Campus and Branch Office Segmentation with Cisco TrustSec

TrustSec version 5.3, November 2015

Contents

Introduction

Cisco TrustSec® technology segments wired, wireless, and VPN networks using security policies. Cisco TrustSec features are embedded in Cisco switching, routing, wireless LAN, and firewall products to protect assets, endpoints, and applications in enterprise and data center networks.

Cisco TrustSec controls improves upon traditional methods, which segment and protect assets using VLANs and access control lists (ACLs). Instead, the solution defines access entitlements by security group policies written in a plain language matrix (Figure 1). Users and assets with the same role classification are assigned to the same security group. These policies are decoupled from IP addresses and VLANs, so resources can be moved without re-engineering the network.

Cisco TrustSec policies are centrally created and automatically distributed to wired, wireless, and VPNs for enforcement. Users and assets thus receive consistent access and protection as they move in virtual and mobile networks. This consistency helps reduce the time needed for network engineering and compliance validation.

Figure 1. TrustSec Policy Management Matrix Example

Destination ▶ Source ▼	Employee	E-Mail	Finance	Internet
Employee	Deny	Permit	Deny	Permit
Executive	Deny	Deny	Permit	Permit
BYOD	Deny	Permit	Deny	Permit
Guest	Permit	Deny	Deny	Permit

For more detail on how Cisco TrustSec works, please refer to the "Quick Start Guide" on http://www.cisco.com/go/trustsec.

TrustSec, an alternative method to VLAN-based segmentation

Security Group Tags

Security Group Tags, or SGT as they are known, allow for the abstraction of a host's IP Address through the arbitrary assignment to a Closed User Group, represented by an arbitrarily defined SGT. These tags are centrally created, managed, and administered by the ISE. The Security Group Tag is a 16-bit value that is transmitted in the Cisco Meta Data field of a Layer 2 MACsec Frame as depicted below. Please note that Security Group Tags can also be carried on Ethernet without MACsec encryption

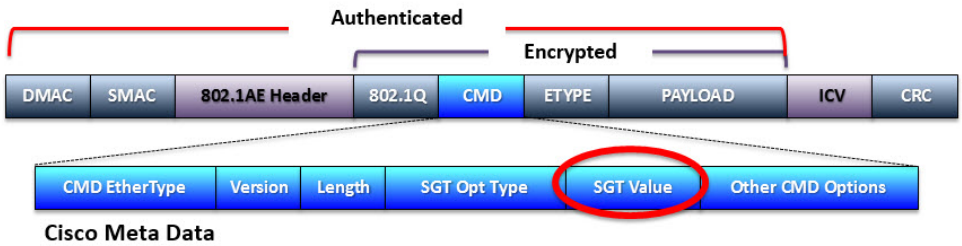

Figure 1. Layer 2 Frame with Cisco Meta Data field containing SGT

The Security Group Tags may be defined locally but are typically created at Cisco Identity Services Engine (ISE) and are represented by a user-defined Security Group name and a decimal value between 1 and 65,535 where 0 is reserved for "Unknown". Security Group Tags allow an organization to create policies based on a user's, device's, or server's role in the network providing a layer of abstraction in security policies based on an SGT as opposed to IP Addresses in ACLs.

TrustSec Fundamental Concepts

Cisco TrustSec technology consists of three fundamental processes known as Classification, Propagation, and Enforcement and can be defined as:

- Classification - The process of assigning the SGT is called Classification. A SGT can be assigned dynamically as the result of an ISE authorization or it can be assigned via static methods that map the SGT to some thing, like a VLAN, subnet, IP Address, or port-profile. Dynamic classification is typically used to assign SGT to users because users are mobile. They could be connected from any location via wireless, wired, or vpn. On the other hand, servers tend not to move, so typically static classification methods are used.
- Propagation – Is the means by which an SGT is either carried or advertised between networking devices. For those platforms that have purpose-built hardware to impose and remove the SGT in a CMD field, inline tagging can be used to carry the SGT embedded in the Ethernet header of a data frame on a hop-by-hop basis in the network. SXP, or Security group tag eXchange Protocol, is a lightweight TCP protocol that can be used to advertise the IP to SGT mapping learned by a device to other arbitrarily defined networking devices. The use of inline tagging or SXP is not mutually exclusive and both are in most cases used together.
- Enforcement – Is the process where a TrustSec or role-based policy which can be either locally defined or more commonly defined within Cisco ISE, is acted upon and traffic between a Source SGT and Destination SGT is either permitted or denied. Enforcement is carried out through the use of a Security Group ACL or SGACL on switches or through security policies on an ASA or IOS firewall commonly referenced as a Security Group Firewall or SG-FW.

Note: It is assumed that the reader is familiar with the basic concepts of TrustSec and has read the Quick Start Guide for TrustSec. This document can be found at {NEED URL FOR QUICK START}.

About This Document

The Cisco TrustSec solution provides simplified and scalable policy enforcement for network traffic between campus and branch-office users, controlling what is commonly referred to as “lateral movement” or “east to west” traffic. Initially Cisco ASR 1000 series and Cisco ISR Series routers supported SXP for tag propagation across the WAN. Today, this has been expanded to include various VPN methods which propagate tags in the data plane: IP Security (IPsec), Dynamic Multipoint VPN (DMVPN), or Group Encrypted Transport VPN (GET VPN) and the use of the EIGRP Over The Top protocol. This document will first discuss what should be considered when implementing Cisco TrustSec. Then common east to west traffic flow examples are used to illustrate how to design TrustSec through the classification, propagation, and enforcement phases.

In general there are three major east-to-west traffic flows:

- **Site to site**: Here we want to block a user at site A from communicating with a user at site B. To accomplish this, the classifications from sites A and B are sent by Security Group Exchange Protocol (SXP) to the headend Cisco Aggregation Services Router (ASR). The Cisco ASR then “reflects” the mappings back to the site-specific Cisco ISRs so that the local firewall can enforce policy.

- **Campus to branch office**: Users connecting to the campus from a branch-office network need to access resources in the campus and in other branch offices. In this case the branch-office access switch assigns the SGT using identity-based features. These SGTs may be used for local enforcement on a security group firewall (SGFW) module on the Cisco ISR or ASR. Additionally, the local Cisco ISR can propagate SGTs into the campus site over the WAN using various VPN methods: IP Security (IPsec), Dynamic Multipoint VPN (DMVPN), or Group Encrypted Transport VPN (GET VPN). Then by using the Then as traffic with the SGT embedded is forwarded across the WAN, the campus network switches may enforce traffic from the remote branch-office users.

- **Within the campus**: Within the same Layer 2 domain on the same switch, Cisco TrustSec policies may be used to prevent user-to-user communication. This is a unique method of controlling malware propagation. Targeted custom attacks quickly comprise the network by attacking one host and moving laterally from peer to peer. Therefore, restricting this type of movement is an important part of Cyber threat containment. A traditional network segmentation method like private vlans would accomplish this however they are not easily managed or modified to adapt to changes in the environment. Segmentation based on TrustSec group-based policies is an alternative which removes the need to architect the network to mitigate the scope of an attack.

Figure 2. Examples of Campus and Branch-Office Segmentation Traffic Flows

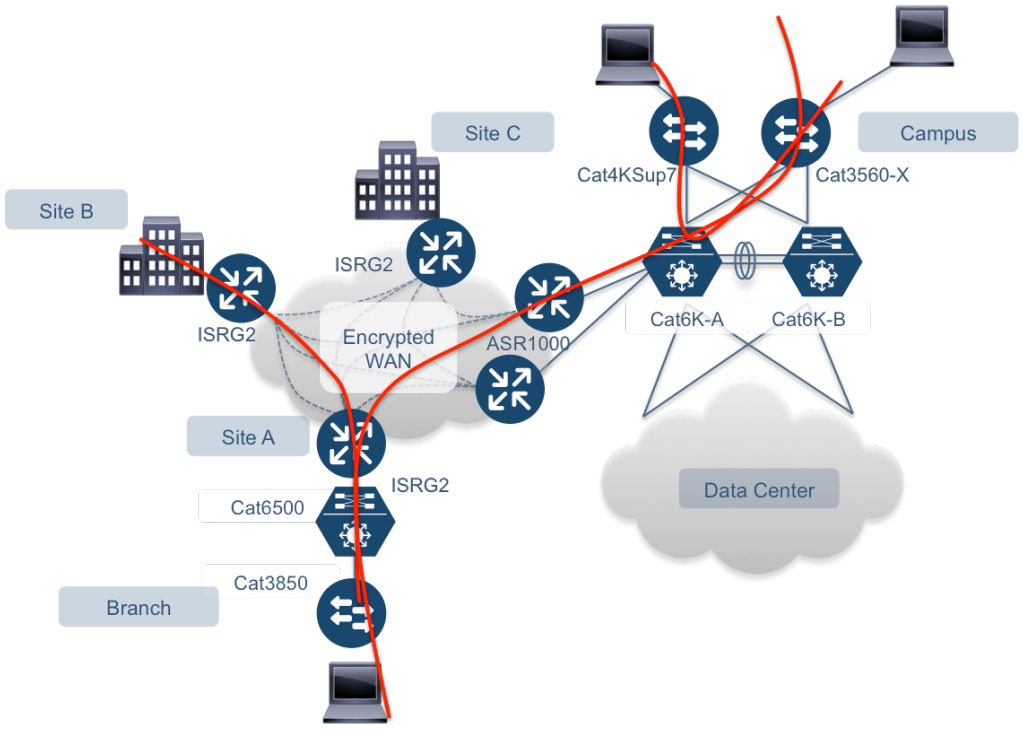

Each of these traffic flows illustrates what should be considered in the network design.

Design Considerations: Classification

In this phase, users (for example, engineers) and servers (for example, development servers) are placed into logical groups. These groups can be manually defined, or they can be predefined from Active Directory or Lightweight Directory Access Protocol (LDAP) servers. The groups are each represented by an SGT.

The SGT is a unique number that is used to represent the role (or group) of a user or server. Every SGT has an associated name (security group name) and value. For example, the employee role can have an arbitrarily assigned value of 101 and the security group name "employee." When Cisco TrustSec devices receive traffic tagged "SGT=101," filtering decisions are made based on policies defined for this tag.

SGTs can be centrally created, managed, and administered by the Cisco Identity Services Engine (ISE). Cisco switches, routers, and firewalls query the Cisco ISE periodically for these SGT-to-role mappings. After the SGT is created, the next step is to assign the SGT to a user or server.

Assigning the SGT at the Access Layer

At the access layer, dynamic classification is the best method of SGT assignment because SGT assignment occurs as the user enters the network. Dynamic classification starts with an authentication method such as IEEE 802.1X,

MAC Authentication Bypass (MAB), or Web authentication (WebAuth) to provide user-specific control. After authentication, the Cisco ISE evaluates the policy, classifies the user, and assigns an SGT that is associated with that classification. The tag is then downloaded to the access device, a Cisco switch or wireless LAN controller (WLC), to be associated with the user's IP and MAC address.

In environments where authentication isn't available, static classification methods are necessary. At the access layer, the recommended classification method is VLAN to SGT. In this case the SGT represents the classification of all of the devices within that VLAN.

Note: The capability to enforce access policies that are based on user identities is lost with VLAN-to-SGT classifications.

For networks with third-party devices or switches that do not support Cisco TrustSec functions, static methods like subnet to SGT or Layer 3 interface to SGT are recommended. These methods summarize traffic from a specific subnet or interface to a security group.

Platform-Specific Considerations

Cisco Catalyst 3560-X and 3750-X Series Switches:

- IP Device Tracking (IPDT) must be enabled before the switch can tag and filter traffic. As of 15.0(2)SE, IPDT is enabled by default when cts manual is configured on a switchport and when IEEE 802.1X, MAB or WebAuth authentication methods or VLAN-to-SGT features are used.
- Cat 3K-X can have Layer 2 adjacent hosts (small WLCs) trunked to Cat3K-X. This is useful in the case where the Cisco Wireless LAN Controller that does not have TrustSec support (pre 7.2 WLC code) You can assign a SGT to the trunked VLAN at the switch.

Note: Cisco TrustSec enforcement is supported on only eight or fewer VLANs on a VLAN-trunk link. If more than eight VLANs are configured on a VLAN-trunk link and Cisco TrustSec enforcement is enabled on those VLANs, the switch ports on those VLAN-trunk links will be error-disabled

- You cannot statically map an IP subnet to an SGT. You can map only IP addresses to an SGT. When you configure IP-address-to-SGT mappings, the IP address prefix must be 32.

Note: For additional details, see http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/appa_cat3k.html#wp1016377.

Cisco Catalyst 4500 Series Switches:

Please consult the Cisco TrustSec Switch Configuration Guide for notes about the Catalyst 4500 Series Switches (http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/appb_cat4k.html).

Design Considerations: Propagation

Inline Tagging Compared with SXP

Which SGT-propagation method is used depends on the platforms in the path. Not all devices are capable of inline SGT, but some devices support both inline tagging and SXP. Inline SGT is better from an operational perspective. Inline SGT occurs within the data plane, so there is no impact on performance. SXP is a control plane function so CPU and memory impact performance.

SXP Scalability

When working with a SXP design it is always important to consider two key factors: the number of SXP peers and the number of IP-SGT mappings supported. The table below lists the scaling numbers for the platforms in this document. For information on other platforms, please refer to the SXP scalability chart on the Cisco TrustSec home page (http://www.cisco.com/go/trustsec).

Table 1. SXP Scaling Numbers

Platform	Max SXP Peers	MAX IP-SGT Bindings
Catalyst 4500 Sup 7E	1000	256,000
Catalyst 3850/WLC 5760	128	12,000
Cisco ASR 1000 Series Routers	1800-unidirectional 900-bidirectional	180,000 750,000 with IOS XE-3.15
Cisco ISR 2900,3900	250-unidirectional 125-bidirectional	180K for unidirectional SXP connections 125K for bidirectional SXP connections
Cisco ISR 4400	1800-unidirectional 900-bidirectional	135,000

Platform-Specific Considerations

Cisco Catalyst 3560-X and 3750-X Switches:

- These switches can be SXP listeners for Layer 2-adjacent traffic only. They cannot be listeners for peers sending aggregated IP-to-SGT bindings, and they cannot take IP-to-SGT bindings from multi-hop SXP connections.
- Additional information is documented here: Considerations for Catalyst 3000 and 2000 Series Switches and Wireless LAN Controller 5700 Series

Cisco Wireless Controllers:

- To have the Cisco Wireless LAN Controller peer with the Cisco Nexus 7000 Series Switch, the Cisco WLC Release 7.4 or later is required.
- To have the Cisco Wireless LAN Controller peer with a Cisco ASA, Cisco WLC Release 7.4 or later is required.
- Cisco wireless access points do not support SXP currently. Therefore, if you are using a Cisco FlexConnect™ solution where the data traffic is switched locally, the local switch must use VLAN-to-SGT mapping for classification.

Cisco Integrated Services Routers Generation 2 and Cisco ASR 1000 Series Routers

- Reference the following links for information on support for bidirectional SXP connections defined with a single pair of IP addresses:
  Cisco TrustSec Configuration Guide, Cisco IOS XE Release 3E
  Cisco TrustSec Configuration Guide, Cisco IOS XE Release 3S (Cisco ASR 1000)

Design Considerations: Enforcement

A general guideline for enforcing policies is to use the device closest to the resources that are being protected. However, in some cases, the closest enforcement device may not be the best choice because of the way the device learned the SGTs, because of device-specific limitations, or because of compliance policies. The traffic flow design section will outline these limitations and list some useful features that may influence where enforcement should occur.

Unknown SGT (SGT=0)

It is unusual to have all the users and servers mapped to an SGT so switches must accommodate for this case. When Catalyst switches receive packets that do not contain a SGT, these switches tag the packet with a "SGT=0", the "Unknown" tag. You can then write a security policy for SGT=0. When Nexus switches receive packets that do not contain a SGT, the Nexus switches will apply the SGT value assigned to the ingress interface and forward the frame with the new value.

Unlike ACLs with an implicit deny at the end, Security Group ACLs (SGACLs) implemented on a switching platform have an implicit permit to Unknown or an implicit permit to all. This policy is not enforced on the Cisco ASA firewall or the Cisco IOS zone-based firewall acting as an SGFW, where an implicit deny is still maintained. On a switch, if no specific tag value is assigned to a server, the destination is considered Unknown and the packet is forwarded by default.

A common error is to create a rule to deny the IP address of the Unknown tag. This, however, means that every packet with an Unknown destination tag will be dropped. It is best to omit a policy for SGT=0 until classifications are fully understood.

Enforcement Priority

If a switch receives SGT mapping information from two classification methods, enforcement is based on the following order of precedence, from lowest (1) to highest (7):

1. **VLAN**: Bindings learned from snooped Address Resolution Protocol (ARP) packets on a VLAN that has VLAN- to-SGT mapping configured.
2. **Command-line interface (CLI)**: Address bindings configured using the IP-to-SGT form of the Cisco TrustSec role-based SGT-map global configuration command.
3. **Layer 3 Interface (L3IF)**: Bindings added using Forwarding Information Base (FIB) entries that have paths through one or more interfaces with consistent L3IF-to-SGT mapping or with identity port mapping on routed ports.
4. **SXP**: Bindings learned from SXP peers.
5. **IPARP**: Bindings learned when tagged ARP packets are received on a Cisco TrustSec capable link.
6. **Local**: Bindings of authenticated hosts that are learned by means of Cisco Identity Services Engine (ISE) and device tracking. This type of binding also includes individual hosts that are learned by means of ARP snooping on Layer 2 ports that are configured for port mirroring.
7. **Internal**: Bindings between locally configured IP addresses and the device's own SGT.

Platform-Specific Considerations

Cisco Integrated Services Routers Generation 2 and Cisco ASR 1000 Series Routers

- Cisco ASR 1000 Series routers: SGTs can be used for the source and destination values in a firewall rule.
- Cisco ISR G2: SGTs can be used for the source values only in a firewall rule.

Cisco Catalyst 3560-X and 3750-X Series Switches:

- For the Catalyst 3750X the maximum number of ACEs in an SGACL is sixty. Upon exceeding that number, the TCAM will not be programmed and the default ACL will remain in effect. On the 3750X platform, TCAM is utilized on the respective switch to which an end device is attached and is not limited to space on the stack master.

Cisco Catalyst 3650 and 3850 Series Switches:

- The 3850 has a shared TCAM design resulting in ~1384k entries available for Layer 3 ACE or 692k for Layer 4 ACE (since these take up 2 spaces per entry)
- The 3850 can apply SGACLs to a maximum of 255 Security Groups

Traffic Flow Design

This section examines each flow in Figure 1 to determine how classification is performed, how the SGT is propagated, and what device provides the enforcement.

Site-to-Site Segmentation

In site-to-site segmentation, the source classification is done at Site A and the SGT is propagated to a Cisco ISR using inline tagging or SXP depending upon what the distribution layer device supports (Figure 3). The Cisco ISR at site A can perform SGFW enforcement with SGTs for source and destination so that traffic destined to Site B does not need to traverse the WAN. In order for Site A Cisco ISR to provide this enforcement, the SGTs from Site B must be propagated to Site A via SXP or by some VPN method.

Please note that if there are multiple sites and if SXP is necessary, the Cisco ISR and ASR both support SXPv4. SXPv4 supports loop detection. Therefore, SXP peering can be done just between the Cisco ISRs and ASRs. You do not have to create an SXP mesh between ISRs.

Table 3 shows the classification, SGT propagation, and enforcement for each device in site-to-site segmentation.

Figure 3. Site to Site Segmentation

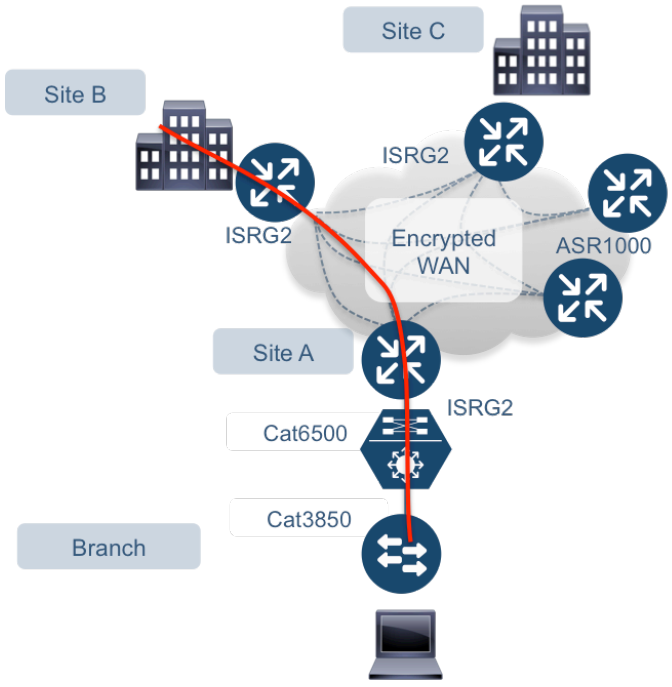

Table 3. Classification, SGT Propagation, and Enforcement Methods for Site-to-Site Segmentation

Classification

Device	SGT Classification	Notes
Access Layer		
Cisco Catalyst 3850Series Switches	dynamic	

Propagation

Device	SGT Propagation	Notes
From Site A		
Cisco Catalyst 6500 Series with SUP2T to site A Cisco ISR G2	inline	
Cisco Catalyst 6500 Series with SUP720 to site A Cisco ISR G2	SXP	

Enforcement

Device	SGT Enforcement	Notes
Site A ISRG2	SGFW	Zone based firewall rules can only use SGTs as the source of the traffic flow. Cisco ISR 4000 or ASR1000 is required to provide SGT based source and destination rules.

Campus to Branch Office

In campus-to-branch-office segmentation, enforcement is possible at the core or at the branch office (Figure 4). The device enforcing the policy must know the DGT, destination group tag. The user is dynamically classified on the Cisco Catalyst 3560-X Series Switch. This classification, resulting in a dynamic IP-SGT mapping at the access switch is then propagated to the Cisco Catalyst 6500 Series Supervisor 2T using inline tagging. The Cisco Catalyst 6500 can propagate the SGT to the Cisco ASR using inline tagging. The Cisco ASR may perform enforcement based on SGFW rules, or it can propagate the tags using SGTover DMVPN, SGT over GETVPN, or SXP to the Cisco Catalyst 6500 in the destination branch office for enforcement. Table 4 shows the classification, SGT propagation, and enforcement method for each device.

Figure 4. Campus-to-Branch-Office Segmentation

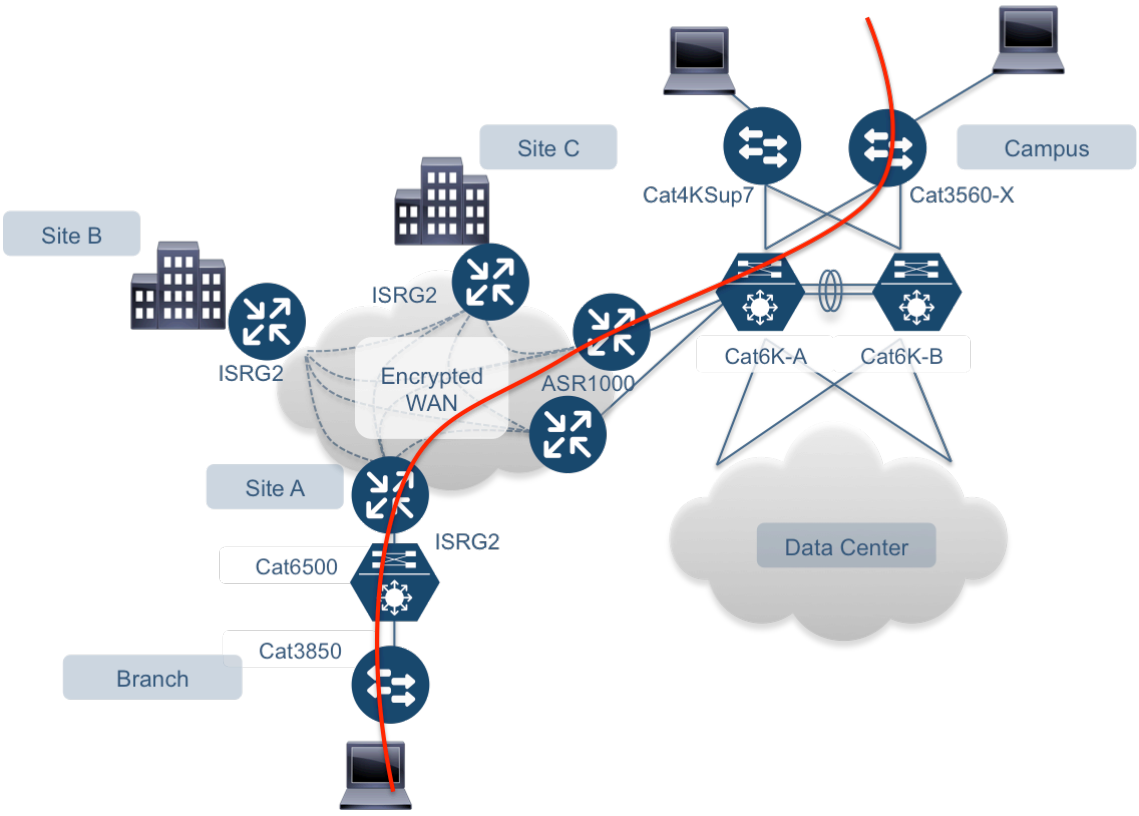

Table 4. Classification, SGT Propagation, and Enforcement Methods for Campus-to-Branch-Office Segmentation

Classification

Device	SGT Classification	Notes
Campus Access Layer		
Cisco Catalyst 3500-X Series Switches	Dynamic	
Branch-Office Access Layer		
Cisco Catalyst 3850 Series Switch	Dynamic	

Propagation

Device	SGT Propagation	Notes
From the Campus Access Layer Toward the Branch Office		
Cisco Catalyst 3560-X to Catalyst 6500 Series Switches with SUP2T to Cisco ASR **Cisco Catalyst 3560-X to Catalyst 6500 Series Switches with SUP720 to Cisco ASR**	Inline SXP	
From the Branch-Office Access LayerToward the Campus		
Cisco Catalyst 3850 to Cisco Catalyst 6500 Series Switches with SUP 2T to Cisco ISR	Inline	

Enforcement

Device	SGT Enforcement	Notes
Cisco ASR	SGFW enforcement	If the ASR is to provide enforcement of site to site communication, the ASR must have the IP to SGT mappings for the devices at the sites that the ASR is providing enforcement for. For example, if the ASR is to enforce traffic between site A and Site b, the ASR must have the IP-SGT mappings from Site B.
Cisco Catalyst 6500 Series Switches	SGACL	Enforcement provided here if traffic originates from Site A.

Within the Campus

In intra-campus segmentation you have two users connected to the Layer 2 switch (Figure 5). These users can be connected to the same VLAN or to different ones. You can use SGTs to tag each user and use SGACLs, rather than ACLs, to enforce traffic between them (Figure 5). Table 5 shows the classification, SGT propagation, and enforcement method for each device.

Figure 5. Example topology with inline switches

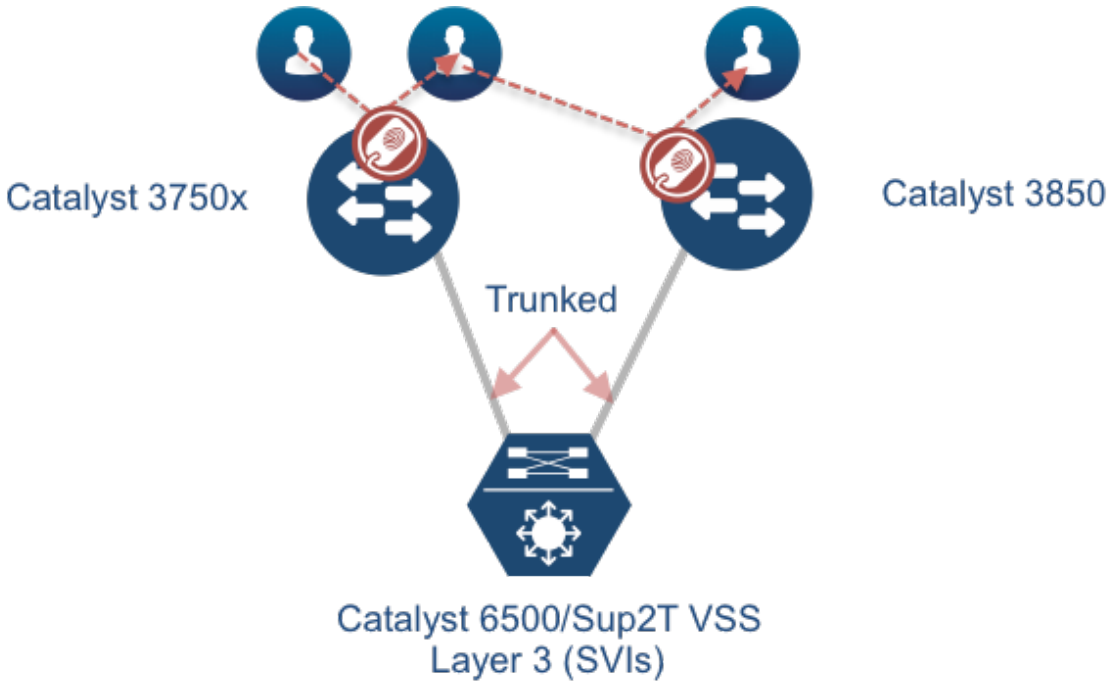

Table 5. Classification, SGT Propagation, and Enforcement Methods for Intra-Campus Segmentation

These traffic flow examples summarize two major East-West Segmentation tests that performed with code supported in the TrustSec 5.3 release.

Use Case	Campus Topology Details	SGT Classification (at Ingress)	SGT Propagation	SGT Enforcement	Solution Tested Devices
East-West Segmentation with Inline Tagging	**Access layer switches trunked to distribution switch. All VLANs are terminated and switched (SVI) on the distribution switch which is inline capable.**	Any	Inline tagging end to end	At the destination access layer	8 stack Cisco Catalyst 3750-X 8 stack Cisco Catalyst 3850 Cisco Catalyst Cat6500-Sup2T/VSS
East-West Segmentation with SXP	**Access layer switches trunked to distribution switch. All VLANs are terminated and switched (SVI) on the distribution switch which is not inline capable. Propagation is achieved by using bi-directional SXP.**	Any	Bidirectional SXP between access and distribution	At the access switch where the source is located if bidir SXP is used or at Dist if one-way to Distribution Switch is used	Cisco Catalyst 3850 (v. 3.6.3E) Cisco Catalyst 4500/Sup8E (v. 3.6.3.E) Cisco ASR 1000 Series (v. 3.15.0)

East-West Segmentation with Inline Tagging

In this example, all switches in the path are capable of inline tagging and enforcement. Therefore intra-VLAN and inter-VLAN communication can be prevented within the access layer.

Figure 6. Example topology with inline capable switches

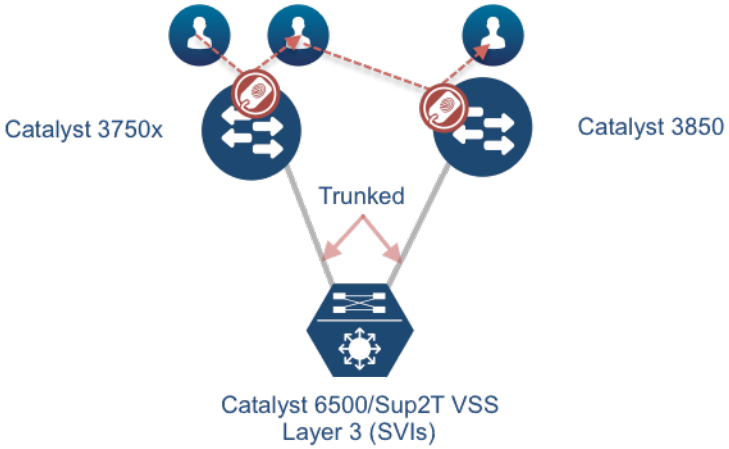

East-West Segmentation with SXP

Note: Refer to Table 1 for SXP scaling numbers. This flow requires bi-directional SXP communication so keep in mind that both the access switches SEND AND LEARN mappings to/from the 6500 (SXP aggregator). Therefore, if the deployment will exceed the scaling numbers, upgrading to a SUP2T is a must.

This use case is common to environments where the customer desires to implement TrustSec segmentation within the same switch and between switches at another campus location. Since the goal is to segment at the access layer, new access switches that support inline tagging and TrustSec group-based enforcement have been added but the existing upstream infrastructure remains untouched.

Figure 7. Example topology mixed with Inline and non-inline capable switches

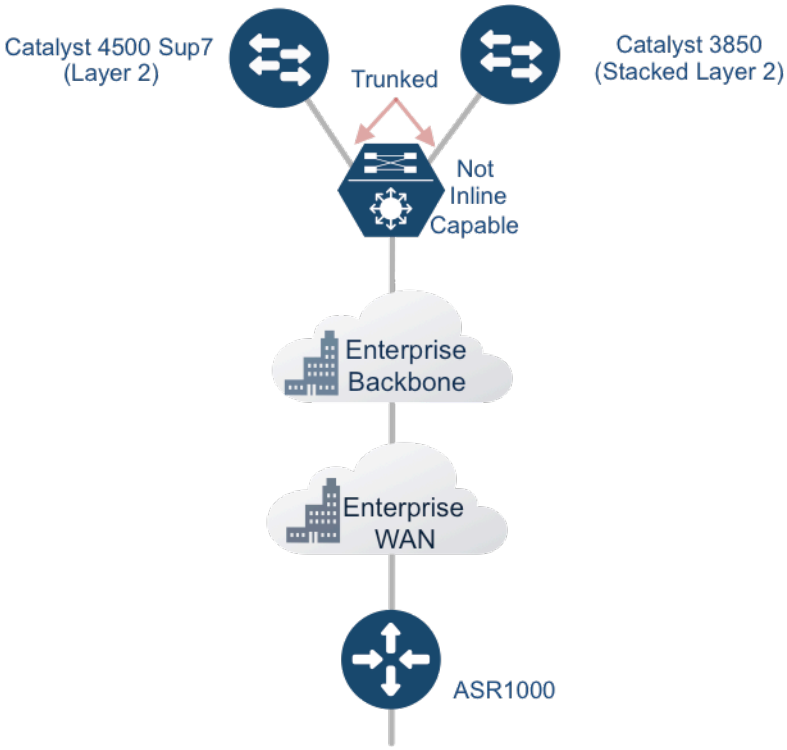

To achieve inter-vlan, switch to switch, enforcement with the above topology, each switch must be aware of the IP to SGT mappings of hosts on the adjacent switches. As stated previously, these IP/SGT

mappings must be propagated via SXPv4. SXPv4, which supports loop detection, is necessary because the entire list of mappings are communicated with other switches which can result in a switch receiving mappings that it was the originator of.

In situations where the number of SXP peers and/or IP/SGT mappings may exceed these limits, introducing another device, such as the ASR 1000 in this example, can serve as a SXP listener, aggregator, and speaker. In other words, the ASR will:

1. Receive (listen) for the IP-SGT mappings from all access layer switches
2. Aggregate all of the IP-SGT mappings into one single table
3. Send (reflect) the list of IP-SGT mapping back into the access layer for enforcement

The diagrams below illustrate this flow.

Figure 8. Access layer switches sending IP-SGT mappings to SXP receiver for aggregation

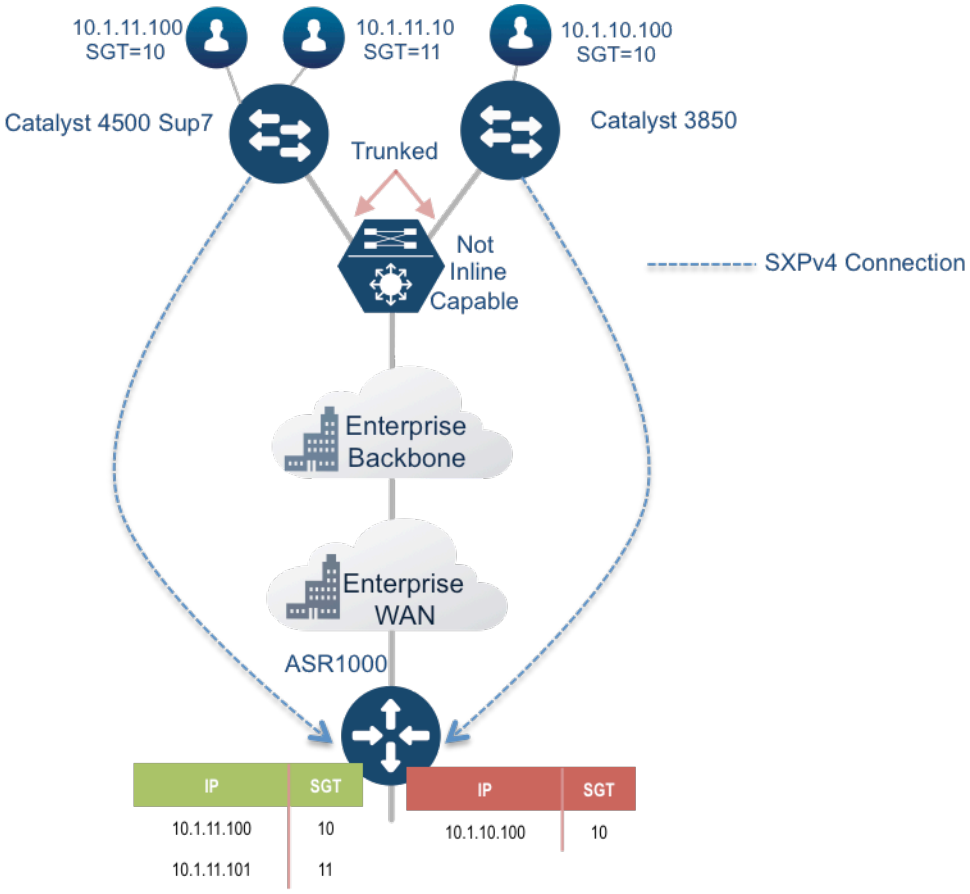

Figure 9. Aggregated IP-SGT mappings reflected back into the access layer for inter-vlan traffic enforcement

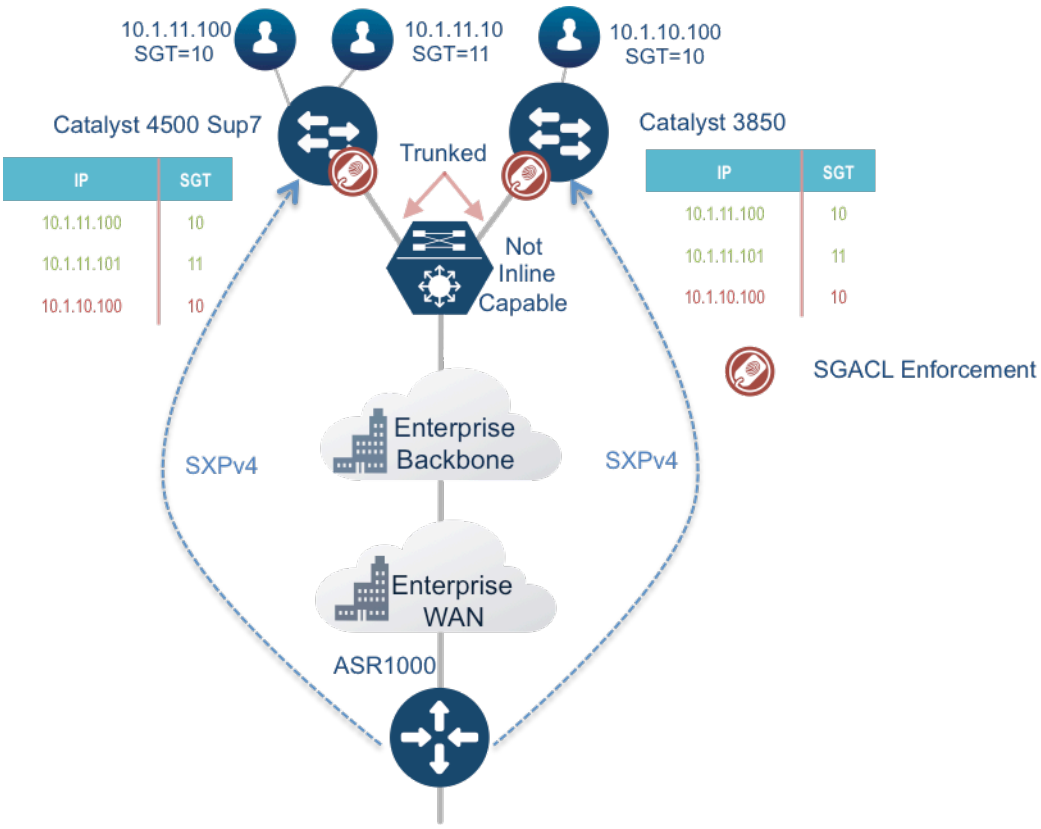

For More Information

Please reference http://www.cisco.com/go/trustsec. The “Cisco TrustSec Quick Start” guide is recommended for reading prior to the How To Guides.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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