Industrial Automation Security Design Guide 2.0

First Published: 2023-01-17

Last Modified: 2023-01-17

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

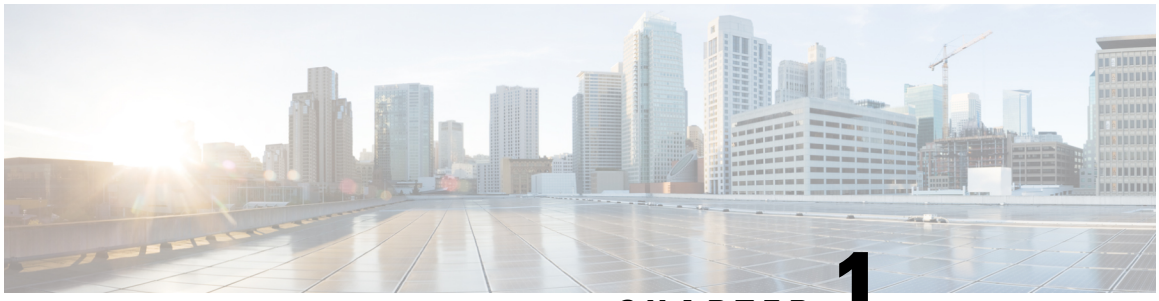

CHAPTER 1

Introduction

Introduction

Protecting industrial automation and control systems (IACS) from cyber threats is top of mind. But converting good intentions to action can be a daunting task. As IACS and underlying networks are often very complex, using legacy technologies and poor security procedures, one could wonder where to start.

For over 15 years, Cisco has been helping industrial organizations around the globe digitize their operations by developing a market-leading networking portfolio that is purpose-built for industrial use cases. Our deep understanding of operational technology requirements plus a comprehensive networking and cybersecurity portfolio is a rare combination.

Cisco's industrial security architecture simplifies complexity across the network by implementing a model that focuses on the use cases an organization must secure. This model treats each use case holistically, focusing on today's threats and the capabilities needed to secure the operational network against those threats. Cisco has deployed, tested, and validated the solution to provide guidance, complete with configuration steps that ensure effective and secure deployments for our customers.

Target Audience

To successfully connect and secure the industrial environment, all stakeholders must work together. Operational technology (OT) teams understand the industrial environment—the devices, the protocols, and the operational processes. Information technology (IT) teams understand the network. The security team understands threats and vulnerabilities. By working together, these specialists can leverage existing networking and security technologies, tools, and expertise to constantly protect the industrial systems without disrupting production safety and uptime.

The Cisco Industrial Threat Defense solution is intended to be used by IT, OT, and security teams and their relevant partners and system integrators. Operations will appreciate the ease of use and simple deployment, as well as the broad support of various IACS vendors and protocols. IT network managers will appreciate the ability to apply skills, technology, and applications already deployed in the enterprise when looking to integrate

production environments. Security teams will have visibility into industrial assets and security events with context enriched by control engineers.

Security Challenges

As highlighted in Figure 1, the first key challenge when securing the IACS is collaboration between IT and OT teams. The view that OT and IT are distinctly separate entities is antiquated. Failing to acknowledge the increasingly interconnected nature of OT and IT can have detrimental consequences for industrial organizations. A lack of trust, understanding, and collaboration between OT and IT departments can have a devastating impact on the security posture of an organization.

Figure 1: Common security challenges in Industrial Networking

IT and OT personnel have different operating procedures and roles to play, and their worldview can differ considerably. However, their goals with respect to securing the company should be identical, and the path forward involves finding common ground. OT personnel are focused on safety, reliability, and productivity. Their role is to protect people, lives, the environment, the operation, and production. Conversely, cybersecurity personnel are focused on maintaining the confidentiality of information and the integrity and availability of IT systems. However, the goals of these entities do overlap. Both are committed to securing the organization, minimizing risk, maximizing uptime, and ensuring that the organization can continue to safely generate revenue.

The second challenge when securing the IACS network is a lack of visibility. As industrial networks can be quite old, widely dispersed, and involve many contractors, operators often do not have an accurate inventory of what is on the network. Without this, they have limited ability to build a secure communications architecture. A lack of visibility also means operators are often unaware of which devices are communicating to each other or even of communications reaching industrial devices from the outside.

The lack of visibility ultimately leads to a lack of segmentation or control. OT networks have been deployed over the years with few or no security policies in place. Networks were not designed with security in mind, updates and patches are harder to deploy, and downtimes are less acceptable. It is critical that OT operations use the visibility to implement the segmentation in their network, as impacts can range from faulty production to lost revenues and even bodily injury, death, or damage to the environment.

Last, but certainly not least, security needs to be top of mind due to threat of malware impacting the environment. Malware must be prevented, when possible, detected when it attempts to breach a network, and

contained to limit potential damage when it infects systems and endpoints. Malware defense calls for a new best-of-breed architectural approach that spans all layers of the industrial network.

Reference Architecture

Plant Logical Framework

To understand the security and network systems requirements of an IACS, this guide uses a logical framework to describe the basic functions and composition of an industrial system. The Purdue Model for Control Hierarchy (reference ISBN 1-55617-265-6) is a common and well-understood model in the industry that segments devices and equipment into hierarchical functions. In addition to the levels and zones, Figure 2 includes an additional demilitarized zone (DMZ) between the enterprise and industrial zones. The purpose of the DMZ is to provide a buffer zone where data and service can be shared between the enterprise and industrial network.

Figure 2: Logical Industrial Cybersecurity Framework for Industrial Automation Networks

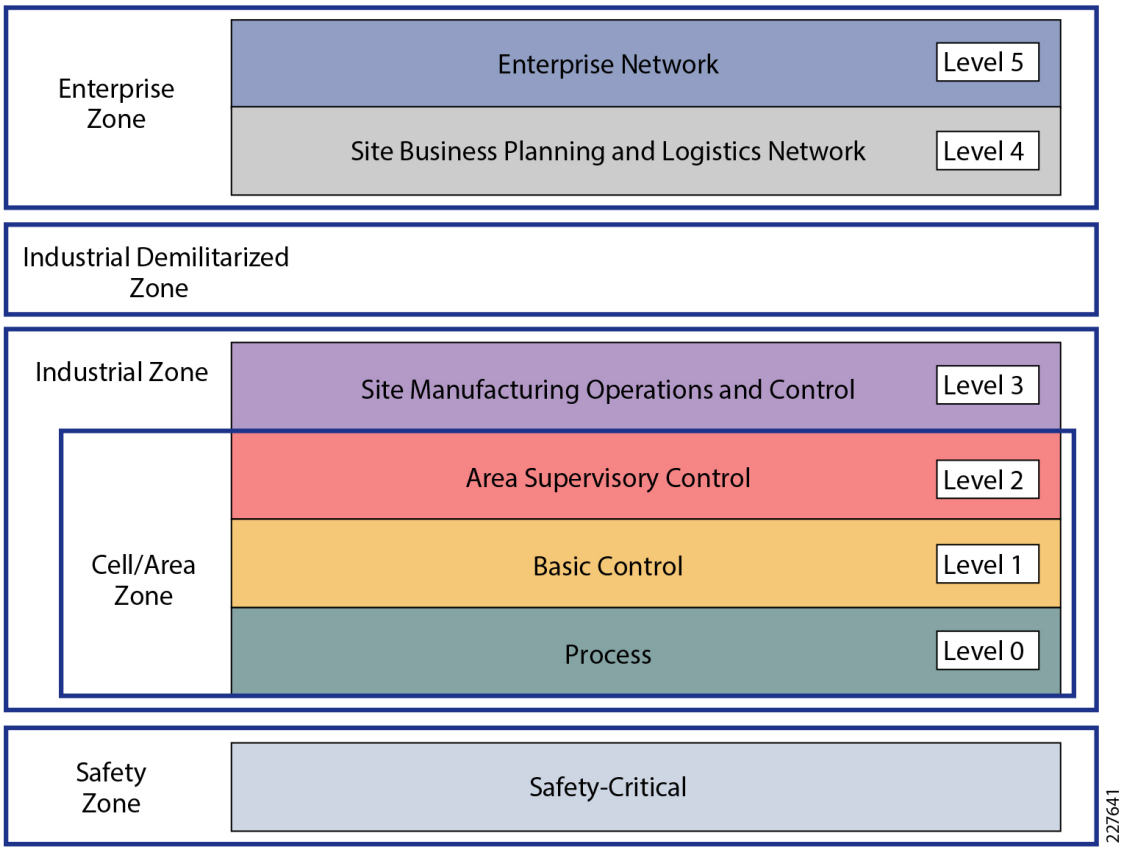

Industrial Zone

The Industrial zone is important because all the IACS applications, devices, and controllers critical to monitoring and controlling the plant floor IACS operations are in this zone. To preserve smooth plant operations and functioning of the IACS applications and IACS network in alignment with standards such as IEC 62443, this zone requires clear logical segmentation and protection from Levels 4 and 5.

The **Safety Zone** may be the most critical zone in an IACS environment. For example, in a manufacturing environment, a robot can cause a fatal impact to personnel if proper safety procedures are not followed. Not only are safety networks isolated from the rest of the IACS (as per Figure 2, positioned below the Industrial Zone), but they typically also have color-coded hardware and are subject to more stringent standards. Industrial automation allows safety devices to coexist and interoperate with standard IACS devices on the same physical infrastructure to reduce cost and improve operational efficiency, resulting in the need for effective security controls to protect from malicious actors looking to cause harm.

The **Cell/Area Zone**, a functional area within a plant or factory, is the foundation of an industrial automation architecture. Most plants will have 10s if not 100s/1000s of functional areas. This is the network that connects sensors, actuators, drives, controllers, robots, machines, and any other IACS devices that need to communicate in real-time (I/O communication). It represents Levels 0-2 of the Purdue model. Most importantly, Cell/Area Zone networks support the critical automation and control functions that keep the plant operating and producing quality products. Fundamentally, the Cell/Area Zone is a Layer 2 access network: a subnet, a broadcast domain, a virtual local area network (VLAN) and/or a service set identifier (SSID). PLCs communicate with their assigned sensors, actuators, and other IACS devices within a Cell/Area Zone. Some industrial traffic is Layer 2 only as there is no IP header attached.

Level 3, the **Site Operations and Control Zone**, represents the highest level of the IACS network and completes the segments of the Industrial Zone. Site operations is generally a “carpeted” space meaning it has heating, ventilation and air conditioning (HVAC) with typical 19-inch rack-mounted equipment in hot/cold aisles utilizing commercial grade equipment. As the name implies, this is where applications related to operating the site reside, where operating the site means the applications and services that are directly driving production. These applications are primarily based on standard computing equipment and operating systems (Unix-based or Microsoft Windows). For this reason, these systems are more likely to communicate with standard Ethernet and IP networking protocols. As these systems tend to be more aligned with standard IT technologies, they may also be implemented and supported by personnel with IT skill sets.

Enterprise Zone

The enterprise zone is where the traditional IT systems exist. These functions and systems include wired and wireless access to enterprise network services such as:

- Internet Access
- Email services
- SAP
- Oracle

Although important, these services are not viewed as critical to the IACS and thus industrial zone operations. Direct access to the IACS is typically not required, but there are applications such as remote access and data collection where traffic must cross the IT/OT boundary. Access to the IACS network from an external zone must be managed and controlled though the industrial demilitarized zone (IDMZ) to maintain the security, availability and stability of the IACS.

Industrial DMZ

Although not part of the Purdue mode, the industrial DMZ is deployed within plant environments to separate the enterprise networks and the operational domain of the plant environment. Downtime in the IACS network can be costly and have a severe impact on revenue, so the operational zone cannot be impacted by any outside influences. Network access is not permitted directly between the enterprise and the plant; however, data and services are required to be shared between the zones, thus the industrial DMZ provides architecture for the secure transport of data. Typical services deployed in the DMZ include remote access servers and mirrored

services. Further details on the design recommendations for the industrial DMZ can be found later in this guide.

Understanding the threats

There are many great resources when learning about the techniques used to infiltrate industrial networks. For example, MITRE ATT&CK for ICS is a knowledge base useful for describing the actions an adversary may take when operating within an IACS network. ATT&CK is short for Adversarial Tactics, Techniques, and Common Knowledge.

There is also the yearly Verizon Data Breach Investigations Report (DBIR) which analyses thousands of incidents and confirmed breaches from around the world so security analysts can understand the most commonly exploited vulnerabilities across industries.

This design guide will use elements of both resources to look at some of the common attack vectors, exploring what they mean and the mitigations that can be put in place to defend against them. Figure 3 shows four common attack techniques described in the MITRE ATT&CK framework.

Figure 3: Typical attack techniques used to exploit the Industrial Network

Initial Access is described by MITRE ATT&CK as an adversary attempting to get into your IACS environment. This is traditionally accomplished by exploiting public facing applications, or the exploitation of remote services. The Colonial Pipeline attack for example, while not an entry into the OT network, was a result of a forgotten Virtual Private Network (VPN) termination point with stolen credentials and no Multi-Factor Authentication (MFA). The 2022 Verizon DBIR stated that over 80% of attacks come from external sources, and with many industrial sites using technologies such as VPN and Remote Desktop Protocol (RDP) for remote access services or implementing Industrial IoT (IIoT) gateways for data collection, it is critical that public facing applications are implementing with security top of mind.

In the case where initial access security was poorly implemented, or an exploit has been found, the first thing an adversary will do on the network is try to discover more information to identify and assess targets in the IACS environment. Triton malware is an example of this where a python script was executed in the network to discover Triconex safety controllers distributed by Schneider Electric. Triconex safety controllers used a proprietary protocol on UDP port 1502, and Triton used this knowledge to scan the network for the devices. If the device exists, the malware can then read the firmware version and use this information in the next phase of an attack. Network segmentation is a great way to combat this threat, as if an attacker does manage to exploit a machine in the network, their reach should not be able to extend beyond the network segment the exploited machine is in. Additionally, being able to detect the presence of network scans enables security analysts to react before an adversary has the chance to use the discovered information in an exploit attempt.

Lateral Movement refers to the adversary attempting to move through the IACS environment. This could involve jumping to engineering workstations using RDP with weak or default credentials, or in the case of e.g., the WannaCry vulnerability, using protocol exploits to hop across machines in the network. Other than

making sure default credentials are not used within the IACS, network segmentation helps solve this problem too, by containing an adversary to the zone in which the initial exploit occurred.

Finally, the adversary will try and communicate with, and control compromised systems, controllers, and applications within the IACS environment. This is known as Command and Control.

Use cases

Common use cases and personas that must be secure in an industrial network include:

- **Cell/Area Zone**: The industrial zone is typically comprised of multiple cell/area zones. All devices located within a given Cell/Area zone should be able to freely communicate with all other assets in this zone. Communication that crosses zone boundaries should be denied unless explicitly allowed as depicted in Figure 4.

Figure 4: Cell/Area Zone to Cell/Area Zone denied by default. No segmentation inside the zone.

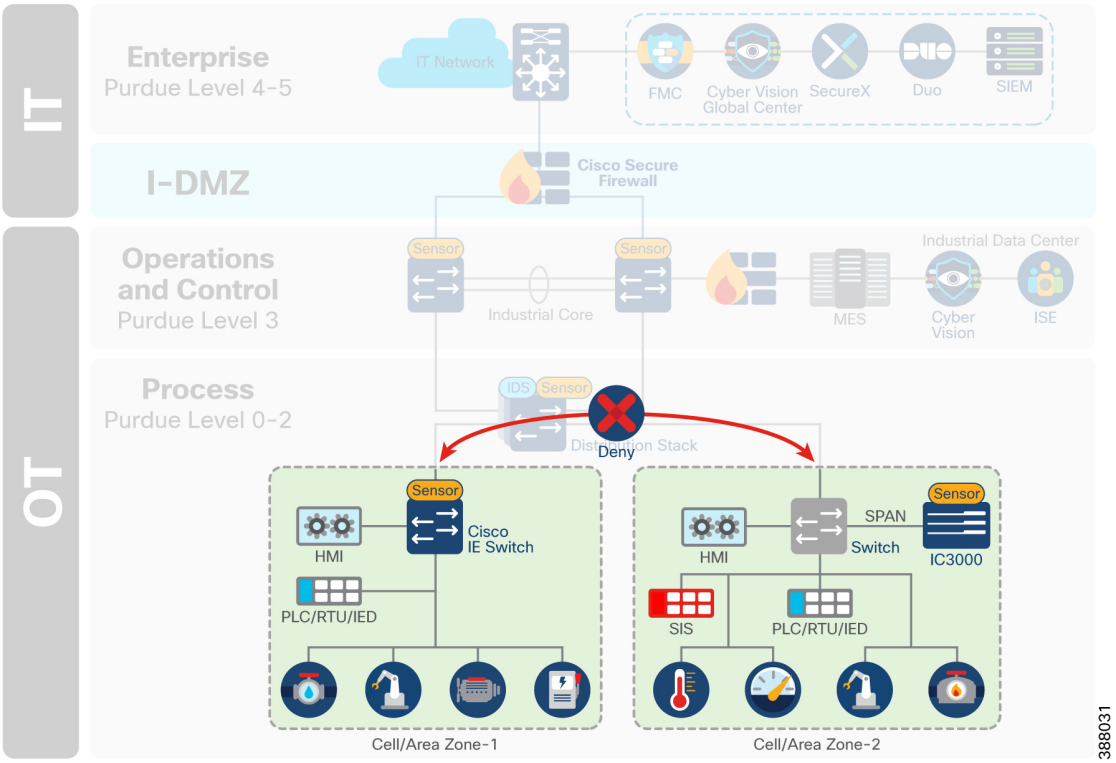

- **Administrative Users**: Figure 5 shows an administrative user who requires access to all zones in the network. They may be responsible for configuration of the network infrastructure, or the application of control logic. Their access should not be limited, but their data should be protected.

Figure 5: Administrative Users need access to all zones

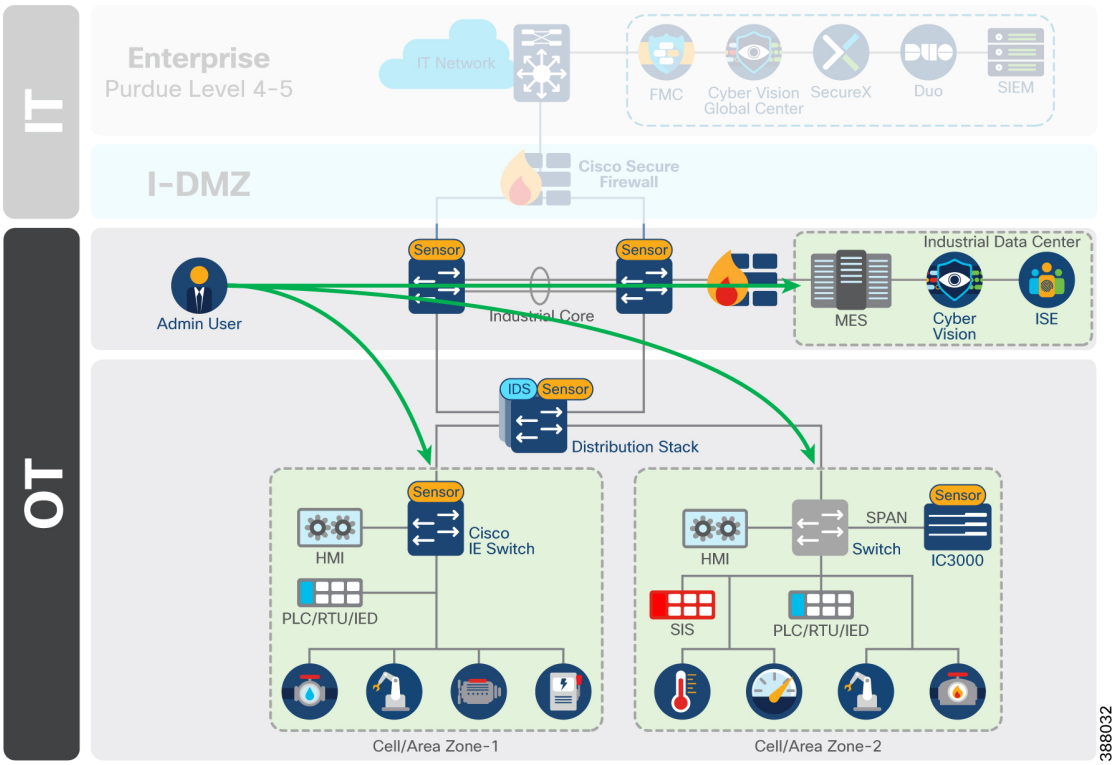

- **Infrastructure Services**: Endpoints that do not have user presence, but still require access to a large chunk of the plant. Services such as DHCP, NTP or LDAP may touch each device on the network.

Figure 6: Infrastructure services that need access to all Cell/Area Zones

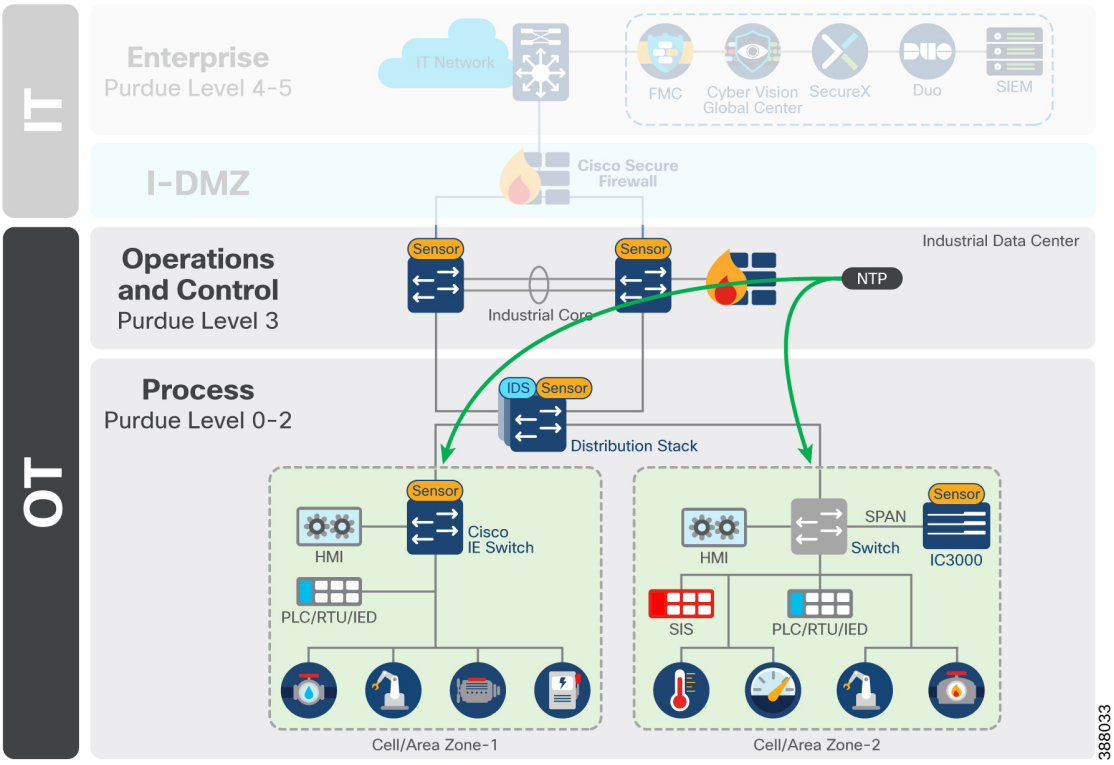

- **Plantwide Applications**: Applications within the industrial data center (IDC) that have specific access requirements. Examples include analytics platforms that require read only access to relevant machinery, or vendor tools used to monitor and maintain plant floor equipment.

Figure 7: Applications that need access to specific services in the cell, but not the full cell

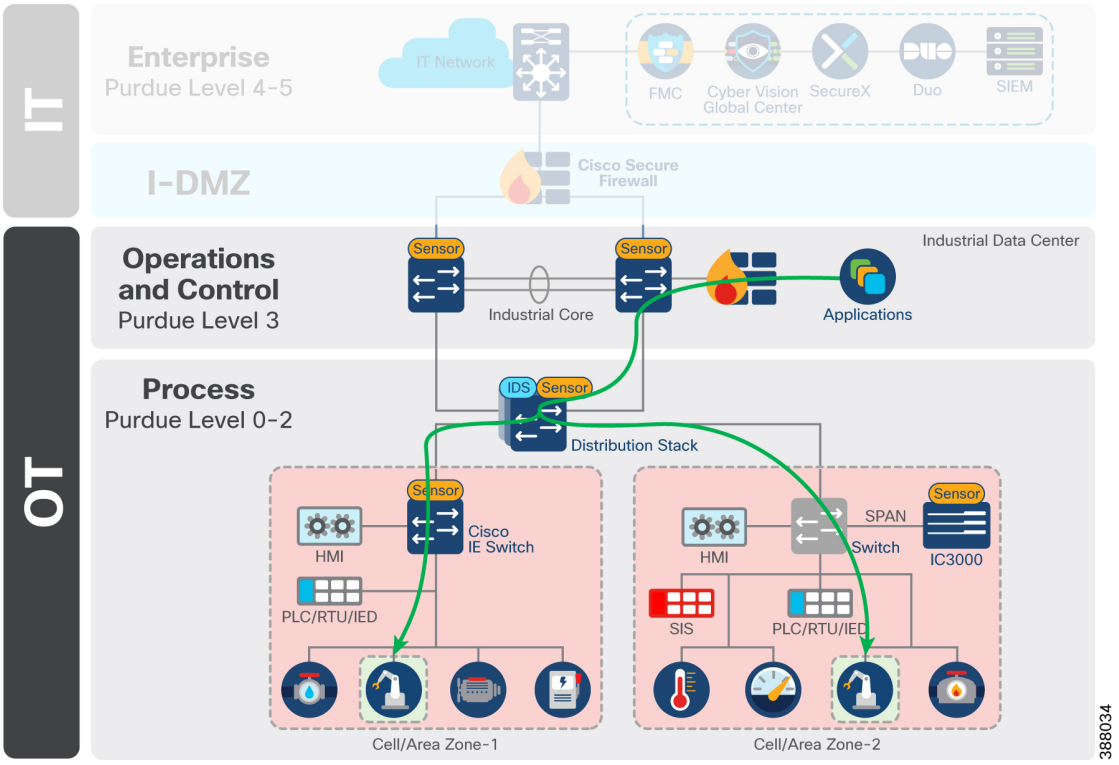

- **Maintenance Workstations**: Maintenance workstations can either reside in a zone outside of the cell/area, and act as the maintenance machine for select zones, or reside within the cell/area zone itself, but require additional privileges when leaving the zone.

Figure 8: Maintenance workstations may reside within the cell or in a dedicated zone outside the cell

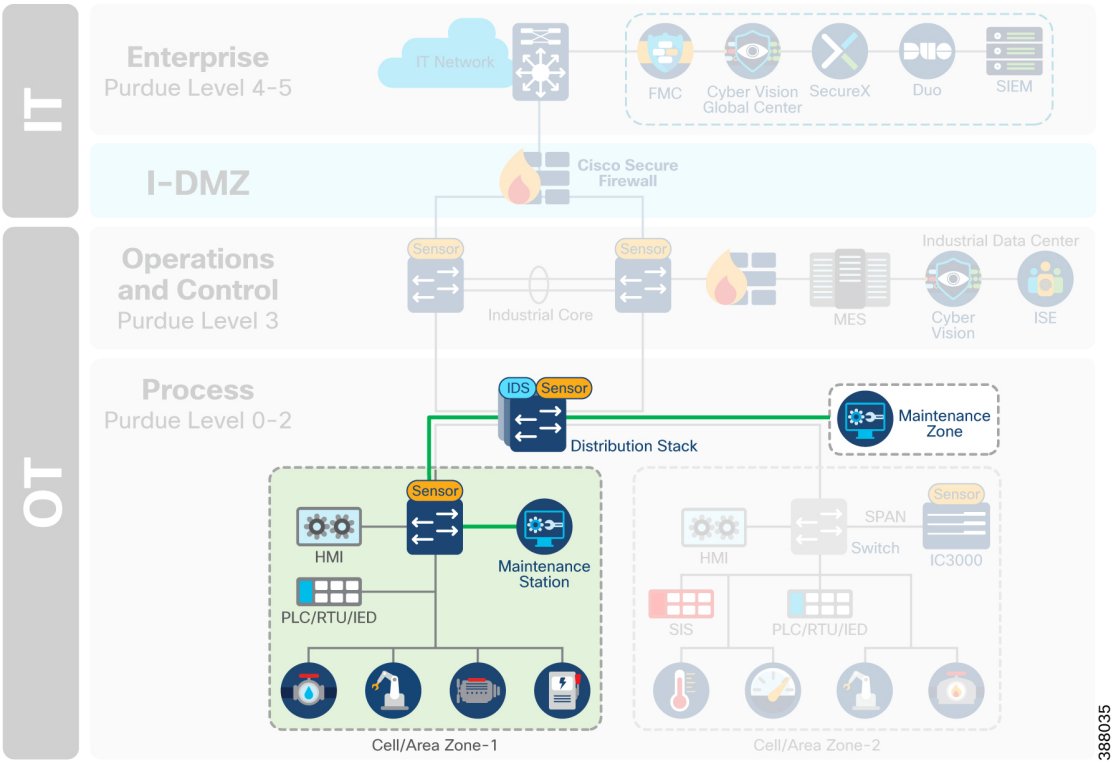

- **Interlocking Programmable logic controllers (PLC) / Interzone communication:** While most control traffic is contained within a cell/area zone, some industrial communications may need to traverse zones for distributed automation functions. A PLC in one zone should not have full access to the services in another zone and least privilege policy should be applied to ensure only valid communication are permitted. If malware was to be introduced into one zone on the network, it is important that it has no automated mechanism to spread to other zones.

Figure 9: Select devices, such as interlocking PLCs, require communication across zones

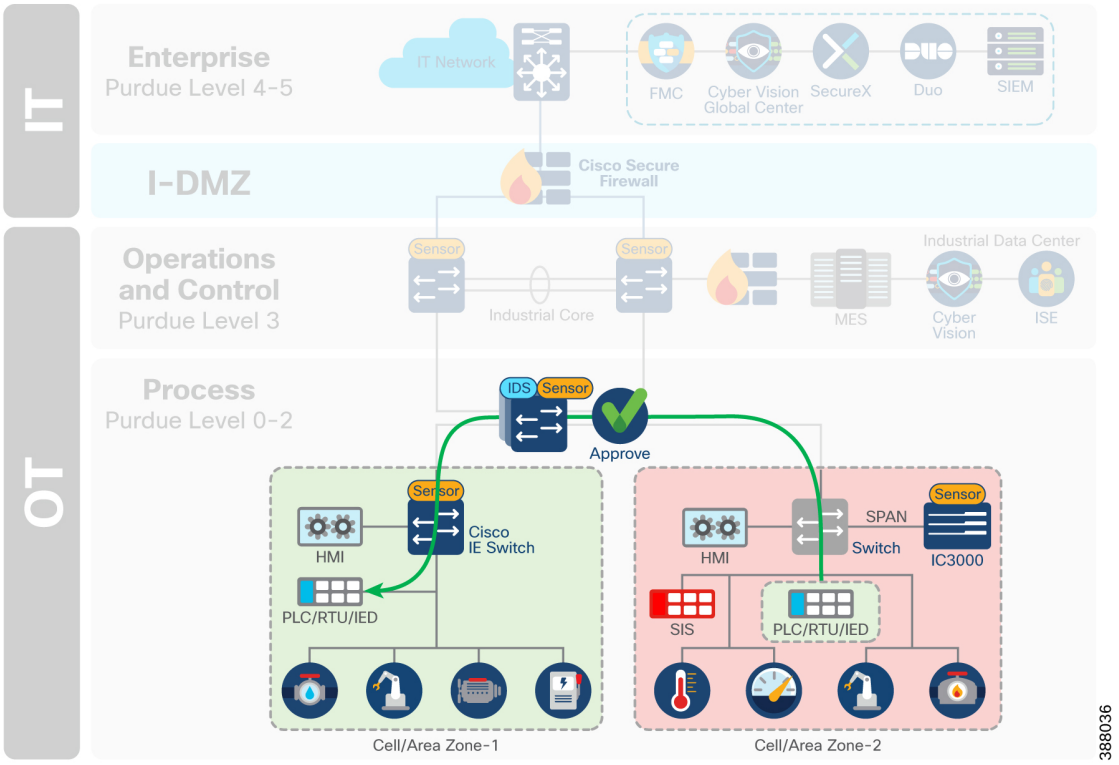

- **Convenience Port:** As operators plug directly into the infrastructure, they will typically bypass all the security checks that have been deployed in the architectural layers above it. Ensuring only authorized users with authorized device posture checks can connect to the network can aid in securing this use case.

Figure 10: Operators plug into a cell using a convenience port and should be able to reach out to extra services

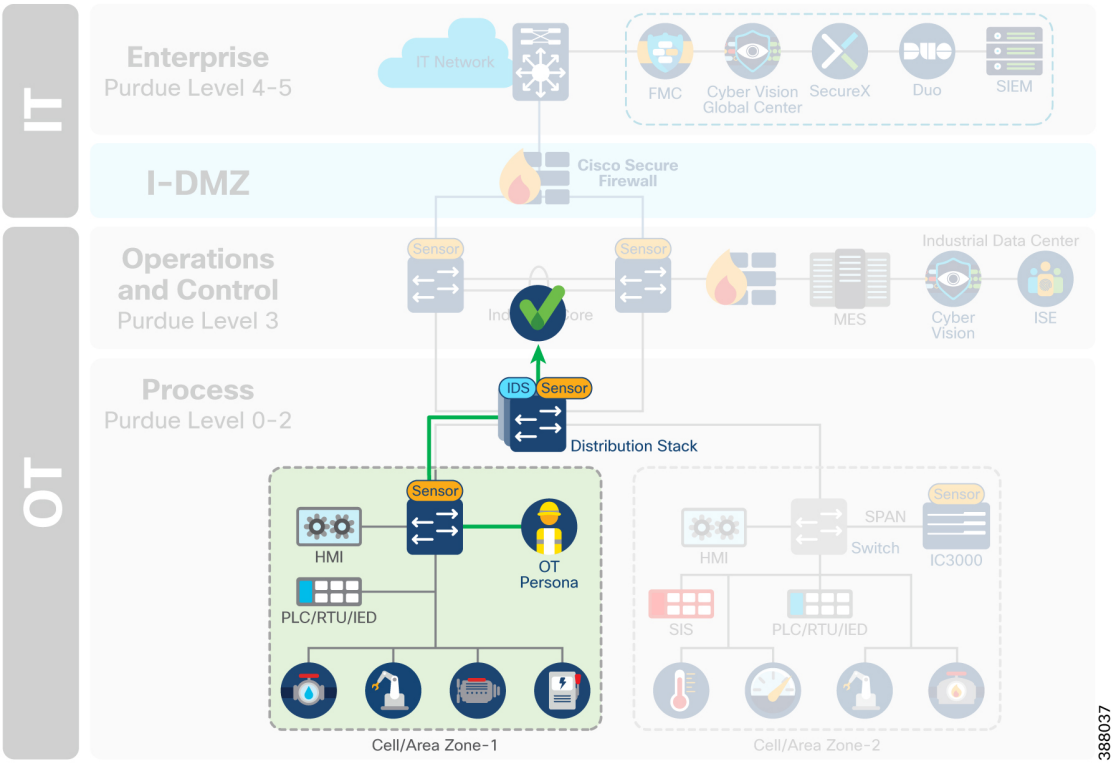

- **Safety Networks:** Safety Instrumented Systems (SIS) are critical to the control network and should either be air gapped from the rest of the network or logically segmented to ensure no data can leak into this zone.

Figure 11: Safety Network could be air gapped, or logically separated from the rest of the network

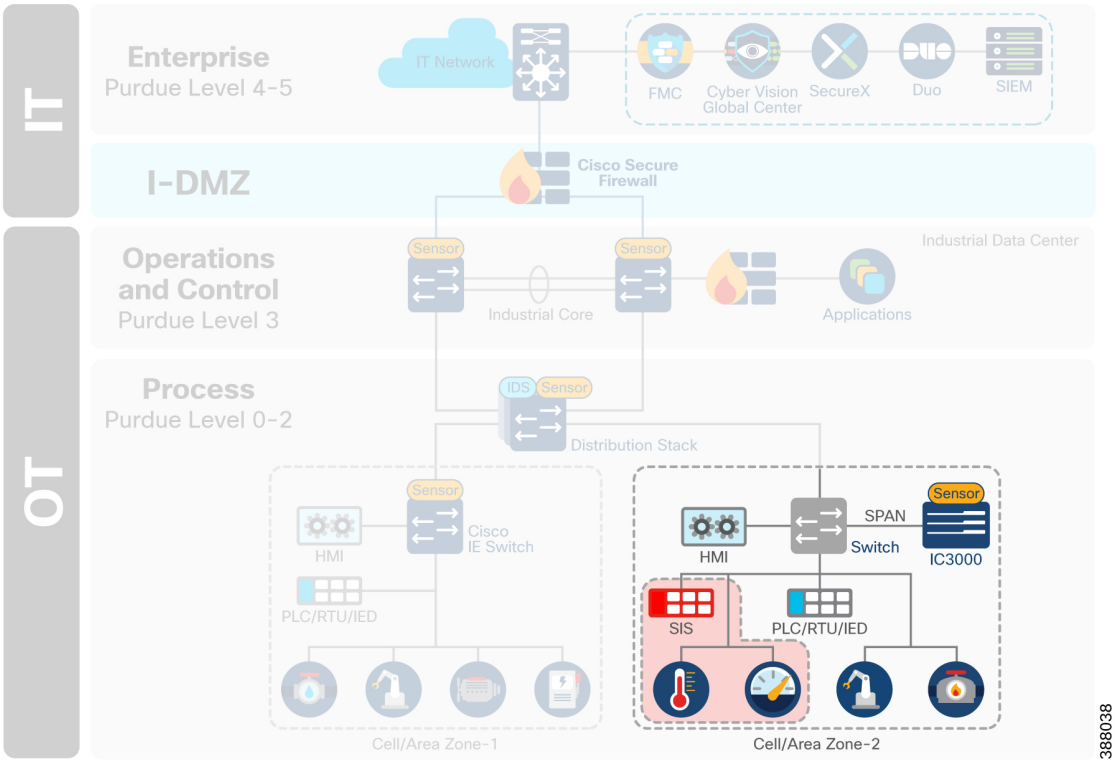

- **Remote Users:** Remote access is commonly granted to personas such as employees, partners and vendors for maintenance, process optimization and troubleshooting purposes. Remote access should be restricted to select devices on the plant floor for a limited amount of time.

Figure 12: Remote users need access to select devices, not a full zone

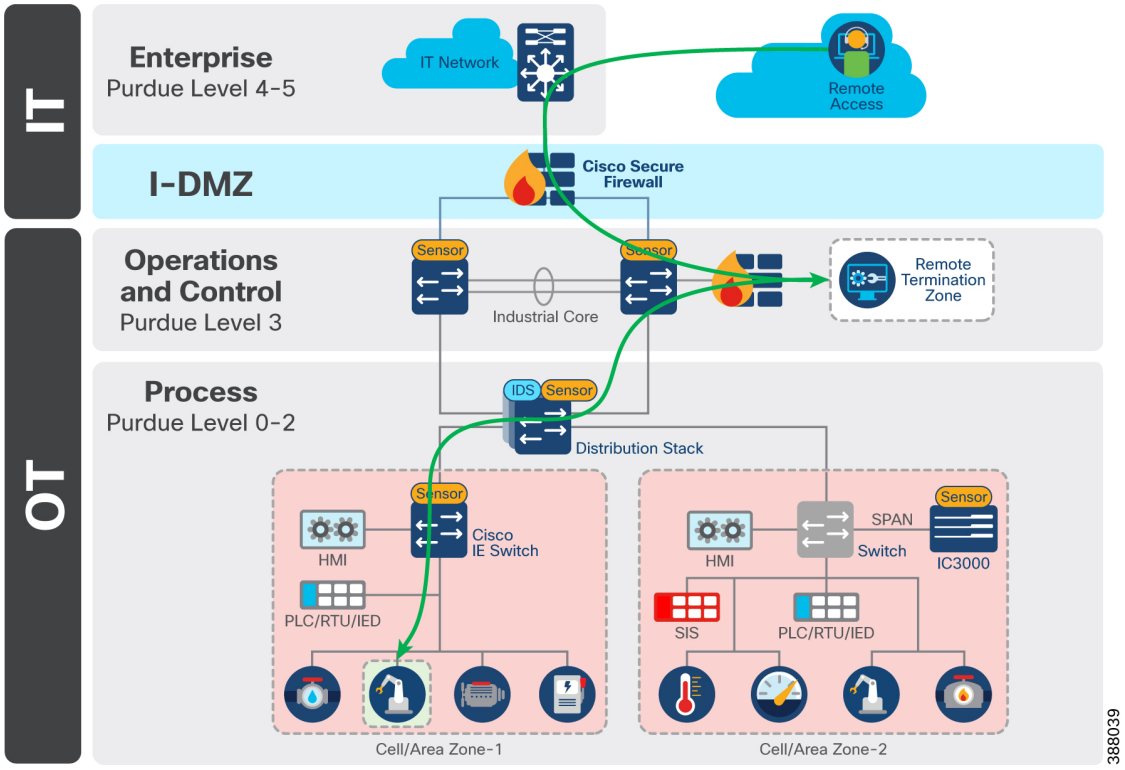

Industrial Security Journey

Industrial Security Journey

Addressing these issues and building a secure industrial network will not happen overnight. To help ensure success, Cisco promotes a phased approach in which each phase builds the foundation for the next, so that you can enhance your security posture at your own pace and demonstrate value to all stakeholders when embarking on this journey.

Figure 13: Industrial Security Journey

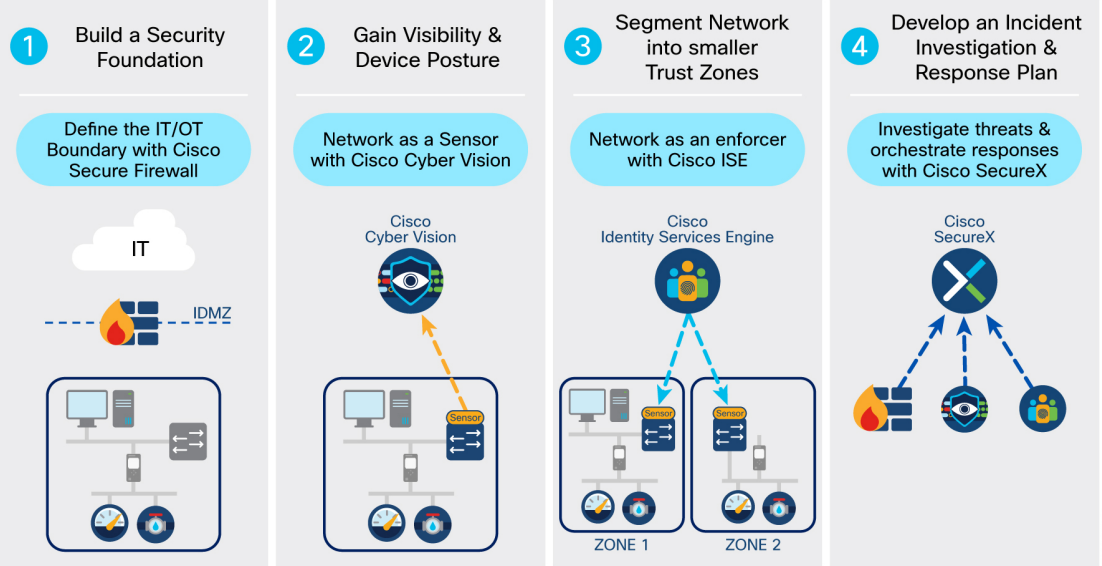

Building a Security Foundation

A solid and flexible network architecture is a key success criterion for robust and certified security. Poor network design creates a huge vulnerability and hinders the concepts of segmentation, extensibility, as well as the integration of cyber security controls and physical security measures.

Security considerations used in this guide are focused around three key networking areas: The Cell/Area Zone supporting the core IACS embedded in the production environment functional zones, the Operations and Control Zone supporting plant-wide applications and services, and the IDMZ providing key segmentation between production and enterprise systems.

More information on the Cisco Industrial Automation reference architecture can be found in the Cisco Solution Brief for Industrial Automation Networks.

Segmenting IT & OT Networks with the Industrial Demilitarized Zone

The first step in the journey to securing your industrial network is to restrict logical access to the OT network. A common deployment method is an Industrial Demilitarized Zone (IDMZ) network with firewalls to prevent network traffic from passing directly between the corporate and OT networks.

Figure 14: Industrial DMZ Architecture

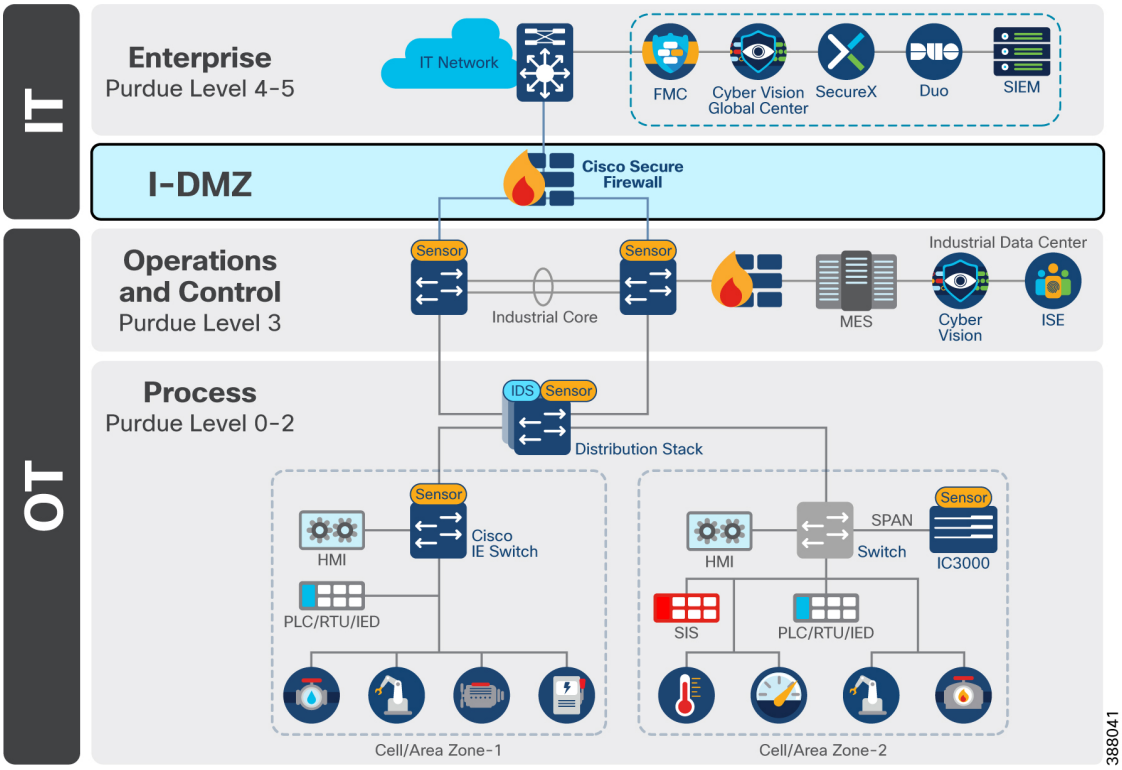

The IDMZ offers a network on which to place data and services to be shared between the Enterprise and Industrial Zones. The IDMZ doesn't allow direct communication between the Industrial and Enterprise Zones but meets the data and service sharing requirement. With the deployment of an IDMZ and Industrial Zone firewall, attacks and issues that arise in one zone cannot easily affect the other zone. In fact, by temporarily disabling the IDMZ and the firewalls, an IACS or IT network administrator can help to protect a zone until the situation is resolved in the other zone.

Cisco Secure Firewall brings distinctive threat-focused next-generation security services. The firewall provides stateful packet inspection of all traffic between the enterprise and OT network and enables intrusion prevention and deep packet inspection capabilities for inspecting application data between the zones designed to identify and potentially stop a variety of attacks. Cisco Secure Firewall is the first line of defense adversaries meet when attempting to breach the network and is the enforcement point for least privilege access for legitimate services to cross the border in a secure way.

Providing design guidance for the IDMZ is out of scope for this design guide but has been extensively covered in another guide. For more information on the IDMZ, see Securely Traversing IACS Data across the IDMZ Using Cisco Firepower Threat Defense.

Moving the IDMZ to the Cloud

Typically, IDMZ designs are architected and deployed at one facility, and replicated across each production site owned by the organization. One of the challenges with an exclusively on-site IDMZ is the limited ability to meet future demand in a world where the growth of Industrial IoT (IIoT) and IT/OT/cloud convergence requires new capabilities. It can also become challenging for operations staff to maintain IDMZ consistency across multiple sites and deliver consistent security policies.

A hybrid cloud IDMZ model can be an alternative. Like an IDMZ deployed on premises, it provides a holistic security strategy, with the benefit of shared resources and assets, allowing for a more repeatable and consistent architecture, as well as easing the operational overhead and complexity. A hybrid cloud IDMZ supports a regional operations center model, which is top of mind for some industrial organizations, especially those with a global footprint.

Design guidance for the hybrid cloud IDMZ will be added later as validation is still in the early stages of development. For an introductory insight into the architecture, see the Hybrid Cloud IDMZ white paper.

Don't Forget about the Hardware

If the hardware is not reliable, any security measures you take on the network and resources that run on that hardware cannot be relied upon. Securing the hardware should be considered fundamental to securing operations.

IEC62443-4-1 describes requirements for the secure development of products used to assemble IACS as well as maturity levels to set benchmarks for compliance. These requisites include requirement, management, design, coding guidelines, implementation, verification and validation, defect management, patch management and product end-of-life. All of these are essential to the security capabilities of a component and the underlying secure-by-design approach of the IACS solution. The overall focus is on continuous improvement in product development and release.

Cisco software and hardware products are developed according to the Cisco Secure Development Lifecycle (CSDL), which enforces a secure-by-design philosophy from product planning through end-of-life.

IEC62443-4-2 contains requirements for components necessary to provide the required security base for 62443-3 and higher levels. In this regard, the standard specifies security capabilities that enable hardware equipment to be integrated into a secure IACS deployment. Part 4-2 contains requirements for four types of components: software application, embedded device, host device, and network device. In essence, a secure IACS solution needs to be built based on secure components.

Various Cisco products have already achieved IEC62443-4-2 certification. In combination with a 62443-certified development process (CSDL), Cisco offers trustworthy communication products which are essential for IACS deployment in critical infrastructures.

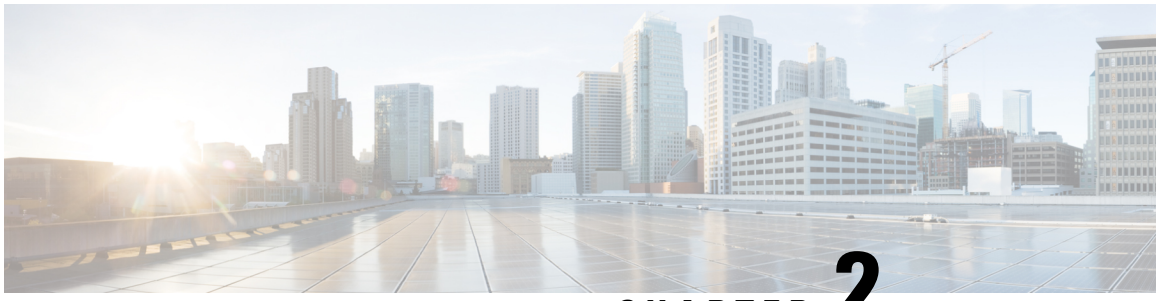

CHAPTER 2

Gain Asset Visibility and Device Posture

Gain Asset Visibility and Device Posture

After defining and securing the network perimeter, the second stage of our security journey is to gain visibility of all assets within the industrial network boundary. As industrial networks can be quite old, widely dispersed, and involve many contractors, operators often don't have an accurate inventory of what is on the network. Organizations may want to understand the normal state of the OT network as a prerequisite for implementing network security monitoring to help distinguish attacks from transient conditions or normal operations within the environment. Whether using a risk-based approach, functional model, or other organizing principles, grouping components into levels, tiers, or zones is a precursor activity before organizations can consider applying policy to protect and monitor communication between zones. Implementing network monitoring in a passive mode and analyzing the information to differentiate between known and unknown communication may be a necessary first step in implementing security policies.

Use Cases

OT visibility is a technology that all personas in OT environments can leverage. OT operators gain benefit of process level visibility to identify and troubleshoot assets residing on the plant floor. IT operators gain insight into device communication patterns to help inform policy and improve network efficiency. Security teams gain insight into device vulnerabilities and deviations from normal device behaviors.

For the purposes of this CVD, nine use cases / personas were identified that required securing. Asset visibility and device posture aids in securing these use cases by:

- **Identifying all assets and grouping them into zones**. It was stated that the Cell/Area Zone would be able to freely communicate within its own zone. Nevertheless, all assets must be identified within the zone to ensure that only intended devices reside in the zone, and critical vulnerabilities can be addressed so exploits cannot occur easily. Visibility also enables IT teams to view when new assets have been onboarded to the zone, or mobile assets have connected to a new location.

- **Visualizing data that flows through the conduits between zones**. While most traffic is contained within a given zone, interlocking PLCs require communication to cross zone boundaries. Before policy is implemented, visibility tools enable IT administrators to view existing dataflows and identify which flows should be removed and which need policy to maintain.
- **Give a clear view of which source data is coming in through external networks**. Network visibility gives a clear view of communication coming from external origins such as the IDMZ or remote access zones, enabling teams to see when a device is attempting unintended communication to external networks, or an unknown entity has breached externally accessible zones and is attempting to communicate deep into the OT network.

How to Gain Visibility into the OT Assets

As most of the communication in an IACS traverses the network (wired or wireless), the network infrastructure is in a good position to act as a sensor to provide visibility of the connected assets. Deep Packet Inspection (DPI) of the IACS communication is a key means to visibility. DPI decodes all communication flows and extracts message contents and packet headers, providing the visibility to understand what devices you need to secure, and the policies required to secure them. DPI allows you to gather device information such as the model, brand, part numbers, serial numbers, firmware and hardware versions, rack slot configurations, and more. It also allows you to understand what is being communicated over the network. For example, you can see if someone is attempting to upload new firmware into a device or trying to change the variables used to run the industrial process.

To achieve complete visibility, all network traffic must be inspected. It is important to note that in an industrial network, most traffic occurs behind a switch at the cell layer, because that is where the machine controllers are deployed. Very little traffic goes up to the central network.

When collecting network packets to perform DPI, security solution providers typically configure SPAN ports on network switches and employ one of three architectures:

- Send all traffic to a central server that performs DPI
- Deploy dedicated sensor appliances on each network switch
- Send traffic to dedicated sensor appliances deployed here and there on the network

While these approaches deliver network visibility, they also create new challenges. Configuring network switches to send traffic to a central server requires duplicating network flows. A new out-of-band network will generally be needed to transport this extra traffic, which can be complex and costly. Although this can be acceptable for a very small industrial site, this cannot be seriously considered in highly automated industries generating a lot of IACS traffic (such as manufacturing), or when devices are widely spread in locations with no or poor network connectivity (oil and gas pipelines, water or power distribution, etc.).

Connecting sensor appliances to network switches addresses the issues associated with duplicating network traffic. The appliance collects and analyzes network traffic locally and only sends data to a server for additional analysis. However, installing, managing, and maintaining dedicated hardware can quickly lead to cost and scalability issues. And because most industrial traffic is local, gaining full visibility requires deploying appliances on each switch on the network, raising cost and complexity to intolerable levels. Some technology providers attempt to address this problem by leveraging remote SPAN (RSPAN). RSPAN allows you to duplicate traffic from a switch that doesn't have a sensor appliance to a switch that has one.

Figure 15: OT visibility using a SPAN network

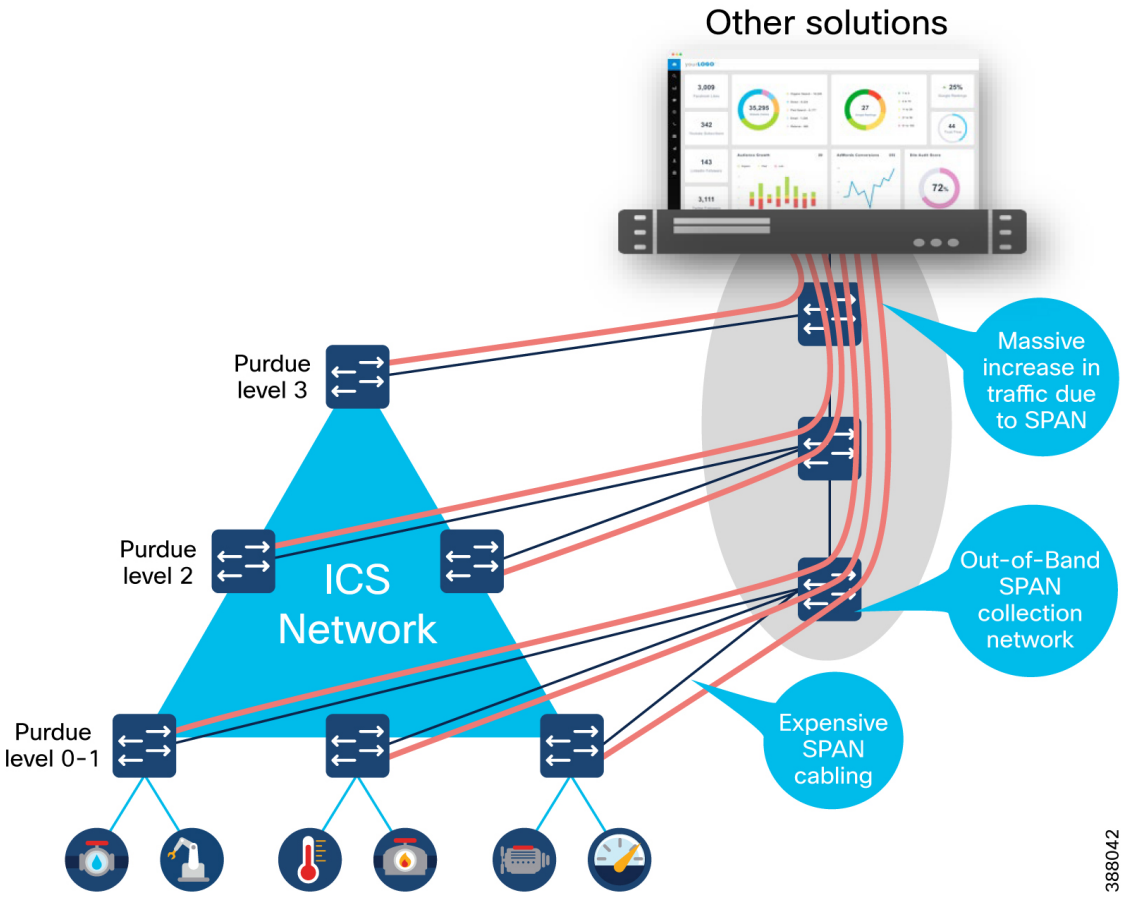

While this approach reduces the number of appliances required to provide full visibility, it still increases the amount of traffic going through the industrial network. Traffic is multiplied because you're duplicating traffic to SPAN it to a remote switch. And the more traffic on the network, the slower it becomes, resulting in jitter — often an unacceptable compromise in industrial networks where processes need to run faster and machines must be timely synchronized.

An Alternative to SPAN

There is a better way to achieve full network visibility: embed DPI capability into existing network hardware. An industrial-grade switch with native DPI capability eliminates the need to duplicate network flows and deploy additional appliances. Obtaining visibility and security functionality is simply a matter of activating a feature within the network switch, router, or gateway. Cost, traffic, and operational overhead are all minimized.

Figure 16: OT Visibility using Sensors embedded in the Switch Infrastructure

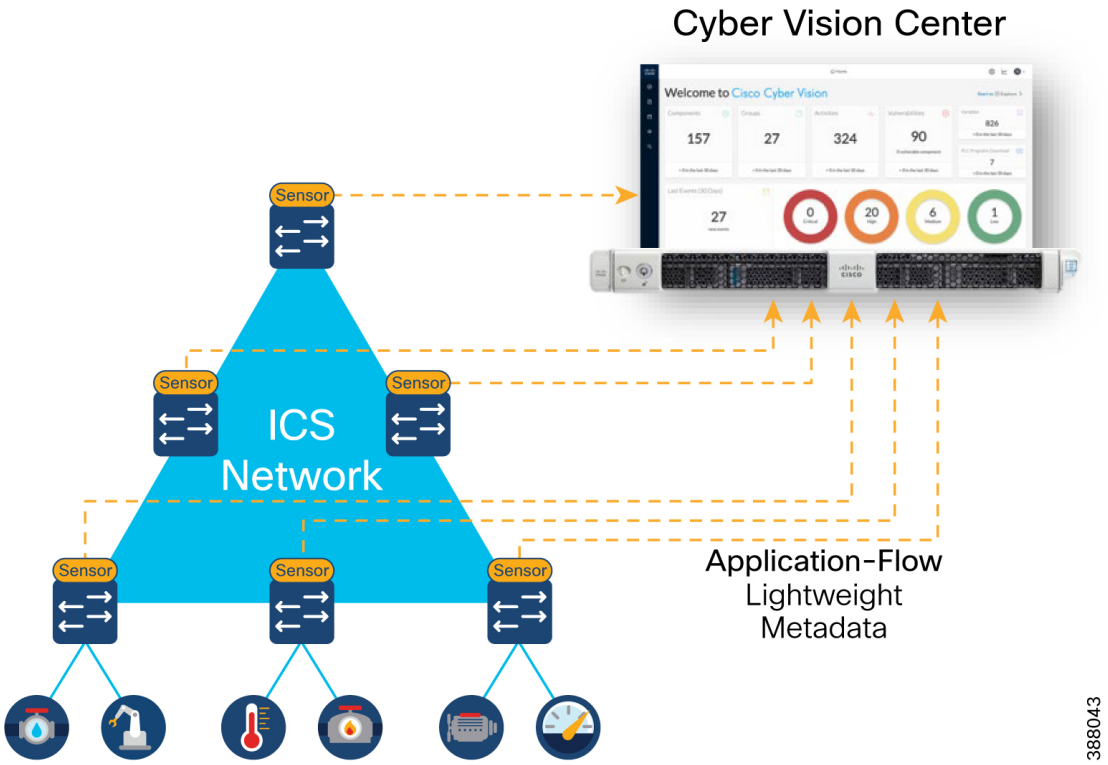

A DPI-enabled switch analyzes traffic locally to extract meaningful information. It only sends lightweight metadata to a central server, which runs the analytics and anomaly detection. That metadata represents about 3-5% of general traffic. The traffic is so lightweight, it can be transferred over the industrial network without causing congestion or requiring extra bandwidth. Embedding DPI in network equipment affords both IT and OT unique benefits. IT can leverage the existing network infrastructure to secure industrial operations without having to source, deploy, and manage additional hardware. Because these network elements see all industrial traffic, embedded sensors can provide analytical insights into every component of the industrial control systems. As a result, OT can obtain visibility into operations that it has never had before.

Active Discovery

The completeness of asset discovery is important for IACS networks to get a complete understanding of all the devices on the network and their associated security risks. For passive discovery to be effective, sensor placement is important and will be discussed later in the document. However, it is difficult to determine how much of the network has adequately been discovered as assets will only be seen as they cross the sensor. Gaining a complete picture takes time and can only determine information that is transmitted by the asset.

Active discovery is an on-demand mechanism for gaining asset visibility. By sending extremely precise and nondisruptive requests in the semantics of the specific IACS protocols, visibility gaps can be filled. However, there are some misconceptions regarding active discovery due to the many ways in which it can be implemented.

Active discovery causes unexpected crashes. The argument made by most vendors is that their solutions only use valid protocol commands supported by the industrial assets. These commands are similar to what the IACS vendor products use for asset management and are hence non-disruptive. In reality, the reason why old IACS devices are susceptible to crashes during active scanning is because they have limited processing power for network functions and get overwhelmed when repeated connection attempts are made for

communication. So, the reason for the crashes has less to do with valid or invalid commands being used but rather a factor of how many connection attempts is being made by the active discovery solution.

From a network hygiene standpoint, it is not uncommon to see industrial networks badly designed with all devices being addressed from a flat /16 IP subnet. Most IACS detection solution available in the market today are based on a centralized architecture where traffic mirroring (SPAN) is used to feed an appliance (or a software VM) located at Level-3 of the Purdue model that does the Passive Discovery.

When the bolt-on Active Discovery capability of these solutions initiate a scan from this central location, they need to cycle through a range of IP addresses within the scan range. Now, one of the first things that needs to happen to establish communication for Active Discovery is to resolve ARP. These ARP requests are seen by all devices within the flat network, and the processing of the barrage of ARP requests can overwhelm the networking stack on legacy IACS devices causing them to crash. While this is not the only reason for legacy devices crashing, it is quite often the primary cause.

In addition, in most multi-vendor IACS environments, centralized discovery solutions sitting at Level-3 of the Purdue model are not aware of the specific protocol being used at the Level 0-2 edge. This requires the scanning process to cycle through a range of IACS protocols (CIP, PROFINET, Modbus, etc.) until the device responds based on the protocol it supports. This results in unnecessary communication attempts that can also overwhelm the processing power of legacy devices causing disruption.

Centralized active discovery solutions cannot penetrate NAT boundaries. Industrial networks are usually built up of units like cells, zones, bays, etc. that are comprised of machines or control systems supplied by machine builders and system integrators. It is common practice for these machines especially in discrete manufacturing to be built in a standardized manner with IACS devices across machines configured in a cookie-cutter approach with repeating IP addresses. Consequently, industrial networks are rife with network address translation (NAT) being used to allow the operations and control systems located in the Level-3 to communicate with IACS devices sitting in the lower levels with duplicate IP addresses.

When it comes to address translation only a small fraction of IACS devices (like PLC, HMI, RTU, etc.) communicate with the site operations layer, and only those devices' IP addresses are translated at the NAT device. The implication of this is that centralized Active Discovery solutions cannot communicate with the vast majority of IACS devices (like IO, drives, safety controllers, relays, IED) sitting below the NAT boundary whose IP addresses are not translated. In the auto manufacturing industry as an example, it is typical for less than 17% of devices in level 0-2 to be visible to a centralized Active Discovery solution. This results in an 83% gap in visibility!

It is recommended that networks use a hybrid approach of active and passive discovery to gain an accurate insight into their OT network.

Intrusion Detection / Prevention Systems Intrusion sensors are systems that detect activity that can compromise the Confidentiality, Integrity or Availability (CIA) of information resources, processing, or systems. An Intrusion Detection System (IDS) has the ability to analyze traffic from the data link layer to the application layer to identify things such as network attacks, the presence of malware, and server misconfigurations.

An Intrusion Prevention System (IPS) can identify, stop, and block attacks that would normally pass through a traditional firewall device. When traffic comes in through an interface on an IPS, if that traffic matches an IPS signature/rule, then that traffic can be dropped by the IPS. The essential difference between an IDS and an IPS is that an IPS can respond immediately and prevent possible malicious traffic from passing. An IDS produces alerts when suspicious traffic is seen but is not responsible for mitigating the threat.

The advantage of IDS deployments is that they create no risk of taking down the IACS. This advantage may be due to "false positives," where the IDS detects a condition that it believes to be an anomaly or attack, when in fact it is business-critical traffic. Because IDS systems are typically not inline, they have no effect on network performance statistics such as propagation delay and jitter (variations in delay). Another risk of IPS

solutions is that a catastrophic failure of the IPS system may cause a complete lack of connectivity. This type of failure is of less concern if solutions are designed with ample redundancy and without single points of failure.

It is recommended that OT networks adopt a hybrid IDS/IPS deployment, where IDS is deployed in the operational zone of the network for security alerting and then deploy an IPS north of the critical zone (for example at the Industrial Data Center) where a false positive would not stop plant operations.

Vulnerability Assessment and Managing Risk

A **vulnerability** is a weakness in a system or its design that can be exploited by a threat. Vulnerabilities are sometimes found in the protocols themselves, as in the case of some security weaknesses in TCP/IP. Often the vulnerabilities are located in operating systems and applications.

A **threat** is any potential danger to assets. A threat is realized when someone or something identifies a specific vulnerability and exploits it, creating exposure. If the vulnerability exists theoretically, but has not yet been exploited, the threat is latent and has not been realized. The entity that takes advantage of a vulnerability is known as the threat agent or threat vector.

A **countermeasure** is a safeguard that mitigates a potential risk. A countermeasure mitigates risk by either eliminating or reducing a vulnerability, or by reducing the likelihood that a threat agent can successfully exploit the risk.

Risk is a function of the likelihood of a given threat source exercising a particular potential vulnerability, and the resulting impact of that adverse event on the organization.

Threat x Vulnerabilities x Impact = Risk

Risk management is the process that balances the operational and economic costs of protective measures and the achieved gains in mission capability by protecting assets and data that support their organizations' missions. For example, many people decide to have home security systems and pay a monthly fee to a service provider to monitor the system for increased protection of their property. Presumably, the homeowners have weighed the cost of system installation and monitoring against the value of their household goods and their family's safety priority. Risk limitation limits a company's risk exposure by taking some action. It is a strategy employing a bit of risk acceptance along with a bit of risk avoidance. It is the most commonly used risk mitigation strategy.

Vulnerability Assessment

The objective of a vulnerability assessment is to ensure that the network and the information systems are tested for security vulnerabilities in a consistent and repeatable manner. Security vulnerabilities will continue to be discovered in technology products and services. These vulnerabilities, regardless of whether they are caused by an unintentional software bug or by design (such as a default administrative password), can be used by malicious persons to compromise the confidentiality, availability, or integrity of your infrastructure.

Hardware and software vendors typically provide software fixes when they announce the vulnerabilities in their products. When there is no fix available, vendors typically provide a workaround or mitigation. There is usually a time period between the announcement of a security vulnerability in a particular technology and the availability of an attack method (an exploit). Within this time period, system administrators should take action to protect their systems against an attack because at this point, the public knows that a flaw exists, but attackers are still trying to find a way to take advantage of that vulnerability. Unfortunately, the vulnerability-to-exploit time period has been steadily decreasing.

With the large quantity of new vulnerabilities from numerous vendors, it can be overwhelming to track all the vulnerabilities. How can the security team analyze any single vulnerability and determine its relevance to

the specific technology architecture? The solution is to have a good process to determine which ones are relevant to your organization.

CVSS Scores

The Common Vulnerability Scoring System (CVSS) provides a way to capture the principal characteristics of a vulnerability and provides a better understanding of the risk that is posed by each vulnerability. CVSS is a free and open industry standard for assessing the severity of computer system security vulnerabilities. CVSS attempts to assign severity scores to vulnerabilities, allowing responders to prioritize responses and resources according to threat. Scores are calculated based on a formula utilizing several metrics that approximate ease of exploit and its impact. Scores range from 0 to 10, with 10 being the most severe.

CVSS provides a standard way to assess and score security vulnerabilities. CVSS analyzes the scope of a vulnerability and identifies the privileges that an attacker needs to exploit it. CVSS allows vendors to better analyze the impact of security vulnerabilities and more clearly define the level of urgency that is required to respond to the vulnerability. While many analysts use only the CVSS base score for determining severity, temporal and environmental scores also exist, and factoring in the likelihood and the criticality to a given network environment.

Cisco Cyber Vision

Cisco Cyber Vision is built on a unique edge architecture consisting of multiple sensor devices that perform deep packet inspection, protocol analysis, and intrusion detection within your industrial network and an aggregation platform known as Cyber Vision Center. The Cyber Vision Center stores data coming from the sensors and provides the user interface, analytics, behavioral analysis, reporting, API, and more. It may be run on a hardware appliance or as a virtual machine.

Components

Cisco Cyber Vision Center can be deployed as a software or hardware appliance depending on your network requirements. Consider the number of sensors, components, and flows to decide the appropriate installation. At the time of writing this guide, a single Cyber Vision Center can support 150 sensors, 50,000 components, and 8 million flows. For the most up to date numbers see the Platform Support page.

For deployments that are too large for a single instance of Cyber Vision Center to handle, or for organizations who wish to aggregate multiple sites into a single dashboard view, a **Cyber Vision Global Center** instance can aggregate up to 20 local Cyber Vision Centers. Cyber Vision Global Center is used for security monitoring across multiple sites, providing a consolidated view of components, vulnerabilities, and events. Nevertheless, sensor operation and management activities can be done only on instances of Cyber Vision Center associated with the sensor.

Cyber Vision sensors passively capture and decode network traffic using DPI of industrial control protocols. Cyber Vision sensors are embedded in select Cisco networking equipment, so you don't have to deploy dedicated appliances or build an out-of-band SPAN collection network. Since Cyber Vision sensors decode industrial network traffic at the edge, they only send lightweight metadata to the Cyber Vision Center, only adding from 2% to 5% load to your industrial network.

Note: Cyber Vision also supports an out-of-band sensor network for environments that require it.

Cyber Vision sensors also have the capability to do active discovery. These active discovery requests originate from the sensor, deep into the IACS network, so these messages are not blocked by firewalls or NAT boundaries.

Key Features

Comprehensive Visibility: Cyber Vision leverages a unique combination of passive and active discovery to identify all your assets, their characteristics, and their communications. The Cisco Cyber Vision unique edge computing architecture embeds security monitoring components within our industrial network equipment. There is no need to source dedicated appliances and think about how to install them. There is no need to build an out-of-band network to send industrial network flows to a central security platform. Cyber Vision enables the industrial network to collect the information required to provide comprehensive visibility, analytics, and threat detection.

Security Posture: Cisco Cyber Vision combines protocol analysis, intrusion detection, vulnerability detection, and behavioral analysis to help you understand your security posture. It automatically calculates risk scores for each component, device and any specific parts of your operations to highlight critical issues so you can prioritize what needs to be fixed. Each score comes with guidance on how to reduce your exposure so you can be proactive and build an improvement process to address risks.

Operational Insights: Cisco Cyber Vision automatically uncovers the smallest details of the production infrastructure: vendor references, firmware and hardware versions, serial numbers, rack slot configuration, etc. It identifies asset relationships, communication patterns, and more. Information is shown in various types of maps, tables, and reports. Cisco Cyber Vision gives OT engineers real-time insight into the actual status of industrial processes, such as unexpected variable changes or controller modifications, so they can quickly troubleshoot production issues and maintain uptime. Cyber experts can easily dive into all this data to investigate security events. Chief information security officers have all the necessary information to document incident reports and drive regulatory compliance.

Incident Investigation and Response: SecureX Threat Response is a security investigation and incident response application. It simplifies threat hunting and incident response by accelerating detection, investigation, and remediation of threats. The threat response application provides your security investigations with context and enrichment by connecting your Cisco security solutions (across endpoint, network, and cloud) and integrating with third-party tools, all in a single console. Abnormal behavior seen in Cyber Vision can be sent to SecureX for further analysis and context from the other security tools deployed on the network such as Cisco Secure Endpoint, Secure Firewall, Umbrella and more. The SecureX ribbon on the Cyber Vision user interface makes it even easier to create a case and launch investigations.

Snort IDS: Cyber Vision integrates the Snort IDS engine in select platforms leveraging Talos subscription rules to detect known and emerging threats such as malware or malicious traffic.

For more information on Cisco Cyber Vision see the Cisco Cyber Vision Datasheet.

Cyber Vision Design Considerations

Cyber Vision Center

The architectural recommendation is to deploy Cyber Vision Center in the Industrial Zone. Cisco Cyber Vision connects to the sensor(s) in the cell/area zone and applications on the industrial zone such as NTP and optionally DNS and ISE. The following figure depicts the communication flows from Cisco Cyber Vision center used in this design guide.

Figure 17: Cyber Vision Communication Flows

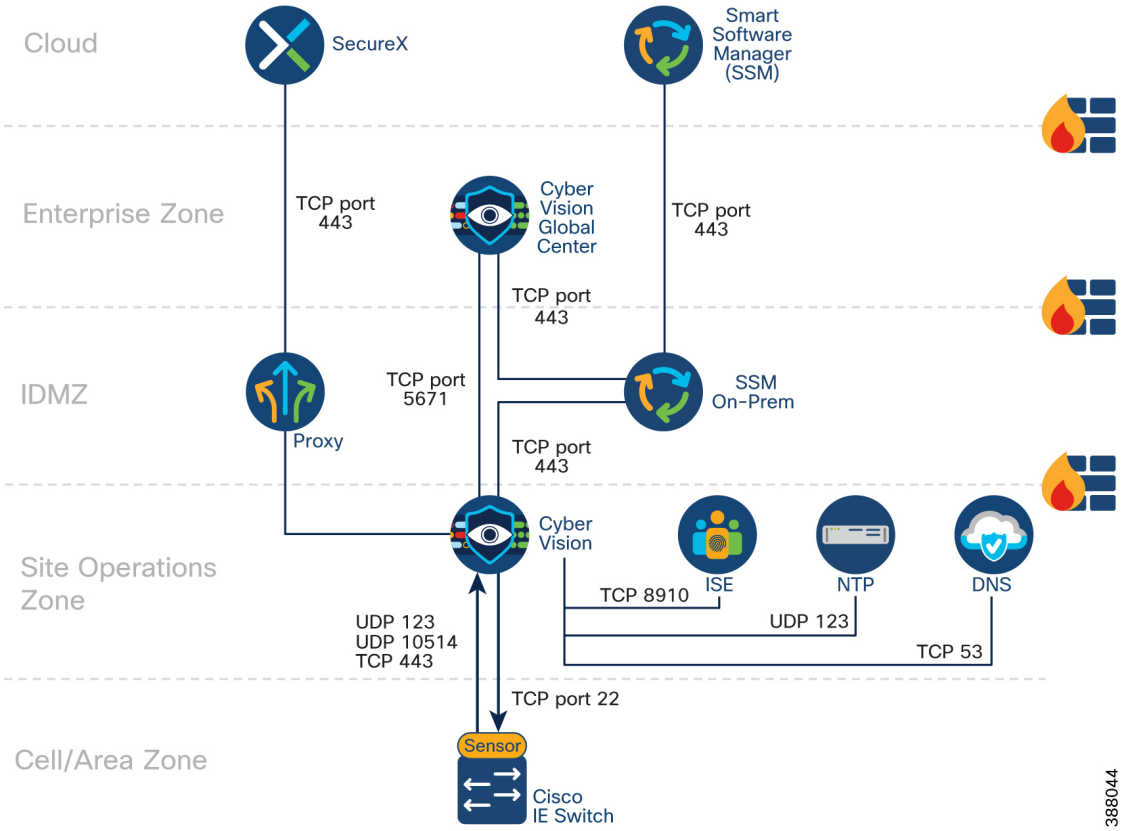

Note: Cisco Cyber Vision Center can operate without any connectivity leaving the industrial zone. The flows in the diagram that meet this condition are optional and their purpose will be explained in this guide.

In Cisco Cyber Vision, the administrator network interface gives access to the graphical user interface (GUI) and the collection network interface connects the Center to the sensors. Ethernet interfaces are allocated in the following way:

- Administration network interface (eth0) gives access to the user interface (GUI or API), to the CLI through SSH and is used for communication with other systems (syslog collector or SIEM, pxGrid, etc.)
- Collection network interface (eth1) connects the Center to the sensors

The Center (physical or virtual appliance) has two preconfigured interfaces—eth0 and eth1—that are allocated to the admin and collection networks respectively by default.

However, if the admin and collection network share the same local area network (LAN), the Center must be configured to use a single interface. In this case the admin and collection interface should share a single IP address on eth0, and eth1 should be reserved as a collection interface for DPI on the Center.

Figure 18: Cyber Vision Center Deployment Modes

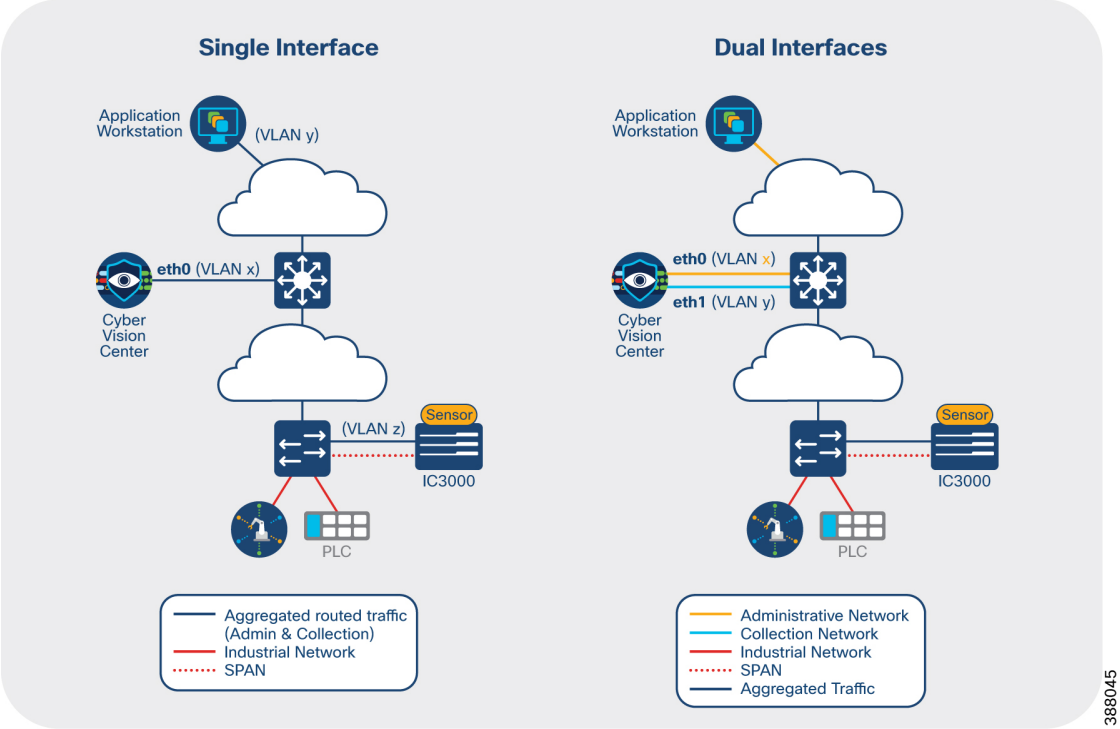

Cisco Cyber Vision Global Center requires only one interface for management and communication with Cisco Cyber Vision Center instances. It uses TCP port 5671 for synchronization and updates to the Center. This port should be proxied in the IDMZ or enabled in the IDMZ firewall to ease communication.

Note: Cisco Cyber Vision Center does not require internet connectivity nor Global Center connectivity to operate. In instances where Cyber Vision Center is not connected to the Internet, upgrades need to be downloaded from Cisco.com and manually uploaded in the appliance.

Cyber Vision Sensor Options

The sensors are supported on the platforms listed in the table below.

Figure 19: Supported Cyber Vision Sensor Platforms

Sensor Type	Platforms Supported
Integrated Network Sensor	Cisco Catalyst IE3400 Rugged Series Switch
	Cisco Catalyst IE3400 Heavy Duty Series Switch
	Cisco Catalyst IE3300 10G Rugged Series Switch
	Cisco Catalyst IR1101 Rugged Series Router
	Cisco Catalyst IR8300 Rugged Series Router
	Cisco Catalyst 9300 Series Switch
	Cisco Catalyst 9400 Series Switch
	Cisco Catalyst IE9300 Rugged Series Switch
Hardware Sensor Appliances	Cisco IC3000 Industrial Compute Gateway

388046

In this design guide, the Catalyst IE3400 is deployed within Cell/Area Zones and the Catalyst 9300 is used as the distribution switch. For the most up to date support information visit the Cisco Cyber Vision Platform Support page.

Effective Sensor Deployment The effectiveness of the Cisco Cyber Vision solution depends on effectively capturing traffic, so deciding the correct location for the sensor(s) in the network is critical.

A sensor deployed on the distribution switch, such as the Catalyst 9300, will capture flows that leave the Cell/Area Zones. This deployment option is helpful for building your macro-segmentation policies (to be discussed later in the document) as you will gain a clear understanding of the zone-to-zone communication patterns. However, significant value of Cyber Vision is lost when it is deployed only on the distribution infrastructure. None of the intra-zone communication traffic would be seen, resulting in missing many devices and the most important communication flows in industrial automation networks.

Visibility inside a Cell/Area Zone

To gain visibility on the cell area zone, the recommended option is to deploy the network sensor on the industrial switches. A sensor is deployed at the edge to capture flows for end devices. Deploying network/hardware sensor at a switch where a controller is attached is an ideal choice to monitor the traffic because all the I/O devices respond to the poll requests initiated by the controller. Note that flows that do not traverse the network sensor will not be visible on Cyber Vision Center. To increase coverage, consider the following options:

- **Dedicated sensor per switch**: to capture all traffic in the Cell/Area Zone, a sensor can be deployed at every switch, resulting in none of the flows being missed. Embedding sensors in the network is the only way to capture all the data in your network at scale
- **Dedicated sensor in aggregation switch**: to capture intra-zone communication between two small sub segments of a Cell/Area Zone, a sensor can be deployed at the aggregation point to capture traffic that crosses the sub-segments
- **Enable SPAN**: deploy a single out-of-band sensor and SPAN the traffic either from all or from select switches, corresponding to option 1 or 2 from this list. This model requires additional cabling from every device to the out-of-band sensor and also a free switch port must be available on those switches from which you want to SPAN traffic.

Note: There are no licensing implications for deploying sensors at every possible location. Cyber Vision licensing is based on the number of endpoints in which it detects and adds additional value to. A sensor can be deployed on every compatible switch in the network.

Visibility for flows leaving a Cell/Area Zone

A sensor deployed on the distribution switch, such as the Catalyst 9300, will capture flows that leave the Cell/Area Zones. This deployment option is helpful to understand zone-to-zone/north-south communication patterns. Keep in mind that this option is not a replacement for sensors on the cell/area zone since few of the intra-zone communication traffic would be seen, resulting in missing the most important communication flows in industrial automation networks.

Caution: Be careful when collecting data at higher levels (distribution level), especially if Internet traffic is being monitored. Monitoring Internet flows in addition to traffic on the industrial network will significantly increase the number of devices (and components) present in the Center database.

Ring Topology Considerations

Visibility of flows in a ring may change depending on sensor positioning and the active traffic path. The following figure illustrates a Resilient Ethernet Protocol (REP) ring with two flows that may not be captured by a sensor. The first flow, between the engineering workstation and IO will never be captured because there are no sensors in the path. The asset will be identified when communicating to the PLC, but if the intent is to capture all communication on the network, beyond just asset identification, sensors need to be placed to capture this information. The second flow, between the application workstation and the PLC may be captured depending on the alternate port configuration. If the traffic navigates the ring travelling anti-clockwise from the distribution switch, a sensor will be in the path. However, if a link fails, there will be no sensors on the path as the data travels clockwise from the distribution switch.

Figure 20: Visibility considerations in a ring topology

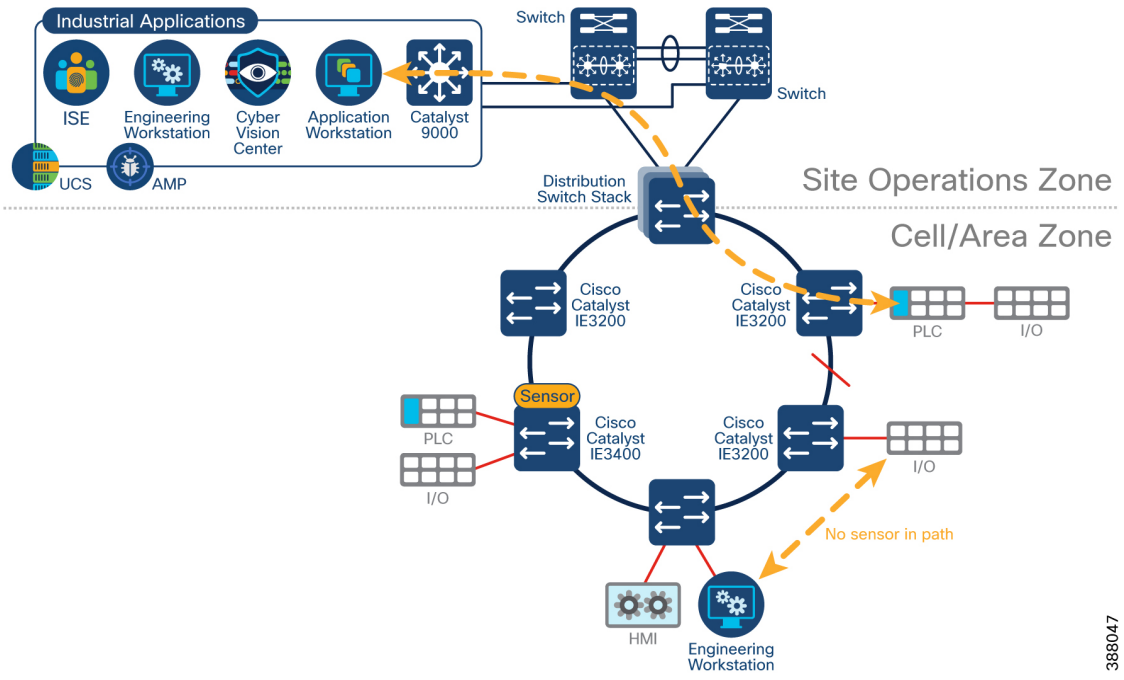

The recommendation is to always install a sensor at the top of a ring. At a minimum, all traffic leaving and entering the zone should be captured, with additional sensors placed within the ring depending on desired visibility levels.

Brownfield Deployment Considerations

If your current industrial network does not have any switches capable of natively running the Cyber Vision sensor, at least one needs to be introduced. To collect the traffic, enable SPAN on switches to an out-of-band monitoring switch that will aggregate traffic andsend it to a sensor such as an IE3400, the IC3000, or the Cisco Cyber Vision Center itself. In this model, additional cabling is required from every switch to the SPAN aggregation point.

Figure 21: Visibility in Brownfield Deployments

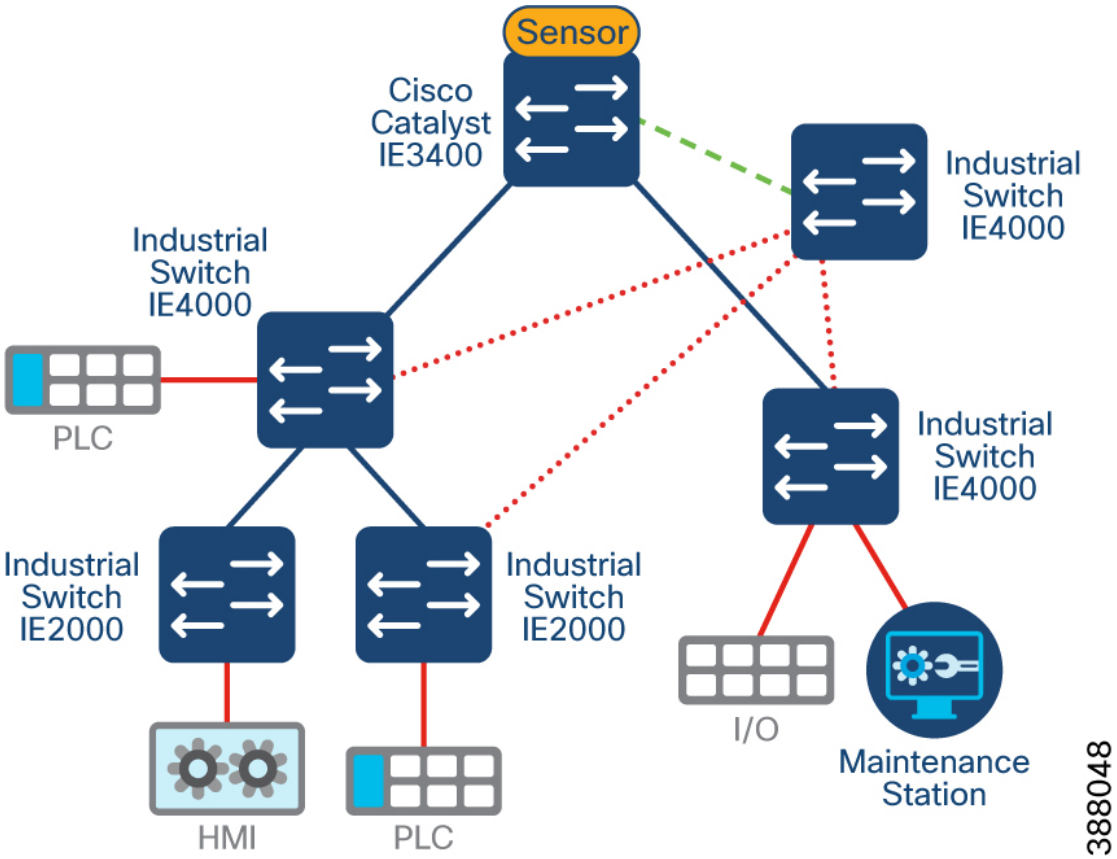

Sensor Considerations

When deploying Cisco Cyber Vision Sensors in the network, the following should be taken into consideration:

- Cisco Cyber Vision sensors are installed as an IOx application. IOx is included with essentials and advanced license of Cisco switches.
- IOx applications need an SD card (Industrial Ethernet switches) or SSD Disk (Catalyst 9300) to be installed. These parts are optional on the switch ordering configuration.
  - Industrial switches (Catalyst IE3400 and Catalyst IE3300 10G) require an SD Card of at least 4GB. SD card should be procured by Cisco to guarantee functionality.
  - Catalyst 9300/9400 switches require an SSD of at least 120GB.

- Sensors need an IP address to communicate with the Cisco Cyber Vision Center (collection interface). For network sensors deployed in IOx, this IP address needs to be different from other IP addresses on the switch. Although it can belong to any VLAN on the switch, it is recommended that the IP address is assigned on the management network. We recommend that the sensor IP address as well as other IP addresses for network management are not NAT'd.
- The sensor also needs a capture interface to reach the monitor session in the switch. This has local significance only, so VLAN used for RSPAN to the sensor should be private to the switch.
- The following ports are needed for communication between Cisco Cyber Vision Center and Cisco Cyber Vision Sensor:
  - From Cisco Cyber Vision Sensor to Cisco Cyber Vision Center
    - NTP (UDP port 123)
    - TLS 1.2 (TCP port 443)
    - Syslog (UDP port 10514)
    - AMQPS (TCP port 5671)
  - From Cisco Cyber Vision Center to Cisco Cyber Vision Sensor
    - SSH (TCP port 22)
    - Network sensor installation (TCP port 443)
    - Hardware sensor installation (TCP port 8443)
- It is possible to install a sensor using CLI, local device manager, or Sensor Management Extension on Cisco Cyber Vision Center. The first two options require getting a provisioning file from the center and copying it to the switch in order to complete installation. When using the Sensor Management Extension, the center connects to the switch directly and provisions the sensor. **Therefore, it is recommended to use Sensor Management Extension on Cisco Cyber Vision Center to simplify sensor installation**.
- If multiple Cisco Cyber Vision sensors discover the same device, Cisco Cyber Vision center combines the information into a single component.
- For networks with devices behind a NAT, IP addresses captured by the sensor will depend on the location of the sensor. If the capture is done before the traffic is translated, Cisco Cyber Vision will show the private IP of the device. If the sensor is installed on the traffic path after translation is done, Cisco Cyber Vision will show the translated IP address. In case of multiple capture points, it is possible to see a component for the private IP and a component for the translated IP on Cisco Cyber Vision Center, in other words, duplicate devices in the inventory.

Cyber Vision Active Discovery

With Cyber Vision, active discovery is initiated by the Cyber Vision Sensor embedded in the Cisco IE switches, that are distributed at the edge of the industrial network. The solution does much more than just distributing the initiator of the discovery. The Active Discovery is a closed-loop system between the Passive and the Active Discovery components. It works by the Passive Discovery first listening to the traffic on the network and then informing the Active Discovery component on which protocols are present on that section of the network. The Active Discovery component then initiates a broadcast hello request in the semantics of specific IACS protocol at play, and the Passive Discovery component decodes the response from the IACS devices.

When needed the Active Discovery component may initiate a unicast command to collect further information from the discovered devices.

Cyber Vision active discovery is non-disruptive. The fact that the Passive and Active components are embedded on the switches at the very point where the IACS devices connect to the network enables Cyber Vision discovery to be extremely precise and non-disruptive. Cisco Cyber Vision does not scan the network, instead it sends hello packets to devices for selected industrial protocols. There is no longer a need to enter IP scan ranges nor is there a need to guess which protocol is being used on a specific machine or process at the edge of the network. The intelligence built into the closed-loop system automates the Active Discovery. The user simply has to enable Active Discovery and has full control to activate the capability on a per switch basis if needed and the ability to configure the frequency which it executes.

Cyber Vision Active Discovery is not handicapped by the presence of NAT. Cisco recognizes the need for NAT in industrial networks and simplifies the process by providing L2 NAT (mapping between inside and outside IPs bound to MAC address) capability at line rate on the Cisco IE switches. This eliminates the need to additional L3 NAT devices. But regardless of whether L2 or L3 NAT is used, by virtue of the Passive and Active components of the Cyber Vision Sensor being embedded in the IE switches, the Active Discovery is distributed and is initiated from below the NAT layer, and results in 100% visibility of the IACS devices on the industrial network.

Vulnerability Assessment in Cyber Vision

Vulnerabilities are detected in Cisco Cyber Vision thanks to rules stored in a Cyber Vision Knowledge Database (DB). These rules are sourced from several CERTs (Computer Emergency Response Team), manufacturers, and partner manufacturers. Technically, vulnerabilities are generated from the correlation of the Knowledge DB rules and normalized device and component properties. A vulnerability is detected when a device or a component matches a knowledge DB rule.

Figure 22: Cyber Vision Vulnerability Dashboard

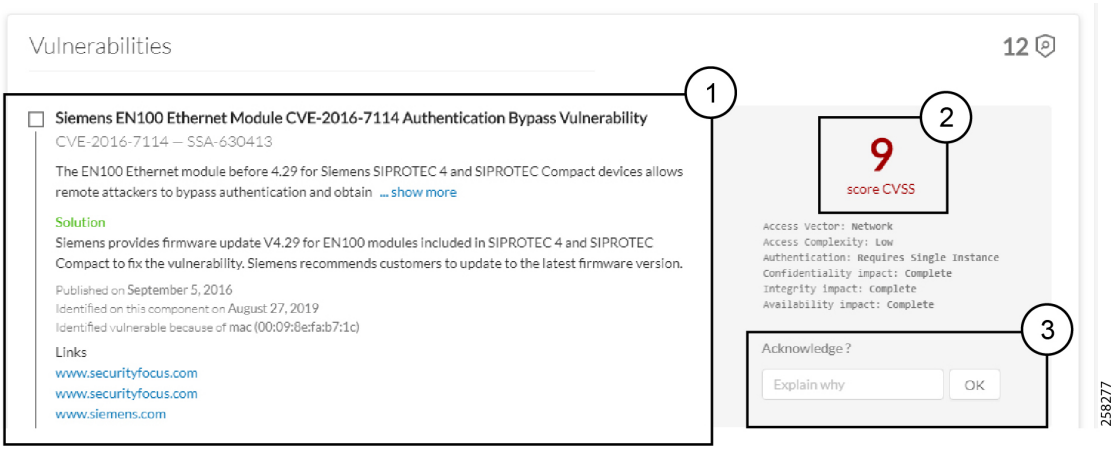

Information displayed about vulnerabilities **(1)** includes the vulnerability type and reference, possible consequences, and solutions or actions to take on the network. Most of the time though, it is enough to upgrade the device firmware. Some links to the manufacturer website are also available for more details on the vulnerability.

A score reports the severity of the vulnerability **(2)**. This score is calculated upon criteria from the Common Vulnerability Scoring System or CVSS. Criteria are for example the ease of attack, its impacts, the importance of the component on the network, and whether actions can be taken remotely or not. The score can go from 0 to 10, with 10 being the most critical score.

You also have the option to acknowledge a vulnerability **(3)** if you do not want to be notified anymore about it. This is used for example when a PLC is detected as vulnerable, but a security policy has been defined to protect against it. The vulnerability is therefore mitigated. An acknowledgment can be canceled at any time. Vulnerabilities acknowledgment/cancelation is accessible to the Admin, Product and Operator users only.

Cyber Vision Risk Score

A risk score is an indicator of the good health and criticality level of a device, on a scale from 0 to 100. It has a color code associated to the level of risk:

Figure 23: Cyber Vision Risk Score by color

Score	Color	Risk Level
0 - 39	Green	Low
40 - 69	Orange	Medium
70 - 100	Red	High

388049

The risk score is meant to help the user easily identify which vulnerable devices are the most critical to mitigate within the overall network. It provides limited and simple information on the cybersecurity of the monitored system. It is intended as a first step in security management to take actions by showing the causes of high scores and providing solutions to reduce them. The goal is to minimize and keep risk scores as low as possible. The solutions proposed can be to:

- Patch a device to reduce the surface of attack
- Remove vulnerabilities
- Update firmware
- Remove unsafe protocols whenever possible (FTP, TFTP, Telnet, etc.)
- Create an access control policy
- Limit communications with external IP addresses

In addition, it is necessary to define the importance of the devices in your system by grouping devices and setting an industrial impact. Thereby, increasing or decreasing the risk score, which will allow you to focus on most critical devices.

The Cyber Vision risk score is computed as follows:

Risk = Impact x Likelihood

Impact answers the question; What is the device “criticality”, that is, what is its impact on the operation? Does it control a small, non-significant part of the network, or does it control a large critical part of the network? To do so, the impact depends on the device tags assigned by Cyber Vision. Is the device a simple IO device that controls a limited portion of the system, or is it a SCADA that controls the entire factory? These will obviously not have the same impact if they are compromised.

Note: A Cyber Vision user can influence the device impact by moving it into a group and setting the group industrial impact (from very low to very high). By default, Cyber Vision may decide the impact a device has on your network is small, because it only communicates with a handful of other devices. However, if you as an administrator decide that these groups of assets are highly critical, the risk score will change based on this manually entered information.

Likelihood answers the question: What is the likelihood of this device being compromised? It depends on:

- Device Activities, or more precisely activity tags. Some protocols are less secure than others. For example, telnet is less secure than SSH.
- The exposure of the device communicating with an external IP subnet.
- Device vulnerabilities, considering CVSS scoring.

Visualize Assets and Flows using Cyber Vision Groups

The first thing to do when using Cisco Cyber Vision is to organize components in a meaningful way. OT networks can contain thousands of components and visibility of these devices can be overwhelming. It is recommended to use Cyber Vision Groups to organize components according to industrial processes or areas. Furthermore, each group can be assigned an industrial impact rating which will have a direct impact on the risk score. Benefits of using groups include:

- Groups can be used as a filter when building a Preset. This allows you to monitor a specific production process or area of the plant.
- Groups simplify network map visualization by aggregating the devices and activities on the Map view. Aggregated activities are called Conduits. The following figure shows a network map for a specific process and the communication conduits.
- Groups identify inter cell/process flows by showing Conduits leaving the area on the Preset Map.
- Groups provide context to ISE for profiling of devices. More information can be found in the segmentation chapter of this design guide.

Figure 24: Cyber Vision Network Maps with Device Groups

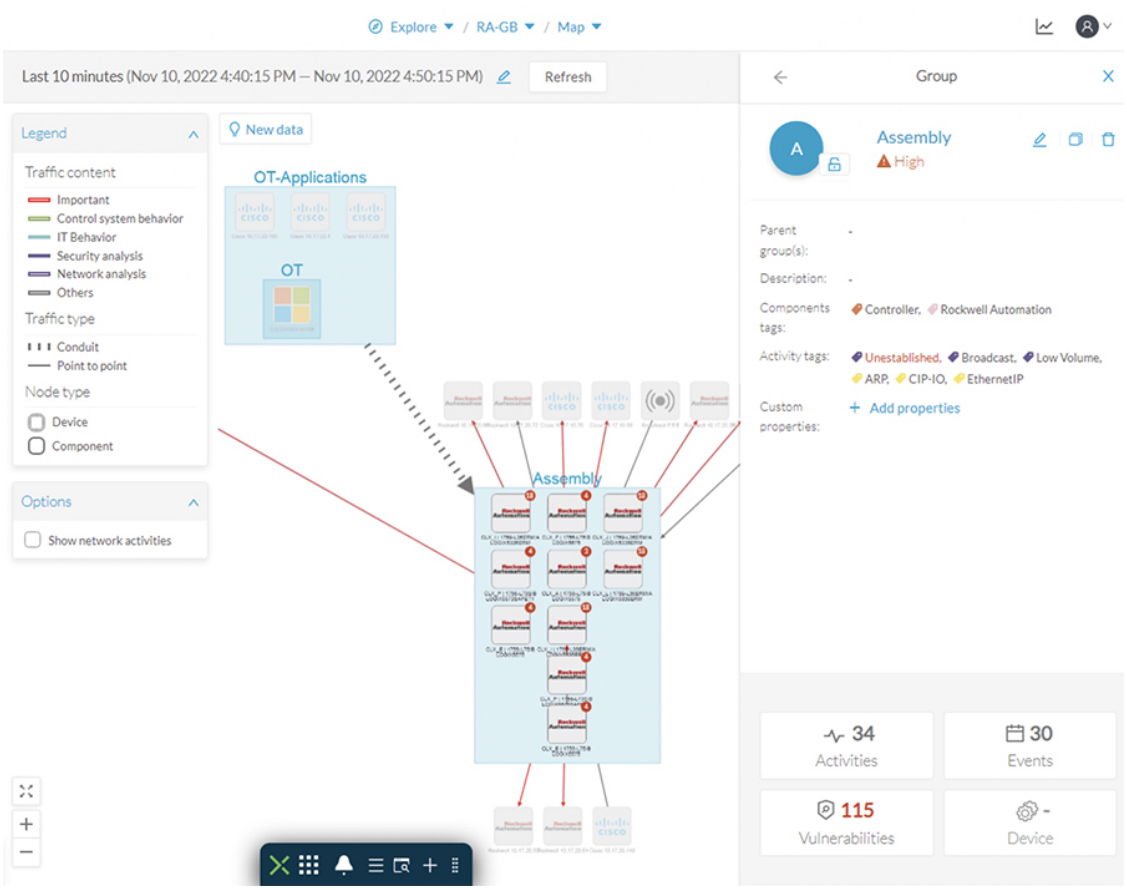

Recommendations when creating device groups in Cyber Vision include:

- Create Parent Groups based on manufacturing areas or processes
- Create Sub-groups based on process groups that span multiple manufacturing areas. This information helps define segmentation policies in the next chapter of the guide
- Assign an industrial impact variable to the group according to group criticality
- If the network is segmented use the subnet filter to identify components to be grouped
- If NAT is used, group devices using the inside IP address

Presets and Baselines

In large networks, it is recommended to use presets to divide the industrial network. A preset is a set of criteria which aims to show a detailed fragment of a network. Cyber Vision data can be filtered to create a preset per device tag, risk score, device groups, activity tags, sensors, network information (e.g., subnet or VLAN), or keyword. It is recommended to use presets to define the processes which should be monitored. For example, a preset could be defined to view all assets and traffic within a given production line, resulting in alerts being generated when a change is detected in production line activity.

Monitor mode in Cisco Cyber Vision is a feature used to detect changes inside industrial networks. The traffic patterns in the industrial network are generally constant and their behaviors tend to be stable over time. To

start monitoring a network, the normal operating state needs to be defined. For example, a normal state of the network can be a typical weekday operating mode, in which numerous processes are performed iteratively. Alternate operating states can also be captured, such as a weekend slow down, or during a holiday period.

After capturing the data (the recommended collection time is 2 weeks), a baseline can be saved and changes, either normal or abnormal, are then noted as differences in the baseline. Deviations from the baseline can either be acknowledged, and included as part of the baseline, or investigated further. The following figure illustrates new components being reported with the following information:

1. How many components are new or have changed
2. List of components
3. Filter criteria for the preset (in this example, the Cisco Cyber Vision Sensor is used)

Figure 25: Cyber Vision deviations from Baseline

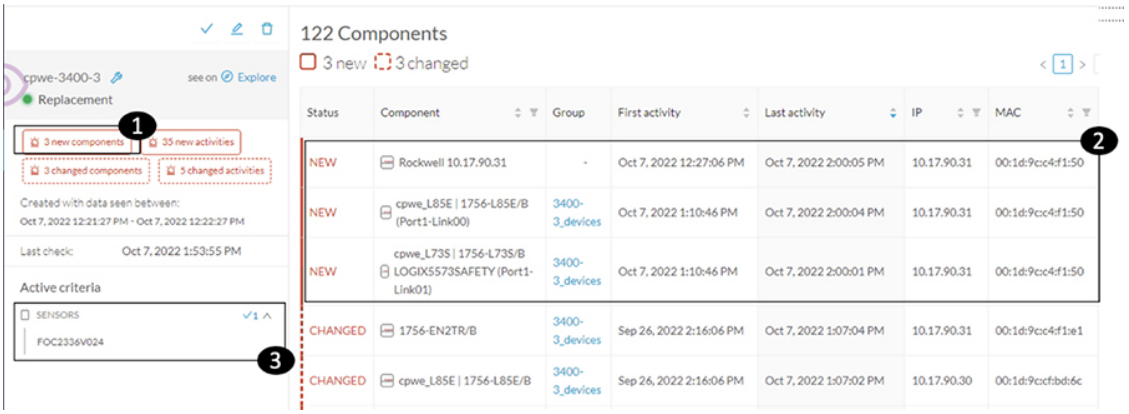

Note: it is not recommended to include the public IP address tag in any baseline creations. This will result in too many alerts as devices that communicate outside the network are typically too dynamic.

Presets containing critical assets are a good candidate for creating baselines. Typically, critical assets are controllers which determine the plant operation. Cisco Cyber Vision can monitor programs and firmware version changes that might cause malfunction or even stop a production line. For this use case a Preset can be created filtering by Group(s) identifying the processes to be monitored and the *Controller* tag. Any changes on the Component will be highlighted as well as any new activities to a controller. Cisco Cyber Vision depicts a changed Component with the following information:

1. How many changes on devices are seen.
2. Detail of changes when selecting a component from the list. In this case a controller mode was changed. It is possible to investigate the change using flows and acknowledge or report differences.
3. Filter criteria for the Preset; in this example Group and controller tag is used.
4. OT user could investigate activity with flows. In this example flow properties show details associated with a program download such as downloaded project, workstation, and user.

Note: It is recommended to include a public IP address preset outside of a baselining activity. Having a preset dedicated to public IP communications will provide clear insight into what devices are trying to reach outside of the industrial network and security should be in place to either block or protect this traffic.

Cyber Vision IDS

Snort IDS is provisioned in some Cisco Cyber Vision sensors such as the Catalyst 9300, IC3000 and Catalyst IR8340. The rules and basic configuration of Snort is packaged in the Cyber Vision knowledge database (KDB) which is updated regularly by Cisco. Rules can be enabled and disabled based on a category and Cyber Vision provides the ability to upload custom rule files to generate specific alerts. For more information about Snort in Cyber Vision see the Cyber Vision GUI User Guide.

Note: Snort IDS is deployed as an IOx application in the Catalyst 9300. The bandwidth is limited to approximately 30,000 packets per second and should be reserved for east / west traffic between cell/area zones only. If a high-performance IDS solution is required, the recommendation is to deploy a dedicated firewall appliance such as the Cisco Secure Firewall alongside the Catalyst to transparently capture the traffic. Design guidance for this approach is currently out of scope for this version of the design guide but will be added during a later release.

Performance

The control system engineer deploying a hardware or network sensor must consider its performance numbers. The critical performance metrics for Cyber Vision Version is documented in the Cisco Cyber Vision Architecture Guide.

Note: In order to reduce the load on Cyber Vision Sensor, avoid monitoring both access and trunk ports as it doubles the number of packets fed to the DPI engine and the bandwidth used if the mirrored traffic is sent over RSPAN. When installing sensors on access switches, monitor the access ports only. If you do not have sensors in the access switches, then SPAN on the aggregation switch trunk ports will be required.

Licensing

Cisco Cyber Vision Center requires a license. Licenses must be available in a smart account to register product instances. The following options are available:

- **Direct cloud access to Cisco Smart Software Manager (SSM)**: Cyber Vision has a direct connection to the SSM cloud.
- **Cloud access via https proxy**: Cyber Vision uses a web proxy such as the Umbrella Secure Internet Gateway to send information to Cisco SSM.
- **Cisco Smart Software Manager On-Prem**: Usage information is sent to a local appliance. Cisco SSM On-Prem would reside in the IDMZ, and information is periodically sent to the SSM cloud.
- **Offline**: Licenses are reserved in SSM and applied manually.

The recommended approach, and the option validated in this design is Cisco Smart Software Manager On-Prem.

Note: Cisco Cyber Vision Global Center does not require an additional license.

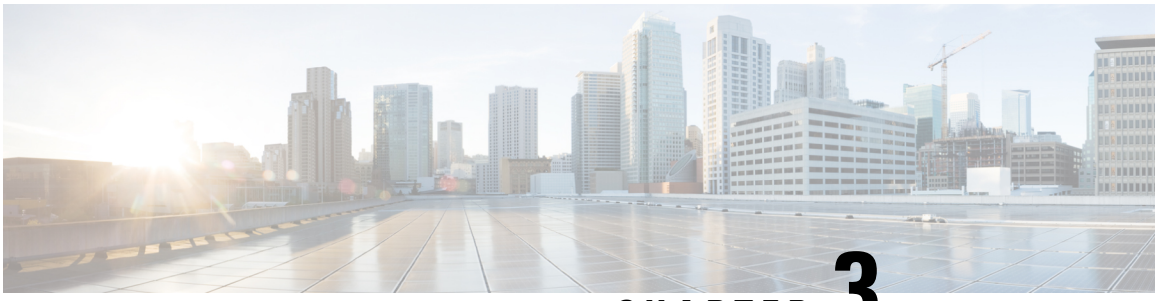

CHAPTER 3

Segment the Network into Smaller Trust Zones

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 form 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 26: Policy Enforcement Across the Distribution Switch

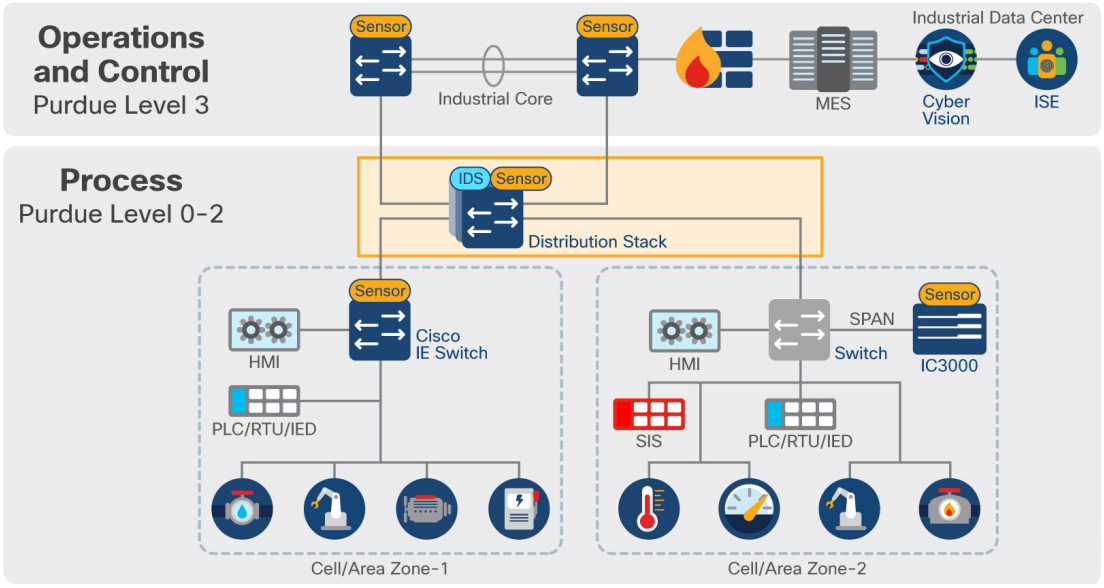

The layer 3 boundary and gateway for devices is in the distribution switch as shown in the following figure. It is recommended that security is first created at this layer of the network, to allow and deny communications for inter-cell/area zone communication such as controller-to-controller communication 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 27: ISE Communication Flows

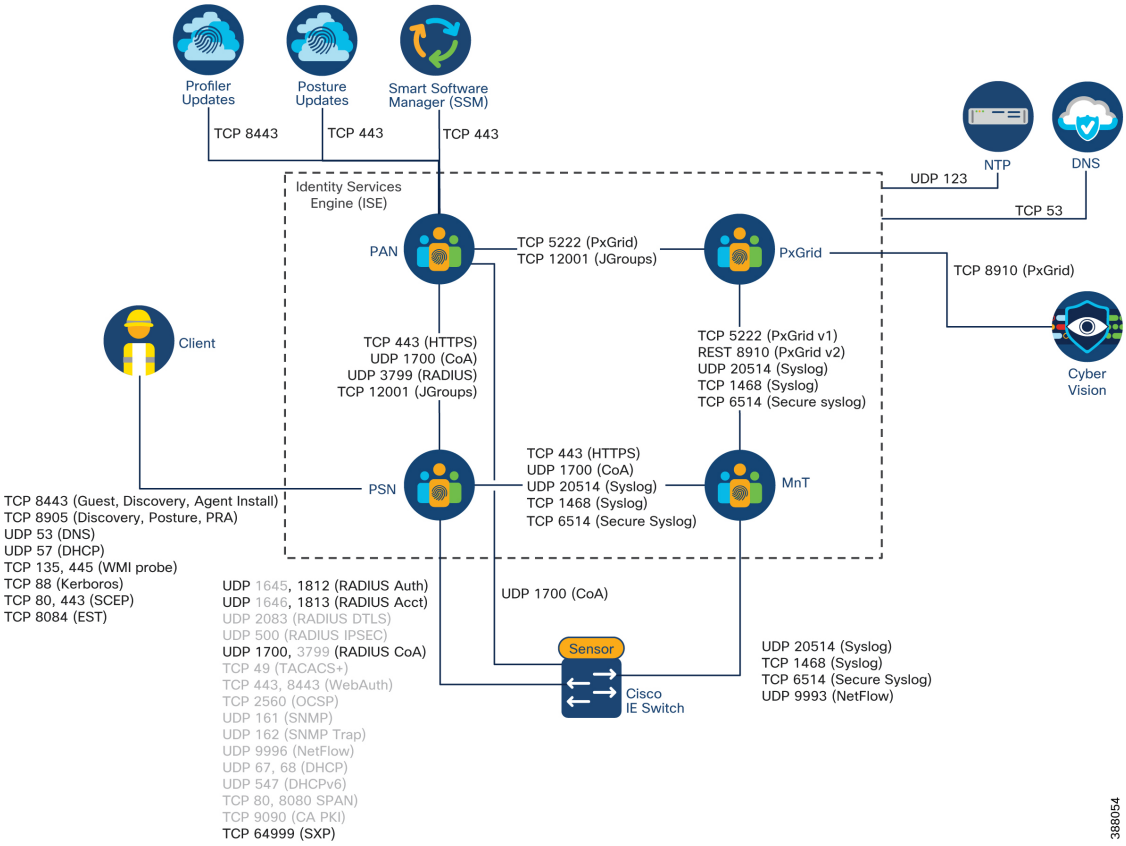

Note: Not all flows depicted in this diagram are used in this design guide. An example is demonstrated in the flow between the ISE PSN and the Cisco switching infrastructure, with greyed out values indicating "not in use". All ports are shown to provide clarity when ISE is portrayed for use beyond this 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 28: Defining the TrustSec Domain

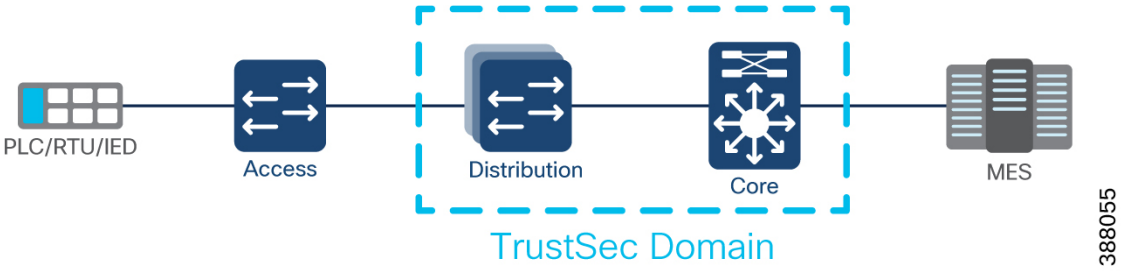

Packets entering the domain are tagged with an SGT containing the assigned security group number of the source device. This packet classification is maintained along the data path within the Cisco TrustSec domain for the purpose of applying security and other policy criteria. The final Cisco TrustSec 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 29: Static vs Dynamic Classification

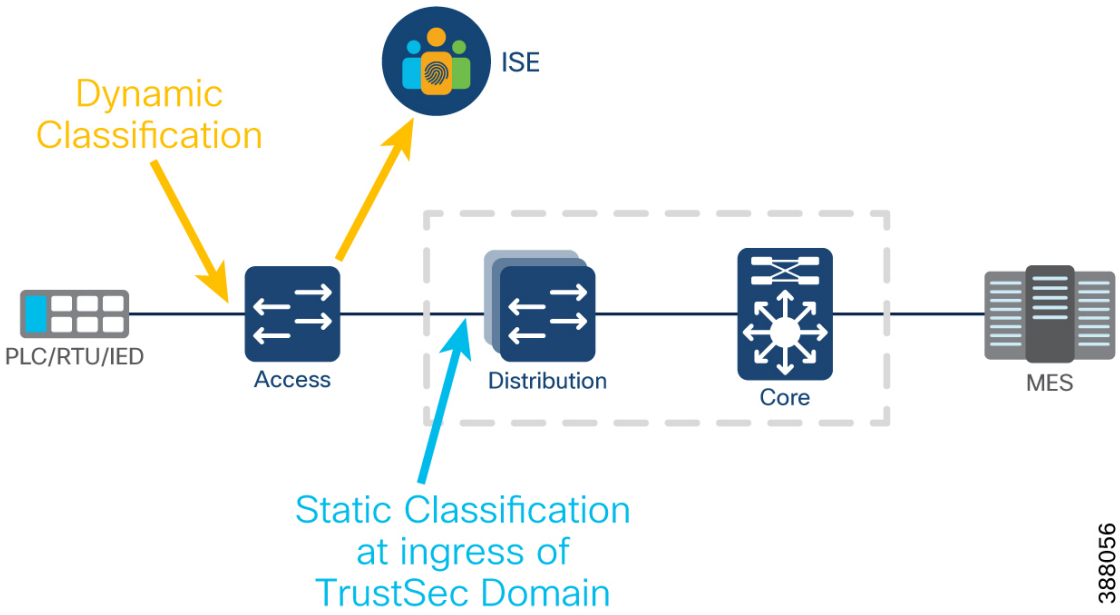

Generally, dynamic classification is done at the access switch and static classification is performed at ingress to the TrustSec domain. The SGT tag that gets inserted into the traffic is known as the source SGT, as it is the group that the source of the traffic belongs to. The destination SGT is the 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 30: SXP vs. Inline Tagging

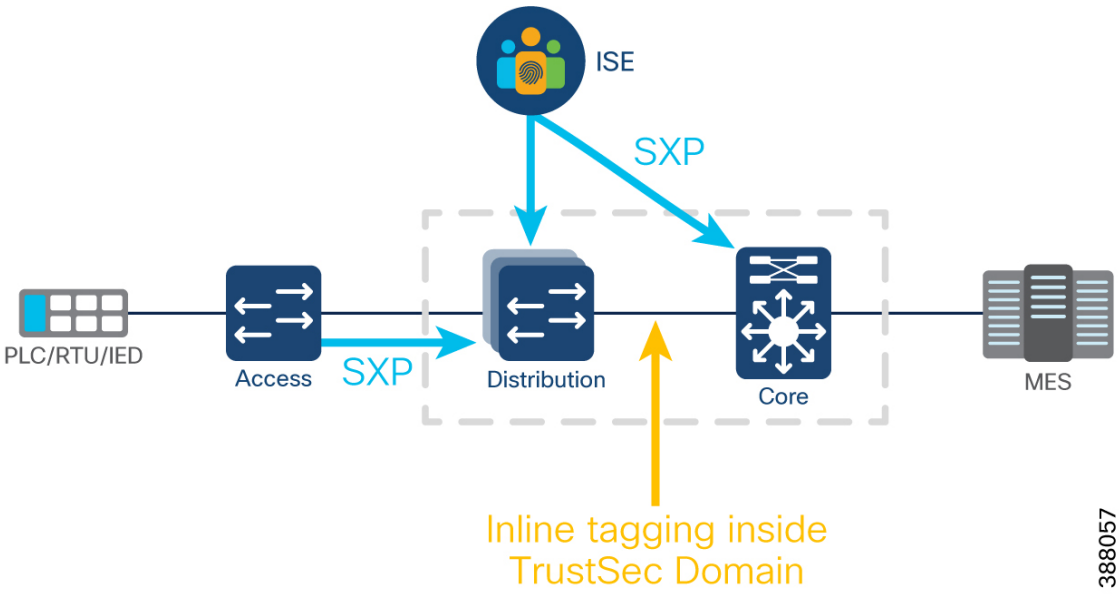

A network device performing the enforcement needs to determine the destination SGT as well as the source for applying the SGACL. The destination SGT can be determined in one of the following ways:

- 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 31: TrustSec Enforcement at egress

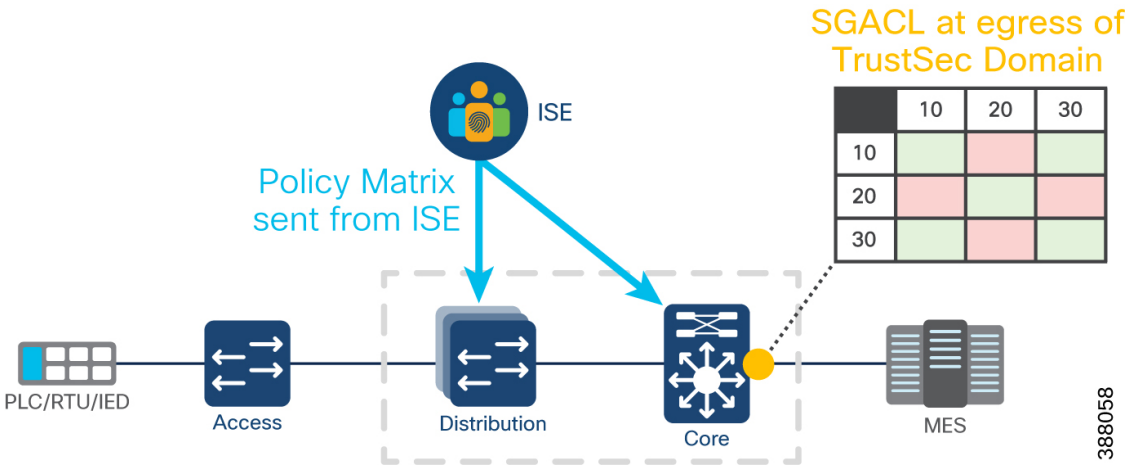

By applying access control between pairs of security groups, Cisco TrustSec achieves role-based, topology-independent access control within the network. Changes in network topology do not normally require a change in the SGACL-based security policy. Some care must be taken to ensure the proper classification 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 32: TrustSec Classification, Transport & Enforcement Example

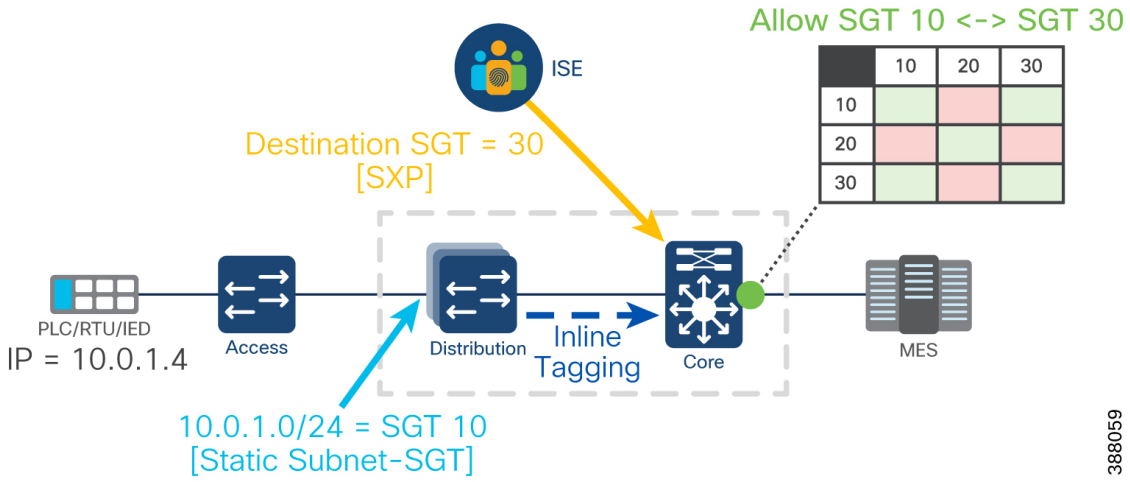

ISE/SGT Design Considerations

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 33: Tree Topology in the Industrial Zone

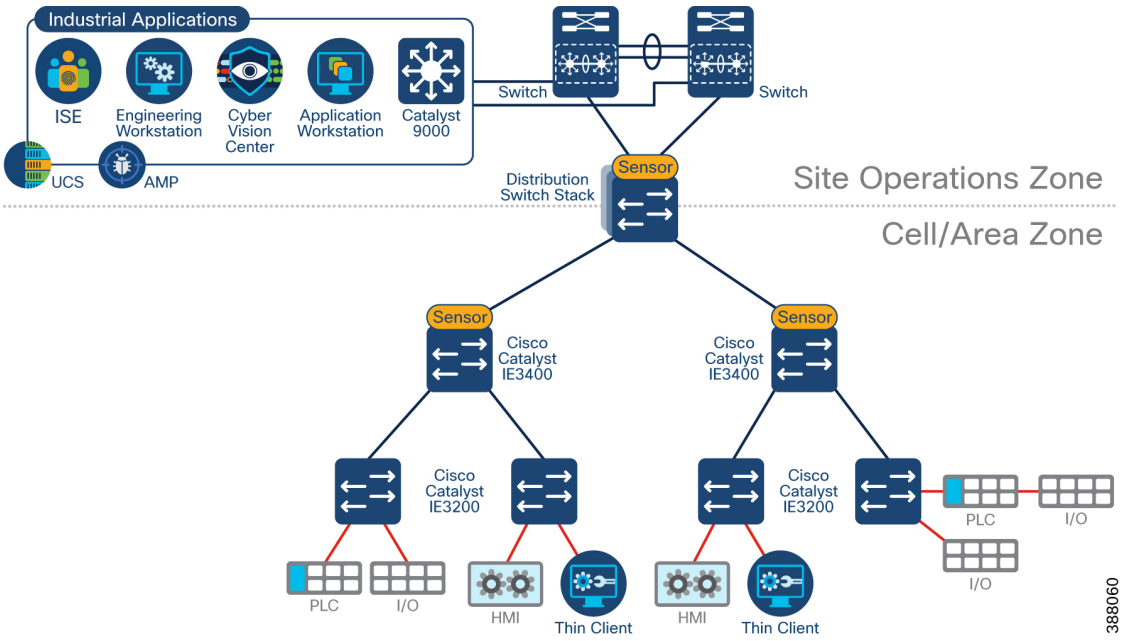

The recommendation for this model is to both classify and enforce at the distribution switch. When creating AA policies on access switches, do not include SGT assignment as part of the policy. Devices will have unrestricted communication within their cell as no PEP exists within the zone. On the distribution 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 34: SGT Classification at Distribution Ingress

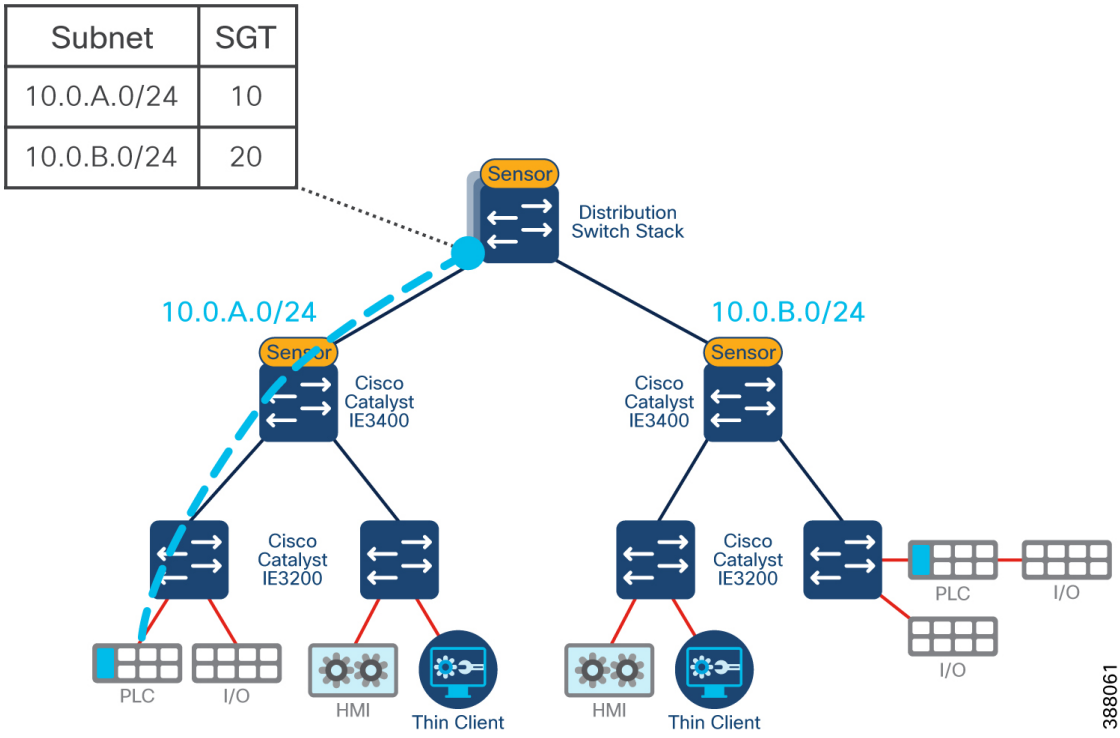

After cells have been defined, the next step is to define policy for communication that must leave the zone. The least privilege approach to security will result in the Cell/Area zones being in a deny by default state (i.e., SGT 10 deny SGT 20) and only select services crossing the zone boundaries. A 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 35: Cyber Vision & ISE pxGrid Integration

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 36: Cyber Vision Interlock Group within Zone2 Parent Group

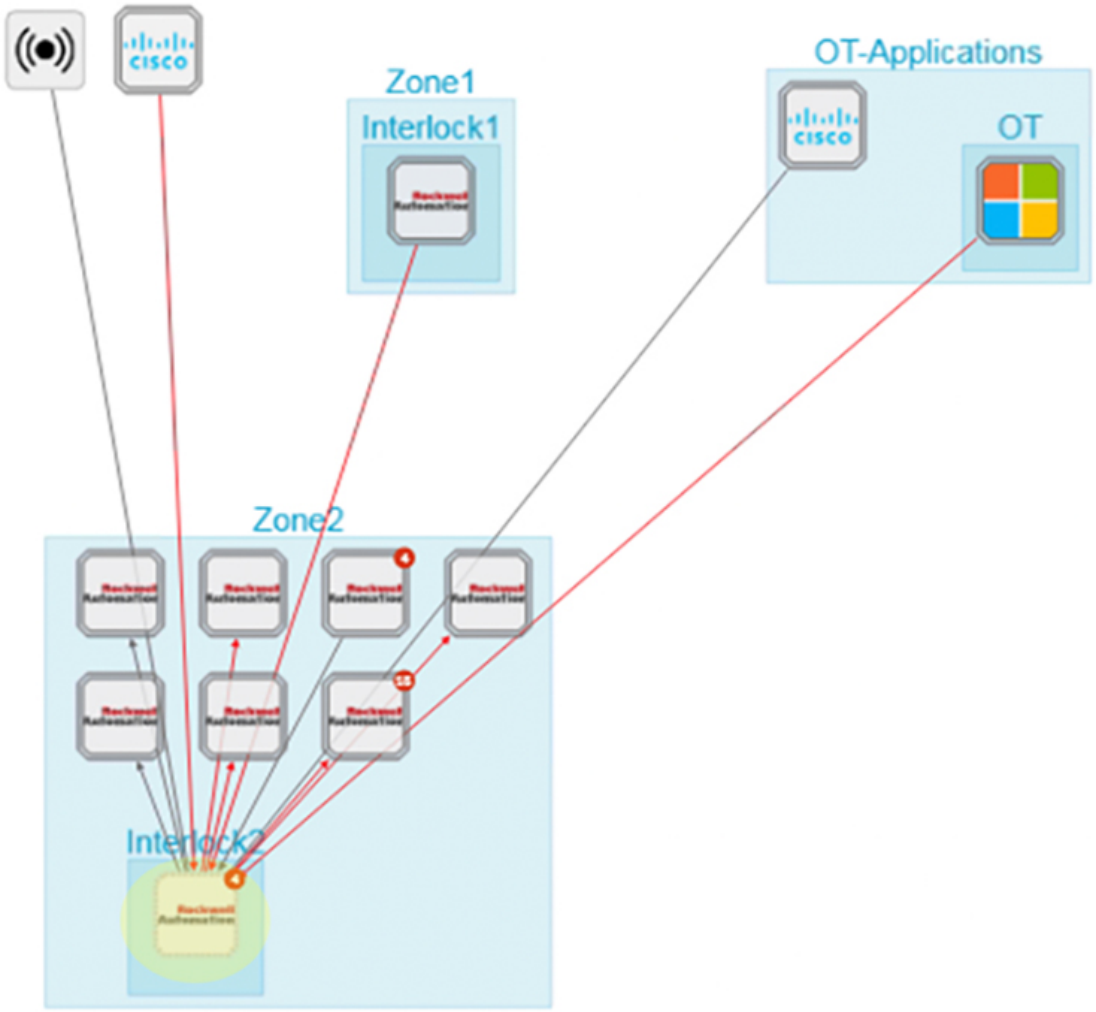

After the group tag has been shared with ISE, a change of authorization (CoA) is sent to the access switch that the PLC is connected to. This results in the PLC reauthenticating with ISE, ultimately matching the new AA policy defined for interlocking PLCs.

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 37: ISE Asset Information after Cyber Vision Integration

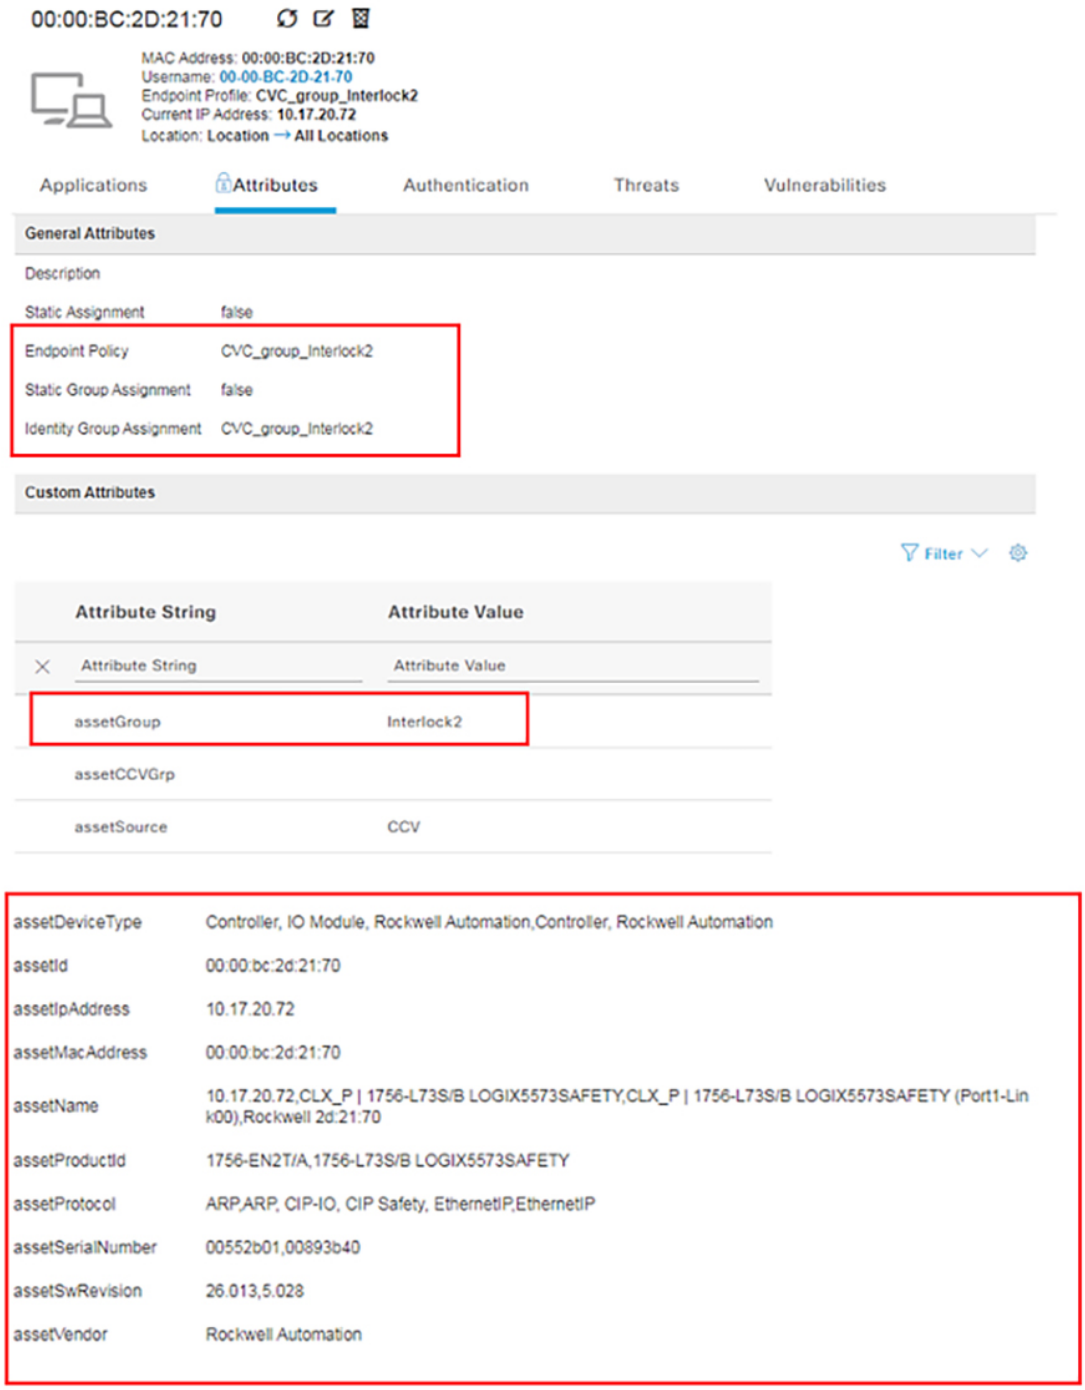

After zones have been defined, and traffic has been classified, the policy enforcement matrix can be defined. The following table shows an example policy enforcement matrix.

Figure 38: Sample TrustSec Matrix

	100	101	102	103	911	4001	4002	5001	9001	9002
100										
101										
102										
103										
911										
4001										
4002										
5001										
9001										
9002										

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

388065

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 39: Policy Enforcement at Distribution Egress to Cell/Area Zone

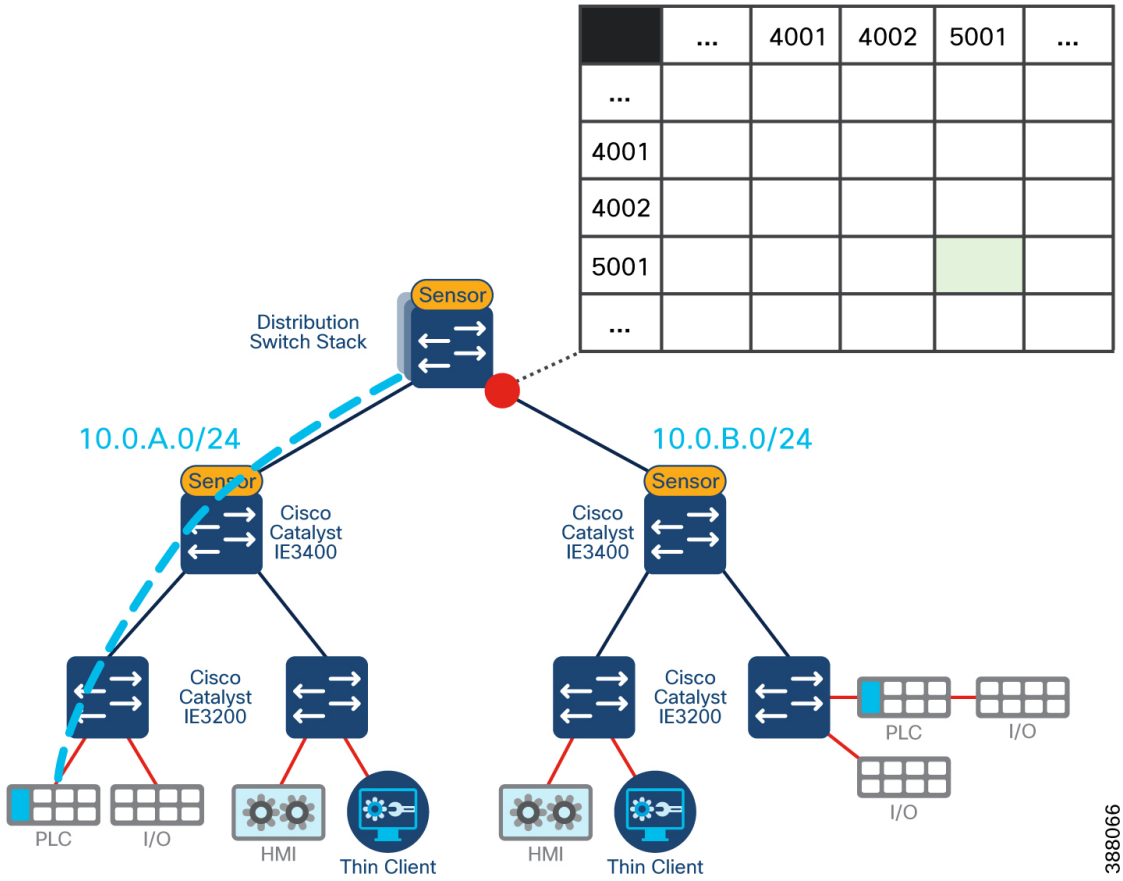

When creating the policy matrix, only think about flows that cross TrustSec domains. If a zone does not enter the TrustSec domain, nor does it intend to, it is okay to deny traffic of the same SGT. Use the learnings from the previous step in the journey with Cisco Cyber Vision to understand traffic that 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 40: Reducing the Policy Matrix Size

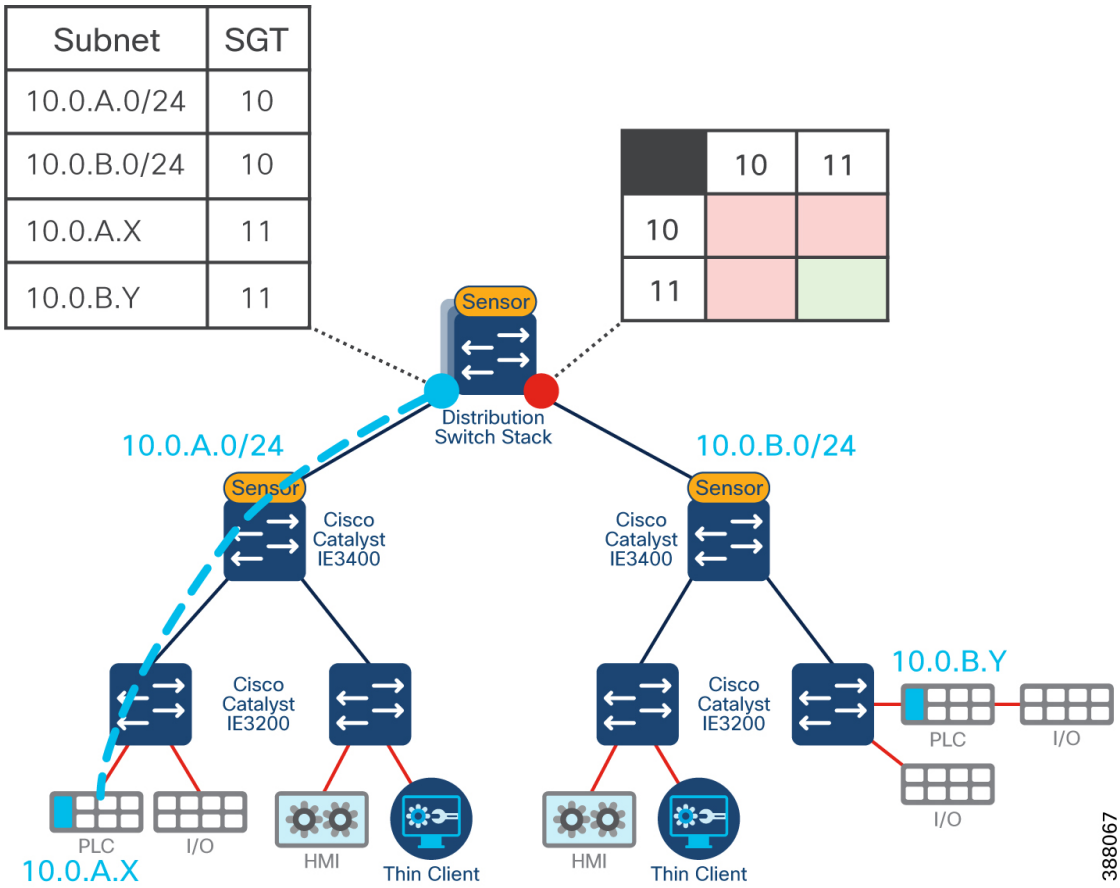

Another consideration in larger networks is the number of SXP connections. The maximum number of ISE SXP peers per PSN is 200. When doing dynamic classification, the IP-SGT binding must be shared from the access switch to the TrustSec domain. One method of doing this is for the access switch to create 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 41: SXP Daisy Chain starting from Access Switch up to Distribution Switch

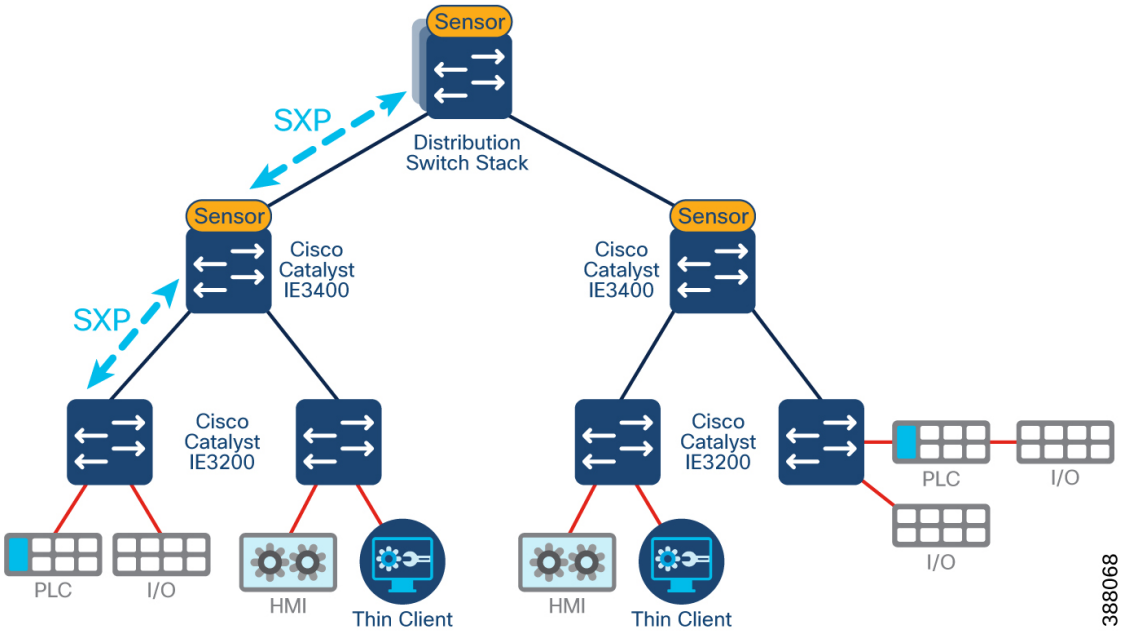

Segmentation when the layer 3 boundary also participates in layer 2 connectivity

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

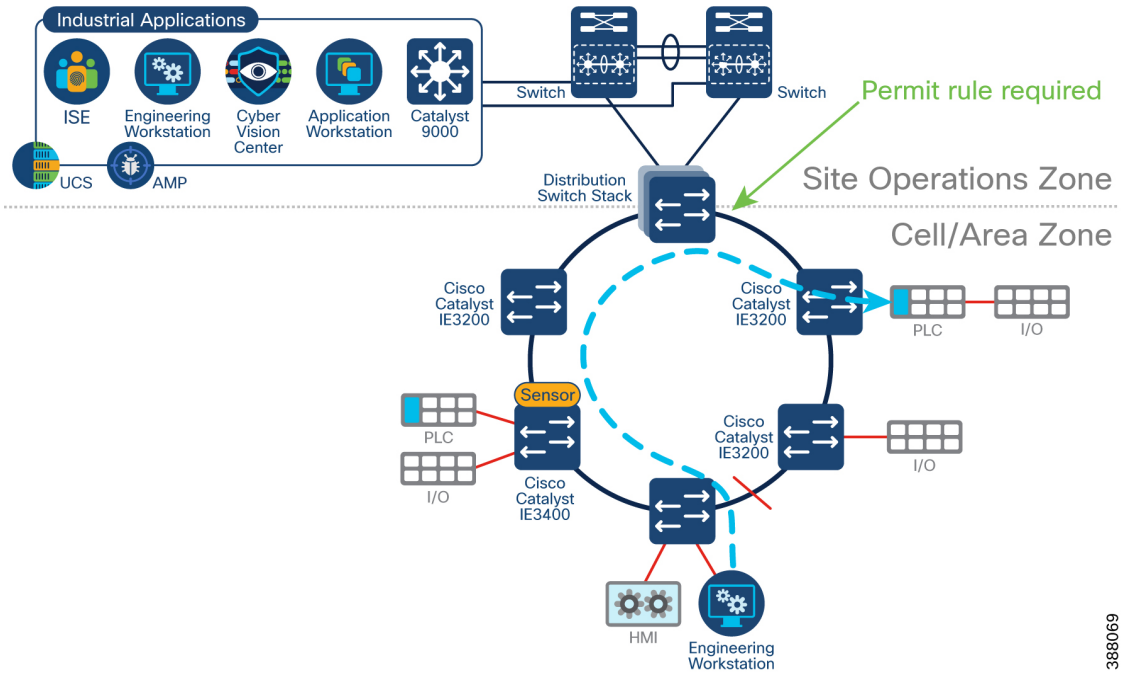

To ensure traffic is not blocked by policy, make sure that each such ring has its own unique SGT and does not share a tag with any other zone in the network as per the design recommendation for large networks.

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.

Figure 43: L2 NAT in the Industrial Zone

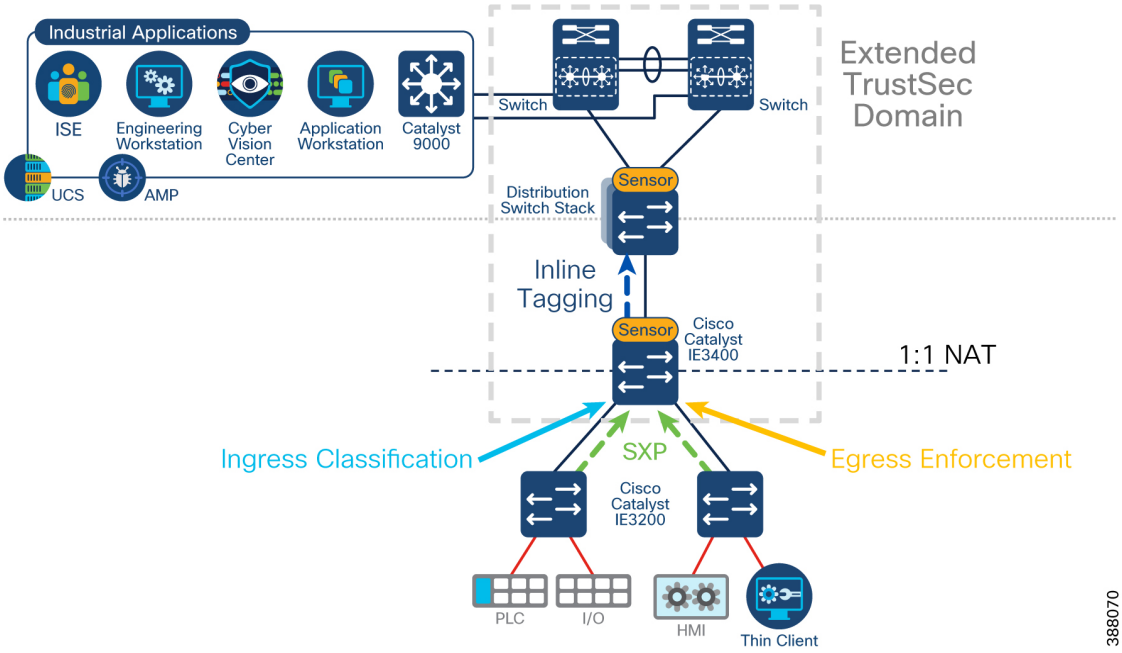

The recommendation in this case is to enable SGT classification on the IE3400 and enable inline tagging between the NAT boundary and the distribution switch. The SGT is not stripped from the traffic during NAT, and since a tag will already exist when entering the distribution switch, it will not be overwritten 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 44: SXP Domain Filtering in the TrustSec Domain

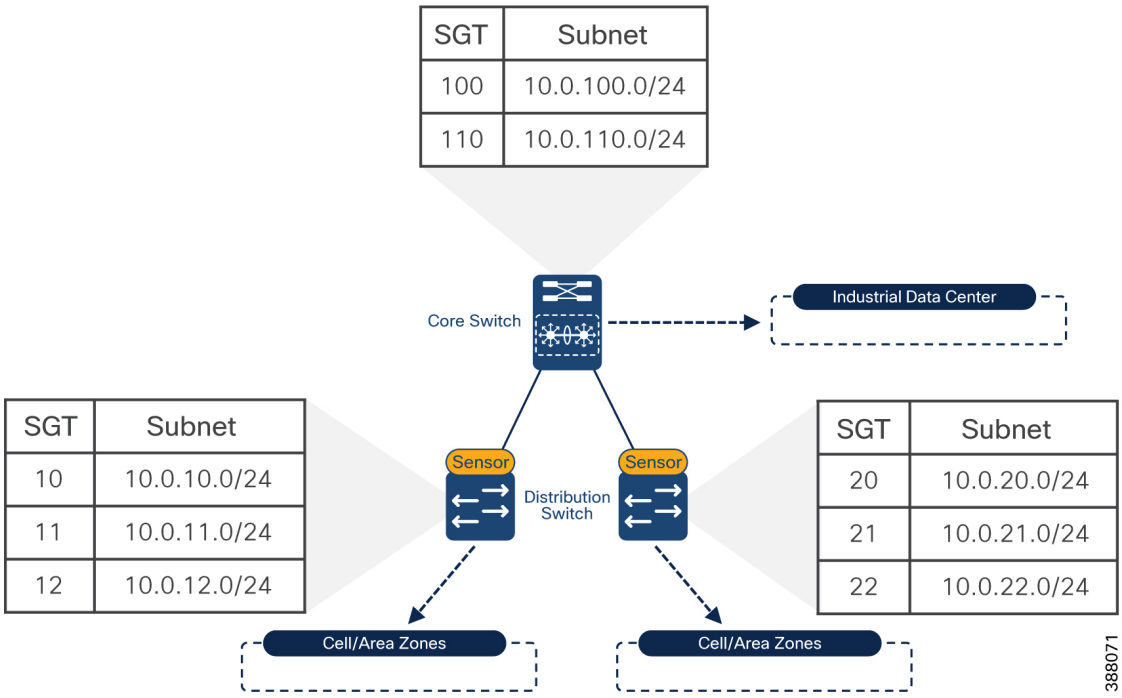

Static Segmentation in the Industrial Zone

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 45: Static SGT Classification on the Physical Ports of the IE3400

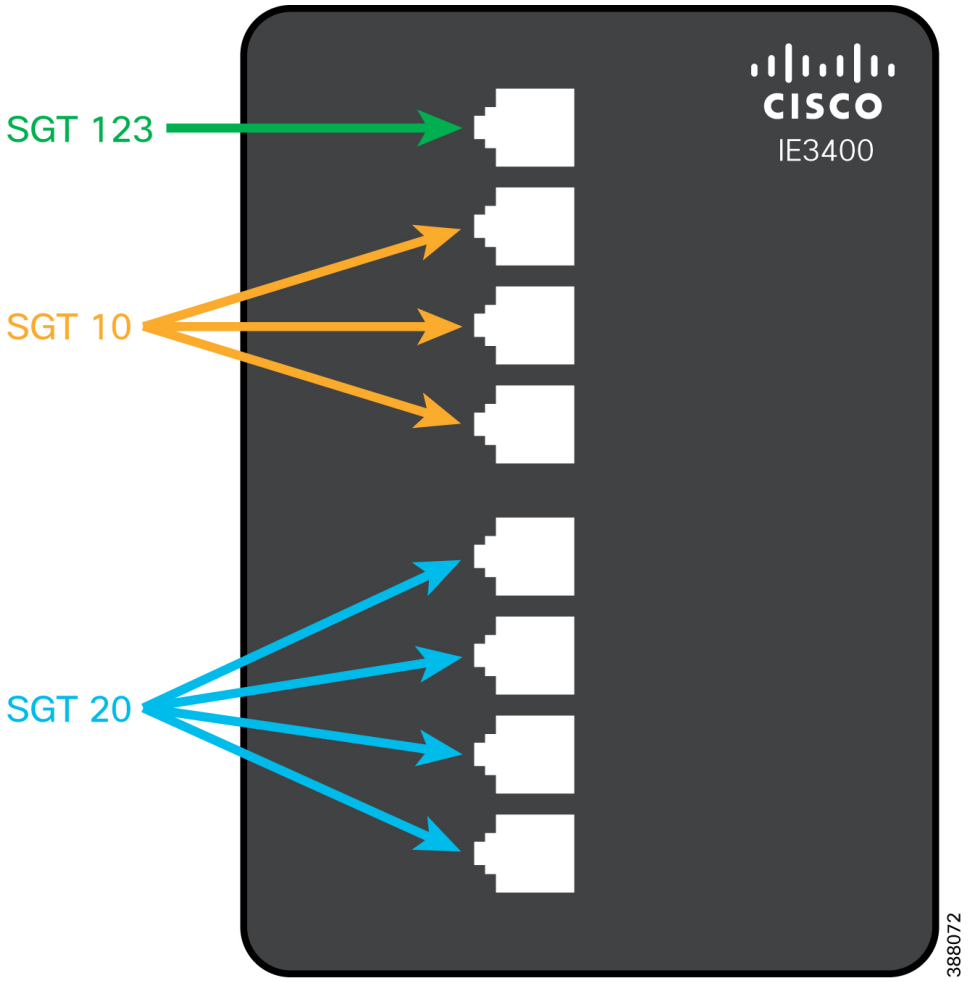

Consider the following when implementing static port assignment on the access switches:

- 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 46: ISE AA Policy with Active Directory User Group

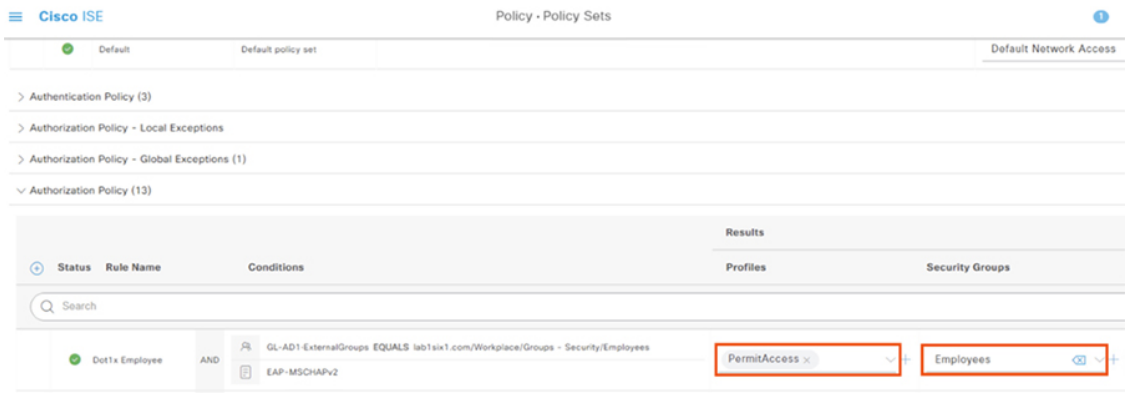

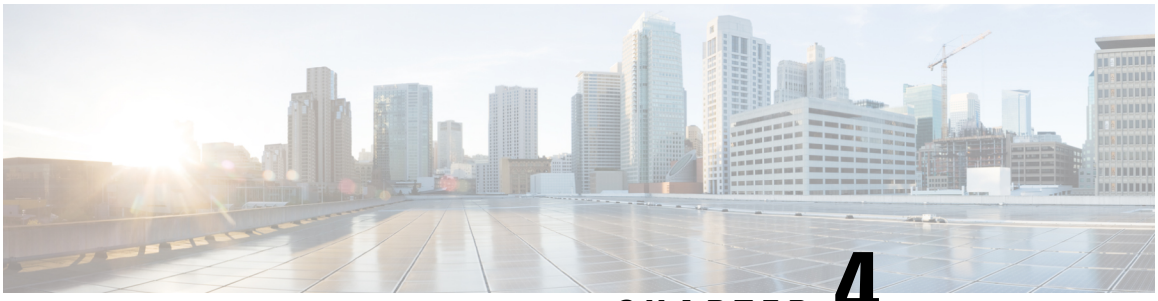

CHAPTER 4

Develop an Incident Investigation and Response Plan

Develop an Incident Investigation and Response Plan

Reducing the mean time to detect (MTTD) and mean time to respond (MTTR) are the end goals of any security operations team. How long does it take to detect an issue, and then how quickly can we respond?

Security information and event management (SIEM) is a well-tested take on log-and-event management solutions. At its core, SIEM is about gathering as much log information as possible from all over an organization. Many SIEM solutions can take log data from IoT security tools, firewall event logs, and everything in between. This kind of solution starts to break down the silo walls, integrating with multiple solutions and centralizing important security information. What SIEM doesn't do is give security engineers a boost in threat response time and efficacy. Seeing the security landscape of your organization is great for many things but responding to threats is just as important.

Security orchestration, automation and response (SOAR) takes a lot of what makes SIEM great and adds extra layers to account for some of the limitations. Like SIEM, SOAR solutions take data from different parts of the security infrastructure and put it in one place. SOAR solutions offer options to automate various auditing, log, and scanning tasks. Automation can't take care of everything, however, and sometimes requires human intervention. The "response" part of SOAR is about organizing and managing the response to a security threat. This feature set utilizes orchestration and automation information to help security staff make decisions and respond to threats. SOAR automation doesn't automate responses to security breaches. It automates simple analysis tasks to reduce security personnel workloads.

While SIEM and SOAR emphasize logs and analysis, Extended Detection and Response (XDR) solutions focus on the endpoints themselves. This is where the action is. This is what the outside parties are attacking.

Cisco SecureX

SecureX is a cloud-native, built-in platform experience within the Cisco Secure portfolio and connected to your infrastructure, which is integrated and open for simplicity, combines multiple otherwise disparate sensor and detection technologies into one unified location for visibility, and provides automation and orchestration capabilities to maximize operational efficiency, all to secure your network, users and endpoints, cloud edge, and applications. With SecureX, security teams can:

Develop an Incident Investigation and Response Plan

- **Radically reduce the dwell time and human-powered tasks** involved with detecting, investigating, and remediating threats to counter attacks or securing access and managing policy to stay compliant – make faster decisions with less overhead and better precision with less error.
- **Enable time savings and better collaboration** involved with orchestrating and automating security across SecOps, ITOps, and NetOps teams, which helps advance your security maturity level using your existing resources and realizes more desired outcomes with measured, meaningful metrics.

  Reduce MTTD / MTTR and reduce costs with real benefits in 15 minutes – even if you start small with a single product and grow as your needs dictate over time to consolidate security vendors without compromising security efficacy.

SecureX Ribbon

Part of the SecureX design philosophy is that you should not have to navigate to multiple different consoles to get all the functions you need for one business task. The SecureX ribbon brings this philosophy to reality across the portfolio. Via the ribbon, a persistent bar in the lower portion of the UI of all ribbon-capable products, you have access to all the functions lent to SecureX by all your deployed SecureX-capable technologies. The ribbon is collapsible and expandable to open ribbon apps, launch integrated applications, and view your account profile. From the ribbon, you can pivot between SecureX or the console of any integrated product, into any other integrated product, and search the current web page for malicious file hashes, suspicious domains and other cyber observables. You can then also add observables to a case or investigate observables in the threat response application.

Figure 47: SecureX Ribbon in Cyber Vision

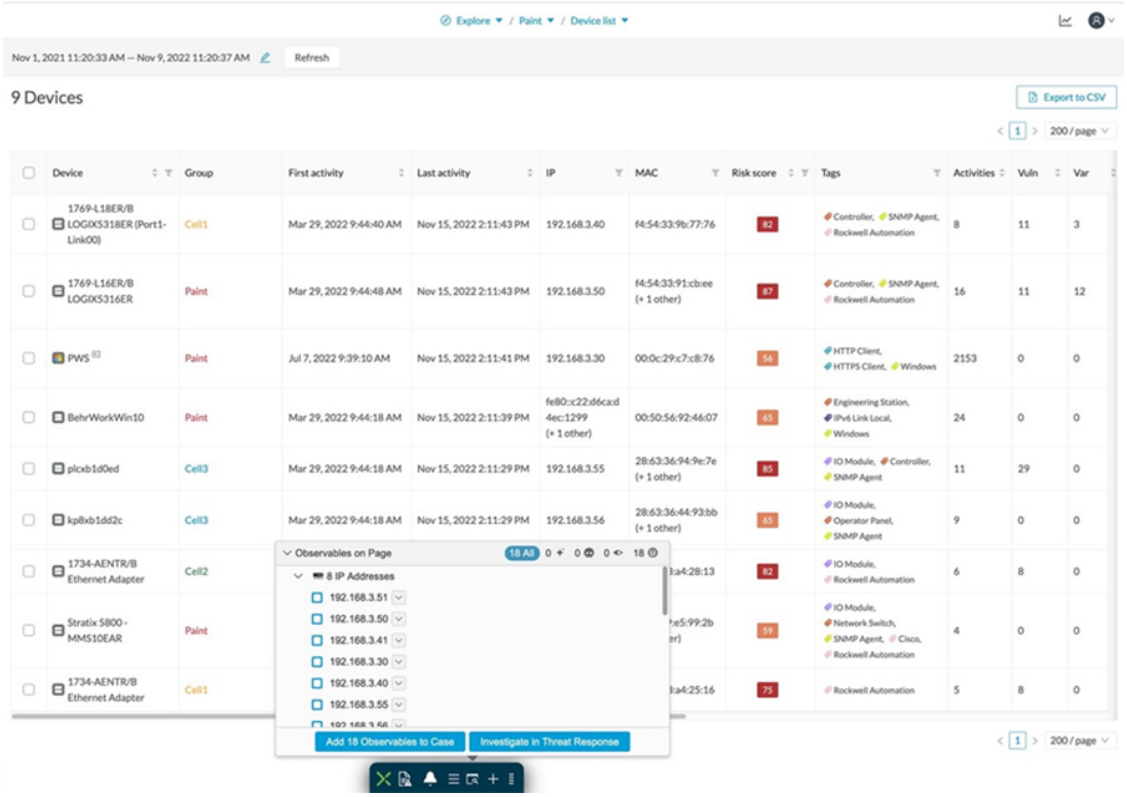

The SecureX ribbon is a feature of Cyber Vision and appears on the bottom of the Cisco Cyber Vision Center user interface.

SecureX Threat Response

SecureX threat response is a security investigation and incident response application. It simplifies threat hunting and incident response by accelerating detection, investigation, and remediation of threats. The threat response application provides your security investigations with context and enrichment by connecting your Cisco security solutions (across endpoint, network, and cloud) and integrating with third-party tools, all in a single console.

To understand whether a threat has been seen in your environment as well as its impact, SecureX threat response aggregates contextual awareness from Cisco security product data sources along with global threat intelligence from Talos® and third-party sources via APIs. Threat response identifies whether observables such as file hashes, IP addresses, domains, and email addresses are suspicious or malicious, and whether you have been affected by them. It also provides the ability to remediate directly from the interface and block suspicious files, domains, isolate hosts, and more without pivoting to another product first. Key features and benefits include:

- **Relations Graph**: visualize all the observables found during the investigation and determine the relationships between them
- **Casebook**: save, share, and enrich threat analysis to enable documentation of all analysis in a cloud casebook so seamlessly work a case across multiple tools, Cisco or otherwise and better collaborate among staff
- **Response Actions**: enforce protective controls without pivoting to other product consoles

Figure 48: SecureX Threat Response Example

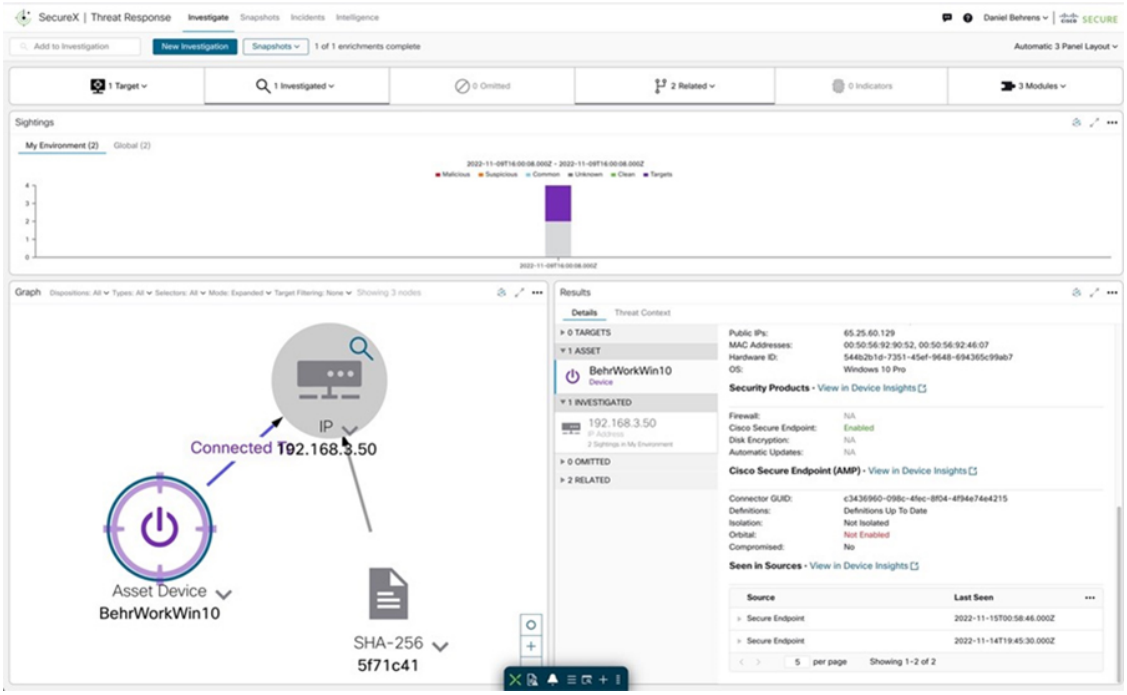

It is possible to launch a SecureX investigation from Cisco Cyber Vision Center. The Cyber Vision baseline feature can help highlight unexpected and potentially malicious activity in the network by monitoring a known good state for any changes. Often, an infected device starts by scanning the network to identify vulnerable components to attack. This traffic anomaly can be easily identified using Cisco Cyber Vision Monitor Mode.

Develop an Incident Investigation and Response Plan

To cross launch an investigation in SecureX Threat Response, click on the *Investigate in Cisco Threat Response* button after clicking on the suspicious component.

SecureX Orchestration

SecureX orchestration automates repetitive and critical security tasks such as threat investigation, hunting, and remediation use cases. SecureX orchestration provides pre-built workflows and response capabilities, or you can build your own with a no/low-code, drag-drop canvas to strengthen operational efficiency and precision, and lower operational costs.

Figure 49: SecureX Orchestration Workflow - Quarantine Device using SGT

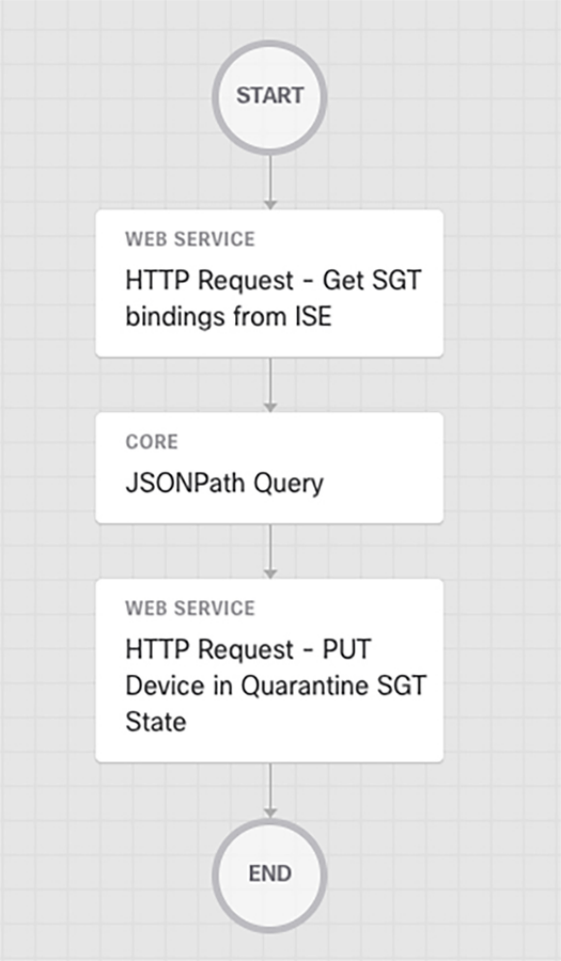

SecureX orchestration enables you to define workflows that reflect your typical security processes; the automation steps (activities), the logic or flow between these steps, and how to flow data from one step to the next. With SecureX, you can leverage Cisco Secure and thirdparty multi-domain systems, applications, databases, and network devices in your environment to create these workflows. An example workflow would be to take an IP address, or hostname and assign that endpoint an SGT in ISE that would ultimately block communication from occurring on the network.

Develop an Incident Investigation and Response Plan

Figure 50: Invoking SecureX Orchestration Workflow from the Ribbon in Cyber Vision

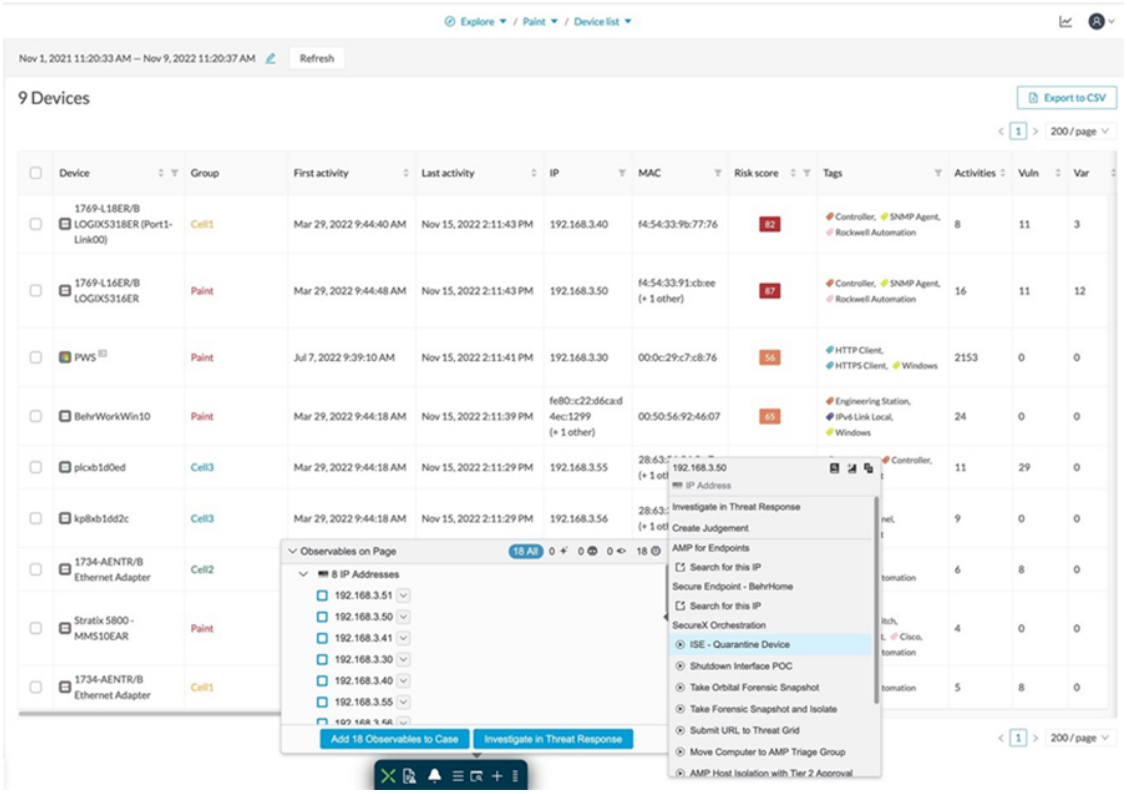

Note: The current version of the Industrial Security Design guide does not provide in-depth design guidance for Incident Investigation and Response. This part of the security will be added in a future release.

Develop an Incident Investigation and Response Plan

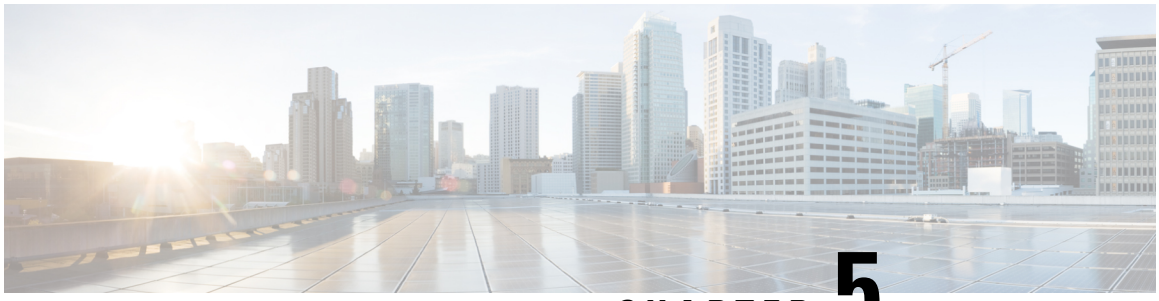

CHAPTER 5

Appendix A

Deployment Guides

IDMZ

- Securely Traversing IACS Data across the IDMZ Using Cisco Firepower Threat Defense - https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/3-5-1/IDMZ/DIG/CPwE_IDMZ_2_CVD/CPwE_IDMZ_2_Chap1.html

Cyber Vision

- Cisco Cyber Vision Center Appliance Installation Guide -https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/CenterAppliance/Release-4-1-202/b_Cisco_Cyber_Vision_Center_Appliance_Installation_Guide.html
- Cisco Cyber Vision Center VM Installation Guide - https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/Center-VM/Release-4-1-2/b_Cisco_Cyber_Vision_Center_VM_Installation_Guide.html
- Cisco Cyber Vision Network Sensor Installation Guide for Cisco IE3300 10G, Cisco IE3400 and Cisco Catalyst 9300 -https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/IE3400/b_Cisco_Cyber_Vision_Network_Sensor_Installation_Guide_for_Cisco_IE3300_10G_Cisco_IE3400_and_Cisco_Catalyst_9300.html
- Cyber Vision Monitor Mode / Baseline Creation - https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/GUI/Release-4-1-2/b_Cisco_Cyber_Vision_GUI_User_Guide_Release-4-1-2/m_monitor.html

ISE

- Cisco Identity Services Engine Installation Guide - https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/install_guide/b_ise_installationGuide32.html
- Cisco Identity Services Engine Segmentation Chapter - https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2/b_ISE_admin_32_segmentation.html
- Integration Cisco Cyber Vision with Cisco Identity Services Engine (ISE) via pxGrid - https://www.cisco.com/c/dam/en/us/td/docs/security/cyber_vision/Integrating-CiscoCyber-Vision-with-Cisco-Identity-Services-Engine-via-pxGrid_3_1_1.pdf

SecureX

- Cisco Cyber Vision SecureX Integration - Cisco Cyber Vision SecureX Integration -https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/GUIAdministration-Guide/b_cisco-cyber-vision-GUI-administrationguide/m_integrations.html#topic_5737

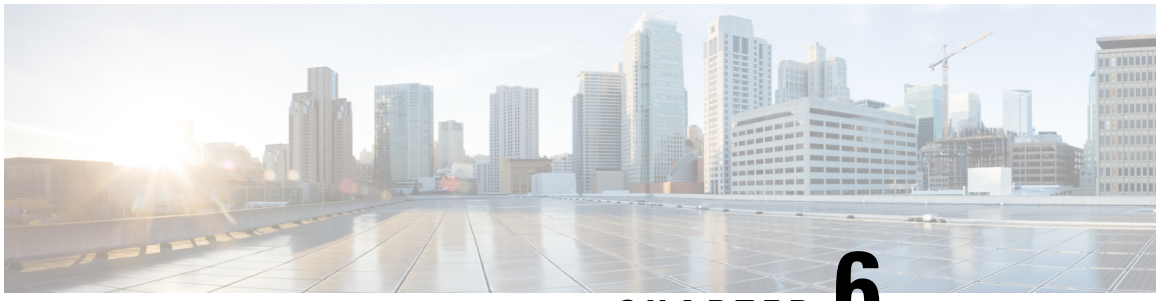

CHAPTER 6

Appendix B

TrustSec Configurations

The following configurations are required to deploy TrustSec on the network:

Figure 51: User Interface for TrustSec Configuration

Configuration Item	Configuration Target	Configuration Tool*
Define and create SGTs and Policies	ISE	Cisco DNA Center or ISE
Define ISE as AAA server on network settings	Industrial switches and ISE	Cisco DNA Center or Industrial switch and ISE
Enable device tracking on access ports	Industrial switches	Cisco DNA Center or Industrial switch
Port-based Authentication	Industrial switches	Cisco DNA Center (templates) or Industrial switch
Fall back policy and static entries	Enforcement switches (Distribution switch or Industrial switch)	Cisco DNA Center (templates) or switch
Propagation (SXP or inline tagging)	Industrial switches, Distribution switch and ISE	Cisco DNA Center (templates) and ISE or switch and ISE
Enable enforcement	Distribution switch or Industrial switch	Cisco DNA Center (templates) or Industrial switch
Profiling and profiling rules	ISE	ISE
Authentication and authorization policies	ISE	ISE
Cyber Vision sensor	Cyber Vision Center and switch	Cyber Vision Sensor Management Extension and Cisco DNA Center (templates) or switch

388078

* Method used for the CVD.

Define and Create SGTs and Policies Using Cisco DNA Center

1. From the Cisco DNA Center web interface, navigate to **Policy** > **Group-Based Access Control**.
2. Click the **Security Groups** tab.
3. Click **Create Security Group**.

4. Fill out **Name** and optional **Tag Value**.
5. Click **Save Now**.
6. Click the **Deploy** link.

After creating the SGTs in Cisco DNA Center, the policy matrix can be updated to suit the enforcement intent. To make changes to the TrustSec policy matrix in DNA Center, do the following:

1. From the Cisco DNA Center web interface, navigate to **Policy > Group-Based Access Control**.
2. Click the **Policies** tab.
3. Click the square of the source and destination pair for which there needs to be a permit or deny contract.
4. On the **Create Policy** slide-in pane, click the **Change Contract** link and choose the appropriate option (**Permit IP, Deny IP**, and so on). Click the **Change** button.
5. Click the **Deploy** link at the top of the matrix.

The following figure shows the TrustSec policy matrix in Cisco DNA Center. The **Create Policies** button (1) is used to create a new policy and the **Default** link (2) allows you to change the default action on the policy. For a default deny policy, choose the **Deny_IP** default action.

Warning: Don't change default action to deny until all TrustSec elements have been configured and the policy has been tested with monitoring mode or log analysis.

Figure 52: TrustSec Policy Matrix in Cisco DNA Center

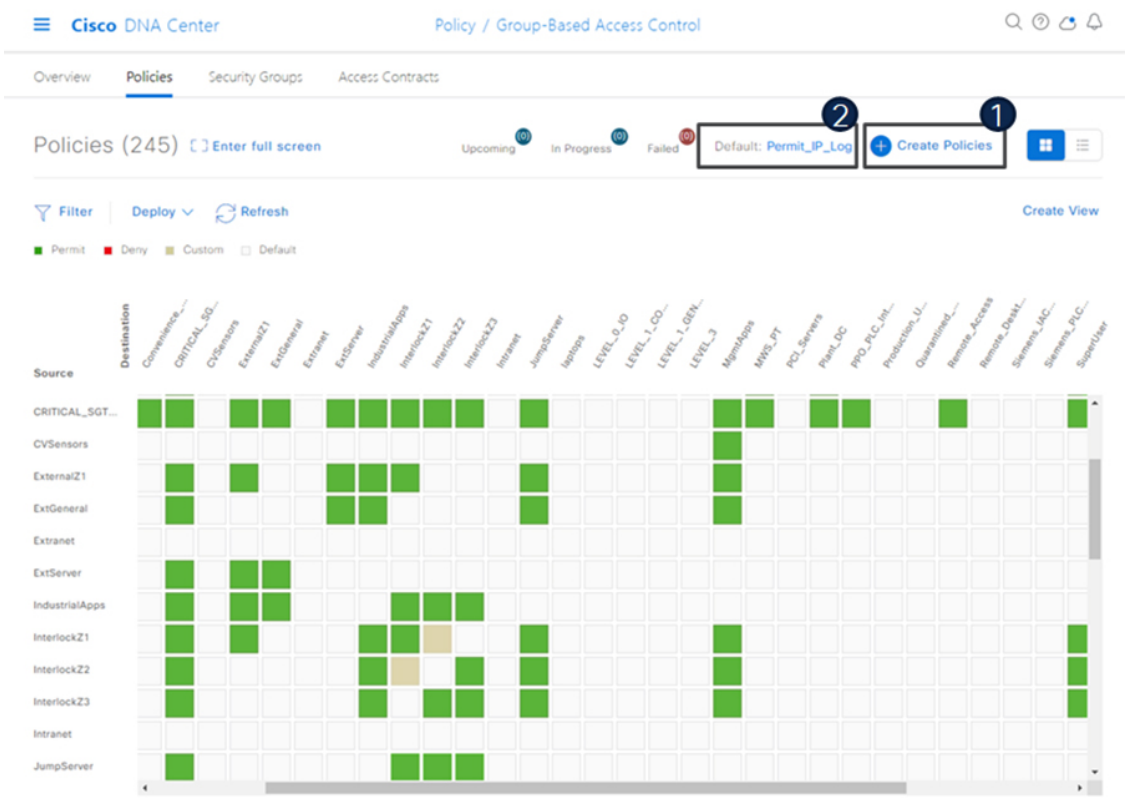

Define ISE as the AAA Server using Cisco DNA Center

When a device is provisioned in the inventory, Cisco DNA Center configures AAA server information, CTS authorization commands, and RADIUS server groups. In addition, Cisco DNA Center configures the device on the ISE PAN and propagates any subsequent updates for the device to the ISE PAN.

*Note: AAA server (ISE) settings for a given area should be configured in **Design > Network Settings > Network.***

1. From the DNA Center web interface, navigate to **Provision > Network Devices > Inventory**.
2. From the device list, check the box for the device to be provisioned.
3. From the **Actions** drop-down list, choose **Provision > Provision device**.
4. If the device is not assigned to a site, the wizard will show the **Assign Site** page. Click the **Choose a site** link and choose the desired Site. Click the **Save** button, then click the **Next** button. (Note that if Site assignment was done previously no action is needed here).
5. On the **Advanced Configuration** step, choose the device from the **Devices** list if there are any template settings to be configured. When finished, or if no template is applied, click the **Next** button.
6. On the **Summary** page, review the configuration to be added to the device. Click the **Deploy** button.

After the device has been provisioned, it will be in the device list of the specified Site.

Note: Provisioning a device that has already been configured with AAA before being discovered will fail. Remove any AAA configuration before pushing AAA using Cisco DNA Center.

Enable Device Tracking on Access Ports using Cisco DNA Center

Cisco DNA Center will automatically configure device tracking when a device is assigned to a site that has the wired client data collection enabled in its Telemetry settings (enabled by default). To verify the current setting, navigate to **Design > Network Settings > Telemetry**.

Configure Port-Based Authentication on the Access Switches

The following CLI output is provided as an example of policy. It can be deployed using Cisco DNA Center templates.

Example AAA Policy

```
class-map type control subscriber match-all
AAA_SVR_DOWN_AUTHD_HOST
match result-type aaa-timeout
match authorization-status authorized
!
class-map type control subscriber match-all
AAA_SVR_DOWN_UNAUTHD_HOST
match result-type aaa-timeout
match authorization-status unauthorized
!
class-map type control subscriber match-all
AI_IN_CRITICALSGT_AUTH
match activated-service-template IA_CRITICAL_SGT
!
class-map type control subscriber match-none
AI_NOT_IN_CRITICALSGT_AUTH
match activated-service-template IA_CRITICAL_SGT
!
class-map type control subscriber match-all DOT1X
match method dot1x
```

```
!
class-map type control subscriber match-all DOT1X_FAILED
match method dot1x
match result-type method dot1x authoritative
!
class-map type control subscriber match-all
DOT1X_MEDIUM_PRIO
match authorizing-method-priority gt 20
!
class-map type control subscriber match-all DOT1X_NO_RESP
match method dot1x
match result-type method dot1x agent-not-found
!
class-map type control subscriber match-all DOT1X_TIMEOUT
match method dot1x
match result-type method dot1x method-timeout
!
class-map type control subscriber match-any
IA_CRITICAL_SGT
match activated-service-template IA_CRITICAL_SGT
!
class-map type control subscriber match-all MAB
match method mab
!
class-map type control subscriber match-all MAB_FAILED
match method mab
match result-type method mab authoritative
!
policy-map type control subscriber IA_DOT1X_MAB_POLICIES
event session-started match-all
10 class always do-until-failure
10 authenticate using mab retries 3 retry-time 0 priority
10
20 authenticate using dot1x retries 3 retry-time 0
event authentication-failure match-first
5 class DOT1X_FAILED do-until-failure
10 terminate dot1x
20 authenticate using mab priority 20
10 class AAA_SVR_DOWN_UNAUTHD_HOST do-until-failure
10 activate service-template IA_CRITICAL_SGT
20 authorize
30 authentication-restart 60
40 pause reauthentication
20 class AAA_SVR_DOWN_AUTHD_HOST do-until-failure
10 authentication-restart 5
20 authorize
30 class DOT1X_NO_RESP do-until-failure
10 terminate dot1x
20 authenticate using mab priority 20
40 class MAB_FAILED do-until-failure
10 terminate mab
20 authentication-restart 60
60 class always do-until-failure
10 terminate dot1x
20 terminate mab
30 authentication-restart 60
event agent-found match-all
10 class always do-until-failure
10 terminate mab
20 authenticate using dot1x priority 10
event aaa-available match-first
10 class AI_IN_CRITICALSGT_AUTH do-until-failure
10 clear-session
20 class AI_NOT_IN_CRITICALSGT_AUTH do-until-failure
```

```
10 resume reauthentication
event violation match-all
10 class always do-until-failure
10 restrict
```

Example Interface Configuration using 'foreach' loops

```
#foreach($interface in $accessInterfaces)
interface $interface.portName
description endpoint
switchport access vlan $dataVlan
switchport mode access
device-tracking attach-policy IPDT_POLICY
#if($netflowPolicy)
ip flow monitor dnacmonitor input
#end
access-session port-control auto
mab
dot1x pae authenticator
spanning-tree portfast
service-policy type control subscriber
IA_DOT1X_MAB_POLICIES
service-policy input CIP-PTP-Traffic
service-policy output PTP-Event-Priority
#if($stormControl)
storm-control broadcast level 3 1
#end
exit
vlan $dataVlan
#end
```

```
#foreach($uplinkInterface in $trunkInterfaces)
interface $uplinkInterface.portName
description trunk
switchport trunk allowed vlan $vlans
switchport mode trunk
#if($cts)
cts manual
policy static sgt $uplinkSGT trusted
exit
exit
#end
vlan $vlans
#end
```

Configure Static Entries and Fallback Policy to Allow Communication in the event of an ISE error

The following configurations are recommended for a default deny policy to guarantee connectivity for critical services:

Change the SGT assigned to switches from "Unknown" to "TrustSec Devices" in ISE

By default, the "Unknown" SGT is configured for network device authorization and changing it to "TrustSec Device" gives more visibility and helps to create SGACLs specifically for switchinitiated traffic.

- From the ISE web UI, navigate to **Work Centers > TrustSec > TrustSec Policy > Network Device Authorization** and click the **Edit** link to update the Security Group.

Create static IP to SGT mappings on the TrustSec domain switches

Having local IP to SGT mappings ensures connectivity is up and connectivity to the critical resources are intact if connectivity to ISE is interrupted. In the example below ISE, DNAC, and the enforcement switch IP addresses are assigned the SGT for TrustSec devices (in this example 9043). Optionally, the subnet for the Cell/Area zone is assigned tag 911 to allow inter-Cell/Area zone communication for all devices when ISE is

not reachable. Once ISE is reachable again, mappings from ISE learned via SXP will take priority. The 911 tag should only be used when ISE is not available.

```
cts role-based sgt-map 10.13.48.132 sgt 9043
cts role-based sgt-map 10.13.48.184 sgt 9043
cts role-based sgt-map 10.17.10.1 sgt 9043
cts role-based sgt-map 10.17.10.0/24 sgt 911
```

Create a Fallback SGACL in the event ISE communication is lost

An SGT mapping is of no use until a relevant SGACL is assigned and hence our next step would be to create an SGACL that acts as a local Fallback in case ISE nodes go down (when ISE services are down, SGACLs and IP SGT mappings are not downloaded dynamically). In the example below we allow communication from the enforcement switch to critical services (ISE and DNA Center). Optionally, policies are created to allow external communication for all devices in the Cell/Area zone (911 tag).

```
ip access-list role-based FALLBACK
permit ip
cts role-based permissions from 9043 to 9043 FALLBACK
cts role-based permissions from 911 to 0 FALLBACK
cts role-based permissions from 0 to 911 FALLBACK
cts role-based permissions from 911 to 911 FALLBACK
cts role-based permissions from 9043 to 911 FALLBACK
cts role-based permissions from 911 to 9043 FALLBACK
```

Propagation on Distribution Switches and Core Switches

To ensure the SGT remains inside the packet throughout the TrustSec domain, configure inline tagging on links between the core and distribution switches.

Note: that this process may be disruptive since the interface bounces when configuring inline tagging. Plan accordingly to disrupt a single link at a time. When using port channels, remove the interfaces from the port channel, add configuration, and then add interfaces to the port channel again.

```
interface $uplinkInterface.portName
description trunk
switchport trunk allowed vlan $vlans
switchport mode trunk
cts manual
policy static sgt $uplinkSGT trusted
```

Each switch within the TrustSec domain must also be configured as an SXP listener. The speaker may be ISE or access switches connected below.

```
cts sxp enable
cts sxp default password 0 $sharedKey
cts sxp connection peer $peerIP source
$sourceIP.ipv4Address password default mode local listener
hold-time 0 0
```

Propagation on Industrial Switches

If using inline tagging in the industrial switches (for example, when using L2NAT), configure inline tagging on both ports of the link.

Note: that this process may be disruptive because the interface bounces when configuring inline tagging. To ensure connectivity is not lost, configure the farther switch first or use out of band connectivity. If configuring a port channel, links need to be removed from the port channel first and add back after configuration is completed.

```
interface $uplinkInterface.portName
description trunk
switchport trunk allowed vlan $vlans
switchport mode trunk
```

```
cts manual
policy static sgt $uplinkSGT trusted
```

If configuring SXP, refer to the following configurations:

- **Trustsec SXP – Speaker role, used when communicating bindings to upstream switches**

```
cts sxp enable
cts sxp default password 0 $sharedKey
cts sxp connection peer $peerIP source
$sourceIP.ipv4Address password default mode local speaker
hold-time 0 0
```

- **Trustsec SXP – Listener role, used when receiving bindings from ISE or access switches**

```
cts sxp enable
cts sxp default password 0 $sharedKey
cts sxp connection peer $peerIP source
$sourceIP.ipv4Address password default mode local listener
hold-time 0 0
```

Configure SXP in ISE

The following configuration creates a domain filter and adds an SXP device.

1. From the ISE web UI, navigate to **Work Centers > TrustSec > SXP > SXP Devices**.
2. Click the **Assign SXP Domain** link, even if no SXP devices are present.
3. On the **SXP Domain Assignment** window, click the **Create New SXP Domain** link.
4. Enter a name for the new domain.
5. Click the **Create** button.
6. Navigate to **Work Centers > TrustSec > SXP > SXP Devices**.
7. Click the **Add** button.
8. Enter the device details: name, IP address, SXP role (speaker), password type, SXP version, and connected PSNs for the peer device. You must also specify the SXP domain to which the peer device is connected.
9. Click the **Save** button.

Add an SXP Domain Filter

By default, session mappings learned from the network devices are sent only to the default group. You can create SXP domain filters to send the mappings to different SXP domains.

1. Navigate to **Work Centers > TrustSec > SXP > All SXP Mappings**.
2. Click the **Add SXP Domain Filter** link.
3. Enter the subnet details. The session mappings of the network devices with IP addresses from this subnet are sent to the SXP domain selected from the SXP Domain drop-down list.
4. From the **SXP Domain** drop-down list, choose the SXP domain to which the mappings must be sent.
5. Click **Save**.

Add IP-SGT Mappings to ISE

1. Navigate to **Work Centers > TrustSec > Components > IP SGT Static Mapping**.

2. Click the **Add** button.
3. Enter the IP address or hostname for a single device or use CIDR notation for subnets.
4. The **Map to SGT individually** radio button is chosen by default.
   a. From the **SGT** drop-down list, choose the SGT name.
   b. From the **Send to SXP Domain** drop-down list, choose the SXP Domain name. If left blank, the default domain is used.
   c. From the **Deploy to devices** drop-down list, select the grouping of devices to which the mapping should be deployed.
5. Click the **Save** button.

Enable TrustSec Enforcement on a Switch

```
cts role-based enforcement
cts role-based enforcement vlan-list $vlanList
```

Disable enforcement on uplink ports

```
interface $uplinkInterface.portName
no cts role-based enforcement
end
```

Create Profiling Rules in ISE

In this procedure, a custom Profiler Policy will be created for devices matching a specific Cyber Vision group.

1. Navigate to **Work Centers > Profiler > Profiling Policies** and click the **Add** button. The **Profiler Policy** page appears.
2. Complete the Profiler Policy form as follows:
   a. Assign a name.
   b. Chaeck the **Policy Enabled** check box
   c. Assign a certainty factor.
   d. Under **Rules**, from the **Conditions** drop-down list choose **Create New Condition** (**Advance Option**).
      1. From the **Expression** drop-down list, choose **Custom Attribute > assetGroup**.
      2. From the logic drop-down list, choose **Contains**.
      3. In the text field, enter the Cyber Vision group value. In this example the Cyber Vision group name is Interlock2.
   e. Enter the Certainty Factor value to be added if the Condition has been met.
3. Click **Submit**.

Figure 53: ISE Profiling Policy using Cyber Vision Group Data

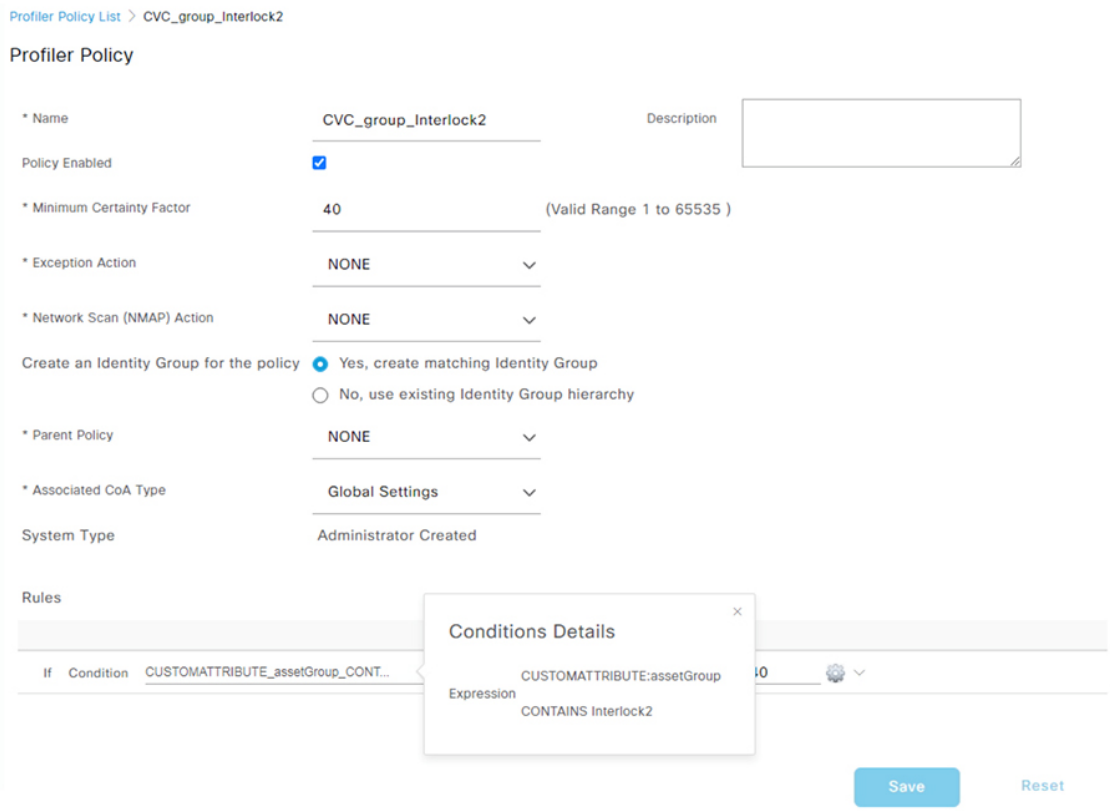

Note: follow the Integrating Cisco Cyber Vision with Cisco Identity Services Engine (ISE) via pxGrid *document to use Cisco Cyber Vision attributes for ISE profiling.*

Create Authentication and Authorization Policies on ISE

To configure the authorization policy in ISE, navigate to **Policy > Policy Sets > Default** and then choose **Authorization Policy**.

The following figure shows examples of authorization policies. The SuperUser rule (1) is an example of a policy that matches a user and assigns an SGT. The Interlock1 rule (2) is an example that matches an endpoint profile and assigns an SGT accordingly. The MABDefault rule (3) shows the default policy, which does not assign an SGT, so endpoints matching this rule will not override the default SGT assigned to the subnet of the Cell/Area zone.

Figure 54: Example Authorization Policies in ISE

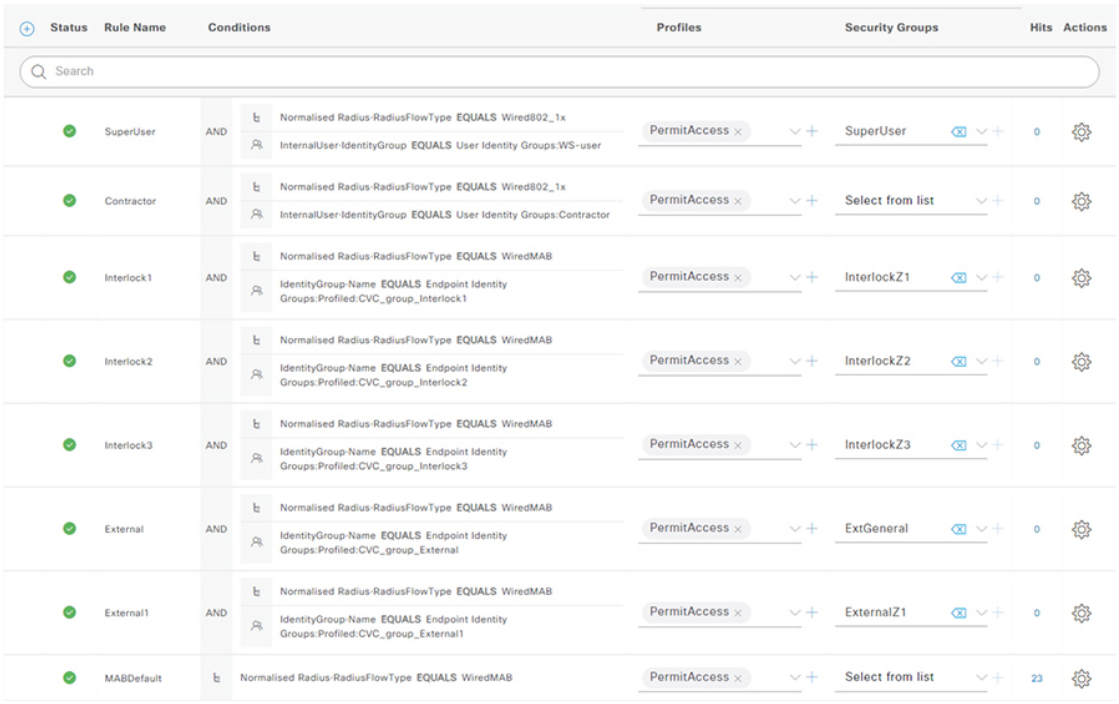

Cyber Vision Sensor

The following template can be used to provision the industrial switch to prepare for Cisco Cyber Vision sensor installation. For actual sensor deployment refer to Cisco Cyber Vision documentation.

```
#if ($enable_iox == 1)

iox

#MODE_ENABLE

terminal shell

sleep 30

sleep 30

terminal no shell

#MODE_END_ENABLE

#end

vlan 2

remote-span

interface AppGigabitEthernet 1/1

switchport mode trunk

exit

monitor session 1 source interface $intRange
```

```
monitor session 1 destination remote vlan 2

monitor session 1 destination format-erspan 169.254.1.2
```

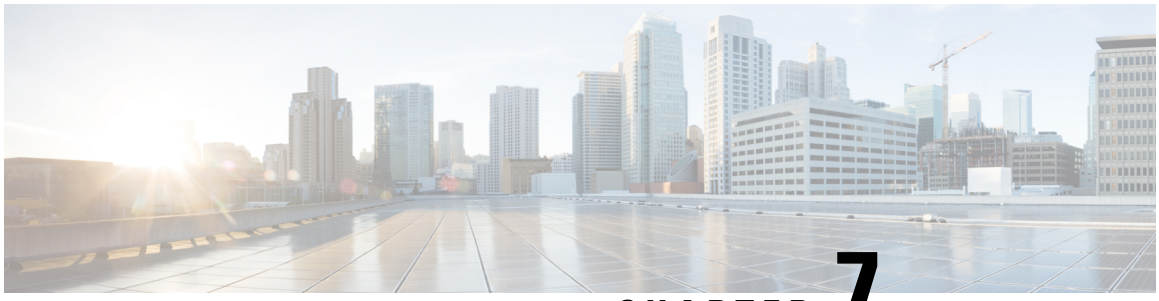

CHAPTER 7

Appendix C

Cisco Cyber Vision vs. Cisco Secure Network Analytics (formerly Stealthwatch)

Cisco Secure Network Analytics and Cisco Cyber Vision are two Cisco security offerings to provide visibility on the network. This section explains their different strengths and recommended role in the industrial network.

Cisco Secure Network Analytics provides enterprise-wide network visibility and applies advanced security analytics to detect and respond to threats in real time. Using a combination of behavioral modeling, machine learning, and global threat intelligence, Secure Network Analytics can quickly, and with high confidence, detect threats such as command-and-control (C&C) attacks, ransomware, distributed-denial-of-service (DDoS) attacks, illicit crypto mining, unknown malware, and insider threats. With a single, agentless solution, you get comprehensive threat monitoring, even if it is encrypted. Secure Network Analytics focuses on Enterprise IT networks and requires the packets to have an IP address. It is recommended for network devices in Levels 3 to 5 in the Purdue model.

Cisco Cyber Vision is an ICS visibility solution specifically designed to ensure continuity, resilience, and safety of industrial operations. It monitors industrial assets and application flows to extend IT security to the OT domain through easy deployment within the industrial network. It focuses on industrial networks and protocols. Cisco Cyber Vision has the capability of detecting Layer 2 flows and is recommended for Levels 0 to 3 in the Purdue model.

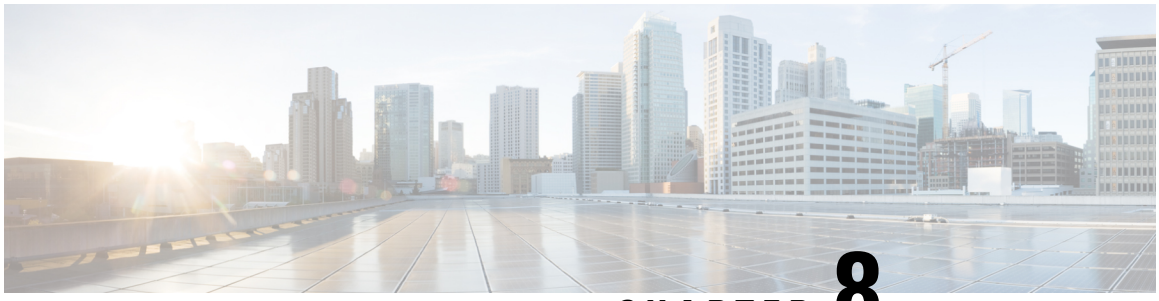

CHAPTER 8

Appendix D

Cisco Cyber Vision vs. Cisco Secure Network Analytics (formerly Stealthwatch)

Cisco Secure Network Analytics and Cisco Cyber Vision are two Cisco security offerings to provide visibility on the network. This section explains their different strengths and recommended role in the industrial network.

Cisco Secure Network Analytics provides enterprise-wide network visibility and applies advanced security analytics to detect and respond to threats in real time. Using a combination of behavioral modeling, machine learning, and global threat intelligence, Secure Network Analytics can quickly, and with high confidence, detect threats such as command-and-control (C&C) attacks, ransomware, distributed-denial-of-service (DDoS) attacks, illicit crypto mining, unknown malware, and insider threats. With a single, agentless solution, you get comprehensive threat monitoring, even if it is encrypted. Secure Network Analytics focuses on Enterprise IT networks and requires the packets to have an IP address. It is recommended for network devices in Levels 3 to 5 in the Purdue model.

Cisco Cyber Vision is an ICS visibility solution specifically designed to ensure continuity, resilience, and safety of industrial operations. It monitors industrial assets and application flows to extend IT security to the OT domain through easy deployment within the industrial network. It focuses on industrial networks and protocols. Cisco Cyber Vision has the capability of detecting Layer 2 flows and is recommended for Levels 0 to 3 in the Purdue model.

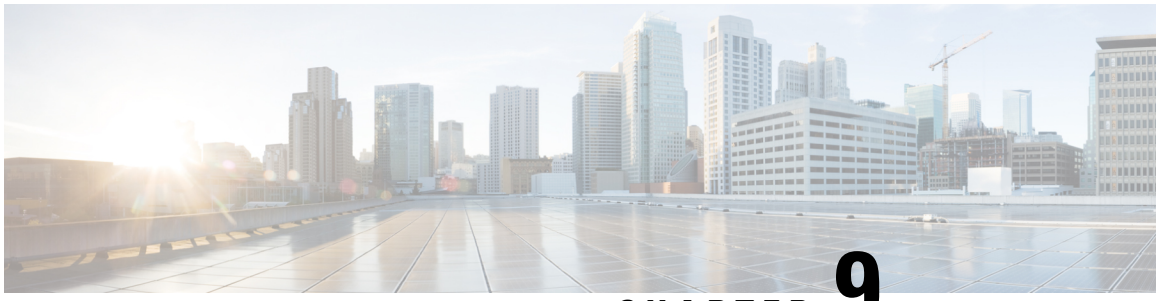

CHAPTER 9

Appendix E

Acronyms and Initialisms

Figure 55: Acronyms and Initialisms

AA	Authentication & Authorization
AAA	Authentication, Authorization, & Accounting
ACE	Access Control Entry
ACL	Access Control List
AD	Active Directory
AWS	Amazon Web Services
CERT	Computer Emergency Response Team
CIA	Confidentiality, Integrity, Availability
CIP	Common Industrial Protocol
CoA	Change of Authorization
CSDL	Cisco Secure Development Lifecycle
CTS	Cisco TrustSec
CVD	Cisco Validated Design
CVSS	Common Vulnerability Scoring System
DACL	Downloadable Access Control List
DBIR	Data Breach Investigations Report
DPI	Deep Packet Inspection
FTP	File Transfer Protocol
GUI	Graphical User Interface
HMI	Human-Machine Interface
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
HVAC	Heating, Ventilations & Air Conditioning
IACS	Industrial Automation and Control System
IDC	Industrial Data Center
IDMZ	Industrial Demilitarized Zone
IdP	Identity Provider
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IPDT	IP Device Tracking
IIoT	Industrial IoT
IO	Input/Output
IoT	Internet of Things
IP	Internet Protocol
IPDT	IP Device Tracking
IPS	Intrusion Prevention System
ISA	International Society of Automation
ISO	International Organization for Standardization
ISE	Cisco Identity Services Engine
IT	Informational Technology
KDB	Knowledge Database
LAN	Local Area Network
MAB	MAC Authentication Bypass
MAC	Media Access Control
MFA	Multi-Factor Authentication
MnT	Monitoring (ISE Node)
MTTD	Mean Time To Detect
MTTR	Mean Time To Respond
NAT	Network Address Translation
NGFW	Next Gen Firewall
NTP	Network Time Protocol
OCI	Oracle Cloud Infrastructure
OSI	Open Systems Interconnection
OT	Operational Technology
PAN	Policy Administration Node (ISE)
PLC	Programmable Logic Controller
PSN	Policy Service Node
PxGrid	Cisco Platform Exchange Grid
RADIUS	Remote Authentication Dial-In User Service
RDP	Remote Desktop Protocol
REP	Resilient Ethernet Protocol
RSPAN	Remote SPAN
RTU	Remote Terminal Unit
SGACL	Security Group Access Control List
SGT	Security Group Tag
SIEM	Security Information and Event Management
SIS	Safety Instrument System
SOAR	Security Orchestration, Automation & Response
SSL	Secure Sockets Layer
SSM	Smart Software Manager
SISF	Switch Integrated Security Features
SXP	SGT Exchange Protocol
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VRF	Virtual Routing Function
XDR	Extended Detection and Response

388082