



## Segment the Network into Smaller Trust Zones

- [Segment the Network into Smaller Trust Zones, on page 1](#)
- [Segmentation Technologies, on page 2](#)
- [Cisco Identity Services Engine, on page 5](#)
- [ISE/SGT Design Considerations, on page 12](#)

## Segment the Network into Smaller Trust Zones

The main goal for segmentation is to minimize the impact of any potential breach. Part one of the security journey provided segmentation between the enterprise and industrial network. However, the risk of breach remains. Malware could be introduced to the network using rogue USBs, or infected devices connecting to plant floor infrastructure. This step of the journey provides guidance to further segment the network into smaller trust zones, so if an adversary does breach the network boundary, their effectiveness can be reduced and contained.

In order to improve interconnection and compatibility between industrial systems, equipment manufacturers are increasingly using standard communication protocols and complying with the requirements of international standards organizations. This is the role of the International Society of Automation (ISA), the International Organization for Standardization (ISO), and the International Electrotechnical Commission (IEC).

ISA/IEC 62443 defines a set of principles to be followed in Industrial environments:

- **Least Privilege:** to give users/devices only the rights they need to perform their work, to prevent unwanted access to data or programs and to block or slow an attack if an account is compromised
- **Defense in Depth:** multiple layered defense techniques to delay or prevent a cyberattack in the industrial network
- **Risk Analysis:** address risk related to production infrastructure, production capacity (downtime), impact on people (injury, death), and the environment (pollution)

Based on these principles, ISA/IEC 62443 recommends segmenting the functional levels of an industrial network into zones and conduits.

A **zone** is a collection of physically and functionally united assets that have similar security requirements. These areas are defined from the physical and functional models of the industrial system control architecture. Some characteristics of a security zone are:

- A zone should have a clear border

- A zone can have other subzones
- The border is used to define access with another zone or outside system
- Access is via electronic communication channels or the physical movement of people or equipment

A **conduit** supports the communication between zones. A conduit supports and defines allowed communication between two or more zones. Some attributes defined within a conduit are:

- The zones interconnected by the conduit
- Type of dataflows allowed
- Security policies and procedures

Partitioning the industrial network into zones and conduits reduces overall security risk by limiting the scope of a successful cyber-attack.

## Segmentation Technologies

### VLAN

A virtual local area network (VLAN) can be created on a Layer 2 switch to reduce the size of broadcast domains. Devices within a VLAN act as if they are in their own independent network, even if they share a common physical infrastructure with other VLANs. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations belonging to the VLAN the packets were sourced from. Each VLAN is considered a separate logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a device that supports routing.

The default Ethernet VLAN is VLAN 1. It is a security best practice to configure all the ports on all switches to be associated with VLANs other than VLAN 1. It is also a good practice to shut down unused switch ports to prevent unauthorized access.

A good security practice is to separate management and user data traffic. The management VLAN, which is VLAN 1 by default, should be changed to a separate, distinct VLAN. To communicate remotely with a Cisco switch for management purposes, the switch must have an IP address configured on the management VLAN. Users in other VLANs would not be able to establish remote access sessions to the switch unless they were routed into the management VLAN, providing an additional layer of security. Also, the switch should be configured to accept only encrypted SSH sessions for remote management.

### VRF-lite

While virtualization in the Layer 2 domain is done using VLANs, a mechanism is required that allows the extension of the logical isolation over the routed portion of the network. Virtualization of a Layer 3 device can be achieved using virtual routing and forwarding lite (VRF-Lite). The use of virtual routing and forwarding (VRF) technology allows you to virtualize a network device from a Layer 3 standpoint, creating different "virtual routers" in the same physical device. A VRF instance consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table.

Technically, there is no difference between a VRF and a VRF-lite. The difference lies in how you use it. VRF is a technology, while VRF-lite is a particular way of using that technology. Both VRF and VRF-lite are built on the same premise: they have separate routing tables (that is, VRFs) created on your router and unique interfaces associated with them. If you remain here, you have VRF-lite. If you couple VRFs with a technology

such as MPLS to communicate with other routers having similar VRFs while allowing to carry all traffic via a single interface and being able to tell the packets apart, you have a full VRF.

To provide continuous virtualization across the Layer 2 and Layer 3 portions of the network, the VRFs must also be mapped to the appropriate VLANs at the edge of the network. The mapping of VLANs to VRFs is as simple as placing the corresponding VLAN interface at the distribution switch into the appropriate VRF.

*Note: For this design guide, VRFs are not utilized, and no design guidance is provided.*

### **Access Control List**

An Access Control List (ACL) is a series of statements that are primarily used for network traffic filtering. When network traffic is processed by an ACL, the device compares packet header information against matching criteria. IP packet filtering can be based only on information found in Open Systems Interconnection (OSI) Layer 3 header or on both Layer 3 and Layer 4 header information. A device extracts the relevant information from the packet headers and compares the information to matching permit or deny rules.

Traffic that enters a routed interface is routed solely based on information within a routing table. However, when an ACL is applied to an interface, the router performs the additional task of evaluating all network packets against the ACL as they pass through the interface to determine if the packet can be forwarded. ACLs can allow one host to access a part of the network and prevent another host from accessing the same part.

### **Stateful Firewall**

A firewall is a network security device that monitors the incoming and outgoing network traffic and decides whether to allow or block the traffic based on a defined set of security rules. Where a stateless packet filter, such as a standard Access Control List (ACL), operates purely on a packet-by-packet basis, a stateful firewall allows or blocks traffic based on the connection state, port, and protocol. Stateful firewalls inspect all activity from the opening of a connection until the connection is closed.

Stateful packet filters are application-aware while additional deeper inspection of transit traffic is being performed, which is required to manage dynamic applications. Dynamic applications typically open an initial connection on a well-known port and then negotiate additional OSI Layer 4 connections through the initial session. Stateful packet filters support these dynamic applications by analyzing the contents of the initial session and parsing the application protocol just enough to learn about the additional negotiated channels. A stateful packet filter typically assumes that if the initial connection was permitted, any additional transport layer connections of that application should also be permitted.

Next-Generation Firewalls (NGFW) are stateful firewalls with additional features such as application visibility and control, advanced malware protection, URL filtering, Secure Sockets Layer (SSL)/Transport Layer Security (TLS) decryption, and IDS/IPS.

Choosing to use a NGFW or ACLs in the OT network will depend on the types of communication that will flow through the network. Device to device communication for example, may use protocols such as Ethernet/IP (TCP port 44818 & UDP port 2222) or Modbus (TCP port 502) which can be filtered on a packet-by-packet basis due to its static network behavior. This is the communication that keeps the plant running, and doing more advanced network inspection between these devices, or implementing an IPS system, may introduce system latency and/or run the risk of OT downtime due to false positives.

It is therefore recommended to introduce NGFW in the network for northbound communication, such as between the IDC and the Cell/Area Zones for advanced threat protection between devices that pose a higher security threat but would not cause production downtime if security was prioritized over connectivity. Having an additional layer of IPS between the IDC and the production floor will ensure advanced threat protection exists not just in the IDMZ. An NGFW could also be deployed for advanced application control such as allowing read-only access to an asset on the plant floor from a vendor application hosted in your IDC.

### **TrustSec**

Cisco TrustSec (CTS) defines policies using logical device groupings known as Security Group Tag (SGTs). An SGT is a 16-bit identifier embedded into the MAC layer of IP traffic. The SGT is a single label indicating the privileges of the group within the entire network. It is in turn propagated between network hops allowing any intermediary devices (switches, routers) to enforce policies based on the group identity tag. The features associated with SGTs on the network devices can be divided into three categories: classification, propagation, and enforcement.

**Classification** is the assignment of SGTs to an IP address. This assignment can be accomplished either dynamically or statically. Generally, dynamic classification is done at the access layer and static classification is done at the egress switch. In OT networks, where devices tend not to have 802.1X capabilities, dynamic classification can be done using MAC Authentication Bypass (MAB). Static classification is configured directly on the switch in which tagging occurs. Options for static classification include the mapping of Subnet, IP address, VLAN, or port to an SGT.

The **transport** of security group mappings can be accomplished through inline tagging or the SGT Exchange Protocol (SXP). With inline tagging, the SGT is embedded in the Ethernet frame header. However, not all network devices support inline tagging. SXP is used to transport SGT mappings across devices that do not support inline tagging.

**Enforcement** is implementing a permit or deny policy based on the source and destination SGTs. This implementation can be accomplished with security group access control lists (SGACLs) on switching platforms and security group firewall (SGFW) on routing and firewall platforms.

*Note: Which method of classification, transport and enforcement to use will be discussed later in the documentation. This section only introduces the technology.*

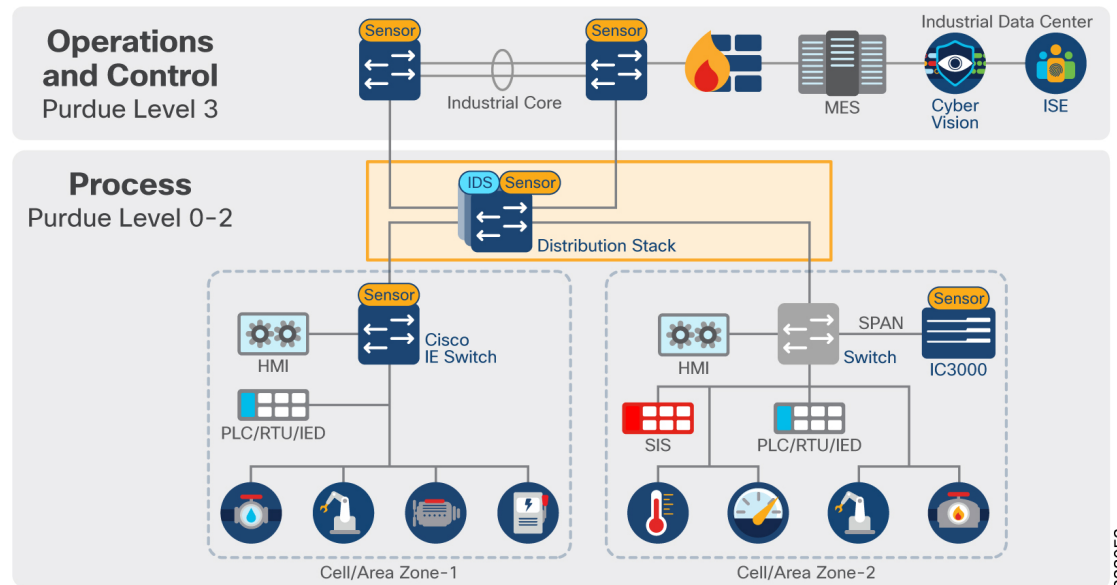
### How to get started with Segmentation

*Note: It is assumed at this point in the document that the first step of the recommended security journey has been completed, and there is segmentation between the Enterprise network and the OT network with the implementation of an IDMZ. The following section provides design guidance for implementing segmentation within the OT environment.*

### Macro-Segmentation

Networks are usually designed in modular fashion where the overall network infrastructure is divided into functional modules, each one representing a distinctive place in the network. Cell/Area zones offer organizations a starting point for segmentation of the control network. If following recommended architecture designs, organizations will make use of a distribution switch stack to transport data to and from different cell/area zones in the network. While organizations are gaining visibility using Cyber Vision and understanding the normal operating state of their networks, policy can be applied to larger functional zones based on subnet, VLAN or other network-based information. This segmentation model is known as macrosegmentation. For example, endpoints in the fabrication shop zone probably require no communication with endpoints in the welding shop zone and can be distinctly identified by the network infrastructure they are physically connected to.

Figure 1: Policy Enforcement Across the Distribution Switch



The layer 3 boundary and gateway for devices is in the distribution switch as shown in the following figure. It is recommended that security is first created at this layer of the network, to allow and deny communications for inter-cell/area zone communication such as controller-to-controller communication across zones or controller-to-site operations zone.

### Micro-Segmentation

For OT environments, micro-segmentation can be thought of as the segmentation within a VLAN segment. Traditionally, private VLANs were used to divide a VLAN into subdomains. This becomes complex and difficult to deploy and manage, so would not be a recommended approach to micro-segmentation.

Cisco TrustSec is a logical grouping framework, and while we recommend its use in macrosegmentation to help define policy between traditional networking boundaries, it can also be decoupled from IP addresses and VLANs. Using a Cisco TrustSec role or SGT as the means to describe permissions on the network allows the interaction of different systems to be determined by comparing SGT values. This avoids the need for additional VLAN provisioning, keeping the access network design simple and avoiding VLAN proliferation and configuration tasks as the number of roles grows. TrustSec SGACLs can also block unwanted traffic between devices of the same role, so that malicious reconnaissance activities and even remote exploitation from malware can be effectively prevented.

While micro-segmentation can be an effective tool for segmenting the OT network, it is a complicated starting point and requires a deep understanding of the OT network. The recommendation is to begin with macro-segmentation across the distribution network and then slowly augment micro-segmentation policies after effective visibility has been gained of the plant floor operations. This will ultimately lead to a hybrid approach, where both macro and micro-segmentation will be implemented using the same TrustSec technology.

## Cisco Identity Services Engine

Cisco Identity Services Engine (ISE) utilizes TrustSec technology to logically segment control system networks. Cisco TrustSec classification and policy enforcement functions are embedded in Cisco switching, routing, wireless LAN, and firewall products.

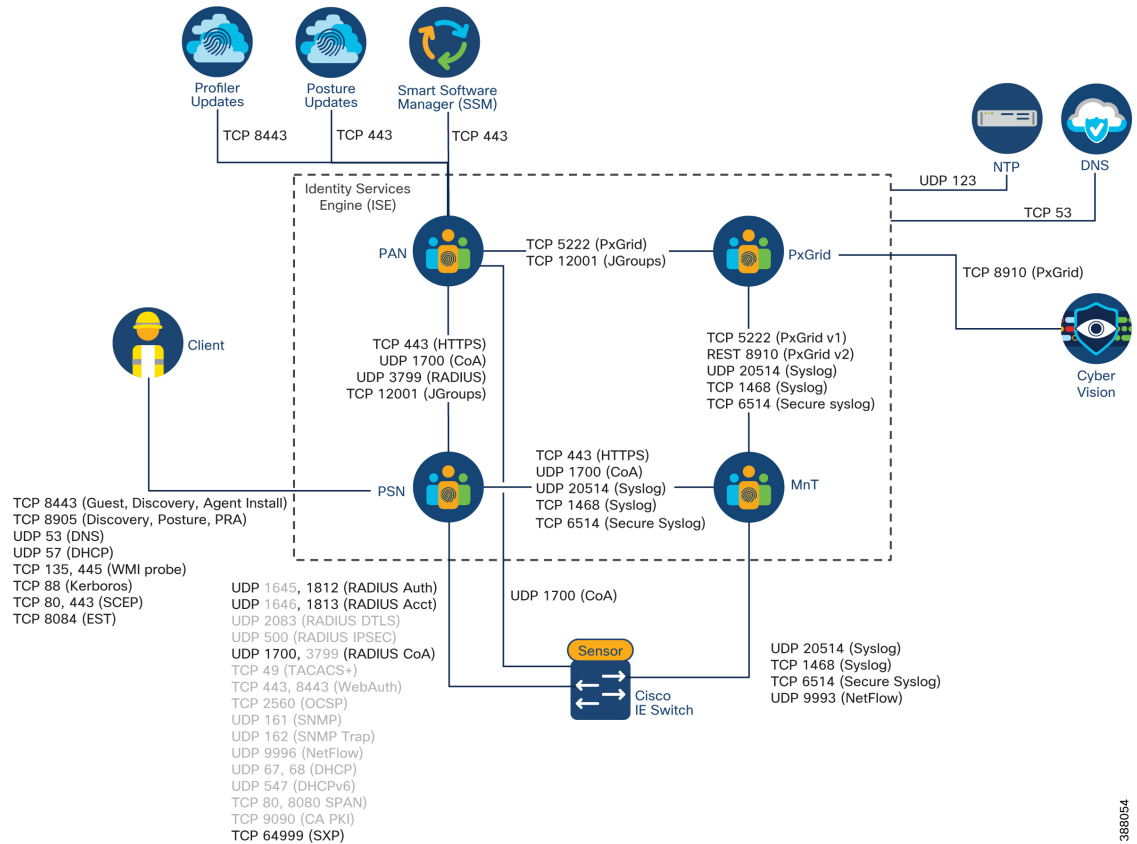
ISE Components / Personas Cisco ISE has four distinct personas/nodes that can either be deployed in one standalone deployment (all personas residing in a single ISE node) or distributed across the network. The personas available in ISE are:

- **Policy Administration Node (PAN):** allows you to perform all administrative operations and configurations on Cisco ISE. It serves as a single pane of glass for viewing all administrative operations, configurations, and contextual data. It synchronizes the configuration to the rest of the nodes in the deployment
- **Policy Service Node (PSN):** provides network access, posture, guest access, client provisioning, and profiling services. This persona evaluates the policies and makes all the decisions
- **Monitoring node (MnT):** stores log messages from all the PANs and PSNs in a network. This persona provides advanced monitoring and troubleshooting tools that you can use to effectively manage the network and resources
- **pxGrid node:** a framework to exchange information between ISE and other Cisco platforms or ecosystem partner systems

Cisco ISE can be deployed as a hardware appliance, virtual appliance, or on public cloud platforms like Amazon Web Services (AWS), Azure Cloud, and Oracle Cloud Infrastructure (OCI). ISE provides a Performance and Scalability Guide to provide sizing guidelines. As an example, a small ISE deployment could be deployed with all personas existing on the same appliance, however, a large deployment recommends that all ISE personas be fully distributed in the network and can support up to 50 PSNs.

For the validation testing within this guide, ISE was distributed, with the PSN and pxGrid node each having their own dedicated instance in the Industrial Zone. The following figure depicts the communication flows required by ISE Cisco ISE.

Figure 2: ISE Communication Flows



*Note: Not all flows depicted in this diagram are used in this design guide. An example is demonstrated in the flow between the ISE PSN and the Cisco switching infrastructure, with greyed out values indicating "not in use". All ports are shown to provide clarity when ISE is portrayed for use beyond this guide.*

### ISE Authentication Policies

Authentication provides a way to identify a user, typically by having the user enter a valid username and password before access is granted. However, most devices in the network are not interactive and therefore do not have the capability to provide a username or password. ISE provides the capability to do MAC Authentication Bypass (MAB), which uses the MAC address of a device to determine the level of network access to provide. Before MAB authentication, the identity of the endpoint is unknown, and all traffic is blocked. The switch examines a single packet to learn and authenticate the source MAC address. After MAB succeeds, the identity of the endpoint is known and traffic from that endpoint is allowed. The switch performs source MAC address filtering to help ensure that only the MAB-authenticated endpoint is allowed to send traffic.

### ISE Authorization Policies

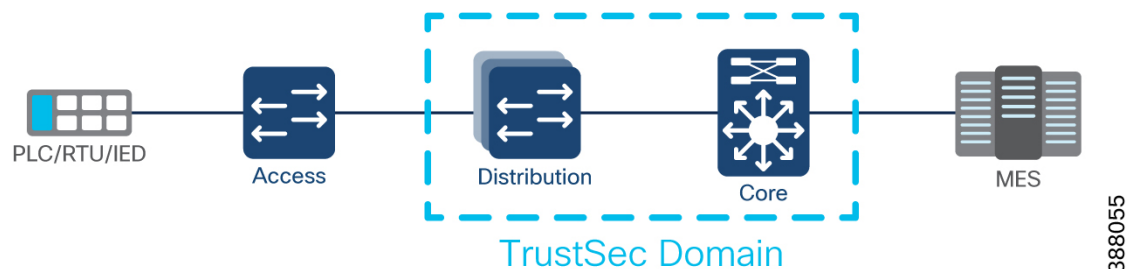
Authorization is the process of enforcing policies and determining what type of activities, resources, or services a user or device is permitted to access. All controlled from a central location, Cisco ISE distributes enforcement policies across the entire network infrastructure. Administrators can centrally define a policy that differentiates vendors from registered users and grant access based on least privilege. ISE provides a range of access control options, such as downloadable Access Control Lists (dACLs), VLAN assignments, and SGTs or Cisco TrustSec.

*Note: Assigning authorization policies in ISE when authenticating to the network should be reserved for special case scenarios which will be described further in this documentation. For readers who are familiar with ISE already at this point in the document, it is recommended that by default, devices will not be assigned an authorization profile (SGT) during authentication, but rather tagged while traversing the network based on the networking information such as subnet.*

### ISE TrustSec Domain

Not all devices in a network are required to be TrustSec capable for TrustSec to be adopted. In fact, even if all switches in the network are TrustSec capable, it is still recommended that not every switch participates. A TrustSec domain for this design guide can be considered as the policy enforcement layer of your network.

**Figure 3: Defining the TrustSec Domain**



Packets entering the domain are tagged with an SGT containing the assigned security group number of the source device. This packet classification is maintained along the data path within the Cisco TrustSec domain for the purpose of applying security and other policy criteria. The final Cisco TrustSec device in the TrustSec domain, enforces an access control policy based on the security group of source device and the security group of the destination endpoint.

### SGT Classification

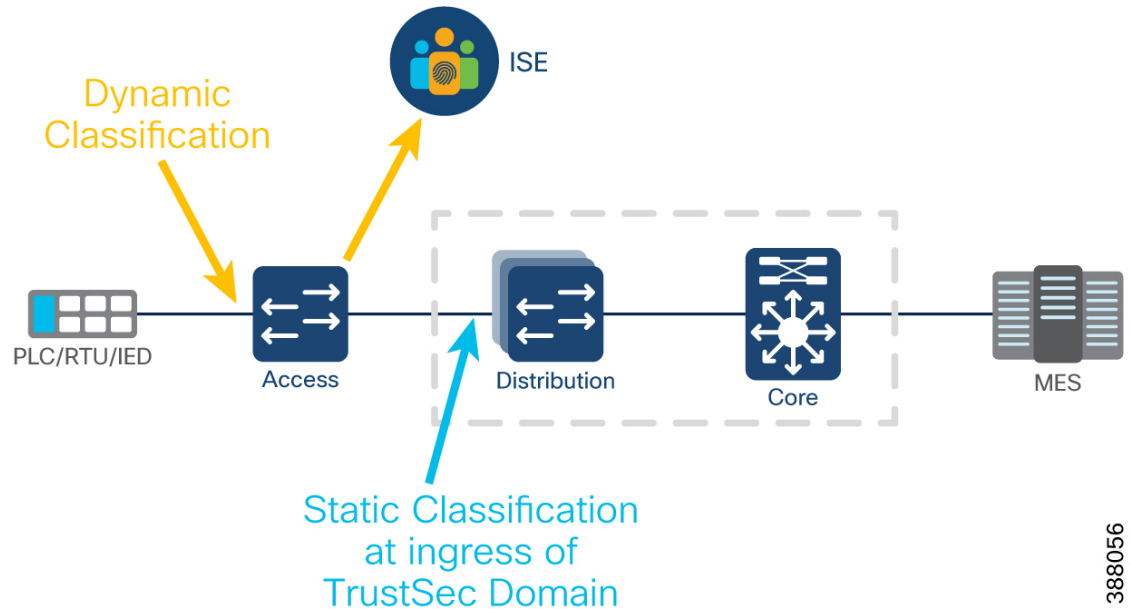
SGT classification, or tagging, can either be dynamic, i.e., obtained from Cisco ISE when network access attempts are made, or static.

**Dynamic tagging** can be deployed with 802.1X authentication, MAB, or web authentication. In these access methods, Cisco ISE can push an SGT to the network access device to be inserted into the client traffic. The SGT is applied as a permission in the ISE authorization policy rules.

**Static tagging** can be configured directly on the networking devices, or statically configured in ISE to be downloaded by the network device. Examples of static tagging include a mapping on an IP host or subnet to an SGT, or the mapping of a VLAN to an SGT.



Figure 4: Static vs Dynamic Classification



388056

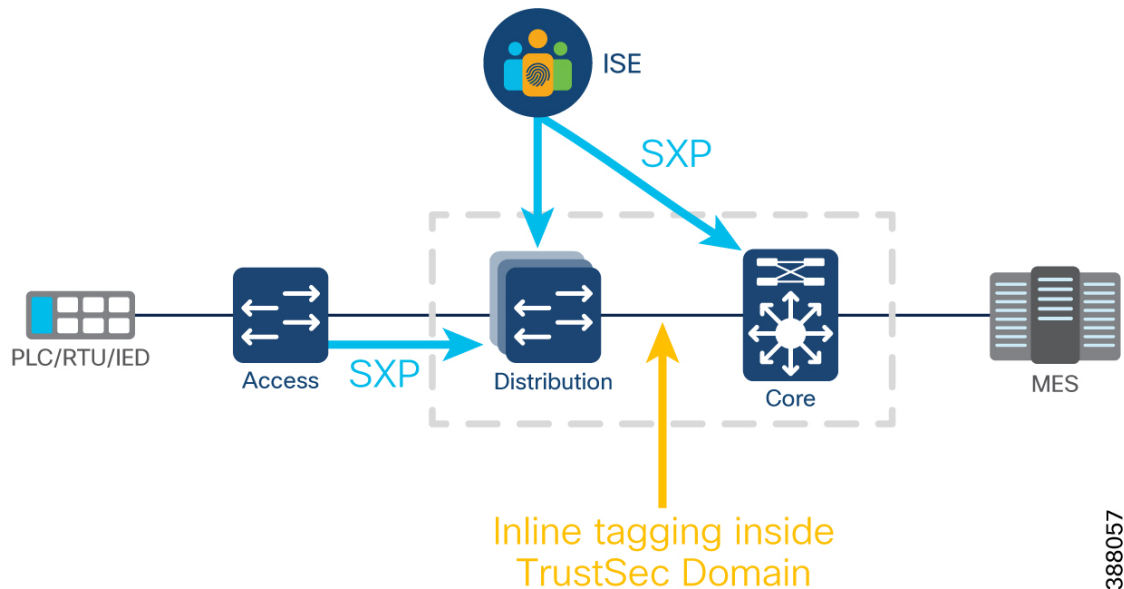
Generally, dynamic classification is done at the access switch and static classification is performed at ingress to the TrustSec domain. The SGT tag that gets inserted into the traffic is known as the source SGT, as it is the group that the source of the traffic belongs to. The destination SGT is the group that is assigned to the intended destination of the traffic. The packet does not contain the security group number of the destination device, but the enforcement point must be aware of this classification.

### SGT Transport

TrustSec has two methods to propagate an SGT, inline and SXP. Cisco TrustSec capable devices have built-in hardware capabilities that can send and receive packets with SGT embedded in the MAC (L2) layer. **Inline tagging** allows Ethernet interfaces on the device to be enabled for SGT imposition so that the device can insert SGT in the packet to be carried to its next hop Ethernet neighbor. The inline propagation is scalable, provides near line-rate performance and avoids control plane overhead. It is recommended that all devices within a TrustSec domain have inline tagging between them when supported.

**SXP** is used to propagate the SGTs across network devices and network segments that do not have support for inline tagging. SXP is a protocol used to transport an endpoint SGT and the IP address from one SGT-aware network device to another. The data that SXP transports is called as IP-SGT mapping. At a minimum, SXP will be enabled between ISE and all devices in a TrustSec domain. However, there are also instances where access switches outside of the domain will establish SXP connections to the first switch within the domain such as sharing the IP-SGT information it stores locally.

Figure 5: SXP vs. Inline Tagging



388057

A network device performing the enforcement needs to determine the destination SGT as well as the source for applying the SGACL. The destination SGT can be determined in one of the following ways:

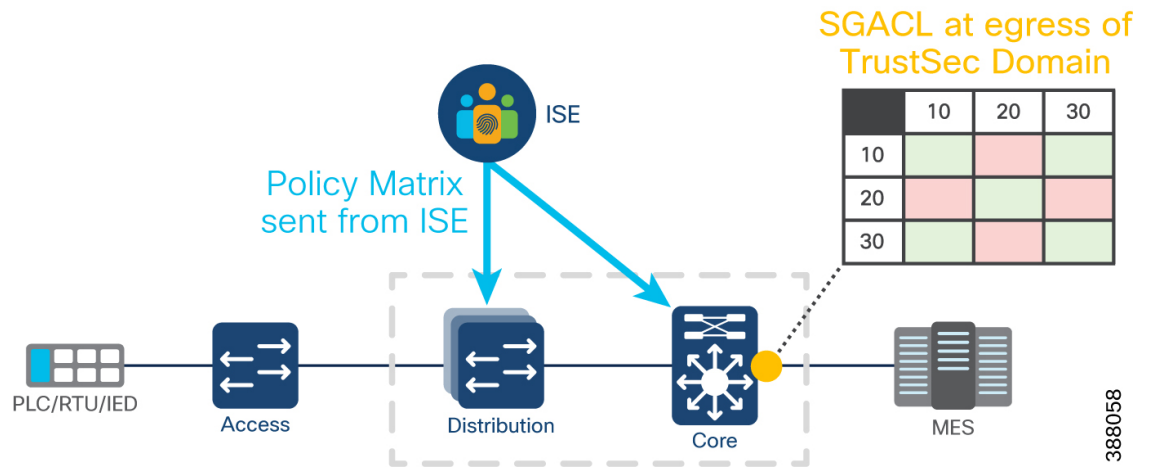
- From ISE using SXP
- SXP from other SGT aware switches (daisy chain SXP communication from access switch to distribution)
- Look up SGT based on destination IP address / subnet
- Look up SGT based on destination physical egress port

### SGT Enforcement

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

It is important to note that the source and destinations are specified in the policy matrix and not in the SGACL. Take, for example, the SGACL entry (ACE) 'deny tcp dst eq 21'. This entry specifies that access from the source to the destination using TCP port 21 is denied. There is no specification of the source or destination group tags in the SGACL. It is the application of the SGACL in the permissions matrix that specifies the source and destination security groups. It is also important to understand that the same SGACL can be applied to multiple source and destination security group pairs within the permissions matrix. Using role-based permissions greatly reduces the size of ACLs and simplifies their maintenance. With Cisco TrustSec, the number of ACEs configured is determined by the number of permissions specified, resulting in a much smaller number of ACEs than when using traditional IP ACLs. Also, only a single copy of an SGACL needs to reside in the TCAM of a device, regardless of how many times the SGACL is used. The use of SGACLs in Cisco TrustSec typically results in a more efficient use of TCAM resources compared with traditional ACLs.

Figure 6: TrustSec Enforcement at egress



By applying access control between pairs of security groups, Cisco TrustSec achieves role-based, topology-independent access control within the network. Changes in network topology do not normally require a change in the SGACL-based security policy. Some care must be taken to ensure the proper classification of new network resources, but the access policy based on business relevant security groups does not change. If the changes do require the creation of a new security group, then the permissions matrix will increase in size by one row and one column. Policy for the new cells is defined centrally in Cisco ISE and dynamically deployed to all SGACL enforcement points.


When using SXP as propagation method from ISE to network devices it is recommended to use SXP domains and domain filters to avoid sending all learned IP-SGT entries to all SXP listeners. This approach will minimize the number of SGT entries on enforcement points and that will ultimately impact the number of SGACLs that the device needs to download from ISE to protect assets in attached cell/area zones.

### SGT Example

The following figure shows an example of traffic entering and exiting a TrustSec domain. The following points apply:

- A PLC with IP address 10.0.1.4 is attempting to communicate outside of the Cell/Area Zone and reach the MES server in the IDC.
- When the traffic hits the ingress of the TrustSec domain (the distribution switch), a static tag is applied using the subnet mapping stored locally on the switch. An SGT 10 is applied to the traffic.
- The link between distribution and core is within the TrustSec domain and inline tagging is enabled. There is no need for an SXP connection between distribution and core because the SGT will remain in the MAC header when it reaches the ingress of the core.
- The core switch is the last point of the TrustSec domain, so the core switch will enforce traffic on its egress port. To apply policy, the core switch must know the SGT of the destination. The IP-SGT relationship is received from ISE via SXP.
- Knowing both the source (SGT 10) and destination (SGT 30) the core switch looks for the corresponding entry in the policy matrix and finds there is a permit any SGACL between the two groups. Traffic will proceed to the IDC.

Figure 7: TrustSec Classification, Transport &amp; Enforcement Example



388059

## ISE/SGT Design Considerations

When using ISE & SGTs for industrial zone enforcement, consider the following:

- While ISE has the capability to apply tags via authentication and authorization (AA) policies, it is not recommended to assign an SGT to every device on the network during the authentication process. Use a hybrid approach between macro and micro segmentation, where the majority of the rules you create are based on the zones in which a device resides, not based on the device type within the zone.
- TrustSec is not optimized to do host-to-host segmentation rules. It is technically possible, but it results in a complex policy matrix as a new SGT will be created for every host-to-host rule required. This results in additional authentication rules, a larger matrix to manage, and can impact the scalability of the deployment.
- The recommendation is to create an SGT based on manufacturing zones and processes and apply a policy to the zone, not the individual devices in the zone. Exceptions to this rule can be made as needed and will be covered later in the guide. In this design guide the following zones are defined:
  - Cell/Area zones (each zone is treated as its own zone, not one large collective zone)
  - Maintenance workstation zone
  - Plantwide application zone
  - Infrastructure zone
- Classifying the zone may differ depending on the network architecture, however, it is recommended that each zone is classified by its own subnet or classless inter-domain routing (CIDR). SGTs can then be assigned statically via a subnet/CIDR to SGT relationship on the ingress of the TrustSec domain.
- TrustSec enforcement should only be enabled on select enforcement points in the network, not on every supported device. In this design guide, the chosen enforcement points were the distribution switch, the core switches and on the IE3400 doing NAT (explained later in this guide).

- On this design, enforcement is applied only on the downlink ports of the TrustSec domain because the objective is to protect traffic on the cell/area zone from unwanted access. To accomplish this, enforcement is enabled globally but disabled on uplink ports.

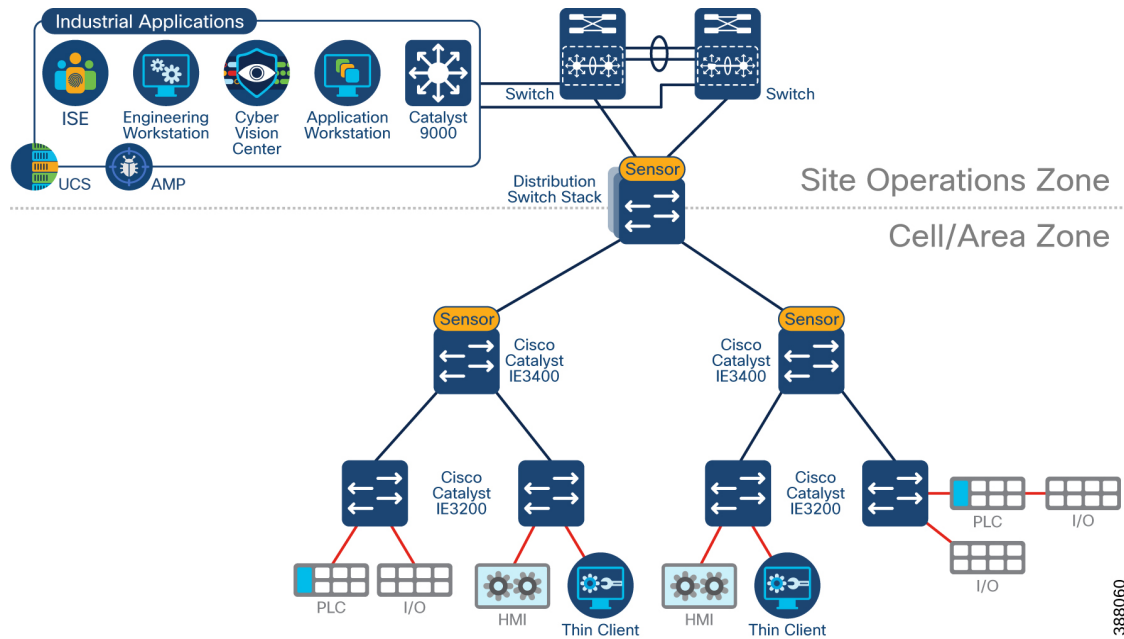
*Note: If using a firewall between the Core switch and the IDC, enforcement on the core switch can be disabled, and the SGT can be used when creating firewall rules. If there is no firewall between the core switch and the IDC, TrustSec enforcement at the core switch egress is recommended.*

- In this design model, the default action is Deny IP and hence the required traffic should be explicitly permitted with the use of SGACLs. This is generally used when the customer has a fair understanding of the kind of traffic flows within their network. This model requires a detailed study of the control plane traffic as well as it has the potential to block ALL traffic, the moment it is enabled. Traffic within the cell will not cross a TrustSec domain, so will be enabled by default in this model.
  - Do not be redundant with policy permissions in the TrustSec matrix. Do not create rules that would ultimately match the default behavior of the matrix. Leave the matrix blank and allow traffic to match the default policy.
  - Use SXP Domain filters to be specific about what entries are needed in each network device. A network device needs only the entries of devices that enter or exit the TrustSec domain through it.
  - Create console access to all enforcement points on the network in case something goes wrong and network connectivity to the devices are lost.
  - When using a deny by default policy the following configurations are recommended for survivability of the site if ISE becomes unavailable:
    - Do not use an unknown SGT tag for switches. Using a dedicated SGT for switches gives more visibility and helps to create SGACL specific for switchinitiated traffic
    - Add static IP-SGT mappings for critical services on core switches and enforcement points. The idea is for Local IP-SGT mapping to be available on the switches even if all ISE goes down
- Configure Fallback SGACL on enforcement points in case ISE nodes go down. When ISE services are down, the SXP connection is lost and hence SGACLs and IP SGT mapping will no longer be downloaded dynamically

Choosing how to tag and where to enforce will depend on how deep in the network you wish to segment, and the network topology deployed in the Industrial Zone.

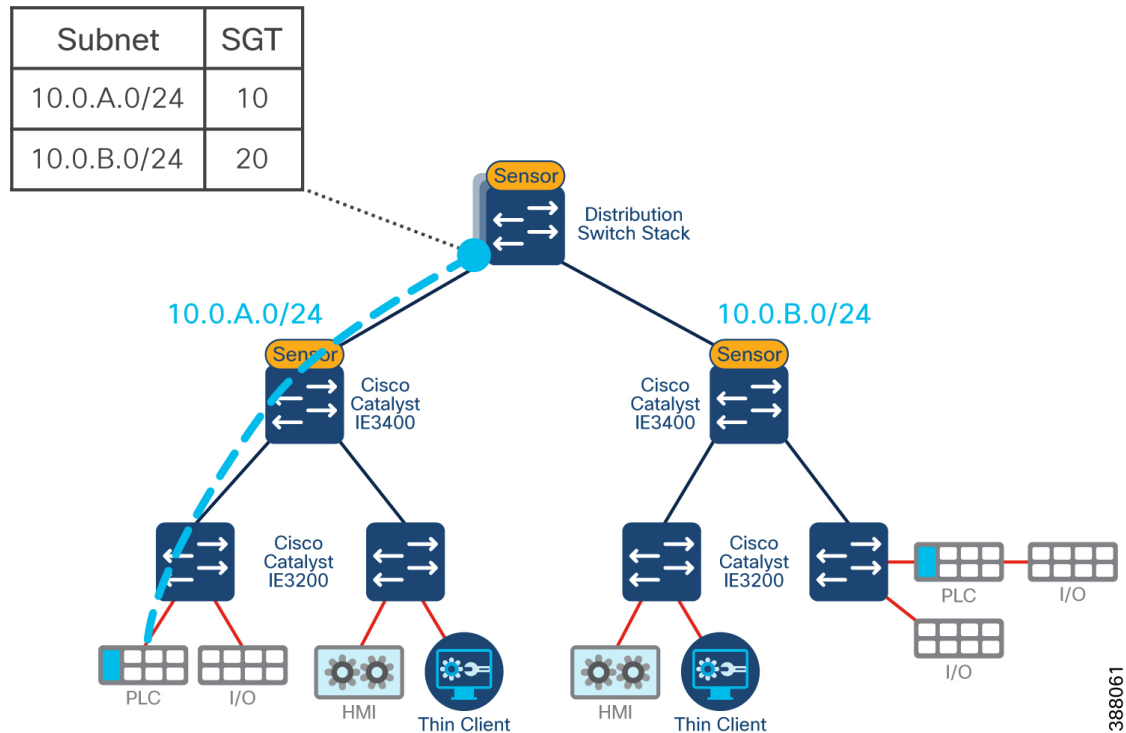
In a tree topology (or any topology where the layer 3 (L3) boundary is outside of the cell/area zone), the distribution switch is used as the L3 gateway between cell/area zones. Each cell/area zone could be a single subnet, or multiple subnets depending on the number of defined VLANs.

Figure 8: Tree Topology in the Industrial Zone



The recommendation for this model is to both classify and enforce at the distribution switch. When creating AA policies on access switches, do not include SGT assignment as part of the policy. Devices will have unrestricted communication within their cell as no PEP exists within the zone. On the distribution switch, define a static subnet to SGT relationship (see Appendix B for switch configuration). When traffic is destined for a service outside the cell, all traffic coming from a select subnet will be tagged on ingress depending on the mapping. The following figure provides an example, where one cell/area zone is tagged with SGT 10, and the other is tagged as SGT 20.

Figure 9: SGT Classification at Distribution Ingress




After cells have been defined, the next step is to define policy for communication that must leave the zone. The least privilege approach to security will result in the Cell/Area zones being in a deny by default state (i.e., SGT 10 deny SGT 20) and only select services crossing the zone boundaries. A common use case is interlocking PLCs, where a PLC in one part of the production facility shares data with another for industrial automation purposes. In this case, PLCs that require interzone communication should be classified with a different SGT to that of the zone it physically resides in so an alternate policy can be enforced across the distribution switch.

There are two methods of assigning a unique SGT to the PLC. The first is a static host to SGT configuration on ISE, where the host to SGT will take precedence over the subnet to SGT relationship and shared via SXP. The second, is by using AA policies in ISE.

The profiling service in ISE identifies the devices that connect to the network. The endpoints are profiled based on the endpoint profiling policies configured in ISE which can subsequently be used in authorization policies and SGT assignment. However, ISE does not natively contain profiling services for IACS devices. To gain visibility of IACS assets, this design uses Cisco Cyber Vision, which provides the context of industrial operations and systems. Cisco Cyber Vision shares endpoints and attributes with ISE using pxGrid.

Figure 10: Cyber Vision &amp; ISE pxGrid Integration



```

assetDeviceType = Controller, IO Module, Rockwell Automation, Controller, Rockwell Automation
assetName = 10.17.10.70,CLX_O | 1756-L73S/B LOGIX5573SAFETY,CLX_O | 1756-L73S/B
LOGIX5573SAFETY (Port1-Link00),Rockwell 3b:55:6f
assetProductID = 1756-L73S/B LOGIX5573SAFETY
assetProtocol = ARP, ARP, CIP-IO, EthernetIP, EthernetIP
assetSerialNumber = 008889a1
assetSwRevision = 26.013
assetVendor = Rockwell Automation

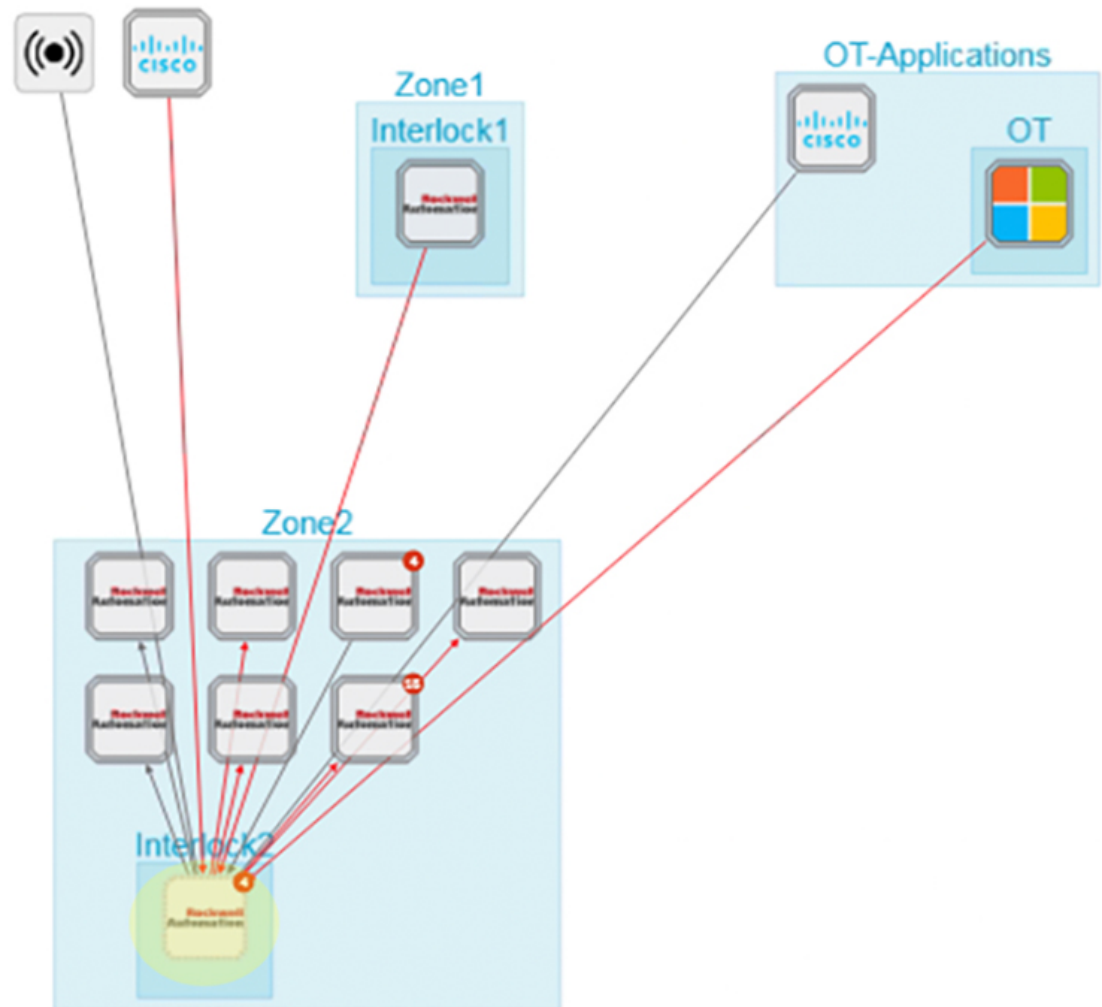
```

388062

In Cyber Vision, when devices and components are placed inside groups, that group tag is shared to ISE via pxGrid. Groups in Cyber Vision can be nested, such that you have a parent group and a child group. It is recommended that a parent group is used to define the production process or areas, such that the visibility groups match the segmentation groups and make logical sense when visualizing network activity. Within the parent group, assign additional group tags to provide context for profiling in ISE, such as an “interlock PLC” group to indicate the device needs to communicate with other control devices in another parent group (cell/area zone for example).






Figure 11: Cyber Vision Interlock Group within Zone2 Parent Group




After the group tag has been shared with ISE, a change of authorization (CoA) is sent to the access switch that the PLC is connected to. This results in the PLC reauthenticating with ISE, ultimately matching the new AA policy defined for interlocking PLCs.

*Note: The CoA does not result in traffic interruption. Traffic will continue to flow as normal until the authentication process is finished and a new SGT can be assigned.*

Figure 12: ISE Asset Information after Cyber Vision Integration

00:00:BC:2D:21:70   


 MAC Address: 00:00:BC:2D:21:70  
 Username: 00-00-BC-2D-21-70  
 Endpoint Profile: CVC\_group\_Interlock2  
 Current IP Address: 10.17.20.72  
 Location: Location → All Locations

Applications    **Attributes**    Authentication    Threats    Vulnerabilities

**General Attributes**

Description



Static Assignment    false

Endpoint Policy    CVC\_group\_Interlock2

Static Group Assignment    false

Identity Group Assignment    CVC\_group\_Interlock2

**Custom Attributes**

 Filter 

Attribute String	Attribute Value
×	Attribute String
×	Attribute Value
assetGroup	Interlock2
assetCCVGrp	
assetSource	CCV

assetDeviceType	Controller, IO Module, Rockwell Automation, Controller, Rockwell Automation
assetId	00:00:bc:2d:21:70
assetIpAddress	10.17.20.72
assetMacAddress	00:00:bc:2d:21:70
assetName	10.17.20.72,CLX_P   1756-L73S/B LOGIX5573SAFETY,CLX_P   1756-L73S/B LOGIX5573SAFETY (Port1-Lin k00),Rockwell 2d:21:70
assetProductId	1756-EN2T/A,1756-L73S/B LOGIX5573SAFETY
assetProtocol	ARP,ARP, CIP-IO, CIP Safety, EthernetIP,EthernetIP
assetSerialNumber	00552b01,00893b40
assetSwRevision	26.013,5.028
assetVendor	Rockwell Automation

After zones have been defined, and traffic has been classified, the policy enforcement matrix can be defined. The following table shows an example policy enforcement matrix.

Figure 13: Sample TrustSec Matrix


	100	101	102	103	911	4001	4002	5001	9001	9002	SGT	Group
100	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	100	Infrastructure
101	Green	Green	White	White	Green	White	White	White	Green	White	101	Management Apps
102	Green	White	White	White	Green	White	White	White	Green	White	102	Plantwide Apps
103	Green	White	White	Green	Green	Green	Green	Green	Green	Green	103	Cyber Vision
911	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	911	911 Tag
4001	Green	White	White	Green	Green	White	White	White	Green	White	4001	Zone 1
4002	Green	White	White	Green	Green	White	White	White	Green	White	4002	Zone 2
5001	Green	White	White	Green	Green	White	White	Green	Green	White	5001	Interlock Zone 1
9001	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	9001	Super User
9002	Green	White	White	Green	Green	White	White	White	Green	Green	9002	TrustSec Devices

First, notice how the matrix is a combination of green and white squares. Up to this point we have been using green and red shades to differentiate an allow rule vs. a deny rule. However, not all squares in the matrix need a value. To save on TCAM space in the switches, only define a rule if it deviates from the default. Since this design uses a default deny rule, any SGT combination that is required to be denied will get no specific policy assigned to it. The decision will fall back to a default rule, which provides the same outcome but with less memory consumption.

*Note: It is recommended to move to a deny by default state only when you are sure that all policies have been accounted for. Policy should be loosely defined to begin with, and the network should be in an allow by default state while gaining visibility with Cyber Vision. Once all communication patterns are understood, enforcement can be fine-tuned. Additionally, when using the Cisco Catalyst 9300, SGACLs can be deployed in monitor mode, so events are created, but no traffic is blocked. For more information see [Configuring SGACL Monitor Mode](#).*

It is important to reiterate the policy enforcement point used in this example is the distribution switch. For example, in the figure below, where 10.0.A.0/24 is assigned the Zone1 tag (4001) from our policy matrix. In our matrix, Zone 1 to Zone 1 communication is denied (default policy). Since the enforcement point is the distribution switch stack, this rule would only take effect if traffic were to leave the zone, and then re-enter the zone. This rule will not stop traffic from flowing within the cell. However, if we changed our enforcement point to the next hop down (Catalyst IE3400), any traffic crossing this boundary would be denied and explicit rules would be required.

Figure 14: Policy Enforcement at Distribution Egress to Cell/Area Zone



388066

When creating the policy matrix, only think about flows that cross TrustSec domains. If a zone does not enter the TrustSec domain, nor does it intend to, it is okay to deny traffic of the same SGT. Use the learnings from the previous step in the journey with Cisco Cyber Vision to understand traffic that flows across L3 boundaries and use that information to inform policy creation.

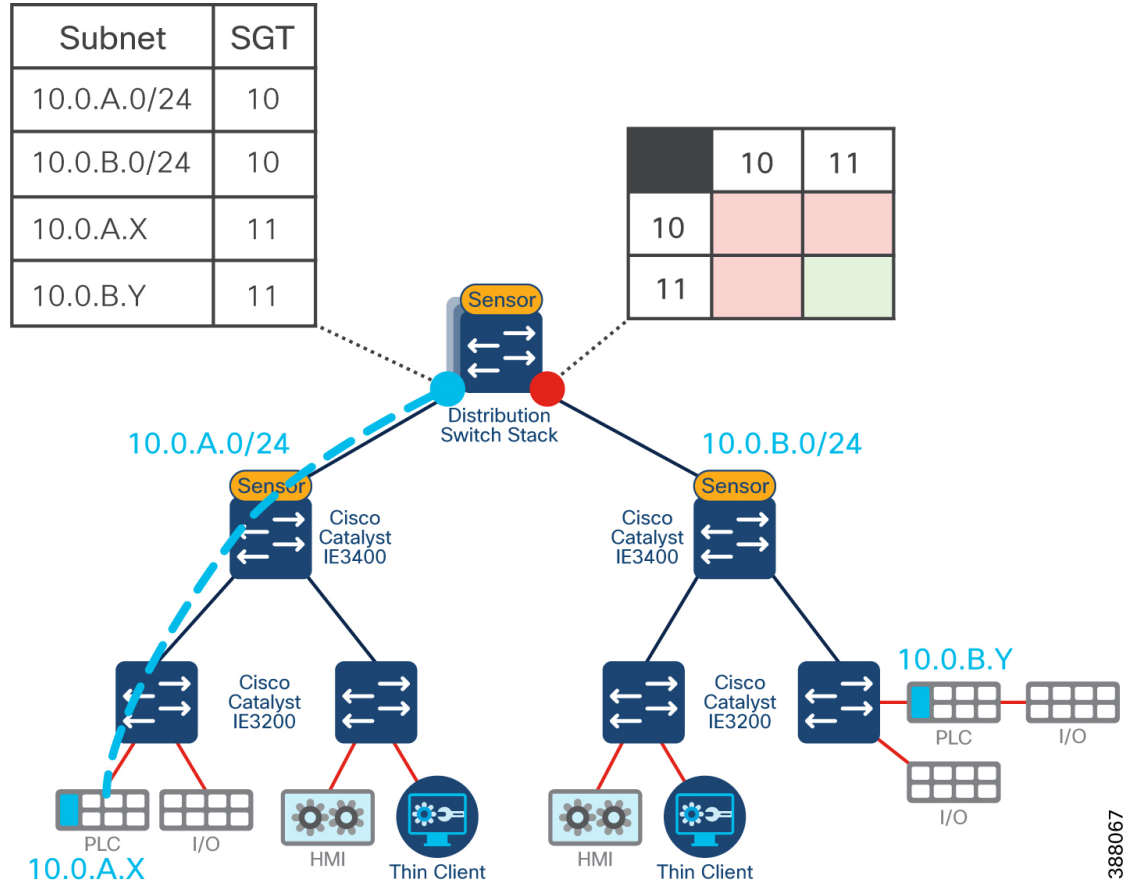
*Note: While discussed as use cases at the start of this guide, safety networks are air-gapped in the validation lab so were not included in the policy matrix. Secure Remote Access is also not part of this guide as it will be addressed in a standalone design guide and linked here upon completion.*

**Scale Considerations in Large Networks**

For security to be effective, it needs to remain simple. For larger networks, where there may be hundreds of zones to manage, it may not be effective to create a unique tag for each zone. Take an example, where 400 cell/area zones exist in the industrial zone. This would result in a matrix that is at least 400 x 400, and even larger if there are multiple VLANs within those zones.

In this scenario, it is recommended to create a single SGT for all zones that do not require any interzone communication, and then deny traffic between zones holding the same SGT. Since tags are classified and enforced in the conduits between zones, no traffic will be denied while it remains in the zone. If traffic were to leave the zone, enter the distribution switch, and come back into the same zone, traffic would be denied. The following figure shows an example where both subnets have been tagged with the same SGT and a deny policy is set between them. Interlocking PLCs are still uniquely tagged, and their communication is enabled. Ultimately, this method leads to a reduction in the matrix size and makes larger networks easier to manage.

Figure 15: Reducing the Policy Matrix Size

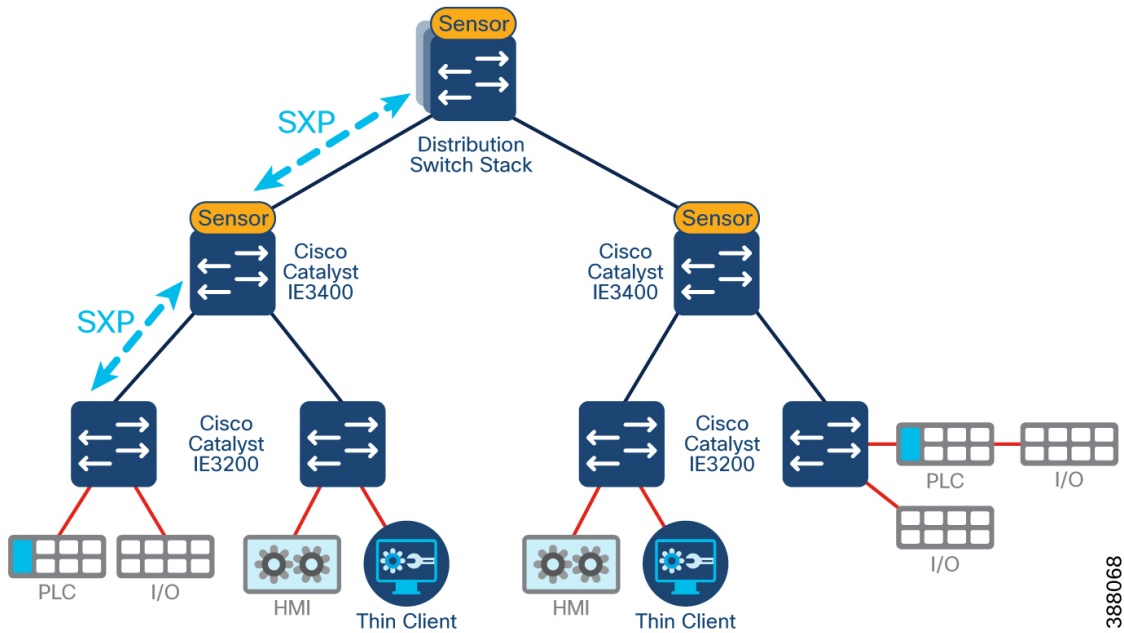


388067

Another consideration in larger networks is the number of SXP connections. The maximum number of ISE SXP peers per PSN is 200. When doing dynamic classification, the IP-SGT binding must be shared from the access switch to the TrustSec domain. One method of doing this is for the access switch to create an SXP session with ISE, and then devices in the TrustSec domain can learn the bindings from ISE. However, in large networks the number of SXP connections may become too much for the ISE nodes to handle.

The recommendation is to daisy chain SXP connections between the access layer and the first layer of the TrustSec domain (in this architecture, the distribution switch). This takes the load off ISE, while still providing the IP-SGT binding. This requires extra switch configuration; though, this could be automated by Cisco DNA Center.

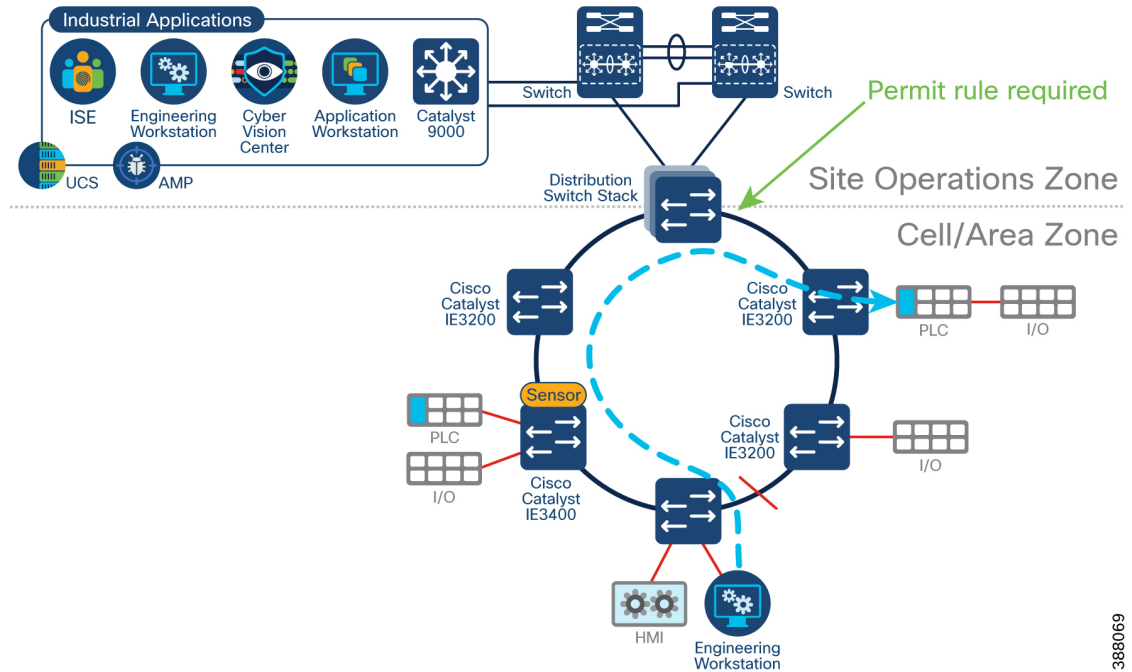
Figure 16: SXP Daisy Chain starting from Access Switch up to Distribution Switch



**Segmentation when the layer 3 boundary also participates in layer 2 connectivity**

Considerations need to be made when the distribution switch is part of the layer 2 communication path such as a ring topology. When the distribution switch is part of a ring, it becomes part of the cell/area zone. Precautions need to be made so that policy will not block communication within the ring. The following figure shows an example where the HMI communicates with two PLCs in a ring. In the case of a link failure between the HMI and a PLC, the alternate path would result in the data crossing the distribution switch to reach its destination.

Figure 17: Inter-Cell/Area Zone traffic traversing the Distribution Switch



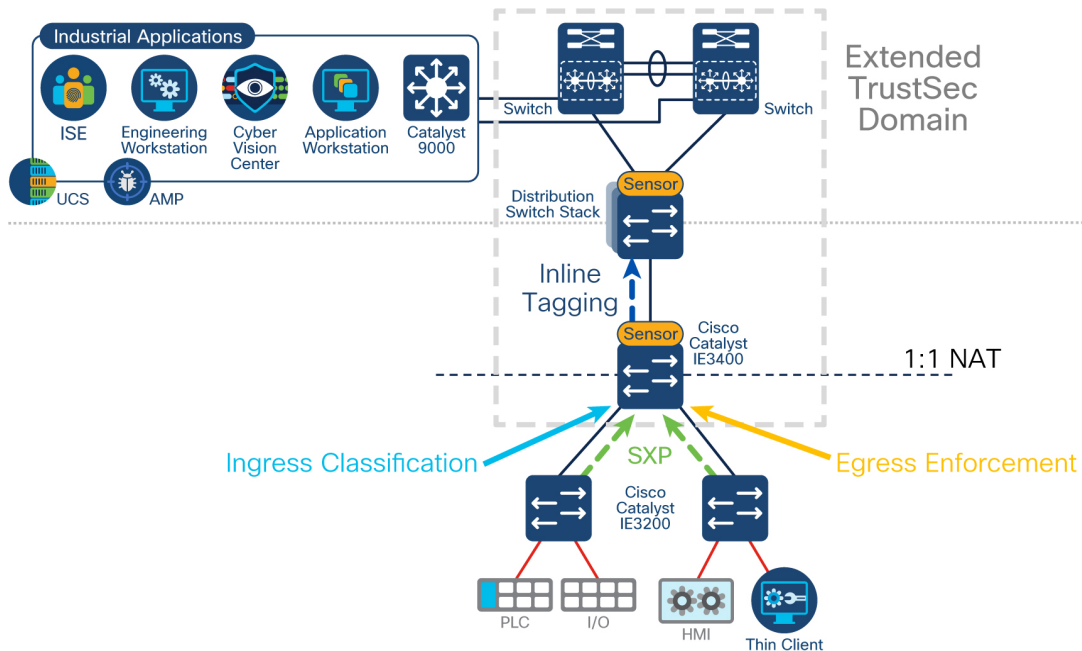
To ensure traffic is not blocked by policy, make sure that each such ring has its own unique SGT and does not share a tag with any other zone in the network as per the design recommendation for large networks.

### NAT Considerations

When doing NAT in the cell/area zone, the IP address of the device when connected will be different to that of the IP the distribution switch will see for tagging. This poses a problem for both static SGT assignment and SGT classification through AA policies. When a device authenticates via ISE, the source IP address is known, and the SGT is assigned to that IP address. However, that IP address will never be seen by the distribution switch and the SGT will never be assigned.

388069

Figure 18: L2 NAT in the Industrial Zone



388070

The recommendation in this case is to enable SGT classification on the IE3400 and enable inline tagging between the NAT boundary and the distribution switch. The SGT is not stripped from the traffic during NAT, and since a tag will already exist when entering the distribution switch, it will not be overwritten by the subnet classification.

In addition to classifying on the IE3400, enforcement is also required. For enforcement, the switch needs to know both the source and destination SGT. When NAT occurs, the distribution switch does not hold the relationship of the true IP address and therefore cannot determine the correct destination SGT when traffic is destined for devices behind the NAT boundary. When enabling enforcement on the NAT boundary, the switch will be able to correctly map the destination SGT and enforce policy as intended.

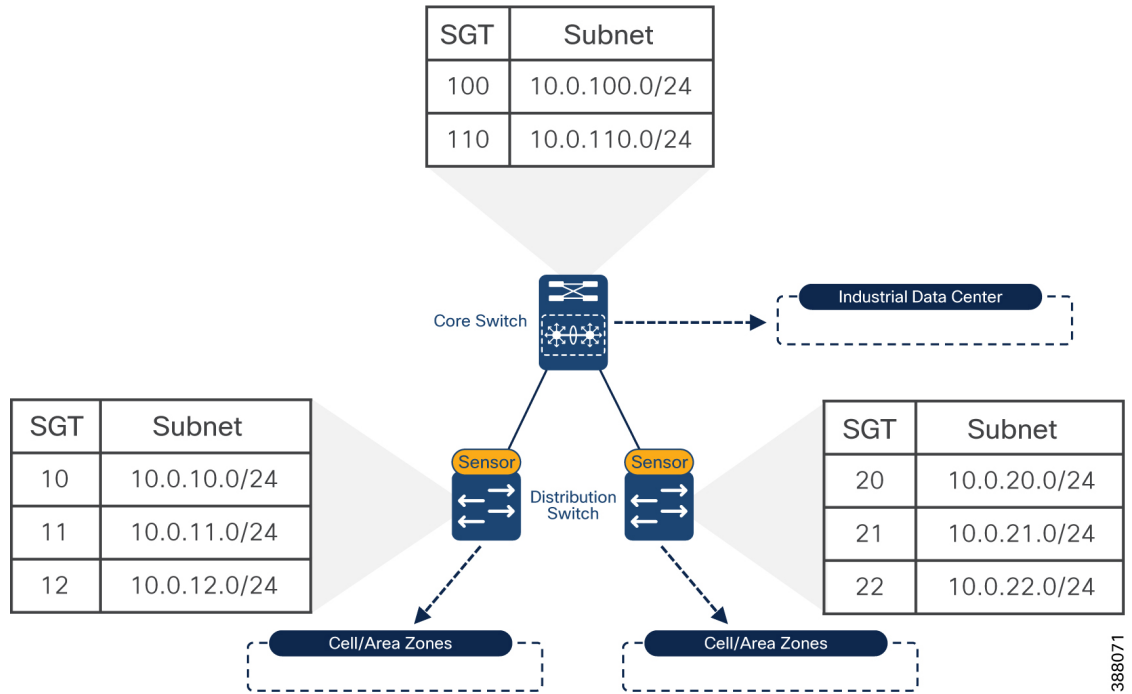
*Note: When using Cyber Vision to assign groups to devices behind the NAT boundary, it is important that you choose the correct device. Depending on sensor placement, Cyber Vision may show two instances of a component when NAT occurs. One will be the instance before the NAT, the second will be after. ISE will only understand the IP address used when authenticating to the network so that is the device in Cyber Vision that should have a group assigned to it for SGT assignment.*

### SXP Domain Filters

An SXP Domain is a collection of SXP devices. There is a default SXP domain that all devices will join when creating SXP sessions. Devices in the default domain will receive all SGT-IP mappings that are known by ISE. SXP domain filters provide a mechanism for SXP peers to deviate from the default, and only receive the IP-SGT mappings that are required for their function on the network. For example, all IP-to-SGT mappings learned through RADIUS authentications are automatically added to the default domain but can be reassigned to a different domain using SXP Domain filters. As a result, any dynamically assigned SGTs can be communicated to the enforcement point that protects assets on that subnet, rather than every switch requiring to store all entries.



Figure 19: SXP Domain Filtering in the TrustSec Domain




388071

**Static Segmentation in the Industrial Zone**

An alternative approach to classifying SGT is static assignment at the access ports on all switches in the network. When assigning SGT directly to ports, authentication with ISE is no longer required. In this case, it is the responsibility of the network administrator to apply the correct SGT to the port, and a process be implemented for local operators to follow.

Figure 20: Static SGT Classification on the Physical Ports of the IE3400



Consider the following when implementing static port assignment on the access switches:

- Static SGT configuration of the physical switch port is only supported on IE3400 and IE9300
- Access switches become part of the TrustSec domain
- SXP is required to propagate the static SGT to the enforcement point
- Switch Integrated Security Features (SISF) needs to be enabled on the access port for the switch to incorporate the static SGT on the IP to SGT bindings
- Do not assign SGTs to any of the physical switch ports that may lead to privileged access of network resources as a local operator could inadvertently open an attack vector by connecting devices to a domain with more freedom simply because it caused the application to work

### Applying Policy to Users

802.1X is an IEEE standard for layer 2 access control, offering the capability to permit or deny network connectivity based on the identity of the end user or device. 802.1X is typically not supported in OT devices, however, it should be a common feature on an employee or contractor end-user device such as a laptop.

802.1X provides a way to link a username with an IP address, MAC address, switch, and port. It also enables you to leverage an authenticated identity to dynamically deliver policy. In ISE, users authenticating via 802.1X can match a Dot1X Authentication rule and be assigned an SGT based on the user group the credentials belong to. A key recommendation is to have all users authenticate to ISE via 802.1X to receive their SGT tag for network entitlements.

It is common for Microsoft Active Directory (AD) to be used as the identity provider (IdP) in industrial networks. When using AD for user authentication, user groups will have already been defined. There may be groups created for administrators, technicians, contractors, etc., all with their own access rights when they connect to the network. Cisco ISE leverages AD for multiple methods of authentication, including 802.1X. When connecting ISE to an AD domain, the user groups configured in AD are imported and can be used when creating authentication and authorization policies.

The design recommendation is to use Microsoft AD for user group definitions and maintenance, and then use those AD defined groups within Dot1X authentication policies to assign a group tag. The figure that follows shows an example of this, where the Employee group in AD is assigned the Employees group tag. This tag can be subsequently used in the TrustSec policy matrix to determine which network zones the employees have access to.

**Figure 21: ISE AA Policy with Active Directory User Group**

