



Cisco TrustSec Overview

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

- [Restrictions for Cisco TrustSec, on page 1](#)
- [Information About Cisco TrustSec Architecture , on page 1](#)
- [Authentication, on page 4](#)
- [Security Group-Based Access Control, on page 7](#)
- [Authorization and Policy Acquisition, on page 13](#)
- [Environment Data Download, on page 14](#)
- [RADIUS Relay Functionality, on page 14](#)
- [Link Security, on page 15](#)
- [SXP for SGT Propagation Across Legacy Access Networks, on page 16](#)
- [Layer 3 SGT Transport for Spanning Non-TrustSec Regions, on page 18](#)
- [Cisco TrustSec Reflector for Cisco TrustSec-Incapable Switching Modules, on page 19](#)
- [VRF-Aware SXP, on page 20](#)
- [Feature History for Cisco TrustSec Overview, on page 20](#)

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).



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

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

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



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


Authentication

Cisco TrustSec and Authentication

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

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

Figure 2: Cisco TrustSec Authentication



Cisco TrustSec Enhancements to EAP-FAST

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

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

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

802.1X Role Selection

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

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

Cisco TrustSec Authentication Summary

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

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

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

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

Device Identities

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

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

Device Credentials

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

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

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

User Credentials

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

Security Group-Based Access Control

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

Security Groups and SGTs

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

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

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

Security Group ACL Support

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

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

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


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

SGACL Policies

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

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

Figure 3: SGACL Policy Matrix Example



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



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


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

Ingress Tagging and Egress Enforcement

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

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

Figure 4: SGT and SGACL in a Cisco TrustSec Domain



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

Determining the Source Security Group

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

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

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

Determining the Destination Security Group

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

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

SGACL Enforcement on Routed and Switched Traffic

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

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

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

SGACL Logging and ACE Statistics

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

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

Device# show ip access-lists sgacl_5_18-01
Role-based IP access list sgacl_5_18-01 (downloaded)
10 deny icmp echo log (1 match)

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

VRF-aware SGACL Logging

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

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

SGACL Monitor Mode

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

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

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

Authorization and Policy Acquisition

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

The authentication server returns the following policy attributes:


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



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

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

Figure 5: NDAC and SAP Negotiation



Environment Data Download

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

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

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

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

RADIUS Relay Functionality

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

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

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

Link Security

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

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

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

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


Configuring SAP-PMK for Link Security

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Device(config)# interface TenGigabitEthernet 1/1/4	Configures an interface and enters interface configuration mode.
Step 4	switchport mode trunk Example: Device(config-if)# switchport mode trunk	Specifies a trunking VLAN Layer 2 interface.
Step 5	cts manual Example: Device(config-if)# cts manual	Enters Cisco TrustSec manual configuration mode.

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

Figure 6: SXP Protocol to Propagate SGT Information



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

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


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

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

Layer 3 SGT Transport for Spanning Non-TrustSec Regions

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

Figure 7: Spanning a Non-TrustSec domain



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

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



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

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 exists for a VRF, IP-SGT mappings for that VRF won't be updated by SXP.
- Multiple address families per VRF is supported. Therefore, one SXP connection in a VRF domain can forward both IPV4 and IPV6 IP-SGT mappings.
- SXP has no limitation on the number of connections and number of IP-SGT mappings per VRF.

Layer 2 VRF-Aware SXP and VRF Assignment

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

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

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

Feature History for Cisco TrustSec Overview

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

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

Release	Feature	Feature Information
Cisco IOS XE Everest 16.6.1	Cisco TrustSec Overview	Cisco TrustSec builds secure networks by establishing domains of trusted network devices.

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

