts-manual)# propagate sgt
Device(config-if-cts-manual)# policy static sgt 77 trusted
```

Feature History for SGT Inline Tagging

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	SGT Inline Tagging	Each security group in a Cisco TrustSec domain is assigned a unique 16 bit tag called the SGT. The SGT is a single label indicating the privileges of the source within the entire network.
Cisco IOS XE Cupertino 17.7.1	Cisco TrustSec and 802.1x Support	Support for Cisco TrustSec manual and 802.1x configurations to coexist was introduced.
Cisco IOS XE Cupertino 17.7.1	SGT Inline Tagging Cisco TrustSec and 802.1x Support	These features were implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 11

Configuring Endpoint Admission Control

This module describes the Endpoint Admission Control (EAC) access methods for authentication and authorization in TrustSec networks.

- [Information About Endpoint Admission Control, on page 105](#)
- [Example: 802.1X Authentication Configuration, on page 106](#)
- [Example: MAC Authentication Bypass Configuration, on page 106](#)
- [Example: Web Authentication Proxy Configuration, on page 106](#)
- [Example: Flexible Authentication Sequence and Failover Configuration, on page 107](#)
- [802.1X Host Modes, on page 107](#)
- [Pre-Authentication Open Access, on page 107](#)
- [Example: DHCP Snooping and SGT Assignment, on page 107](#)
- [Feature History for Endpoint Admission Control, on page 108](#)

Information About Endpoint Admission Control

In TrustSec networks, packets are filtered at the egress, not the ingress to the network. In TrustSec endpoint authentication, a host accessing the TrustSec domain (endpoint IP address) is associated with a Security Group Tag (SGT) at the access device through DHCP snooping and IP device tracking. The access device transmits that association (binding) through SXP-to-TrustSec hardware-capable egress devices, which maintain a continually updated table of Source IP to SGT bindings. Packets are filtered on egress by the TrustSec hardware-capable devices by applying security group ACLS (SGACLs).

Endpoint Admission Control (EAC) access methods for authentication and authorization can include the following:

- 802.1X port-based Authentication
- MAC Authentication Bypass (MAB)
- Web Authentication (WebAuth)

All port-based authentication can be enabled with the authentication command. Each access method must be configured individually per port. The flexible authentication sequence and failover features permit the administrator to specify the failover and fallback sequence when multiple authentication modes are configured and the active method fails. The 802.1X host mode determines how many endpoint hosts can be attached per 802.1X port.

Example: 802.1X Authentication Configuration

The following example shows the basic 802.1x configuration on a Gigabit Ethernet port:

```
Device> enable
Device# configure terminal
Device(config)# dot1x system-auth-control
Device(config)# interface GigabitEthernet2/1
Device(config-if)# authentication port-control auto
Device(config-if)# dot1x pae authenticator
```

Example: MAC Authentication Bypass Configuration

MAC Authentication Bypass (MAB) enables hosts or clients that are not 802.1X capable to join 802.1X-enabled networks. It is not required to enable 802.1X authentication prior to enabling MAB.

The following example is of a basic MAB configuration:

```
Device> enable
Device# configure terminal
Device(config)# interface GigabitEthernet2/1
Device(config-if)# authentication port-control auto
Device(config-if)# mab
```

For additional information on configuring MAB authentication, see the configuration guide for your access device.

Example: Web Authentication Proxy Configuration

Web Authentication Proxy (WebAuth) allows the user to use a web browser to transmit their login credentials to the Cisco Secure ACS through a Cisco IOS web server on the access device. WebAuth can be enabled independently. It does not require 802.1X or MAB to be configured.

The following example shows a basic WebAuth configuration on a Gigabit Ethernet port:

```
Device(config)# ip http server
Device(config)# ip access-list extended POLICY
Device(config-ext-nacl)# permit udp any any eq bootps
Device(config-ext-nacl)# permit udp any any eq domain
Device(config)# ip admission name HTTP proxy http
Device(config)# fallback profile FALLBACK_PROFILE
Device(config-fallback-profile)# ip access-group POLICY in
Device(config-fallback-profile)# ip admission HTTP
Device(config)# interface GigabitEthernet2/1
Device(config-if)# authentication port-control auto
Device(config-if)# authentication fallback FALLBACK_PROFILE6500(config-if)#ip access-group POLICY in
```

Example: Flexible Authentication Sequence and Failover Configuration

Flexible Authentication Sequence (FAS) allows the access port to be configured for 802.1X, MAB, and WebAuth authentication methods, specifying the fallback sequence if one or more of the authentication methods are not available. The default failover sequence is as follows:

- 802.1X port-based Authentication
- MAC Authentication Bypass
- Web Authentication

Layer 2 authentications always occur before Layer 3 authentications. That is, 802.1X and MAB must occur before WebAuth.

The following example specifies the authentication sequence as MAB, dot1X, and then WebAuth:

```
Device> enable
Device# configure terminal
Device(config)# interface gigabitEthernet 2/1
Device(config-if)# authentication order mab dot1x webauth
Device(config-if)# ^Z
```

For additional information on FAS, see [Flexible Authentication Order, Priority, and Failed Authentication](#).

802.1X Host Modes

Four host classification modes can be configured per port:

- Single Host—Interface-based session with one MAC address
- Multi Host—Interface-based session with multiple MAC addresses per port
- Multi Domain—MAC + Domain (VLAN) session
- Multi Auth—MAC-based session with multiple MAC address per port

Pre-Authentication Open Access

The Pre-Authentication Open Access feature allows clients and devices to gain network access before port authentication is performed. This process is primarily required for the PXE boot scenario, where a device needs to access the network before PXE times out and download a bootable image that may contain a supplicant.

Example: DHCP Snooping and SGT Assignment

After the authentication process, authorization of the device occurs (for example, dynamic VLAN assignment, ACL programming, etc.). For TrustSec networks, a Security Group Tag (SGT) is assigned per the user

configuration in the Cisco ACS. The SGT is bound to traffic sent from that endpoint through DHCP snooping and the IP device tracking infrastructure.

The following example enables DHCP snooping and IP device tracking on an access device:

```
Device> enable
Device# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Device(config)# ip dhcp snooping
Device(config)# ip dhcp snooping vlan 10
Device(config)# no ip dhcp snooping information option
Device(config)# ip device tracking
```

Feature History for Endpoint Admission Control

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Endpoint Admission Control	In Cisco TrustSec networks, packets are filtered at the egress, not the ingress to the network. In Cisco TrustSec endpoint authentication, a host accessing the Cisco TrustSec domain (endpoint IP address) is associated with a SGT at the access device through DHCP snooping and IP device tracking.
Cisco IOS XE Cupertino 17.7.1	Endpoint Admission Control	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.