Cisco TrustSec Configuration Guide, Cisco IOS XE Gibraltar 16.12.x (Catalyst 3650 Switches)

First Published: 2019-07-31

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

CONTENTS

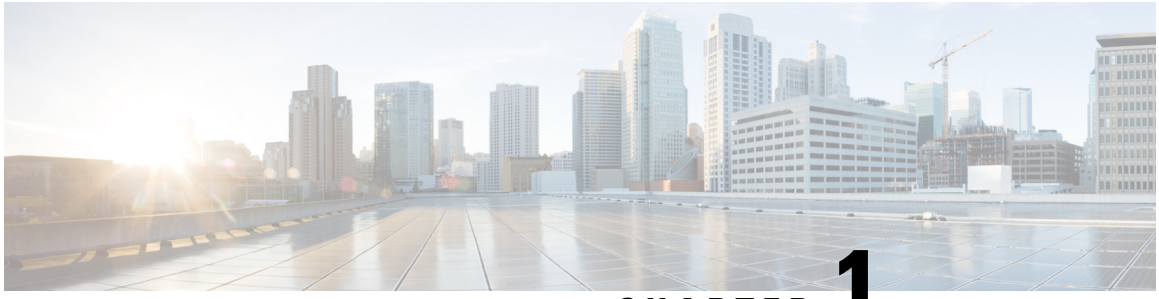

CHAPTER 1

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

- 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.

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

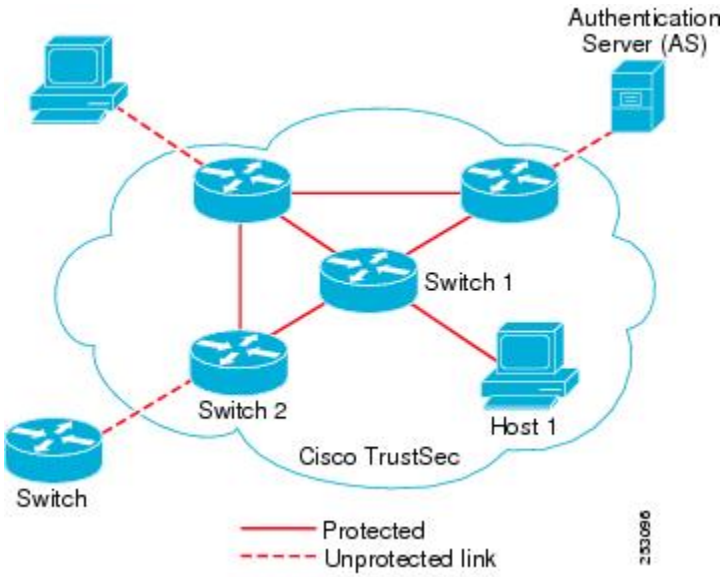

Each participant in the Cisco TrustSec authentication process acts in one of the following roles:

- 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).

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

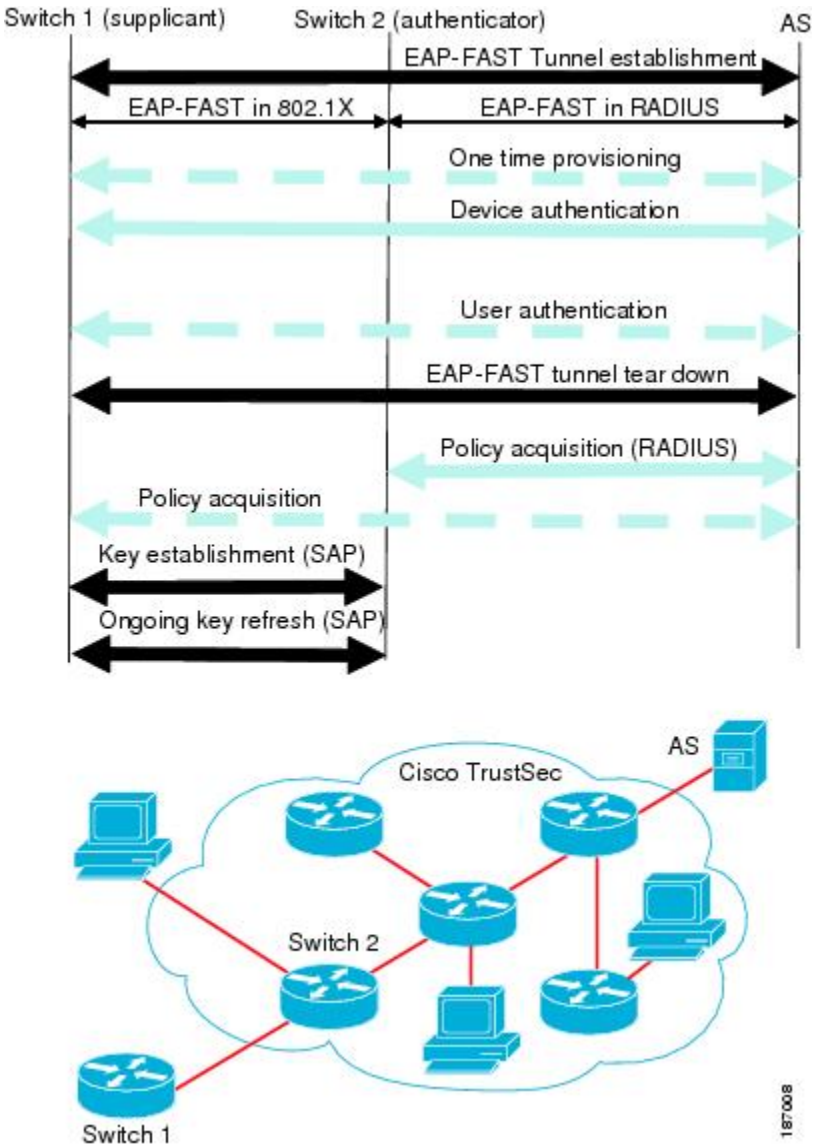

Cisco TrustSec Enhancements to EAP-FAST

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

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 switch 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.

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

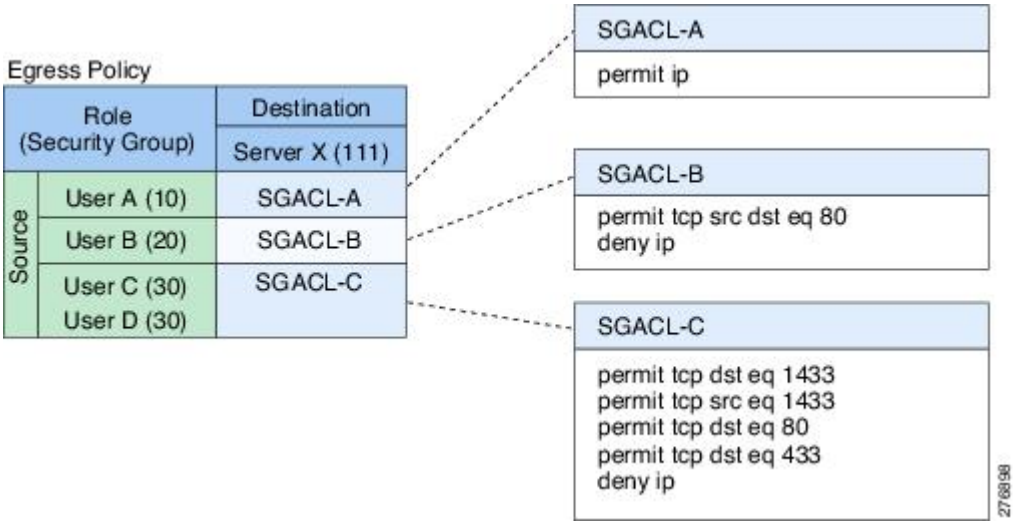

By assigning users and devices within the network to security groups and applying access control between the security groups, Cisco TrustSec achieves role-based topology-independent access control within the network. Because SGACLs define access control policies based on device identities instead of 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 switch, 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 switch 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 total of 1380 SGACL entries are supported on the TCAM.

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

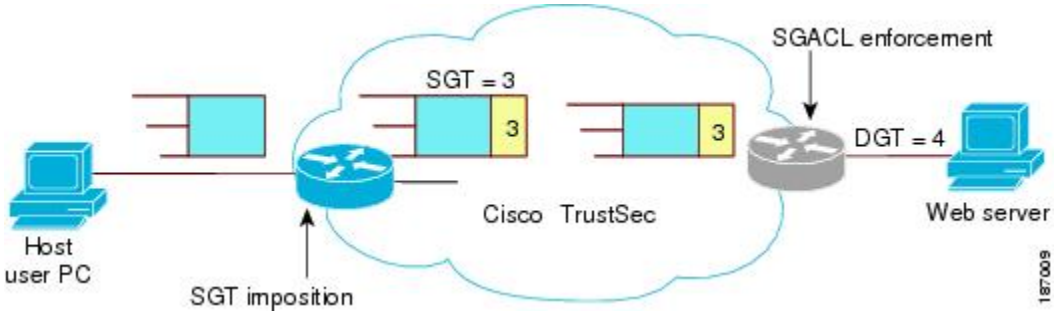

1. The host PC transmits a packet to the web server. Although the PC and the web server are not members of the Cisco TrustSec domain, the data path of the packet includes the Cisco TrustSec domain.
2. The Cisco TrustSec ingress switch modifies the packet to add an SGT with security group number 3, the security group number assigned by the authentication server for the host PC.
3. The Cisco TrustSec egress switch 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 switch 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**.

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:

```
Switch# 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.

```
Switch# 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

Switch# 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)

Switch# 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
```

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

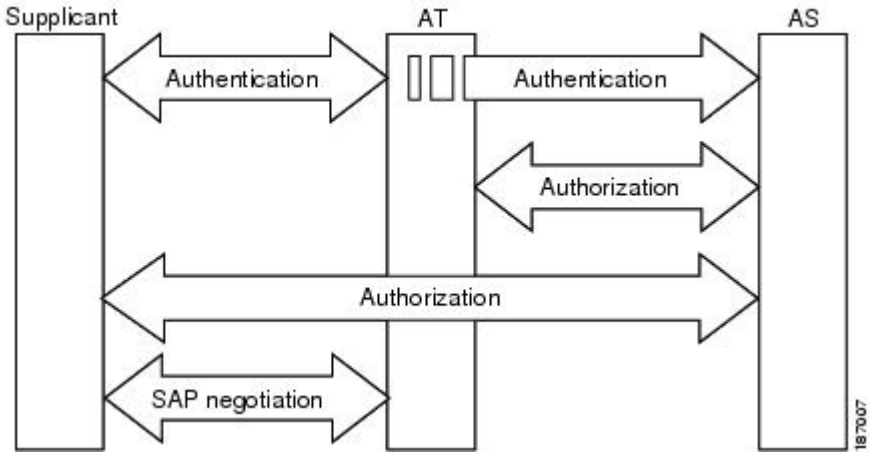

Environment Data Download

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 switch that plays the role of the Cisco TrustSec authenticator in the 802.1X authentication process has IP connectivity to the authentication server, allowing the switch 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.

Using Cisco TrustSec-Incapable Devices and Networks in a Cisco TrustSec Network

SXP for SGT Propagation Across Legacy Access Networks

Tagging packets with SGTs requires hardware support. You might have devices in your network that, while capable of participating in Cisco TrustSec authentication, lack the hardware capability to tag packets with SGTs. By using the SGT Exchange Protocol (SXP), these devices can pass IP-address-to-SGT mappings to a Cisco TrustSec peer device that has Cisco TrustSec-capable hardware.

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 switches. Distribution switches 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

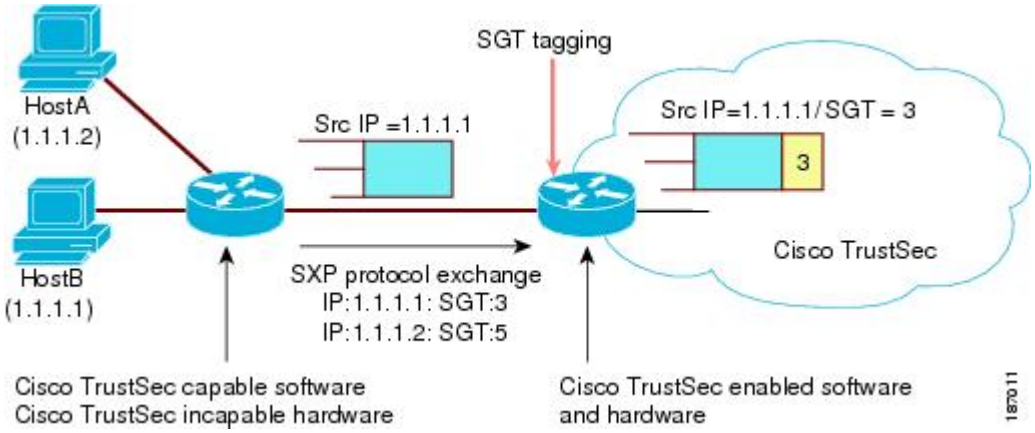

You must manually configure an SXP connection between a peer without Cisco TrustSec hardware support and a peer with Cisco TrustSec hardware support. The following tasks are required when configuring the SXP connection:

- 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

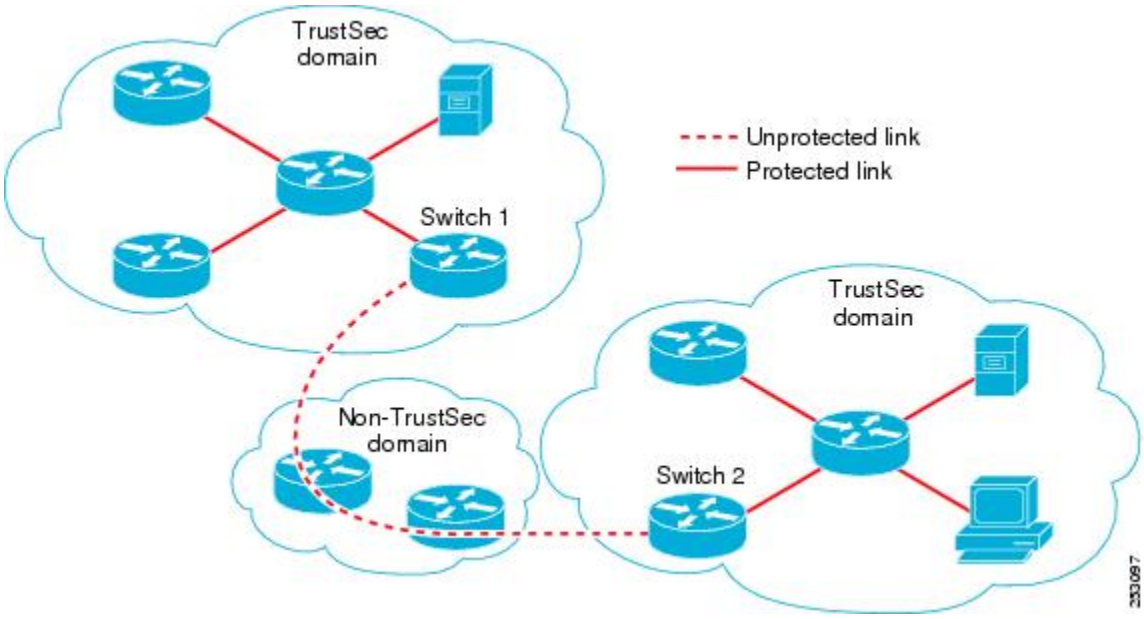

To support Cisco TrustSec Layer 3 SGT Transport, any device that will act as a Cisco TrustSec ingress or egress Layer 3 gateway must maintain a traffic policy database that lists eligible subnets in remote Cisco TrustSec domains as well as any excluded subnets within those regions. You can configure 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.

Cisco TrustSec Reflector for Cisco TrustSec-Incapable Switching Modules

A Cisco device in a Cisco TrustSec domain may contain any of these types of switching modules:

- Cisco TrustSec-capable—Hardware supports insertion and propagation of SGT.
- Cisco TrustSec-aware—Hardware does not support insertion and propagation of SGT, but hardware can perform a lookup to determine the source and destination SGTs for a packet.
- Cisco TrustSec-incapable—Hardware does not support insertion and propagation of SGT and cannot determine the SGT by a hardware lookup.

If your switch contains a Cisco TrustSec-capable supervisor engine, you can use the Cisco TrustSec reflector feature to accommodate legacy Cisco TrustSec-incapable switching modules within the same switch. Cisco TrustSec reflector uses SPAN to reflect traffic from a Cisco TrustSec-incapable switching module to the supervisor engine for SGT assignment and insertion.

Two mutually exclusive modes, ingress and egress, are supported for the Cisco TrustSec reflector. The default is pure mode, in which neither reflector is enabled. A Cisco TrustSec ingress reflector is configured on an access switch facing a distribution switch, while a Cisco TrustSec egress reflector is configured on a distribution switch.

Ingress Reflector

A Cisco TrustSec ingress reflector is implemented on an access switch, where the Cisco TrustSec-incapable switching module is on the Cisco TrustSec domain edge and the Cisco TrustSec-capable supervisor engine uplink port connects to a Cisco TrustSec-capable distribution switch.

The following conditions must be met before the Cisco TrustSec ingress reflector configuration is accepted:

- The supervisor engine must be Cisco TrustSec-capable.
- Any Cisco TrustSec-incapable DFCs must be powered down.
- A Cisco TrustSec egress reflector must not be configured on the switch.
- Before disabling the Cisco TrustSec ingress reflector, you must remove power from the Cisco TrustSec-incapable switching modules.

Egress Reflector

A Cisco TrustSec egress reflector is implemented on a distribution switch with Layer 3 uplinks, where the Cisco TrustSec-incapable switching module faces an access switch. The Cisco TrustSec egress reflector is

supported only on Layer 3 uplinks, and is not supported on Layer 2 interfaces, SVIs, subinterfaces, or tunnels, and is not supported for NAT traffic.

The following conditions must be met before the Cisco TrustSec egress reflector configuration is accepted:

- The supervisor engine or DFC switching module must be Cisco TrustSec-capable.
- Cisco TrustSec must not be enabled on non-routed interfaces on the supervisor engine uplink ports or on the Cisco TrustSec-capable DFC switching modules.
- Before disabling the Cisco TrustSec egress reflector, you must remove power from the Cisco TrustSec-incapable switching modules.
- A Cisco TrustSec ingress reflector must not be configured on the switch.

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 exits 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 Information for Cisco TrustSec Overview

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Cisco TrustSec Overview

Feature Name	Release	Feature Information
Cisco TrustSec Overview	Cisco IOS XE Denali 16.1.1	Cisco TrustSec builds secure networks by establishing domains of trusted network devices.

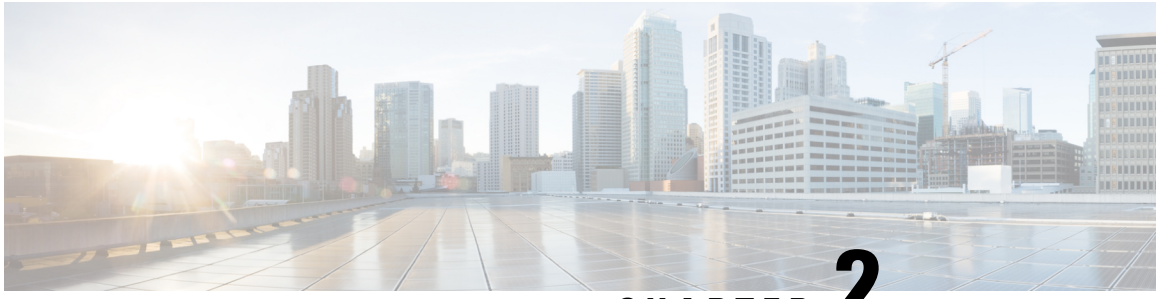

CHAPTER 2

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 Security Group Access Control Lists (SGACL) Policies

The following restriction apply to the Cisco Catalyst 3650 Series Switches and Cisco Catalyst 3850 Series Switches:

- CTS SGACLs cannot be enforced for punt (CPU bound) traffic due to hardware limitations.

How to Configure SGACL Policies

The following sections provide information on various SGACL policy configurations.

SGACL Policy Configuration Process

Follow these steps to configure and enable Cisco TrustSec Security Group ACL (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 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 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.

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.

	Command or Action	Purpose
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 to privileged EXEC mode.
Step 6	**show cts interface** **Example:** `Device# show cts interface`	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.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts role-based enforcement vlan-list** *vlan-list* **Example:** `Device(config)# cts role-based enforcement vlan-list 31-35,41`	Enables Cisco TrustSec SGACL policy enforcement on the VLAN or VLAN list.

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# 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 monitor enable** **Example:** `Device(config)# cts role-based monitor enable`	Enables device level monitor mode. • By default device level monitor mode is enabled. If device monitor mode is disabled, monitor mode information is still downloaded from ISE but not applied on device until this configuration is turned on.
Step 4	**cts role-based monitor permissions from** {*sgt_num*} **to** {*dgt_num*} [**ipv4** | **ipv6**] **Example:** `Device(config)# cts role-based permissions from 2 to 3 ipv4`	Enables monitor mode for IPv4/IPv6 Role Based Access Control List (RBACL) (Security Group Tag (SGT)- Destination Group Tag (DGT) pair).
Step 5	**end** **Example:** `Device(config)# end`	Exits to privileged EXEC mode.
Step 6	**show cts role-based monitor permissions from** {*sgt_num*} **to** {*dgt_num*} [**ipv4** | **ipv6**] [**details**] **Example:** `Device# show cts role-based permissions from 2 to 3 ipv4 details`	Displays the SGACL policies and details about the monitor mode feature for each pair. The command output displays monitored 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`	Displays all 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 the Cisco ISE or the Cisco Secure ACS. To manually (that is, locally) configure SGACL policies, do the following:

1. Configure a role-based ACL.
2. Bind the role-based ACL to a range of SGTs.

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

Manually Configuring and Applying IPv4 SGACL Policies

Note When configuring SGACLs and Role-Based access control lists (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** *rbacl-name* **Example:** `Device(config)# ip access-list role-based allow_webtraff`	Creates a Role-based ACL and enters Role-based ACL configuration mode.
Step 4	{ [*sequence-number*] | **default** | **permit** | **deny** | **remark**} **Example:** `Device(config-rb-acl)# 10 permit tcp dst eq 80 dst eq 20`	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. Press Enter to complete an ACE and begin the next.

	Command or Action	Purpose
		The following ACE commands or keywords are not supported: • **reflect** • **evaluate** • **time-range**
Step 5	**exit** **Example:** `Device(config-rb-acl)# exit`	Exits role-based ACL configuration mode and returns to global configuration mode.
Step 6	**cts role-based permissions {default** | [**from** {*sgt_num* | **unknown**} **to** {*dgt_num* | **unknown** }] {*rbacls* | **ipv4** *rbacls*} **Example:** `Device(config)# cts role-based permissions from 55 to 66 allow_webtraff`	Binds SGTs and DGTs to the RBACL. The configuration is analogous to populating the permission matrix configured on the Cisco ISE or the Cisco Secure ACS. • Default—Default permissions list • *sgt_num*—0 to 65,519. Source Group Tag. • *dgt_num*—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 following RBACL is IPv4. • *rbacls* —Name of RBACLs
Step 7	**end** **Example:** `Device(config)# end`	Exits global configuration mode and returns to privileged EXEC mode.
Step 8	**show cts role-based permissions** **Example:** `Device# show cts role-based permissions`	Displays permission to RBACL configurations.
Step 9	**show ip access-lists** {*rbacls* | **ipv4** *rbacls*} **Example:** `Device# show ip access-lists allow_webtraff`	Displays ACEs of all RBACLs or a specified RBACL.

Configuring IPv6 Policies

To manually configure IPv6 SGACL policies, perform this task:

Note IPv6 SGACL is not supported on Cisco IOS XE Everest 16.8.1.

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** *sgacl-name* **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** {*doh-number* | *doh-type*}] [**dscp** *cp-value*] [**flow-label** *fl-value*] [**mobility** | **mobility-type** {*mh-number* | *mh-type*}] [**routing** | **routing-type** *routing-number*] [**fragments**] [**log** | **log-input**] [**sequence** *seqno*]	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 commands or keywords are not supported: • **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 SGACLs. 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 a source security group (SGT) and destination security group (DGT). Values for source-sgt and dest-sgt range from 1 to 65533. By default, SGACLs are considered to be IPv4. • **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. **Note** An SGACL policy downloaded dynamically from the ACS will override any conflicting manual policy.

Displaying SGACL Policies

After configuring the Cisco TrustSec device credentials and AAA, you can verify the Cisco TrustSec SGACL policies downloaded from the authentication server or configured manually. Cisco TrustSec downloads the SGACL policies when it learns of a new SGT through authentication and authorization on an interface, from SXP, or from manual IP address to SGT mapping.

Using the 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(config)# 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(config)# show cts role-based permissions from 3`	Specifies the SGACLs to be applied for a source security group (SGT) and destination security group (DGT). Values for source-sgt and dest-sgt range from 1 to 65533. By default, SGACLs are considered to be IPv4. • **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. **Note** An SGACL policy downloaded dynamically from the ACS will override any conflicting manual policy.

Refreshing the Downloaded SGACL Policies

Procedure

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

	Command or Action	Purpose
	`Device# enable`	
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts refresh policy** {**peer** [*peer-id*] | **sgt** [*sgt_number* | **default** | **unknown**]} **Example:** `Device(config)# cts refresh policy peer my_cisco_ise`	Performs an immediate refresh of the SGACL policies from the authentication server. • If a *peer-id* is specified, only the policies related to the specified peer connection are refreshed. To refresh all 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 security group tag policies, press Enter without specifying an SGT number. Select default to refresh the default policy. Select unknown to refresh unknown policy.

Configuration Examples for SGACL Policies

The following sections provide information on various SGACL policy configuration examples.

Example: Enabling SGACL Policy Enforcement Globally

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

Example: Enabling SGACL Policy Enforcement Per Interface

```
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

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

Example: Configuring SGACL Monitor Mode

```
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

```
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 show cts role-based permissions from 50 to 70
```

Example: Manually Applying SGACLs

```
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# 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 Information for SGACL Policies

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 2: Feature Information for SGACL Policies

Feature Name	Release	Feature Information
SGACL Policies	Cisco IOS XE Denali 16.1.1	This feature was introduced.

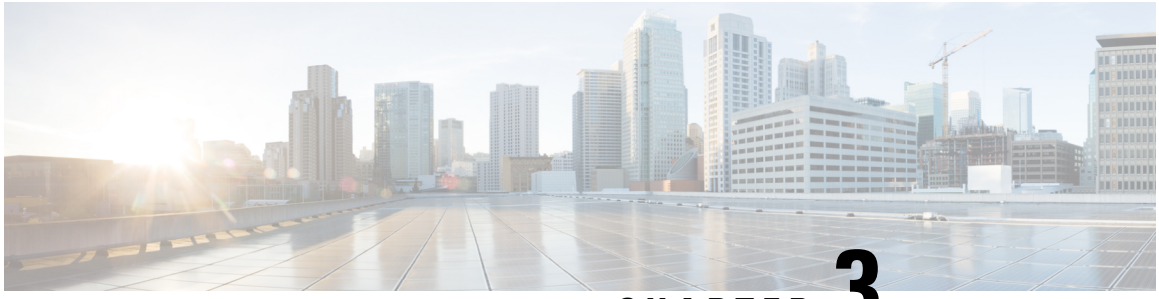

CHAPTER 3

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

This document assumes the following:

- An understanding of Cisco TrustSec and the Security Group access control lists (SGACL) configuration.
- Switches are configured to function as a stack.
- All the switches 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:
        multple_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:
        multple_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
F23623C470511E0000001356DEA96C00093A80538898D40F633C368B053200D4C9D2422A7FEB4837EA9DBB89D1
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:
        multple_ace-14
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
```

Additional References for Configuring Cisco TrustSec SGACL High Availability

Related Documents

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/support

Feature Information for SGACL High Availability

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 3: Feature Information for Cisco TrustSec SGACL High Availability

Feature Name	Release	Feature Information
Cisco TrustSec SGACL High Availability	Cisco IOS XE Denali 16.2.1	Cisco TrustSec Security Group access control lists (SGACLs) support the high availability functionality available on the switch stack manager. There is no Cisco TrustSec-specific configuration to enable this functionality. This functionality is only available on switches that have the stack manager architecture and use Cisco IOS XE Denali 16.2.1 and later releases.

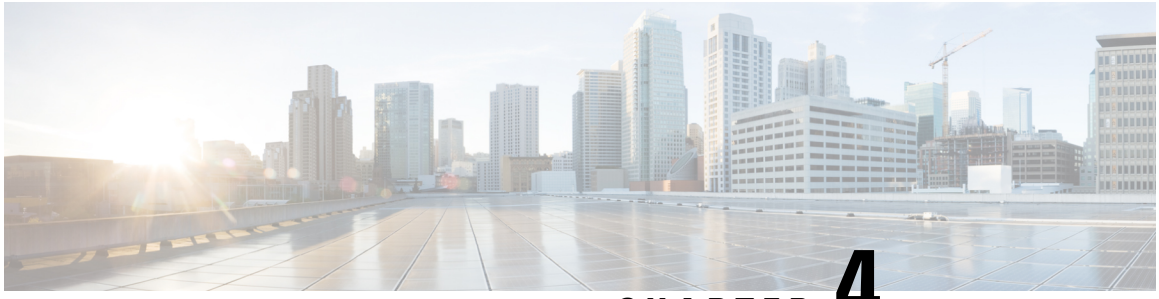

CHAPTER 4

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

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 supported only on physical interfaces and not on logical 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

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 switch 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

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.
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.

	Command or Action	Purpose
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 switch.

To configure an SXP peer connection, 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 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: • **default**—Use the default SXP password you configured using the **cts sxp default password** command. • **none**—Do not use a password. The mode keyword specifies the role of the remote peer device:

	Command or Action	Purpose
		• **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. The optional **vrf** keyword specifies the VRF to the peer. The default is the default VRF.
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. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts sxp default password** [**0** | **6** | **7**] *password* **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.

	Command or Action	Purpose
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. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts sxp default source-ip** *src-ip-addr* **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. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts sxp reconciliation period** *seconds* **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. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**cts sxp retry period** *seconds* **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.
Step 4	**exit** **Example:**	Exits global configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
	`Device(config)# exit`	

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. • 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

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 Switch A, the speaker, and Switch 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 Switch B, the listener, and Switch 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:

```
Switch# 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:

```
Switch# 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 Information for SGT Exchange Protocol

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 4: Feature Information for SGT Exchange Protocol

Feature Name	Release	Feature Information
SGT Exchange Protocol	Cisco IOS XE Denali 16.1.1	The SGT Exchange Protocol (SXP) propagates the Security Group Tags (SGTs) across network devices that do not have hardware support for Cisco TrustSec.

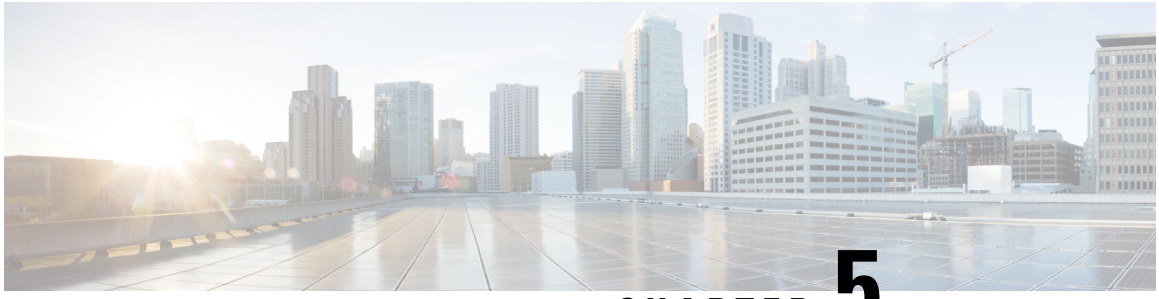

CHAPTER 5

Cisco TrustSec VRF-Aware SGT

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

Information About Cisco TrustSec VRF-Aware SGT

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 exits 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 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. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Device(config)# interface vlan 101`	Enables an interface and enters interface configuration mode.
Step 4	**vrf forwarding** *vrf-name* **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

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-Layer2-VLAN Assignments

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

Additional References for Configuring Cisco TrustSec VRF-Aware SGT

Related Documents

MIBs

MIB	MIBs Link
All the supported MIBs for this release.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use the Cisco MIB Locator found at: http://www.cisco.com/go/mibs.

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	https://www.cisco.com/c/en/us/support/index.html

Feature Information for Cisco TrustSec VRF-Aware SGT

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 5: Feature Information for Cisco TrustSec VRF-Aware SGT

Feature Name	Release	Feature Information
Cisco TrustSec VRF-Aware SGT	Cisco IOS XE Denali 16.1.1	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.

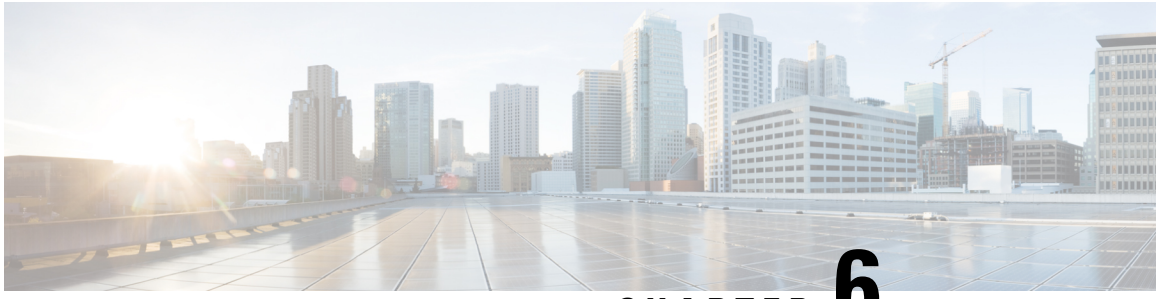

CHAPTER 6

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

- No high availability support for the stateful synchronization of IP-Security Group Tag (SGT) bindings in an 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

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 ruled is defined.

Procedure

	Command or Action	Purpose
Step 1	**enable**	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**cts sxp filter-list** *filter-name*	Configures a Cisco TrustSec filter list and enters filter-list configuration mode.
Step 4	*sequence-number* **permit ipv4** *ip-address/prefix* **deny sgt** *sgt-value*	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** *filter-name*	Configures a Cisco TrustSec filter list and enters filter list configuration mode.
Step 7	[*sequence-number*] **deny sgt** *sgt-value* **permit ipv6** *ipv6-address/prefix*	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** *filter-name*	Configures a Cisco TrustSec filter list and enters filter list configuration mode.
Step 10	[*sequence-number*] **permit ipv6** *ipv6-address/prefix* **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. • Enter your password if prompted.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**cts sxp filter-group listener** *listener-name*	Configures an SXP filter-group listener, and enters filter-group configuration mode.
Step 4	**filter** *filter-list-name*	Configures a filter list rule.
Step 5	**peer** *ipv4-address*	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** *speaker-name*	Configures a voice VLAN on a multiple VLAN access port.
Step 8	**filter** *filter-list-name*	Configures a filter list name.
Step 9	**peer** *ipv4-address*	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. • Enter your password if prompted.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**cts sxp filter-group listener global** *filter-list-name*	Configures a global listener filter group.
Step 4	**cts sxp filter-group speaker global** *filter-list-name*	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. • 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** *filter_name*	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. • Enter your password if prompted.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**cts sxp filter-list** *filter-name*	Configures a Cisco TrustSec filter list and enters filter-list configuration mode.
Step 4	**permit ipv4** *ip-address/prefix*	Permits access if the conditions are matched.
Step 5	**deny ipv6** *ipv6-address/prefix*	Denies access if the conditions are matched.
Step 6	**permit sgt all**	Permits bindings corresponding to all SGTs.
Step 7	**end**	Exits filter-list configuration mode and returns to privileged EXEC mode.

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

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 6: Feature Information for IP-Prefix and SGT-Based SXP Filtering

Feature Name	Release	Feature Information
IP-Prefix and SGT-Based SXP Filtering	Cisco IOS XE Denali 16.1.1	The IP-Prefix and SGT-Based SXP Filtering feature provides a filtering mechanism to solve the high IP-SGT bindings scale issue. These commands were introduced: **debug cts sxp filter events**, **cts sxp filter-list**, **cts sxp filter-group**, **cts sxp filter-enable**, **show cts sxp filter-group**, **show cts sxp filter-list**.

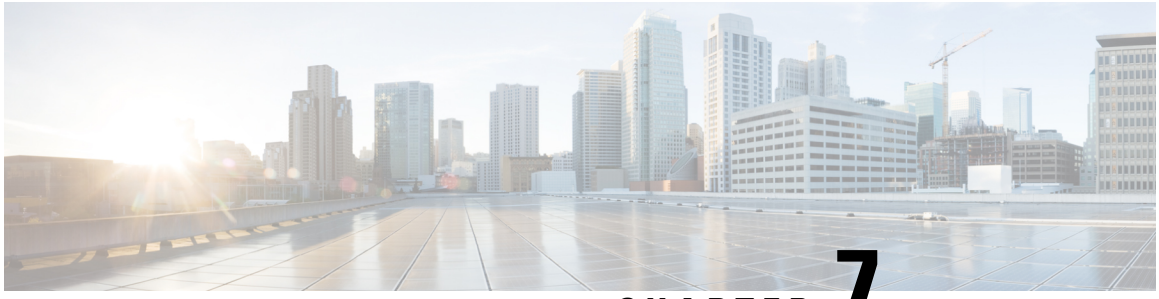

CHAPTER 7

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

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 a basic MAB configuration on a Catalyst switch:

```
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 switch.

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 though 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 switch:

```
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 Information for Endpoint Admission Control

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 7: Feature Information for Endpoint Admission Control

Feature Name	Release	Feature Information
Endpoint Admission Control	Cisco IOS XE Denali 16.1.1	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.