Overview of TrustSec

January 2014

Contents

What You Will Learn

If you have read the Cisco Secure Access® How-To Guides, you have been exposed to many ways of controlling network access based on the context of user and device. VLAN assignment controls network access at the Layer 3 edge or by isolating that VLAN into a segmented virtual network. Additionally, access control list (ACL) assignment, which can be a local ACL, can be called into action by a RADIUS attribute or a downloaded ACL (dACL). These ACLs are applied at the switch port ingress or at the virtual port in the case of the wireless LAN controller (WLC).

These are all very good access-control methods, but controlling access only at the point of network ingress can leave room for a more desirable and scalable solution. In this guide, we will discuss a Cisco innovation that makes access control more scalable and powerful Cisco TrustSec.

Cisco TrustSec is defined in three phases: classification, propagation, and enforcement. This guide will focus on these fundamentals as well as the configuration of the many devices available for use in a Cisco TrustSec environment. Basic use cases will be presented where scalable security policy can be implemented with switches, Security Group ACLs (SGACLs), and Security Group Firewalls (SGFWs).

Ingress Access Control Challenges

VLAN assignments and dACLs are fantastic ways of controlling access to a network. However, when a network grows, so do the challenges of keeping up with the ingress access controls. Let's take a look at each one of these standard use cases individually and discuss the challenges.

VLAN Assignment

VLAN assignment based on the context of a user or device is a very common way to control access to a network. Let's use the hypothetical scenario of controlling access to servers that contain credit card data. This access falls under Payment Card Industry (PCI) compliance standards.

1. A user is a member of the Retail Managers group in Windows Active Directory.
2. The posture of the system is compliant.
3. Therefore the Cisco® Identity Services Engine (ISE) assigns the user to the PCI-allowed VLAN on the switch or WLC.

Now, an ACL must be applied somewhere to enable VLAN assignment to control access to the servers that house the PCI data (Figure 1). Let's assume that the ACL is applied at a firewall between the campus or branch network and the data center.

1. The ACL on the data center firewall must be updated to include all the source IP addresses of PCI-allowed VLANs throughout the entire network infrastructure.

Figure 1. Controlling Access with VLANs on a Single Switch
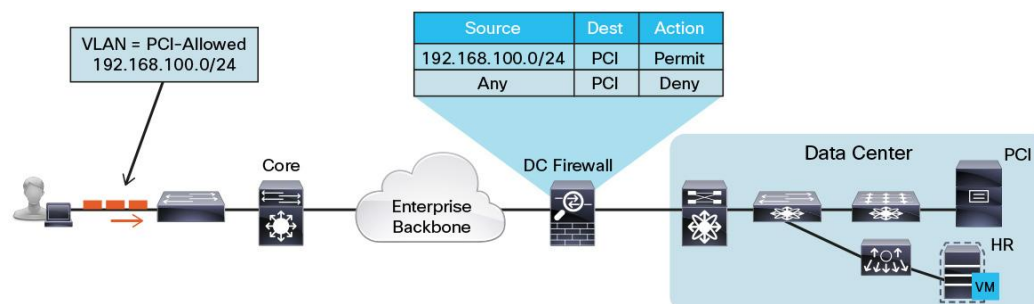
Next, the company has decided to control access to the human relations (HR) department server, so that only members of that department may talk to HR servers (Figure 2). Another set of rules will need to be built that assign the HR VLAN, and another set of entries must be made in the ACL.

Figure 2.   Controlling Access with Two VLANs on a Single Switch
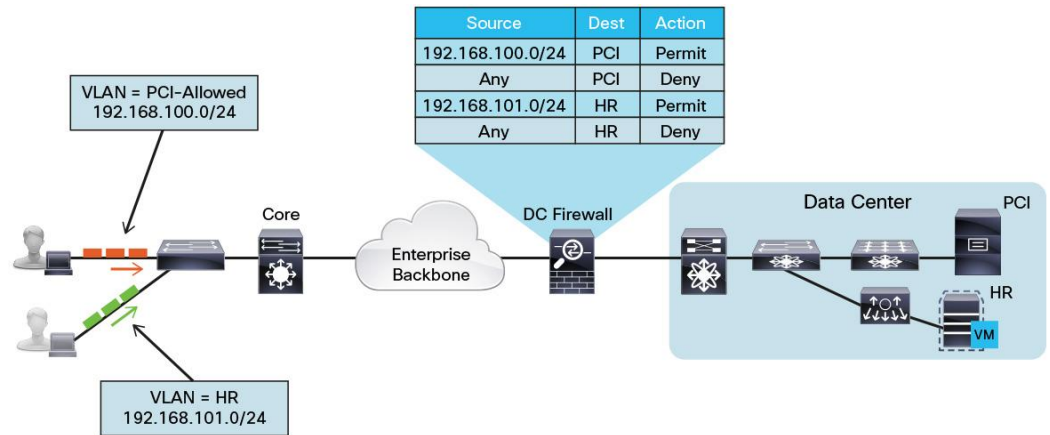
Now, consider how this can scale as we continue to add VLANs and we continue to add switches and WLCs to the equation. One of our large customers has more than 50,000 switches in the access layer. That is a tremendous number of VLANs to create and addresses to maintain in an access list on a firewall. That same customer had 15 full-time employees managing the firewall rules. The company needed to find a better mechanism to control access that would lower its operational expenses tremendously.

What if you had 100 remote sites? A hundred new IP subnets could easily modify your existing route strategy. When that is the case, the route summarization alone can cause a network redesign, which will add even more operational cost (Figure 3).

Figure 3.   Operationally Taxing VLAN Control
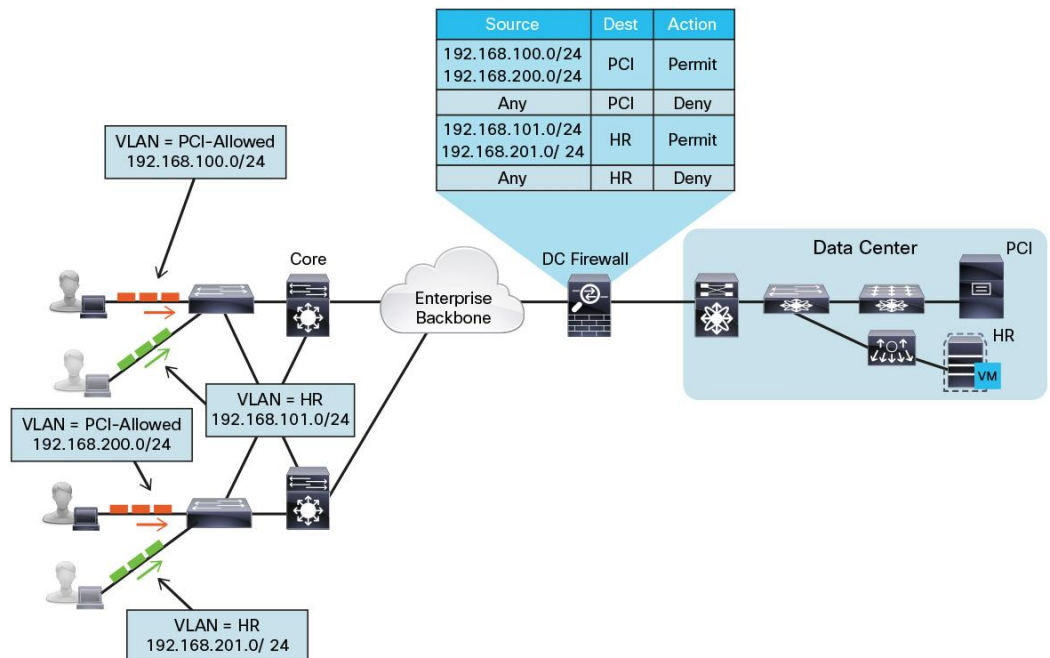
The number of access control entries (ACEs) in an ACL can be determined by a formula. The formula takes the number of sources multiplied by the number of destinations multiplied by the permissions of the ACL.

```
(sources) * (destinations) * permissions = ACEs
```

So with the environment depicted in Figure 3, we would need 32 ACEs for only four sources, two destinations, and four permissions. Now with 100 remote sites… it is easy to visualize the explosion of ACEs.

Ingress ACLs

Another way to control access is to use access lists applied at ingress (inbound) at the port (or virtual port) that the user or device is using to access the network (Figure 4). These could be locally defined ACLs that are called by using the Filter-ID RADIUS attribute, or they could be dACLs, in which the entire ACL is defined on the Cisco ISE and downloaded to the port.

Obviously, dACLs provide a better operational model, because you have to update an ACL only once. Additionally, the number of ACEs required is lower when the ACL is applied to a switch port than it would be if the ACL were applied to a centralized location. Because the ACL is being applied at the point of ingress, there is only a single source IP address (theoretically). Cisco switches perform source substitution on these ACLs to make it even easier. With source substitution, the "any" keyword in the source field of an ACL is replaced with the actual IP address of the host on the switch port.

Using the same formula for six destinations and four permissions, we have:

```
1 source * 6 destinations * 4 permissions = 24 ACEs
```
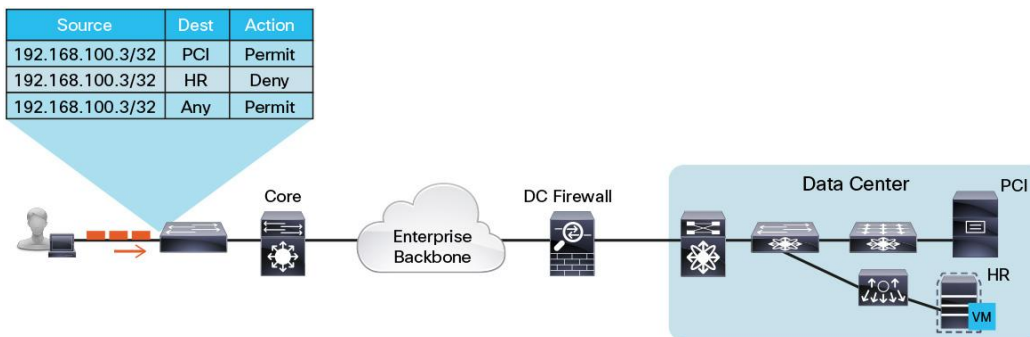
However, there are a few complications with using ACLs on access layer devices. Two major complications are the size of the access lists and the need to regularly maintain them.

If ACLs are used to explicitly defend host networks, they must be updated regularly for all the new destinations that get added to the network. This maintenance can cause an exorbitant amount of operational expense. Additionally, a switch will be able to apply a limited number of ACEs.

ACLs get loaded into and executed from ternary content addressable memory (TCAM). Access layer switches have a limited amount of TCAM, which is usually assigned per application-specific integrated circuit (ASIC). Therefore the number of ACEs that can be loaded depends on a number of factors, such as the number of hosts per ASIC and the amount of free TCAM space.

Due to that limited amount of TCAM, ACLs cannot be overly large, especially when the access layer may be a mixture of different switches, each switch having a different level of TCAM per ASIC. The best-practice recommendation is to keep the ACEs less than 64 per dACL. This figure may need to be adjusted for your specific environment, but it is a good place to start.

Figure 4.   Ingress ACLs

Source	Dest	Action
192.168.100.3/32	PCI	Permit
192.168.100.3/32	HR	Deny
192.168.100.3/32	Any	Permit

What Is Cisco TrustSec Technology?

The Cisco TrustSec solution simplifies the provisioning and management of highly secure access to network services and applications. Unlike access control mechanisms that are based on network topology, Cisco TrustSec policies use logical groupings. Highly secure access is consistently maintained even as resources are moved in mobile and virtualized networks. Decoupling access entitlements from IP addresses and VLANs simplifies security policy maintenance tasks, lowers operational costs, and allows common access policies to be consistently applied to wired, wireless, and VPN access. Cisco TrustSec classification and policy enforcement functions are embedded in Cisco switching, routing, wireless LAN, and firewall products. By classifying traffic according to the contextual identity of the endpoint instead of its IP address, the Cisco TrustSec solution enables more flexible access controls for dynamic networking environments and data centers.

The ultimate goal of Cisco TrustSec technology is to assign a tag (known as a Security Group Tag, or SGT) to the user's or device's traffic at ingress (inbound into the network), and then enforce the access policy based on the tag elsewhere in the infrastructure (in the data center, for example). This SGT is used by switches, routers, and firewalls to make forwarding decisions. For instance, an SGT may be assigned to a Guest user, so that Guest traffic may be isolated from non-Guest traffic throughout the infrastructure. Here is a list of some very common security groups:

- Network Infrastructure: This SGT gets assigned to all the switches, routers, WLCs, and firewalls within the organization
- Network Services: This SGT is assigned to the servers providing common services (Domain Name System, Dynamic Host Configuration Protocol, Network Time Protocol, etc.) that most everyone should be able to reach
- Executive: Many organizations classify their executives with a separate SGT, simply to ensure that Executives will never be denied access to anything
- Sales
- Finance
- HR
- Line of Business 1: SGTs are used quite often when an umbrella company has many lines of business and those lines of business cannot have access to each other's data
- Line of Business 2, 3, and so on: See the previous entry

Note: Each end user or end device may be assigned only one SGT.

So, What Is a SGT?

An SGT is a 16-bit value that the Cisco ISE assigns to the user or endpoint's session upon login.

The network infrastructure views the SGT as another attribute to assign to the session and will insert the Layer 2 tag to all traffic from that session. Figure 5 shows how the tag is inserted into the Layer 2 frame

Figure 5.   Layer 2 Ethernet Frame with SGT
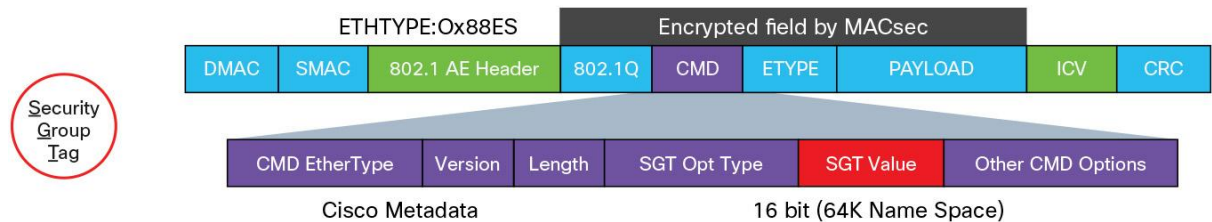
Note: Layer 2 tag insertion requires hardware support. Please refer to the section on Propagation for further details.

An SGT can represent the context of the user and device (Figure 6). Here is an example:

The customer is a retail organization, and it accepts credit cards from its customers, which places it under the domain of PCI compliance. Access to any server housing credit card data has to be protected as strictly as any technology will allow.

In this customer's case, you can define a rule in the Cisco ISE that looks for machine and user authentication (Extensible Authentication Protocol Chaining). It also verifies that the user is a member of a PCI group in Active Directory and that the machine's posture is compliant. If the user and machine meet all three conditions, then an SGT named "PCI" is assigned. No access is granted to PCI servers without the PCI SGT.

So, as you can see, SGTs can be based on the full context of the authentication or simply based on a single condition, such as "Guest."

Note: The endpoint itself is not aware of the tag. It is known in the network infrastructure. Within the network infrastructure, only "trusted," or authenticated, peers can apply the tag.

Note: For additional details on EAP Chaining, please reference:
http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns744/docs/howto_80_eapchaining_deployment.pdf.

Figure 6.   SGT Applied to a Session

```
C3750X#sho authentication sess int g1/0/2
            Interface:  GigabitEthernet1/0/2
          MAC Address:  0050.5687.0004
           IP Address:  10.1.10.50
            User-Name:  employee1
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
      Oper control dir:  both
        Authorized By:  Authentication Server
           Vlan Group:  N/A
              ACS ACL:  xACSACLx-IP-Employee-ACL-(
                  SGT:  0002-0
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0A01300200000022DC6C328F
      Acct Session ID:  0x00000033
               Handle:  0xCC000022

Runnable methods list:
       Method    State
       dot1x     Authc Success
```

Defining the SGTs

In the following example, the Cisco ISE will serve as the single source of truth for which SGTs exist, and the Cisco ISE will consider an SGT a policy result. Therefore, you will create one SGT result for each SGT you wish to define in the environment.
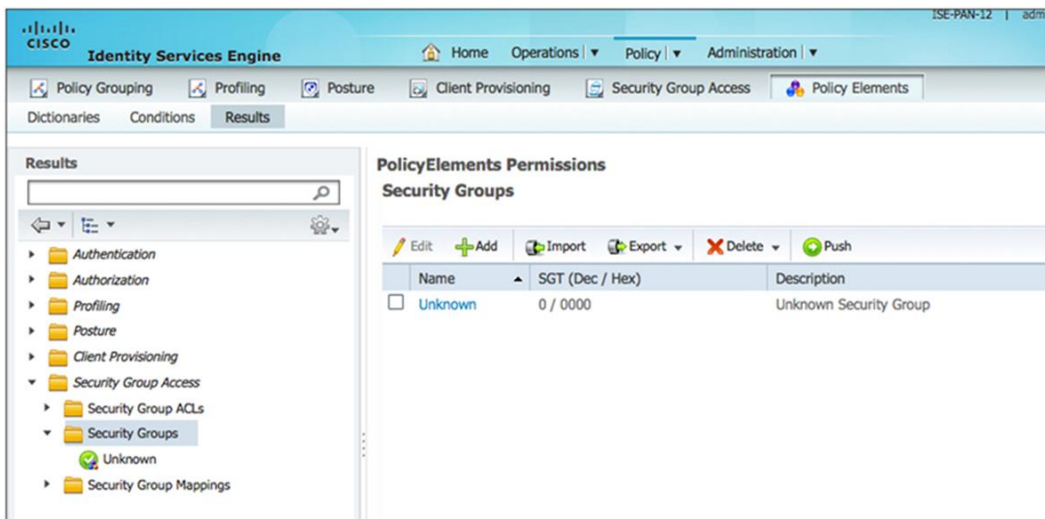
Note: It is recommended that all SGTs be created on the Cisco ISE first to avoid any sort of overlap for SGTs that are statically created on the switches.

Create the SGTs in the Cisco ISE (Figure 7). From within the Cisco ISE GUI:

Step 1. Navigate to Policy → Policy Elements → Results.

Step 2. Select Security Group Access → Security Groups.
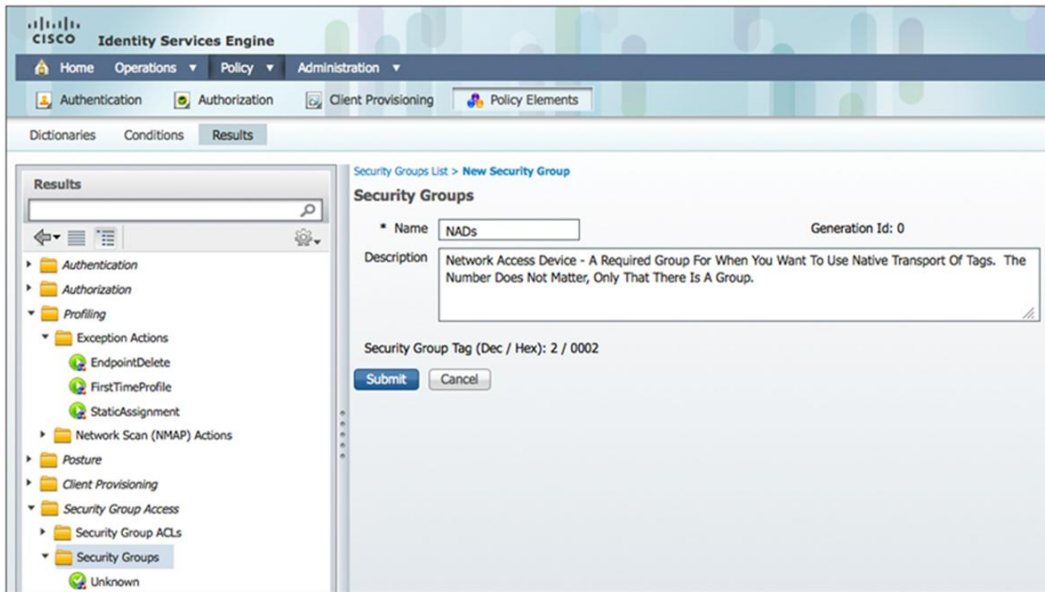
Figure 7.   Security Groups

Notice in Figure 7 that there is a default SGT of 0, "Unknown." This is the tag that will be used if traffic arrives that is untagged. In other words, even the lack of an SGT can be used in the security policy.

Step 3. Click Add.

Step 4. Give the new SGT the name "NADs."

We will begin by creating a security group for network access devices (Figure 8). A dedicated security group will be required for network devices that use native tagging.

Figure 8.   Adding a Security Group for Network Access Devices

Notice in Figure 8 that the SGT value is predetermined. The Cisco ISE can automatically assign the value in order from2 to 65535.
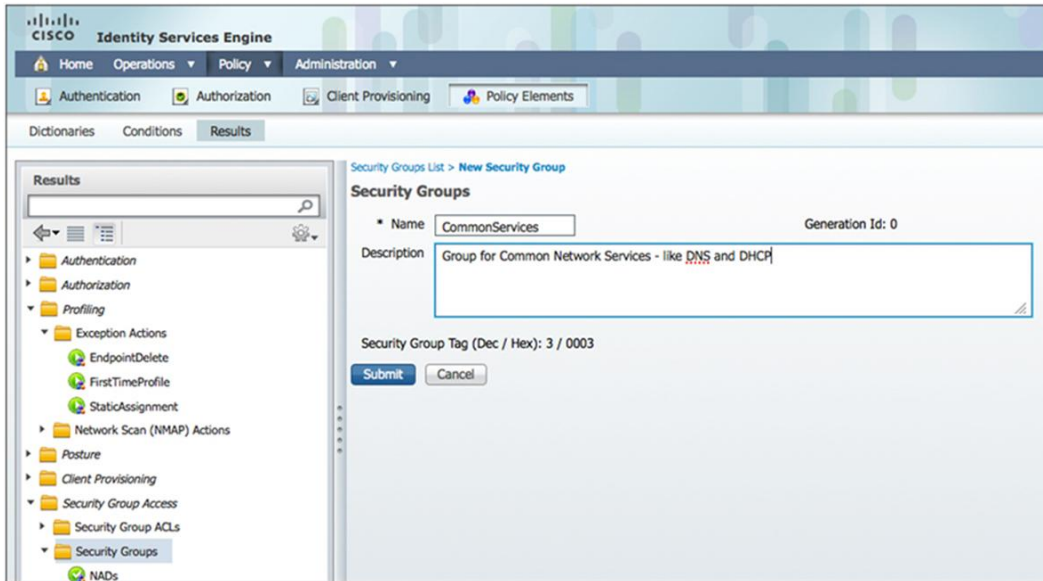
Step 5. Click Submit to save.

Best Practice: A deployment should also have a security group for all the common network services that will exist on a network. These are services, like DNS and DHCP, that should always be accessible by any device.

Step 6. Click Add.

Step 7. Name the new group "Common Services."

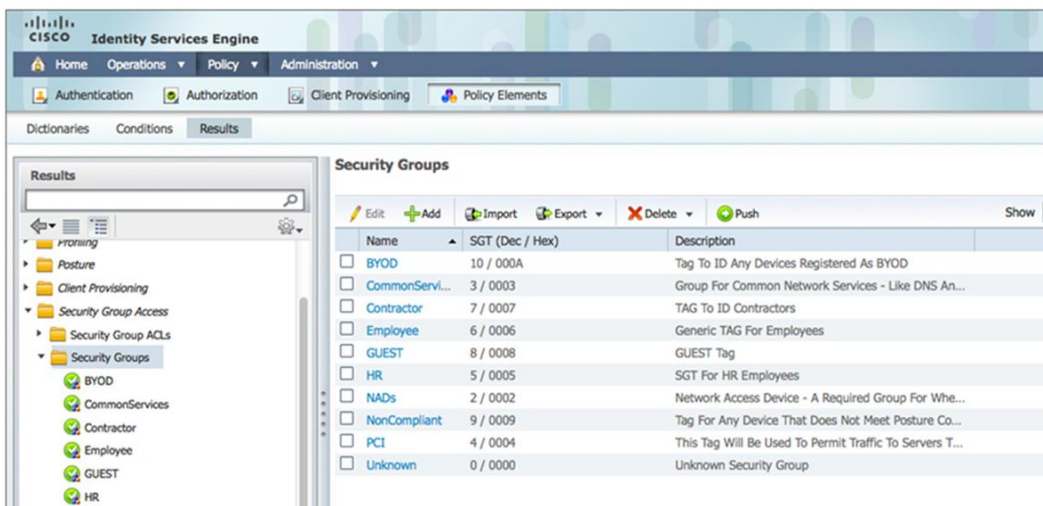**Figure 9.**   Adding a Security Group for Common Services

Step 8. Click Submit to save.

Step 9. Repeat Steps 6 through 8 until you have the appropriate groups created.

Figure 10 shows an example set of security groups.

Figure 10.   Security Groups (Sample List)

Classification

In order to use SGTs within your infrastructure, your devices must support SGTs. All Cisco switches and wireless controllers embedded with Cisco TrustSec technology support the assignment of SGTs. An SGT can be assigned dynamically or statically. Dynamic classification occurs via an authentication sequence, via 802.1x, MAB, or web authentication. When authentication isn't available, static classification methods are necessary. In static classification the tag maps to some thing (an IP, subnet, VLAN, or interface) rather than relying on an authorization from the Cisco ISE. This process of assigning the SGT is defined as "classification." These classifications are then transported deeper into the network for policy enforcement.

Dynamic SGT Classification

Assigning a tag is as simple as adding it as another "permission" or "result" of an authorization in an authorization policy (Figures 11 and 12). When viewing the authorization policy:

Step 1. Edit your existing authorization rule.

Step 2. Click the plus sign (+) under permissions.

Step 3. Click the plus sign (+) next to Authorization Profile.

Step 4. Choose Security Group.

Step 5. Select the appropriate security group to apply.

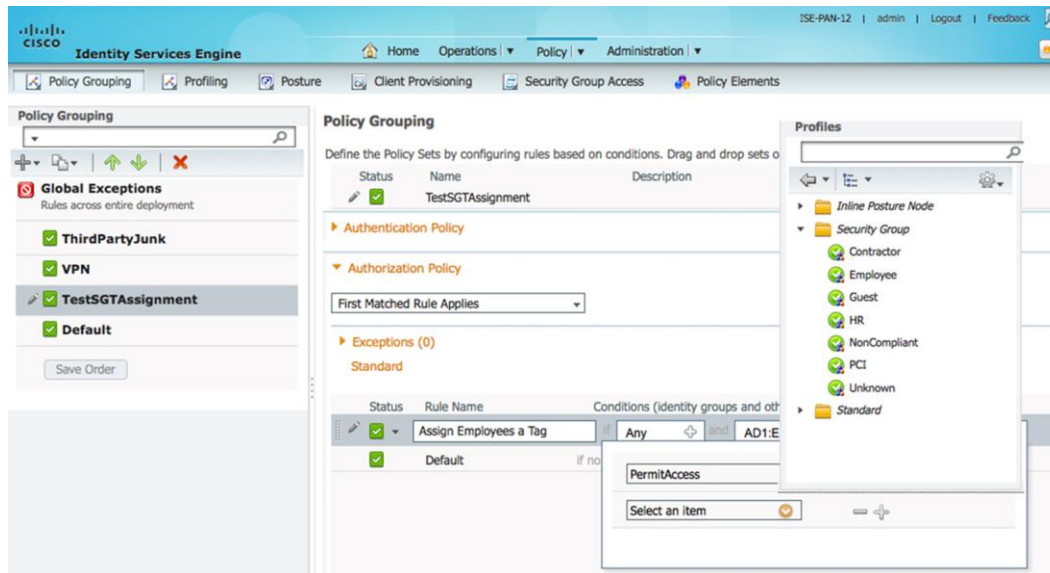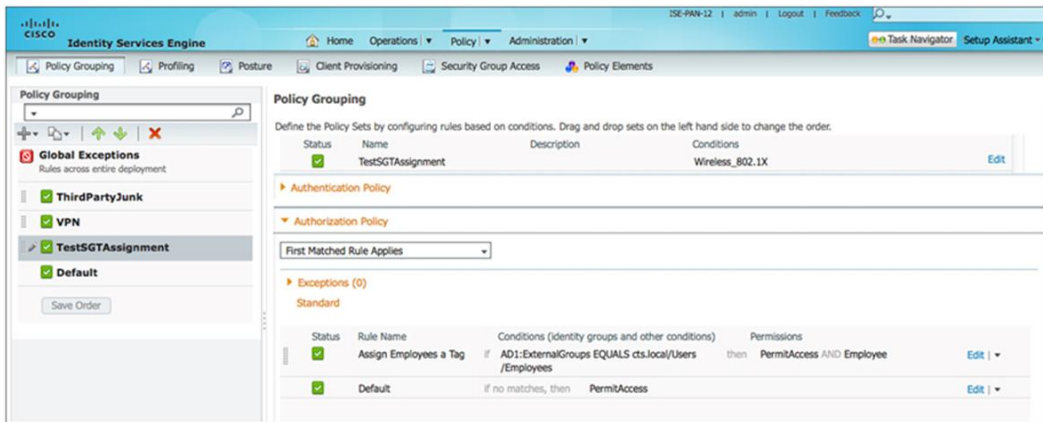**Figure 11.** Adding the SGT to the Results

Figure 12.   Final Rule Showing SGT Assignment

Static Classification

Static classifications are commonly used for static devices, such as data center servers, or topology based policies, such as a subnet based policy. There are several classification methods to meet various deployment needs.

Note: Static classification methods are platform dependent, please see http://www.cisco.com/go/trustsec for further details.

Note: CLI syntax and particular command options vary by platform. The commands shown below represent common syntax. Please consult platform-specific configuration guides for more information.

IP to SGT

This is a method to manually bind the SGT to a particular IP without using the Cisco ISE

```
N7K(config)# cts role-based sgt-map ?
  A.B.C.D  IP Address in format A.B.C.D
```

Subnet to SGT

Subnet-to-SGT mapping binds an SGT to all host addresses of a specified subnet. Cisco TrustSec functions impose the SGT on an incoming packet when the packet's source IP address belongs to the specified subnet.

VLAN to SGT

VLAN to SGT mapping binds packets from a specified VLAN. This is a very useful feature in networks that:

- Do not have authentication enabled
- Use third-party switches
- Have Cisco devices that don't support the Cisco TrustSec solution

Layer 2 Interface to SGT (L2IF-SGT)

This mapping directly maps an SGT to a Layer 2 interface.

Layer 3 Logical Interface to SGT (L3IF-SGT)

L3IF-SGT directly maps an SGT to a Layer 3 interface. Supported interfaces are:

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

```
C6K-DIST(config)#cts role-based sgt-map ?
  A.B.C.D            IPv4 host address
  A.B.C.D/nn         IPv4 prefix <network>/<length>, e.g., 35.0.0.0/8
  X:X:X:X::X         IPv6 host address x:x::y
  X:X:X:X::X/<0-128> IPv6 prefix <network>/<length> (x:x::y/<z>)
  interface          Layer 3 interface
  vlan-list          VLANs to be mapped to SGT
  vrf                Select VPN Routing/Forwarding instance for the binding
```

Port to SGT

Port-to-SGT mapping binds the SGT to a port. Any device connecting to this port assumes the classification represented by tag, as illustrated below:

```
N5K(config)# int <interface>
N5K(config-if)# cts manual
N5K(config-if-cts-manual)# policy static sgt 0x3 ?
  <CR>
  trusted  Specify trust state of the link
```

Port Profile to SGT

SGT is mapped to a port profile. Any devices connected to the port with the same profile assume the same tag value:

```
N1KV
port profile type vethernet Servers
  vmware port-group
  switchport access vlan 100
  cts sgt 3
```

Manually Binding IP Addresses to SGTs Using the Cisco ISE

As an alternative to assigning the SGT to the port itself, the Cisco ISE added the ability to centrally configure a database of IP addresses and their corresponding SGTs. SGT-capable devices may then download that list from the Cisco ISE (Figures 13 and 14).

Note: This capability is supported on Cisco Nexus® 7000 and Nexus 5000 Series Switches.

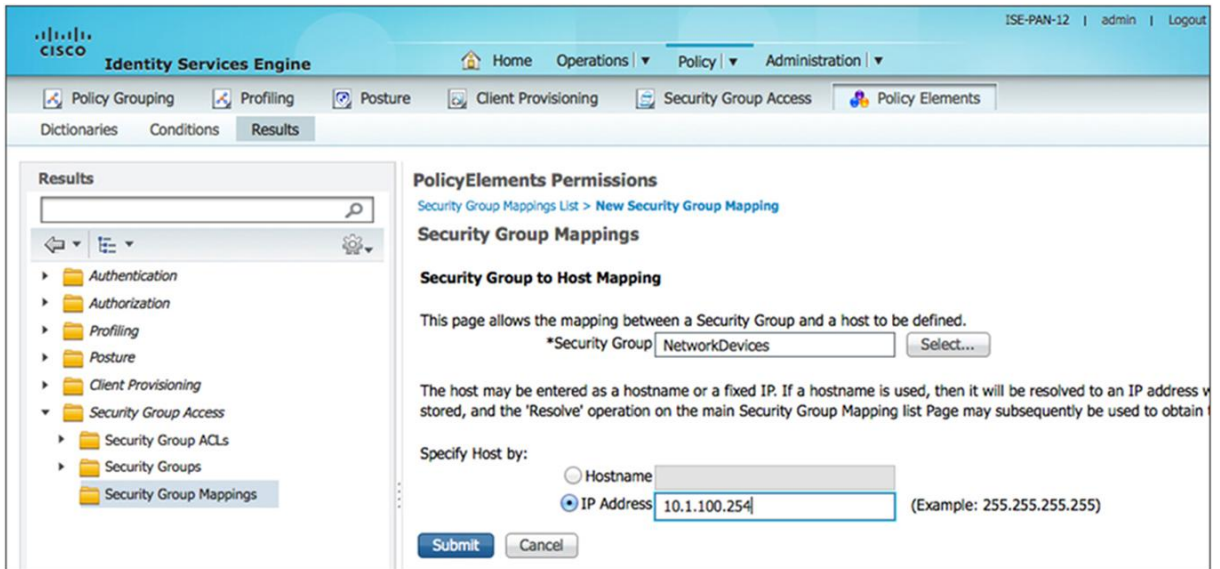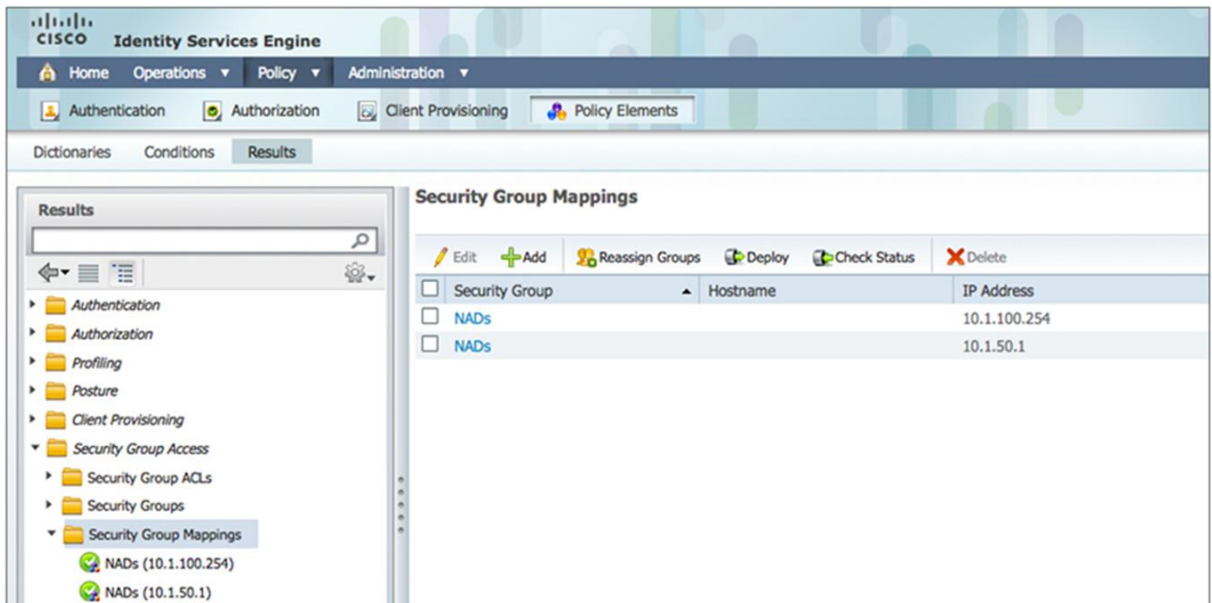**Figure 13.** Mapping an IP Address in the Cisco ISE

Figure 14. Security Group Mappings

Propagation

Now that the SGT is assigned to the user's session, the next step is to communicate the tag upstream to TrustSec devices that enforce policy based on SGTs. This communication process is defined as "propagation". Cisco TrustSec has two methods to propagate a SGT, inline and SXP. Figure 123 shows an example of one access switch that has native tagging. The packets get tagged on the uplink port and through the infrastructure. It also shows a non-inline capable switch, which uses a peering protocol to update the upstream switch. In both cases, the upstream switch continues to tag the traffic throughout the infrastructure.
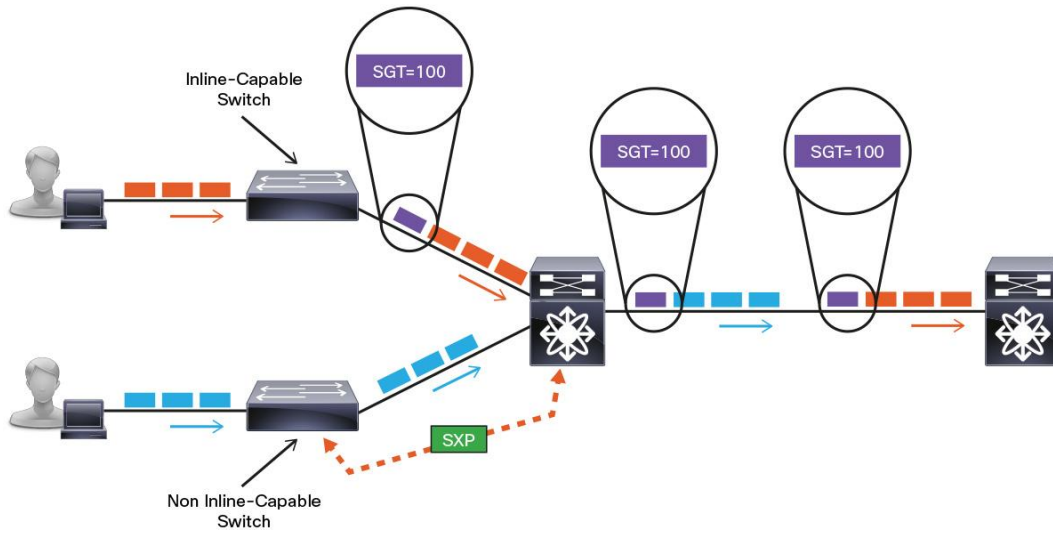
Propagation: Inline

Inline tagging is the ultimate goal. With this approach, the access layer is capable of applying the SGT to the Layer 2 frame as it is sent across the wire to the upstream host. The upstream host will continue that and make sure the tag is applied, and so on. The tag is present throughout the entire infrastructure.

Figure 15. Layer-2 Frame Format with SGT

Native tagging allows the technology to scale virtually endlessly, and it remains completely independent of any layer-3 protocol. In other words, architecturally speaking: if the traffic is IPv4 or IPv6, it does not matter. The tag is completely independent.

Figure 16. Pervasive Tagging
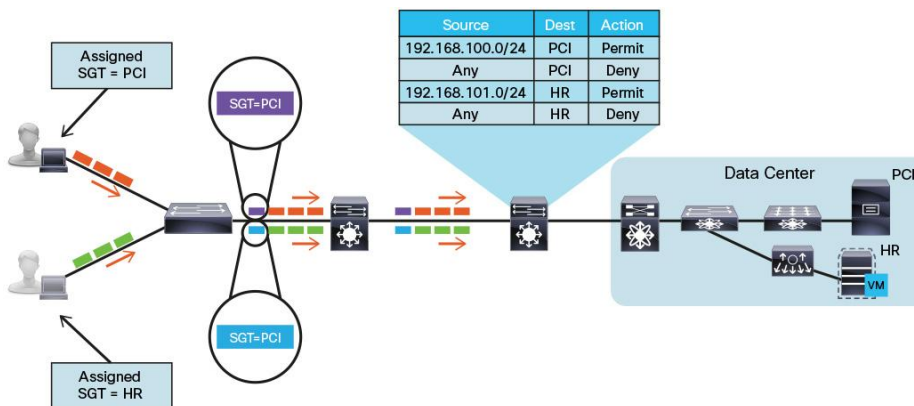
As you can see in Figure 16, when native tags are supported pervasively within the infrastructure, the SGT is communicated hop by hop. This feature provides for end-to-end segmentation and tremendous scale. With the tag being applied to the traffic at every Layer 2 link, you can enforce policy at any point in the infrastructure, and an IP-to-SGT mapping database can be any size, since the database is not being used at all.

Configuring Native SGT Propagation (Tagging)

The next few procedures will show the enabling of inline tagging on three types of switches: a Cisco Catalyst 3500-X Series access layer switch, a Cisco Catalyst 6500 Series distribution layer switch, and a Cisco Nexus data center switch. Figure 17 shows the logical network layout used in the configuration examples to follow.

Note: These three switches were chosen for illustrative purposes. For detailed TrustSec Use Case based configuration, please refer to http://www.cisco.com/go/trustsec.

Figure 17. SGTs from Access to Distribution and Distribution to Data Center
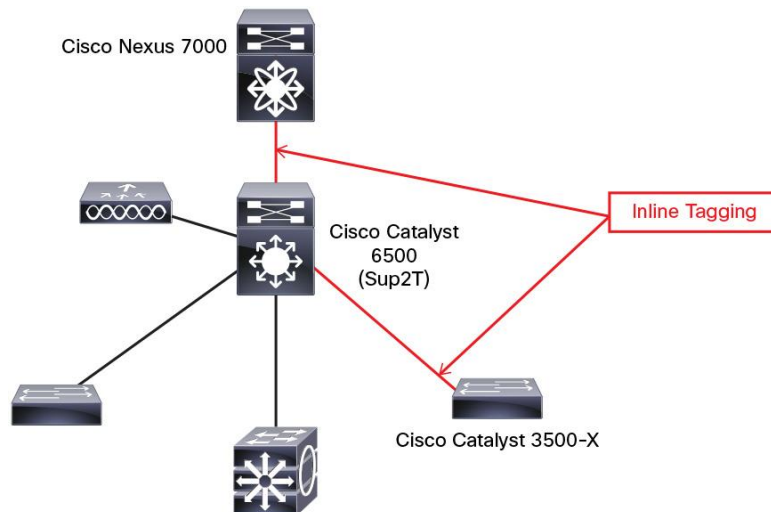
Procedure 1 Configuring SGT Propagation on Cisco IOS Switches

This procedure will focus on the configuration of SGT propagation on access layer switches such as the Cisco Catalyst 3560-X and 3750-X Series that have the ability to use native tags. The Cisco Catalyst 6500 and Cisco Nexus switches will be covered in separate procedures.

From global configuration:

Step 1. Type **cts role-based enforcement**.

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs (discussed in the Enforcement section). However, without this command in the global configuration, the switch will not tag the Layer 2 traffic.

Step 2. Enter into the interface configuration mode of the tagging-capable port by typing interface [interface-name]

Step 3. Type **cts manual**.

We are using **cts manual**, because we are not using NDAC in this guide. The Cisco TrustSec manual mode of operation will allow us to apply the tag to the Layer 2 frame, without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 4. Type **policy static sgt** [sgt-value] **trusted**.

When we created the security groups earlier in this guide, we created a special group for network access devices. We called that security group NADs, and the value of that group was 2 (0x02). That is the value we are applying here with this **policy static sgt 2 trusted** command (Example 4). The trusted keyword in this command helps ensure that no changes are made to the incoming tags, as they are from a trusted source (Example 5).

Note: The SGT assigned to a NAD is called a device SGT. The device SGT is useful for managing control-plane traffic.

Example 4 Enabling Tagging on a Cisco Catalyst 3750-X Series Access Switch

```
C3750X(config)#cts role-based enforcement
C3750X(config)#interface Ten 1/1/1
C3750X(config-if)#cts manual
C3750X(config-if-cts-manual)#policy static sgt 2 trusted
```

Example 5 Verifying Tagging on a Cisco Catalyst 3750-X Series Access Switch

```
C3750X#sho cts interface Ten 1/1/1
Global Dot1x feature is Enabled
Interface TenGigabitEthernet1/1/1:
    CTS is enabled, mode:    MANUAL
    IFC state:               OPEN
    Authentication Status:   NOT APPLICABLE
        Peer identity:       "unknown"
        Peer's advertised capabilities: ""
    Authorization Status:    SUCCEEDED
        Peer SGT:            2
        Peer SGT assignment: Trusted
    SAP Status:              NOT APPLICABLE
        Configured pairwise ciphers:
            gcm-encrypt
            null

        Replay protection:      enabled
        Replay protection mode: STRICT

        Selected cipher:

    Propagate SGT:           Enabled
    Cache Info:
        Cache applied to link : NONE

    Statistics:
        authc success:           0
        authc reject:            0
```

```
        authc failure:          0
        authc no response:      0
        authc logoff:           0
        sap success:            0
        sap fail:               0
        authz success:          3
        authz fail:             0
        port auth fail:         0


    L3 IPM:   disabled.
```

Procedure 2 Configuring SGT Propagation on a Cisco Catalyst 6500 Series Switch

The Cisco Catalyst 6500 Series Switch is a special case. This switch is sometimes used in the access layer, but it is most often used in the distribution layer or even in the data center. There are also a tremendous number of line cards possible for this chassis-based switch, some of which can support native tagging and some of which cannot. Because of the possibility of multiple locations and multiple line cards, the Cisco Catalyst 6500 requires the administrator to set whether the switch should be used for egress (receiving the tag from other devices) or ingress, which would place it at the access layer. These modes are referred to as "reflector" modes.

Note: This switch is unable to be configured for both ingress and egress modes simultaneously.

Ingress Reflector Mode
The ingress reflector mode should be used only in the access layer. This mode allows line cards that do not support Cisco TrustSec features to be used along with a supervisor that does. (An example of this would be a Cisco Catalyst 6504-E chassis populated with a Supervisor 2T and a Cisco Catalyst 6500 Series 6148 line card.) With this mode all packet forwarding will occur on the Supervisor 2T policy feature card. Line cards that use distributed forwarding (such as the 6748-GE-TX) are not supported in ingress reflector mode.

With this mode of operation, the Cisco ISE is able to assign an SGT to a device entering the access layer with any supported line card, but that tag is applied only to network traffic leaving one of the ports physically on the Supervisor 2T. In other words, the switch can apply the tag on an uplink port, but not any of the downlink ports. Additionally, the switch will not be able to read the incoming tag on any ports except the ones physically on the Supervisor 2T module itself.

Note: Using a Supervisor 2T in the access layer is not normally recommended and is also not part of Cisco TrustSec systems testing.

Egress Reflector Mode
The egress reflector mode is normally associated with the Cisco Catalyst 6500 Series being deployed in the distribution layer or data center. With this mode, Cisco TrustSec propagation and encryption (MACsec) may be enabled on the Supervisor 2T and Cisco 6900 Series line cards. These are the line card models most often seen in the distribution layer and thus make for a nice Cisco TrustSec aggregation design. The switch will be able to read all incoming SGT tagged packets and apply that tag to the traffic leaving the switch as well. This is the model of tagging that one normally thinks of when discussing the topic. Additionally, if the Cisco Catalyst 6500 Series is an SXP peer, it is capable of applying the SGT to Layer 2 traffic based on the IP-to-SGT bindings learned through SXP.

From the global configuration on the Cisco Catalyst 6500 Series Switch:

Step 1. Choose the Cisco TrustSec reflector mode by typing **platform cts** {**egress** | **ingress**}.

Since this is a distribution layer deployment of the Cisco Catalyst 6500, we would choose egress mode. If this were an access layer deployment, where end users would be authenticated, then we would have chosen ingress mode.

Step 2. Type **cts role-based enforcement**.

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs (discussed in the Enforcement section). However, without this command in the global configuration, the switch will not tag the Layer 2 traffic.

Step 3. Enter into the interface configuration mode of the tagging-capable port by typing **interface** [interface-name].

Step 4. Type **cts manual**.

We are using **cts manual**, because we are not using NDAC at this point. The Cisco TrustSec manual mode of operation will allow us to send the apply the tag to the Layer 2 frame, without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 5. Type **policy static sgt** [sgt-value] **trusted**.

When we created our security groups earlier in this guide, we created a special group for network access devices. We called that security group NADs, and the value of that group was 2 (0x02). That is the value we are applying here with the **policy static sgt 2 trusted** command (Example 6). The trusted keyword in this command helps ensure that no changes are made to the incoming tags, as they are from a trusted source (Example 7).

Example 6 Enabling Tagging on the Cisco Catalyst 6500 Series Supervisor 2T

```
C6K-DIST(config)#platform cts egress
C6K-DIST(config)#cts role-based enforcement
C6K-DIST(config)#interface Ten1/5
C6K-DIST(config-if)#cts manual
C6K-DIST(config-if-cts-manual)#policy static sgt 2 trusted
```

Example 7 Verifying Tagging on the Cisco Catalyst 6500 Series Supervisor 2T

```
C6K-DIST#show cts interface Ten1/5
Global Dot1x feature is Enabled
Interface TenGigabitEthernet1/5:
    CTS is enabled, mode:    MANUAL
    IFC state:              OPEN
    Authentication Status:  NOT APPLICABLE
        Peer identity:      "unknown"
        Peer's advertised capabilities: ""
    Authorization Status:   SUCCEEDED
        Peer SGT:           2
        Peer SGT assignment: Trusted
    SAP Status:             NOT APPLICABLE
        Configured pairwise ciphers:
```

```
        gcm-encrypt
        null


    Replay protection:      enabled
    Replay protection mode: STRICT


    Selected cipher:

Propagate SGT:          Enabled
Cache Info:
    Cache applied to link : NONE


Statistics:
    authc success:              0
    authc reject:               0
    authc failure:              0
    authc no response:          0
    authc logoff:               0
    sap success:                0
    sap fail:                   0
    authz success:              1
    authz fail:                 0
    port auth fail:             0


    L3 IPM:   disabled.
```

Procedure 3 Configuring SGT Propagation on a Cisco Nexus Switch

From global configuration on the Cisco Nexus Switch:

Step 1. Type **feature dot1x.**

The Cisco Nexus Series Switches require the feature dot1x to be enabled before enabling Cisco TrustSec features.

Step 2. Type **cts enable.**

This command enables security group access, MACsec and NDAC features to be enabled and configured.

Step 3. Type **cts role-based enforcement.**

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs. However, without this command in the global configuration, the switch will not tag the Layer 2 traffic.

Step 4. Enter into the interface configuration mode of the tagging-capable port by typing interface [interface-name].

Step 5. Type **cts manual.**

We are using **cts manual**, because we are not using NDAC at this point. The Cisco TrustSec manual mode of operation will allow us to apply the tag to the Layer 2 frame, without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 6. Type **policy static sgt** [sgt-value] **trusted**.

When we created our security groups earlier in this guide, we created a special group for network access devices. We called that security group NADs, and the value of that group was 2 (0x02). That is the value we are applying here with this **policy static sgt 2 trusted** command (Example 8). The trusted keyword in this command helps ensure that no changes are made to the incoming tags, as they are from a trusted source.

Example 8 Enabling Tagging on the Cisco Nexus 7000 Switch

```
NX7K-CORE(config)# feature dot1x
NX7K-CORE(config)# cts enable
NX7K-CORE(config)# cts role-based enforcement
NX7K-CORE(config)# int eth1/26
NX7K-CORE(config-if)# cts manual
NX7K-CORE(config-if-cts-manual)# policy static sgt 0x2 trusted
```

Propagation: SGT Exchange Protocol (SXP)

In a perfect world, all of your network devices would support inline SGT. This is not reality, because ASIC support is required. So for those platforms that do not support the native SGT capabilities, SXP was created to advertise IP-to-SGT mappings.

SXP is a TCP-based peering protocol. Devices that speak SXP may be a speaker or a listener or both. The definition of a "speaker" is a device that sends the IP-address-to-SGT bindings. The definition of a "listener" is a device that receives the IP-address-to-SGT bindings.

SXP is primarily used to propagate SGTs in two situations:

- The device is not capable of SGT inline tagging. The WLC is one such device. It can function only as an SXP speaker. In Figure 18, the WLC is peered with a device where SGT inline tagging is supported. As the IP-to-SGT bindings arrive, the source IP address is identified and either the associated SGT can be added to the packet and forwarded or an applicable SGACL is enforced.

**Figure 18.** SXP Propagation from the Access Layer (Example 1)
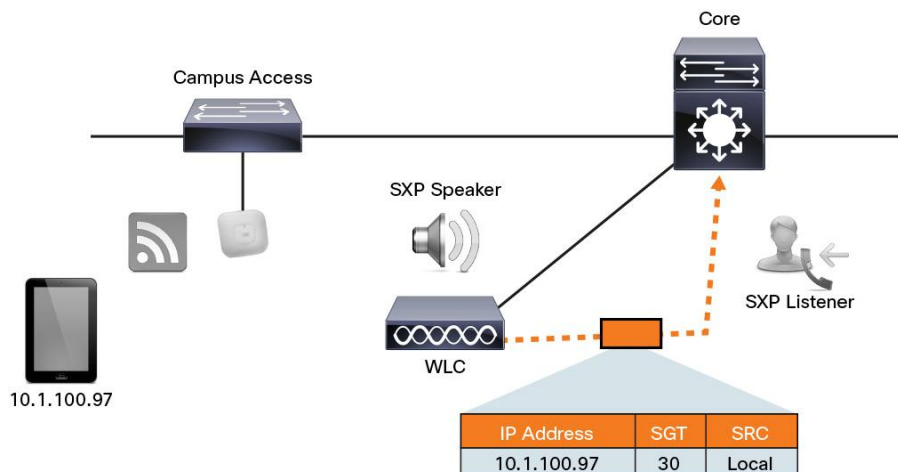
- The adjacent device is a non-SGT-capable device. An SXP peer can be an adjacent device or one that is multiple Layer 2 or Layer 3 hops away. Figure 19 shows an example in which peering to a nonadjacent device allows SGT propagation to continue by skipping over the non-SGT-capable device.

Figure 19. SXP Propagation from the Access Layer (Example 2)
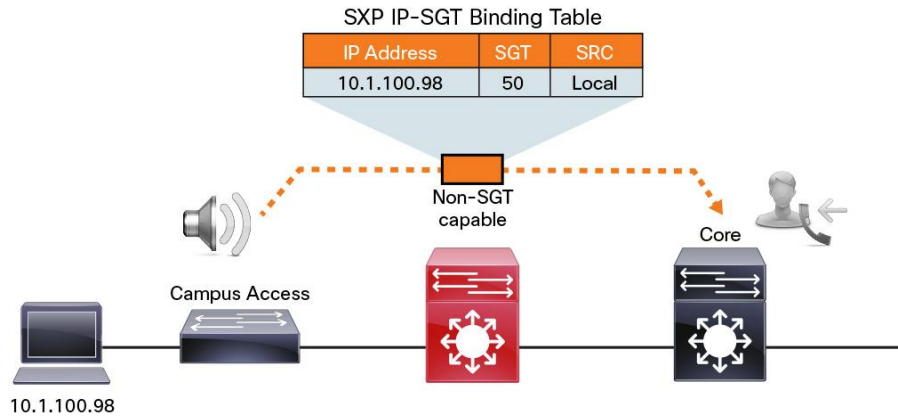
Routing protocols are limited in the number of neighbors they can scale to, and so is SXP. Due to these limitations of scale, SXP design may be architected to be multihop, which allows for aggregation points. Devices like the Cisco Catalyst® 6500 Series Supervisor Engine 2T, the Cisco ASR Aggregation Services Router, or the Cisco Adaptive Security Appliance (ASA), are solid choices for SXP aggregation due to the number of peer connections and IP-to-SGT mappings supported.

Figure 20. SXP Multihop
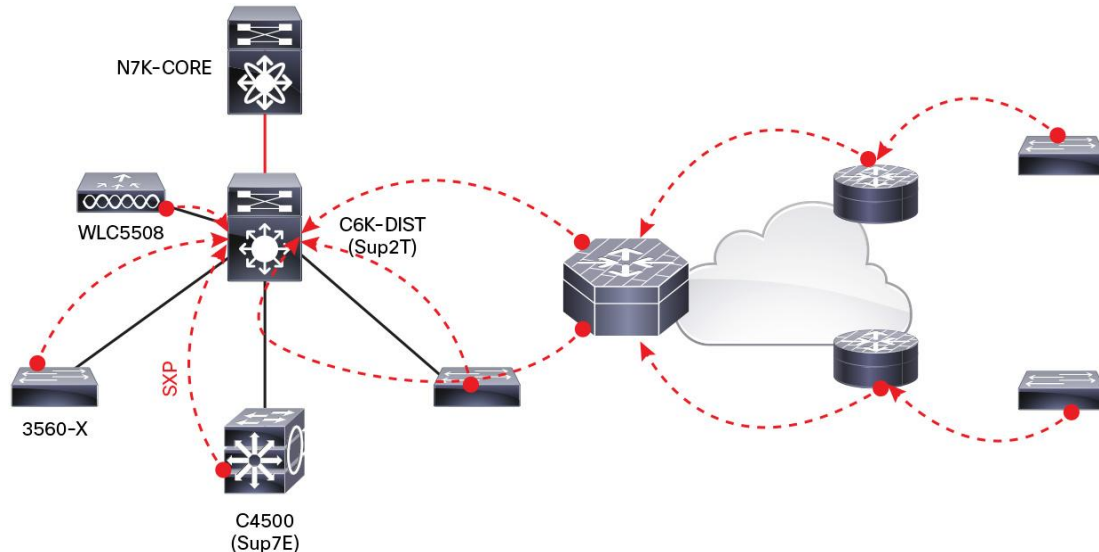
This design has numerous benefits. It does not require an SXP-aware infrastructure along every hop in the network path, and it also provides a deterministic scalable design.

SXPv4: Loop Detection

SXP versions prior to version 4 required careful attention to SXP traffic flow. For example, in Figure 20, SXP traffic flows in one direction (access layer to data center) and from the data center to the distribution layer. This unidirectional traffic pattern is done on purpose, because if SXP traffic were to flow in the opposite direction, an SXP loop could be created. SXP version 4 prevents a loop from occurring.

For example, in Figure 21, there are three switches. An endpoint connects to switch 1, and its IP-to-SGT mapping is entered into the local database. This mapping is communicated to switch 2. Switch 2 communicates the mapping to switch 3. If switch 3 is peered with switch 1, a SXP loop is created. To prevent this, SXPv4 introduces a node ID to identify the switch and the path attribute. These are communicated with the mappings so the receiving switch knows the ID of the source of the mapping and where the binding is forwarded from. When switch1 receives the same mapping, switch1 realizes that it has a local mapping already, and therefore it does not enter the information from switch 3 into its database.

Figure 21. SXP v4 Loop Detection
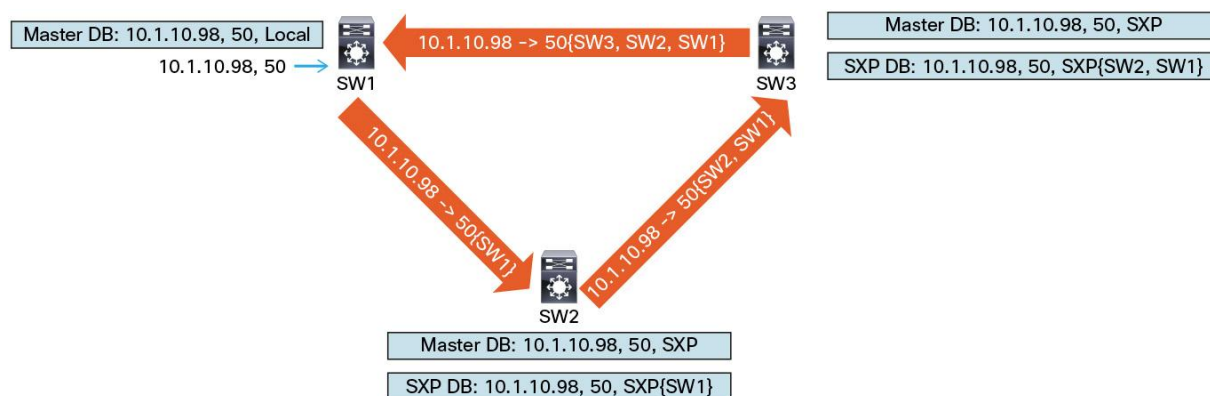
Configuring SXP

In this section, we will configure SXP on Cisco IOS® Software devices, Cisco Wireless LAN Controllers, and the Cisco ASA Next-Generation Firewalls.

Figure 22. Sample SXP Design from Access Layer to Data Center with Cisco ASA Next-Generation Firewall
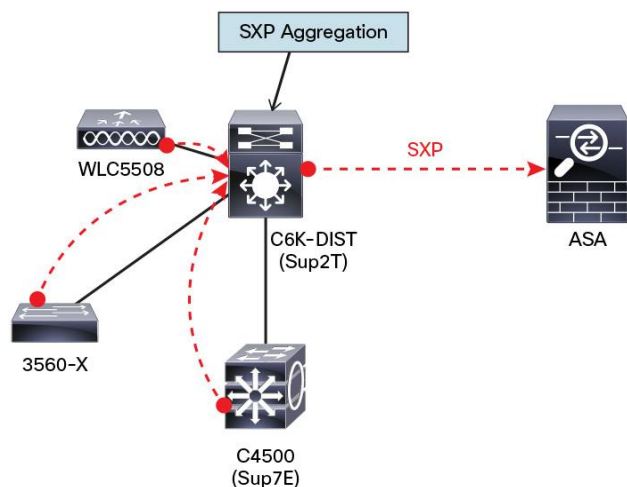
Configuring SXP on Cisco IOS Software-Based Switches

Procedure 4 Configuring SXP on Cisco IOS Devices

From global configuration:

Step 1. Type **cts sxp enable.**

This has turned SXP on globally. Each peer will need to be added individually, as well as setting a global default SXP password.

Step 2. Type **cts sxp connection peer** [peer-ip-address] **password** [default | none] **mode** [local | peer] [listener | speaker].

This command is used to define the **sxp** peer. The options are as follows:

- **password default:** States that the password is defined globally for all SXP connections (at the current time it is not possible to have different SXP passwords per peer)
- **password none:** States that no password is used with this SXP peer
- **mode local:** States that the following **sxp** argument is defining the local side of the connection
- **mode peer:** States that the following **sxp** argument is defining the peer's side of the connection
- **listener:** Defines that the specified device (local or peer) will receive SXP updates through this connection
- **speaker:** Defines that the specified device (local or peer) will send SXP updates through this connection

Step 3. (Optional) **cts sxp default password** [password]

Step 3 is an optional step when your connections will use the globally defined password instead of no password.

Figure 23 illustrates setting up the SXP connection between a Cisco Catalyst 4500 Series Supervisor Engine 7-E (an access-layer device that does not support native tagging) and a Cisco Catalyst 6500 Series Supervisor Engine 2T (a distribution-layer device that supports native tagging). Examples 1 and 2 display the steps.

Figure 23. SXP Between the Cisco Catalyst 4500 and 6500 Series Supervisor Engines
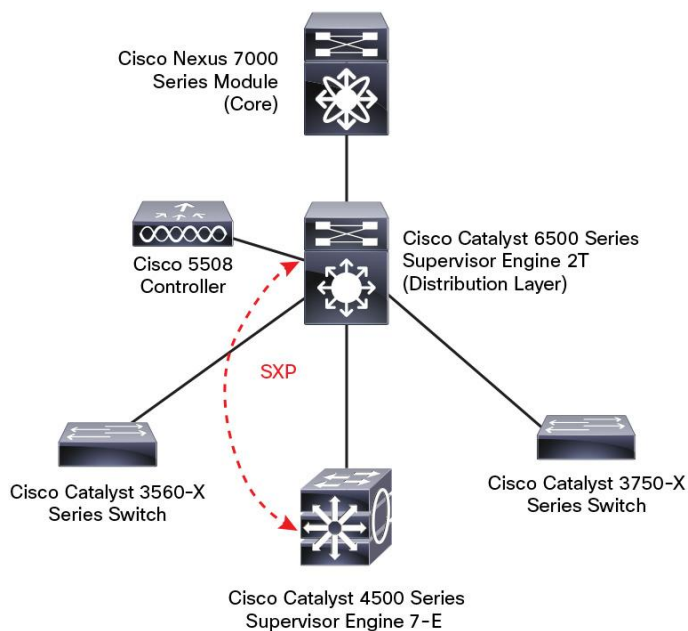
Example 1 Enabling SXP on the Cisco Catalyst 4500 Series Supervisor Engine 7-E

```
4503(config)#cts sxp enable
4503(config)#
*Aug  9 06:51:04.000: %CTS-5-SXP_STATE_CHANGE: CTS SXP enabled


4503(config)#cts sxp connection peer 10.1.40.1 password default mode peer listener
*Aug 10 09:15:15.564: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
*Aug 10 09:15:15.565: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.1, 10.1.40.2>-
1 state changed from Off to Pending_On.
*Aug 10 09:15:15.566: %CTS-3-SXP_CONN_STATE_CHG_OFF: Connection <10.1.40.1,
10.1.40.2>-1 state changed from Pending_On to Off.


4503(config)#cts sxp default password TrustSec123
*Aug 10 09:17:20.936: %CTS-5-SXP_DFT_PASSWORD_CHANGE: CTS SXP password changed.
```

Example 2 Enabling SXP on the Cisco Catalyst 6500 Series Supervisor Engine 2T

```
C6K-DIST(config)#cts sxp enable
Aug 10 16:16:25.719: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.


C6K-DIST(config)#cts sxp default password TrustSec123
C6K-DIST(config)#cts sxp connection peer 10.1.40.2 password default mode peer
speaker
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Off to Pending_On.
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Pending_On to On.
```

Configuring SXP on Wireless LAN Controllers

The Cisco Wireless LAN Controller (WLC) added support for SGT classification and SXP transport in Release 7.2. To enable the use of SXP, you must first enable it globally on the WLC and then add the individual SXP peers (Figure 24).

Procedure 1 Enabling SXP Globally on the WLC

From the Cisco WLC User Interface:

Step 1. Using the top menu navigation, select Security.

Step 2. Along the left side, choose TrustSec SXP (second from the bottom).

Configure the settings on this page to be:

Step 3. SXP State = Enabled.

Step 4. Default Password = the same default password you configured on the switches. All passwords in the SXP "domain" will need to be the same.

**Figure 24.** WLC Global SXP Settings
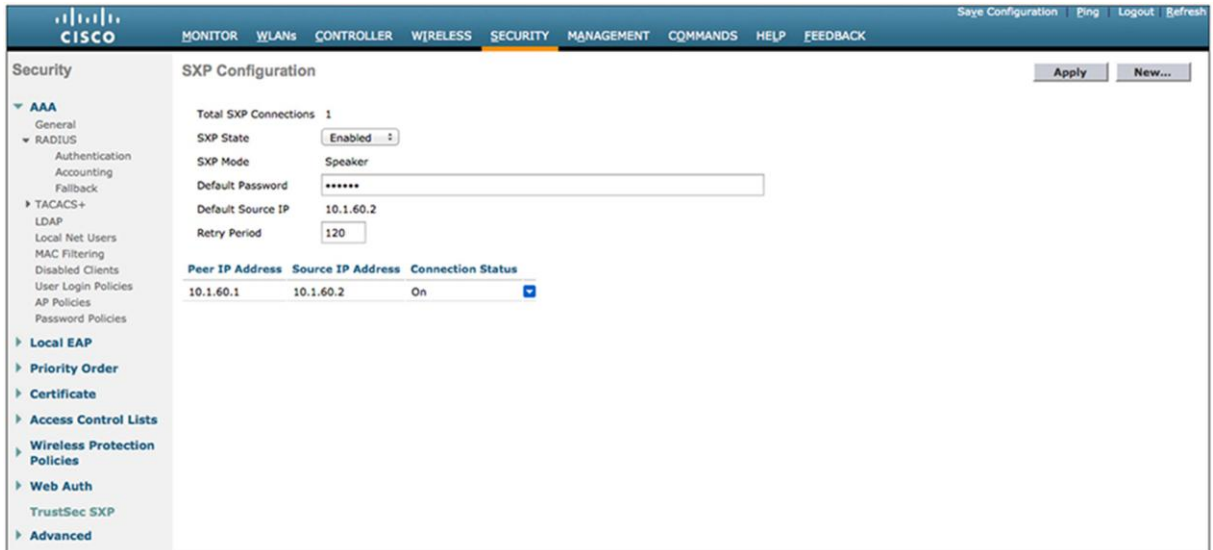
Procedure 2 Adding SXP Peers on the Cisco WLC

The previous procedure turned SXP on globally. Each peer will need to be added individually (Figure 25). To add a new SXP peer (a listener):

Step 1. Click New… (button in the upper-right corner).

Step 2. Type the IP address of the listener peer.

Step 3. Click Apply (upper-right corner).

Figure 25. New SXP Connection on the Cisco WLC
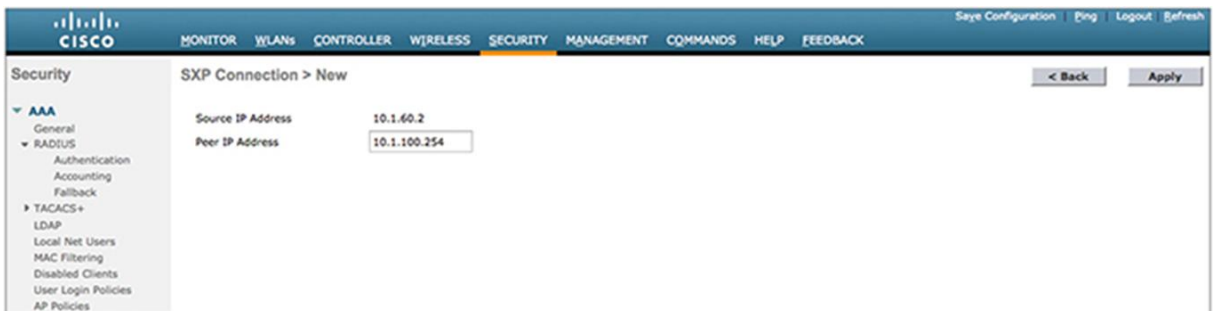
The added peers will be displayed on the Cisco TrustSec SXP page (Figure 26).

Procedure 3 Verifying SXP Connections on the Cisco WLC

From the Cisco WLC GUI:

Step 1. Navigate back to Security → TrustSec SXP.

An SXP peer's status will be listed next to the IP address. Once the peer is configured on the other side, the status should change from "off" to "on."

Figure 26. Cisco TrustSec SXP Page: Peer Status
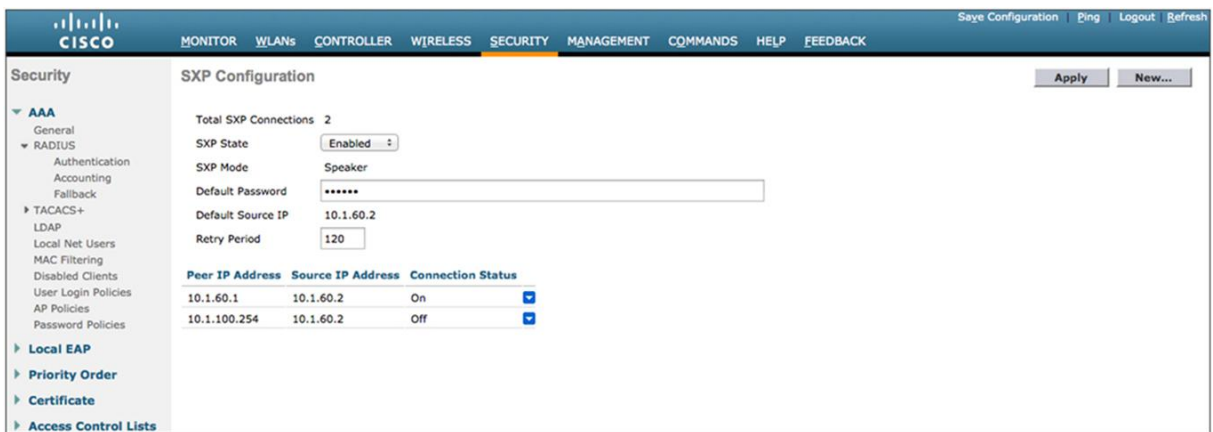
It is also possible to verify the SXP connection from the other side, as shown in Example 3.

From the Cisco IOS Software CLI:

Step 2. Type **show cts sxp connections brief.**

Example 3 Verifying the Connection Between the Cisco WLC and the Cisco Catalyst 6500 Series Supervisor Engine 2T

```
C6K-DIST#sho 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
-------------------------------------------------------------------------
10.1.40.2        10.1.40.1          On                 4:06:36:24 (dd:hr:mm:sec)
10.1.60.2        10.1.60.1          On                 0:00:03:31 (dd:hr:mm:sec)


Total num of SXP Connections = 2
```

Configuring SXP on the Cisco ASA Firewall

Cisco ASA added support for SGT enforcement in its next-generation firewalls. While the Cisco ASA software does not currently support native tagging, it does support SXP for the transport of IP-to-SGT bindings.

It is important to note that the Cisco ASA has multiple functions. These functions include deep packet inspection firewalling and remote access VPN (among many others). At the time this guide was written, only the firewalling functions of SGTs are supported in Cisco ASA Software Release 9.0.1, not the VPN. So the Cisco ASA will **enforce** SGTs, it will receive (**propagate**) SGTs, but it will not assign (**classify**) SGTs.

Procedure 1 Enabling SXP Globally Within the Cisco ASDM

From the Cisco ASA Device Manager (ASDM):

Step 1. Navigate to Configuration → Firewall → Identity by TrustSec.

Figure 27.   Cisco ASDM: Identity by Cisco TrustSec Policy
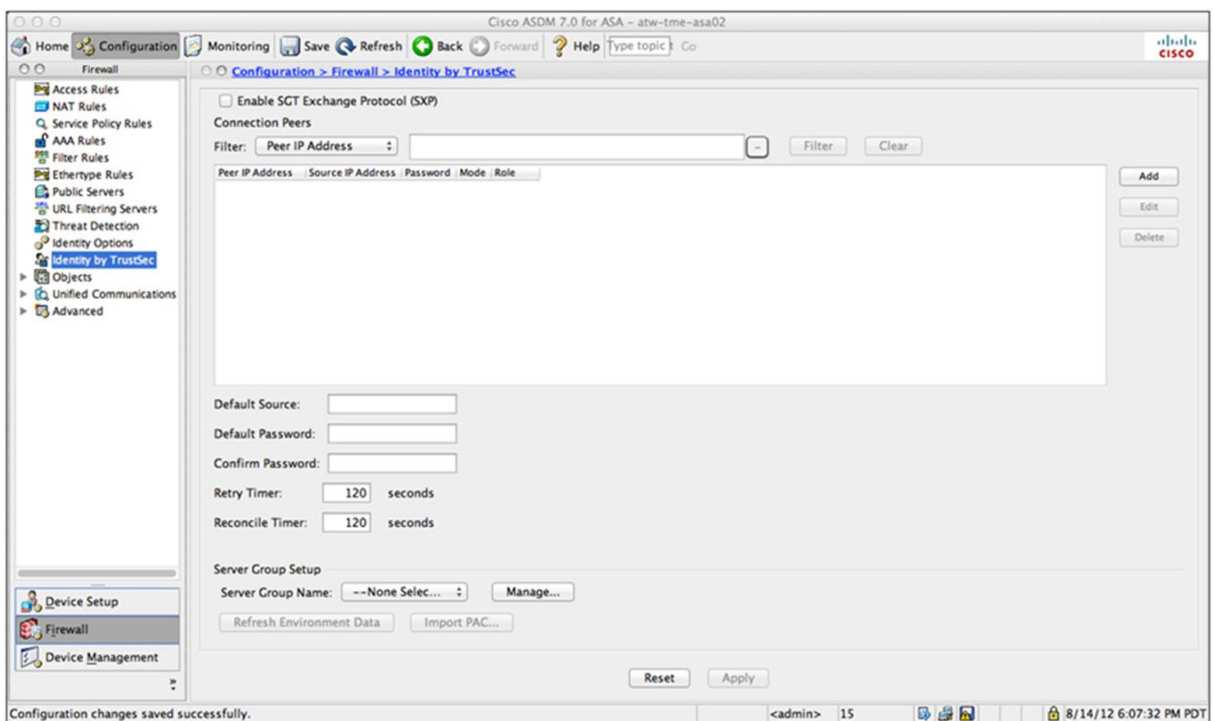
Step 2. You can globally enable SXP by checking the Enable SGT Exchange Protocol (SXP) check box in the upper left (Figure 27).

Procedure 2 Adding SXP Peers in the Cisco ASDM

Step 1. Click Add to add a new SXP peer.

Step 2. In the Add Connection Peer pop-up window, add the IP address of the remote peer.

Step 3. Choose Default for the password (unless you will not be using passwords).

Step 4. Set the mode to Peer.

Step 5. Set the role to Speaker.

Step 6. Click OK.

Figure 28. Adding an SXP Peer
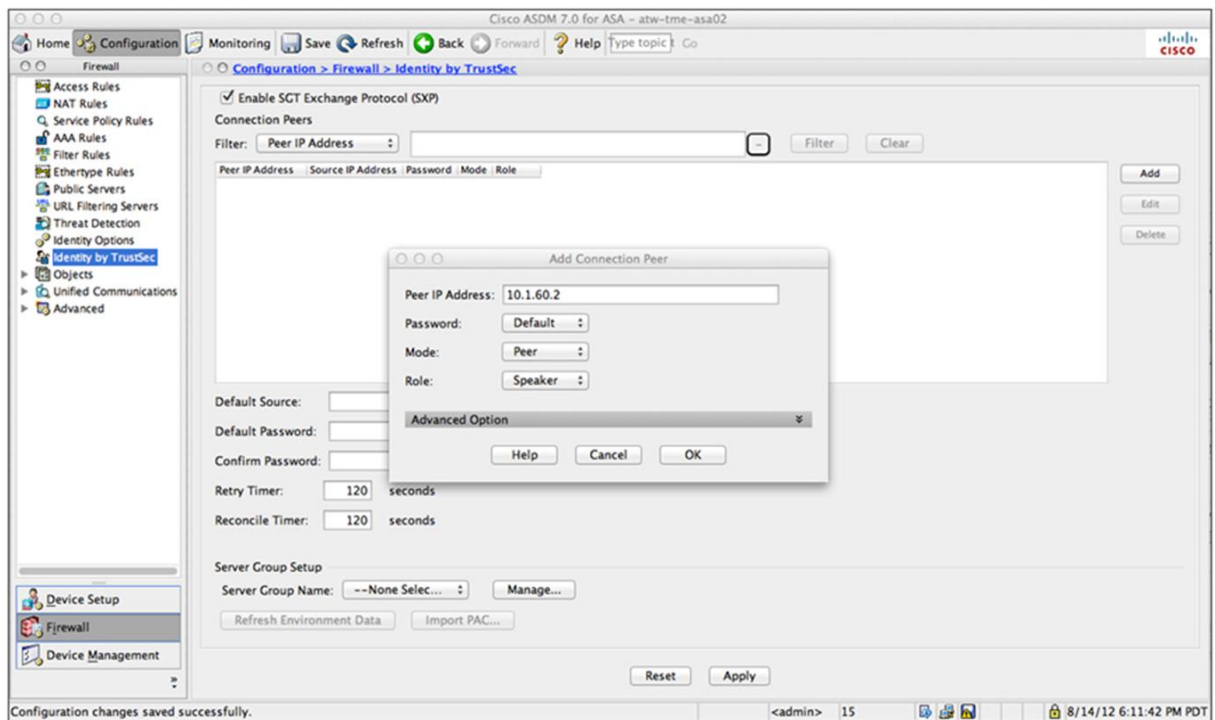
After clicking OK, you are returned to the main Identity by TrustSec page. At this point, you will have SXP enabled, and a single peer defined, but no default password yet (Figure 28).

Step 7. (Optional) If you will be specifying the source IP address of the Cisco ASA, you may configure that source in the Default Source field.

Step 8. Type the default password for your entire SXP deployment.

Procedure 3 Verifying SXP Connections in the Cisco ASDM

From the Cisco ASDM GUI:

Step 1. Navigate to Monitoring → Properties → Identity by TrustSec.

Step 2. Click SXP Connections to see the configured peers and their status (Figure 29).

Figure 29. Monitoring SXP in the Cisco ASDM

Step 3. Click IP Mappings to see any IP-to-SGT mappings that the Cisco ASA has learned (Figure 30).

Figure 30. Monitoring IP-to-SGT Mapping in the Cisco ASDM
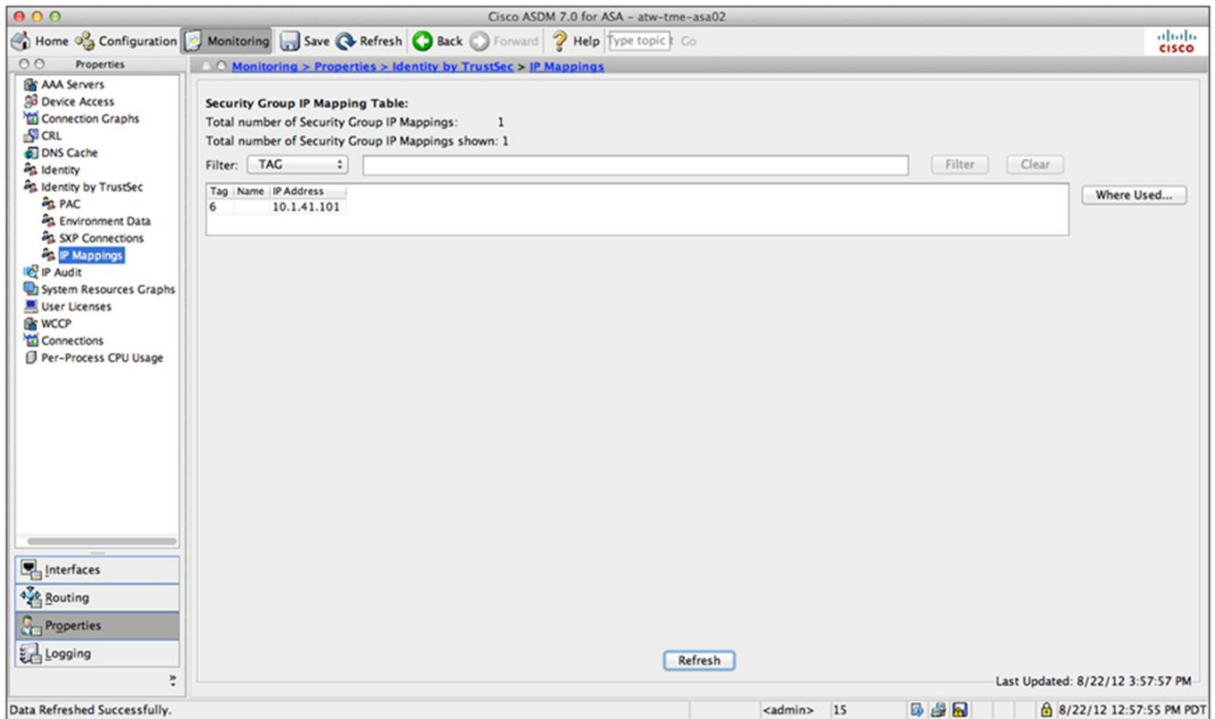
Enforcement

Now that we have security groups assigned (classification), and they are being transmitted across the network (propagation), it is time to focus on the third staple of Cisco TrustSec configurations: enforcement.

There are multiple ways to enforce traffic based on the tag, but ultimately, we can divide them into two major types:

- Enforcement on a switch (SGACL)
- Enforcement on a firewall (SGFW)

SGACL

Historically, enforcement with SGACL was the only option available. It started with the Cisco Nexus 7000 Series Switches and has expanded to the Cisco Nexus 5000 Series, Cisco Catalyst 6500 Series Supervisor 2T, and the Cisco Catalyst 3500-X Series Switches. A major benefit to SGACL use is the consolidation of access ACEs and the operational savings involved with maintenance of those traditional access lists.

One way an SGACL can be visualized is as a spreadsheet. It is always based on a source tag to a destination tag. Figure 31 shows an SGACL policy on the Cisco ISE, which represents the SGACLs in columns and rows. The box that is highlighted in gray shows that when traffic tagged Employee (SGT value 6) attempts to reach the HR destination (SGT value 5), an SGACL named "Permit WEB" and a catch-all of "Deny IP" will be applied. The contents of the Permit WEB ACL are displayed in Figure 32, where you can see that only HTTP and HTTPS are permitted.

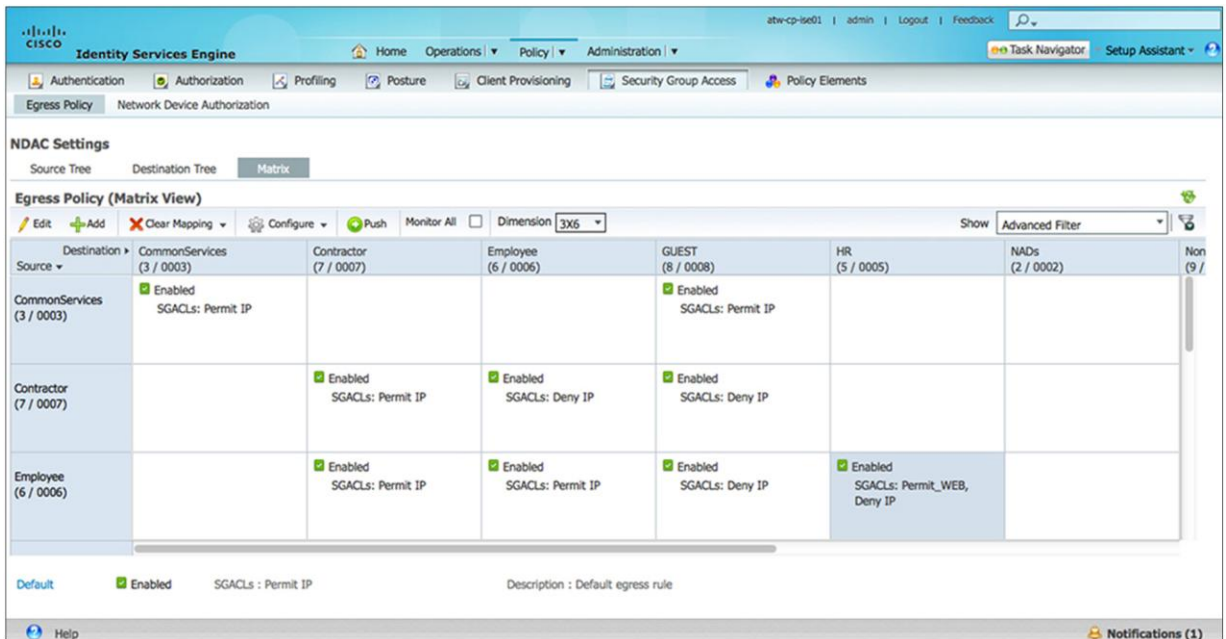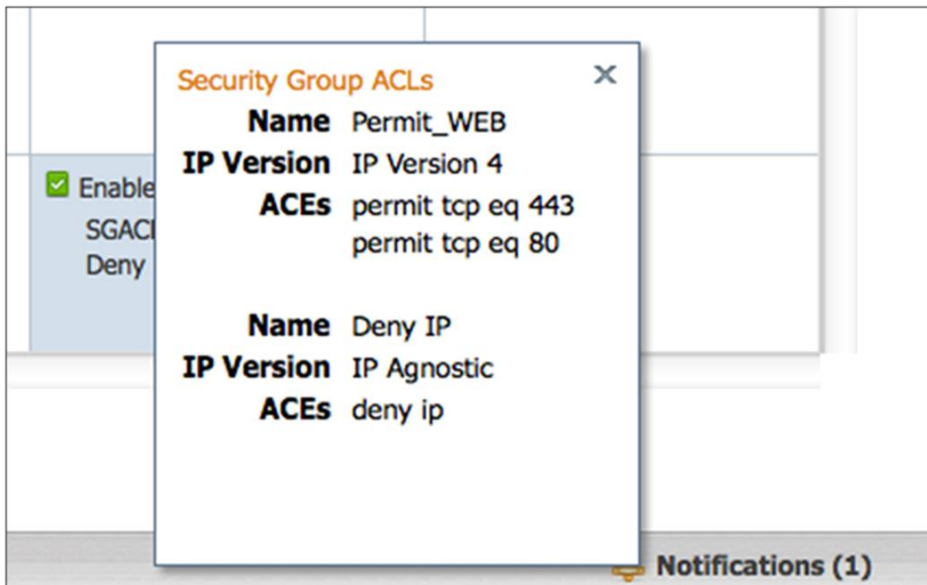**Figure 31.** SGACL Egress Policy: Matrix View

Figure 32. Permit WEB SGACL Contents

As you can see with Figures 31 and 32, the resulting ACL would be to permit HTTP (TCP 80) and HTTPS (TCP 443) and deny all other traffic. This traffic is applied at the egress of the switch where the SGACL is configured. In this case it is applied at the Cisco Nexus 7000 in the data center to traffic attempting to reach the HR server.

This form of traffic enforcement can provide a tremendous savings on the complexity and number of ACEs to maintain. A general formula for the savings is as follows:

```
(sources) * (destinations) * permissions = ACEs
```

With a traditional ACL on a firewall:

```
4 VLANs (sources) * 30 (destinations) * 4 permissions = 480 ACEs
```

Per source IP on a port using dACL:

```
1 group (source) * 30 (destinations) * 4 permissions = 120 ACEs
```

With SGACLs the number of ACEs is a magnitude smaller:

```
4 SGT (sources) * 3 SGT (destinations) * 4 permissions = 48 ACEs
```

There are two main ways to deploy SGACLs: north-south and east-west. "North-south" refers to the case of a user or device being classified at the access layer, but enforcement with the SGACL occurring at the data center. For example, a guest entering the access layer is assigned a Guest SGT. Traffic with a Guest SGT will be dropped if it tries to reach a server with financial data.

"East-west" refers to the case of an SGACL protecting resources that exist on the same switch. For example, if a development server and a production server are on the same Cisco Nexus 5000 Series Switch in the data center, an SGACL may be deployed to prevent the development server from communicating with the production server. Another east-west example is a guest and an employee using the same access layer switch. Traffic may be filtered between these two devices so the guest cannot communicate to the employee who is in the same VLAN on the same switch.

Note: East-west control is supported on switches that support native tagging and enforcement.

Figure 33. North-South and East-West Traffic
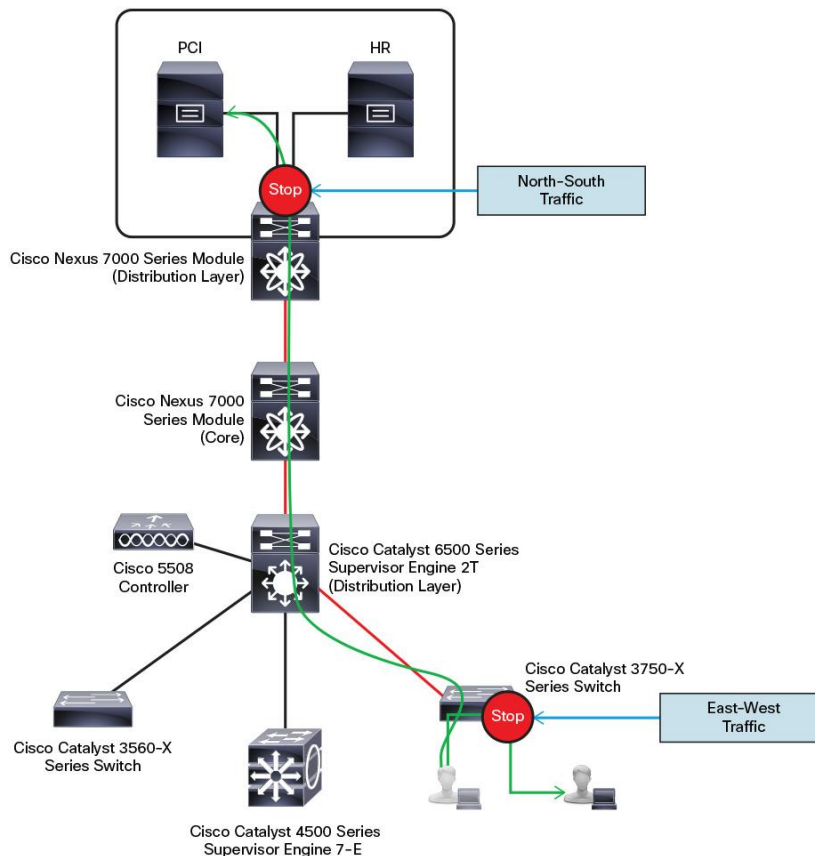
Creating the SGACL in the Cisco ISE

The Cisco ISE provides three views in which you can create SGACLs: two tree views (Source Tree and Destination Tree), and a Matrix view. The Matrix view is the one that looks and acts like a spreadsheet. That view will be our focus in this document.

Procedure 1 Creating a Simple Egress SGACL

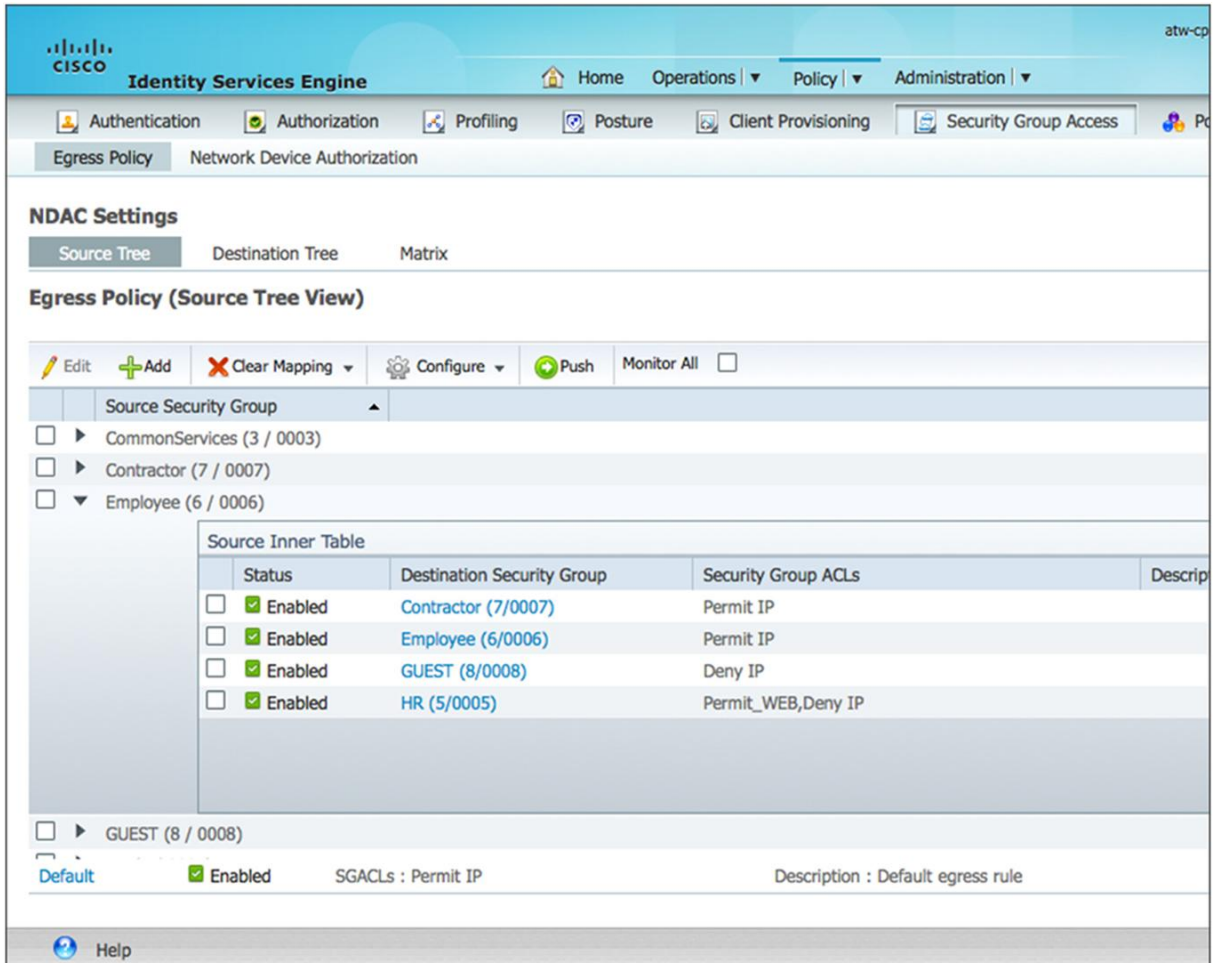From the Cisco ISE Administration GUI:

Step 1. Navigate to Policy → Security Group Access → Egress Policy (Figure 34).

Figure 34. Navigate to the Egress Policy
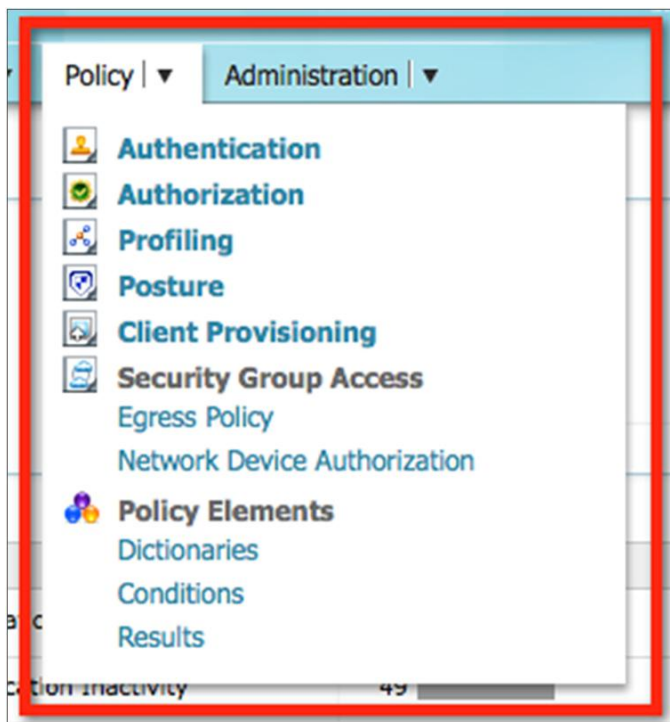
As you can see in Figure 35, the Cisco ISE's default view is the Source Tree view.

Figure 35.　Egress Policy: Source Tree View

Step 2.　Click the Matrix button.

Step 3.　Click the square for the intersection of a Source SGT and a Destination SGT.

In our example we are using the square where the Contractor SGT is trying to reach a device with the HR SGT (Figure 36).

Figure 36. Matrix View: Selecting Contractor (SGT Value 7) to HR (SGT Value 5)

Step 4. Double-click the square.

Step 5. The Edit Permissions screen is displayed.

From this screen, you are able to select an SGACL and pick a final Catch All rule. For the purposes of this example, we will simply deny all traffic from Contractor (value 7) to HR (value 5).

Step 6. Click the drop-down menu for Final Catch All Rule and choose Deny IP (Figure 37).

Step 7. Click Save.

Figure 37. Deny IP from Contractor to HR

Next, we will follow similar steps but create an SGACL that permits employees to access common services like DNS and DHCP. This would normally be handled by a default permit-traffic rule for any undefined "boxes," but we are performing this as an example.
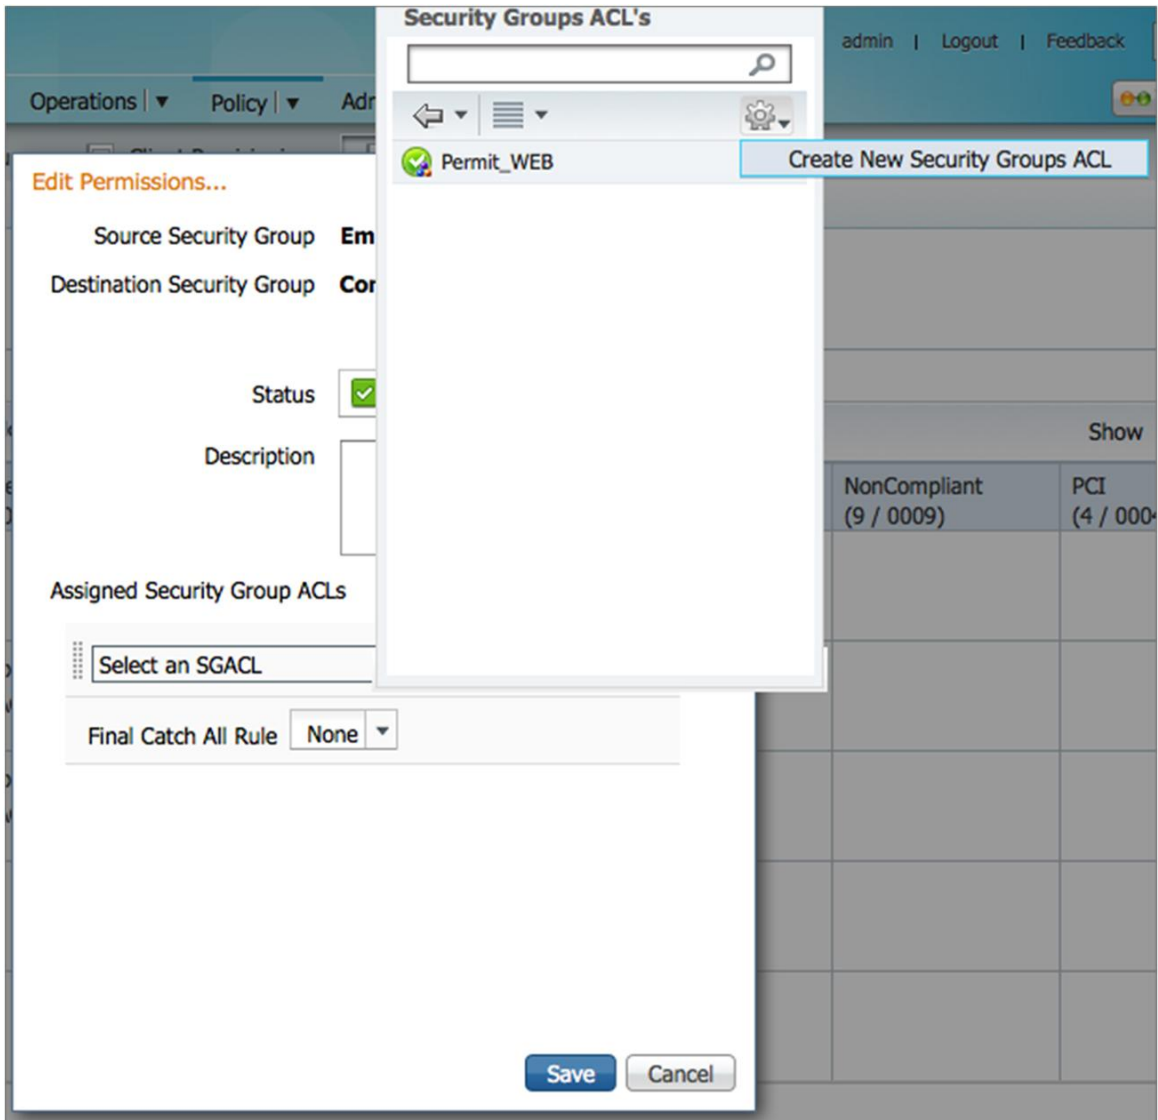
Procedure 2 Creating a More Complex Egress SGACL

Step 1. Double-click the box where the two SGTs intersect.

Step 2. Click the drop-down menu for Select an SGACL.

Step 3. Click the silver cog in the upper-right corner, and select Create New Security Group ACL as shown in Figure 38.

Figure 38. Create a New Security Group ACL

The Create New Security Group ACL window appears. [[**NOTE:** These are called windows or screens elsewhere.]]

Step 4. Type a name and a description in the appropriate fields.

Step 5. IP Version may be IPv4, IPv6, or Agnostic.

The Security Group ACL content field is where the ACEs belong. Notice the SGACL contents in Figure 39. These ACLs are egress only, so you are building an ACE only for the traffic going to the destination tag. There is no source in the ACL at all, only the destination.

Figure 39.   Security Group ACL Contents

Step 6. Click Save.

Step 7. Make sure the new SGACL is selected in the Edit Permissions window.

Step 8. Click the Final Catch All Rule menu. In this example, we are setting the final catch-all to Deny IP (Figure 40).

Step 9. Click Save.

Figure 40. Final Settings in Edit Permissions Window

Configuring the Cisco ISE to Allow the SGACLs to Be Downloaded

Now that the SGACLs are created in the Cisco ISE, we still need download them to the switches that will be enforcing them. This is the first time in this guide that we have needed the data center switches to communicate directly with the Cisco ISE. We are adding the switches to the Cisco ISE so we can download the IP-to-SGT bindings that we created earlier. It can then download the SGACLs that we created along with the egress policies.

Procedure 1 Procedure 1 Adding Enforcement Switches to Cisco ISE Network Devices

The configuration of this should be very familiar to you at this point, with a few minor differences. Let's take a look from the Cisco ISE GUI:

Step 1. Navigate to Administration → Network Resources → Network Devices.

Step 2. Click Add

Step 3. Type the device name, IP address, and any NDGs as you have done in previous Cisco How-To Guides.

Step 4. Add the RADIUS shared secret under Device Authentication Settings.

Now begins the new part. Most likely, you have never clicked the Advanced TrustSec Settings section, unless it was just to look at what might be there. Every device that downloads SGACLs will perform a mutual authentication with the Cisco ISE. As you can see in Figure 41, the device ID and the password must match on both sides.

Step 5. Expand Advanced TrustSec Settings.

Step 6. Enter the device ID and password to match what is configured on the Cisco TrustSec capable switch.

Step 7. Under the SGA Notifications and Updates section, the default settings will most likely suffice, but they can be tuned for your environment.

Step 8. Under Device Configuration Deployment, enter the EXEC Mode username and the two passwords as requested.

This setting will push the IP-to-SGT bindings defined on the Cisco ISE to the switch.

Step 9. Click Save.

Step 10. Repeat steps 1 through 9 for all SGACL-capable switches that will need to download the environment data.

Figure 41. Network Device Settings

Procedure 2 Configuring the Switches to Download SGACLs from the Cisco ISE

Now that the Cisco ISE is configured to allow the switch to download the Cisco TrustSec environment data, we must configure the switch to allow the same.

From global configuration:

Step 1. Configure the device ID and password to match what is configured on Cisco ISE with the **cts device-id** [name] **password** [password] command.

Step 2. Add the Cisco ISE as a RADIUS server by typing **radius-server host** [ip-address] **key** [shared-secret-key] **pac** command.

Note: The proxy autoconfiguration (PAC) keyword is needed before the device can request a protected credential from the Cisco ISE.

Step 3. Create a RADIUS server group using the **aaa group server radius** [group-name] command.

Step 4. Add the server to the group by entering **server** [ip-address].

Step 5. If there is VPN routing and forwarding (VRF) in use, configure which VRF to use to communicate to the Cisco ISE with the **use-vrf** [vrf-name] command.

Step 6. Type **exit** to return to global configuration mode.

Step 7. Configure AAA authentication for 802.1X using the command **aaa authentication dot1x default group** [group-name].

Step 8. Configure AAA accounting for 802.1X using the command **aaa accounting dot1x default group** [group-name].

Step 9. Now do same for authorization: **aaa authorization cts default group** [group-name].

Step 10. Now enter the Cisco TrustSec device ID and password again, to immediately kick-start the PAC file download (Example 9): **cts device-id** [name] **password** [password].

Example 9 Configuring the Cisco Nexus 7000 to Communicate with the Cisco ISE

```
NX7K-DIST(config)# cts device-id NX7K-DIST password TrustSec123
NX7K-DIST(config)# radius-server host 10.1.100.231 key TrustSec123 pac
NX7K-DIST(config)# aaa group server radius ise-radius
NX7K-DIST(config-radius)# server 10.1.100.231
NX7K-DIST(config-radius)# use-vrf default
NX7K-DIST(config-radius)# exit
NX7K-DIST(config)# aaa authentication dot1x default group ise-radius
NX7K-DIST(config)# aaa accounting dot1x default group ise-radius
NX7K-DIST(config)# aaa authorization cts default group ise-radius
NX7K-DIST(config)# cts device-id NX7K-DIST password TrustSec123
```

Procedure 3 Validating the PAC File and Cisco TrustSec Data Downloads

Examples 10 through 14 demonstrate commands on the switch to validate that the PAC files and Cisco TrustSec environment data have been downloaded successfully, while Figure 43 shows the view from the Cisco ISE Live Log.

Example 10 The sho cts pac Command

```
NX7K-DIST(config)# sho cts pac
PAC Info :
==============================
  PAC Type           : Trustsec
  AID                : e9e44428fc9c3fc6be59d35784bb285f
  I-ID               : NX7K-DIST
  AID Info           : Identity Services Engine
  Credential Lifetime : Mon Nov 26 22:57:21 2012

  PAC Opaque         : 000200b80003000100040010e9e44428fc9c3fc6be59d35784bb285f
0006009c00030100d068ae1f1d873e923e2c317e9852bd91000000135034e4aa00093a807e635ba2
e5ae451bdddbd9b17cdcf000dd4516f55324eca75a8dae4786d5e33d669a19d41a62fc9116962c58
b208cac2537eccd2aff08e4b6de47965e69d76e5b16d214030c91f5ebc15ac23e9d5356d60e69cbe
90e9cfa9ee756d259c200dd1afd7abe66c694e0649475665cad145191ac140234d78158e7ceca829
```

Example 11 The sho cts environment-data Command

```
NX7K-DIST(config)# sho cts environment-data
CTS Environment Data
==============================
  Current State         : CTS_ENV_DNLD_ST_ENV_DOWNLOAD_DONE
  Last Status           : CTS_ENV_SUCCESS
  Local Device SGT      : 0x0000
  Transport Type        : CTS_ENV_TRANSPORT_DIRECT
  Data loaded from cache : FALSE
  Env Data Lifetime     : 86400 seconds after last update
  Last Update Time      : Tue Aug 28 22:22:52 2012
  Server List           : CTSServerList1
     AID:e9e44428fc9c3fc6be59d35784bb285f IP:10.1.100.231 Port:1812
```

Example 12 The sho cts role-based access-list Command

```
NX7K-DIST(config)# sho cts role-based access-list
rbacl:Deny IP
        deny ip
rbacl:Permit IP
        permit ip
rbacl:CommonServices
```

Example 13   The sho cts role-based policy Command

```
NX7K-DIST(config)# sho cts role-based policy


sgt:3
dgt:3   rbacl:Permit IP
        permit ip
sgt:6
dgt:3   rbacl:Deny IP
        deny ip


sgt:any
dgt:any rbacl:Permit IP
        permit ip
```

Example 14   The sho cts role-based sgt-map Command

```
NX7K-DIST(config)# sho cts role-based sgt-map
IP ADDRESS          SGT        VRF/VLAN        SGT CONFIGURATION
10.1.100.254        3          vlan:100        Learned on interface:Ethernet1/3
10.1.50.2           2          vrf:1           Learned on interface:Ethernet1/1
```

Figure 42.   Cisco ISE Live Log Showing PAC and Cisco TrustSec Downloads

| Time	Status	Details	Identity	Endpoint ID	IP Address	Network Device	Event	Device Port
Aug 28,12 10:57:58.433 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:58.432 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:58.431 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:57.422 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:47.862 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:47.862 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:47.861 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:37.260 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:26.940 PM	✓	🔍	#CTSREQUEST#			NX7K-DIST	CTS Data Download Succeeded	
Aug 28,12 10:57:22.512 PM	✓	🔍	NX7K-DIST			NX7K-DIST	PAC provisioned	

Security Group Firewalls

Some organizations prefer to do the traffic enforcement on the switching infrastructure and prefer it to be one on a firewall, a device that is purpose-built to do traffic filtering. Cisco has added the ability to enforce traffic on firewalls by implementing the security group firewall (SGFW). Two types of SGFWs exist: the Cisco ASA-based SGFW and the router-based SGFW. This makes sense, since the routers use a zone-based firewall (ZBFW) and the Cisco ASA does not.

Security Group Firewall on the Cisco ASA

Beginning with Cisco ASA Release 9.0.1, the Cisco ASA firewall gained SGFW functionality. The Cisco ASDM supports the full configuration, and therefore the Cisco ASA is the only SGFW that has a GUI (as of the writing of this document).

The SGFW in the Cisco ASA operates on a very simple concept. The powerful firewall policy has been expanded to include source and destination security groups. As you can see in Figure 44, there is a new Security Group column in the Source Criteria and Destination Criteria sections.

**Figure 43.**   Cisco ASDM Firewall Policy
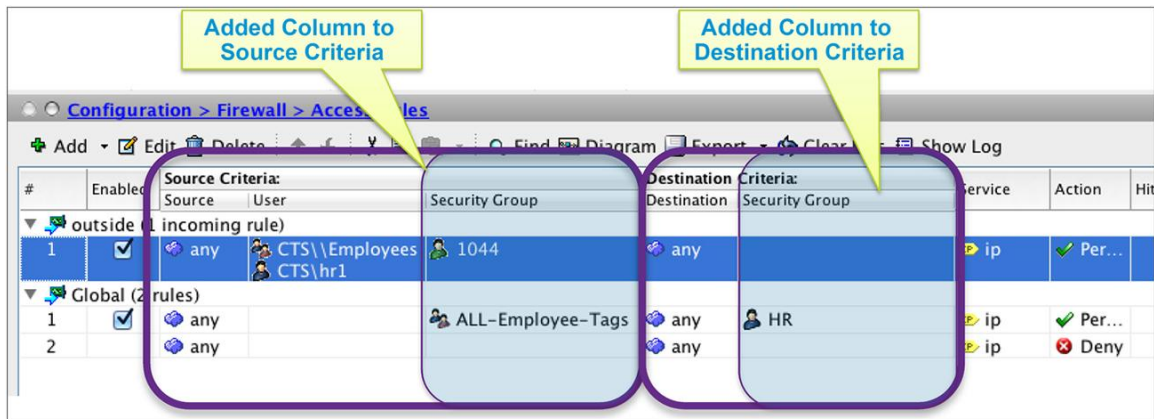
Configuring Cisco TrustSec Downloads from the Cisco ISE Through the Cisco ASDM

Unlike switches, an SGFW does not download the SGACLs from the Cisco ISE. Firewalls tend to have their own security policies. However, the SGFW must still be able to download the list of SGTs that exist as well as the static IP-to-SGT mappings that were created from the Cisco ISE.

Procedure 1 Adding an Authentication Server Group

From within the Configuration → Firewall section of the Cisco ASDM, there is a new option called Identity by TrustSec (Figure 44). This is the primary location for any and all Cisco TrustSec configurations from within the GUI.
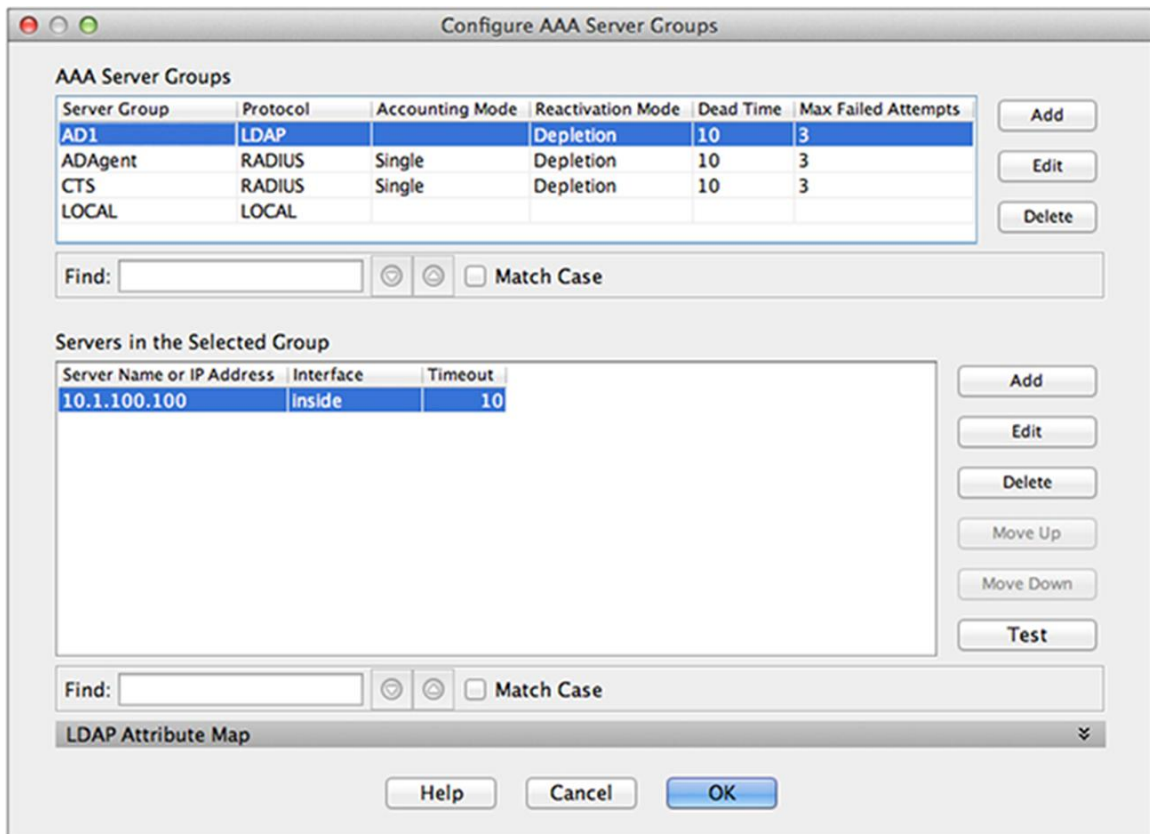
Figure 44. Identity by TrustSec
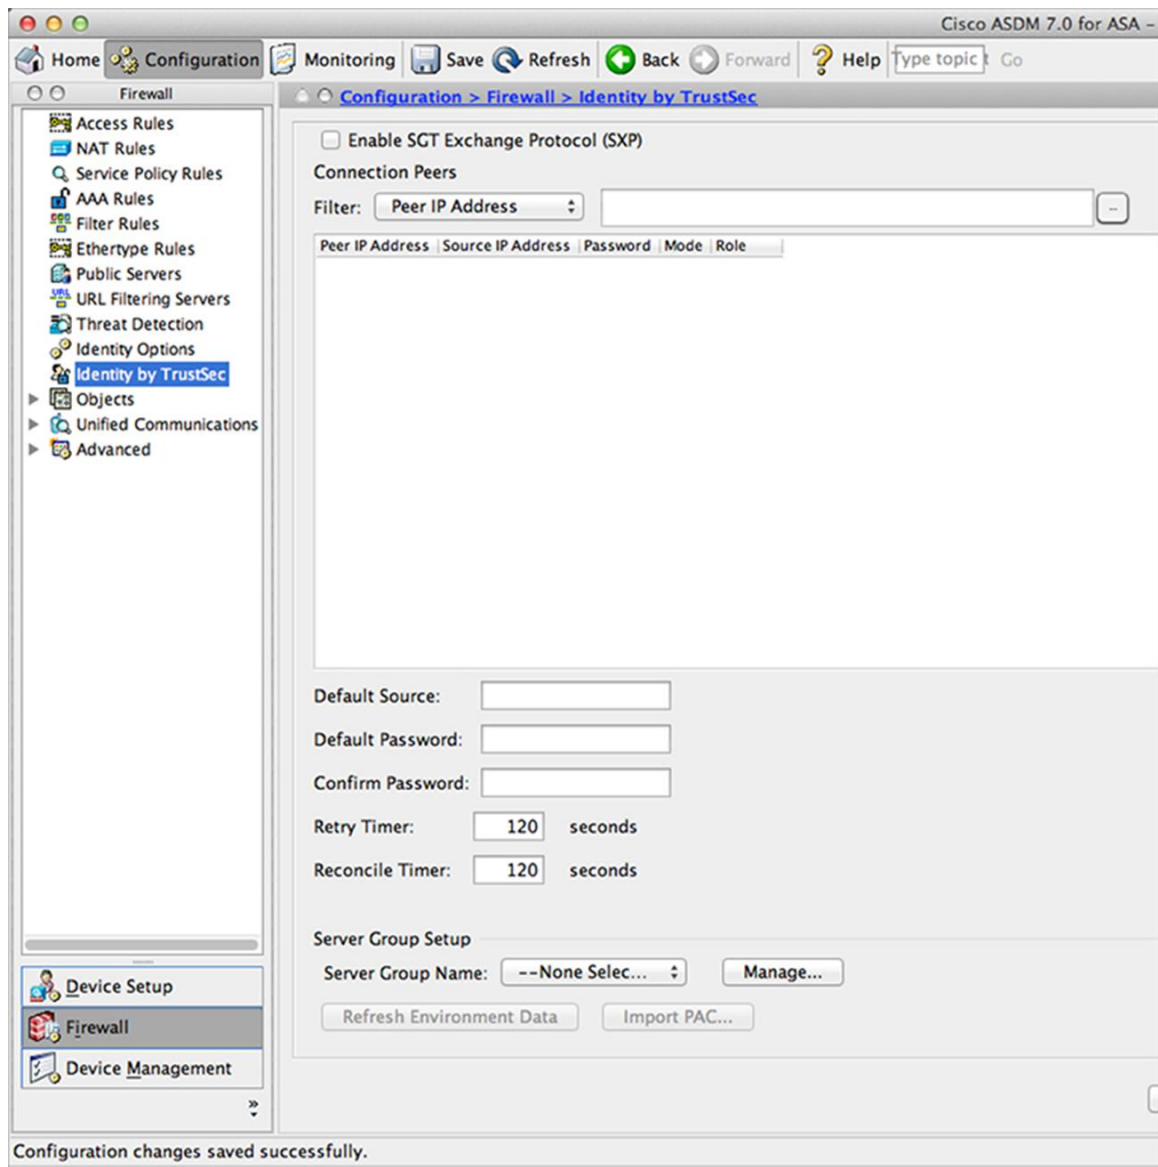
Step 1. To start, we need to add an authentication server. At the bottom of the screen, in the Sever Group Setup section, click Manage.

Step 2. The Configure AAA Server Groups window will appear (Figure 45).

Figure 45. Configure AAA Server Groups

Step 3. Click Add at the top right to create a new server group (Figure 46).

Step 4. Type a name for the new group in the AAA Server Group field.

Step 5. Make sure the chosen Protocol is RADIUS.

Step 6. Click OK.

Figure 46. Add AAA Server Group

Procedure 2 Adding Cisco ISE as a New AAA Server

Step 1. Back in the Configure AAA Server Groups window, make sure the new server group is highlighted, as in Figure 47.

Figure 47. The New AAA Server Group

Step 2. Now click Add on the lower half of the window to add a new authentication server to the group (Figure 48).

Figure 48. The New AAA Server

Step 3. Choose the correct interface to reach the Cisco ISE.

Step 4. Add the IP address of the Cisco ISE.

Step 5. The Cisco ISE uses port 1812 for authentication and 1813 for accounting.

Note: Cisco ASDM default values are 1645 and 1646.

Step 6. Type the shared secret key for the RADIUS communication.

Step 7. Click OK to save the server.

Step 8. Click OK to save the server group.

Step 9. Make sure your new server group is selected in the Server Group Setup section and click Apply.

Procedure 3 Adding the Cisco ASA to Cisco ISE Network Devices

Now that the Cisco ASA is configured to communicate with the Cisco ISE, you must ensure that the Cisco ISE has the Cisco ASA added as a network device. Unlike the Cisco Nexus switches, the Cisco ASA will not automatically download the PAC file. It must be manually generated in the Cisco ISE GUI and imported from the Cisco ASDM.

From within the Advanced TrustSec Features section of the Network Device:

Step 1. Click Generate PAC.

The Generate PAC pop-up screen appears, with the Identity field already populated (Figure 49).

Figure 49. Generate PAC

Step 2. Type in an encryption key that you will remember, and click Generate PAC.

Step 3. The download should begin automatically (Figure 50).

Figure 50. Save the PAC File

Step 4. Save the file to a location that you will remember.

Step 5. In the Cisco ASDM Identity by TrustSec window, click Import PAC (Figure 51).

Figure 51. Import PAC
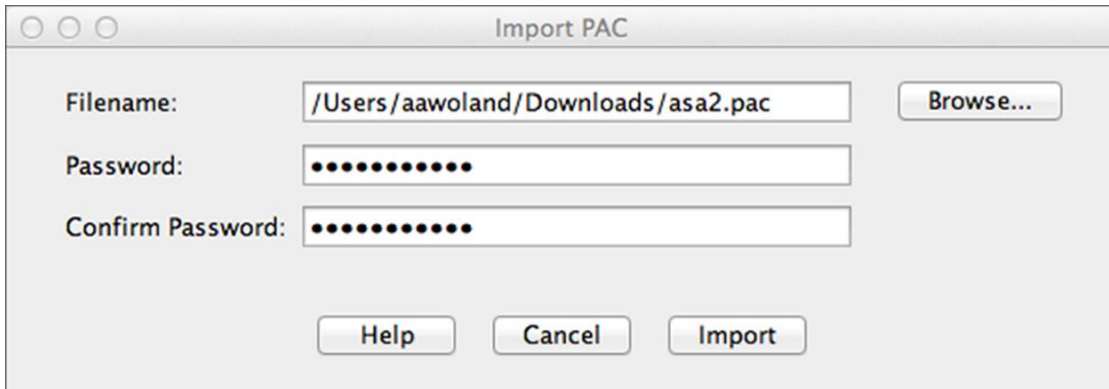
Step 6. Click Browse and choose the saved PAC file.

Step 7. Type the encryption key into the Password and Confirm Password fields.

Step 8. Click Import.

Step 9. You should receive a "PAC Imported Successfully" message.

Procedure 4 Validating the Cisco TrustSec Communication

Within the Cisco ASDM:

Step 1. Navigate to Monitoring → Properties→ Identity by TrustSec.

Step 2. Choose Environment Data.

The screen will show you the status of your communication with ISE, the last successful download, and the list of SGT-to-Security-Group mappings (Figure 52).

Figure 52. Cisco ASDM: Identity by TrustSec (Environment Data)

○ ○ **Monitoring > Properties > Identity by TrustSec > Environment Data**

Environment Data:

```
Status:                    Active
Last download attempt:     Successful
Environment Data Lifetime: 86400 secs
Last update time:          19:24:03 PDT Aug 28 2012
Env-data expires in:       0:23:59:37 (dd:hr:mm:sec)
Env-data refreshes in:     0:23:49:37 (dd:hr:mm:sec)
```

Security Group Table:

```
Valid until:               19:24:03 PDT Aug 29 2012
Total entries:             10
```

Name	Tag	Type
ANY	65535	unicast
CommonServices	3	unicast
Contractor	7	unicast
Employee	6	unicast
GUEST	8	unicast
HR	5	unicast
NADs	2	unicast
NonCompliant	9	unicast
PCI	4	unicast
Unknown	0	unicast

Step 3. To verify the imported PAC file, navigate to Identity by TrustSec → PAC (Figure 54).

Figure 53. Cisco ASDM: Identity by TrustSec (PAC)
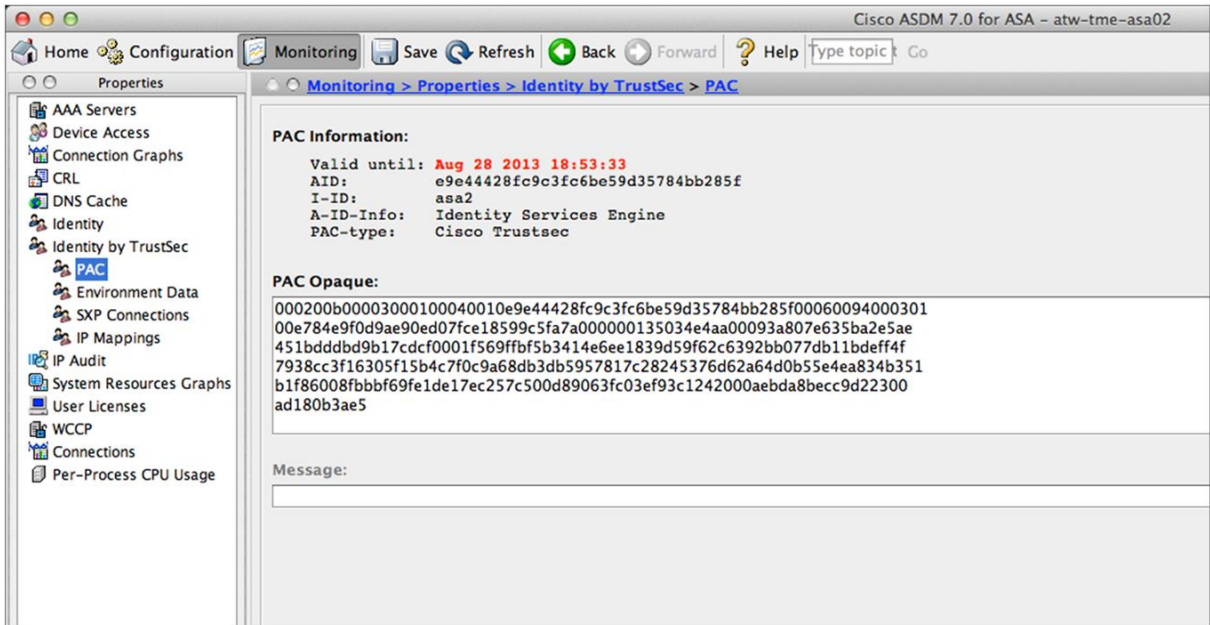
Configuring SGFW Policies Through the Cisco ASDM

Now that all the communication between the Cisco ASA and Cisco ISE is configured and working, the Cisco ASA knows about all the security groups that were configured on the Cisco ISE. Those security groups may now be used in the firewall policy on the Cisco ASA.

Procedure 1 Configuring Firewall Access Rules

From within the Configuration → Firewall section of the Cisco ASDM:

Step 1. Select Access Rules.

Step 2. Click Add to add a new rule, or Edit to modify an existing rule.

Step 3. Configure a source IP (if needed). In the example shown in Figure 55, we are basing the rule only on the source tag, not on the source IP or AD user.

Step 4. Choose the source security group.

Step 5. Configure a destination IP (if needed). In the example shown in Figure 54, we are choosing the destination DataCenter.

Step 6. Choose a destination security group. In the example in Figure 54, notice that the destination Security Group is PCI.

Figure 54. Adding a Firewall Rule to Permit Traffic to PCI Servers

This rule has been created to allow any source with the SGT PCI to reach any server in the data center that also has the PCI tag. The next rule in the firewall policy would be to deny all other traffic to any server with the PCI tag, as seen in Figure 55.

Step 7. Click OK.

Step 8. Repeat steps 2 through 7 for the remainder of your rules.

Figure 55 shows a sample firewall policy.

Figure 55. Sample Firewall Policy

#	Enabled	Source Criteria: Source	User	Security Group	Service	Action	Destination Criteria: Destination	Security Group	Description
\multicolumn Global (7 rules)									
1	✓	any		PCI	ip	✔ Permit	DataCenter	PCI	This rule permits all traffic from a Source with a PCI Tag to ANY Server in the DataCenter with the PCI Tag.
2	✓	any		ANY	ip	✗ Deny	any	PCI	Deny anyone else from reaching a PCI Server
3	✓	any		HR	ip	✔ Permit	any	HR	
4	✓	any		ALL-Employee-Tags	http https	✔ Permit	any	HR	
5	✓	any		ALL-Employee-Tags	ip	✗ Deny	any	HR	
6	✓	any		ALL-Employee-Tags	ip	✔ Permit	DataCenter	ANY	Permit Employees to reach everything else in the DC
7		any			ip	✗ Deny	any		Implicit rule

Security Group Firewall on the Cisco ISR and ASR

The Cisco ASA is not the only SGFW available. Both the Cisco Integrated Services Routers Generation 2 (ISR G2) and the Cisco Aggregation Services Routers (ASR) have a very powerful ZBFW capability.

The Cisco ISR G2 (the 3900 Series, 2900 Series, and 2901, 1941, and 890 Routers) began support of SGFW as of Cisco IOS Software Release 15.2(2)T. The Cisco ASR 1000 Series Aggregation Services Routers added support of the SGFW as of Cisco IOS XE 3.4.

As of the time this guide is written, neither the Cisco ISR nor the Cisco ASR support the downloading of the environmental data from the Cisco ISE, or the download of IP-to-SGT bindings from the Cisco ISE. So they are not able to take advantage of any of that data today, and all tags must be created manually within the Cisco IOS CLI.

However, both routers are capable of running SXP, and the Cisco ASR has native tagging capabilities.

Configuring SGFW on the Cisco ASR and ISR

The configurations on the Cisco ASR and ISR are identical with one exception: The Cisco ASR is capable of using the SGT in both the source and destination of the ZBFW, while at this time the Cisco ISR is capable of using only the source SGT.

For this configuration example, we will keep things rather simple and create only a few security groups, and then add those groups to the ZBFW policy (Figure 56).

Note: With a Cisco IOS ZBFW, any traffic not explicitly permitted with a "pass" or "inspect" command will be dropped.

Figure 56.   Simple Branch ZBFW with SGFW

Procedure 1 Adding Security Groups to the Router

From the global configuration mode, each security group will need to be manually created. For the purposes of this example, we will be adding three of the security groups previously added to the Cisco ISE.

Step 1. Create a new security object group for each SGT (Examples 14 through 16).

Example 14 Adding the PCI Security Group

```
atw-asr1k(config)#object-group security PCI
atw-asr1k(config-security-group)#description SGT for PCI users and servers
atw-asr1k(config-security-group)#security-group tag 4
```

Example 15 Adding the NADs Security Group

```
atw-asr1k(config)#object-group security NADs
atw-asr1k(config-security-group)#description Network Devices
atw-asr1k(config-security-group)#security-group tag 2
```

Example 16 Adding the CommonServices Security Group

```
atw-asr1k(config)#object-group security CommonServices
atw-asr1k(config-security-group)#description Group for things like DNS
atw-asr1k(config-security-group)#security-group tag 3
```

Step 2. Add an inspection class map for the security groups.

Example 17 Creating a Class Map for the Source and Destination (SGT = PCI)

```
atw-asr1k(config)#class-map type inspect match-all pci-sgt
atw-asr1k(config-cmap)#match group-object security source PCI
atw-asr1k(config-cmap)#match group-object security destination PCI
```

Example 17 has configured a class map that looks for traffic with PCI source and destination tags (ssgt = PCI and dsgt = PCI). Only traffic matching both of those conditions will match this particular class map.

Example 18 Creating a Class Map Only for Destination SGT = PCI

```
atw-asr1k(config)#class-map type inspect match-all pci-sgt-dest
atw-asr1k(config-cmap)#match group-object security destination PCI
```

Example 18 has configured a class map that looks for traffic with a destination SGT of PCI (dsgt = pci) and is not looking for a source SGT at all.

Step 3. Create the policy map to inspect, pass, or drop traffic.

In this example we will allow only PCI-tagged traffic to reach PCI servers and will permit all other traffic.

Example 19 Creating the Policy Map

```
atw-asr1k(config)#policy-map type inspect branch-policy
atw-asr1k(config-pmap)#class type inspect pci-sgt
atw-asr1k(config-pmap-c)#inspect
atw-asr1k(config-pmap)#class type inspect pci-sgt-dest
atw-asr1k(config-pmap-c)#drop
atw-asr1k(config-pmap)#class class-default
atw-asr1k(config-pmap-c)#pass
```

The action of "inspect" in Example 19 enables deep packet inspection for all traffic with a source SGT of PCI traveling to a PCI server (ssgt = pci to dsgt = pci). The next inspection of "pci-sgt-dest" has an action of drop, to deny all other traffic destined to a server that has a PCI SGT. Lastly, the class default with an action of pass is allowing all other traffic to pass through the ZBFW.

Procedure 2 (Optional) Setting Up a ZBFW

This procedure is optional and would be configured only if you did not already have the security zones and service policy configured. In other words, you would not follow this procedure if your ZBFW were already configured and you were simply adding an SGFW to it.

Step 1. Create the security zones (Example 20).

Example 20 Creating the Security Zones

```
atw-asr1k(config)#zone security lan
atw-asr1k(config)#zone security servers
```

Step 2. Create the zone pair (Example 21).

Example 21 Creating the Zone Pair

```
atw-asr1k(config)#zone security lan-servers source lan destination servers
```

Step 3. Apply the service policy to the zone pair (Example 22).

Example 22 Applying the Service Policy to the Zone Pair

```
atw-asr1k(config-sec-zone-pair)#service-policy type inspect branch-policy
```

Step 4. Assign the interfaces to the correct zones (Example 23).

Example 23 Assigning Interfaces to Zones

```
atw-asr1k(config)#int g0/0/2
atw-asr1k(config-if)#zone-member security lan
atw-asr1k(config)#int g0/0/3
atw-asr1k(config-if)#zone-member security servers
```

Conclusion

In this guide you have learned all about Cisco TrustSec fundamentals, and at this point you should (hopefully) be able to articulate why it is so valuable, and how much operational expense it can save your organization.

You have learned that there are three foundational pillars of Cisco TrustSec technology: classification, propagation, and enforcement. Classification is the ability to accept the tag for a particular network authentication session. Propagation, or transport, is the ability to send that assigned tag to upstream neighbors through either native tagging or SXP. Enforcement may be on switches using SGACLs or on an SGFW.

Additionally, we have covered the basic configurations of all of these features across the many supported platforms.

For More Information

Reference http://www.cisco.com/go/trustsec

Printed in USA

C07-730151-00 01/14