




Industrial Security 3.1

Design Guide

June 2025



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS DESCRIBED IN THIS DOCUMENT ARE SUBJECT TO CHANGE WITHOUT NOTICE. THIS DOCUMENT IS PROVIDED "AS IS."

ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED, OR STATUTORY INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, PUNITIVE, EXEMPLARY, OR INCIDENTAL DAMAGES UNDER ANY THEORY OF LIABILITY, INCLUDING WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OF OR INABILITY TO USE THIS DOCUMENT, EVEN IF CISCO HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

©2025 CISCO SYSTEMS, INC. ALL RIGHTS RESERVED

Table of Contents

Introduction.....	8
Target Audience.....	8
Security Challenges	9
Understanding the threats	10
Industrial Security Journey.....	11
Chapter 1. Gaining Visibility and Understanding your Security Posture	14
Why understanding security posture is necessary	14
How to gain Visibility of the OT assets	15
An Alternative to SPAN.....	16
Active Discovery	17
Intrusion Detection / Prevention Systems	19
Vulnerability Assessment and Managing Risk.....	19
Vulnerability Assessment	20
CVSS Scores	20
Cisco Cyber Vision	21
Components	21
Key Features	21
Cyber Vision Design Considerations	22
Cyber Vision Center	22
Cyber Vision Sensor Options	24
Effective Sensor Deployment	25
Ring Topology Considerations	26
Brownfield Deployment Considerations	27
Sensor Considerations	27
Cyber Vision Active Discovery	29
Vulnerability Assessment in Cyber Vision.....	29
Cyber Vision Risk Score.....	30
Visualize Assets and Flows using Cyber Vision Groups.....	31
Presets and Baselines	33
Cyber Vision IDS.....	35
Performance	35

Licensing	35
Chapter 2. Preventative Controls in Plant Networks	36
Reference Architecture	36
Industrial Zone.....	37
Enterprise Zone	38
Industrial DMZ	38
Use cases	39
How to get started with Segmentation	48
IT/OT Boundary	48
Industrial Data Center	50
Preventing lateral movement on the plant floor	52
ISA/IEC 62443 Zones and Conduits model for OT Segmentation	52
Architecting a network with secure components	54
Segmentation Technologies	55
VLAN	55
VRF-lite	55
Access Control List.....	56
Stateful Firewall.....	56
Next-Generation Firewall	56
TrustSec	57
Segmentation Design Guidance with Cisco Secure Firewall	57
Cisco Secure Firewall portfolio	58
Management options	58
Deployment modes	59
Key Capabilities	60
Industrial Data Center Segmentation with Cisco Secure Firewall	69
East/West Segmentation with Cisco Secure Firewall	71
Using a firewall to route between OT VLANs	71
Transparent firewalls at the Cell/Area zone boundary	74
Network Access Control with Cisco Identity Services Engine	75
ISE Components / Personas.....	75
ISE Authentication Policies	76
ISE Authorization Policies	77
ISE TrustSec Domain	77
SGT Classification.....	77

SGT Transport.....	78
SGT Enforcement.....	79
SGT Example.....	80
Macro-Segmentation with ISE.....	81
Micro-Segmentation using ISE.....	82
ISE / SGT Design Considerations.....	83
Scale Considerations in Large Networks.....	92
Segmentation when the layer 3 boundary also participates in layer 2 connectivity	94
NAT Considerations	95
SXP Domain Filters.....	96
Static Segmentation in the Industrial Zone	97
Applying Policy to Users.....	98
<i>Chapter 3. Preventative Control in Distributed Field Networks</i>	<i>100</i>
Reference Architecture	100
Use Cases.....	104
Cyber Resiliency vs. Cyber Security	108
Cisco Industrial Router	109
Portfolio.....	109
Management Options.....	110
Application Firewall.....	111
Network Segmentation	112
Authentication, Authorization, and Accounting.....	114
Port Security	114
TrustSec	114
Denial of Service Protection	118
Application Hosting with IOx	120
NGFW Add-On	120
TLS Decryption.....	123
Secure Access Service Edge (SASE).....	123
Plug and Play	124
Cisco Industrial Router Design Guidance	125
Understanding defense-in-depth	125
Deny by default firewall configuration.....	127
When to use Snort.....	128
Cloud Connectivity.....	129

End-to-end segmentation with TrustSec	129
Multiple IOx applications in a single deployment	132
Chapter 4. Secure Remote Access for Industrial Networks	138
Remote Access Technologies	138
Virtual Private Networks	138
Accessing Jump Servers with the Remote Desktop Protocol	140
Zero Trust Network Access	141
Secure Equipment Access Design Guidance	145
Cisco Secure Equipment Access	145
Cisco SEA Architecture Guidance	148
Cisco SEA Authentication Options	157
Cisco SEA Policy Creation Guidance	161
Cisco SEA Documentation	163
Chapter 5. Cross-Domain Detection, Investigation and Response	164
Introducing Splunk	165
Components of Splunk Enterprise	166
Data Collection Tier	167
Indexing Tier	167
Search Tier	167
Management Tier	168
Splunk Enterprise Concepts and Features	169
Search Processing Language	169
Apps	172
Alerts, Dashboard & Reports	174
Splunk Design Considerations	176
Install Considerations	176
Getting Data In	177
Index Design and Life-Cycle Management Recommendations	183
Using the Search & Reporting app	186
Splunk's Alert Framework	198
Building Custom Dashboards and Visualizations	200
Using Dashboards to build Reports	205
Splunk Enterprise Security	207
OT Security Add-On for Splunk	207
Appendix A – Deployment Guides	209

IDMZ.....	209
Cyber Vision.....	209
ISE	209
XDR	210
Appendix B – Example TrustSec Configuration in Plant Networks	211
Define and Create SGTs and Policies Using Cisco Catalyst Center.....	212
Define ISE as the AAA Server using Cisco Catalyst Center	213
Enable Device Tracking on Access Ports using Cisco Catalyst Center	214
Configure Port-Based Authentication on the Access Switches	214
Example AAA Policy	214
Example Interface Configuration using ‘foreach’ loops	216
Configure Static Entries and Fallback Policy to Allow Communication in the event of an ISE error. 218	
Change the SGT assigned to switches from “Unknown” to “TrustSec Devices” in ISE.....	218
Create static IP to SGT mappings on the TrustSec domain switches.....	218
Create a Fallback SGACL in the event ISE communication is lost	218
Propagation on Distribution Switches and Core Switches.....	219
Propagation on Industrial Switches	219
Configure SXP in ISE	220
Add an SXP Domain Filter	220
Add IP-SGT Mappings to ISE	221
Enable Trustsec Enforcement on a Switch	221
Disable enforcement on uplink ports	221
Create Profiling Rules in ISE	221
Create Authentication and Authorization Policies on ISE.....	223
Cyber Vision Sensor	223
Appendix C – Installing custom Snort rules in Cisco Catalyst SD-WAN	225
Upload and use file in Cisco Catalyst SD-WAN Manager.....	227
Appendix D – Cisco Cyber Vision vs. Cisco Secure Network Analytics.....	228
Appendix E – Cisco SecureX.....	229
SecureX Ribbon.....	229
SecureX Threat Response.....	230
SecureX Orchestration	231
Appendix F – Cisco Cyber Vision and Splunk Quick Deployment Guide	234

<i>Appendix G – Cisco ISE and Splunk Quick Deployment Guide</i>	<i>246</i>
<i>Appendix H - References</i>	<i>250</i>
<i>Appendix I - Acronyms and Initialisms</i>	<i>251</i>

Introduction

The Cisco [2024 State of Industrial Networking Report](#), which had over 1000 respondents across 17 countries operating in more than 20 sectors, found cybersecurity to be the biggest reported challenges in running and maintaining industrial networks.

There is a clear sense that artificial intelligence (AI) will boost business growth for those who can successfully use it to run better industrial networks. Leaders will ensure their operational technology can capture the required data to fuel their AI models. However, increased network connectivity leads to a continually expanding attack surface. In 2023, the [worlds critical infrastructure suffered 13 cyber-attacks every second](#), and in 2024, cyber-attacks on critical infrastructure [surged by 30%](#).

Safeguarding industrial automation and control systems (IACS) from cyber threats is a critical priority, but transforming these intentions into effective actions can be challenging. Given the complexity of IACS and their networks, which often rely on outdated technologies and inadequate security measures, it can be difficult to determine the best starting point.

For over 20 years, Cisco has been helping industrial organizations around the globe digitize their operations by developing a market-leading networking and security 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.

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, and the information technology (IT) teams understand the network, and the security team understands threats and vulnerabilities. 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.

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 organization should be identical, and the path forward involves finding a 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.

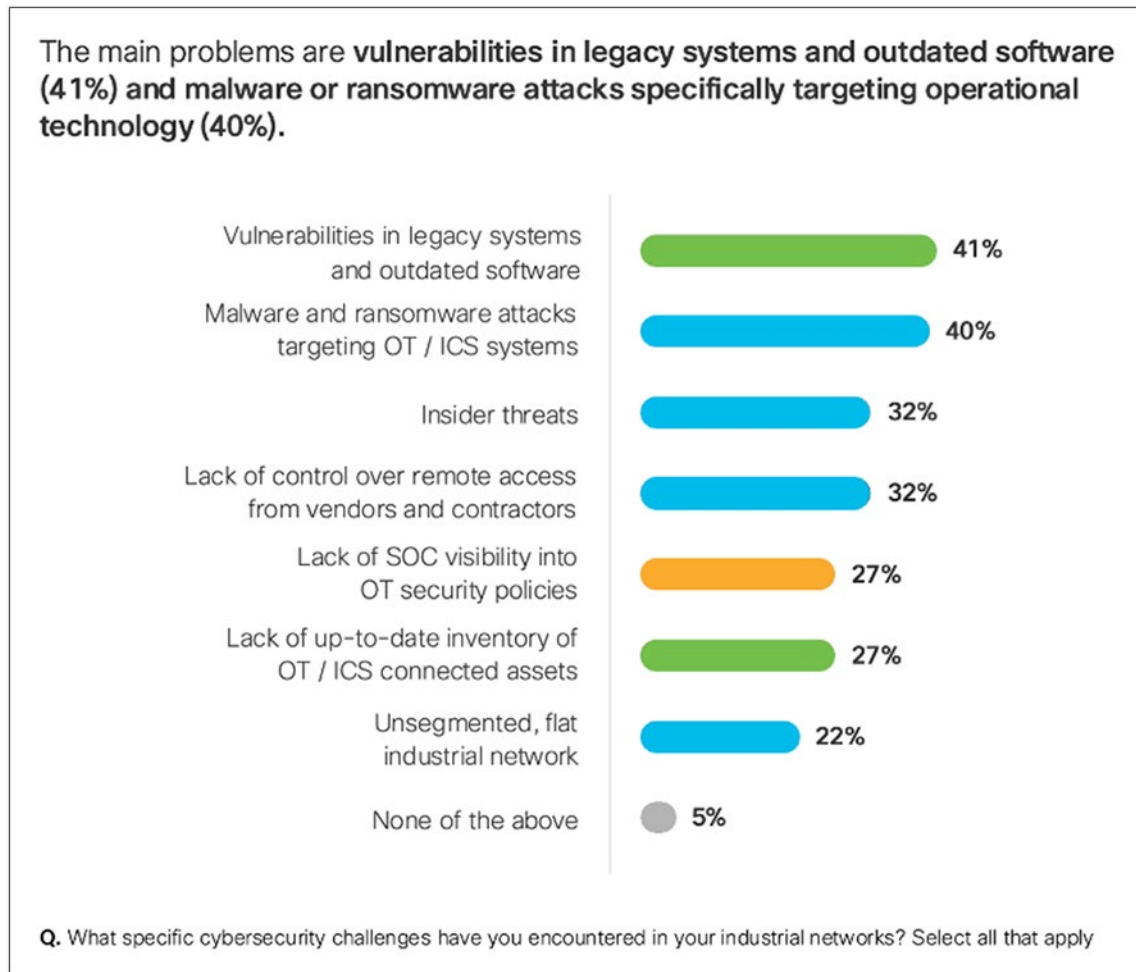
The [Cisco Industrial Threat Defense](#) solution is intended to be used by IT, OT, and security teams and their relevant partners and system integrators. Operations will appreciate the ease of use and simple deployment, as well as the broad support of various IACS vendors and protocols. IT network managers 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 part of the State of Industrial Networks Report, Cisco asked operational leaders what specific cybersecurity challenges they have encountered in their industrial networks.

Figure 1 Cybersecurity challenges faced in industrial networks



The first challenge faced when securing the IACS network is vulnerabilities in legacy systems and outdated software (41%). As industrial networks can be quite old, and patching is almost non-existent, vulnerabilities exist. However, without visibility into what vulnerabilities our networks are exposed to its very hard to contain them. Firewalls are deployed at the edge of our critical infrastructure to “virtually patch” our devices. If the vulnerabilities are unknown, then firewall rules cannot be optimized, or patching cannot be prioritized.

Additionally, operators often don’t have an accurate inventory of what is on the network (27%). 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. This leads to the second biggest concern for operational leaders; legacy networks that are not equipped to prevent the spread of malware and ransomware attacks targeting OT systems (40%). 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 also telling that operational leaders are equally concerned with insider threats (32%) and a lack of control over remote access from vendors and contractors (32%). Although the threat of remote users is more commonly exploited, operational leaders recognize that once any user gets a foothold in the network there is a lack of control on the damage they could potentially cause.

Last, but certainly not the least, is the lack of SOC visibility into OT security policies (27%). Compounding the cybersecurity challenges faced by industry, we face a workforce challenge such as employee retention and a shortage of cybersecurity practitioners. Organizations cannot afford to deploy solutions in silos, bringing new screens to every new capability deployed. Any solution brought to the OT network must have linkages back to a security operations centre (SOC) where events and logs can be analysed in the existing infrastructure already deployed by IT personnel. Isolated screens and scattered event logs lead to missed incidents and additional work to correlate data across the entire ICS attack chain.

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 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 2 shows four common attack techniques described in the MITRE ATT&CK framework.

Figure 2 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 as top of mind.

In the case where initial access security is 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 the 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 on. 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 environment, network segmentation helps solve this problem too, by containing an adversary to the zone in which the initial exploit occurs.


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](#).

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 3 Industrial Security journey for plant networks



The first step is to gain **asset visibility and security posture** with Cisco Cyber Vision. Cisco Cyber Vision provides asset owners full visibility into their industrial networks and their OT security posture so they have the information they need to reduce the attack surface, segment the industrial network, and enforce cybersecurity policies. Cyber Vision helps answer questions such as; what vendors exist on the network? Are there devices that I do not recognise? What vulnerabilities can be exploited in the environment? What devices are communicating with external networks? Do these devices with heightened exposure also have a path to the critical services in the network?

Combining a unique edge architecture that embeds deep packet inspection (DPI) into your industrial network, and integration with the Cisco leading security portfolio, Cisco Cyber Vision can be easily deployed at scale to enable IT and OT teams to work together in building innovative industrial operations while securing the global enterprise.

While visibility is important, taking preventative measures to secure your operations is required. Cisco recommends taking a **zero trust approach to securing the industrial network**. Secure Equipment Access (SEA) combines all the benefits of a Zero-Trust Network Access (ZTNA) solution with a network architecture that makes it simple to deploy at scale in operational environments. There is no dedicated hardware to install and manage. No complex firewall rules to configure and maintain. The Cisco industrial switches or routers that connect your OT assets now also enable remote access to them. And it features comprehensive security capabilities, with advanced cybersecurity controls and easy-to-build policies based on identities and contexts.

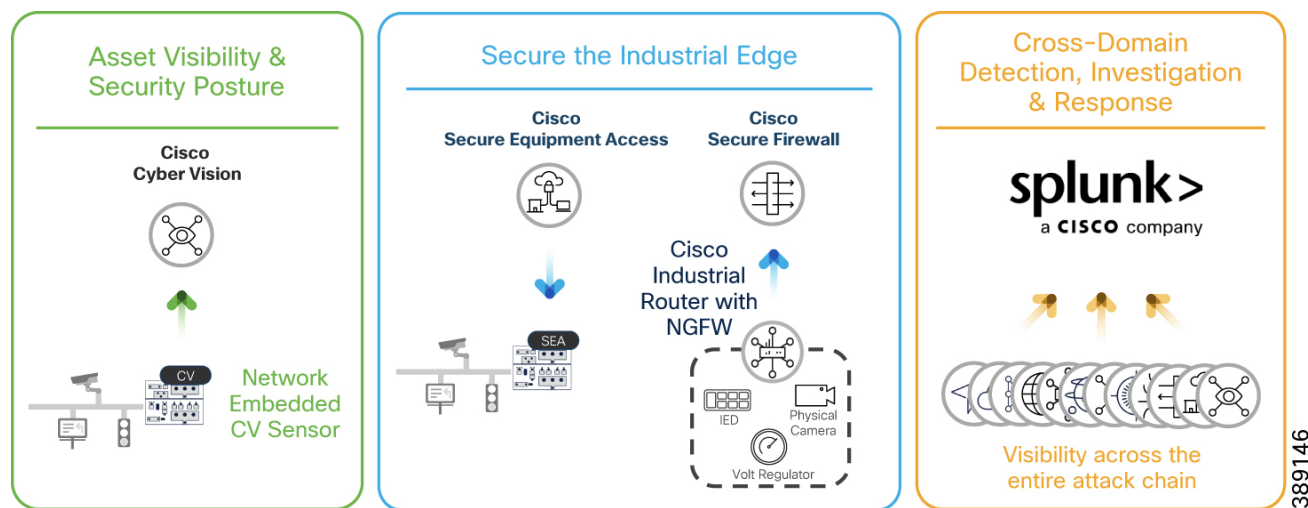
When implementing a zero-trust model within an OT network, Cisco recommends building zones of trust, aligning to the IEC 62443 framework for zones and conduits. The details of this will be explored later in the guide, but in short, different use cases require different technology stacks. For example, data centre (DC) modernisation in plant networks is resulting in more compute, more connectivity, and more virtualisation of the plant floor. While legacy systems will continue to operate as normal, new systems are being introduced in the plant data centre than require additional connectivity to the cloud, resulting in more exposure to threats. The industrial DC needs to be firewalled from the plant floor, and the Cisco Secure Firewall, using an integration with Cisco Cyber Vision enables contextual policies to be created between OT devices and DC workloads.

For plant floor segmentation, a network access control solution is better suited for implementing policy. Using Identity Services Engine (ISE), logical segmentation can be implemented in the network infrastructure. The same boxes that provide connectivity and visibility into the network can also provide the control needed for critical infrastructure protection. Together, Cisco ISE and Cisco Cyber Vision offer an ideal solution for operations and IT teams to work together in implementing policies that will limit communications between industrial assets without having to modify network setups or impacting production. Malicious traffic and cyberattacks can now be contained, and network resources are optimized to improve production efficiency.

To solve the challenge of scattered event logs across the ecosystem, **Splunk provides cross-domain detection, investigation and response**. Splunk is a leading security information and event management (SIEM) solution that provides the detection, analytics, case management, incident response, and orchestration platforms all in one interface. Splunk ingests data from Cyber Vision to provide visibility into OT, and correlates that with other data sources like network access control, NGFW, among others, to provide a holistic view of the entire OT Network from endpoints in the LAN, to egress/ingress points and all the way to the data centre by across a multi-vendor environment.

This technology stack works well for industrial automation networks, however, when dealing with distributed field assets, a modified approach must be taken.

Figure 4 Industrial security journey for distributed field networks



The fundamentals remain the same; visibility, control and cross-domain detection. The change comes from the control points we have available in field networks. Across industries such as roadways and utilities, organizations need advanced, agile, and secure Wide Area Network (WAN) infrastructures to connect distributed OT assets to control centres and unlock the potential of digitization. Whether it is about connecting roadways assets, first responder or public transport vehicles, water, oil, or gas infrastructures, renewable energy resources, power substations, EV charging stations, or any critical remote assets, you need rugged routers with cutting-edge cybersecurity capabilities. Not only do Cisco Catalyst Industrial Routers offer unconditional connectivity for all your remote assets but come with a comprehensive next-generation firewall (NGFW) features and many more cybersecurity capabilities to block modern threats.

Chapter 1. Gaining Visibility and Understanding your Security Posture

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 analysing the information to differentiate between known and unknown communication may be a necessary first step in implementing security policies.

Note: While much of the guidance supplied is vertical agnostic, this version of the design guide focused on the deployment of Cyber Vision in plant networks. Future iterations of the design guide will include specific guidance to verticals such as utilities and transportation.

Why understanding security posture is necessary

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 behaviours.

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 what 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 of 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:

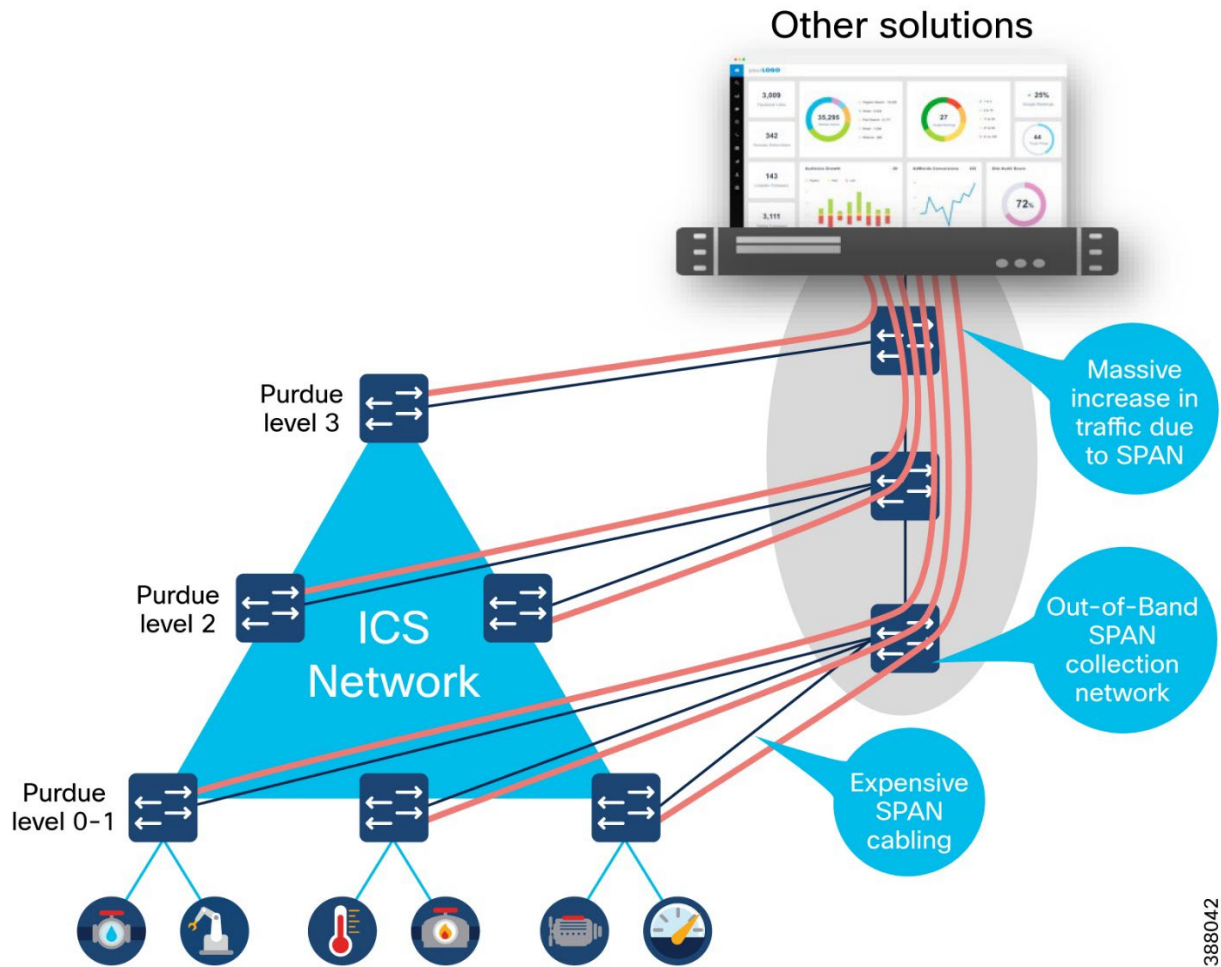
- Send all traffic to a central server that performs DPI
- Deploy dedicated sensor appliances on each network switch
- Send traffic to dedicated sensor appliances deployed here and there on the network

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 5 OT visibility using a SPAN network

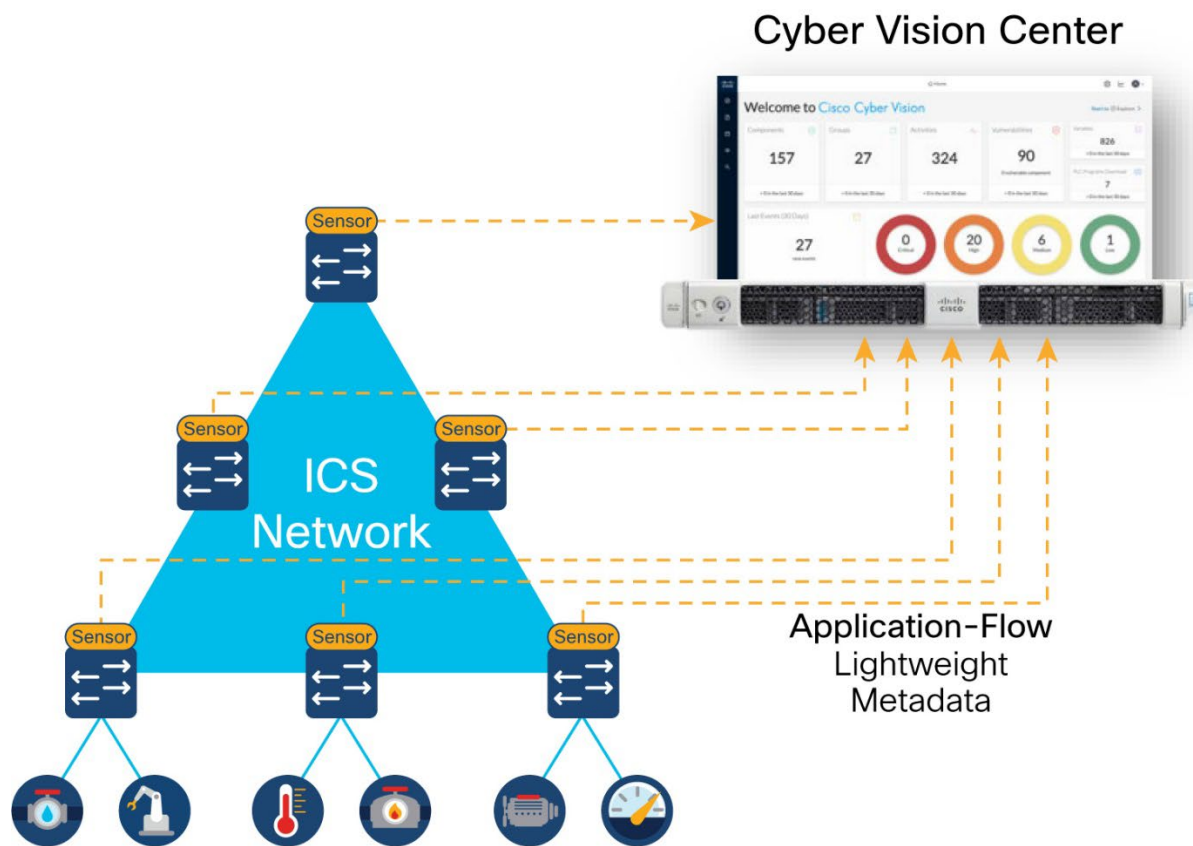


388042

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 6 OT Visibility using sensors embedded in the switch infrastructure

388043

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, and so on) 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, and so on.) 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.

$$\text{Threat} \times \text{Vulnerabilities} \times \text{Impact} = \text{Risk}$$

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 safety priority. Risk limitation limits a company 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 between the announcement of a security vulnerability in a particular technology and the availability of an attack method (an exploit). Within this time, 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 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. 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 2-5% 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.

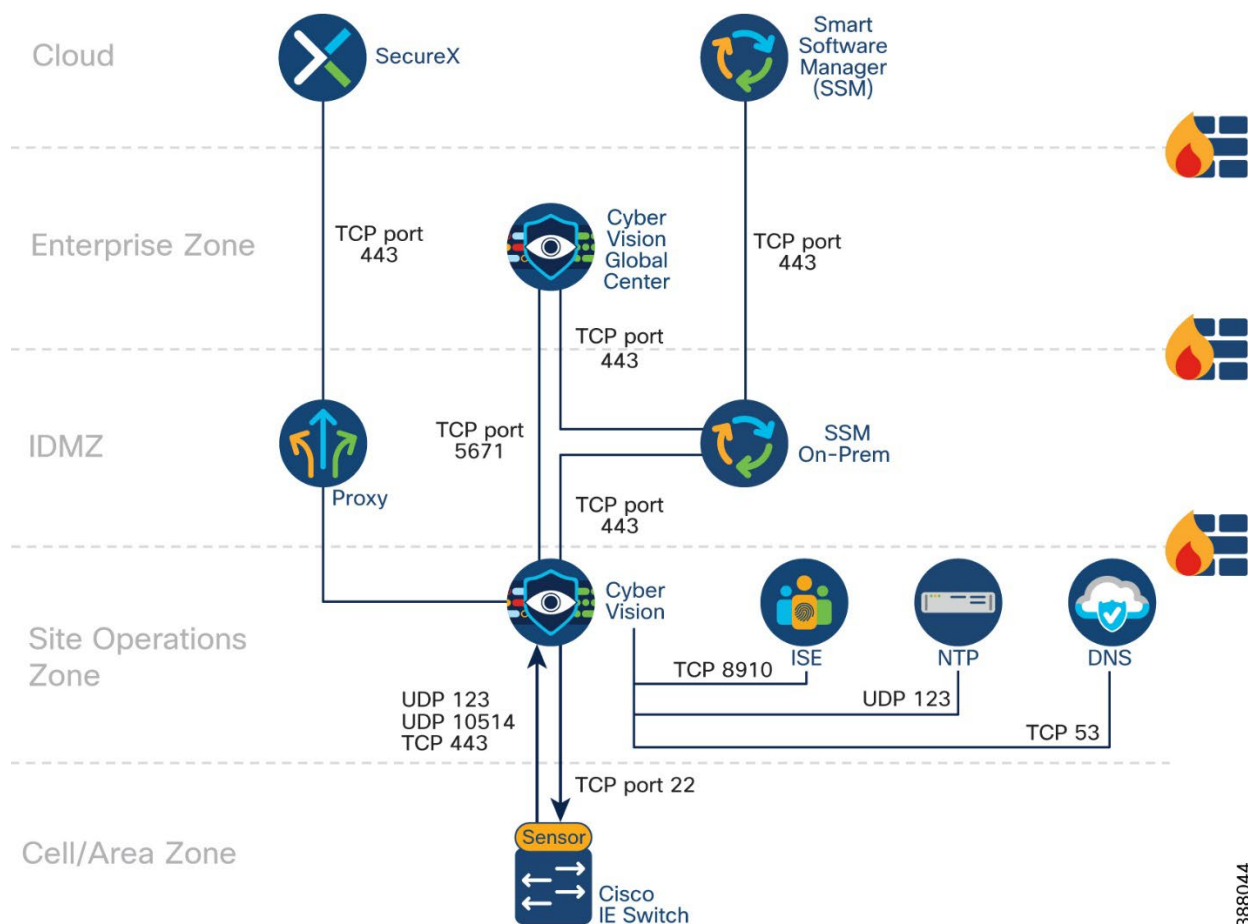
For more information on Cisco Cyber Vision see the [Cisco Cyber Vision Datasheet](#).

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. Figure 7 depicts the communication flows from Cisco Cyber Vision center used in this design guide.

Figure 7 Cyber Vision communication flows



388044

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.

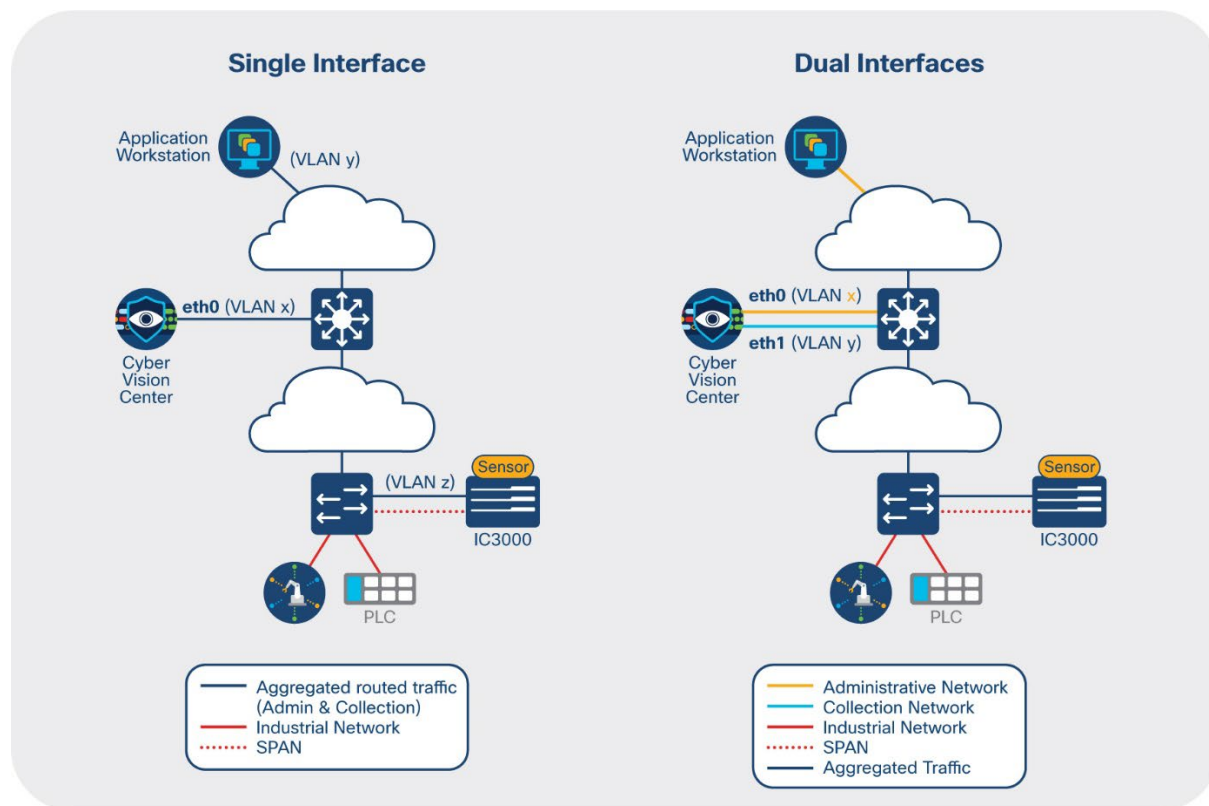
Ethernet interfaces are allocated in the following way:

- 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, and so on)
- 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 respectively allocated to the admin and collection networks 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 8 Cyber Vision Center deployment modes



388045

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 that follows.

Table 1 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

Note: 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, a lot of the value of Cyber Vision is lost when deployed only on 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 IO 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 traffic either from all or select switches depending on if you want option 1 or 2 from this list. This model requires additional cabling from every device to the out of band sensor and a free switch port must be available on the switches you wish to SPAN traffic from

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.

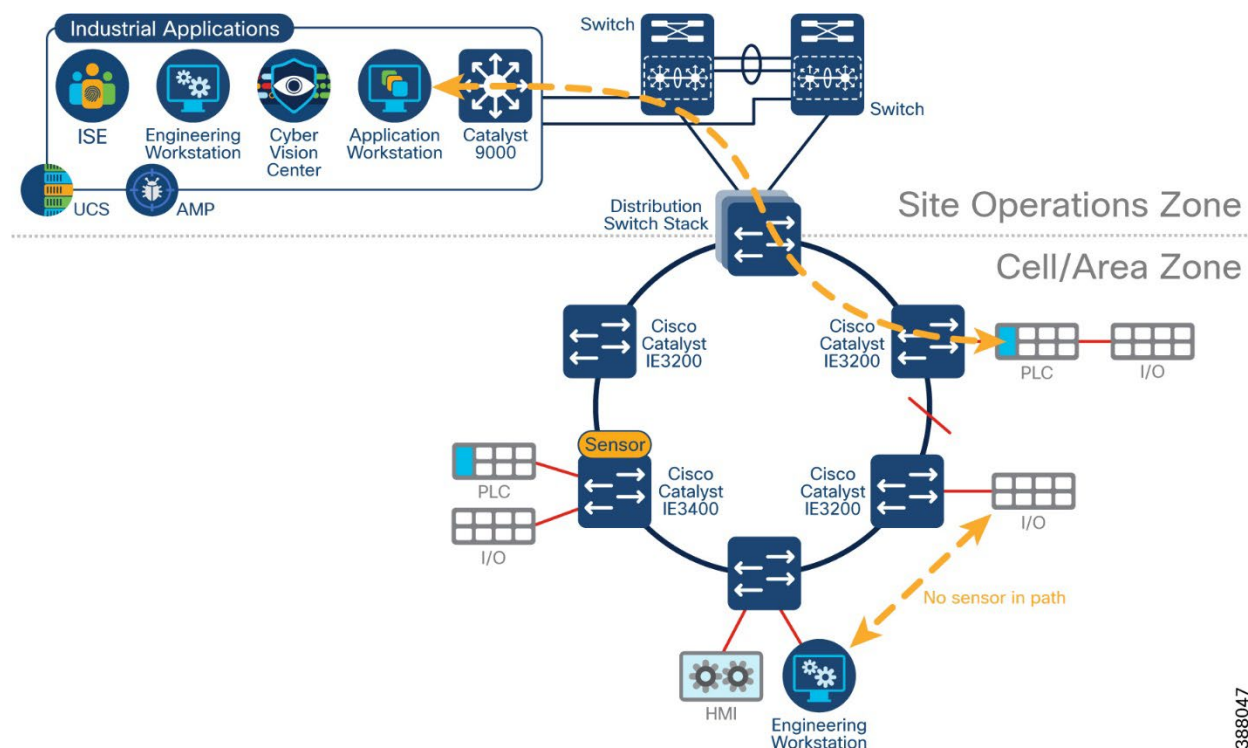
Warning: Be careful when collecting data at higher levels (distribution level), especially if Internet traffic is being monitored. Monitoring Internet flows in addition to traffic on the industrial network will significantly increase the number of devices (and components) present in the Center database.

Ring Topology Considerations

Visibility of flows in a ring may change depending on sensor positioning and the active traffic path. Figure 9 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 9 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 and send 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 10 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.

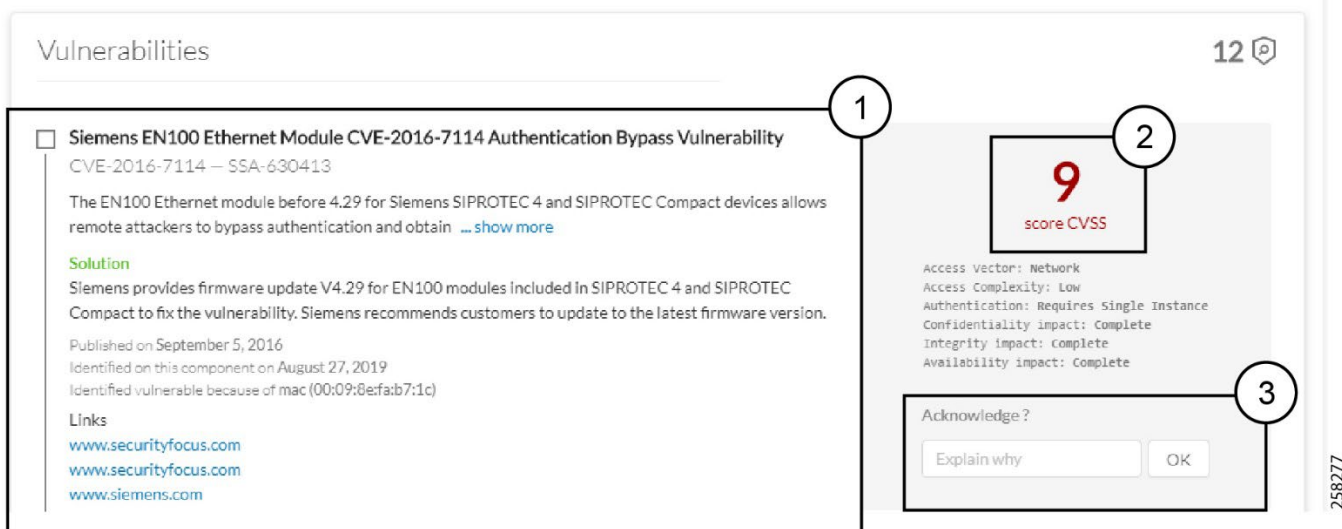
Cyber Vision Active Discovery is not handicapped by the presence of NAT.

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 11 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 don't 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/cancellation 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:

Table 2: Cyber Vision Risk Score by Color

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

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:

$$\text{Risk} = \text{Impact} \times \text{Likelihood}$$

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 has the possibility to act on 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.

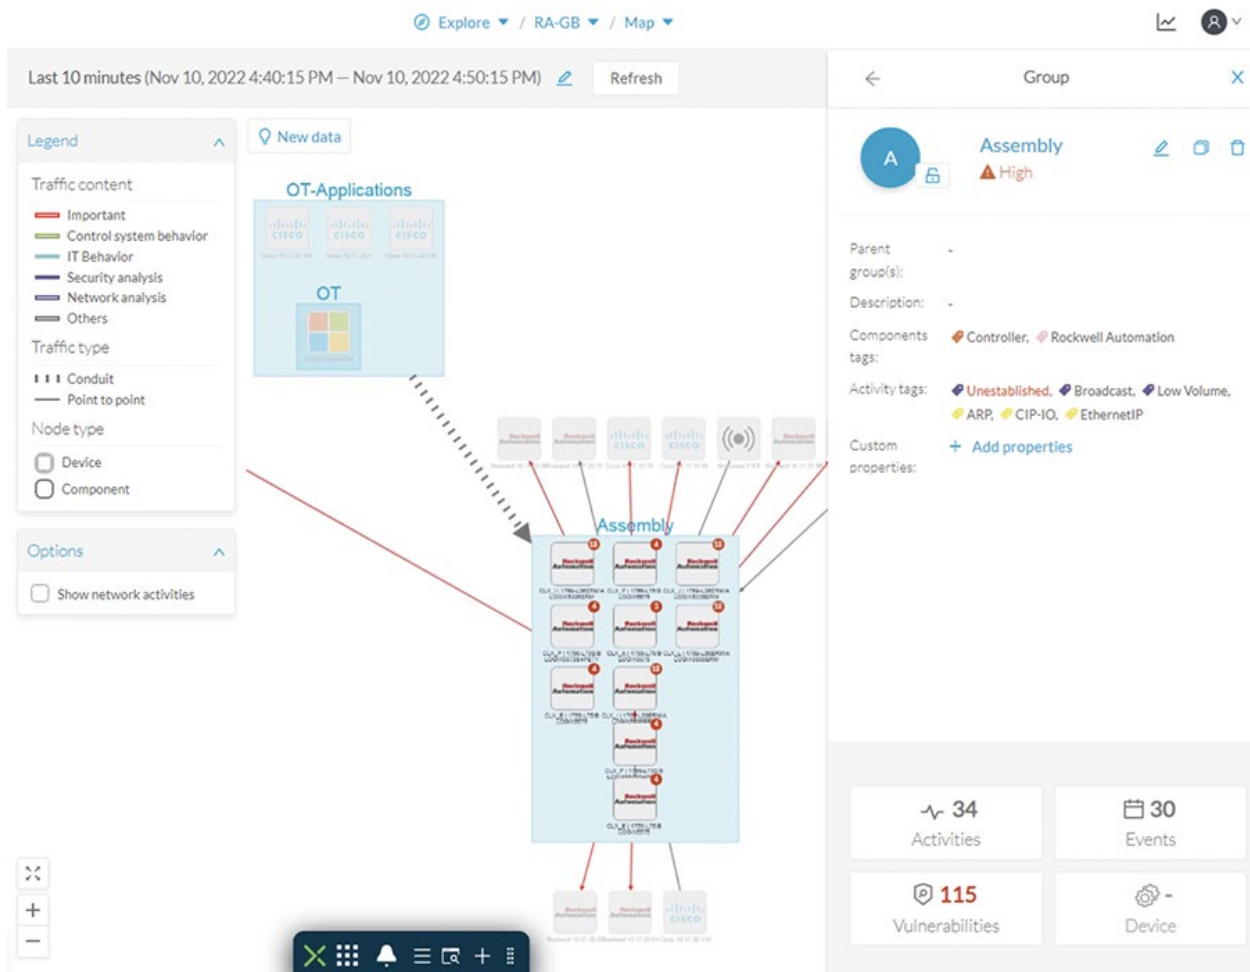
Visualize Assets and Flows using Cyber Vision Groups

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. Figure 12 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 12 Cyber Vision Network Map 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. Figure 13 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 13 Cyber Vision deviations from Baseline

Status	Component	Group	First activity	Last activity	IP	MAC
NEW	Rockwell 10.17.90.31	-	Oct 7, 2022 12:27:06 PM	Oct 7, 2022 2:00:05 PM	10.17.90.31	00:1d:9cc4f1:50
NEW	cpwe_L85E 1756-L85E/B (Port1-Link00)	3400-3_devices	Oct 7, 2022 1:10:46 PM	Oct 7, 2022 2:00:04 PM	10.17.90.31	00:1d:9cc4f1:50
NEW	cpwe_L73S 1756-L73S/B LOGIX5573SAFETY (Port1-Link01)	3400-3_devices	Oct 7, 2022 1:10:46 PM	Oct 7, 2022 2:00:01 PM	10.17.90.31	00:1d:9cc4f1:50
CHANGED	1756-EN2TR/B	3400-3_devices	Sep 26, 2022 2:16:06 PM	Oct 7, 2022 1:07:04 PM	10.17.90.31	00:1d:9cc4f1:e1
CHANGED	cpwe_L85E 1756-L85E/B	3400-3_devices	Sep 26, 2022 2:16:06 PM	Oct 7, 2022 1:07:02 PM	10.17.90.30	00:1d:9cc4f1:d6

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 too dynamic in nature.

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 as depicted in figure 14. 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.

Figure 14 Cyber Vision Baseline Investigation

The screenshot displays the Cisco Cyber Vision interface for baseline investigation. The main table shows one activity with the following details:

Status	Component	Component	First activity	Last activity	Tags	Flows	Packets	Volume
NEW	Vmware 10.13.48.175	cpwe_L735/1756-L85E/B	Nov 10, 2022 5:12:26 PM	Nov 11, 2022 10:38:50 AM	Program Download, Start CPU, Read Var, EthernetIP	~10	482	

The 'Properties' panel for the selected activity shows the following details:

Property	Value
enip-class	unknown-8x84
enip-general-status	PartialTransfer
enip-request	true
enip-couname	cpwe_L735
enip-device-type	ProgrammableLogicController
enip-event	Download
enip-event-failed	true
enip-location	Endpoint
enip-log-description	Downloaded project [C:\USERS\ADMINISTRATOR\DOCUMENTS\STUDIO1580\PROJECTS\cpwe_L735.ACD] to [\\AB_ETH-2110.17.00.31\Backplane\1\cpwe_L735]
enip-log-extended-info	Property List: Inhibit state (un-Inhibited)
enip-log-txid	WIN-339TIVCV308\ADMINISTRATOR
enip-log-username	WIN-339TIVCV308\ADMINISTRATOR
enip-log-workstation	WIN-339TIVCV308
enip-name	1756-L735/B LOGIX5573SAFETY
enip-productcode	8x84
enip-serial	60ad62fb
enip-status	AtLeastOneIOConnectionEstablishedAllIdle,ReservedBits12-15:0x3
enip-status-major	REH
enip-status-minor	PROG
enip-value	STOP
enip-vendor	Rockwell Automation/Allen-Bradley
enip-version	34.011
ethertype	IPv4
protocol	TCP

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: 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.

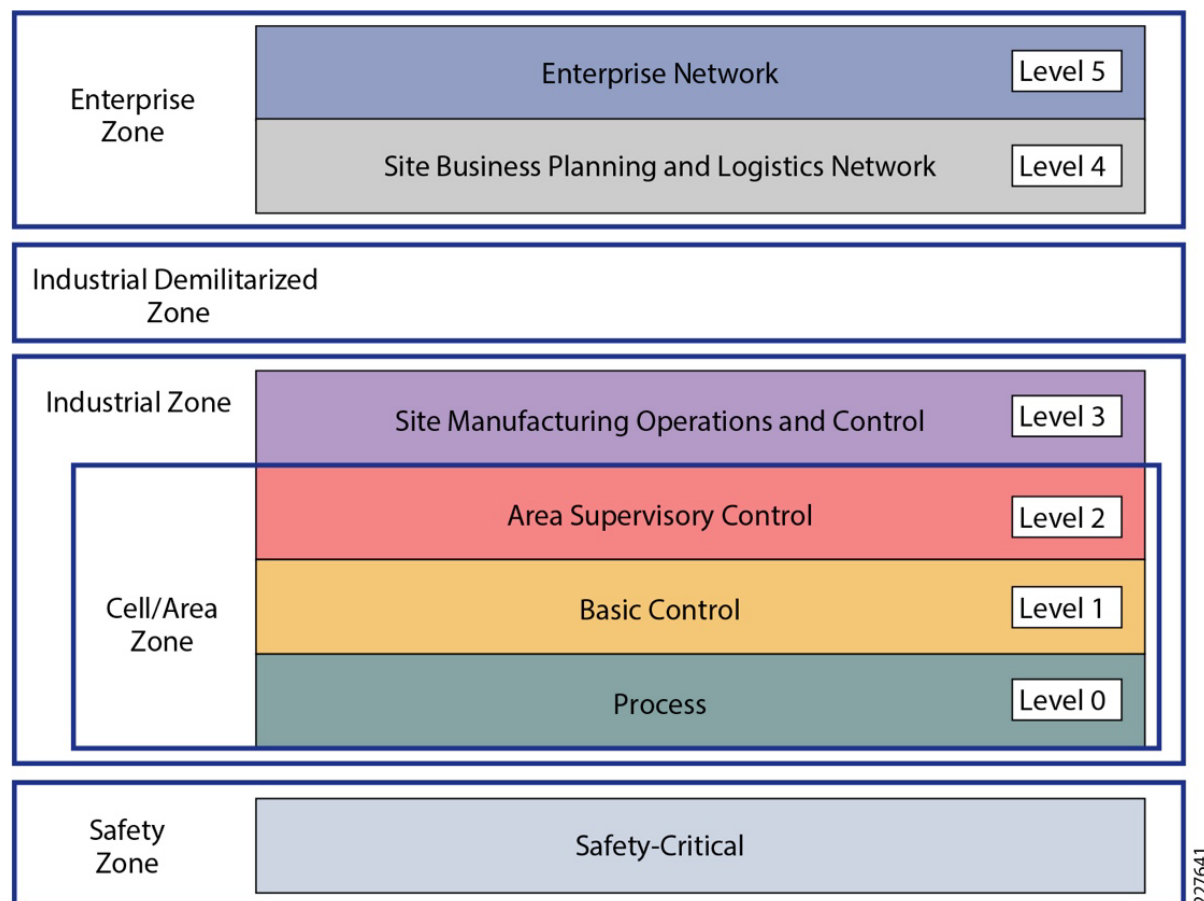
Note: Cisco Cyber Vision Global Center does not require an additional license.

Chapter 2. Preventative Controls in Plant Networks

Reference Architecture

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 15 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 15 Logical Industrial Cybersecurity Framework for Industrial Automation Networks



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](#).

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 shown in Figure 15, 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 through the industrial demilitarized zone (IDMZ) to maintain the security, availability and stability of the IACS.

Industrial DMZ

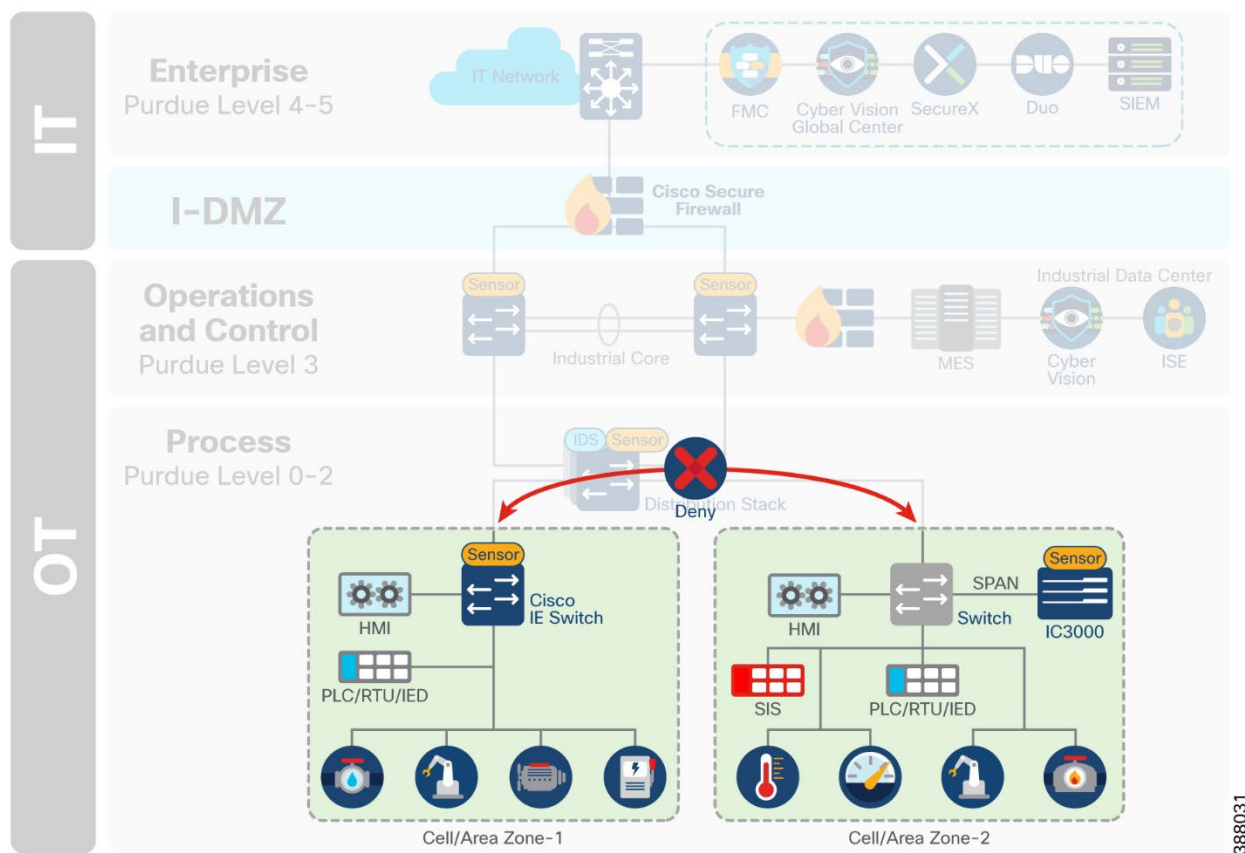
Although not part of the Purdue model, the industrial DMZ is deployed within plant environments to separate the enterprise networks and the operational domain of the plant environment. Downtime in the IACS network can be costly and have a severe impact on revenue, so the operational zone cannot be impacted 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.

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 16.

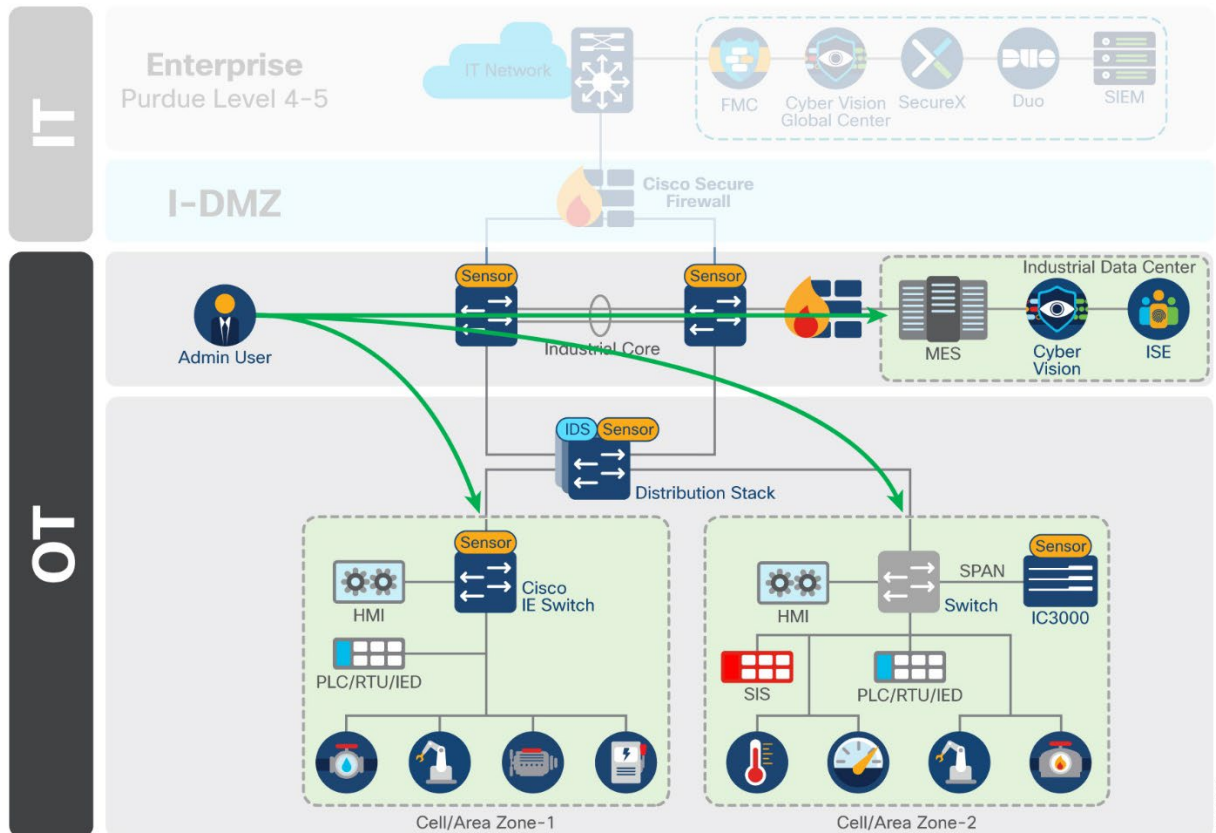
Figure 16 Cell/Area Zone to Cell/Area Zone denied by default. No segmentation inside the zone



388031

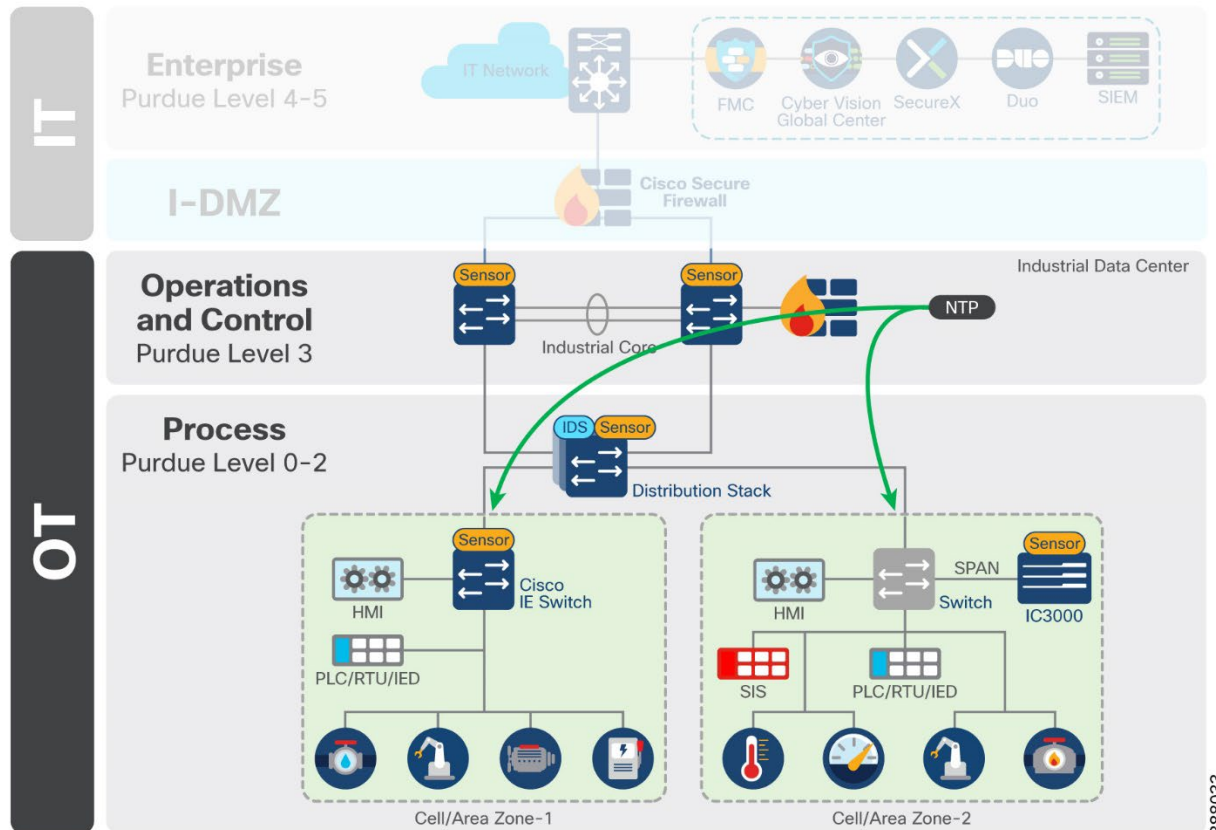
- **Administrative Users:** Figure 17 shows an administrative user who requires access to all zones in the network. They may be responsible for configuration of the network infrastructure, or the application of control logic. Their access should not be limited, but their data should be protected.

Figure 17 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.

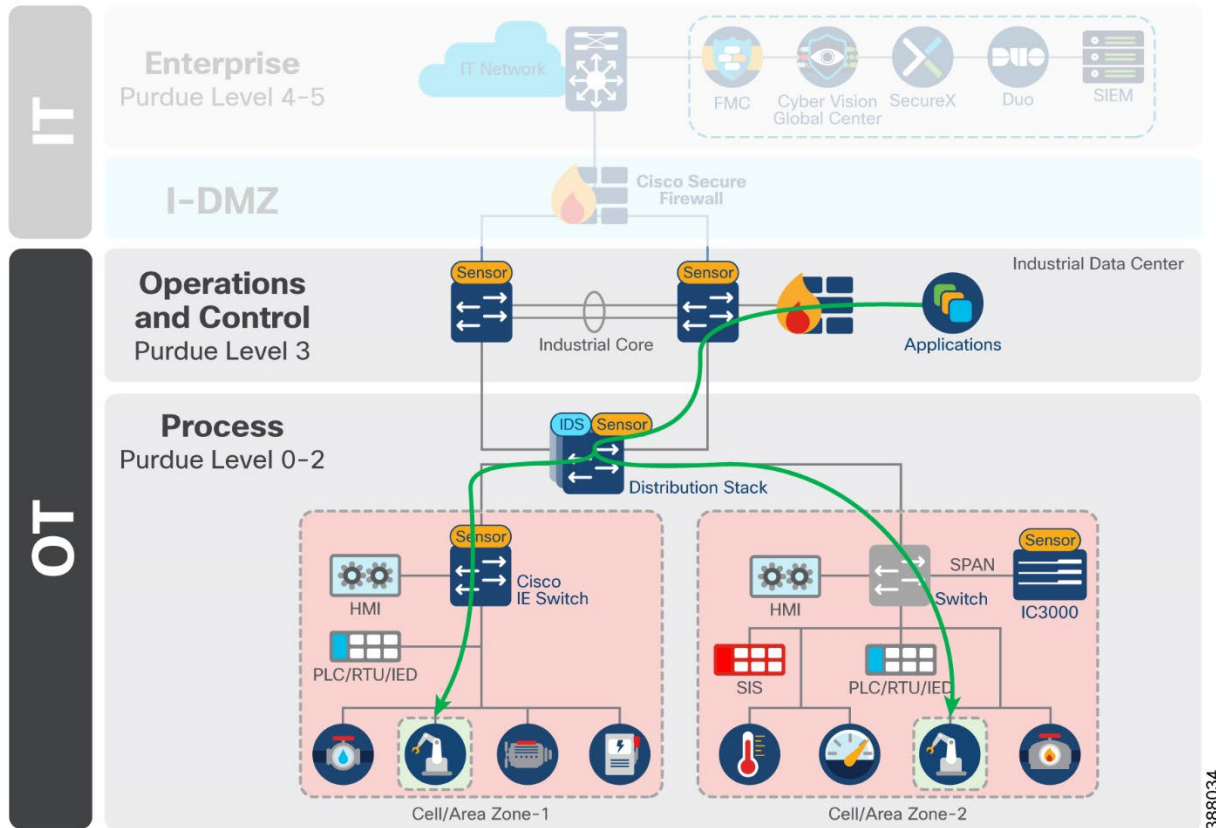
Figure 18 Infrastructure services that need access to all Cell/Area zones



388033

- **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.

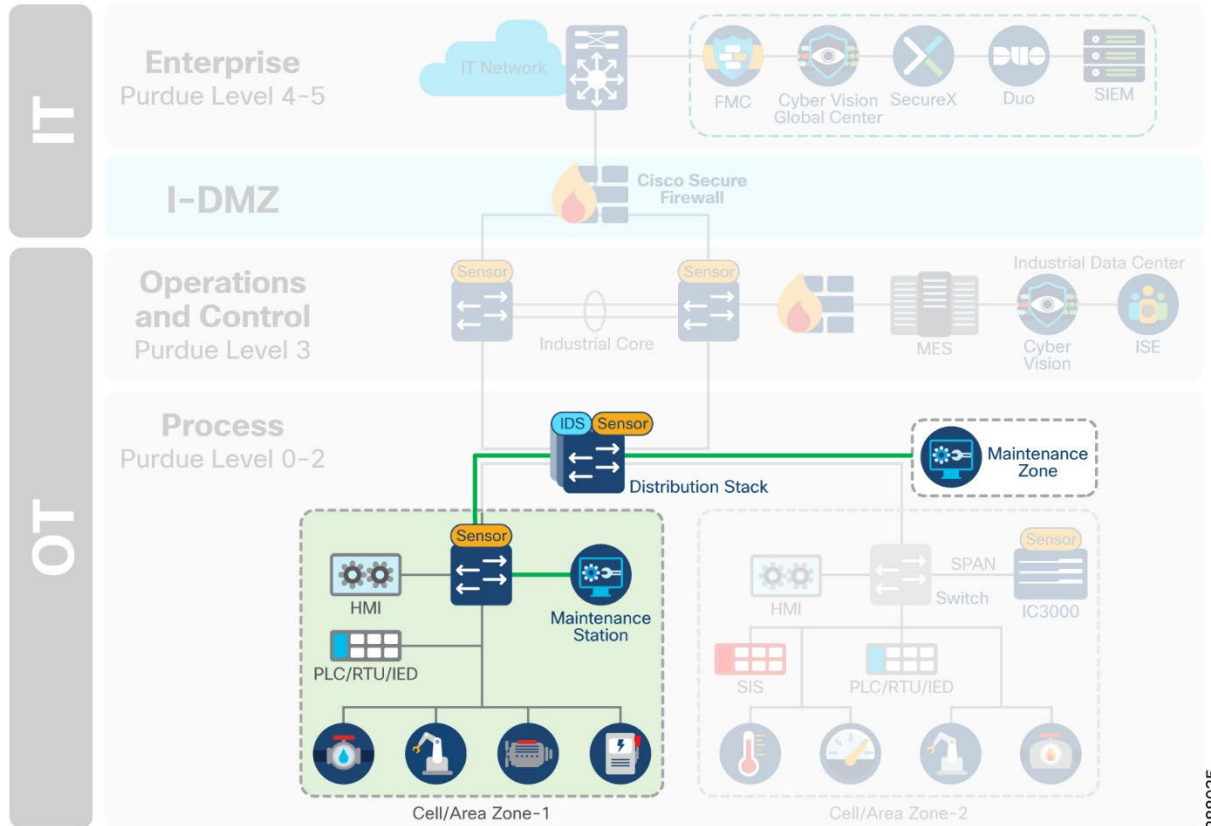
Figure 19 Applications that need access to specific services in the cell, but not the full cell



388034

- **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.

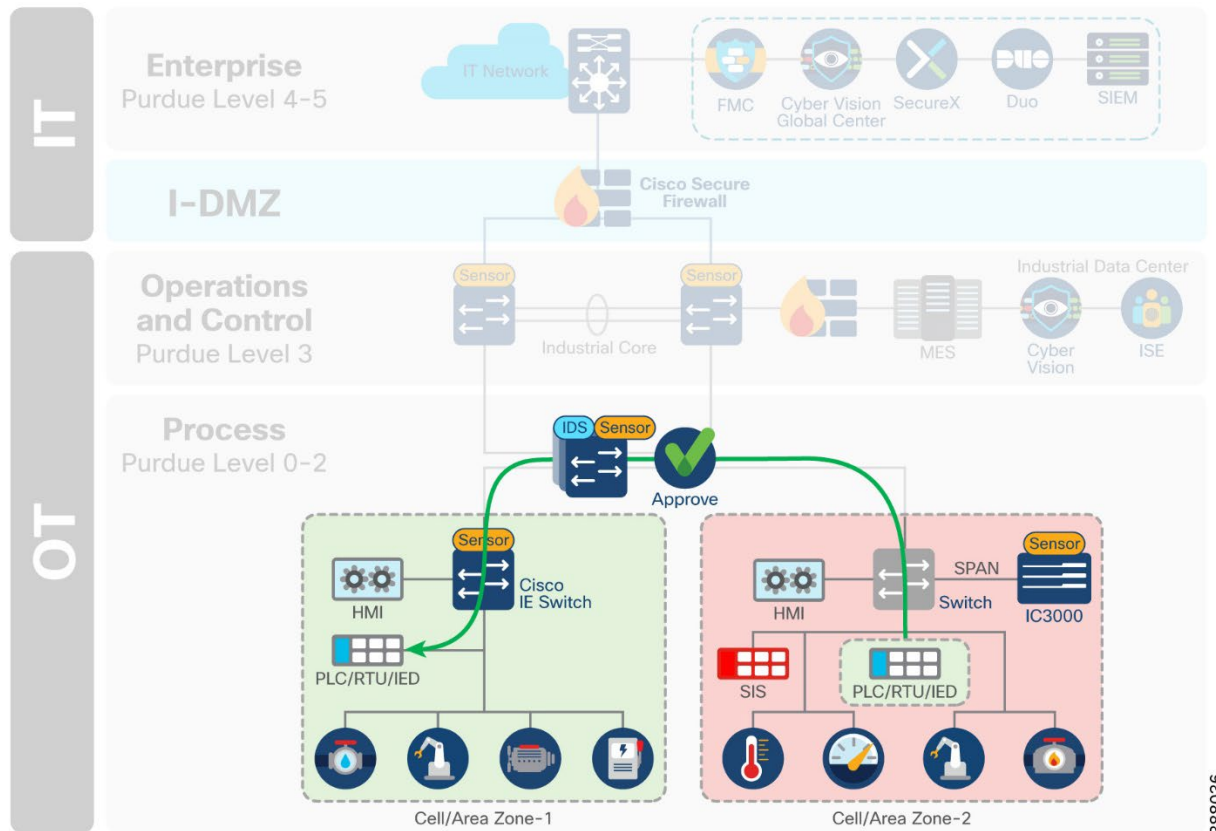
Figure 20 Maintenance workstations may reside within the cell or in a dedicated zone outside the cell



388035

- 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.

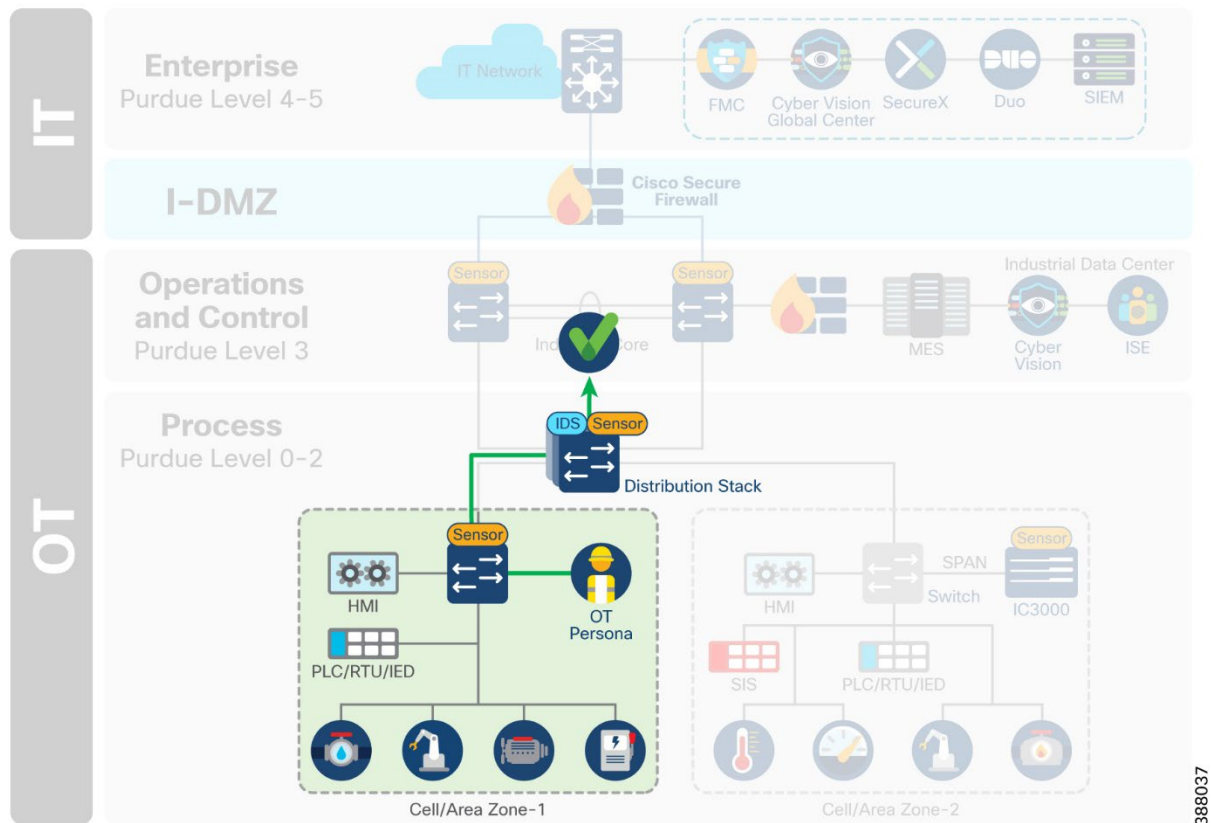
Figure 21 Select devices, such as interlocking PLCs, require communication across zones



388036

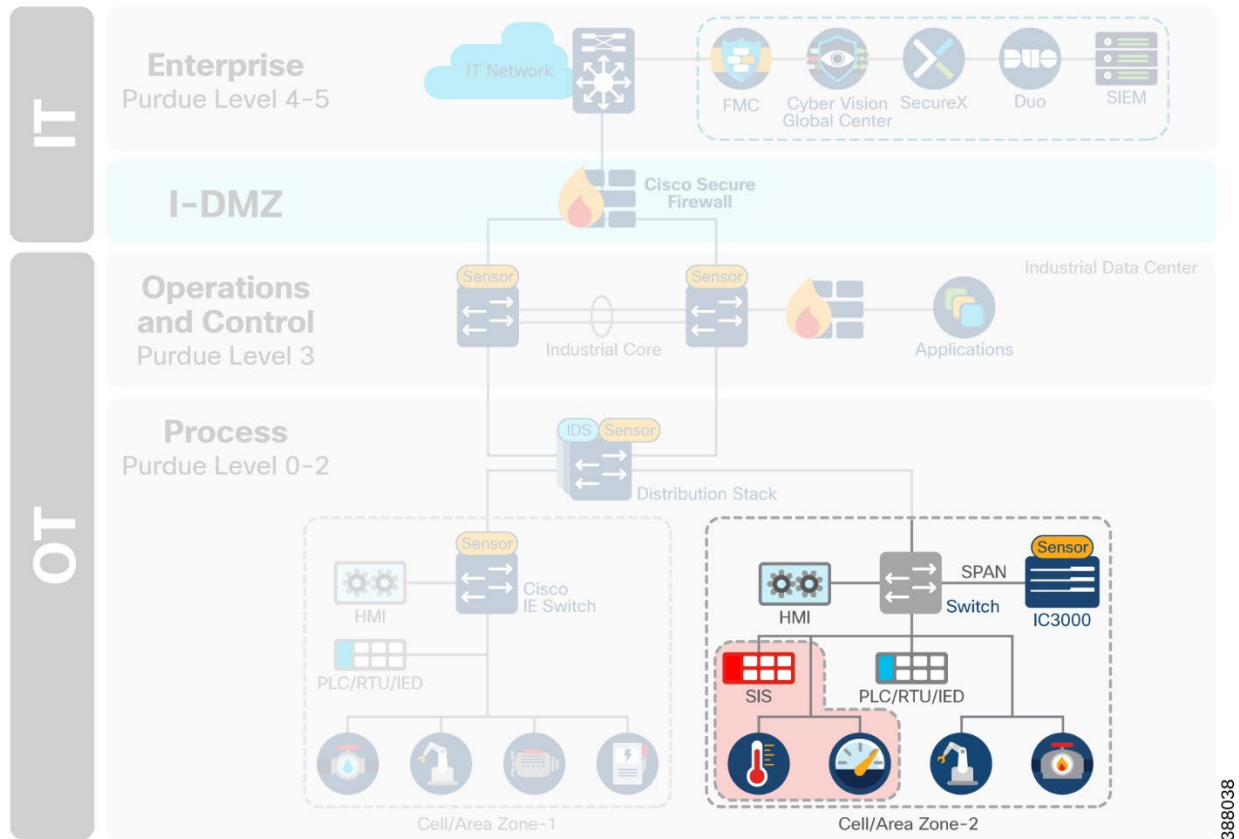
- **Convenience Port:** As operators plug directly into the infrastructure, they will typically bypass all the security checks that have been deployed in the architectural layers above it. Ensuring only authorized users with authorized device posture checks can connect to the network can aid in securing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

Figure 22 Operators plug into a cell using a convenience port and should be able to reach out to extra services



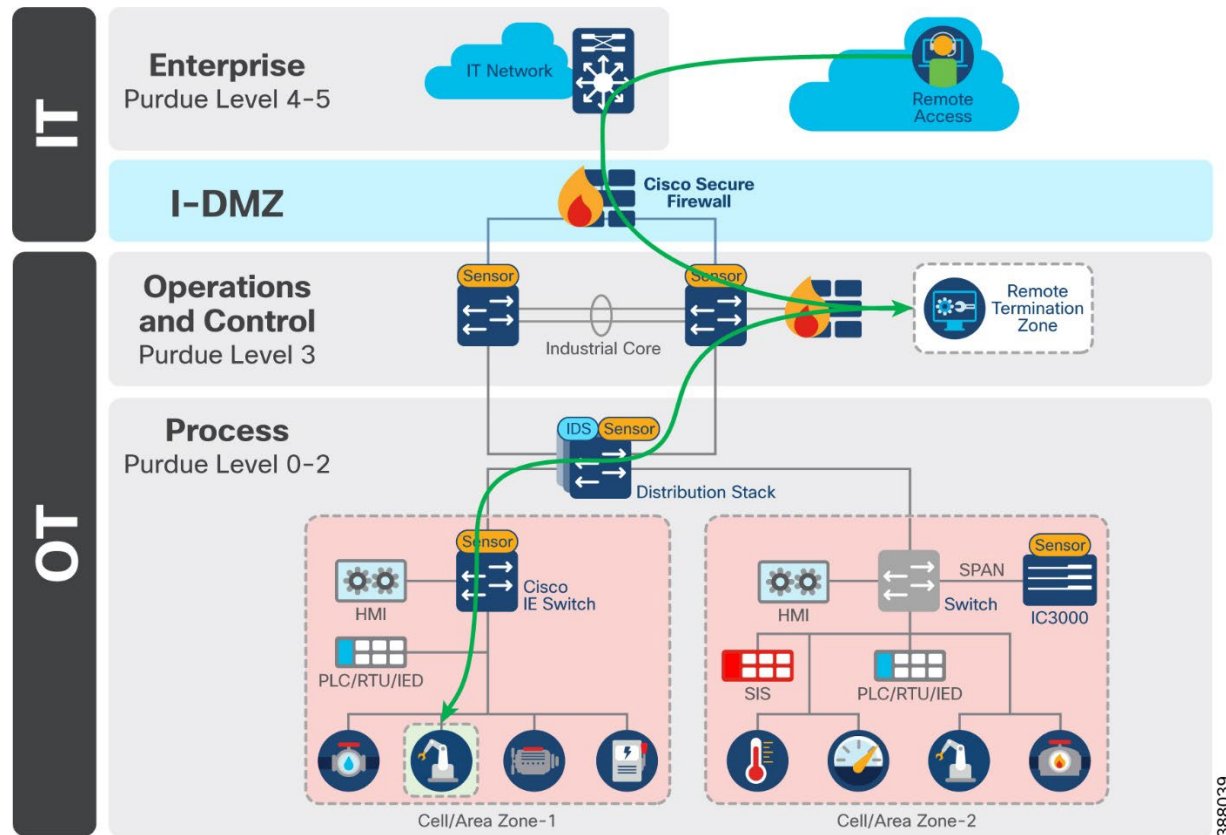
- **Safety Networks:** Safety Instrumented Systems (SIS) are critical to the control network and should either be air gapped from the rest of the network or logically segmented to ensure no data can leak into this zone.

Figure 23 Safety Network could be air gapped, or logically segmented 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 24 Remote Users need access to select devices, not a full zone



How to get started with Segmentation

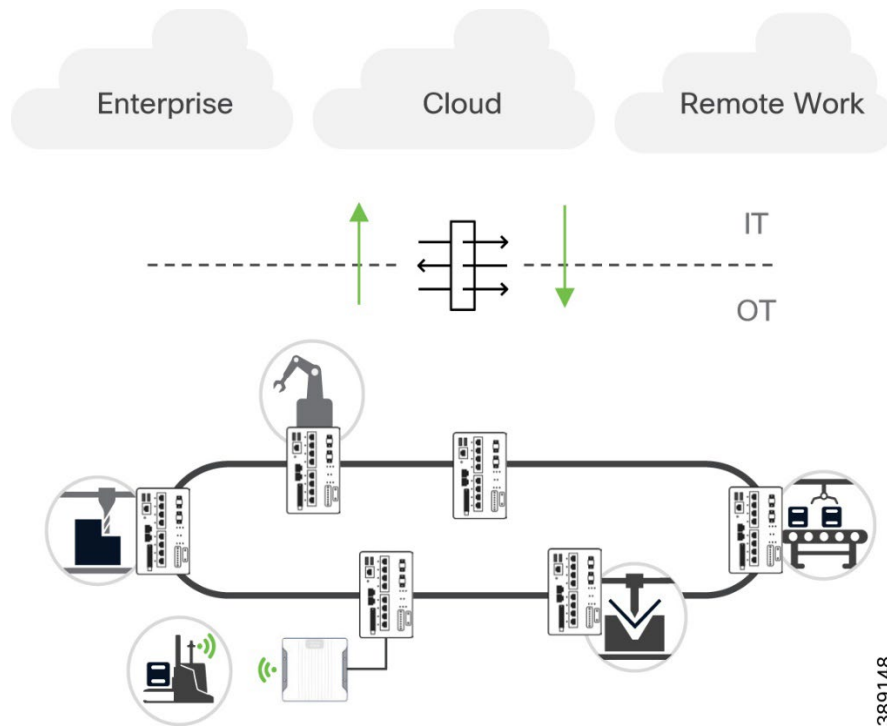
Defense in depth is crucial to protecting the network. The promise of zero trust and the use of micro segmentation technologies sounds great in principle, but in practice, is complex to implement. Enterprise networks have been implementing security for years, decades even. As new security technologies surface, IT teams can increment their security posture by building on the foundation of what has come before. When implementing a segmentation strategy in operational networks, micro segmentation cannot be the starting point. Protecting the plant comes in stages, and it is important to understand where the threats are coming from and how they make their way to the control systems. The guide recommends focusing on, in order, three main control points in plant networks:

- IT/OT Boundary
- Industrial Data Center
- Plant Floor

IT/OT Boundary

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 25 Enforce policy at the IT / OT boundary



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.

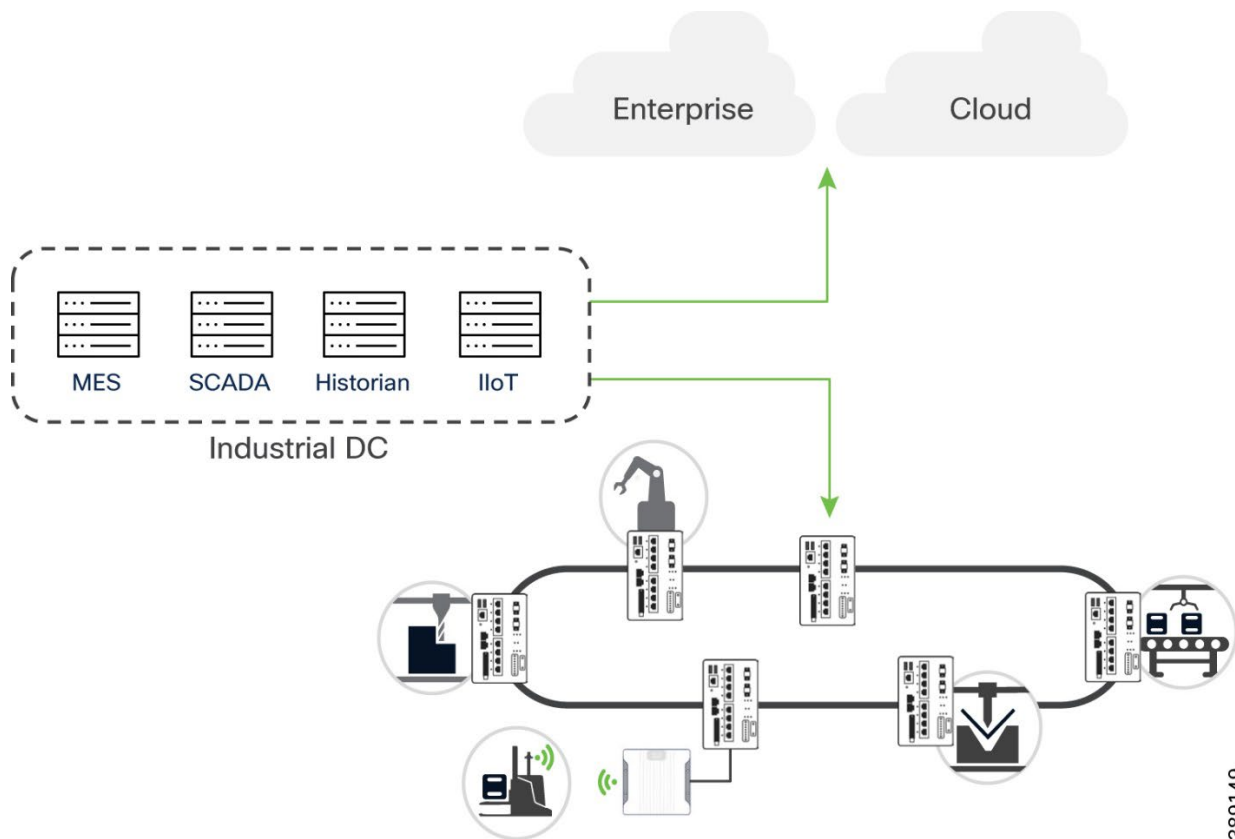
Design guidance for the hybrid cloud IDMZ will be added later as validation is still in the early stages of development. For an introductory insight into the architecture, see the [Hybrid Cloud IDMZ white paper](#).

Industrial Data Center

Rising investments into AI and the virtualization of the plant floor is resulting in the industrial data center (IDC) becoming a critical component of operational networks. [Virtual PLCs](#) are an example of this shift, where virtual controllers allow for a more flexible and modular design of production plants.

In a traditional Purdue model architecture, the IDC would reside in level 3, the industrial operations zone. This is important to distinguish, as many operational networks who have implemented some levels of control have done so at the IDMZ, or level 3.5. As the IDC becomes more modern, it also becomes more connected, relying on cloud connectivity for services to run as intended. As already established, more connectivity leads to an expanded attack surface, and if an attack was to breach the boundary firewall, more protection is needed.

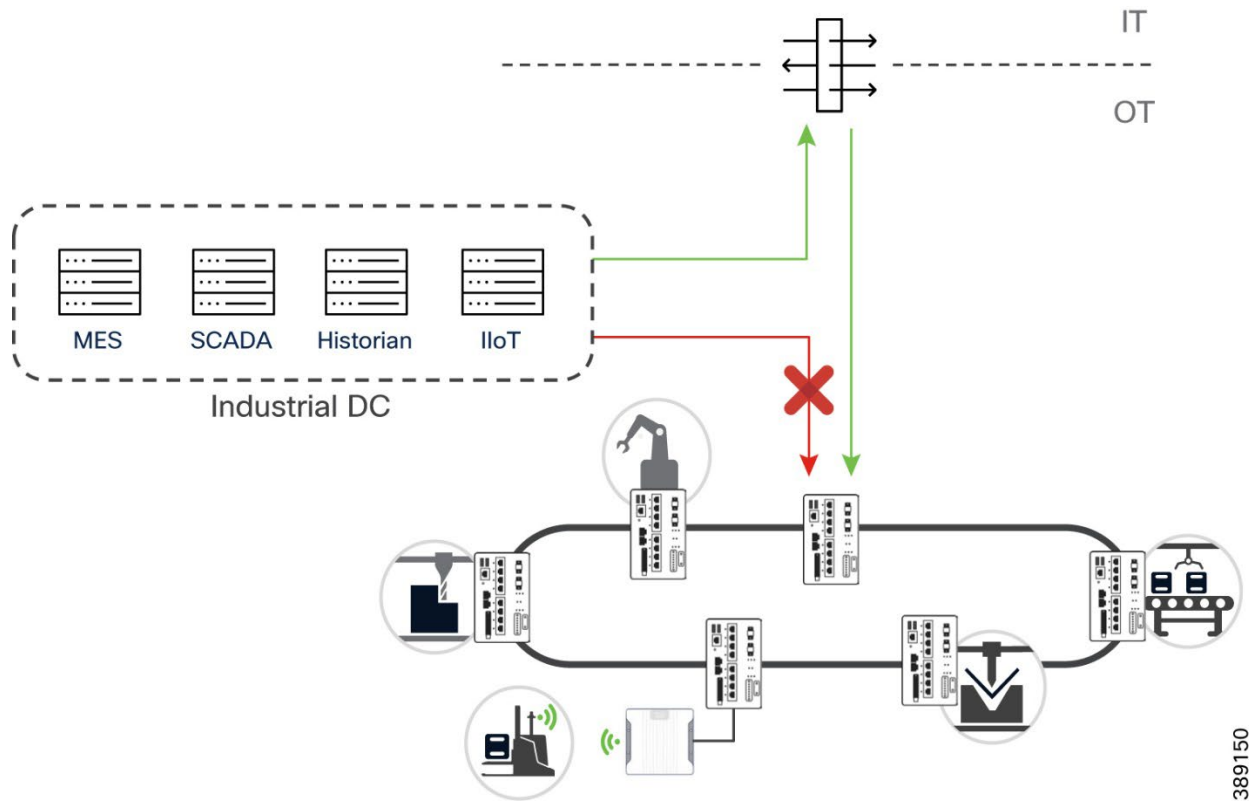
Figure 26 Access control at the industrial data center boundary



The IDC must have a segmentation point between it and the plant floor. Micro segmentation (see [Cisco data center security](#)) of the data center is out of scope for this design guide, but in scope is the firewall that should be placed at the IDC boundary. As modern systems are deployed in the operational network, they will ultimately co-exist with legacy systems that need protected from the newly exposed attack vectors.

Ideally, there is a separate firewall deploying dedicated to the IDC boundary, even if that firewall is virtual. However, in small plant environments, the same firewall used at the IT/OT boundary could be shared with the IDC or any other “macro zones” in the operational network.

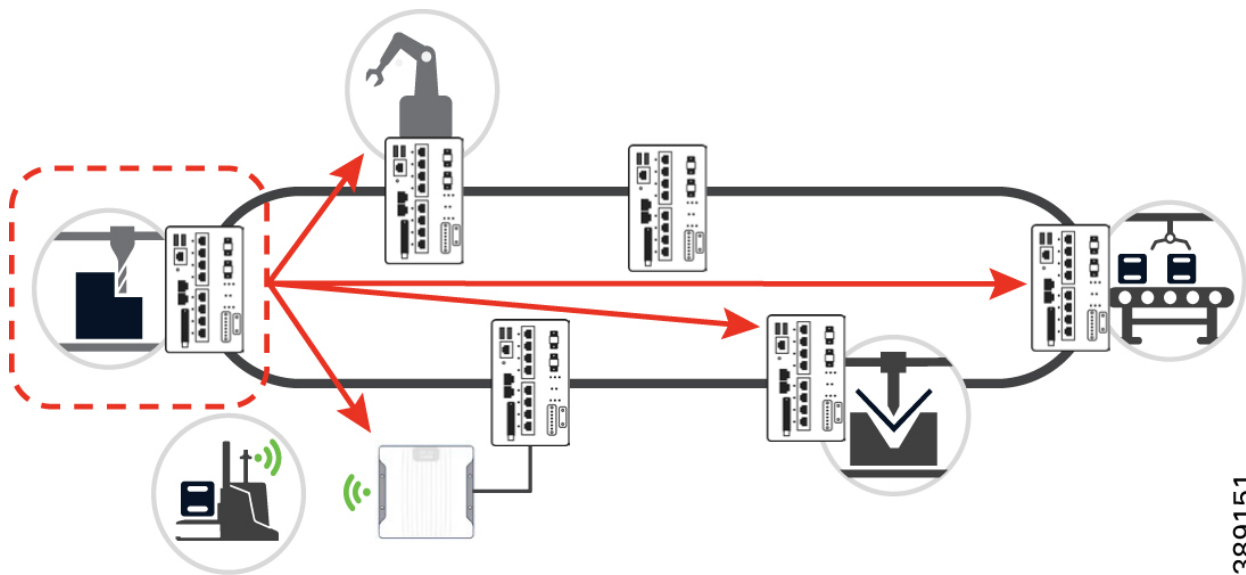
Figure 27 Single firewall pair for macro segmentation of plant network



Preventing lateral movement on the plant floor

It may seem counter intuitive to leave this last, as this is the main reason operators need a segmentation strategy. How do we ensure that the most critical parts of our networks will not be taken down by a malicious attack?

Figure 28 Prevent lateral movement within the OT



Security is an architecture, not a product. Defense in depth is a strategy, it is not deploying the same policies everywhere. Internet policies should have been covered by the IT/OT boundary firewall, threats from modern assets should be covered by the IDC firewall, and remote users will be covered later in [this guide](#). Preventing lateral movement in the plant floor is the last line of defense. If threats were to make its way into a system, either through poorly implemented controls, cellular back doors, or a simple USB stick, the blast radius will be reduced to only the system(s) that has been affected. The following sections will offer two ways to do this; Cisco Secure Firewall, or Cisco Identity Services Engine.

ISA/IEC 62443 Zones and Conduits model for OT Segmentation

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.

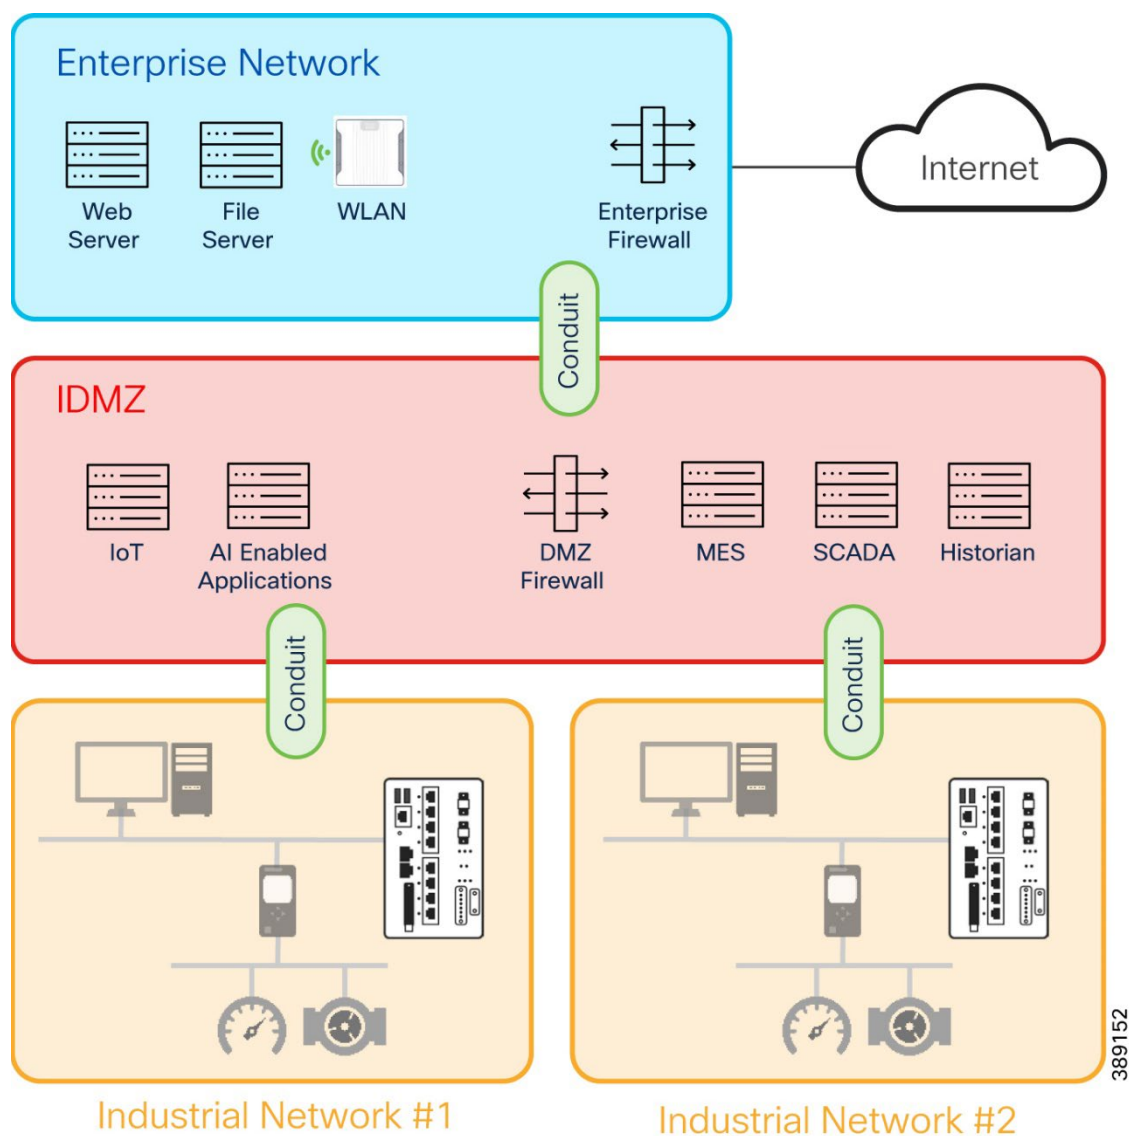
To improve interconnection and compatibility between industrial systems, equipment manufacturers are increasingly using standard communication protocols and complying with the requirements of international standards organizations. This is the role of the International Society of Automation (ISA), the 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 cyber-attack 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.

Figure 29 ISA/IEC 62443 Zones and Conduits Model



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

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

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

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

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

Architecting a network with secure components

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

IEC-62443-4-1 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.

IEC-62443-4-2 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.

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

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.

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 Firewall

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.

Segmentation Design Guidance with Cisco Secure Firewall

Note: The purpose of this section of the document is to provide firewall design guidance for use cases below the IT/OT boundary. For further references to securing the IT/OT boundary see [IT/OT boundary](#).

Threats have become more sophisticated, and networks have become more complex. Very few, if any, organizations have the resources to dedicate to staying up to date and successfully fending off these constantly emerging and evolving threats.

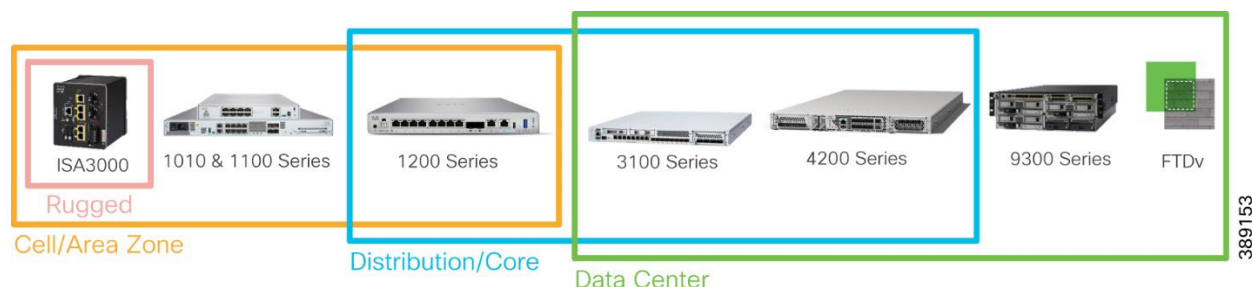
As threats and networks become more complex, it is imperative to have the right tools to protect your data, applications, and networks. Cisco Secure Firewalls have the power and flexibility that you need to stay one step ahead of threats. They offer a dramatic 3x performance boost over the previous generation of appliances, in addition to unique hardware-based capabilities for

inspecting encrypted traffic at scale. As well, the human-readable rules of Snort 3 IPS help simplify security. Dynamic application visibility and control is available through the Cisco Secure Workload integration, for consistent protection for today's modern applications across the network and workload.

Cisco Secure Firewall portfolio

The Secure Firewall brand encompasses the ASA and Firepower solutions. For the purposes of this document, any reference to the Cisco Secure Firewall will be referring to the Firepower Threat Defense (FTD) portfolio, otherwise known as Cisco Secure Firewall Threat Defense.

Figure 30 Cisco Secure Firewall portfolio



The choice of firewall will ultimately be determined by the throughput requirements for a given use case. A dedicated firewall for securing data in and out of the cell/area zone will have much smaller throughput requirement than the data center appliance. Use cases will be discussed later in the document, but for more information on each firewall specifications, such as performance metrics, see [Cisco Secure Firewall](#).

Management options

With the Secure Firewall portfolio, you gain a stronger security posture, equipped with future-ready, flexible management. Cisco offers a variety of management options tailored to meet your business needs:

- **Cisco Secure Firewall Device Manager (FDM):** Manages a single firewall locally; this is an on-device management solution to Firewall Threat Defense (FTD).
- **Cisco Secure Firewall Management Center (FMC):** Manages a large-scale firewall deployment. It is available in all form factors, such as on-premises, private cloud, public cloud, and Software as a Service (SaaS).
- **Cisco Security Cloud Control (SCC):** A cloud-based manager that streamlines security policies and device management across multiple Cisco products, such as Cisco Secure Firewall, Meraki® MX, and Cisco IOS® devices.

Cisco also offers Cisco Security Analytics and Logging for scalable log management. It enhances threat detection and meets compliance mandates across the organization with longer retention and behavioural analysis capabilities.

While many of the Cisco Secure Firewall capabilities can be achieved with local management, all design guidance in this document assumes the use of Cisco FMC to manage the devices.


Security guideline would suggest that the Cisco FMC is installed within the operational network, however, in reality security operators wish to consolidate the firewall management for both OT and IT networks. The connectivity needs for Cisco FMC can be found in the [Administration Guide](#).

Deployment modes

Routed Firewall Mode

In routed mode, the firewall is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet.


Figure 31 Cisco Secure Firewall in routed mode



Transparent Firewall Mode

A transparent firewall is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. However, like any other firewall, access control between interfaces is controlled, and all the usual firewall checks are in place.

Figure 32 Cisco Secure Firewall in transparent mode



Bridge Groups

Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the firewall uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. You can have multiple bridge groups for multiple networks.

Bridge groups are supported in both transparent and routed firewall modes.

- In transparent mode, bridge groups are the mechanism in which layer 2 connectivity is achieved. Since the firewall is in transparent mode, there is no mechanism for bridge groups to communicate with each other.
- In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces for a mixed deployment.

More information about firewall deployment modes can be found in the [Cisco Secure Firewall Management Center Device Configuration Guide: Transparent or Routed Firewall Mode](#).

Inline Sets

Inline sets is an IPS-only mode that bypass many firewall checks and only support the IPS security policy. This may be beneficial if you are already deploying another technology for access control policies, such as Cisco Identity Services Engine, and you want to supplement the deployment with IPS without the overhead of firewall functions.

Additionally, inline sets can be deployed with passive interfaces, where a copy of the traffic is sent to Snort for intrusion detection only. Inline sets would not be a recommended best practice for OT networks, as it is more appropriate to use a dedicated OT visibility tool such as Cyber Vision to passively inspect the network.

Key Capabilities

Application Recognition for OT

In its most basic form, a firewall provides access control policies that inspect and/or control network traffic across a boundary. Cisco Secure Firewalls can control traffic based on:

- Simple, easily determined transport and network layer characteristics - source and destination, port, protocol, and so on
- The latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application uses, or URL visited
- Realm, user, user group, or ISE attribute
- Security Group Tag (SGT)
- Characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- Whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- Time and day (on supported devices)

One common way to control the flow of traffic through a modern firewall is the use of application aware policies. The Cisco Secure Firewall currently identifies nearly [6,000 specific applications](#). Using Access Control rules, application traffic can be trusted, blocked or allowed but passed on for deep packet analysis and threat inspection.

Applications can be identified whether or not they are operating on standard network ports. In some cases, the presence of specific applications operating over non-standard ports may indicate a policy violation or an attempt to evade firewall controls. The Secure Firewall can identify these non-standard connections and generate alerts or block traffic as needed.

More information on application control features in Cisco Secure Firewall can be found in their [online documentation](#). This design guide highlights the ability to detection SCADA protocols throughout the plant network. Whether it Modbus messages, or Ethernet/IP between Rockwell devices, the Cisco Secure Firewall not only understands that the application is present in the communication, but knows the specific function codes being used by the system. For example,

understanding if a newly deployed IIoT device is reading data from a PLC (i.e. Modbus Read) or if they are attempting to manipulate data in the control loop (i.e. Modbus Write). The Cisco Secure Firewall can do per packet inspection, and per packet control to ensure only the relevant permissions are enabled between devices, which is especially important as modern systems look to retrieve critical information for data analytics and operators need peace of mind that read only access is enforced over the network.

Snort Intrusion Detection/Prevention System

Typically, an intrusion detection/prevention system (IDS/IPS) sits behind an access control engine. Whereas the access control engine blocks or permits network traffic based on layer 2-7 attributes of a traffic flow, IDS/IPS detects exploit attempts (attacks) within the traffic flows allowed by access control. IDS/IPS use protocol decoding engines and specific traffic flow characteristics (sometimes via regular expression pattern matching) to detect and block incoming attacks. The traffic characteristics used for identifying vulnerability exploitation are known as signatures and are the workhorses of IDS/IPS.

Snort is an open-source IDS/IPS implementation. It was initially developed in 1998 and has been available for free since then. Snort technology is the IDS/IPS engine used in Cisco Secure Firewall.

Note: that in the Snort world, signatures are called “Snort rules.” This document will exclusively use the term “signature” to ensure consistency in the exposition.

Virtual Patching

Traditional patching methods, although effective, may not always be feasible due to operational constraints and the risk of downtime. When a zero-day vulnerability is discovered, there are a few different scenarios that play out.

Consider two common scenarios:

- A newly discovered CVE poses an immediate risk and in this case, the fix or the patch is not available and:
- The CVE is not highly critical so it is not worth patching it outside the usual patch window because of the production or business impact. In both cases, one must accept the interim risk and either wait for the patch to be available or for the patch window schedule.

Virtual patching, a form of compensating control, is a security practice that allows you to mitigate this risk by applying an interim protection or a “virtual” fix to known vulnerabilities in the software until it has been patched or updated. Virtual patching is typically done by leveraging the Intrusion Prevention System of Cisco Secure Firewall. A key capability is the Cyber Vision ability to discover CVE information in the operational environment so the relevant IPS policies can be enabled in the Cisco Secure Firewall.

SnortML for Zero-Day vulnerabilities

IDS/IPS rely on “signatures” to detect exploits in progress. However, the signatures may not detect exploits that utilize previously unknown variants of known vulnerabilities. We call these exploits “zero-day variants.” Network security researchers and developers have long investigated techniques that would enable IDS/IPS to detect the zero-day variants. Machine learning (ML) –

a branch of AI – has finally provided a mechanism to extend IDS/IPS detection to these zero-day variants.

IDS/IPS signatures are usually crafted by humans. After creation, they undergo an extensive testing cycle to build confidence in their efficacy. Depending on the signature provider, signatures may also be deployed in limited live settings before broader deployment. Again, the limited deployments aim to increase confidence in the signatures.

Signatures are generally written to tightly fit known exploitation of a known vulnerability. This is done deliberately to lower the probability of matching legitimate traffic, thereby keeping false positives low. The flip side of the endeavour to maintain a low false positive rate is that signatures may miss zero-day variants.

Thus, the human signature writer must manually tune the generalizability of a signature. If the signature is too tight, it will not catch even modest variations of a known attack, let alone zero-day variants. If the signature is too general, it will result in false positives. Finding the right balance is a tedious process that requires frequent trial and error. Frequently, false positives are enough of a concern that only tightly matching signatures are deployed.

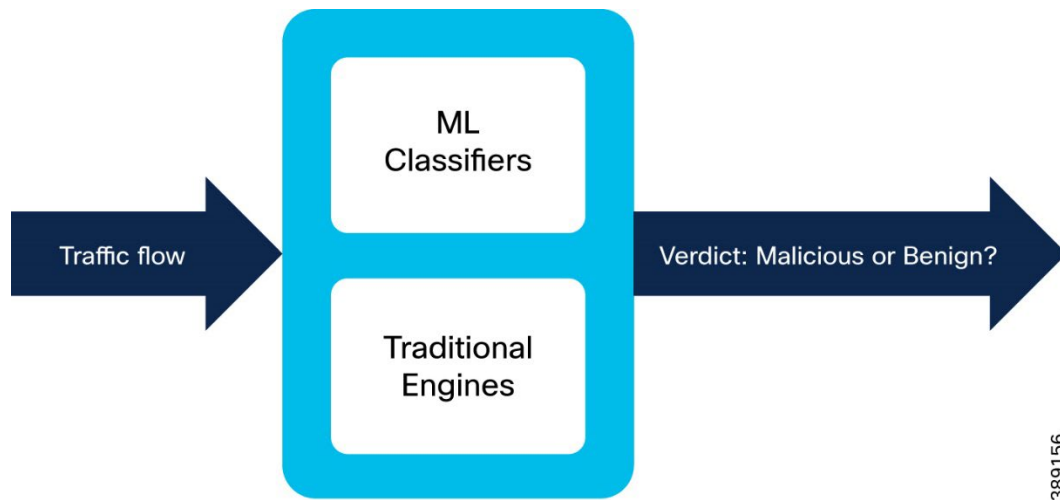
SnortML provides an automated mechanism to find the right balance between generalizability and false positives. As mentioned above, SnortML uses machine learning – in particular, a deep neural network – to detect exploits.

Machine learning techniques are an alternate way to “learn” the signature of a class of related exploits. Here, a deep neural network is trained on malicious and benign traffic corpora. The neural network infers generalized versions of the exploit patterns in the malicious corpora and learns to distinguish between malicious and benign traffic. For example, the malicious traffic may be an SQL injection exploit, whereas the benign traffic may be legitimate SQL queries.

Note that a neural network has many parameters that are tuned during the training process. Effectively, during training, generalized inferred patterns of attacks are embedded in the parameters of the neural network. These generalized patterns enable the neural network to detect zero-day variants. Continuing the SQL example above, a neural network can learn the pattern of related SQL injection exploits and detect a new exploit even if it has never seen it before.

SnortML has two components. The first is the machine learning engine, which loads machine learning classifiers (trained over malicious and benign traffic as discussed above) and makes them available for detection. The second is an inspector, which subscribes to data provided by the underlying Snort architecture, passes the data to classifiers, and then acts on the classifiers’ output. The SnortML classifiers run in parallel to traditional signature-matching engines within Snort, as shown in Figure 33.

Figure 33 Machine learning classifiers and traditional signature-based detection in parallel




SCADA Inspectors

While the Cisco Secure Firewall provides thousands of preconfigured Snort rules for use in access control rules, Snort also has the capability to load custom signatures. The [Snort 3 rule writing guide](#) provides documentation on this rule writing process, detailing each option available to users to create their own detections.

To make rule creation easier, especially rule options that require payload detection, Snort offer “inspectors”. Inspectors decode applications and provide custom rule options for that application. For example, Modbus is a protocol used in SCADA networks, and its traffic is typically seen on TCP port 502. The Modbus service inspector decodes the Modbus protocol and provides three rule options that rule writers can use to evaluate Modbus traffic.

Those three options are **modbus_data**, **modbus_func**, and **modbus_unit**.

Figure 34 Modbus TCP packet structure



Snort 3 provides inspectors for the following SCADA protocols:

- [Distributed Network Protocol 3 \(DNP3\)](#)
- [Common Industrial Protocol \(CIP\)](#)
- [IEC 60870-5-104 \(IEC 104\)](#)
- [Manufacturing Message Specification \(MMS\)](#)
- [Modbus](#)
- [S7 Communication \(S7Comm and S7CommPlus\)](#)

Port Scan Detection

Port scanning is a common attacker reconnaissance activity. Malicious actors use port scanning to discover open ports on a firewall as well as port sweeps to discover listening ports/services on a host. Previous releases of Secure Firewall include port scan detection in the Intrusion policy as part of the Snort 2 and Snort 3 detection engines. However, this detection is lacking because Snort distributes connections across threads or CPU cores.

The system uses a hash of several connection attributes, including source/destination IP and port, to assign a connection to a CPU or thread. Due to the nature of a port scan, the hashing algorithm causes this activity to be split across several CPUs on the device. This means no single CPU/thread has access to all the port scan activity from a given scanner, thus reducing the effectiveness of this feature.

To address this, release 7.2 moves the port scan detection capability from Snort to the operating system of the firewall. By moving this capability, the device can now detect port scans more effectively as the port scan detection process has visibility of all the scan traffic for a given scanner. This visibility also holds true for distributed port scans where there are multiple scanners and single or multiple targets, and port sweeps.

As part of this move, the port scan detection configuration parameters are now in the Access Control policy. This new port scan detection capability is only available on 7.2 or higher devices using the Snort 3 detection engine.

For more information, including the configuration details, see [port scan detection](#).


Cisco Secure Dynamic Attribute Connector

The Cisco Secure Dynamic Attributes Connector (CSDAC) was initially created for the firewall policy to adapt in real-time to the changes in public and private cloud workloads and business-critical SaaS applications.

The Firewall Management Center and CSDAC integration simplify management through firewall policy automation by keeping the rules up to date without tedious manual updates and policy deployment. With CSDAC, you can centrally manage workload attribute feeds obtained from multiple public and private cloud environments, enabling firewalls to adapt to changes instantaneously to help accelerate Cisco Secure Firewall integration with your complex and dynamic environment. CSDAC significantly improves network security with automatic endpoint attribute and contextual awareness propagation simultaneously, preventing the build-up of outdated firewall rules over time.

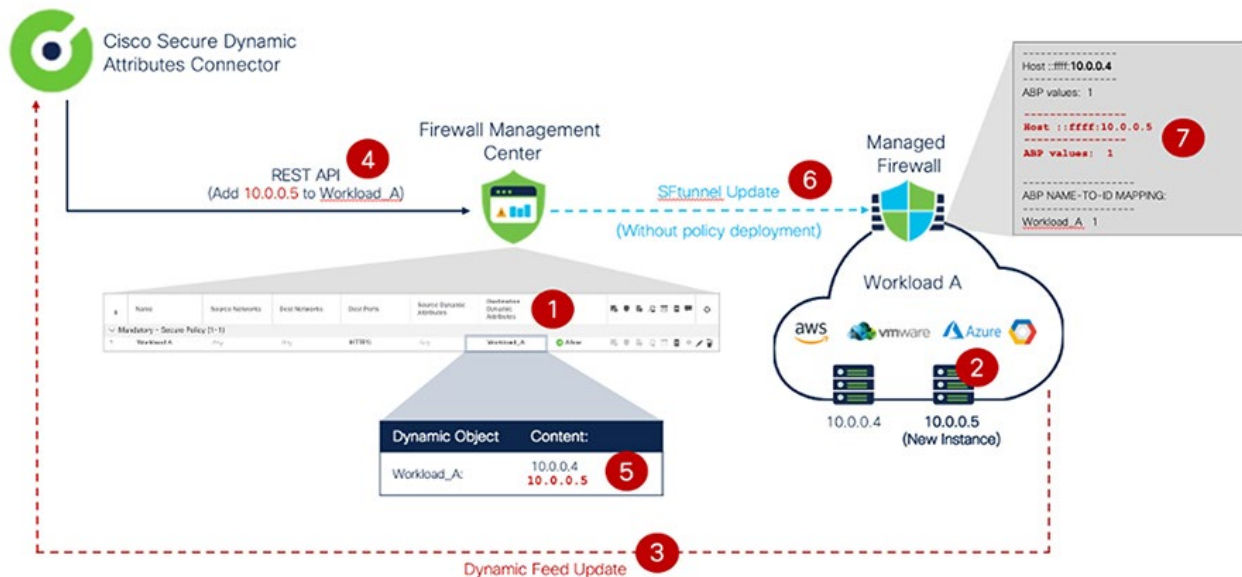
With CSDAC the firewall policy becomes more dynamic, secure, and much easier to manage. As illustrated in the figure below, the CSDAC discovers resources and IP addresses in the cloud and translates this information to dynamic network objects consumed by firewall deployment. The CSDAC keeps track of the changes on the cloud side and updates dynamic objects accordingly in near real time.

Figure 35 CSDAC connector examples



CSDAC maps IP addresses of cloud resources to Dynamic Objects, which are then used in Access Control Policy rules. Changes in the cloud detected by CSDAC are cascaded in real-time to the FMC, and in turn, to the managed firewalls without any administrator action. The figure below illustrates step-by-step, how CSDAC dynamically updates firewall policy.

Figure 36 CSDAC dynamic object update



1. The firewall protects a Workload and is configured with an Access Control Policy containing the dynamic object Workload_A representing cloud resources.
2. CSDAC monitors changes to the workload constantly and detects when a new instance is spun up.
3. CSDAC detects the workload change and evaluates the user-created attribute filters.
4. Then CSDAC triggers a REST request to update the Workload_A dynamic attribute with the 10.0.0.5 IP address of the new server.
5. The Firewall Management Center adds the new IP address to the dynamic object.
6. Immediately after the Workload_A object change, the FMC pushes an update to all the managed firewalls using that object in deployed Access Control Policies. The dynamic object update happens automatically and does not require a policy deployment.
7. The firewall updates the new IP address in the Snort identity memory and its policy to allow the new server access.

The CSDAC connector is a software interface that interacts with a public or private cloud provider to retrieve up-to-date network information, categories, and tags. CSDAC translates information provided by the connectors to Dynamic Objects used in firewall access control policies on the FMC. Architecturally the connectors are software plug-in modules installed in CSDAC, which allows the straightforward addition of new connectors in future releases. Up to this point, all examples have demonstrated how the firewall can be kept up to date with changes in cloud or data center workloads, however, when deploying policy in OT environments, it is important to have context for both ends of the connection. This is where the Cyber Vision connector can be used.


Cyber Vision Connector

Cyber Vision and Firewall Management Center allow OT segmentation groups defined by the OT team in Cyber Vision to be used for firewall enforcement. This level of automation helps reduce manual workloads, streamlines your security management process by enabling IT/OT collaboration, and helps ensure that your firewall policies remain in lockstep with your OT industrial processes.

- Cisco Cyber Vision inventories industrial assets and maps their communication activities.
- Operations managers leverage the Cyber Vision maps to group assets into industrial zones.
- Cisco FMC pulls asset group information from Cyber Vision using CSDAC.
- Each Cyber Vision group becomes a dynamic object in FMC, to which the IP addresses of the assets in the group are mapped in real time.
- IT and OT managers work together to define access policies to be applied to each dynamic object.
- Policies defined in FMC are enforced by Cisco Secure Firewalls.

- Any modification to Cyber Vision groups is reflected in FMC dynamic objects in real time and is automatically enforced by Cisco Secure Firewalls, without the need to redeploy policies.

Figure 37 Cyber Vision integration to Cisco FMC via CSDAC



Using CSDAC, OT asset groups created in Cyber Vision are automatically made available to the Firewall Management Center as dynamic objects. IP addresses of OT assets are continuously imported and mapped to dynamic objects, helping ensure that objects are always aligned with the industrial processes defined by the OT team.

The dynamic nature of this integration eliminates the need for manual policy deployment each time there is a change to the Cyber Vision map. Adding an asset to a group in Cyber Vision or moving it to another group will automatically modify the corresponding object in FMC. The access policy configured for this object will apply to this asset in a matter of just a few seconds.

Encrypted Visibility Engine

Traffic inspection is one of the most important tools used by today's firewall systems. However, encryption technologies like Transport Layer Security (TLS) present challenges to traditional deep packet inspection technologies. The inability to inspect packet payloads for encrypted traffic means key firewall capabilities are not available. Decryption technologies offer one solution to this issue. However, the process of decrypting traffic for inspection is resource intensive. In addition, enabling decryption significantly reduces next-generation firewall throughput. There are also operational concerns, for example, certificate pinning makes decrypting some applications impossible. There may also be privacy or compliance issues for certain types of traffic, making decryption difficult or otherwise undesirable.

Cisco Secure Firewall Threat Defense offers several technologies to enhance encrypted traffic inspection without the need to implement full main-in-the-middle (MITM) decryption. The most recent of these is the Encrypted Visibility Engine (EVE).

About EVE

EVE is a new means of identifying client applications and processes utilizing TLS encryption. It enables visibility and allows administrators to take actions and enforce policy within their

environments. EVE works by fingerprinting the Client Hello packet in the TLS handshake. By identifying specific application fingerprints in TLS session establishment, the system can identify the client process and take appropriate action (allow/block).

Currently, EVE can identify over 5,000 client processes. Secure Firewall, maps a number of these processes to Client Applications for use as criteria in Access Control rules. This gives the system the ability to identify and control these applications without enabling TLS decryption.

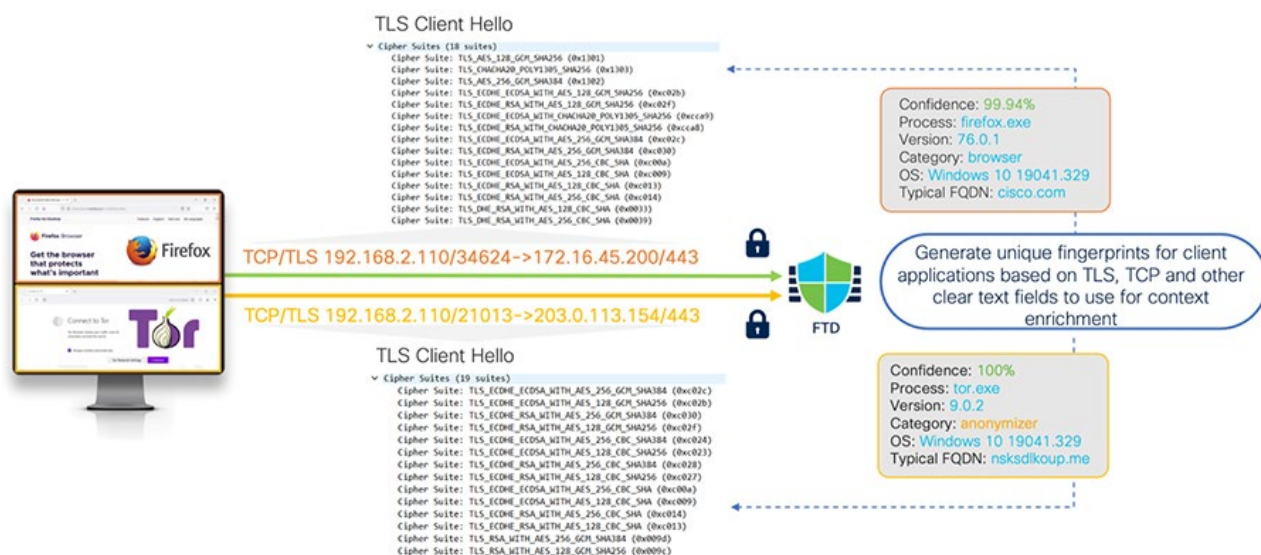
This capability also extends beyond just identifying Firefox or a TOR browser. By using fingerprints of known malicious processes, EVE technology can also be used to identify and stop malware. The malware identification feature is currently under development, and Cisco is constantly expanding the number of known fingerprints. As this capability matures, future releases will use EVE to identify and block encrypted malicious traffic without outbound decryption.

EVE brings another benefit, enhanced operating system identification. The system uses EVE to improve data accuracy in the FMC host database built from passive host discovery. By leveraging EVE fingerprint technology, the system uses encrypted host communications to identify host operating systems better. This leads to better application and vulnerability data, thereby improving the accuracy of features such as Snort Recommendations.

How does EVE work?

EVE inspects the Client Hello portion of the TLS handshake to identify client processes. The Client Hello is the initial data packet sent to the server following the three-way handshake. As it turns out, the Client Hello gives a good indication of the client process on the host. This fingerprint, combined with other data such as destination IP address, provides the basis for EVE application identification.

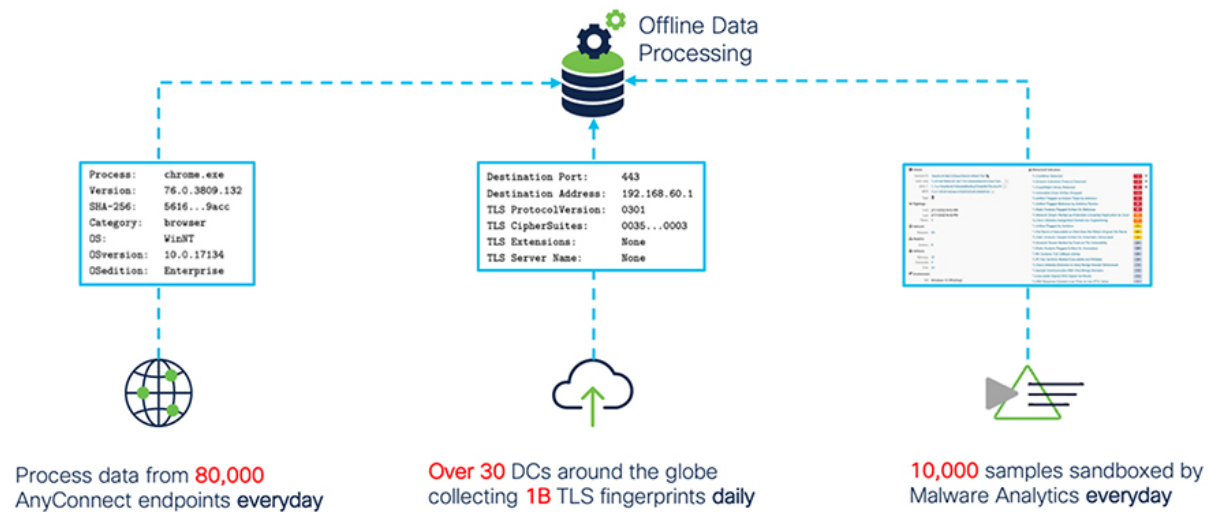
Figure 38 EVE communications



Through machine learning technology, Cisco processes over one billion TLS fingerprints and over 10,000 malware samples daily to create and update EVE fingerprints. These updates are then delivered to customers via vulnerability database (VDB) packages.

Figure 39 EVE machine learning

Machine Learning Requires a Comprehensive Data Set

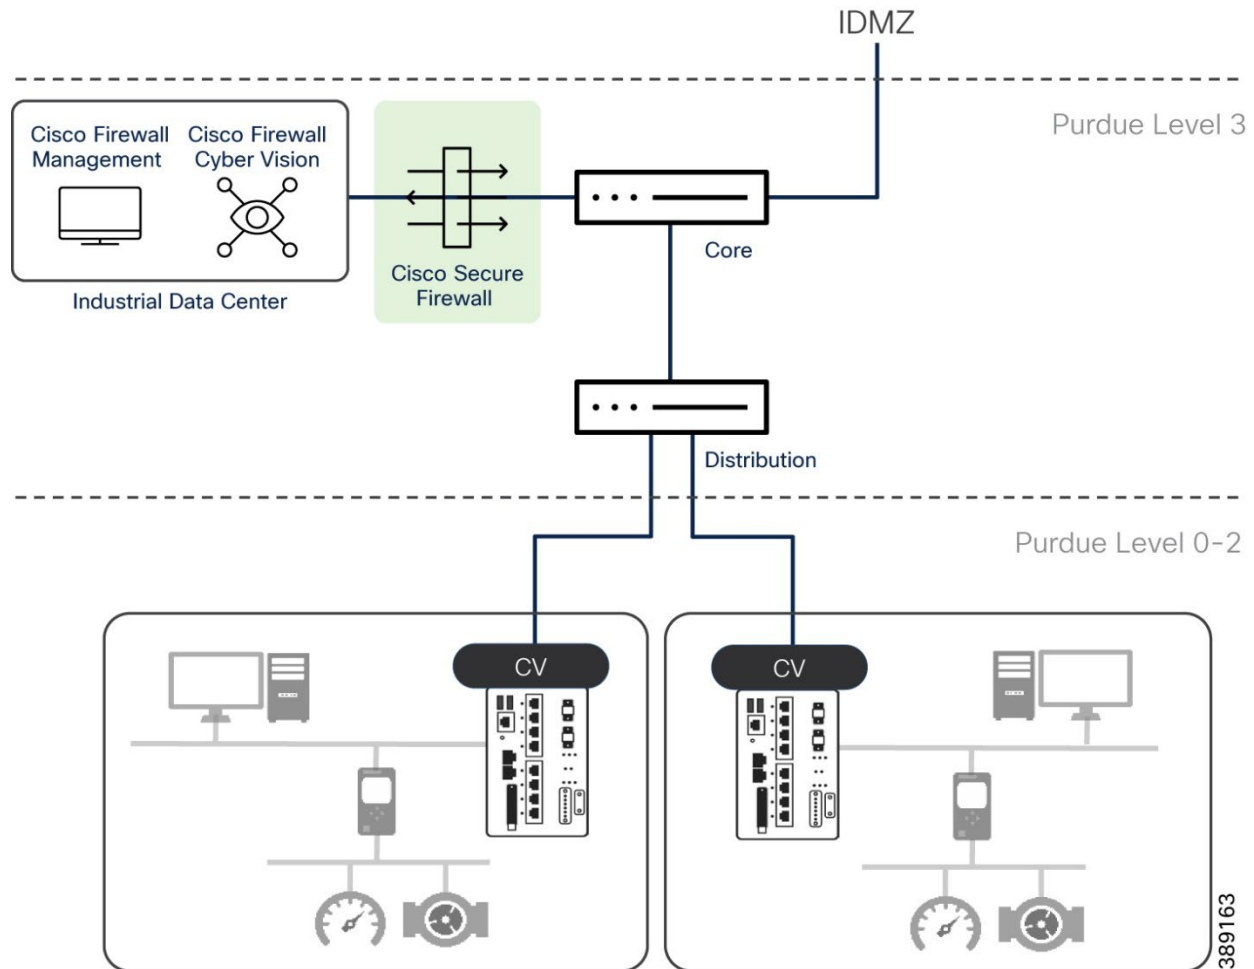


For the latest information see [Encrypted Visibility Engine](#).

Industrial Data Center Segmentation with Cisco Secure Firewall

As discussed earlier in the document, there is a clear sense that AI will boost business growth for those who can successfully use it to run better industrial networks. Leaders will ensure their operational technology can capture the required data to fuel their AI models. This will ultimately lead to an investment in the data center, bringing more technology, more connectivity within the legacy operational perimeter.

Figure 40 Cisco Secure Firewall at industrial DC boundary



When using a Cisco Secure Firewall for protecting IDC flows in OT environments consider the following:

- Start with visibility to map out the existing network topology and understand traffic flows and potential bottlenecks. Firewalls cannot be deployed using trial and error policies within an operational environment. Without understanding the flows that will cross an enforcement point, it will lead to unexpected downtime.
- Choose an appropriate model keeping in mind the load on the network. Sizing of the firewall is critical to avoid an increase in latency and jitter which is detrimental to OT traffic. Ensure the firewall can handle the peak loads of the network traffic.
- Deploy access control policies in monitor mode before pushing them to production. When using the **Monitor** action in an access control policy, policy matches will be logged, but the system will continue to match traffic against additional rules to determine whether to permit or deny it. This allows administrators to test their policies in production without the risk of erroneous rules causing downtime.
- For increased flexibility and ease of use, use **objects** in access control policies which are reusable configurations that associate a name with a value. For example, when referencing a list of IP addresses in an access control policy (e.g. list of engineering

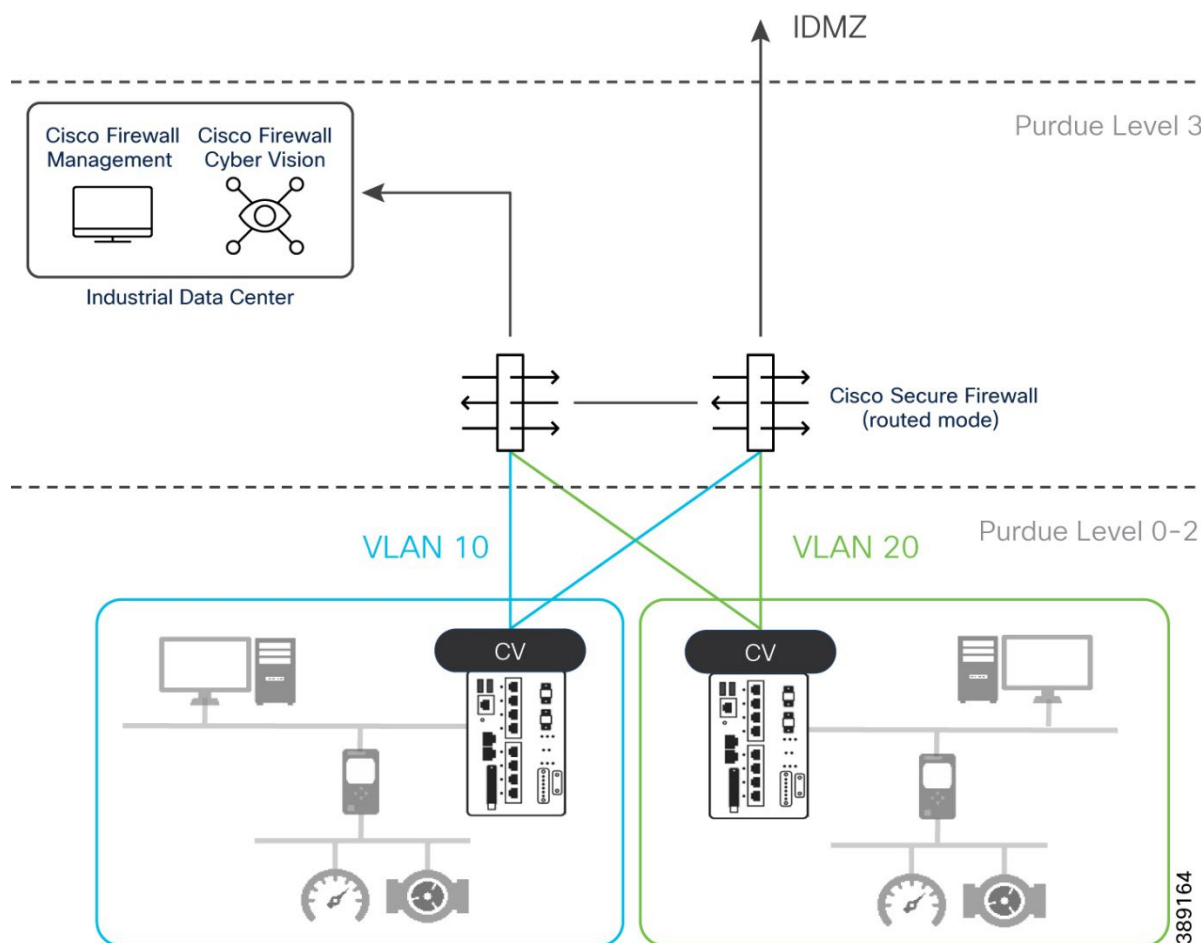
workstations on the plant floor), manage the list with an access rule. That way, anytime the engineering workstations are referenced in policy, they can be called by name. If IP addresses where to be added or removed, by changing the object, all policies will remain up to date.

- Use application control policies to provide read-only access to assets on the shop floor. Introducing ML applications into OT networks will typically require some level of visibility into control operations. To accomplish this MQTT brokers will be introduced to the network, but those brokers often have OT protocol connectors to collect data before publishing to a message bus. Use application control policies in the network to make sure these brokers, or anything else in the IDC that requires access to data over insecure OT protocols, is restricted to read-only rules.
- Consider using a virtual firewall to protect virtualized infrastructure in plant networks. Organizations are looking to reduce their physical footprint, and the firewall is another element of the network that can be digitized while offering the same level of protection we are accustomed to with physical appliances.
- Implement fail-safe mechanisms to maintain operations if the firewall fails. Clustering allows multiple Cisco Secure Firewalls to function as a single logical firewall. As of FTD 7.2, clustering is also supported on Cisco Secure Firewall Threat Defense virtual (FTDv).

East/West Segmentation with Cisco Secure Firewall

Using a firewall to route between OT VLANs

The first consideration is to use the Cisco Secure Firewall as the termination point for VLANs in OT networks.

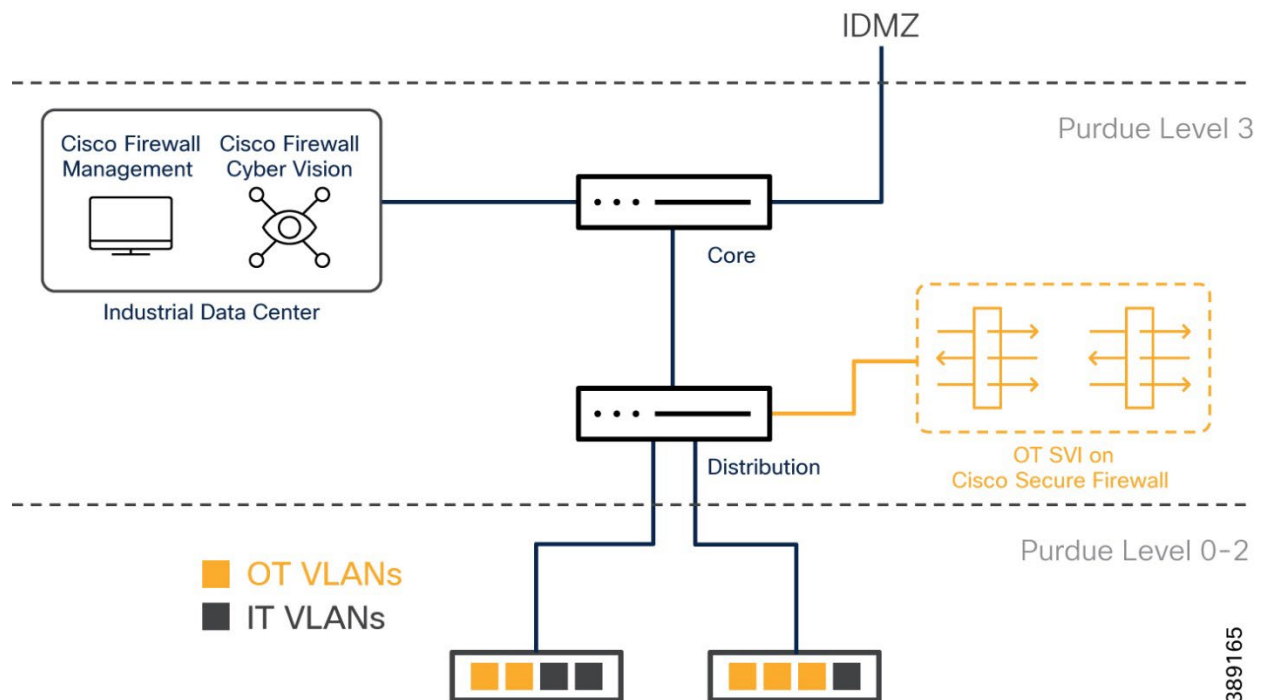
Figure 41 Cisco Secure Firewall as routing point for OT networks

As mentioned [earlier in the document](#), VLANs are a great way to separate devices into their own network. The “problem” with VLAN segmentation, is that each segment is not hidden from the others. Devices can simply communicate from one VLAN to the next through a routing point, which in many operational networks is an L3 switch. Many deployment often rely on the L3 boundary to implement policies in the network, and while a switch based enforcement option will be discussed later in the document (see [Network Access Control with Identity Services Engine](#)), security architects should consider a firewall at this layer of the network.

By using a Cisco Secure Firewall as the point of routing in an OT network, the blast radius will be reduced to an individual VLAN. This approach works well in greenfield environments, where network architects have the luxury of implementing a network design from scratch, and the VLAN structure can be well thought out. It is also useful for transitioning a plant with an existing VLAN structure to a more secure state.

Using a firewall to terminate the VLANs does not have to be an all or nothing approach. For many organizations, the operational network may share IT and OT resources and already using a switch network to route between subnets. In this case, some of those VLANs could be migrated to a firewall, where the gateway of the VLAN is removed from the distribution switch and instead given to the firewall.

Figure 42 Cisco Secure Firewall for OT VLANs only




In this design, non-OT (or traffic deemed less critical) would continue to use the switching infrastructure to route between VLANs, and selective traffic would be subject to firewall enforcement between networks.

Transparent firewalls at the Cell/Area zone boundary

The reality of OT networks is that many organizations are dealing with a large flat network, and creating VLAN segments is not always an option. In this scenario, the Cisco Secure Firewall can be deployed in transparent mode to act as a bump in the wire for all traffic that traverses the cell/area zone boundary.

Figure 43 Cisco Secure Firewall in transparent mode



Using bridge groups, firewall ports can be grouped together to pass traffic between the interfaces. Each directly-connected network must be on the same subnet. This deployment mode should be considered for brownfield deployments, where the network is relatively flat, and cannot be changed to accommodate VLAN separation. The distribution and core switches will continue to be responsible for passing packets throughout the network, allowing firewall rules can be added without any changes to the network.

Note: There will need to be some period of downtime, as cables will be unplugged to position a firewall in between. To reduce unnecessary outages in production, it is recommended to provision the firewall(s) in a lab environment before moving into the plant network.

When inserting a firewall transparently in an OT network, the firewall will interfere with traffic communicating in the same network domain. OT protocols expect deterministic network latency and jitter (for example, Profinet). When inserting a firewall device into an existing brownfield environment, or even when planning a new deployment, it is of paramount importance to test the

end-to-end latency in a staging or proof of concept (POC) area to ensure the firewall can handle data with minimal latency to avoid disrupting real-time operations.

Network Access Control with 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

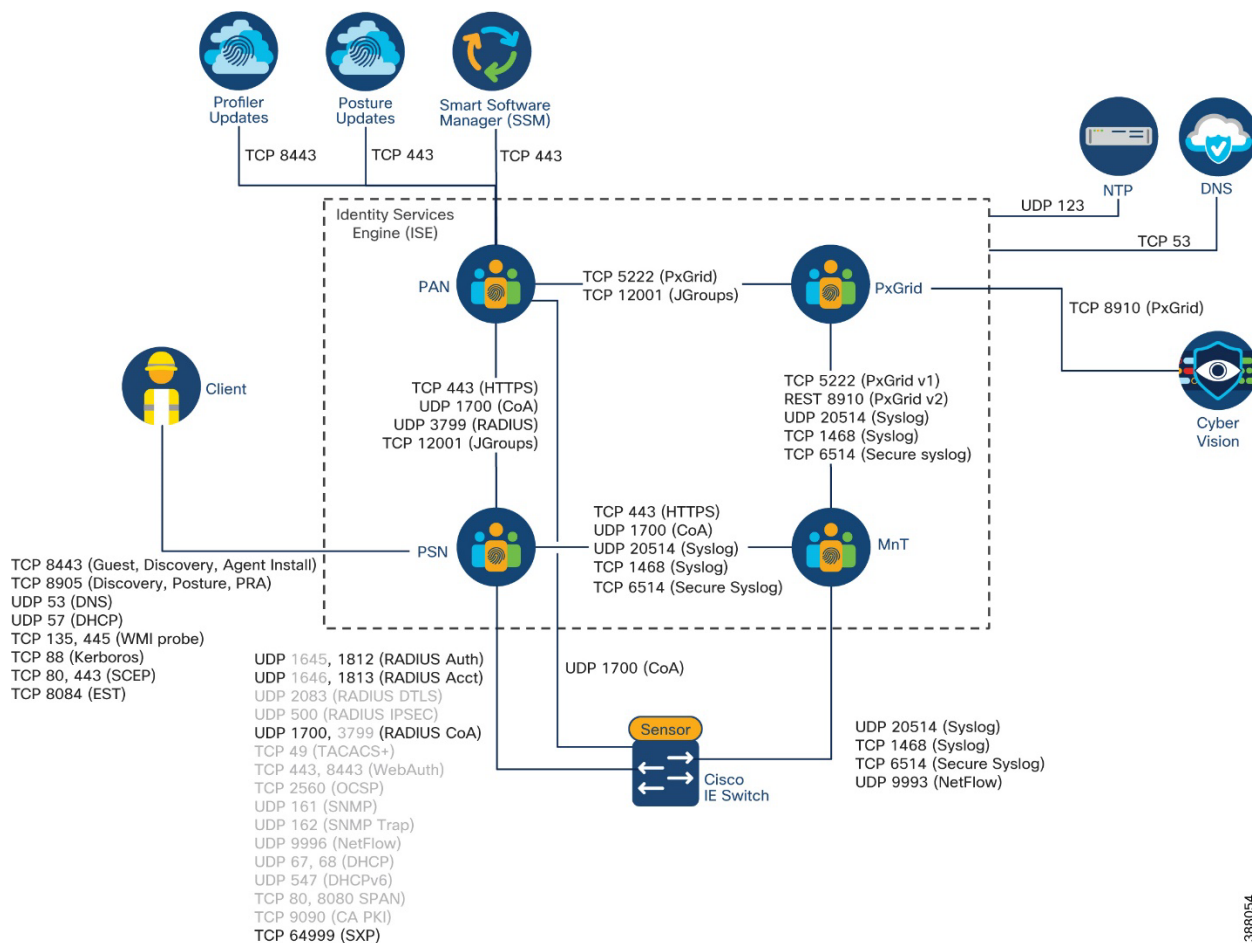
Cisco ISE has four distinct personas/nodes that can either be deployed in one standalone deployment (all personas residing in a single ISE node) or distributed across the network. The personas available in ISE are:

- **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. Figure 44 depicts the communication flows required by ISE Cisco ISE.

Figure 44 ISE Communication Flows



Note: Not all flows depicted in this diagram were used in the creation of the 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 being used for more than what is portrayed in 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

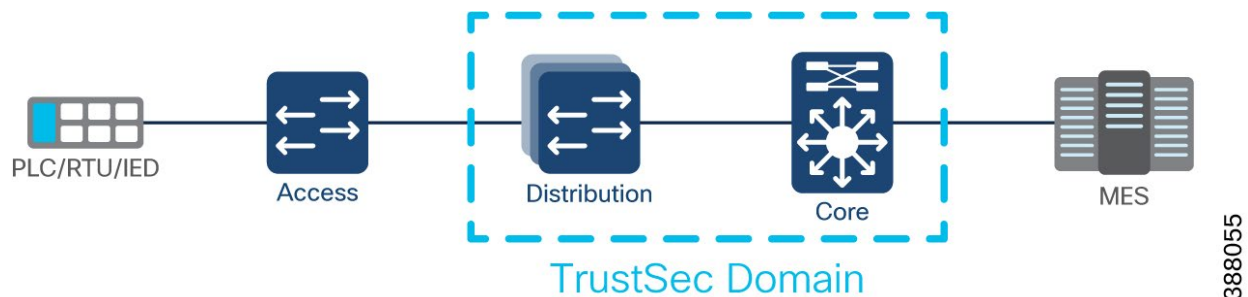
Authorization is the process of enforcing policies and determining what type of activities, resources, or services a user or device is permitted to access. All controlled from a central location, Cisco ISE distributes enforcement policies across the entire network infrastructure. Administrators can centrally 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 45 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.

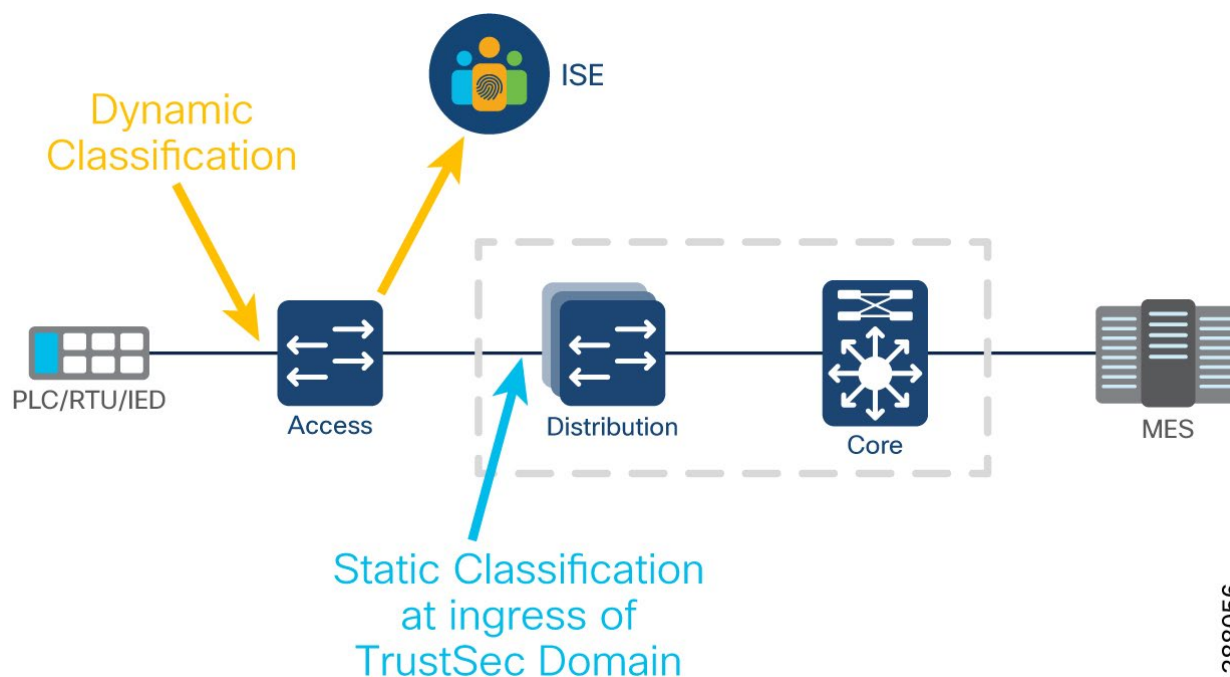
SGT Classification

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 46 Static vs Dynamic Classification



388056

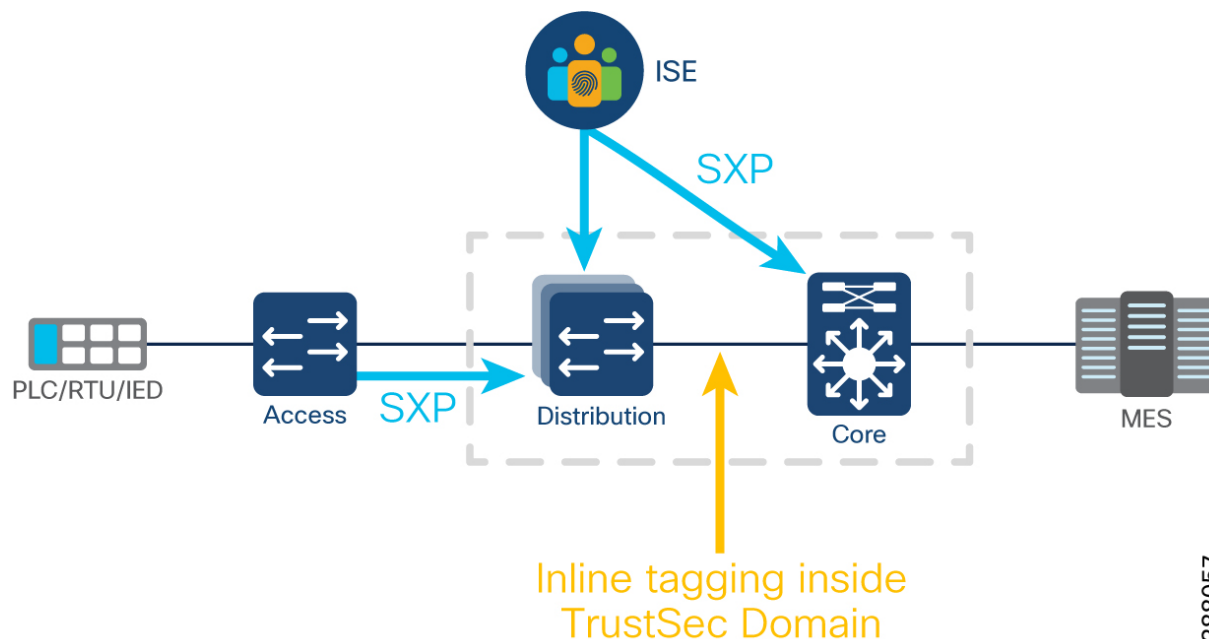
Generally, dynamic classification is done at the access switch and static classification is performed at ingress to the TrustSec domain. The SGT tag that gets inserted into the traffic is known as the source SGT, as it is the group that the source of the traffic belongs to. The destination SGT is the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

SGT Transport

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

SXP is used to propagate the SGTs across network devices and network segments that do not have support for inline tagging. SXP is a protocol used to transport an endpoint SGT along with 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 47 SXP vs. Inline Tagging



388057

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

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

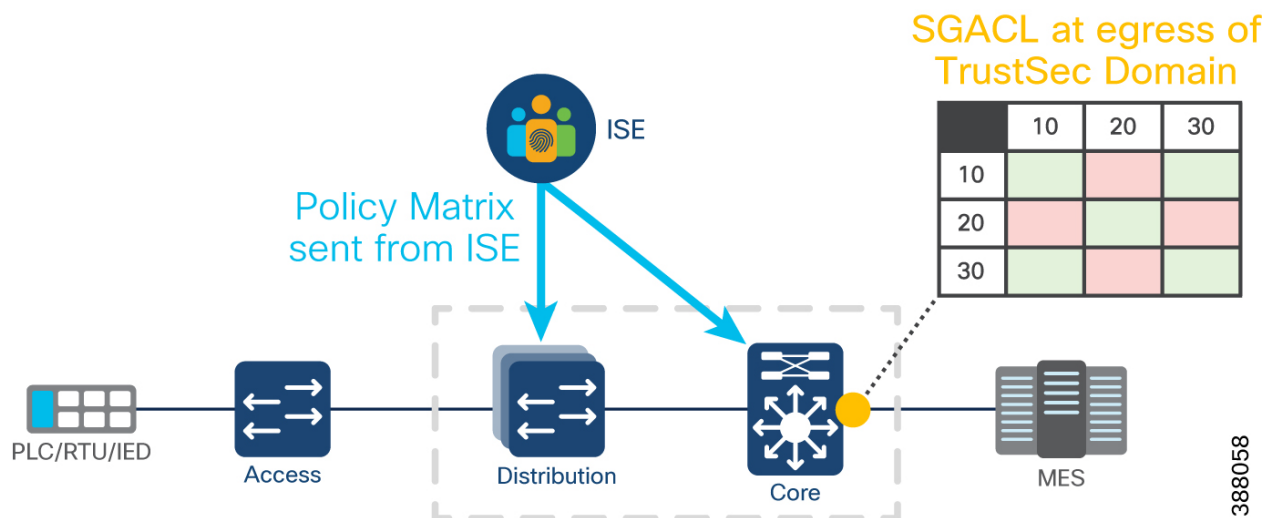
SGT Enforcement

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

It is important to note that the source and destinations are specified in the policy matrix and not in the SGACL. Take, for example, the SGACL entry (ACE) *'deny tcp dst eq 21'*. This entry specifies 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 48 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

Figure 49 shows an example of traffic entering and exiting a TrustSec domain.

- 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 49 TrustSec Classification, Transport & Enforcement Example

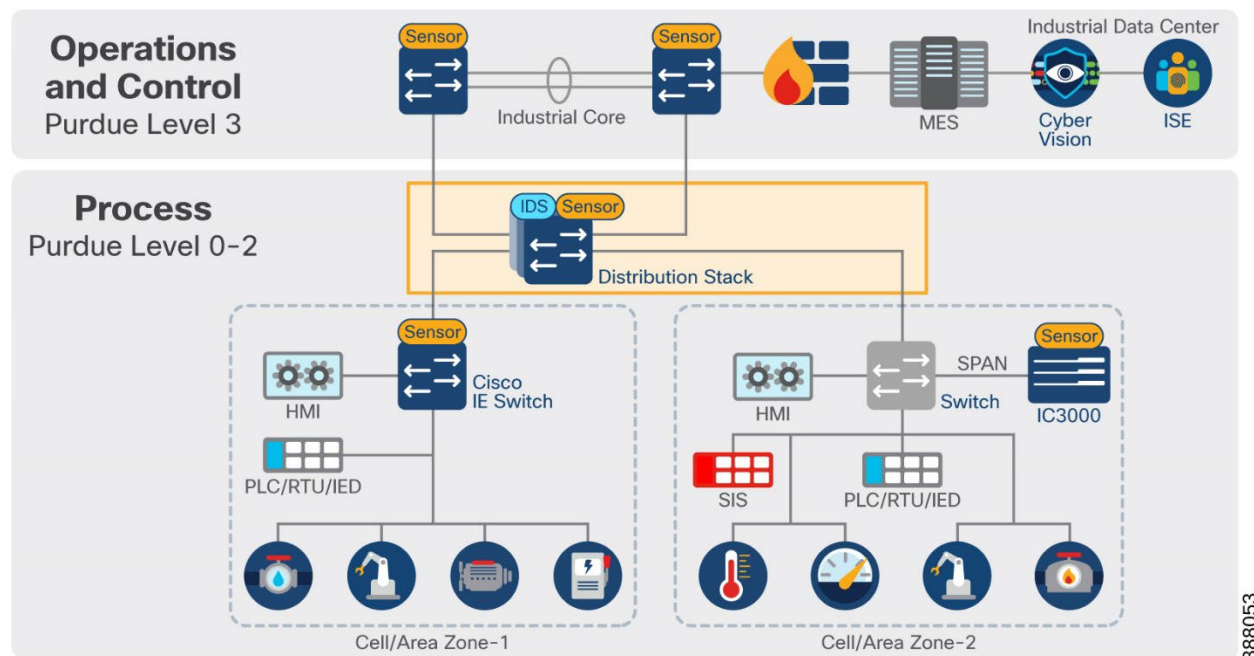


Macro-Segmentation with ISE

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 **macro-segmentation**. 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 50 Policy Enforcement Across the Distribution Switch



The layer 3 boundary and gateway for devices is in the distribution switch as shown in figure 50. 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 using ISE

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 macro-segmentation 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.

ISE / SGT Design Considerations

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

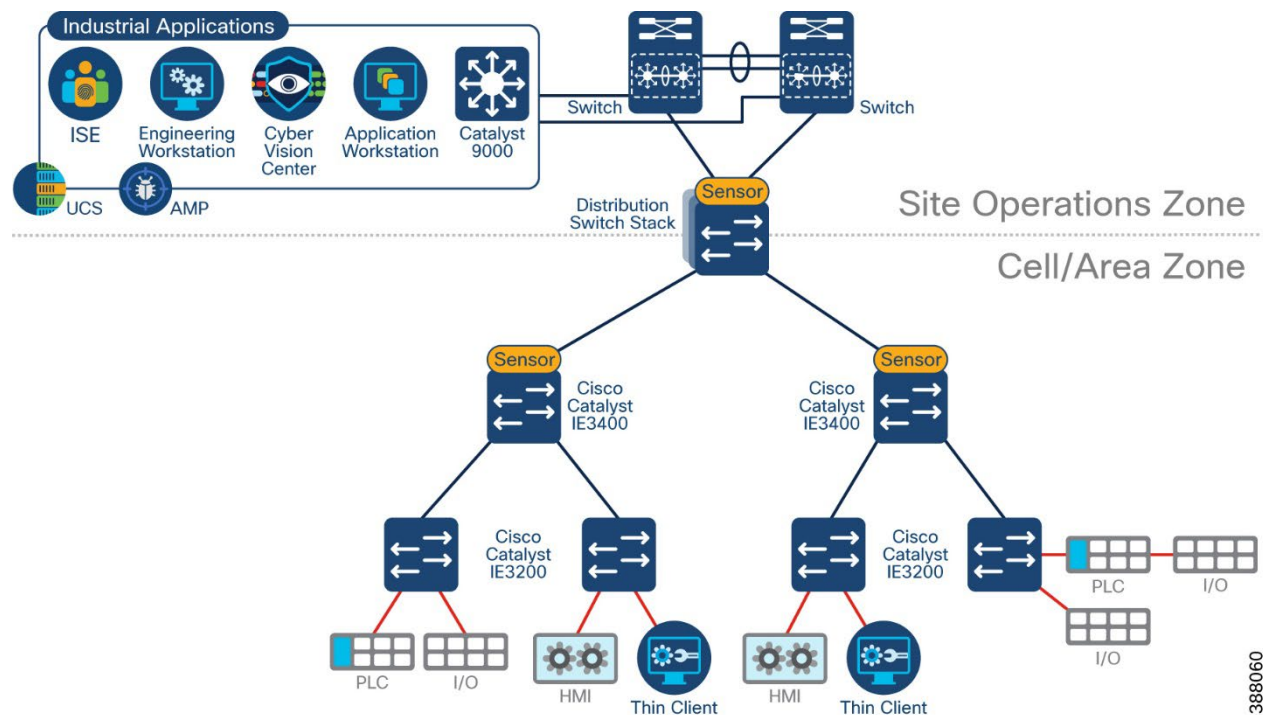
- While ISE has the capability to apply tags via authentication and authorization (AA) policies, it is not recommended to assign an SGT to every device on the network during the authentication process. Use a hybrid approach between macro and micro segmentation, where most of the rules you create are based on the zones in which a device resides, not based on what the device is 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 were 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 (to be explained later in the document).
- 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.

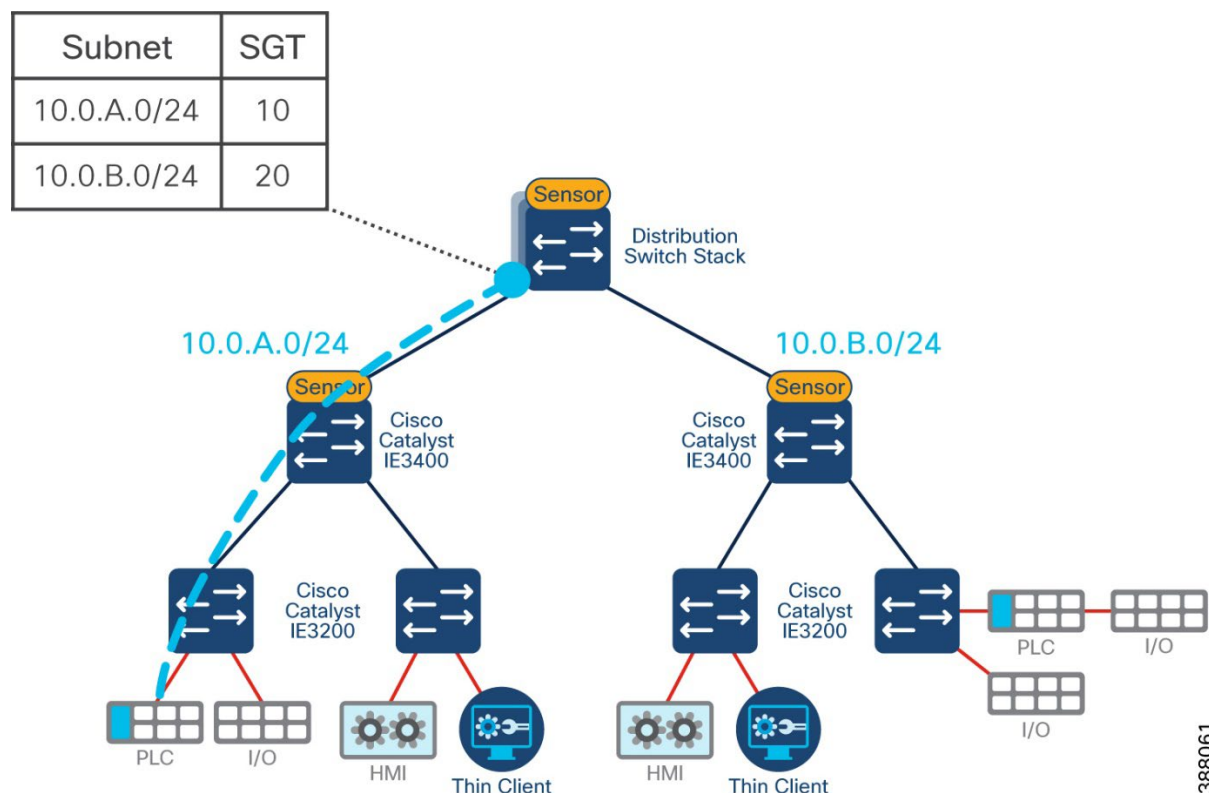
- In this design model, the default action is Deny IP and hence the required traffic should be explicitly permitted with the use of SGACLs. This is generally used when the customer has a fair understanding of the kind of traffic flows within their network. This model requires a detailed study of the 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 switch-initiated 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 51 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. Figure 52 provides an example, where one cell/area zone is tagged with SGT 10, and the other is tagged as SGT 20.

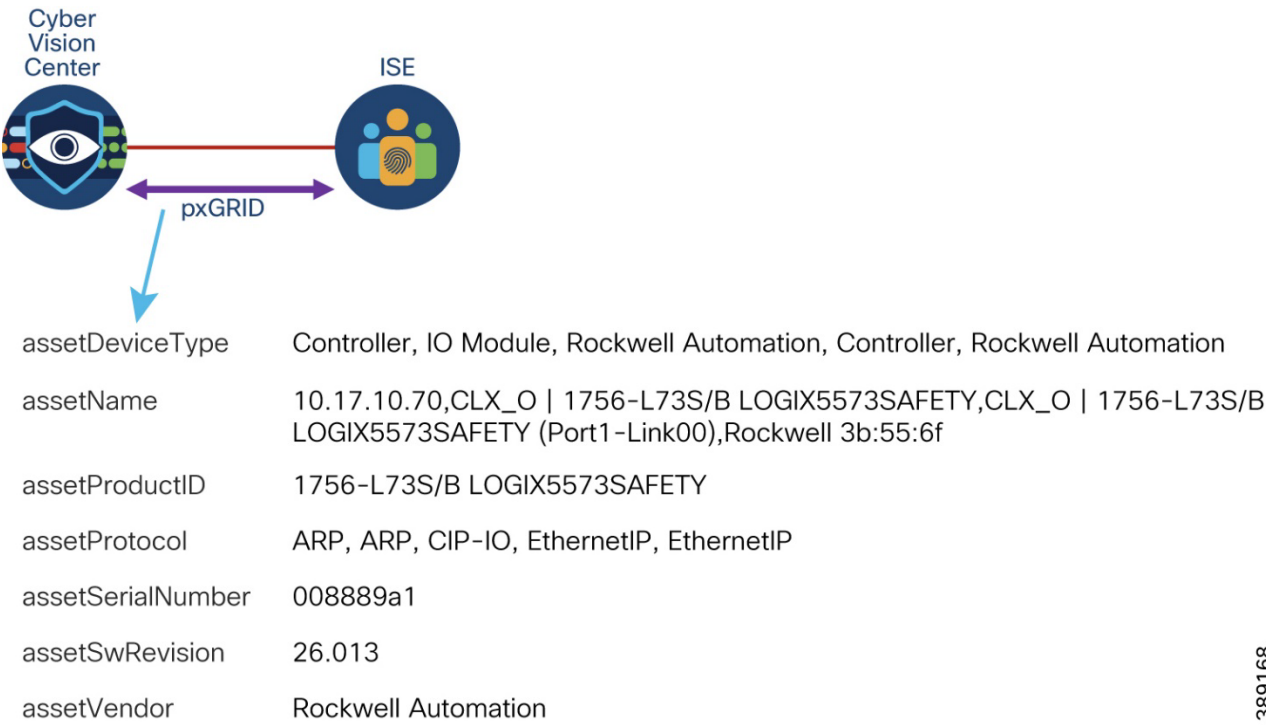
Figure 52 SGT Classification at Distribution Ingress

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 (such as, 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.

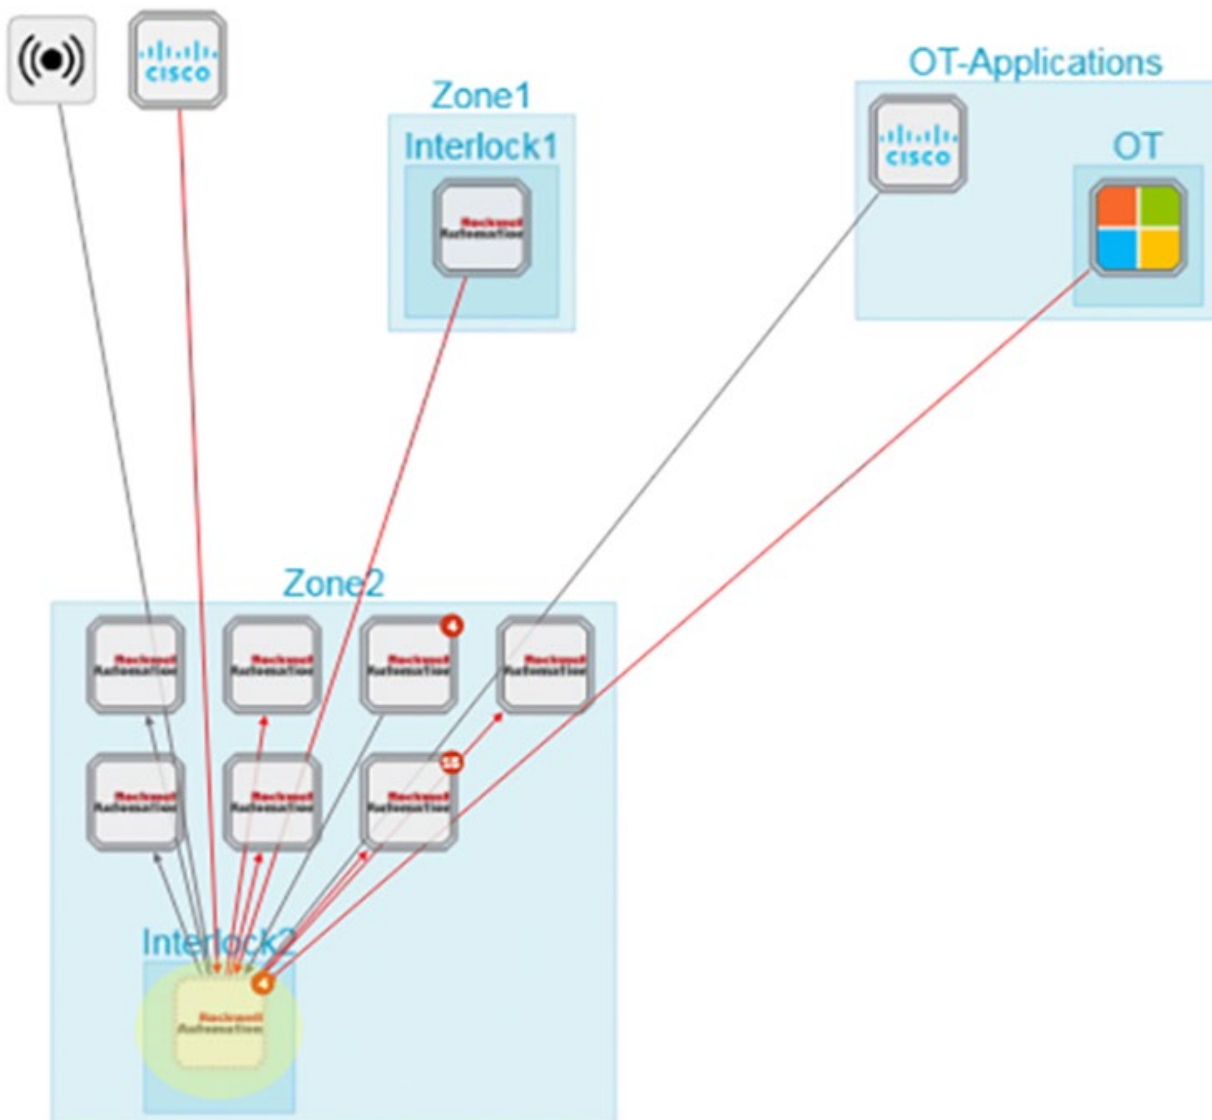
Figure 53 Cyber Vision & ISE pxGrid Integration



389168

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 54 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 55 ISE Asset Information after Cyber Vision Integration

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

MAC Address: 00:00:BC:2D:21:70
Username: 00-00-BC-2D-21-70
Endpoint Profile: CVC_group_Interlock2
Current IP Address: 10.17.20.72
Location: Location → All Locations

Applications

Attributes

Authentication

Threats

Vulnerabilities

General Attributes

Description

Static Assignmentfalse

Endpoint PolicyCVC_group_Interlock2

Static Group Assignmentfalse

Identity Group AssignmentCVC_group_Interlock2

Custom Attributes

Filter

Attribute String	Attribute Value
assetGroup	Interlock2
assetCCVGrp	
assetSource	CCV

assetDeviceType

Controller, IO Module, Rockwell Automation,Controller, Rockwell Automation

assetId

00:00:bc:2d:21:70

assetIpAddress

10.17.20.72

assetMacAddress

00:00:bc:2d:21:70

assetName

10.17.20.72,CLX_P | 1756-L73S/B LOGIX5573SAFETY,CLX_P | 1756-L73S/B LOGIX5573SAFETY (Port1-Lin k00),Rockwell 2d:21:70

assetProductId

1756-EN2T/A,1756-L73S/B LOGIX5573SAFETY

assetProtocol

ARP,ARP, CIP-IO, CIP Safety, EthernetIP,EthernetIP

assetSerialNumber

00552b01,00893b40

assetSwRevision

26.013,5.028

assetVendor

Rockwell Automation

Once zones have been defined, and traffic has been classified, the policy enforcement matrix can be defined. Figure 56 shows an example policy enforcement matrix.

Figure 56 Sample TrustSec Matrix

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

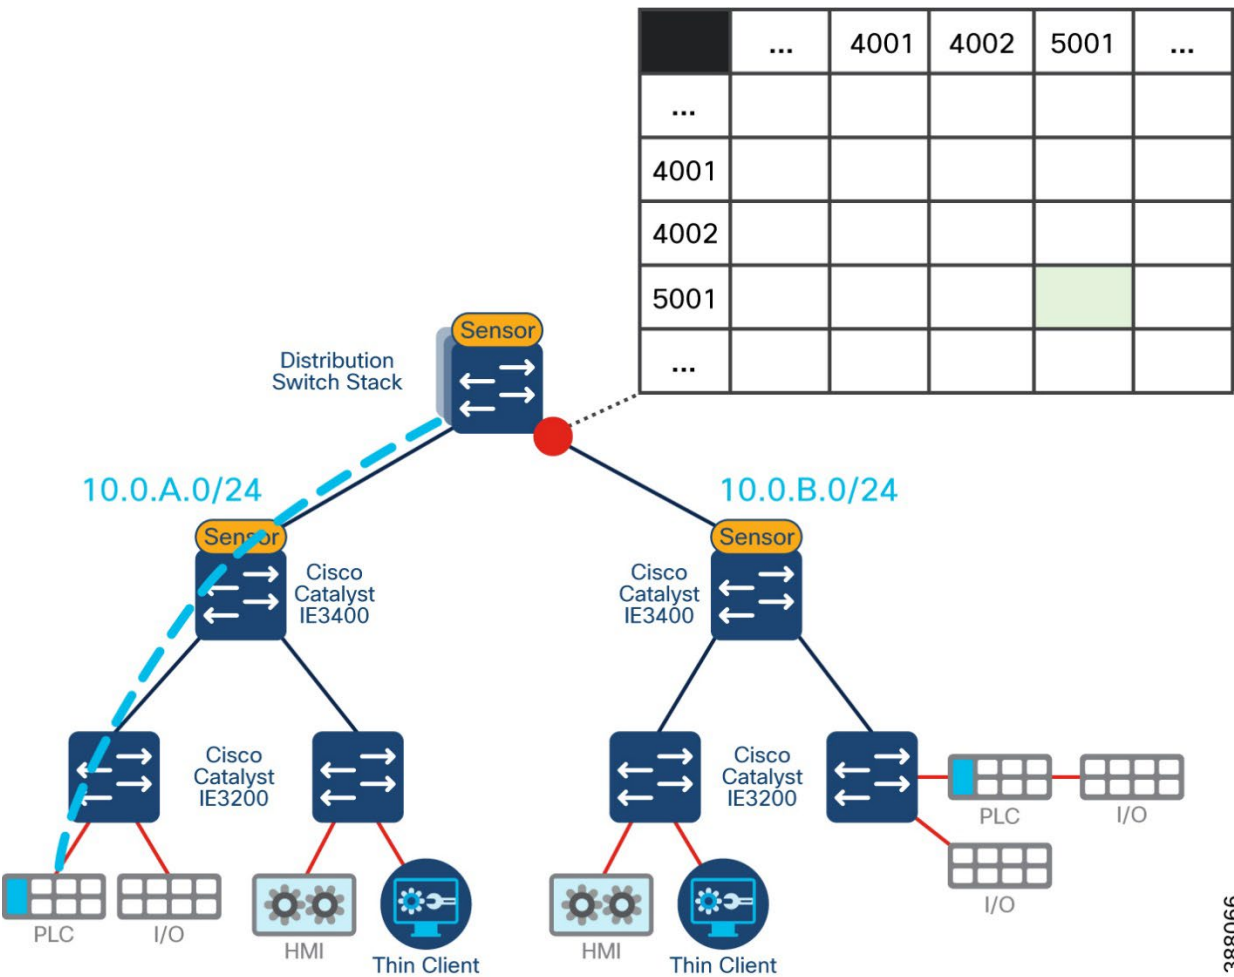
SGT	Group
100	Infrastructure
101	Management Apps
102	Plantwide Apps
103	Cyber Vision
911	911 Tag
4001	Zone 1
4002	Zone 2
5001	Interlock Zone 1
9001	Super User
9002	TrustSec Devices

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 re-iterate the policy enforcement point used in this example is the distribution switch. Let us take Figure 57 as an example, 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 57 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

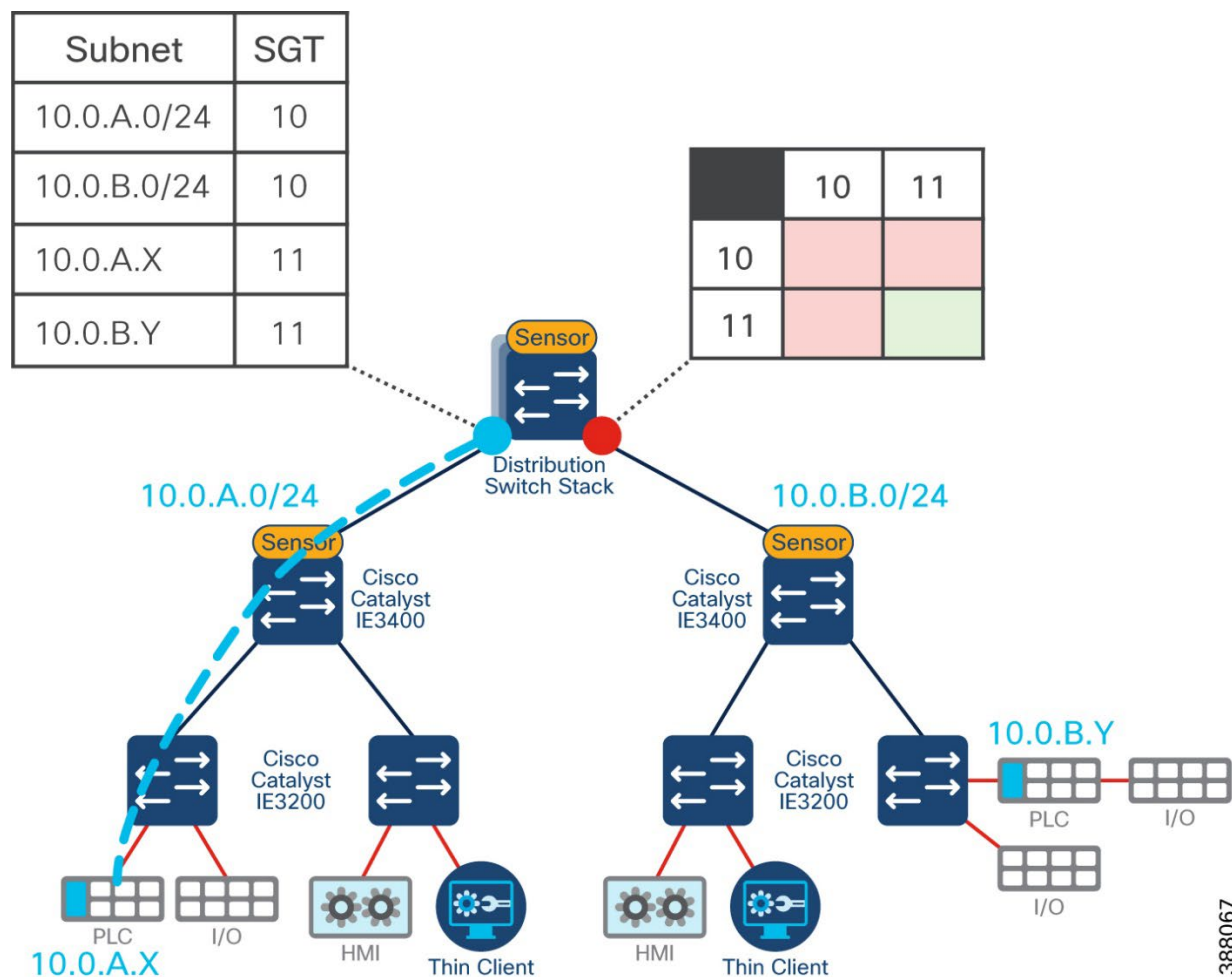
understand traffic that flows across L3 boundaries and use that information to inform policy creation.

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.

Scale Considerations in Large Networks

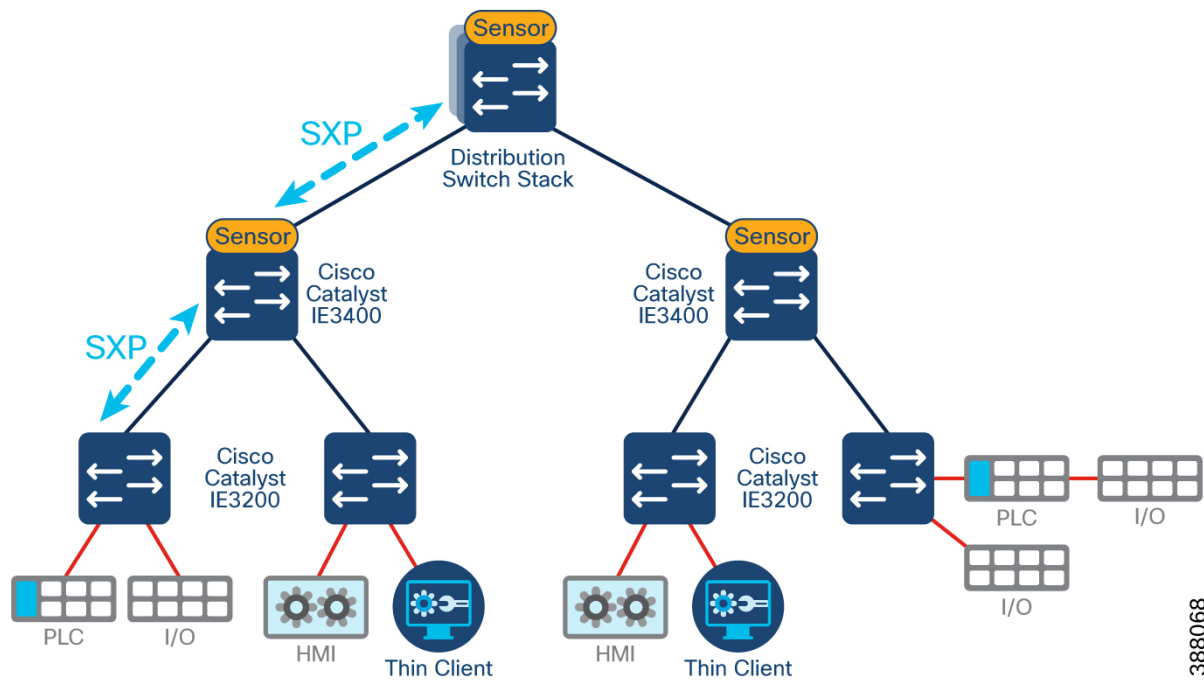
For security to be effective, it needs to remain simple. For larger networks, where there may be hundreds of zones to manage, it may not be effective to create a unique tag for each zone. Take an example, where 400 cell/area zones exist in the industrial zone. This would result in a matrix that is at 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. Figure 58 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 58 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 Catalyst Center.

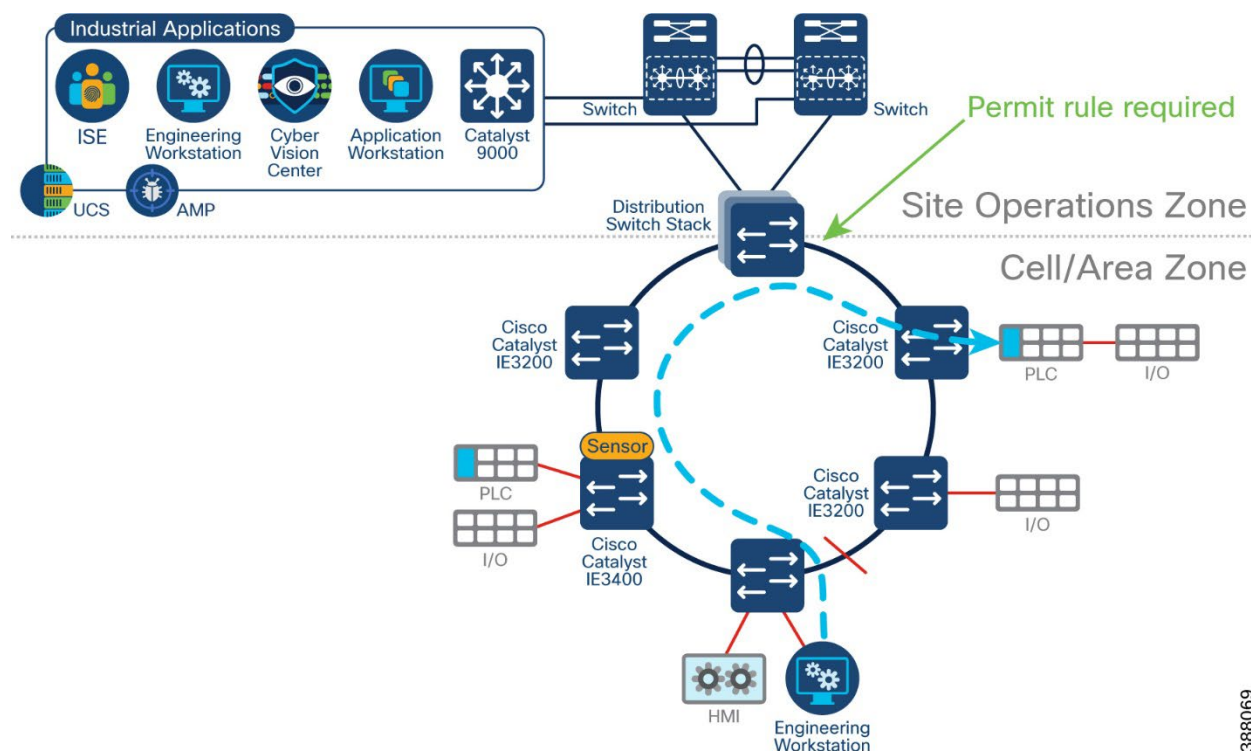
Figure 59 SXP Daisy Chain starting from Access Switch up to Distribution Switch

388068

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. Figure 60 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 60 Inter-Cell/Area Zone traffic traversing the Distribution Switch



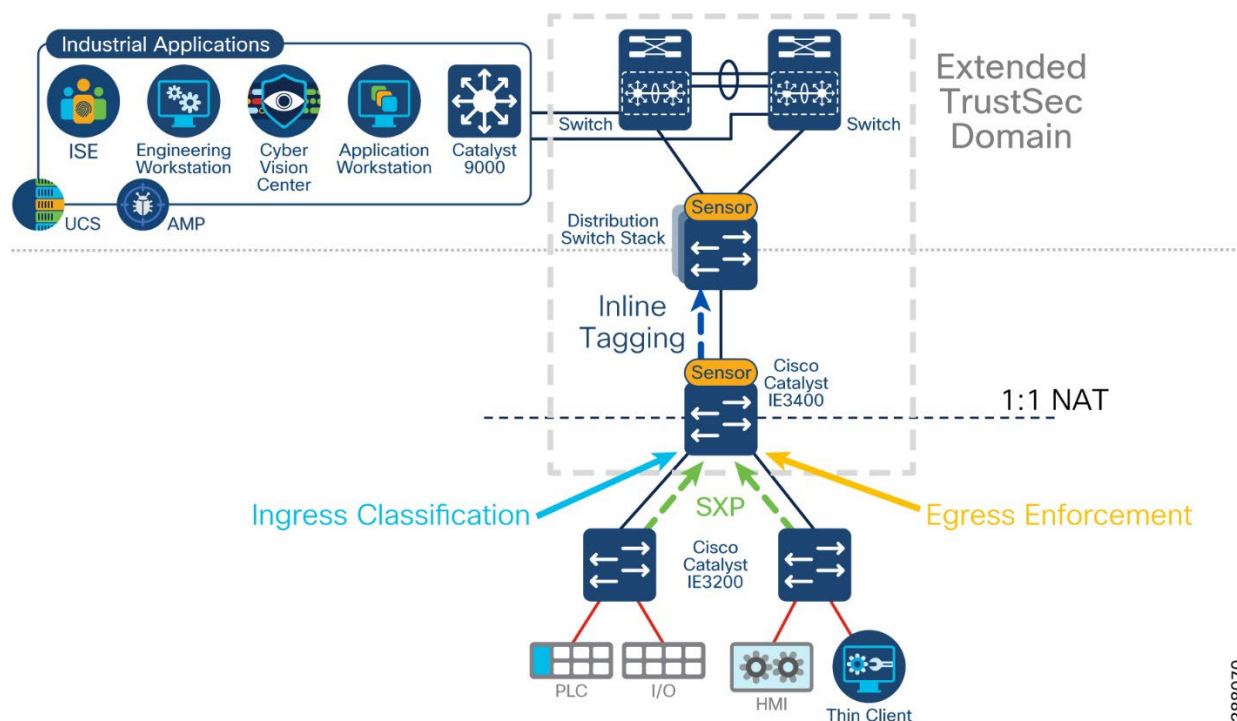
To ensure traffic is not blocked by policy, make sure that each such ring has its own unique SGT and does not share a tag with any other zone in the network as per the design recommendation for large networks.

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.

69069

Figure 61 L2 NAT in the Industrial Zone



388070

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

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

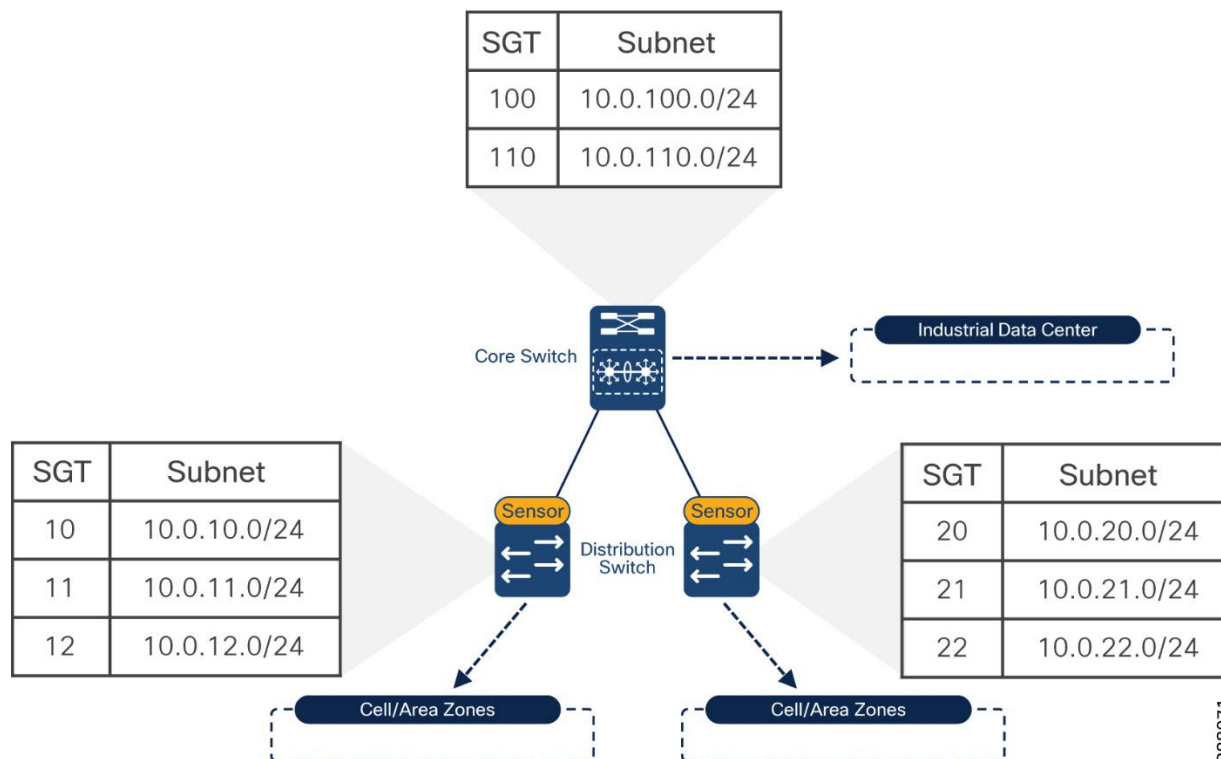
Note: When using Cyber Vision to assign groups to devices behind the NAT boundary, it is important that you choose the correct device. Depending on sensor placement, Cyber Vision may show two instances of a component when NAT occurs. One will be the instance before the NAT, the second will be after.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

SXP Domain Filters

An SXP Domain is a collection of SXP devices. There is a default SXP domain that all devices will join when creating SXP sessions. Devices in the default domain will receive all SGT-IP mappings that are known by ISE. SXP domain filters provide a mechanism for SXP peers to deviate from the default, and only receive the IP-SGT mappings that are required for their function on the network. For example, all IP-to-SGT mappings learned through RADIUS authentications are automatically added to the default domain but can be reassigned to a different

domain using SXP Domain filters. As a result, any dynamically assigned SGTs can be communicated to the enforcement point that protects assets on that subnet, rather than every switch requiring to store all entries.


Figure 62 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 63 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

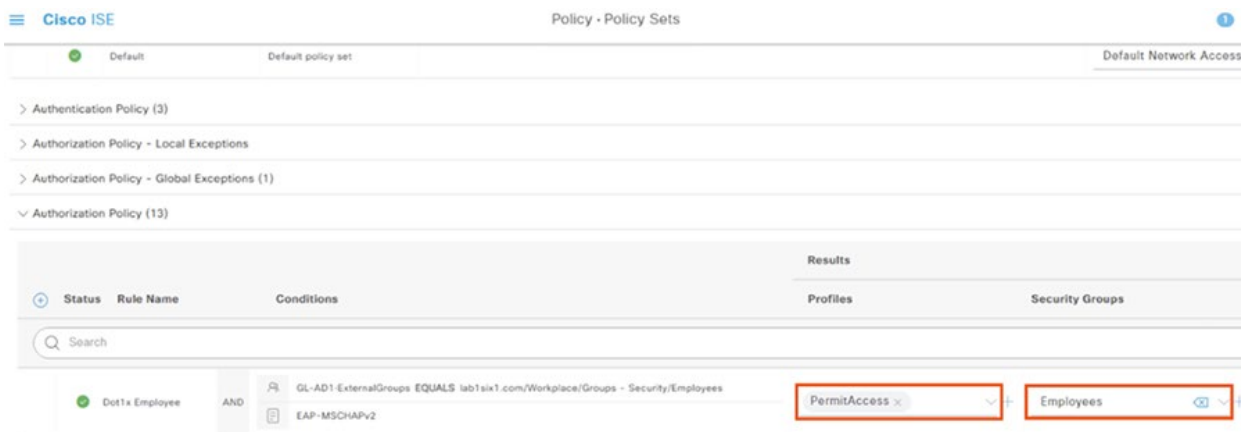
802.1X is an IEEE standard for layer 2 access control, offering the capability to permit or deny network connectivity based on the identity of the end user or device. 802.1X is typically not supported in OT devices, however, it should be a common feature on an employee or contractor end-user device such 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. Figure 64 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 64 ISE AA Policy with Active Directory User Group



Chapter 3. Preventative Control in Distributed Field Networks

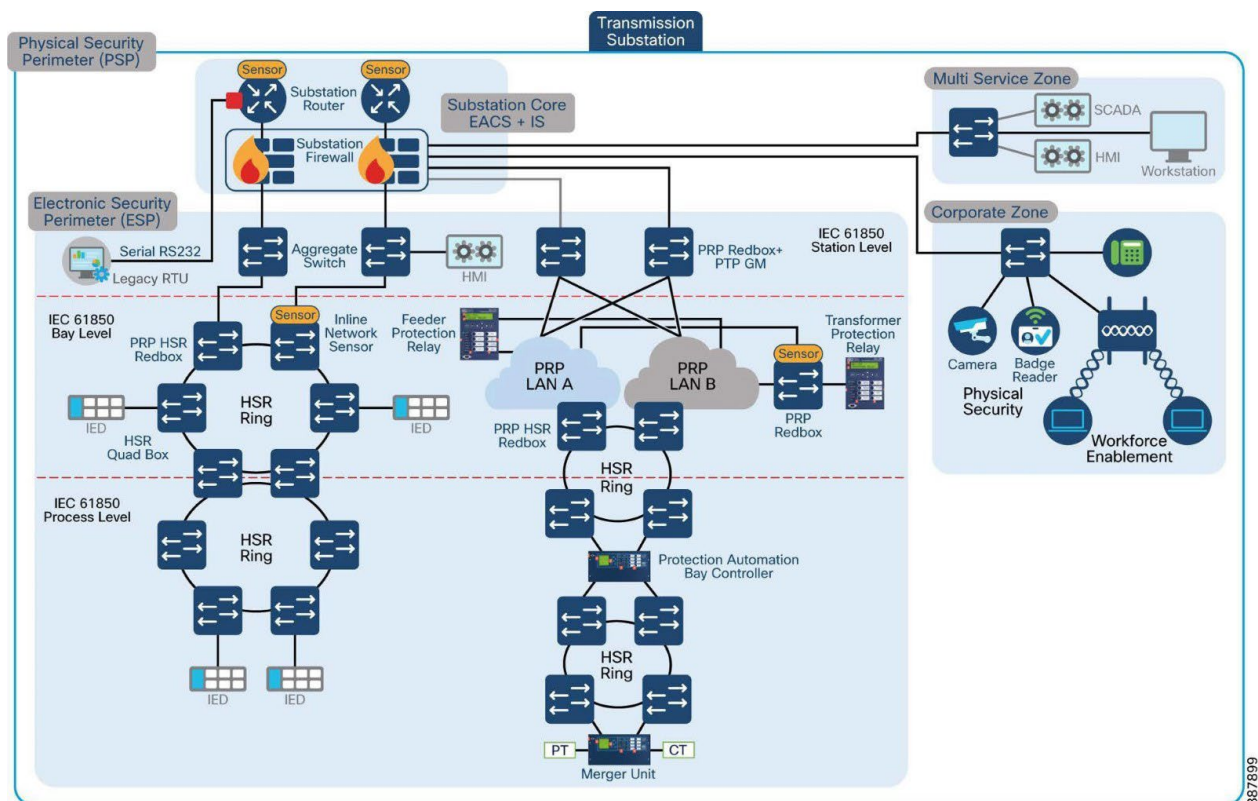
Across all industries, organizations need advanced, agile, and secure Wide Area Network (WAN) infrastructures to connect distributed OT assets to control centers and unlock the potential of digitization. Whether it is about connecting roadways assets, first responder or public transport vehicles, water, oil, or gas infrastructures, renewable energy resources, power substations, EV charging stations, or any critical remote assets, you need rugged routers with cutting-edge cybersecurity capabilities.

As we define the networking standards of the future, Cisco believes industrial routers must become a platform to easily deploy advanced OT security capabilities at scale. In addition to enabling smarter and simpler WAN infrastructures, Cisco industrial routers come with next generation firewall capabilities, malware protection, cloud security, and threat intelligence feeds to help you build secure distributed networks so you can run modern industrial operations with peace of mind.

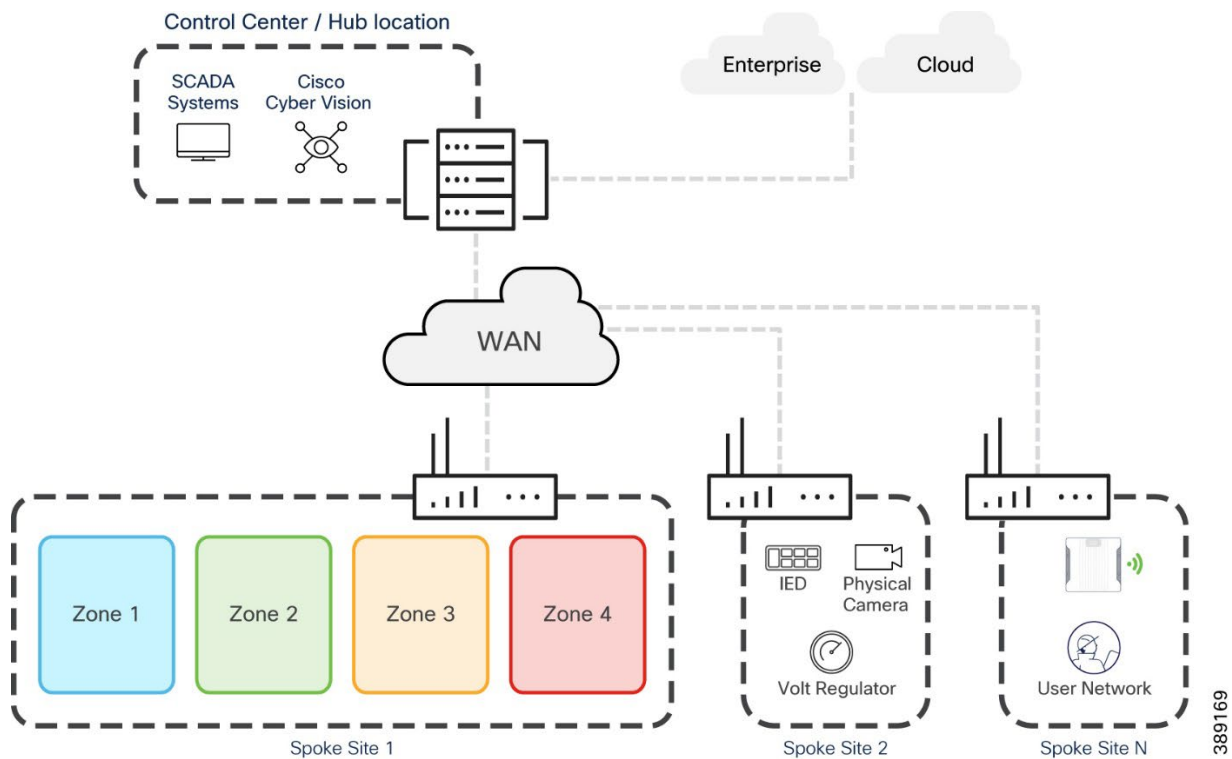
Reference Architecture

This design guide is written to be vertical agnostic, providing design guidance for securing critical infrastructure of all kinds across the WAN. To accomplish this, a generic architecture diagram will be used to capture the architectural concepts found in industries such as utilities (see figure below) and transportation. Details of these architectures can be found at cisco.com/go/iotcvd.

Figure 65 Transmission Substation architecture



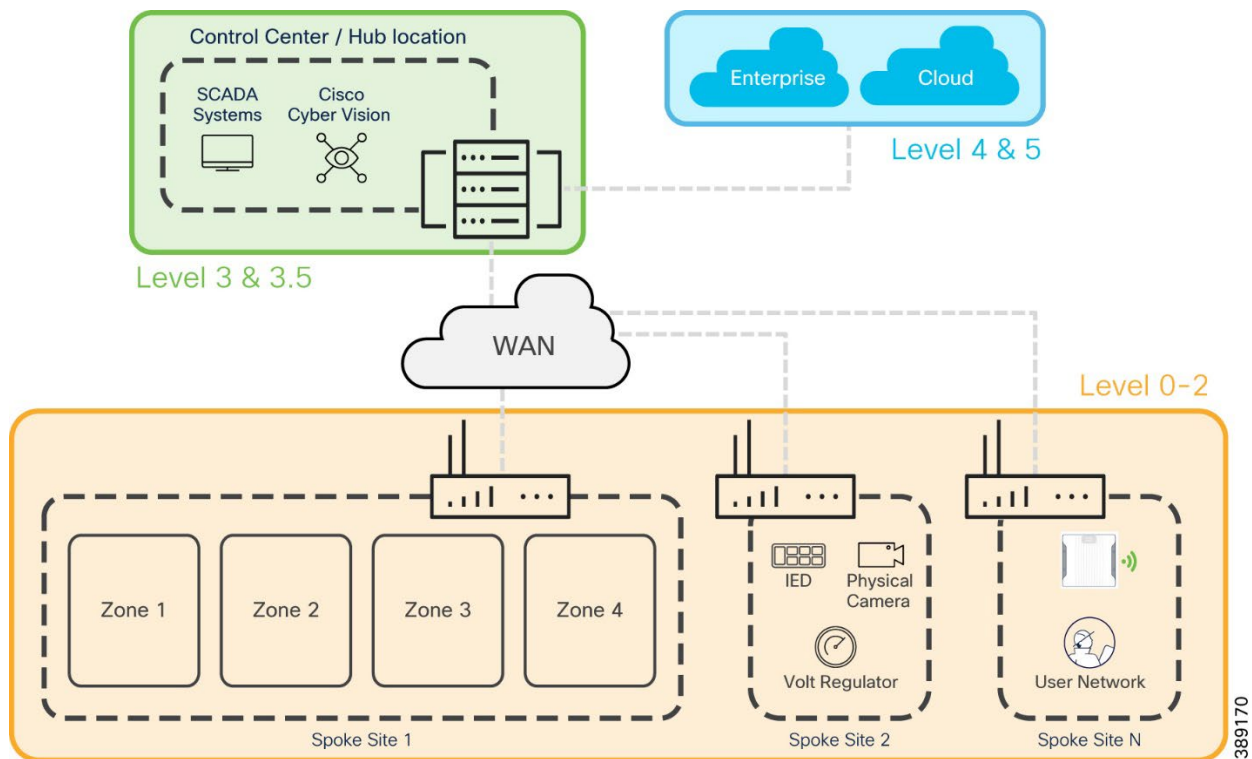
In this design, the reference architecture for the field network will be split into four zones, with “Zone 1” depicting the most critical part of the network. For example, Zone 1 could be the electronic security perimeter in a transmission substation, or the traffic controller in an intersection. The other three zones will be of lesser importance, and may represent things such as physical security, users or telemetry networks.

Figure 66 Reference architecture for distributed field networks

The field networks are connected over a WAN in a hub and spoke architecture to a control center. Each site may also have different security requirements. The primary site is depicted with four zones, but smaller spoke sites may exist in a real world deployment offering less functional zones. The transport details of this are not in scope for this design guide (i.e. cellular, MPLS, etc.). As will be mentioned later in the document, the design guide will primarily focus on SD-WAN design guidance, which is transport agnostic.

Lastly, the IT/OT boundary is located at the hub site. If this architecture was to be referenced using the Purdue model, the field sites would be level 0-2 and the control center would be level 3 and 3.5, which represent the DMZ between IT and OT networks.

Figure 67 Reference architecture mapped to the Purdue model

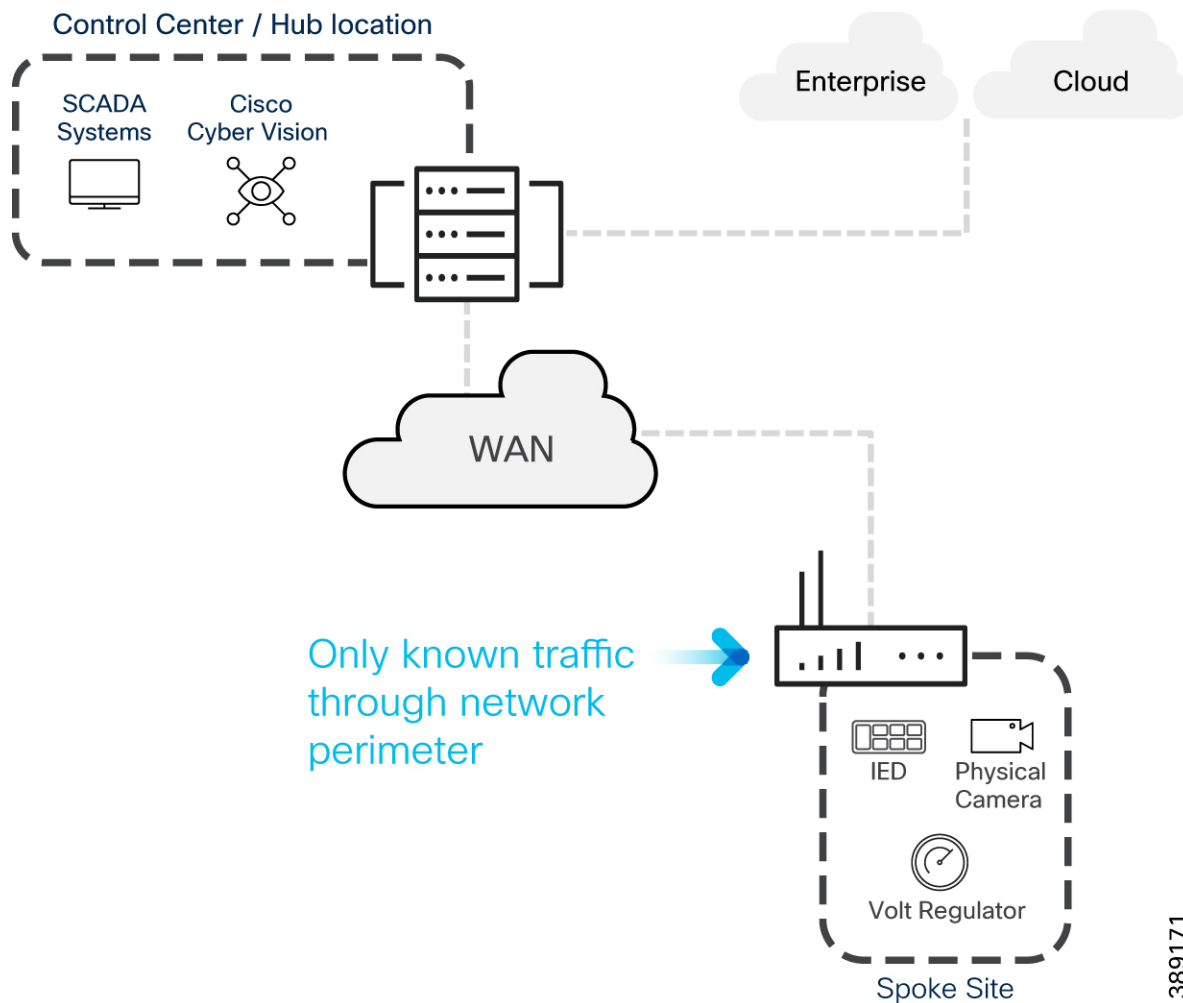


Use Cases

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

- **Deny by default on traffic leaving the field network.** As critical infrastructure gets connected to an IP network, only trusted and verified traffic should be enabled to cross the WAN.

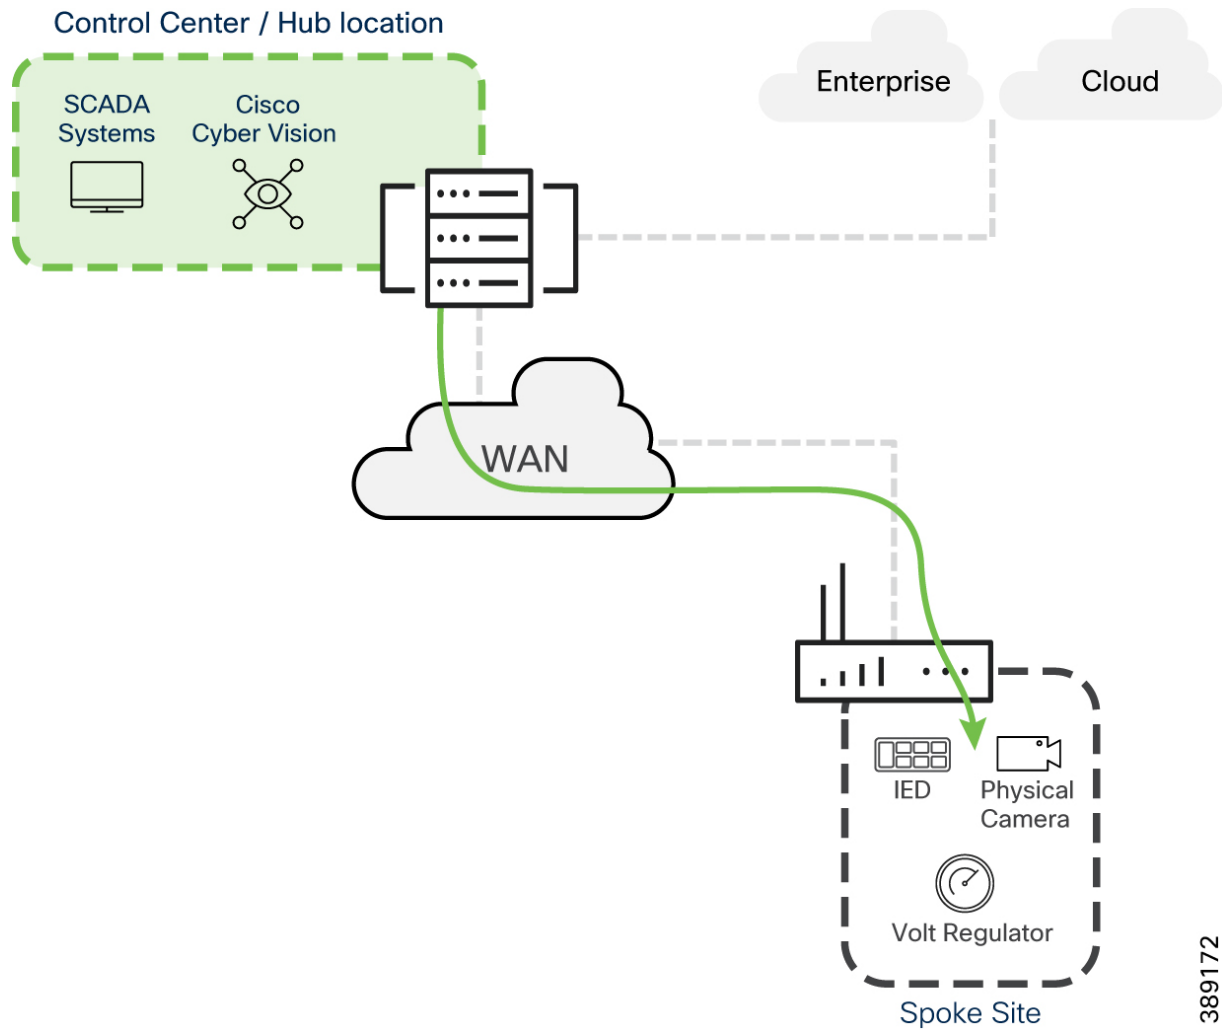
Figure 68 Deny by default firewall at the field edge



389171

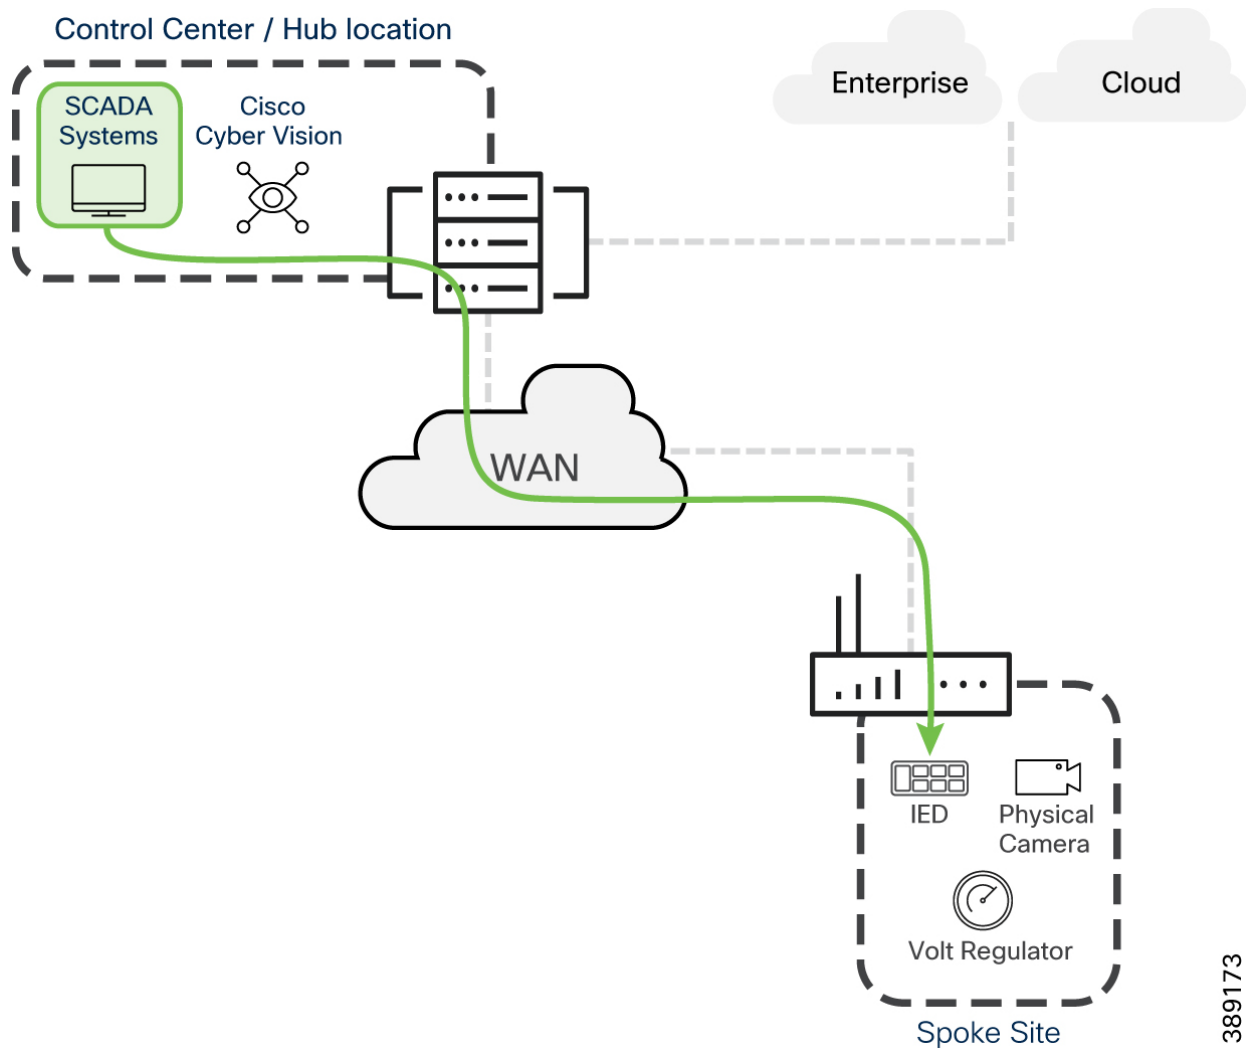
- **Applications hosted in a centralised data center accessing field assets.** The industrial data center for field assets typical resides in a central hub location and accesses field assets across a WAN network. Operators need a secure way to control these flows so only trusted servers have access to the critical assets.

Figure 69 Hub location has access to assets in the field



- **Read only access across the WAN for SCADA systems.** Not all servers need control over assets in the field. Some SCADA or IoT systems need read only access to the controllers to gain telemetry information and by reducing their privilege to read only access it reduces risk of an attack due to the server being compromised.

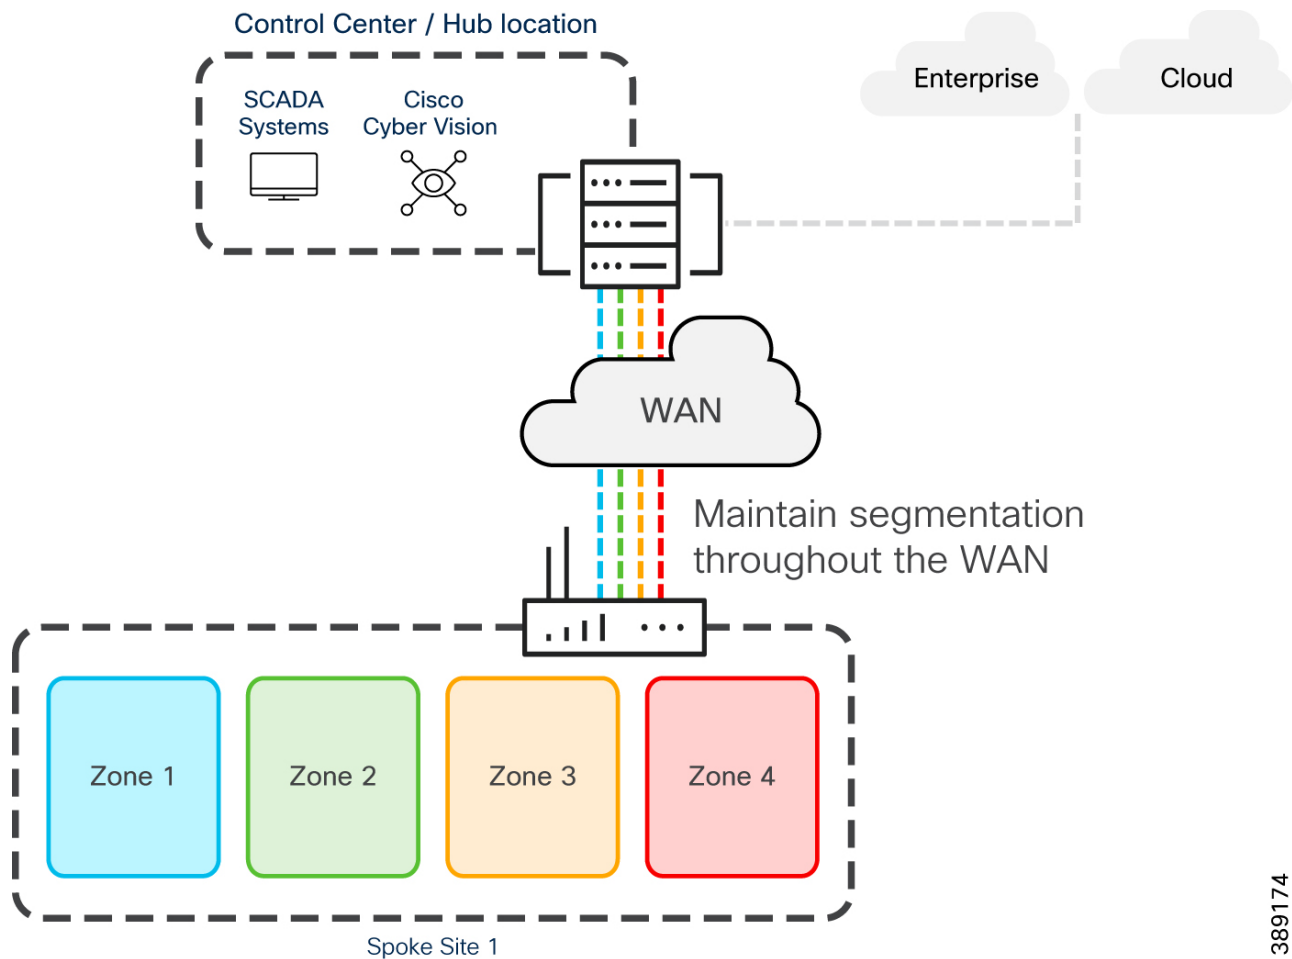
Figure 70 SCADA system with read only access to assets in the field network



389173

- **Segmentation of assets at the site.** Cabinets in the field are often space constrained and run many different services, from traffic controllers, to cameras, to badge readers. If one domain was to be compromised in an attack we must ensure the others remain operational.


Figure 71 Segmentation in the field site continued across the WAN



389174

- **Remote Access.** Enabling remote access at the field could save hours, maybe even days of a travel time for maintenance operators depending on the vertical. Offshore windfarms is a good example of this. Downtime is costly, and the quicker a resolution can be implemented the better. [Securing remote access](#) has a dedicated chapter in this design guide.

Figure 72 Remote access to field assets



389175

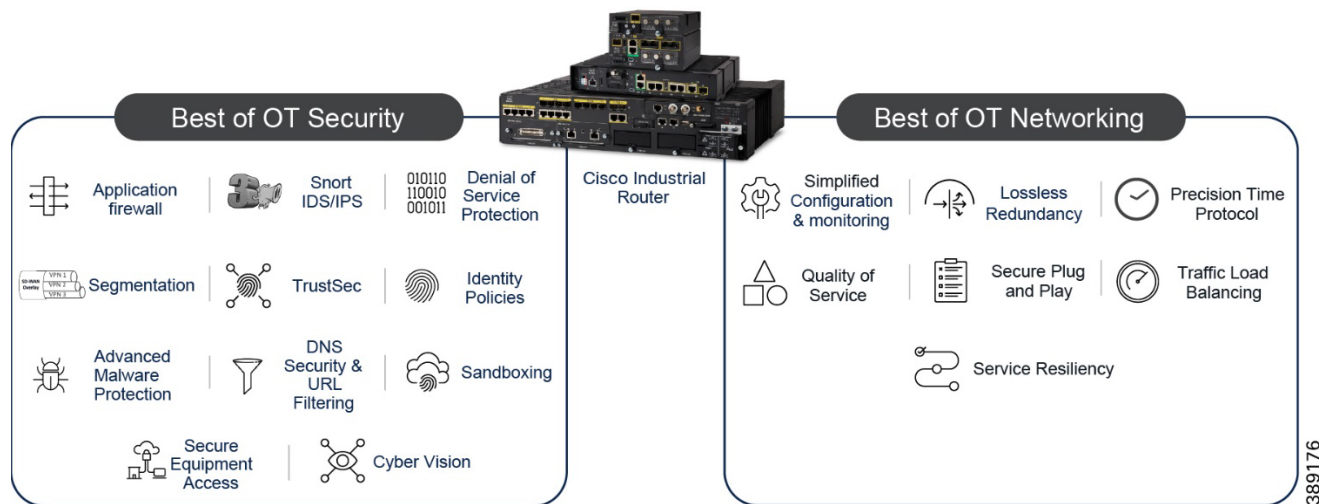
Cyber Resiliency vs. Cyber Security

While we often talk a lot about cyber security, which refers to the robust tools and policies implemented to prevent attacks from occurring in operational networks, we often overlook cyber resiliency. Cyber resiliency refers to an organizations ability to maintain its critical operations even in the face of cyber-attacks.

Cyber security is of course part of a cyber resiliency architecture. Capabilities such as firewalls, segmentation and the implementation of a zero trust model means that if an attacker does get a foothold in the network, their reach is limited and both reconnaissance and lateral movement can be prevented. However, cyber security practitioners and networking teams often make the mistake of treating themselves as siloed entities in the organization. The network configuration is

just as important as the security appliances deployed in the network. Quality of Service (QoS) ensures that critical traffic always has priority when the network is in a degraded state, lossless redundancy protocols ensure that critical traffic meets latency metrics when network paths go down, management plane security ensures only trusted users get access to the network infrastructure and cannot be taken down by malicious actors, and plug and play ensures that new network devices are onboarded with a secure configuration out of the box. While all these features are typically considered part of networking, it is the combination of networking and security that results in a cyber resilient architecture.

Figure 73 Best of OT networking and OT security with Cisco Industrial Router



Cisco Industrial Router

This design guide focuses on the design components, considerations, working and best practices of using the Cisco Industrial Router to secure critical infrastructure. However, the document is not meant to exhaustively cover all capabilities in the devices.

Portfolio

[Cisco® Catalyst Industrial Routers](#) offer unconditional connectivity for all your remote assets. They can withstand extreme temperatures, humidity, and dust. They offer a variety of WAN connectivity options, including 5G/ LTE cellular, MPLS, Ethernet, and fiber, through pluggable interface modules that can be easily replaced when needs or technologies evolve. In addition, Cisco Catalyst SD-WAN simplifies deploying and managing a large and complex WAN infrastructure from a central location.

Figure 74 Catalyst Industrial Router portfolio



Catalyst industrial routers also come with comprehensive Next-Generation Firewall (NGFW) features and many more cybersecurity capabilities to block modern threats:

- Standard firewall capabilities like stateful inspection,
- Application awareness and control to block application-layer attacks,
- Integrated intrusion prevention (IDS/IPS),
- Continuously up-to-date threat intelligence,
- Malware protection and sandboxing, URL filtering,
- Integration with a Secure Services Edge (SSE).

Building a modern industrial WAN infrastructure requires advanced routing capabilities such as only Cisco can offer. Having state-of-the-art cybersecurity features built into your industrial routers not only is vital to keep the organization safe, but it is also key to simplify and scale deployment and management tasks. Converging industrial networking and cybersecurity helps ensure unified security policies are enforced across sites, eliminating gaps in defenses due to cost and complexity of integrating many point products together.

Management Options

The primary management option used in this design guide is Cisco Catalyst SD-WAN Manager. Cisco IOS-XE, can be deployed in two operational modes: **Autonomous** and **Controller** modes.

For the purpose of this design guide, Autonomous mode refers to the Cisco Industrial Router being managed standalone (that is, not using Cisco Catalyst SD-WAN Manager). Some administrators may decide to manage the router using the CLI or the web interface, or through Ansible playbook.

Controller mode is when the Cisco Industrial Router is part of the Cisco Catalyst SD-WAN solution. This is the recommended deployment mode and while this design guide will provide a

feature comparison between the two modes, the primary design guidance assumes the use of SD-WAN.

Cisco Catalyst SD-WAN Manager can also be deployed to manage Cisco Catalyst routers in Autonomous mode. This is referred to as software defined routing (SD-Routing), where Cisco Catalyst SD-WAN Manager provides configuration management without the controller infrastructure for WAN orchestration.

Application Firewall

Cisco Industrial Routers offer a stateful firewall with application recognition that organizes the network into zones and enables policy creation for traffic flowing between those zones. Traffic can be evaluated based on:

- Physical and virtual interfaces
- Subnet / IP address
- VLAN
- UDP / TCP port
- Objects
- User ID via ISE
- FQDN
- SGT
- Application

Using network-based application recognition version 2 (NBAR2), Cisco Industrial Routers can examine the data portion of packets, enabling it to identify applications regardless of the port numbers they use. Nevertheless, OT protocols typically use a static port mapping when communicating on the network. For example:

- Modbus: 502
- DNP3: 20000
- Ethernet/IP: 44818
- OPC UA: 4840
- MMS: 102

In instances where NBAR2 may not recognise the protocol by name, security administrators can choose to open/close ports to achieve the same level of control. Alternatively, custom application can be created. IP address and port-based custom protocol enables NBAR2 to recognise traffic based on IP addresses and port numbers and to associate an application ID to that traffic. This is achieved using the `ip nbar custom transport` command. For example:

```
ip nbar custom OPCUA transport tcp port 4840
```

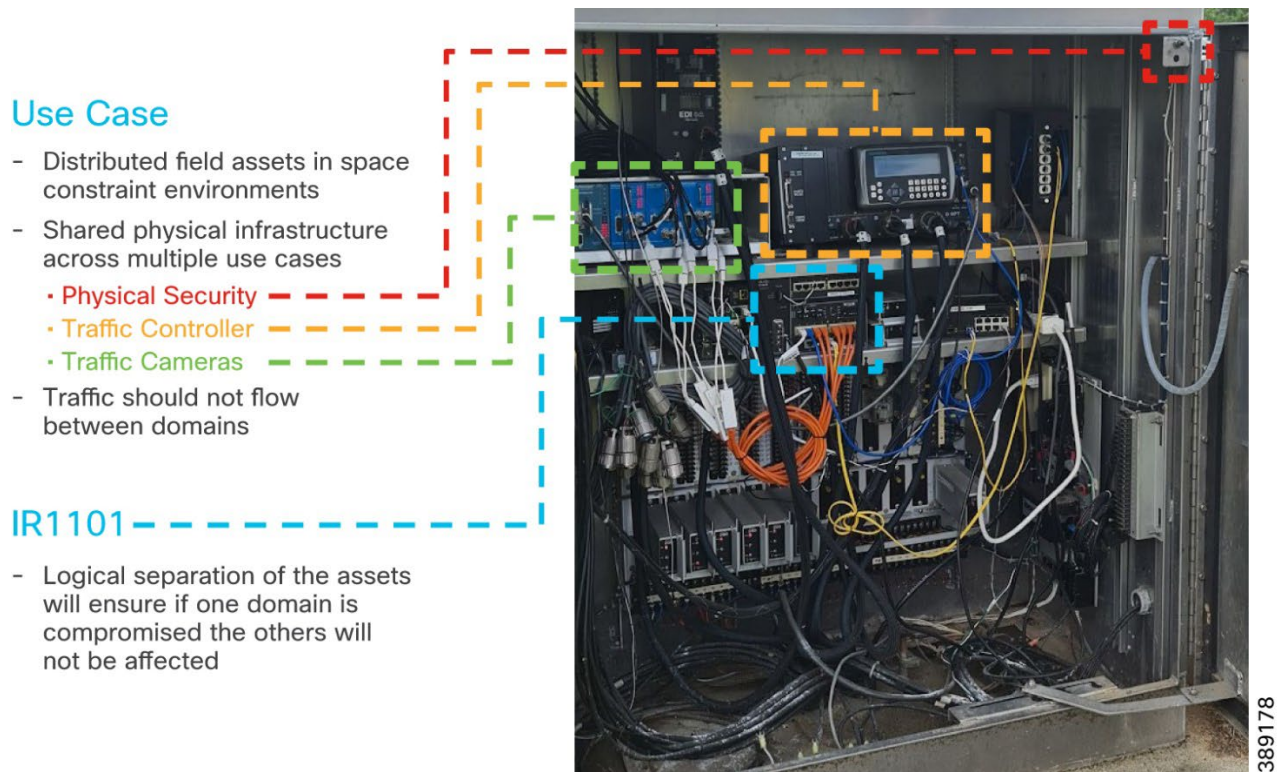
This command would create a custom protocol detector called OPCUA that will look for TCP packets that have a destination or source port of 4840.

Note: NBAR2 is not only used for creating firewall rules but also can be used for application aware routing and QoS in an SD-WAN deployment.

Network Segmentation

Beyond traditional packet filtering techniques, critical infrastructure can be further protected from noncritical assets that share the same physical infrastructure by segmenting traffic flows into separated virtual networks. [Virtual Routing and Forwarding \(VRF\)](#) allows a Cisco industrial router to run more than one routing table simultaneously. The routing tables are completely independent and fully segmented by default. For traffic originating in one domain to reach another domain, it must be explicitly routed through a firewall, reducing the possibility that an administrative error will lead to a wide-open network. The figure below shows an example of a traffic cabinet where this requirement is important. If one domain was to be compromised, the others should not be affected.

Figure 75 Example segmentation requirement in a traffic cabinet



In the context of Cisco Catalyst SD-WAN, the term VPN and VRF are used interchangeably. This is because Cisco Catalyst SD-WAN routers, such as the Cisco Industrial Routers, use VRFs for segmentation and network isolation and the Service VPN profile is used to configure them.

There are three different types of VPNs in Cisco Catalyst SD-WAN:

- **Transport VPN (VPN 0)** carries control traffic via the configured WAN transport interfaces. Initially, VPN 0 contains all device interfaces except for the management interface, and all interfaces are disabled. VPN 0 is synonymous with the global routing table.
- **Management VPN (VPN 512)** carries out-of-band network management traffic among the Cisco IOS XE Catalyst SD-WAN devices in the overlay network. The interface used for management traffic resides in VPN 512.
- **Service VPNs (VPN 1-511, 513-65530)** are used to carry device traffic across the overlay.

By default, no traffic will cross service VPNs, but if there is reason to do so, all traffic will be subjected to firewall rules.

Note: If using the Cisco Industrial Router in autonomous mode, it is still recommended to deploy a segmentation architecture. VRF-lite can still be used, but it may be more manageable to use a combination of VLANs and TrustSec.

Authentication, Authorization, and Accounting

Most field networks have no port authorization enabled today; that any device can be connected into an available Ethernet port and will be given network service. This is problematic because distributed networks such as intelligent transportation systems or the utility grid are uniquely exposed amongst critical infrastructure, with weak physical security and often being readily accessible to the public – and bad actors. Instead of having a “default open” posture we move to the best-practice of “default closed”; any connected devices by default have no network service and are given network service once the network can establish their identity as a trusted device.

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 802.1X, or if not supported, 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.

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.

Accounting involved tracking and recording the activities of users and devices on the network. This includes logging access time, the duration of resource usage, and the actions performed. Accounting helps in auditing and monitoring for compliance.

Port Security

The port security feature in IOS-XE enables administrators to restrict input to an interface by limiting and identifying MAC addresses of the endpoints allowed to access a port. When using port security, the port does not forward packets with source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the endpoint attached to that port is the only device that will have access to that port. For more information see [Port Security](#).

TrustSec

As detailed in a previous chapter, [TrustSec](#) defines policies using logical device groupings known as SGTs. The SGT is a single label indicating the privileges of the source within the entire network. SGTs are an important technology when deploying security across an expansive architecture. As will be discussed further in the [defense-in-depth](#) section, most use cases will traverse multiple policy enforcement points when traversing the network, and as a security architect, we need to co-ordinate both policy creation and log collection to understand the traffic

flows throughout our networks. SGTs provide a common identity that all enforcement points can use. Rather than relying on IP addresses, which may change depending on what part of the network you're on, and how many times you did NAT, an SGT provides business context to that device. I need to create policies for my employee managed device (for example, SGT 10) at a roadside cabinet, at the branch, at the data center. I need to create a policy for a traffic signal controller (for example, SGT 20) at the cabinet, across the WAN, in the control center, and across the data center. Regardless of the IP address, policy can be consistent across the architecture, and it becomes much easier to correlate log information across the entire network.

Figure 76 Importance of TrustSec



TrustSec has three main functions; classification, propagation, enforcement. The level of support for each function with the Cisco Industrial Router platform can be found below.

TrustSec Classification

Subnet to SGT mapping binds an SGT to all host addresses of a specified subnet. Once this mapping is implemented, Cisco TrustSec imposes the SGT on any incoming packet that has a source IP address which belongs to the specified subnet. To create a subnet to SGT map (or an IP to SGT map), use the following configuration:

```
cts role-based sgt-map <ip-address> sgt <number>
```

Interface to SGT mapping binds all traffic on a Layer 3 ingress interface to an SGT. Once this mapping is implemented, Cisco TrustSec can use the SGT to segregate traffic from various logical Layer 3 ingress interfaces; routed Ethernet interfaces, 802.1Q sub interfaces, tunnels. To create an interface to SGT map, use the following configuration:


```
interface <type slot/port>
cts role-based sgt-map sgt <number>
```

Subnet to SGT, or Interface to SGT are examples of static mappings that don't consider device context. Dynamic mapping can be implemented during authentication to ISE. When doing an 802.1X or MAB authentication request on a Cisco Industrial Router, Cisco ISE can apply an SGT dynamically based on device profiling. Design guidance will be discussed in a [later section](#).

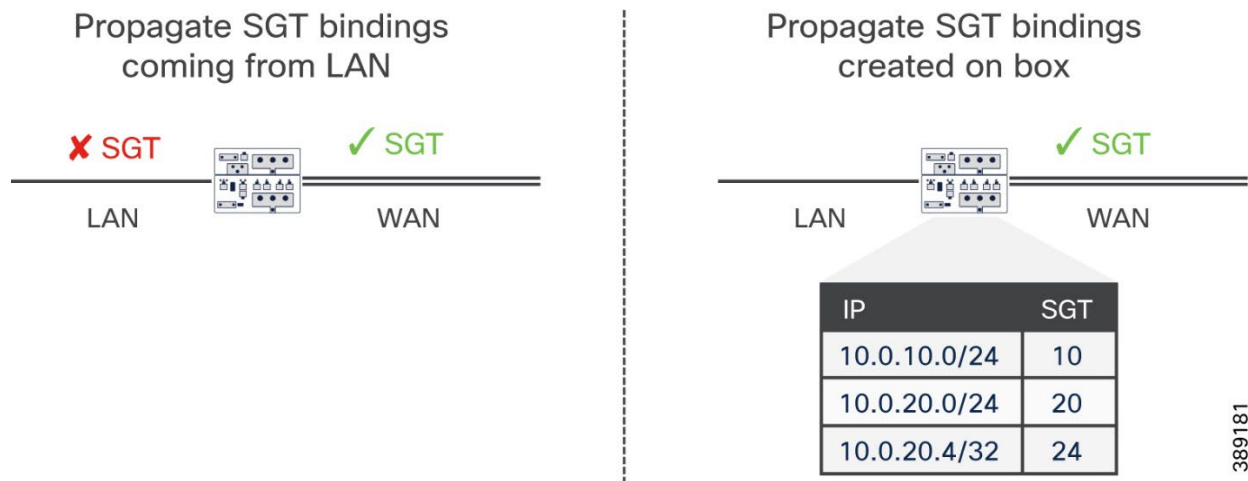
TrustSec Propagation

Inline tagging, when the SGT is embedded in the Ethernet frame header, is supported on the layer 3 interfaces of all Cisco Industrial Routers.

Figure 77 Inline tagging across IPSec tunnels

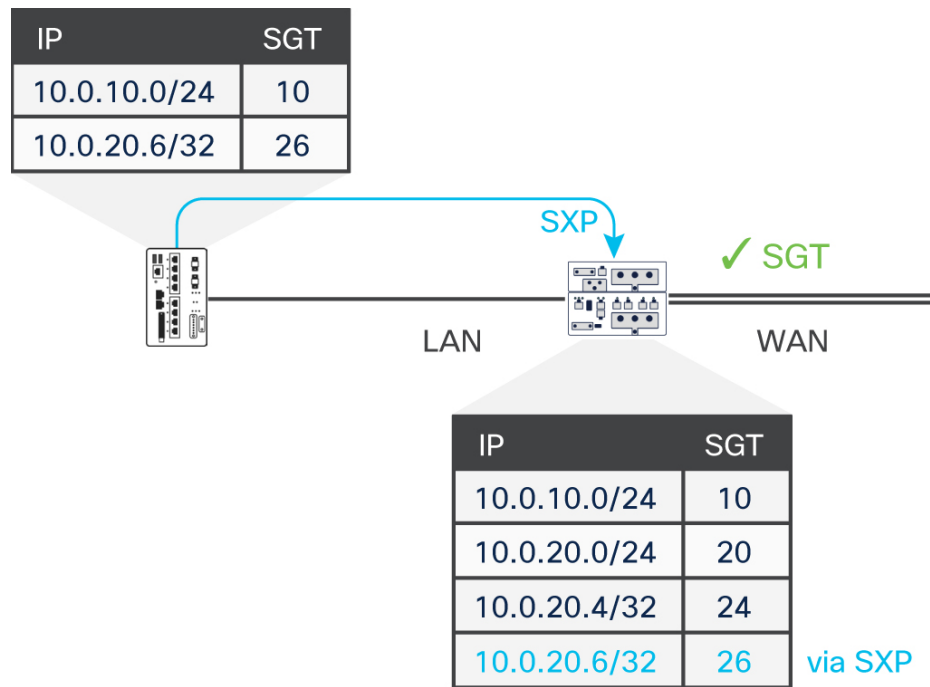


Inline tagging provides the capability to tag packets locally at one site, and carry that across the WAN to be enforced at another. By extending inline tagging on the layer 3 ports, TrustSec can be distributed across both routers and switches to extend segmentation capabilities beyond the reach of the router. However, inline tagging is not supported on layer 2 interfaces, for example, the LAN ports of an IR1101.

Figure 78 Inline tagging limitations

389181

To overcome this, all Cisco Industrial Routers support SXP (both speaker and listener), which is the protocol used to exchanged IP to SGT bindings across network links that do not support inline tagging.

Figure 79 Using SXP to propagate IP to SGT bindings

389182

After SXP has populated the table, SGT propagation across the WAN can be done as per previous section. This will also enable SGTs that have come from the WAN to be transferred across to enforcement points in the LAN.

TrustSec Enforcement

An SGT can be used in source or destination data prefixes in a firewall policy across all Cisco Industrial Routers. Using the SGT instead of other attributes allows administrators to better

understand policies and logs in an end-to-end security system, as while attributes such IP address and VLAN may change throughout the life of a packet, the SGT should remain the same.

Device Support

The table below represents each feature per platform in the Cisco Industrial Routers.

Table 3 TrustSec support in Industrial Router portfolio

Device	Classification		Propagation		Enforcement	
	Static	Dynamic	Inline	SXP	ZBFW	SGACL
IR1101	IP to SGT, Subnet to SGT, switch virtual interface (SVI) to SGT	Yes	Tunnel interface, L3 ports	Speaker, Listener	Yes	No
IR1800	IP to SGT, Subnet to SGT, SVI to SGT	Yes	Tunnel interface, L3 ports	Speaker, Listener	Yes	No
IR8340	IP to SGT, Subnet to SGT, SVI to SGT	Yes	Tunnel interface, L3 ports	Speaker, Listener	Yes	No*

* This is under investigation and may be added in a later release.

Denial of Service Protection

Denial of service (DoS) attacks are malicious attempts to disrupt the normal functioning of a targeted server or network by overwhelming it with a flood of illegitimate requests or traffic. The goal is to make the target unavailable to legitimate data flows.

An **ICMP flood attack** is a type of DoS attack where the attacker overwhelms the target system with a high volume of Internet Control Message Protocol (ICMP) Echo Request packets, commonly known as “ping” packets. The objective is to exhaust the network resources, including the CPU and memory, causing it to become slow or completely unavailable to process legitimate packets.

A **Smurf attack** is an example of a distributed denial of service (DDoS) that also uses ICMP, but in a different way. The attacker sends ICMP Echo Request packets with a spoofed source IP address, set to the IP address of the target victim. The ICMP packets are sent to the broadcast network, but since the source IP has been spoofed, the target now receives a large volume of ICMP Echo Replies from multiple devices, is overwhelmed by the flood of traffic, leading to network congestion.

To protect against DoS attacks it is recommended to protect both the data plane and the control plane packets.

Quality of Service

Quality of Service (QoS) refers to a set of technologies and techniques used to manage and prioritise network traffic to ensure the performance of critical applications, minimize latency, reduce packet loss, and provide a predictable and reliable network experience. QoS is particularly important in environments where bandwidth is limited and where critical infrastructure applications have sensitivities to latency and dropped packets.

Traditional QoS does have its limitations because it can't predict the changing bandwidth on the link. Adaptive QoS, the shapers at the edge can adapt to the available WAN bandwidth, including across long term evolution (LTE).

For more information on adaptive QoS in Cisco SD-WAN see [Adaptive QoS](#).

Control Plane Policing

Control Plane Policing (CoPP) improves security on the Cisco Industrial Router by protecting the CPU from unnecessary traffic and DoS attacks. It can also protect control traffic and management traffic from traffic drops caused by high volumes of other, lower priority traffic.

A router is typically segmented into three planes of operation, each with its own objective:

- The data plane, to forward data packets.
- The control plane, to route data correctly.
- The management plane, to manage network elements.

CoPP can be used to protect most of the CPU-bound traffic and ensure routing stability, reachability, and packet delivery. Most importantly, CoPP can be used to protect the CPU from a DoS attack.

CoPP uses the modular QoS command-line interface (MQC) and CPU queues to achieve these objectives. Different types of control plane traffic are grouped together based on certain criteria, and assigned to a CPU queue. CPU queues can be managed by configuring dedicated policers in hardware. For example, the policer rate can be modified for certain CPU queues (traffic-type), or the policer can be disabled for a certain type of traffic.

Although the policers are configured in hardware, CoPP does not affect CPU performance or the performance of the data plane. But since it limits the number of packets going to CPU, the CPU load is controlled. This means that services waiting for packets from hardware may see a more controlled rate of incoming packets (the rate being user-configurable).

For more information on CoPP in IOS-XE see [Configuring Control Plane Policing](#).

Application Hosting with IOx

Cisco IOx is an application enablement platform that provides Cisco Industrial Routers to host applications at the network edge. The platform is designed to meet the growing demand for edge computing, allowing organizations to deploy applications closer to where data is generated, and is the mechanism for hosting application such as [Cisco Cyber Vision](#) and [Cisco Secure Equipment Access](#) which have been discussed in this design guide already.

NGFW Add-On

To expand upon the stateful firewall capabilities within the zone based firewall, the Cisco IR1835 and IR8340, due to the expanded memory, can host an NGFW add-on within the IOx infrastructure for advanced threat protection embedded in the router. The NGFW add-on, often referred to as Unified Threat Defense (UTD), brings the following capabilities:


- Snort IDS/IPS detects and blocks malicious activities by analyzing network traffic patterns and identifying known threats
- Advanced Malware Protection (AMP) offers file analysis and sandboxing to detect, block and remediate malware across the network
- URL filtering to control access to websites based on categories, reputation, and custom policies to prevent exposure to malicious sites

Note: File detection and sandboxing relies on cloud connectivity.

Snort IDS/IPS

Snort is the network intrusion detection and prevention system used by Cisco across many products within its security portfolio. Snort has been covered extensively in this guide in a [previous section](#) when detailing design guidance on the Cisco Secure Firewall. All of the same features exist on the Industrial Routers including the SCADA pre-processors for creating OT protocol rule violations. An example Snort configuration for Cisco Industrial Routers is in Appendix C.

Figure 80 SCADA preprocessors in Snort 3



Note: SnortML is unique to the Cisco Secure Firewall and is not yet available on the Cisco Industrial Router.

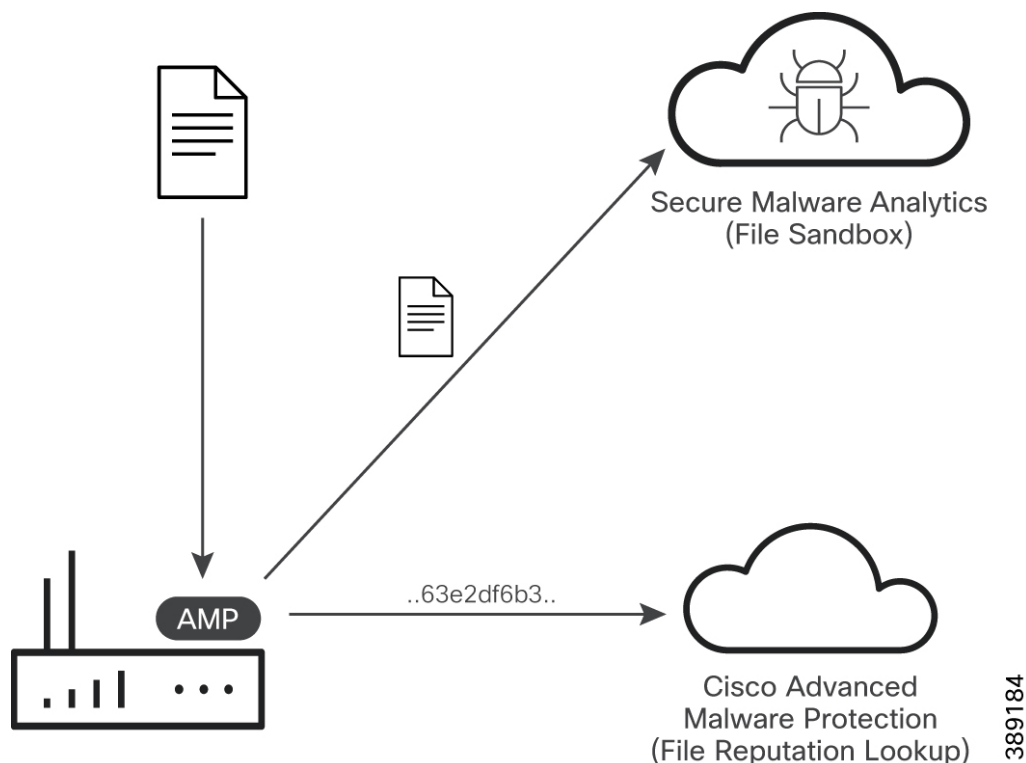
URL Filtering

Use cases such as predictive maintenance or IoT applications often require connections to cloud resources, increasing the attack surface. To enable such innovation, URL filtering in the Cisco industrial routers allows control access to trusted cloud resources by configuring domain-based or URL-based policies. Although we recommend that access to cloud and internet resources be disabled by default, and that you explicitly allow only trusted domains, security administrators have peace of mind that the network is protected by reputation-based filtering. Each URL has a web reputation score associated with it to help ensure that users or applications are not communicating with high-risk parts of the internet.

Advanced Malware Protection

Malware is one of the most common cyber threats. Detecting and removing malicious files before they enter your network is key to prevent breaches. Cisco Advanced Malware Protection (AMP) integrated into Cisco industrial routers equips the platform to provide protection and visibility from malware. Before letting a file enter the network, your Cisco industrial router generates a 256-bit Secure Hash Algorithm (SHA256) signature and compares it against a database curated by Cisco Talos, the largest collection of file reputation intelligence in the industry.

Figure 81 File inspection by Cisco AMP and Secure Malware Analytics clouds



Files with an unknown disposition can be sent to the Cisco Secure Malware Analytics cloud for further analysis within a sandbox. During detonation, the sandbox captures artifacts and observes the behavior of the file, then gives the file an overall score of abnormal behaviors. Based on the observations and score, Secure Malware Analytics will define the file as clean or malicious so your Cisco industrial router will let it pass or block it.

Supported Platforms

All Cisco Catalyst industrial routers have security built in. Cisco IOS® XE, the software that powers all Cisco networking infrastructure, provides stateful packet inspection, application visibility and control, VPN, segmentation, DoS mitigation, and FQDN matching.

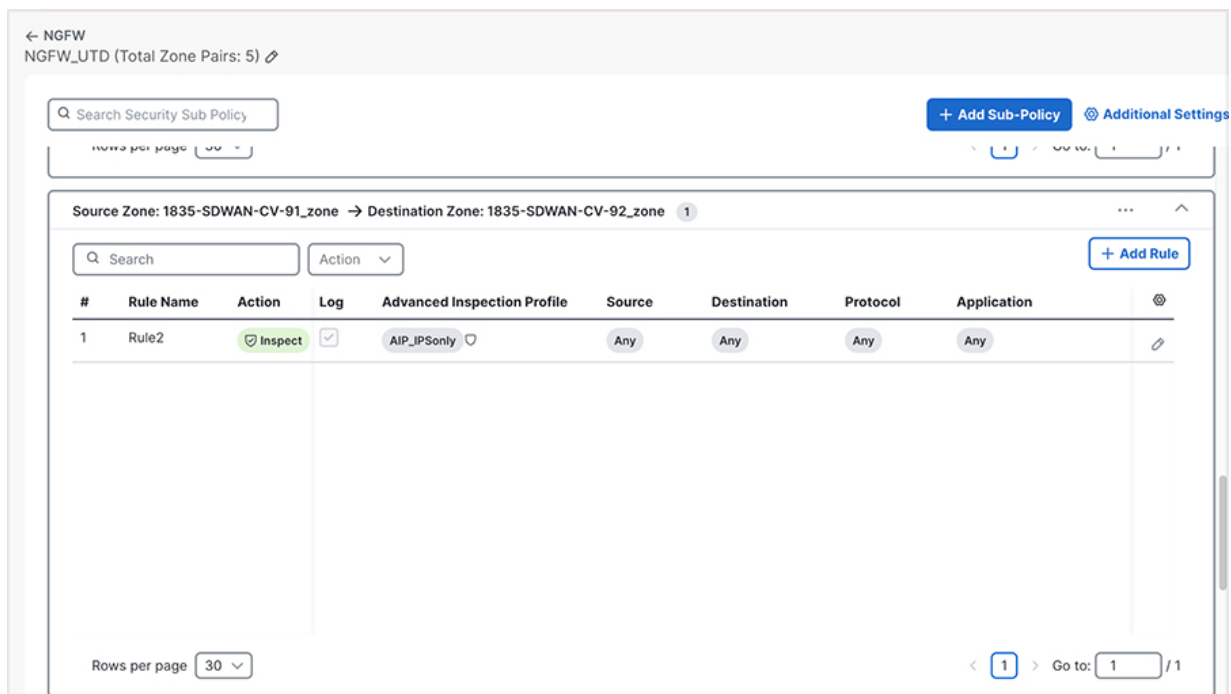
The remaining features come from the NGFW add-on that can be deployed in devices with 8 GB of memory. The NGFW add-on for industrial routers provides Snort IDS/IPS, reputation-based URL filtering, and malware protection.

Table 4 Security capabilities per routing platform

Features	IR1101, IR18xx	IR1835	IR8340
Stateful Packet Inspection	✓	✓	✓
Application Visibility & Control	✓	✓	✓
VPN	✓	✓	✓
Segmentation	✓	✓	✓
DoS Protection (QoS, CoPP, etc.)	✓	✓	✓
AAA	✓	✓	✓
Port Security	✓	✓	✓
TrustSec	✓ *	✓ *	✓
IDS/IPS	✗	✓	✓
Malware Protection	✗	✓	✓
File Sandboxing	✗	✓	✓
FQDN Filtering	✓	✓	✓
Reputation and Category Web Filtering	✗	✓	✓
TLS Decryption	✗	✗	✓

*TrustSec enforcement in the ZBFW only. SGACL is not supported

Most features are available regardless of the management platform used. The only exception is the NGFW add-on. As of version 20.16, the firewall policy manager in Cisco Catalyst SD-WAN Manager has been redesigned to be consistent with the user experience of other firewall managers in the market.

Figure 82 Firewall policy manager in Cisco Catalyst SD-WAN Manager

Security administrators create a set of policies between zone pairs on the Cisco Industrial Router, and traffic matching options like objects, subnets, applications and protocols are available to create allow, deny or inspect rules. When using the inspect action, any traffic that matches this policy will be sent the NGFW add-on for advanced threat inspection. If using Catalyst SD-WAN Manager, all of the features of the NGFW add-on will be available. However, if deploying the Cisco Industrial Router without it, there are limitations:

- Snort IDS/IPS can still be used, but there is no support for loading a custom rules file.
- Advanced Malware Protection is not supported
- Web filtering is fully supported

TLS Decryption

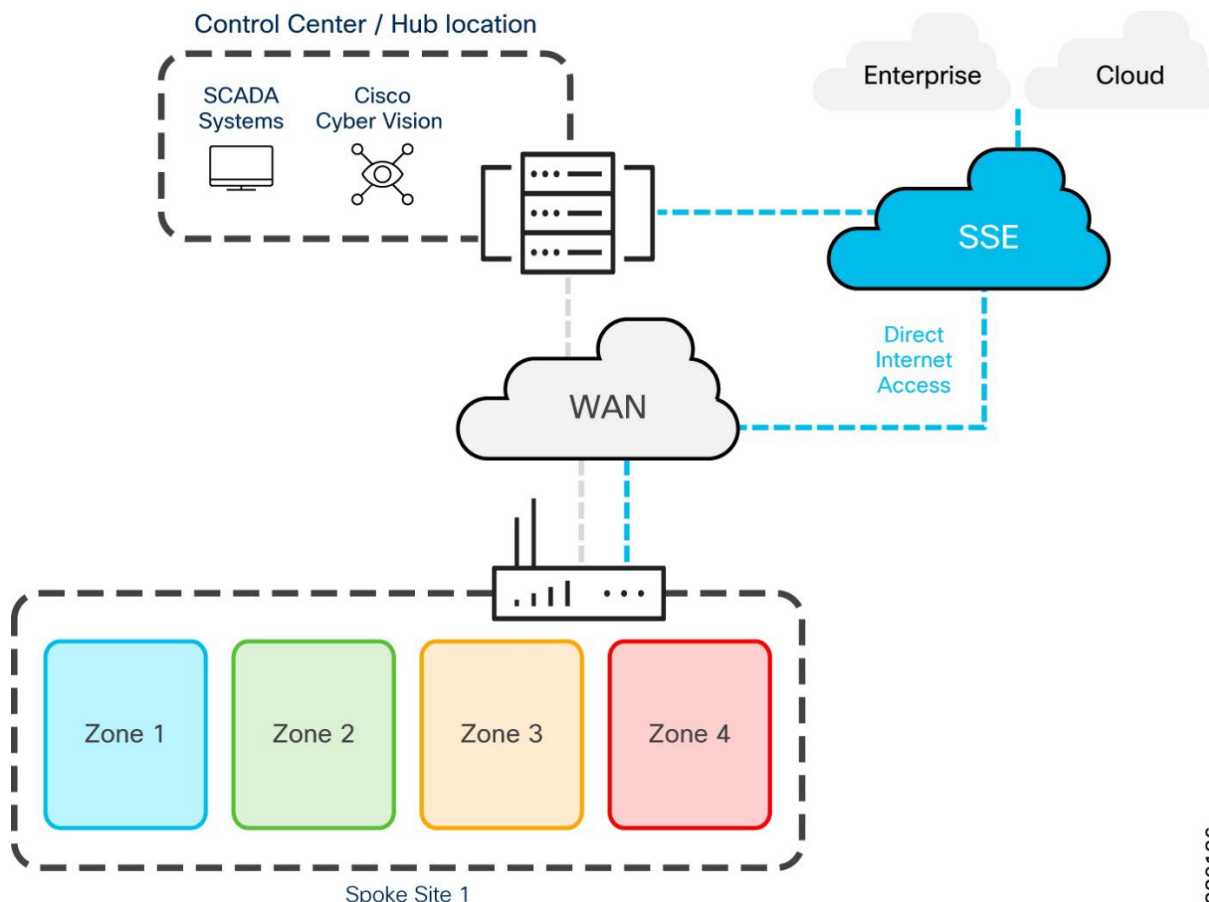
The majority of Internet traffic is encrypted, which may lead to malware remaining hidden and a lack of control over security. The TLS proxy in a Cisco IR8340 (decryption not supported on other devices) acts as a man-in-the-middle to decrypt encrypted TLS traffic traveling across WAN and to inspect the contents for malicious activity. The data is re-encrypted post inspection before being sent to its final destination.

Secure Access Service Edge (SASE)

Whether it is rail-side deployments spanning hundreds of miles or traffic intersections that are distributed across a whole city, critical infrastructure is often widely distributed. When deploying Cisco industrial routers, network architects have a choice of where advanced security policies will be deployed. A common deployment model is to centralize the most advanced policies in the network, alleviating the burden that may exist on edge nodes. Cisco industrial routers can

leverage Cisco Secure Access or any third-party Security Service Edge (SSE) via IPsec tunnels to centralize policy enforcement across sites or toward the cloud.

Figure 83 SASE for distributed field networks



With Cisco Secure Access, network administrators enable segmentation and prioritization to the most critical traffic on the network, while security administrators maintain granular control of data that comes into and out of each remote site with a single set of policies, so that only known, trusted traffic flows throughout the infrastructure.

Plug and Play

According to a Gartner *State of the Firewall* report “99% of firewall breaches will be caused by misconfigurations, not firewall flaws”. While network firewalls have a plethora of capabilities at their disposal to deal with cyber threats, they first must be configured. Out of the box, most firewalls are deployed with a default ‘any any’ rule, which allows all traffic to pass through without interruption. What often happens is an IT team has a limited time window to complete a network upgrade or install, and the priority is uptime. So firewalls are deployed in their most basic form to complete a job, and configuring policies is seen as a day 2 operation.

Plug and play (PnP) is a technology that allows devices to automatically be provisioned on the network without requiring manual intervention from the user. Not only does the concept simplify the process of adding new gateways to the network, but it does so in a way that pulls the latest configuration designed by the networking and security teams.

By leveraging Cisco Catalyst SD-WAN, when a Cisco Industrial Router is provisioned, or perhaps replaced in a field network, the device configuration is pulled into the router which will contain all the security policies it needs to secure operations from day one. All of the capabilities listed above such as denial of service protection, port security and firewall rules do not need to be configured as an afterthought, reducing the risk of misconfiguration which could lead to a breach.

Cisco Industrial Router Design Guidance

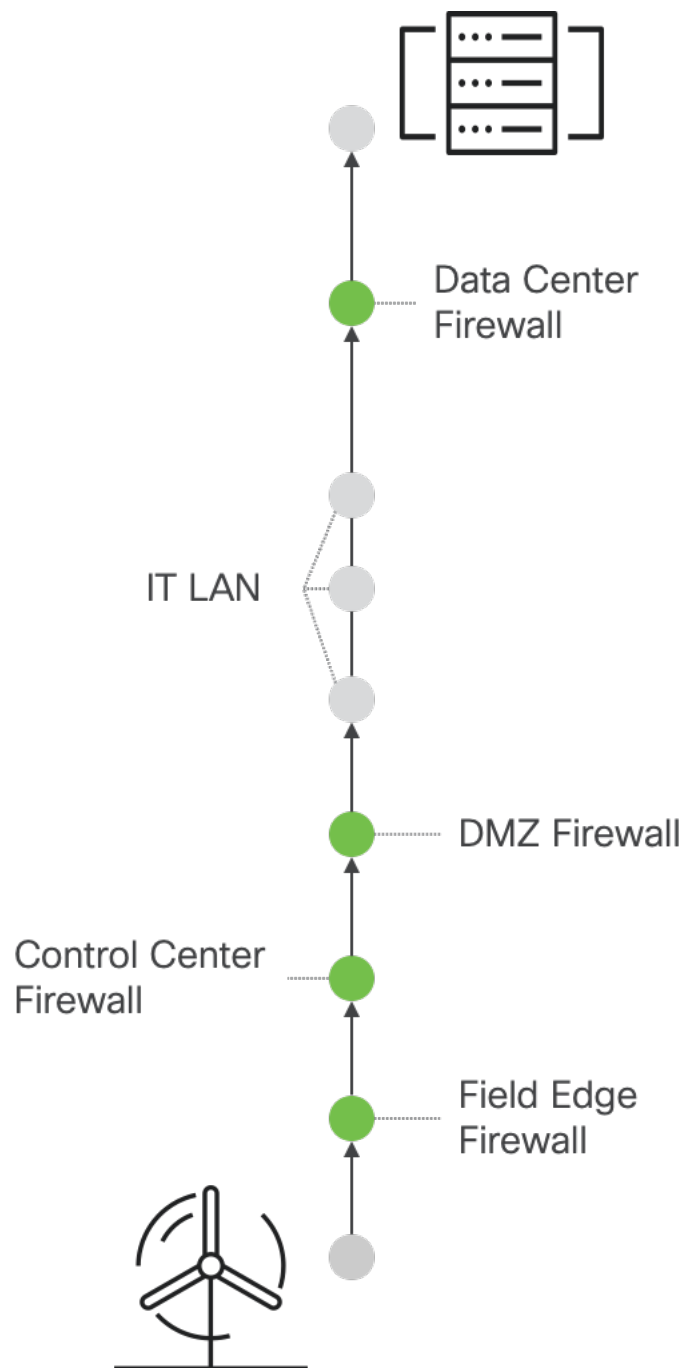
Just because a capability can be deployed at the edge, it does not necessarily mean that it should. The purpose of this section is to guide security architects on when, where and why the capabilities available on the Cisco Industrial Router should be utilised.

Understanding defense-in-depth

NIST special publication 800-82, *Guide to OT Security*, describes a defense-in-depth strategy as a multifaceted strategy to establish variable barriers across multiple layers and dimensions of an organization. It is considered a best practice across numerous standards and regulatory frameworks, as the basic concept is to prevent single points of failure in cybersecurity defences and to assume no single origin of threats.

What defense-in-depth is not, is implementing the same set of security controls across multiple parts of the network. For example, a firewall is an important tool used in cybersecurity architectures to control traffic across boundary points of the network. However, defense-in-depth does not mean multiple layers of firewalling. If an attacker can bypass one firewall, then maybe they can also bypass the next.

Figure 84 A typical data flow if using firewalls for defense-in-depth architectures



Organizations should implement appropriate technology for the use cases in which they need to protect. When implementing cyber security at the field edge, administrators should ask themselves;

- Is this traffic going direct to the Internet, or should those policies instead be deployed at the Internet breakout?
- Is TLS decryption required at every hop in the network, or should computationally expensive actions be reserved at strategic points in the network?

- Is IPS required across all boundaries, subjecting the traffic to the same set of signatures that have already been bypassed?

This design guide will not have the scope to cover all possible use cases that industry will face, but rather act as a guiding set of principles to build a cyber security architecture for protecting critical infrastructure. When designing a defense in depth architecture, consider the following:

- Not all security controls need to be enabled at every hop in the network. The router deployed at the edge of the field network may seem like the most important box in the cyber security strategy, but there are often multiple enforcement points between field assets and their intended destinations. Use core security practices like AAA, port security, segmentation and basic firewalling at the edge to ensure only trusted devices cross the boundary using trusted ports, and supplement that with advanced inspection at the hub location. It is much easier to manage, rules are likely kept up to date and as technology advances it is easier to swap out central firewalls with the latest and greatest security technology than to roll out hardware to thousands of micro-sites.
- Try not to duplicate the same policies across multiple hops. This is easier said than done, as firewalls will often have conflicting rules, but reducing the number of overlapping rules will help maintain policies in the long term. Additionally, computationally expensive actions like TLS decryption only need to be done once. It is important this is done at the Internet boundary, so make that a mandatory capability in your architecture. However, most OT devices don't communicate with the Internet, so before deploying TLS decryption at the edge, ask yourself if it can be done higher up the chain where a more powerful box can be used. Rugged devices are smaller and lack fans, so making TLS decryption a priority at the edge is an expensive design decision.
- Connect logs to a security information and event management (SIEM) product. A SIEM is a solution that collects, analyzes, and responds to security data from various data sources. It is important that logs are correlated across all the policy enforcement points that critical infrastructure is subject to. More information can be found in [Cross-Domain Detection, Investigation and Response](#).

Deny by default firewall configuration

As mentioned [earlier in this document](#), 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. In its most basic form, firewalls allow security administrators to restrict traffic based on UDP and TCP ports. But innovations over the last decade have added additional features such as intrusion detection, encrypted traffic analysis, file analysis, web filtering, among other capabilities needed to protect organizations. However, it is important to understand why these features are necessary.

Implementing firewall policies for enterprise networks is difficult. Users access the Internet every day, and the destinations are unknown. Security administrators implement DNS policies to try and stop users going to malicious websites. File analysis policies are in place to make sure that the file downloaded from their personal cloud storage does not contain malware. Decryption is performed on web traffic to ensure employees can safely browse the web without being

restricted on the sites they can visit. Essentially, the default rule of the firewall is to allow all traffic to leave the enterprise, but with exceptions and threat intelligence providing protections.

The industrial edge is not where these innovations are required. Building on the principles of defense-in-depth, it is advised to apply the appropriate policy at different levels of the network. The most important design decision to make at the industrial edge is to follow the principles of least privilege. Least privilege means that users and assets that cross the network boundary will only be able to do so if it is necessary. This concept is a key philosophy in NERC CIP-005-7, where all inbound and outbound access permissions across the critical substation boundary must be documented with a reason for granting access and all other access is denied by default.

When implementing firewall policies at the industrial edge, consider the following:

- Identity which devices have permission to cross the boundary, deny the rest. Advanced security capabilities are not required to make these decisions, a simple deny rule will stop any traffic coming from devices who have not been explicitly given permission to cross.
- For devices that can cross the firewall, limit the ports and applications that can be used. Video cameras need to stream to a server? Open the appropriate ports, block the rest so the web browser cannot be access. Users need to SSH to a server? Make sure only trusted devices can initiate the connection. A SCADA system needs to read data from a controller using Modbus? Make sure only that trusted server can communicate over Modbus, and use application policies to make it read only. By only opening the specific ports that are needed between known devices will make it very difficult for an attacker to compromise the network as they are limited to what they can do with legitimate traffic flows.

When to use Snort

Now that the basics have been covered, when should an [NGFW](#) be deployed at the industrial edge? Revisiting NERC CIP-005-7, there is a requirement for substation networks to have one or more methods for detecting known or suspected malicious communications for both inbound and outcoming communications. Additionally, the [TSA security directives](#) also ask for continuous monitoring and detection policies that are designed to prevent, detect, and respond to cybersecurity threats affecting critical cyber systems.

IDS/IPS plays an important in role keeping the network safe from malicious attacks. However, as packets will traverse multiple network hops before reaching the intended destination, traffic should not be subject to inline detection multiple times, as it causes unnecessary latency to the traffic. Additionally, when deploying security policies across the network, the more distributed the policies are, the harder it is to keep them up to date. If an administrator needs to make a change to policy, they must ensure that the change has been successfully deployed at every enforcement point in the network. When deploying Snort, or any IDS/IPS system consider the following:

- If Snort is only required for traffic entering/leaving the remote site, consider doing the inspection in a central location. The firewall rules deployed at the industrial edge will ensure that only known traffic will cross WAN (north/south), and the Snort inspection at the hub site will give an additional layer of protection to ensure there are no malicious packets hidden within legitimate traffic.

- If intrusion detection is needed east/west (i.e. between devices at the site), then consider deploying Snort on either the IR1835 or the IR8340, configuring IDS/IPS policies in addition to the firewall rules deployed at the site.
- Use custom Snort signatures to block specific commands on OT protocols. For example, allowing DNP3 in the firewall rules, but denying all function codes other than the read commands in the IPS rules.
- Alternatively (or additionally), deploy a network visibility tool such as Cisco Cyber Vision to do out of band inspection. Cisco Cyber Vision also has the capability to run Snort as a detection engine and will give you much more visibility into the industrial network than a traditional IDS/IPS system will. If IPS is not a requirement for east/west traffic, it is recommended to deploy Cyber Vision instead of Snort on the industrial router.

Cloud Connectivity

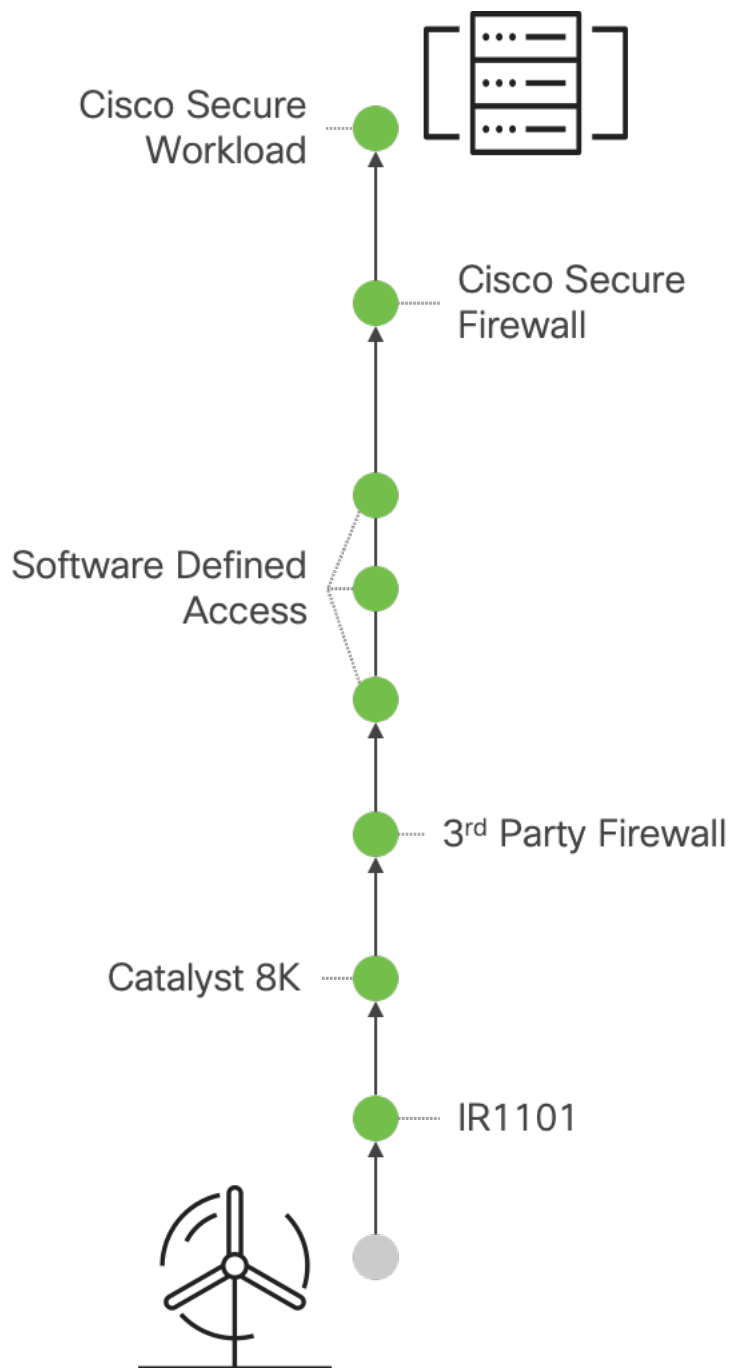
With the rise in artificial intelligence and the goal to digitise a lot of our operations, cloud connectivity is becoming more accepted in industrial networks. Saying that, the demarcation point between field assets and the internet should not be the industrial router. Following the theme from previous guidance, the organization should have a defined Internet boundary, a central choke point where policy is defined and kept up to date. Generally, the IT/OT boundary is at the control center, and for Internet connectivity, the control center will leverage the IT network. When connecting field assets to cloud resources, consider the following:

- While Cisco Industrial Routers have the capability to do URL filtering and scan malicious files, they should not be the devices used for Internet connectivity. It is recommended that Internet traffic is backhauled to a central location and all traffic is subject to policies that are maintained by the security team.
- If direct Internet access (DIA) is required, consider a security service edge (SSE) such as Cisco Secure Access to protect users and devices when accessing cloud resources. This model may already be used by the enterprise, so rather than backhauling traffic to a hub only to be subject to policies from the SSE, an IPsec tunnel can be created from the Cisco Industrial Router.
- For minor use cases, such as providing a guest network, and an SSE does not exist, using the inbuilt web filtering is advised. In this case, segmentation is highly recommended.

End-to-end segmentation with TrustSec

The TrustSec section in this chapter primarily focused on TrustSec capabilities within the Cisco Industrial Router. However, there is a larger design focus on TrustSec when discussing [network access control with Cisco Identity Services Engine](#) that have not been replicated here. TrustSec is an important technology when implementing a defense-in-depth security architecture. A previously discussed misconception of defense-in-depth is using the same technology stack in multiple parts of the network. Figure 85 depicts a more realistic view of the security stack. Taking a flow originating from the OT network destined for the data center:

- An OT device must pass through a router/firewall/gateway at the field network edge. A policy decision must be made here to determine if that traffic is permitted.
- The traffic traverses the WAN, whether that is a private circuit or an overlay over the Internet and terminates at a hub router. Another policy decision point.
- Typically, the router hosted at the hub location is not the main firewall, and all traffic traversing the hub boundary will be subject to firewall enforcement before reaching the IT network.
- Modern IT networks have implemented some level of network access control and segmentation. In this diagram Cisco software defined access (SDA) is depicted, which comes with its own policy engine.
- Upon reaching the data center, there is typically a firewall protecting it.
- Data center security is also a hot topic, with products such as Cisco Secure Workload or Cisco Hypershield potentially being used here.

Figure 85 An example policy enforcement stack between IT and OT networks

For a single use case, traffic originating from the field and reaching the data center has crossed six policy decision points in the network. As a security administrator this can be difficult to keep track of. To alleviate some of the complexity, use a common identity. By using SGTs, all policy decision points can use the same identifier when applying policy to traffic. This helps correlate logs and troubleshoot policies end to end.

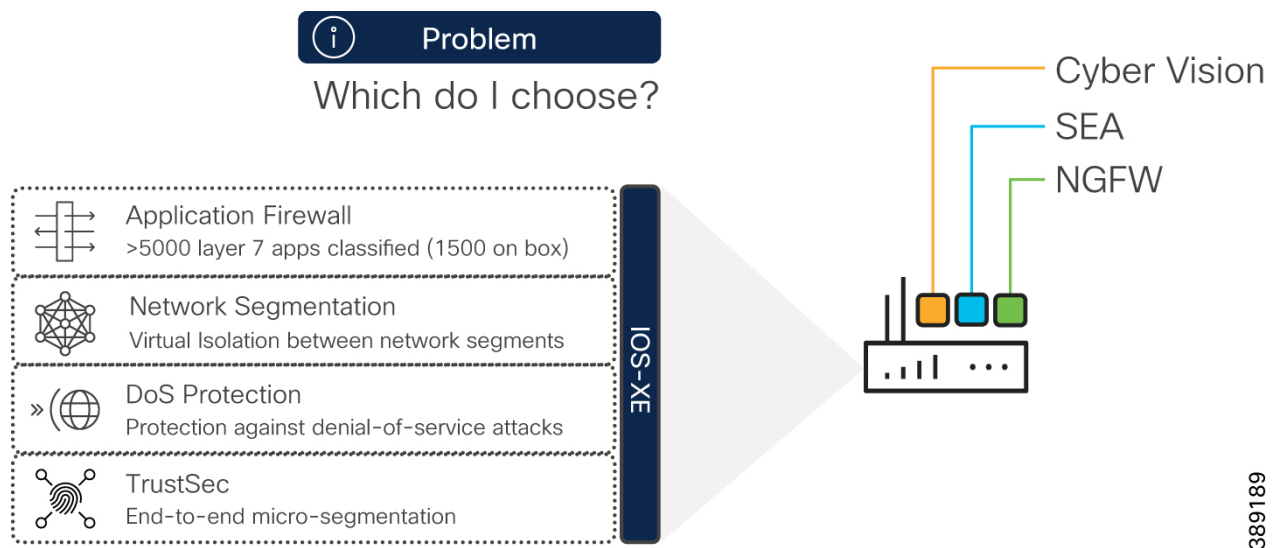
In addition to the design guidance already discussed in [ISE / SGT Design Considerations](#), consider the following at the field network:

- Build a dedicated management network between the network infrastructure and ISE. In an SD-WAN deployment, this would be a dedicated service VPN. See [ISE Components / Personas](#) for the communication details between network infrastructure and ISE.
- Use static classification as a starting point for segmentation. Not all endpoints connected to the network need to be dynamically profiled to receive policy. If there is a subnet dedicated to a traffic controllers, statically assign the traffic controller SGT to that subnet.
- Combine TrustSec with port security. Building on the static classification workflow, if there is a concern of rogue devices connecting to the network, such as a malicious device joining the traffic controller subnet, use port security to limit the connections to only known devices. There are two approaches here. The first is to only allow trusted MAC addresses to join the network. In this approach, a MAC address must be registered to ISE before being authorized on the network. The second approach is to limit the port to a single MAC address. This is a more dynamic approach, where the first MAC address to connect to the port is now the only MAC address that can use that port until a network administrator intervenes.
- While the previous step reduces the risk, the next design choice may be to dynamically profile the endpoint to verify its trustworthiness. There is an advanced scenario where an attacker could mimic the MAC address of a trusted device. To overcome this, use endpoint profiling in ISE. When using MAB or 802.1x to authenticate to the network, ISE can assign an SGT based on the endpoint profile. This does not need to be a one-time operation. Every time ISE learns something new about the device, a decision can be made to re-authenticate the asset. This is known as change of authorization (CoA). Use Cyber Vision to constantly monitor the network and feed device properties to ISE. If a device property results in an endpoint profile that does not belong to the edge network, ISE can automatically remove it from the network. See [ISE / SGT Design Considerations](#) for more information on the integration between Cyber Vision and ISE.
- Use SGTs in the firewall rules. There may be a misconception that all TrustSec policies are generated from ISE. While ISE does push policy to the network infrastructure, it does so through the switches. ISE creates SGACLs, which are stateless packet filters for segmentation in switch networks. Stateful packet inspection, and the use of NGFW control options are out of scope of ISE. The policy engine for the routers is the firewall, and just like how an SGT can be used in the Cisco Secure Firewall, or other third party firewalls, the SGT can, and should, be used as an identifier for policy creation.

Multiple IOx applications in a single deployment

The design guidance found in this document has referenced multiple applications that can be hosted in IOx – Cyber Vision, Secure Equipment Access (to be detailed in [this guide](#)) and the NGFW add-on for Cisco Industrial Routers. However, a single Cisco Industrial Router can only host one application at a time. So which one should be chosen? The purpose of this section is to help make a decision on which application should be hosted depending on the security need. It is important to note, a single device can only be hosted on one device, but security can be spread across an architecture.

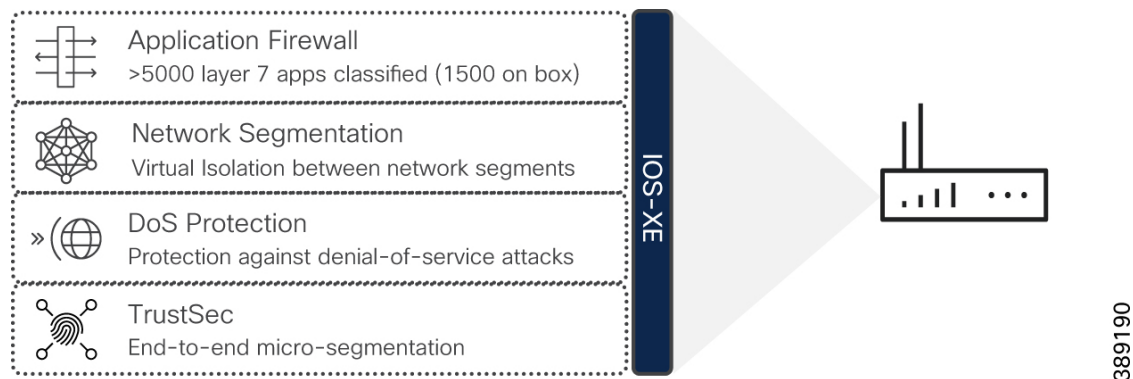
Figure 86: A single Cisco Industrial Router can only host one application



Problem statement: I need a firewall to protect my critical assets.

The first question that should be asked is “what is a firewall?” Is the zone-based firewall enough to meet the needs for the use case? Any Cisco Industrial Router can be deployed with a deny by default firewall configuration and specific ports can be opened between known IP addresses. TrustSec could also be used, where endpoints are authenticated and authorized depending on their attributes and an SGT is assigned to be used across various enforcement points in the network. The details on what is required by the firewall is important, as in this scenario, no application may be necessary.

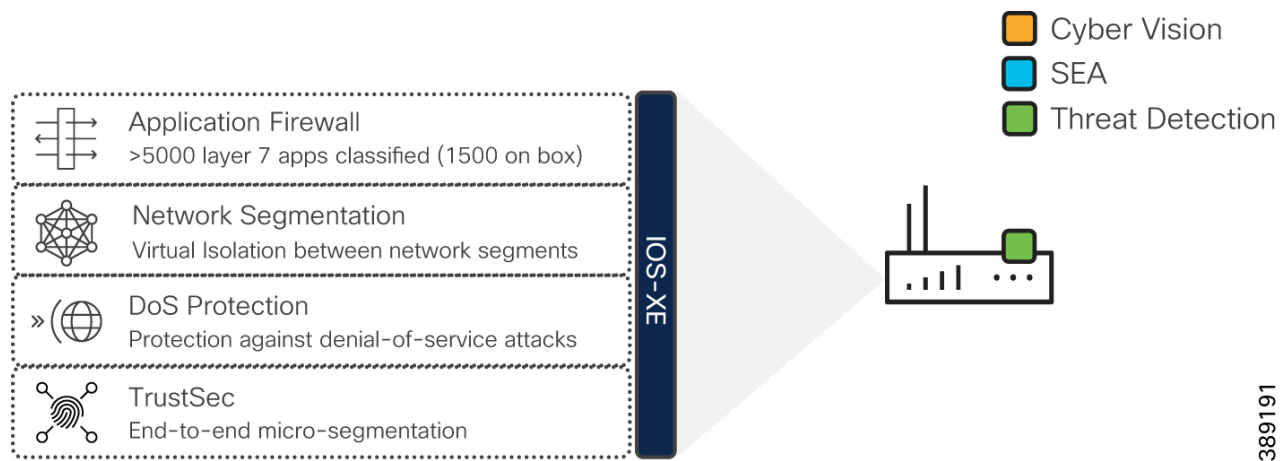
Figure 87: A highlight of the firewall capabilities in IOS-XE



Problem statement: I need a firewall to protect my critical assets, and I need IDS/IPS capabilities.

In this scenario the NGFW add-on for either the IR1835 or the IR8340 can be used. If only one application is required, the solution can be quite simple. Nevertheless, follow the design description in [this guide](#) to understand if IDS/IPS is necessary at the edge. It is important to note that IOS-XE contains the firewall for Cisco Industrial Routers, the NGFW add-on is just an advanced threat detection module.

Figure 88: Cisco Industrial Router with the NGFW add-on




389191

Problem statement: I need a firewall to protect my critical assets, I need IDS/IPS capabilities, but I also want Cyber Vision for visibility.

In this scenario, because multiple applications are required, it can be handled in two ways. The first method is to use the principles of defense-in-depth and divide responsibility across the architecture. Ask yourself the question; is IPS required for east/west traffic (for example, the traffic located at the site). If the answer is no, and IPS is only a requirement for traffic entering and leaving the field network, then architect a solution with a firewall at the hub location. When deploying the field network in a hub and spoke model, all traffic must traverse the hub location when crossing its boundary. Deploy your IPS solution at the hub using a powerful enterprise firewall and deploy Cyber Vision on the Cisco Industrial Router.

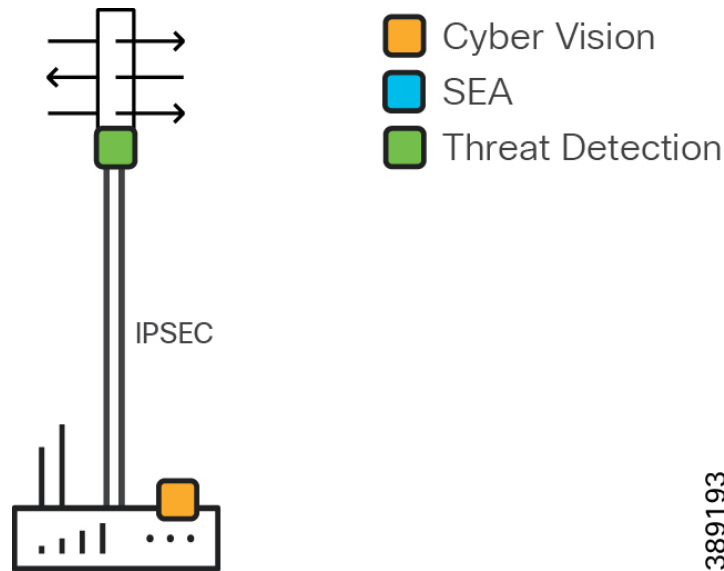
Figure 89: Cisco Industrial Router with Cyber Vision at the spoke site, IPS deployed in the hub firewall



389192

If IPS is a requirement at the edge, then an appropriate router must be deployed. In this scenario, use the IR1835 or the IR8340 to host Snort, and then deploy Cyber Vision on a neighbouring switch. If there is no room at the location for additional hardware, then a choice must be made between the applications.

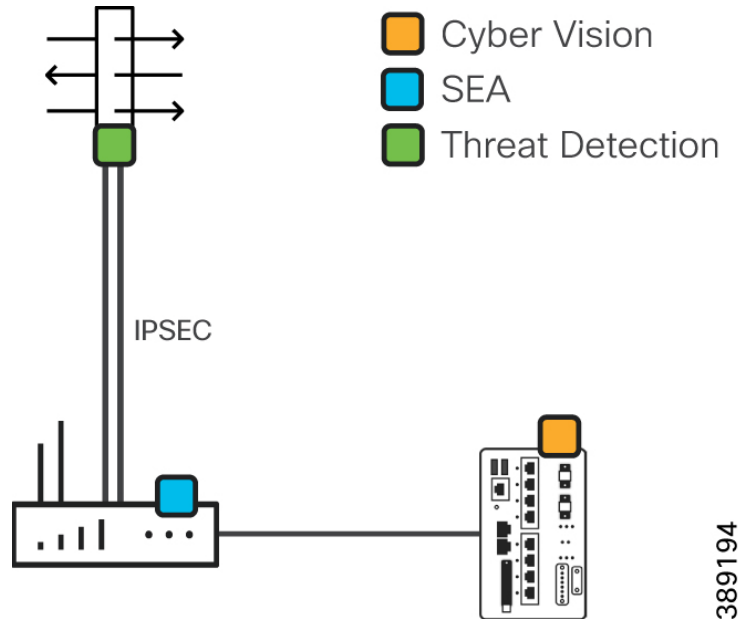
Figure 90: Cisco Industrial Router with the NGFW addon, Cisco Cyber Vision deployed on a neighboring switch



Problem statement: I need a firewall to protect my critical assets, I need IDS/IPS capabilities, I also want Cyber Vision for visibility, and I want Cisco SEA for remote access.

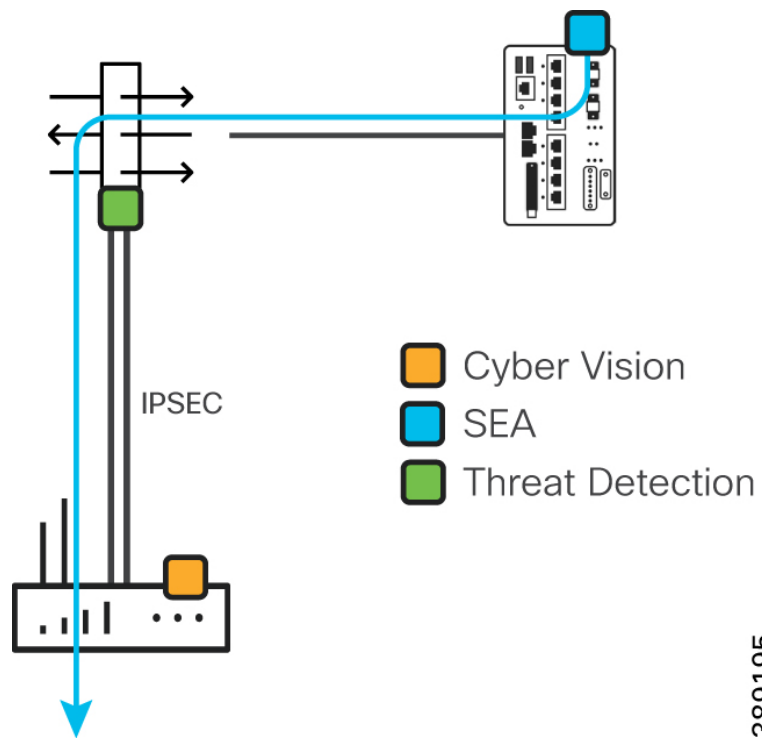
In the ideal scenario, all capabilities can be deployed across various points in the architecture. The IPS solution can be deployed at the hub location protecting all north/south traffic from the field networks, Cyber Vision can be deployed on the access switches used by OT endpoints, and SEA can be deployed at the router for remote access.

Figure 91: Cisco Industrial Router with SEA, Cyber Vision on a neighbouring switch, and IPS deployed at the hub firewall



However, similar to the last scenario, this may not always be available. All Cisco SEA needs to operate is a connection northbound to the cloud, and a routed connection to the OT assets. To accomplish this, Cisco SEA can be deployed at the hub location, alleviating the burden to host the SEA agent on the Industrial Router. This design decision then brings you back to the decision tree for deploying both Cyber Vision and Snort which was discussed in the previous scenario.

Figure 92: Cisco SEA deployed at the hub site with IP reachability to the spoke



For more information on deploying Cisco SEA at the hub location, see [Hosting SEA Agent\(s\) at a hub site](#).

Chapter 4. Secure Remote Access for Industrial Networks

The pandemic normalized remote working; and at the end of the pandemic, with a large section of employees still favouring the flexibility offered by remote working, the work model in currency is hybrid work. It allows employees, partners, and vendors to work on site as well as at home, co-working spaces, and anywhere in between—wherever and however they work most productively.

This flexibility comes with a cost. Cyberattacks today have become more sophisticated and multipronged. At the same time, the proliferation of user endpoints—both corporate and personal devices, whether managed or unmanaged—has expanded the attack surface, leaving organizations and their end users vulnerable to malware and ransomware attacks.

In a bid to combat these security challenges, organizations often purchase new tools to solve specific problems, ending up with more security tools than they can effectively orchestrate or manage. This leads to tool sprawl which can create too many alerts, reduce threat response time, and open security vulnerabilities.

Organizations need a solution that provides trusted hybrid / remote workers, contractors and vendors with secure access to the organization network, applications, data, machines and other services. Access should be granted to specific devices, only when needed, under flexible constraints to meet compliance needs.

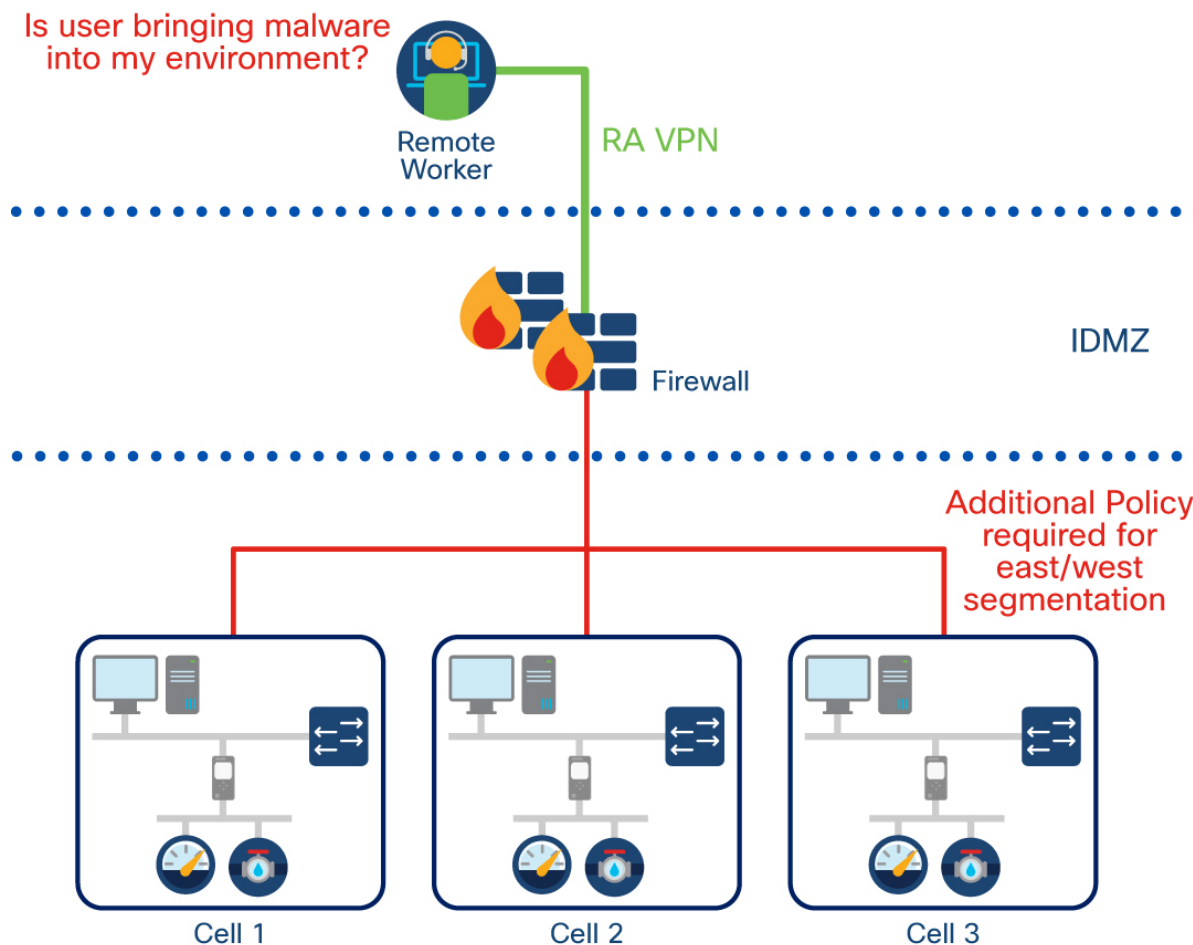
Remote Access Technologies

Remote access solutions come in many forms, and it can often be confusing to understand which one will meet business needs. This design guide will talk a little about virtual private networks, the remote desktop protocol, and the evolution towards zero trust network access.

Virtual Private Networks

Virtual private networks, or VPNs, are a way to give remote users an encrypted connection over an Internet network. The encrypted connection helps ensure that sensitive data is safely transmitted. It prevents unauthorized people from eavesdropping on the traffic and allows the user to conduct work remotely. VPN technology is widely used in corporate environments. While logging into these networks helps users securely and remotely access work resources and applications, they can be exploited by hackers seeking to steal login credentials.

Figure 93 Using a VPN to connect to the ICS network



388476

VPNs do offer security benefits, but if not configured correctly, can give an adversary unrestricted access to the OT (Operational Technology) network. With what is known about the risk of stolen credentials, using an MFA (Multi Factor Authentication) solution should be a given for any remote access solution. However, the additional risk with VPN solutions is the potential for even a legitimate user to unknowingly bring malicious software into the environment. Using a VPN, the network is extended to the remote user, which brings their machine as a remote client to the network. If a device is not scanned before access is given, malicious software could unintentionally be introduced to the OT network.

Additionally, with a VPN solution, additional access control needs to be placed on the user to limit their actions to a pre-determined scope of work. In the case of a vendor who is performing maintenance in a production environment, the vendor needs access to a single machine, during the duration of a maintenance window, and should not have the ability to laterally move to any other machine, neither intentionally nor unintentionally.

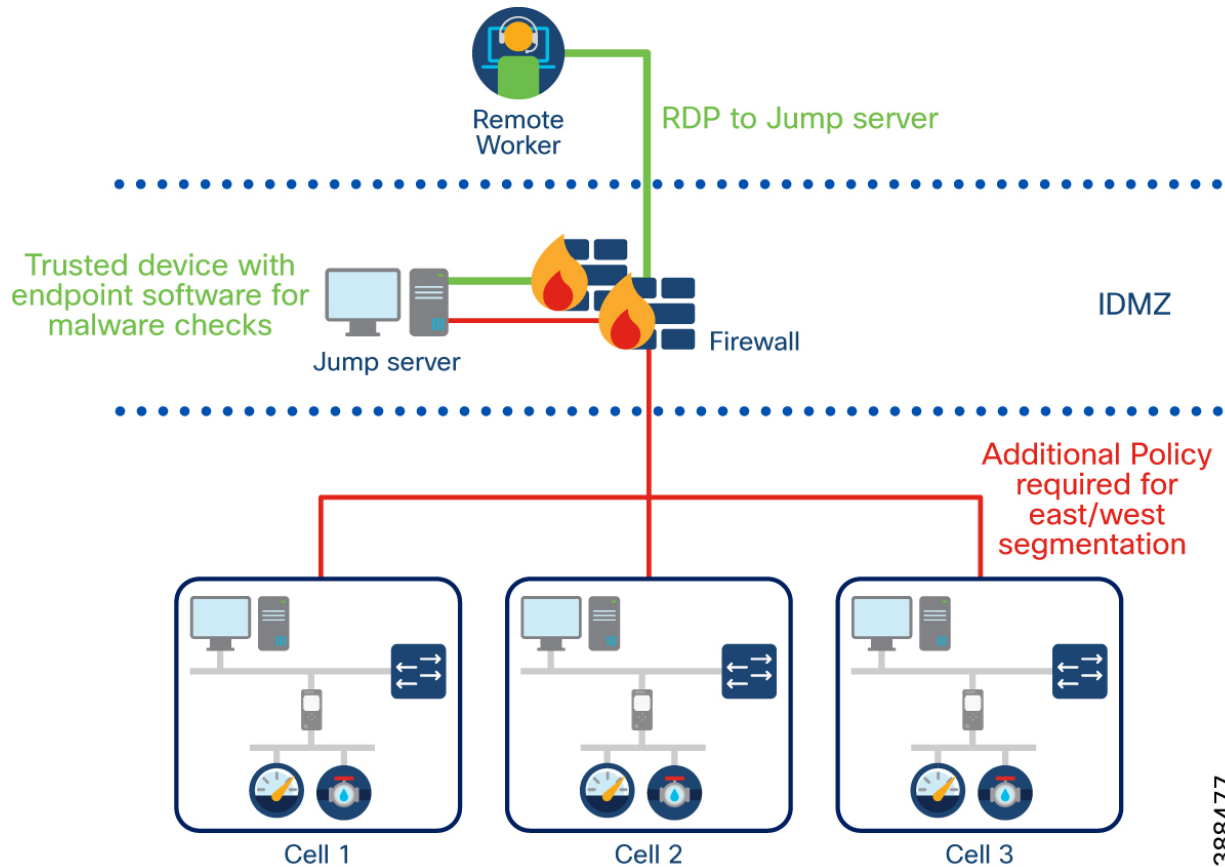
While all the necessary security checks and controls are possible with a traditional VPN based solution, the operational overhead often leads to lackluster policies and wide-open access policies after remote access has been granted. Additionally, VPN access is commonly maintained by a separate entity in the organization, which causes delays for vendor connectivity and slows

down the line of business. As a result, even in the IT world, there has been a shift towards looking for VPN alternatives.

Accessing Jump Servers with the Remote Desktop Protocol

Remote Desktop Protocol (RDP) is a protocol that lets you take control of a computer remotely. For example, an employee can access all their engineering workstations, project files and network resources from their home computer using RDP. It is also often used by third party support to remotely access machines that need repair.

Figure 94 Using a Jump Server to connect to the ICS network



In OT networks, the server which grants access to the plant remotely is often known as a jump server (or a jump box / jump host). The jump server solves the challenge an OT network can face with malicious software being introduced to the network by a remote device. After a VPN connection is established, policy is placed on the firewall to only allow users to access a set of jump servers that have been assigned to them. All activities performed on the OT network must originate from the jump server, which is a trusted device fully controlled by the networking team.

While jump servers solve one challenge, they do not help solve the challenge of controlling what a user can do once they have access to the jump server. Best practices would leave jump servers in a quarantined state, where they are denied any access to the OT network until called upon. As necessary, security administrators will open specific policies to control what a user can and

cannot do from that server. Once again, the operational burden can be overwhelming, and, in reality, jump servers expose themselves to the same frailties as the VPN solution.

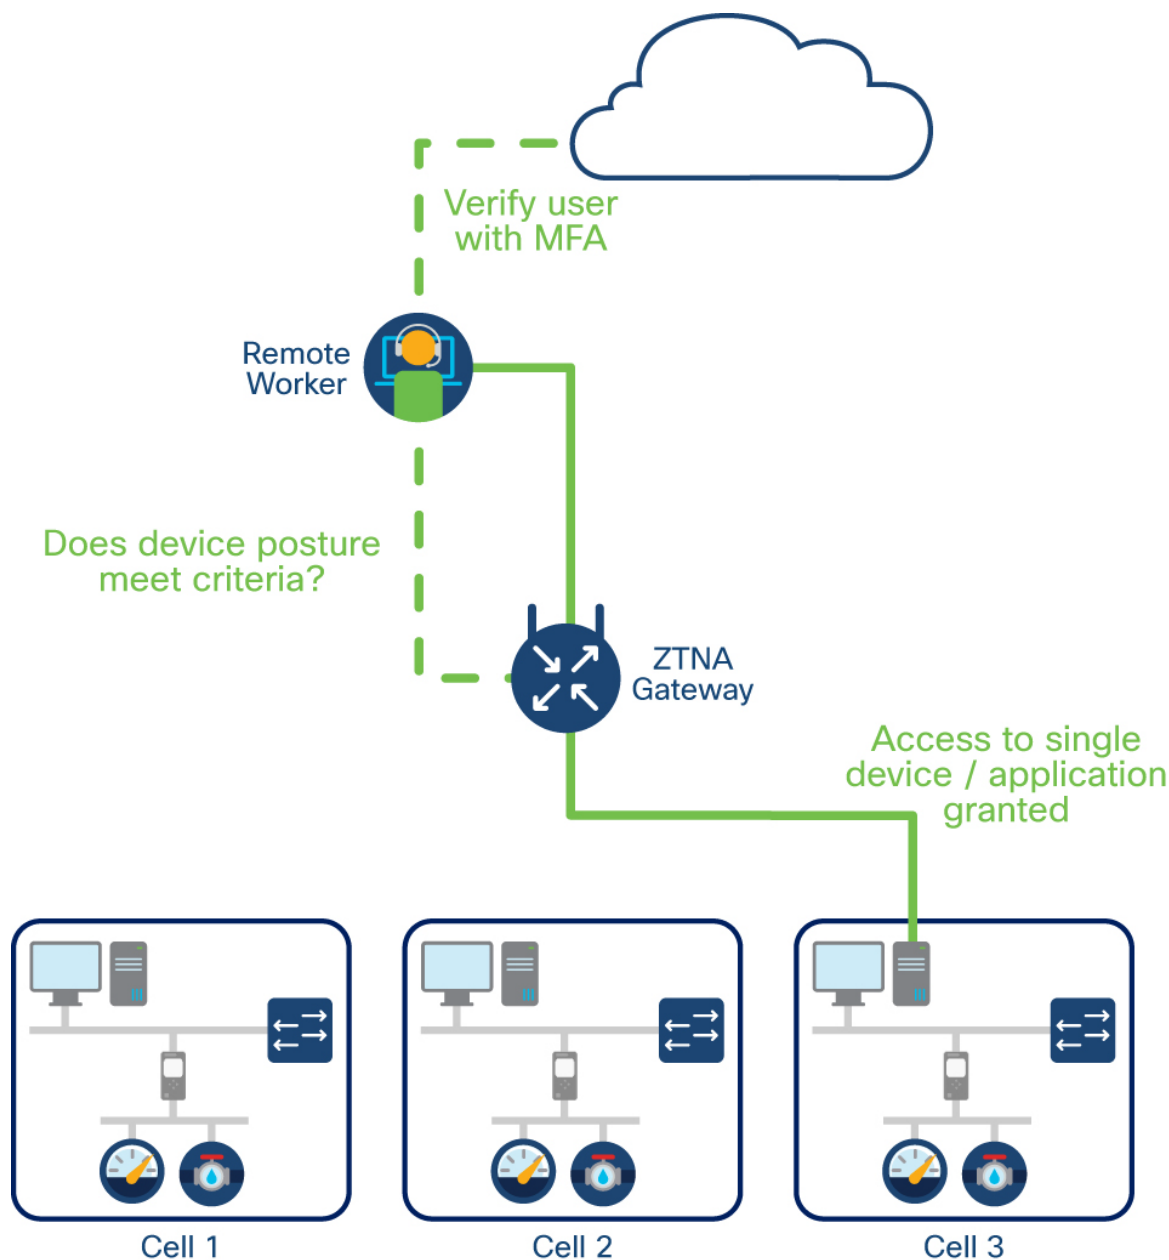
In 2020 [hackers harvested and sold as many as 250,000 RDP server credentials in an underground marketplace](#), xDedic. These credentials gave buyers access to all the data on the servers and the ability to launch future attacks using the servers. [Attacks against RDP grew by 768% in 2020](#), according to ESET.

Unrelated to security, jump servers are assets that will need to be managed, consuming power, and requiring ongoing maintenance. Additionally, if licensed software is needed to perform a task such as the PLC programming software, licensing must be purchased, that software must be maintained, and the cost and complexity increases for the consumer.

Zero Trust Network Access

Zero Trust Network Access (ZTNA) is a security service that verifies users and grants access to specific applications based on identity and context policies. Zero trust can be summed up as “never trust; always verify.” Often when users log in to a VPN, they are granted complete access to the entire network. Alternatively, ZTNA solutions connect authorized users directly to applications rather than to the network—and only to those applications they are authorized to access on need-to-know-based policies.

Figure 95 Using ZTNA to connect to a single device on the ICS network



388478

Adoption of zero trust can help address common security challenges in the workforce, such as phishing, malware, credential theft, remote access, and device security (BYOD). This is done by securing the three primary factors that make up the workforce: users, their devices, and the applications they access.

Verify Users

Ensuring the trust of your users whenever they attempt to access applications remotely is the first step toward secure remote access. The following capabilities aid in mitigating the threat of Initial Access exploits:

- **Multi-Factor Authentication (MFA)** - Authentication based on usernames and passwords alone is unreliable since users may have trouble storing, remembering, and managing them across multiple accounts, and many reuse passwords across services and create passwords that lack complexity. Passwords also offer weak security because of the ease of acquiring them through hacking, phishing, and malware. MFA requires extra means of verification that unauthorized users will not have. Even if a threat actor can impersonate a user with one piece of evidence, they will not be able to provide two or more.
- **Single Sign On (SSO)** - SSO is an authentication process that provides users with one easy and consistent login experience across all applications, eliminating the need to supply user credentials with every application or access request. Using SSO, user experience is streamlined across multiple applications, while security administrators can enforce strict user policies in a centralised location. MFA and user policy can be applied during SSO and eliminates the need to duplicate and maintain authentication policies across multiple applications, such as remote access software.

Device Posture

If users are logging into your company applications with outdated devices, there is a chance they could also be unwittingly spreading malware and using keyloggers to record your keystrokes. Meaning any data you type, including your username or password, can be recorded, and sent to an attacker command and control servers. As a result, your company data could be at risk if just one out-of-date device logs in, potentially spreading malware throughout your environment. Or worse, spreading ransomware that will keep your files hostage until a ransom is paid to decrypt them.

To protect the OT network from introducing malware to the environment, the remote access solution must have the following capabilities:

- **Device Posture Assessment** - The device posture assessment analyses the device, assesses its security posture, and reports it to the policy decision management system. Organizations need to enable secure and direct access to business applications for a diverse set of users (remote workers, vendors, and contractors) and their devices that typically reside outside of the control of corporate EMM (enterprise mobility management) and MDM (mobile device management) solutions. Enforcing consistent security policies across managed devices, bring your own devices (BYOD), corporate owned, personally enabled, and third-party (contractor or partner) devices poses a significant challenge. IT security teams often lack insight and an enforcement mechanism when making an access decision on endpoints, particularly among unmanaged devices. This is when device trust is important to establish.
- **Anti-Malware** - Advanced malware goal, in general, is to penetrate a system and avoid detection. Once loaded onto a computer system, advanced malware can self-replicate and insert itself into other programs or files, infecting them in the process. Anti-malware protection should be implemented in both the network (to prevent initial infection and detect attempts of spread) and in the endpoint (to prevent endpoint infection and remove unwanted threats). For the purposes of this design guide, this capability represents endpoint anti-malware.

Least Privilege Access Control

Zero trust requires that a user be given access only to the applications they truly need to do their job — and no more. Application access should be governed by adaptive access policies, created based on the sensitivity of the data in the application. This granularity ensures that access is provided only to users or groups of users who need it, from locations and devices that are trusted. To protect the network from discovery, lateral movement and impairing process control, a remote access solution must provide the following capabilities:

- **Identity Authorization** - Establish trust by verifying user and device identity at every access attempt. Least privilege access should be assigned to every user and device on the network, meaning only the applications, network resources and workload communications that are required should be permitted. Access to the full network should never be granted.
- **Time-based Access** – Remote access should not be an always-on feature. Access should be granted only when needed and restricted to the resources required for a given access attempt. A remote access solution for the OT network should be off by default, and access is granted at time of need, for a specified period before being turned off by the system. If a session expires beyond the allocated window, a new session should be created.
- **Session Request** – Just because a session has been scheduled, does not always mean that it is safe to perform the remote access request. By having a session approval flow, a remote user makes a request for a session to be open and only when an administrator clicks approve will the user be granted access.

Auditing

Many compliance standards will require that an audit trail be maintained for all activity that occurs from remote networks:

- **Authentication Logs** – Authentication logs show you where and how users authenticate, with usernames, location, time, device posture and access logs.
- **Administration Logs** – Administration logs show you the sessions that were created, who created the sessions and what access control measures were put in place for the end users.
- **Session Monitoring** – Enable an administrative user to supervise a remote session and view in real-time what is happening during the session. For example, when an external technician delivers remote support for an asset, an internal OT operators may want to overlook the actions taken during the remote access session
- **Session Termination** – The ability for an administrator to terminate an active session that either should never have become active in the first place, or while monitoring a session, the remote user attempts to deviate from their permitted actions
- **Session Recording** – The ability for remote sessions to be recorded and stored for use in an audit trail. If a breach were to occur, having the ability to watch back what remote users did to a system aids incident investigation

- **Flow Analytics.** Network Detection and Response (NDR) solutions leverage pre-existing infrastructure to offer enterprise-wide, contextual visibility of network traffic. Flow information can be used to conduct forensic analysis to aid in lateral threat movement investigations, ensure ongoing zero trust verification is provided, and modern tools can even detect threats in encrypted traffic

Secure Equipment Access Design Guidance

To address the secure remote access requirements outlined in the previous section, this Cisco Validated Design (CVD) guide covers the following platforms and software to accomplish a comprehensive and secure remote access for industrial networks design:

- Cisco Secure Equipment Access
- Cisco Duo
- Cisco Secure Endpoint

Cisco Secure Equipment Access

Note: Not every feature of Cisco Secure Equipment Access will be explored in this design guide. The purpose of this document is to provide architectural and design recommendations when deploying and operating the product. For full feature documentation see the [product documentation](#).

Cisco Secure Equipment Access (SEA) is a hybrid cloud service that enable operations teams to easily connect to remote assets or machines for configuration, monitoring, and troubleshooting. SEA provides granular access controls that can easily be managed by an operations administrator, and secure connectivity for authorized users, including internal employees and external workers. The SEA service help organizations to improve efficiency by decreasing the time and cost required for travel to remote sites for equipment maintenance or to respond to an emergency. For example, operations teams can configure equipment such as traffic signal controllers, in-vehicle dispatch systems, cameras and other systems deployed in the field and connected using [Cisco Industrial Routers](#) and/or [Cisco Industrial Ethernet Switches](#).

Using SEA, a worker can access a remote asset from anywhere simply by using a browser, without needing to install any additional software on their laptop. The remote equipment can be accessed using either GUI- or CLI-based methods. Supported protocols are HTTP(S), SSH, RDP for Windows-based systems, VNC, and Telnet.

Note: Although HTTP and Telnet are not recommended forms of communication due to their use of cleartext, when used under the SEA construct they are wrapped in an encrypted session, and therefore the communication only becomes cleartext at the other end of the proxy (SEA Agent). Pre-cautions still need to be made between the SEA agent and the target endpoint.

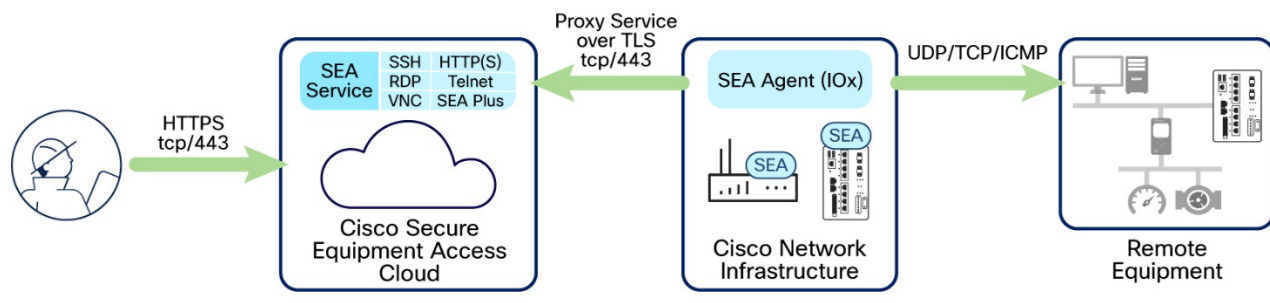
SEA Plus, a client-based capability for ZTNA in OT, provides further flexibility by enabling users to configure any type of equipment that supports IP connectivity. With SEA Plus, a direct, secure data connection is created between client software on the user's computer and the remote asset, enabling the user to easily interact with and exchange files with the asset. SEA Plus supports IPv4 TCP, UDP, and ICMP based protocols. The feature provides users with the

advanced ability to define specific channels for communications between a user and the remote system and block everything outside that.

Cisco SEA Components

SEA comprises a few elements that must be considered when architecting a secure remote access solution.

Figure 96 Components of Cisco SEA



389196

Secure Equipment Access Cloud

Cisco SEA is a hybrid-cloud solution, where all users interact with a SaaS portal, and access to remote equipment is provided via on-premises proxies. The SEA cloud is the main component of the solution. It is where SEA administrators configure access, define policies, monitor sessions and SEA users access the resources they have permissions to. There are many misconceptions about the cloud, especially its use in operational environments. However, remote access is a solution that cannot avoid connectivity. Remote users originate outside of the network, and if not using a cloud broker, operators must maintain policy across every site individually. If a security administrator wants to make an organisational wide policy change, they must make sure each firewall boundary is configured correctly to enforce this change.

Figure 97 Security policies must be managed and maintained across all firewall boundaries

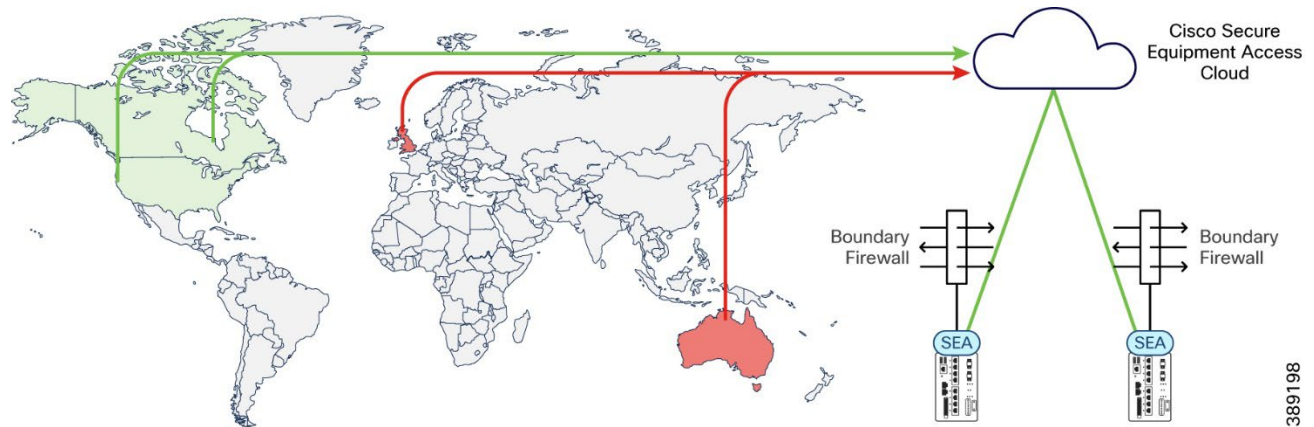


389197

By first bringing all remote users to a trusted cloud resource it enables security administrators to have a single point of control for those users. Identity can be verified, device posture can be

checked, geolocation can be looked up, and any other policy controls that we will discuss later in the document. On the OT, operators simply need to allow connectivity between the SEA agent that exists behind the firewall, rather than open multiple connections from different users, potentially all over the world.

Figure 98 Consistent policy across all locations



Secure Equipment Access Agent

IOx is a Cisco-developed end-to-end application framework that provides application hosting capabilities for different application types on Cisco network platforms. The SEA agent is an IOx application that is installed on supported network devices. The SEA agent creates a secure connection (TLS 1.3 over TCP port 443) to the SEA cloud and provides a reverse proxy function from the SEA cloud to the devices in an OT network. The SEA agent is supported on the following network devices:

- Cisco Catalyst Industrial Routers (IR1101, IR1800, IR8340)
- Cisco Catalyst Industrial Ethernet Switches (IE9300, IE3400, IE3300, IE3100)
- Cisco Catalyst 9300


Cisco SEA Architecture Guidance

The position that SEA components are placed in the architecture will depend on the architecture that remote access is being granted.

Direct Access to the Cloud

Some use cases, such as distributed assets in a transportation network, industrial routers will have a direct connection over a cellular interface.

Figure 99 Connecting the SEA Agent to IoT OD over a cellular channel




In this architecture, there is no interference between the SEA Agent sitting on the Industrial router and the SEA cloud and no additional architectural considerations are required.

Hosting SEA Agent(s) at a hub site

Hosting the SEA Agent at distributed sites is not always an option. Perhaps the hardware does not support hosting the agent, or [other applications have been chosen to run at the edge that is consuming all compute resources](#). The SEA agent does not need to be directly connected to the network device of target devices. The SEA agent just needs reachability so network packets can reach its destination.

Figure 100 Hosting SEA Agent(s) at a hub site



If the SEA agent cannot be deployed at the target site, it can be run at a hub location with IP accessibility to its target devices. In a typical hub-spoke model, applications and services running in the hub will have IP accessibility to the devices in the spokes. The following design considerations will be needed when deploying in this model:

- If there are firewall rules between the hub and spoke locations, which is the most likely outcome, firewall rules will need to be created to enable connectivity from the SEA agent to the remote targets. It is important to note that SEA has all the policies needed to provide least privilege access to remote targets. It is not recommended that security administrators duplicate these granular policies on the firewall between hub and spoke sites. Instead, security administrators should have a single firewall policy, either allow or

deny, with security administrators only opening the firewall connection during known and scheduled time windows. SEA will already restrict which devices and over what protocols users can access the network, the additional firewall rule just gives an additional layer of protection, but we do not want to over complicate our design.

- The SEA agent has a maximum of 10 concurrent remote access sessions (depending on the device chosen to host it). This is normally not a concern when running at the edge, but when running at a hub location, it may be a limiting factor. If more than 10 concurrent sessions are required, it is recommended to deploy more agents to handle the scale.
- Deploy a dedicated SEA agent for each site. While not a strict requirement, troubleshooting and management of connections will be much easier as the scale of the operations grows.


SEA in Plant networks

Cisco SEA provides benefits to more than just distributed field assets. Industrial Automation networks, or any network that has a Purdue model structure (OT - IDMZ - IT), typically deploy VPN solutions, accompanied by Jump Servers in the ICS network. When pivoting to a ZTNA solution, a decision needs to be made on where the SEA Agent will be installed.

The SEA Agent needs two connections. On the northside interface, the SEA Agent needs IP connectivity to the SEA cloud. On the southside interface, the SEA Agent needs IP accessibility to the target endpoints.

The first architecture option is to place the SEA Agent(s) within the same subnet that the remote targets reside.


Figure 101 Connecting the SEA Agent through IT controlled network infrastructure



By deploying the SEA Agent within the same subnet as the target devices, the SEA Agent becomes part of the same layer 2 domain and does not require any additional routes or security policies to be placed on the network to reach the target devices. A secure connection is created from the SEA Agent in the process network (for example, a Cell/Area Zone) and sent through IT controlled infrastructure to the SEA cloud. The SEA Agent then proxies remote connections into the process network, and additional policy can be put in place on the network device to contain all remote access sessions to the layer 2 domain the SEA Agent resides.

It is important to note, in this design, the OT network itself does not need connectivity to the cloud. The SEA agent supports up to 10 virtual interfaces, which can all reside in different VLANs.

Figure 102 SEA agent with multiple virtual interfaces




With this approach, the SEA agent can have one interface with a route to the cloud; either a dedicated subnet just for SEA agents or deployed within an existing cloud facing subnet within the OT network. Additional interfaces are then created for its southbound access, where OT VLANs remain unroutable from cloud connected subnets.

SEA in IDMZ

Regardless of the security controls, some architects may have discomfort deploying SEA agents with an interface in an OT VLAN. In this case, SEA can be deployed fully in the IDMZ.

Figure 103 SEA agent in the IDMZ

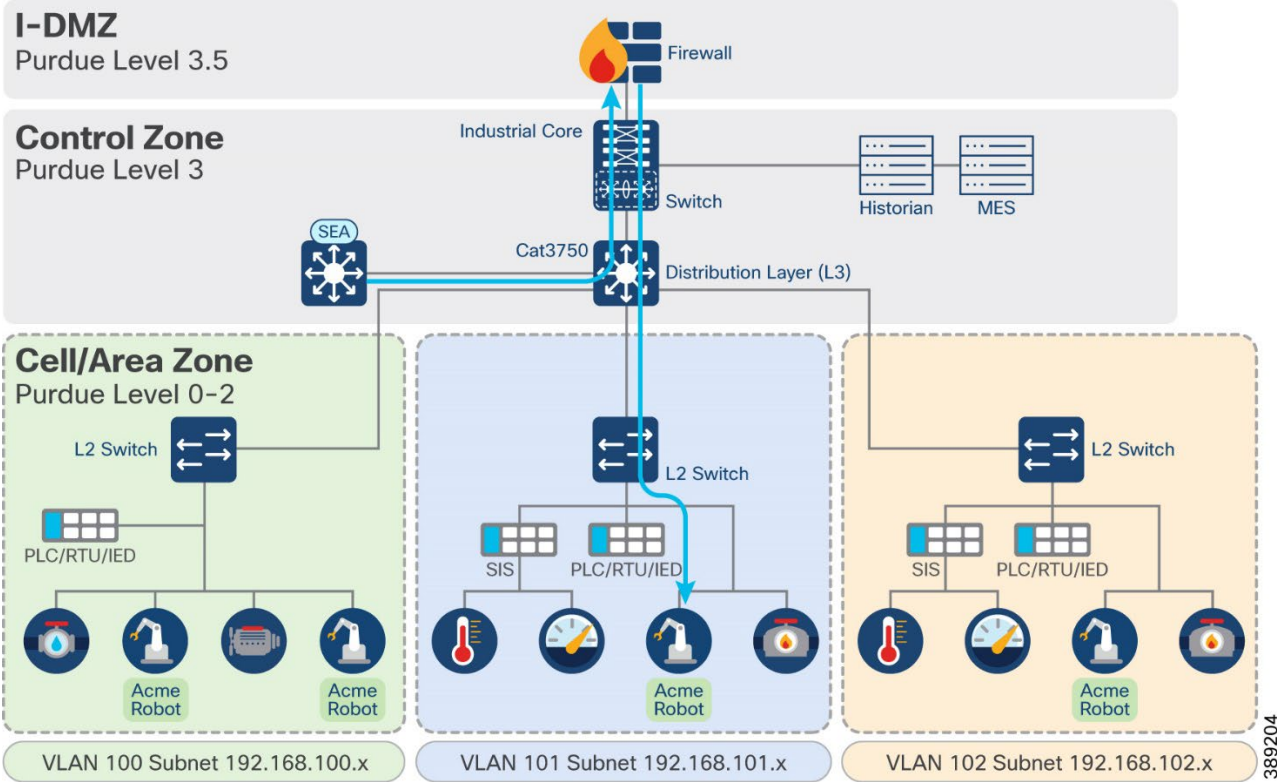


The IDMZ is designed to be the boundary between the IT and the OT network. Services that reside within the IDMZ have interfaces into both parts of the network, which is fitting for the SEA Agent. The advantage of putting the SEA Agent in the IDMZ is it tightly aligns with existing mechanisms that may already exist in the network; nothing in the OT domain should have a direct connection through the IDMZ. It is worth noting, the SEA Agent is a security appliance, and since it does follow the principles of zero trust, the presence of the agent within the OT network would not pose an additional security risk to the organization.

Nevertheless, if the SEA Agent is to be placed in the IDMZ, all the same design considerations will apply as already discussed in the chapter [Hosting SEA agent\(s\) at a hub site](#), where firewall rules and scale considerations must be considered.

Note: The SEA agent does not need to physically reside in the IDMZ. If the SEA agent is in its own subnet, the gateway for that subnet can be the IDMZ firewall, logically placing it as an IDMZ application. This deployment may be considered if the architecture has a device that can technically host the SEA agent, but architects still want to route all its traffic, including southbound, through the IDMZ firewall.

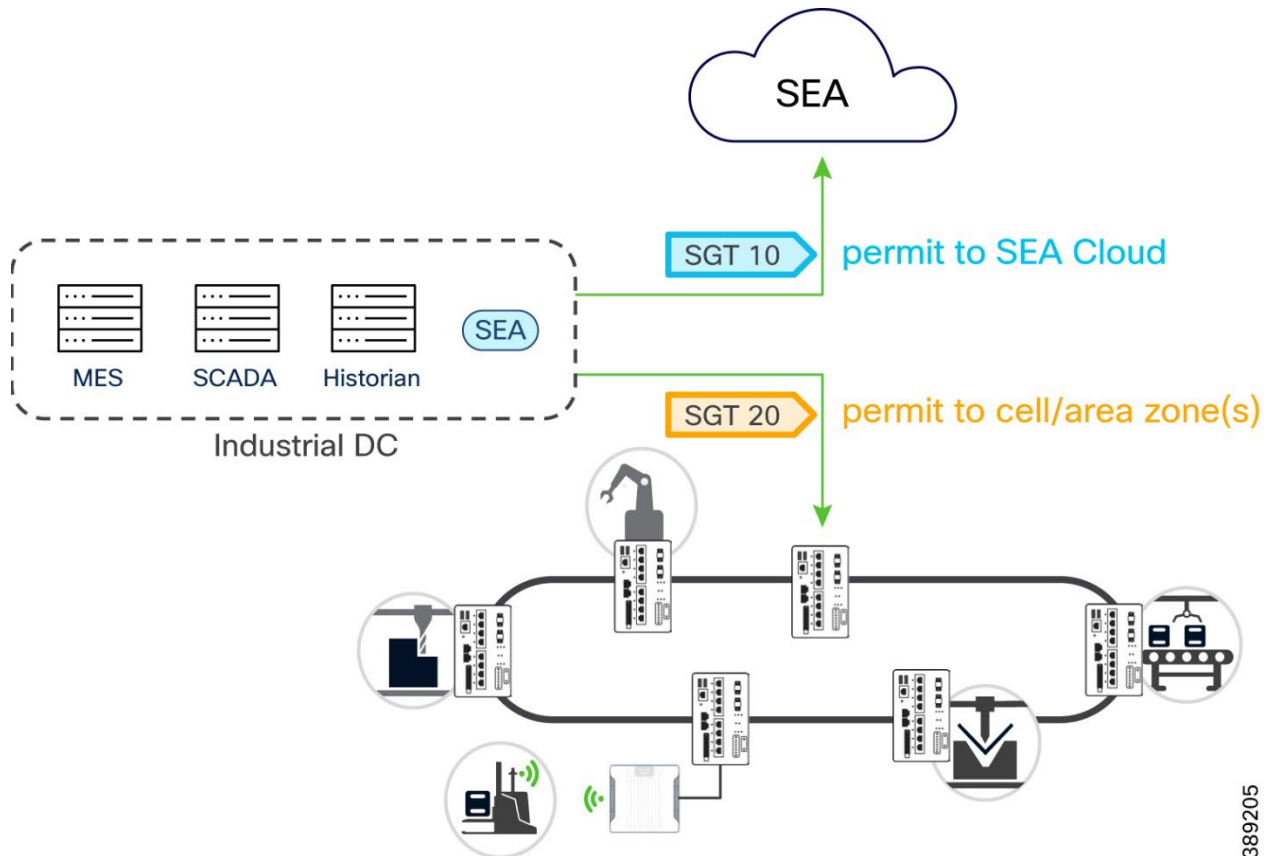
Figure 104 SEA agent logically in the IDMZ



Using TrustSec with SEA

If your organization is following segmentation guidance from [this guide](#), then a rule may be needed in the switches to allow traffic coming from the SEA agent to reach its intended destination. The recommendation is to provide a static SGT assignment to the SEA Agent IP addresses and to create SGACLs that allow the chosen tag to communicate with all other SGTs in the OT network. A different SGT can be assigned to the different interfaces for the agent. For example, the subnet defined for accessing the SEA cloud can be assigned a different set of permissions than the interface created for accessing the control network.

Figure 105 Using SEA in a TrustSec protected network




There is no real security benefit in creating complex policies in TrustSec to control communications coming from SEA. SEA in itself is the security appliance for remote access. Only trusted and verified users should have access to the system, they will only have access to resources allocated to them, they must undergo a session approval flow, and only during the time window allocated to them. All of the necessary controls are already in place and the additional of TrustSec policies would only complicate the deployment for minimal security gain.

If a security administrator did want to disconnect SEA through TrustSec, the SGT could act as a “kill switch” for the system, where all SEA agents could quickly be quarantined by swapping the static SGT assignment to one that has no connectivity.

Connecting to the SEA Cloud via web proxy

In some instances, your organization may be using a web proxy in the IDMZ to proxy all web connection from the OT network to the cloud.

Figure 106 Connecting SEA Agent to SEA cloud via a web proxy



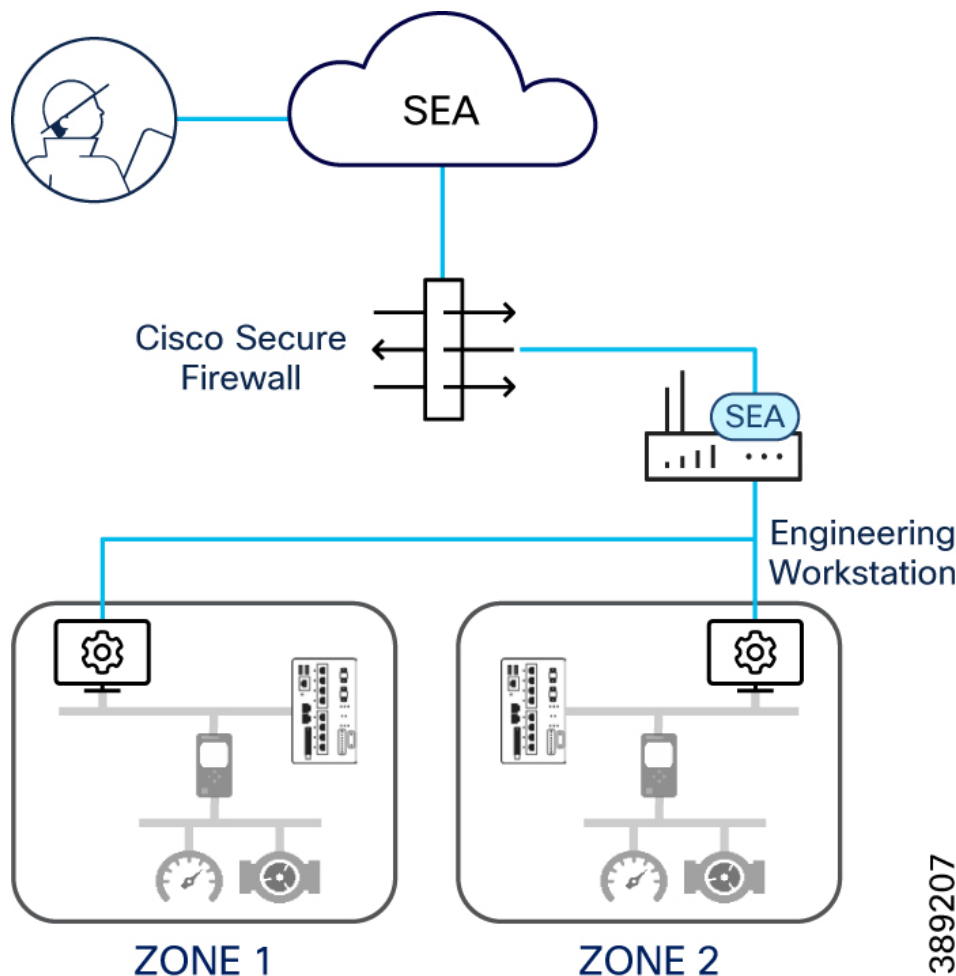
When deploying SEA agents an explicit proxy can be configured instead of requiring direct access to the cloud. There is no recommendation on which proxy to use, the only requirement is that the proxy supports the HTTP CONNECT method. The HTTP CONNECT method requests that a proxy establish an end-to-end encrypted session between the SEA agent and the SEA cloud service, meaning the data will be unreadable to the proxy. The only purpose of this proxy is to manage a single firewall rule between the trusted web proxy and the SEA cloud.

Using a web proxy is not a requirement. If one already exists, it is recommended to use this feature. If a proxy does not yet exist, administrators could create a dedicated subnet for SEA agents, provide reachability to the SEA cloud from that subnet, and create a firewall rule that ensures IP addresses in that subnet can reach out to SEA, and SEA only. That way, if another device was to establish itself on the SEA subnet, its only outbound reachability would be to the SEA cloud, where additional security measures would be in place to prevent malicious activity.

Migrating from existing remote access solution to SEA

Finally, if the organization is not yet ready to deviate from their traditional remote access solution, or need a transition path to adopt ZTNA, there may still be challenges with controlling access to engineering workstations that SEA can help solve.

Figure 107 Using SEA to complement existing traditional remote access solution



Engineering workstations are placed throughout the ICS network to bring remote vendors down to the cell/area zones to perform remote maintenance of their machines. One of the problems that arises from this is the possibility of an always on connection, or the administrative overhead of managing both VPN credentials and RDP access. By placing the SEA agent in the network, the product can be used to control vendor access to engineering workstations and gain all the user authentication, access control, and auditing benefits the SEA solution provides. Additional security measures will still be required between the workstations and the rest of the network, but the key challenge of making sure only trusted users have access to the infrastructure at determined times can be solved using SEA.

Connectivity requirements for SEA

The following TCP/UDP network ports and IP protocols must be opened on the network firewall to allow the edge devices to communicate with Cisco SEA cloud.

Table 5 Firewall Rules for SEA Cloud access

Port	Protocol	Destination	Description
53	UDP	IP of assigned DNS Server	The network device must have access to DNS resolution service
443	TCP	https://us.ciscoiot.com or https://eu.ciscoiot.com	HTTPS connection to access SEA cloud and for devices to register

Note: 1. You only need to choose one between EU and US in your firewall rules depending on which cluster is relevant to your deployment.

2. If you require IP address filters instead of domain name filters see [Firewall Rules](#) for an up-to-date list of IP addresses for each of the URLs specified in the table above.

Cisco SEA Authentication Options

As discussed previously, a ZTNA solution is broken down into three key components: verifying users, verifying devices, and least privilege access control. When verifying users in a remote access solution, it is recommended to use an SSO to avoid password sprawl and apply consistent policy control across applications.

Cisco Customer Identity

By default, all users created in Cisco SEA will use Cisco Customer Identity (CCI), an identity provider (IdP) managed and used by Cisco to enable users to navigate across multiple Cisco applications and websites with one set of login credentials ensuring seamless operation.

By default, SEA organizations will be configured to require MFA for all users accessing the dashboard. After the user has an account created in CCI, they must navigate to id.cisco.com to enroll in MFA. Cisco Duo and Google Authenticator are the only supported options at time of writing this guide.

External IdP

While organizations can protect their SEA deployments with MFA, they do not own a CCI administration portal and cannot define additional authentication policies on users accessing the SEA dashboard. For more control, it is recommended to bring your own SSO solution. SSO integration into SEA is done via Security Assertion Markup Language (SAML) 2.0. For this integration, SEA is the service provider, and your organization identity server is the IdP. While any IdP with SAML 2.0 support will work, the remainder of the design guide will demonstrate value using Cisco Duo.

Duo integrates with SEA in two ways:

- **Single Sign On:** SSO allows users to log in to SEA using their corporate account credentials. When a user enters their email ID, they are redirected to your organization IdP authentication page. After authentication, they are redirected back to SEA and logged in. By integrating Duo with the IdP, all users who log in to SEA via SSO will be protected by Duo policy.
- **SEA Plus:** Separately, Duo is natively integrated into SEA for the SEA Plus access method. SEA Plus enables client-based access to the OT network, which requires extra caution to be taken on which devices can be brought into the network.

Note: Both methods can be protected by the same Duo administration account but are two separate integrations. The SSO is a function of SEA login, while SEA Plus is a feature within SEA. The Duo integration that is embedded in the SEA user interface is for enabling Duo protection for SEA Plus enablement. To bring your own SSO to SEA, such as Duo, contact your Cisco representative.

Cisco Duo Components

Cisco Duo is a cloud-based solution but does comprise of some on-prem elements depending on the use case for your Duo deployment. For this design guide only two components are required:

- Duo Administration Panel
- Duo Device Health Application

Duo Administration Panel

The Duo Administration Panel is a software as a service (SaaS) where a Duo administrator can manage all aspects of a Duo subscription such as enforce policy on protected applications, manage user and their devices or monitor access activity.

An application in Duo is the mechanism to bind Duo protection to one or more of your services. For example, a Duo administrator may add the [Microsoft RDP](#) application to add MFA to their remote desktop logins or may add the [Microsoft Azure Active Directory](#) application to add protection to Azure Active Directory logons.

For this design guide, integration is between two cloud services (for example, Duo and SEA) and therefore no distinctive design considerations will be required.

Duo Device Health Application

Duo helps you control access to your applications through the policy system by restricting access when devices do not meet particular security requirements. Users are prompted to download the Duo Device Health app during the first log in attempt to a protected application with the Device Health application policy set to require the app. After installing the Device Health application, Duo blocks access to applications through the Duo browser-based authentication prompt if the device is unhealthy based on the Duo policy definition and informs the user of the reason for denying the authentication.

For more information see [Duo Device Health](#).

Validating User Trust with Cisco Duo

Cisco Duo is built on the foundation of zero trust. Duo verifies the identity of users and protects against breaches due to phishing and other password attacks with an advanced MFA solution, verifying trust in multiple ways before granting access.

Multifactor Authentication adds a second layer of trust that your users are who they say they are. After completing primary authentication (usually by entering a username and password), users verify their identity a second time, through a different channel. This reduces the likelihood that someone else can log in, since they would need both the password and their second factor to pose as the original user.

Duo provides flexible authentication options to fit a broad range of users, security profiles, and technical backgrounds such as employees, frequent travellers, contractors, vendors, customers, and partners. For more secure access to high-risk applications, require the use of:

- Easy to use, out-of-band mobile push notifications
- Phishing-proof Universal 2nd Factor security keys
- Biometric-based WebAuthn

Other MFA methods support diverse user login scenarios:

- Phone call-back for users who cannot receive texts
- Mobile one-time passcodes for travellers while offline
- Text message passcodes for users without Internet connectivity
- Temporary bypass codes for users who temporarily cannot use their enrolled devices

With Duo, you can enforce contextual access policies allowing access to your applications with user-, device-, and location-based controls. The context includes several aspects of their login attempt such as where they are located, what role they have in your organization, what type of device, they are using, etc. With these policies, you can limit access to only what your users need to do their jobs and add stricter controls for access to more sensitive applications without negatively impacting use workflows.

Validating Device Trust when using the SEA Plus Access Method

Duo verifies the security posture and management status of endpoints before granting access to your applications, and blocks access if the device is unhealthy or does not meet your security requirements.

Before enabling client-based access, Duo can check the security health of all user devices attempting to access your applications. By leveraging the visibility of devices connecting to your applications, you can establish device-based access policies to prevent any risky or untrusted devices from accessing your applications. For access to high-risk applications, you may require a device to be corporate-owned or managed by your organization IT team.

Duo allows you to establish mobile device trust with or without the use of Mobile Device Management (MDM) software. Users may object to installing MDMs on their personal devices due to privacy concerns, resulting in lower overall adoption and reduced insight into their device

security. And sometimes it is outside of your IT team's control to install an agent on the personal devices of third-party provider that may need access to your applications.

Whether or not you have an MDM solution, Duo can allow you to block devices from accessing your applications based on:

- OS, browser, and plug-in versions and how long they have been out of date
- Status of enabled security features (configured or disabled)
- Full disk encryption
- Mobile device biometrics (face ID/touch ID)
- Screen lock
- Tampered (jailbroken, rooted, or failed Google SafetyNet)

Cisco Secure Endpoint

Cisco Secure Endpoint integrates prevention, detection, threat hunting, and response capabilities in a unified solution leveraging the power of cloud-based analytics. Secure Endpoint will protect your Windows, Mac, Linux, Android, and iOS devices through a public or private cloud deployment.

Note: SEA Plus is only supported on Windows (minimum version Windows 10 64-bit).

Duo and Secure Endpoint work together to provide stronger access security. With Duo and Secure Endpoint, you have the tools in place to establish trust in users' devices connecting to protected applications. Together, they offer the ability to:

- **Prevent:** Duo evaluates risk conditions, device health, and security status on every access attempt. Secure Endpoint strengthens defenses using the best global threat intelligence and automatically blocks known fileless and file-based malware.
- **Detect:** Duo verifies whether the Secure Endpoint agent is running on the user's device before granting application access. Secure Endpoint detects stealthy threats by continuously monitoring file activity, while allowing you to run advanced search on the endpoint.
- **Respond:** Duo automates the response by leveraging telemetry from Secure Endpoint to block access from devices infected with malware. Secure Endpoint rapidly contains the attack by isolating an infected endpoint and accelerating remediation cycles.

The Cisco Secure Endpoint integration verifies endpoint status and blocks access from Duo trusted endpoint client systems that Cisco Secure Endpoint identifies as "compromised". By restricting SEA Plus access to only those devices with a clean Secure Endpoint install, you can gain peace of mind that use cases such as file transfer from remote users to OT devices will not contain hidden malware and cause a breach by an unwitting threat actor.

Cisco Secure Endpoint Capabilities

In the rapidly evolving world of malware, threats are becoming harder and harder to detect. The most advanced 1% of these threats, those that will eventually enter and wreak havoc in your network, could potentially go undetected. However, Secure Endpoint provides comprehensive

protection against that 1%. Secure Endpoint prevents breaches, blocks malware at the point of entry, and continuously monitors and analyses file and process activity to rapidly detect, contain, and remediate threats that can evade front-line defenses.

Stopping threats at the earliest point in time ensures minimal damage to endpoints and less downtime after a breach. Secure Endpoint employs a robust set of preventative technologies to stop malware, in real-time, protecting endpoints against today's most common attacks as well as emerging cyberthreats.

- **File Reputation:** known malware is quickly and easily quarantined at the point of entry without any processor-intensive scanning.
- **Antivirus:** endpoints benefit from custom signature-based detection, allowing administrators to deliver robust control capabilities and enforce blocklists.
- **Polymorphic malware detection:** Malware actors will often write different variations of the same malware to avoid common detection techniques. Secure Endpoint can detect these variants, or polymorphic malware through loose fingerprinting. Loose fingerprinting will look for similarities between the suspicious file content and the content of known malware families, and convict if there is a substantial match.
- **Machine Learning analysis:** Machine learning capabilities in Secure Endpoint are fed by the comprehensive data set of Cisco Talos to ensure a better, more accurate model. Together, the machine learning in Secure Endpoint can help detect never-before-seen malware at the point of entry.
- **Exploit Prevention:** Memory attacks can penetrate endpoints, and malware evades security defenses by exploiting vulnerabilities in applications and operating system processes.
- **Script protection:** provides enhanced visibility in Device Trajectory into scripts executing on your endpoints and helps protect against script-based attacks commonly used by malware. Script control provides additional protection by allowing the Exploit Prevention engine to prevent certain DLLs from being loaded by some commonly exploited desktop applications and their child processes.
- **Behavioral protection:** enhanced behavioral analysis continually monitors all user and endpoint activity to protect against malicious behavior in real-time by matching a stream of activity records against a set of attack activity patterns which are dynamically updated as threats evolve.
- **Device Control:** lets you control the usage of USB mass storage devices and prevent attacks from these devices.

Cisco SEA Policy Creation Guidance

After users have been authenticated in the system, least privilege policies should be created to ensure users only have access to the systems they need to perform their job, and only for the limited period needed to complete their task. The following best practices should be taken into consideration when configuring SEA for remote access into critical infrastructure.

SEA Roles

SEA has the following built-in roles:

- **SEA System Admin** can add SEA Agents to network devices, configure connected clients, and define the access methods to reach those clients
- **SEA Access Admin** has permissions to both access remote sessions and manage user permissions to those sessions
- **SEA Access Manager** can approve access requests and launch remote sessions
- **SEA Users** are the end users of the remote access deployment. These users have no configuration privileges, but simply consume the product within the boundaries set by the admins.

Do not fall into the trap of making all internal users admin users of the system. Administrator roles should be reserved for a limited set of users who are responsible for provisioning the system. The more SEA System Admin accounts that exist, the more users who can be targeted in a social engineering attack. If in doubt, give new users the SEA User role until they need permission to perform administrative functions.

Access Control Groups

An access control group is a way to manage remote access for SEA users, allowing administrators to control who can access what and when. There are three types of access control groups:

- **Always Active:** users in this group can always access the assets defined in this group
- **Scheduled Access:** users in this group can access remote assets during the scheduled time window
- **Request Access:** users must explicitly request access to the remote assets, and the connection will only be granted after an access manager has approved the request

Note: The request access group also has an option for scheduled access to limit the time window in which a service can be requested.

It is recommended that access control groups are created using a schedule. Create groups in an “always active” state increases the risk to the network if an account was to be compromised.

SEA Plus Protocol Definitions

SEA Plus has the capability to allow any IP based protocol from traversing a remote device to a remote target. While this flexibility is key to enabling application from vendor laptops, it should not be deployed in an open configuration. SEA Plus does have out of the box definitions, but they are a starting point, not a recommended deployment model.

The three predefined definitions: **Allow all Protocols**, **TCP All Ports+ICMP**, and **UDP All Ports+ICMP** should be used with caution. Cisco does not recommend using them because they offer less protection. Once you are familiar with setting up the SEA Plus Definitions, Cisco recommends configuring your own protocols and ports to fit your needs and security requirements.

Protocol definitions should be defined on a per access basis. SEA Plus should only be used for applications that require it, such as RSLogix connecting to a PLC, and the protocol definition list should be restrictive to only what is needed for that particular application.

Cisco SEA Documentation

For the latest information on Cisco Secure Equipment Access, see the [documentation](#) hosted on DevNet.

Chapter 5. Cross-Domain Detection, Investigation and Response

The main purpose of this design guide is to provide the blueprint and best practices for securing critical infrastructure, and the operational technology is considered the most important part of the network. However, it is not always direct attacks to the OT networks that is the cause of operational impacts. According to the [2025 OT Cyber Threat Report by Water Security Solutions](#), which focuses on cyber-attacks which had physical consequences to the OT network, almost 90% of impacts were caused by indirect attacks to the OT networks. *“Indirect attacks are those where there was no direct impact on OT systems, but physical operations still suffered physical consequences because operations depended on access to or services from impaired IT systems.”*

An OT security strategy is not complete without coverage across both IT and OT networks. There does not necessarily have to be a convergence of these worlds, but context sharing and a unified view across the entire network is critical for protecting operations. In the same OT cyber threat report it was noted that the average time it takes between a ransomware criminal gaining remote control of some machine on a network and encryption activity starting was 16 hours – meaning half of attacks are even faster than that. It is essential security analysts find indicators of compromise as quickly as possible to reduce the impact of a cyber-attack, which may ultimately be an IT compromise that would have adverse effect on OT if left unnoticed.

Mean time to detect (MTTD) refers to the amount of time it usually takes for a security team to find and confirm that a security breach has happened. Lowering MTTD is more than just analysing an alert. Detecting a threat or ongoing attack involves looking at multiple interlocking systems and coordinating communication between teams in charge of those systems. For example, if a security analyst notices activity that looks like a port scan, a potential reconnaissance attack, the activity needs to be cross-referenced with the operational teams to see if the activity is part of regular operations. The challenge is coordinating detection quickly, while not losing accuracy.

Mean time to recovery (MTTR) then refers to the average measure of how long it takes security teams to stop or fix the problem after it has been identified and confirmed. The big questions for MTTR are what has been compromised and how is it fixed? This question poses several challenges, as the environment is typically so large that it’s difficult to see all the damage that was done.

- The initial access vector may have successfully been found, but where else did the attack go after initial entry?
- If a database was breached, it can be difficult to see exactly which tables were accessed
- The attack may be multipronged with the attacker breaching several points of entry

Each problem may require a different solution. Do the operational teams need to push a quick software update or fix a vulnerability? Does the network team need to block an IP address or a

port to prevent further attacks? Deciding on the right response is just as important as finding the breach in the first place.

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 IIoT 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.

Introducing Splunk

For the tenth consecutive time, Splunk was named a leader in the 2024 Gartner Magic Quadrant for SIEM. Splunk is a software platform primarily used for searching, monitoring, and analysing machine-generated data in real time. Thousands of organizations worldwide rely on Splunk Enterprise Security as their SIEM to rapidly detect and respond to critical events so they can stay ahead of emerging threats and maintain cyber resilience. With the ability to perform a wide range of security analytics and operations use cases, organizations can remain flexible and agile in the face of evolving threats and business needs. Splunk however, is not just one product. For example:

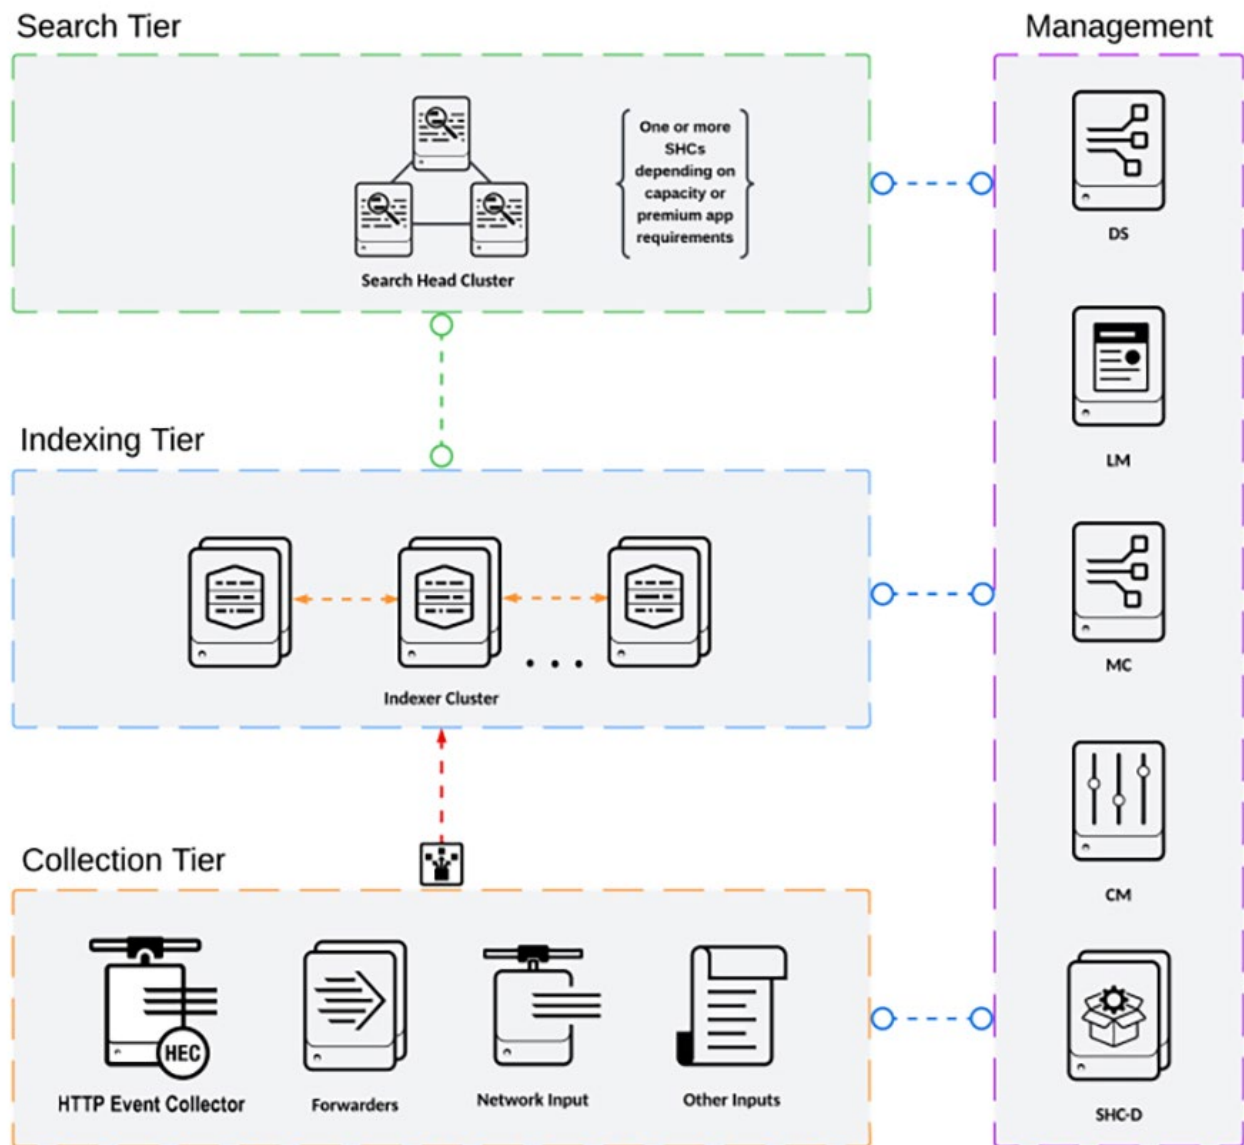
- **Splunk Enterprise** is the core data platform for collecting, indexing and analysing machine data. It can be installed on-prem or in a private cloud instance.
- **Splunk Cloud Platform** delivers the benefits of Splunk Enterprise as a cloud-based service.
- **Splunk SOAR** is used for automating security operations through playbooks and integrations. There is support for 300+ third part tools and supports 2800+ automated actions. It is available as a standalone version but can also be offered in a cloud version.
- **Splunk Observability Cloud** is a purpose-built suite of tightly integrated solutions that provide complete visibility across your data and systems for accurate, fast and in-context troubleshooting.

Splunk also provides products such as Splunk Enterprise Security, Splunk Asset & Risk Intelligence and Splunk Attack Analyzer which sit on top of platforms like Splunk Enterprise. The current iteration of this design guide only focuses on Splunk Enterprise, and some of the apps that reside live within the ecosystem. For further details on other products in the Splunk portfolio, see the [Splunk webpage](#).

Components of Splunk Enterprise

A Splunk Enterprise deployment consists of several key components that work together to collect, index, search and visualise machine data. Depending on the size and complexity of the environment, some of these components may be combined or separated for scalability and performance.

Figure 108 Splunk Enterprise components



Data Collection Tier

Forwarders and other data collection components

A Forwarder is a Splunk instance that forwards data to remote indexers for data processing and storage. There are two main types of forwarders:

- **Universal forwarder:** contains only the components that are necessary to forward data.
- **Heavy forwarder:** a full Splunk Enterprise instance that can index, search and change data as well as forward it. The heavy forwarder does have some features disabled to reduce system resource usage.

For more information on forwarders, see [Forwarding Data](#).

Splunk Connect for Syslog

Syslog serves as a universal protocol for transmitting log data across an enterprise. Its simplicity, extensibility, and compatibility with various devices and applications make it a key component of scalable and reliable enterprise data collection frameworks.

Many of the most common data sources that power Splunk product use cases require a syslog server for data collection. Most administrators do not possess the specific expertise required to successfully design, deploy and configure a syslog server to properly work with Splunk at scale. Additionally, the traditional Universal Forwarder or Heavy Forwarder approach to syslog collection have several issues with scale and complexity. Some customers send syslog events directly to Splunk to avoid architecting a syslog server, which introduces further problems. To help customers address these issues, Splunk Connect for Syslog (SC4S) was developed, a Splunk open-source community developed product.

Indexing Tier

Indexes and the Indexer

When data is added, Splunk software parses the data into individual events, extracts the timestamp, applies line-breaking rules and stores the events in an *index*. By default, data is stored in the “Main” index, but new indexes can be created for different data inputs. An indexer is the Splunk instance that indexes data. The indexer transforms the raw data into events and stores the events into an index. The indexer also searches the indexed data in response to search requests. The search peers are indexers that fulfil search requests from the search head.

For more information on indexers, see [Managing Indexers and Clusters of Indexers](#).

Search Tier

Search Head

Search is the primary way users interact with data in Splunk. Searches can be executed to retrieve events from an index, use statistical commands to calculate metrics and generate reports, search for specific conditions with a rolling time window, identify patterns in data, predict future trends, and so on. Searches can also be saved as reports and used to power dashboards.

The search head is a Splunk instance that directs search requests to a set of search peers (indexers) and merges the results back to the user. If the instance does only search, and not indexing, it is referred to as a dedicated search head.

Search Head Cluster

A search head cluster is a pool of at least three clustered Search Heads. It provides horizontal scalability for the search head tier and transparent user failover in case of outages. Search head clusters require dedicated servers of ideally identical system specifications.

Management Tier

Deployment Server

The deployment server is the tool for distributing configurations, apps, and content updates to groups of Splunk Enterprise instances. It is used to distribute updates to most types of Splunk Enterprise components – forwarders, non-clustered indexers, and search heads.

Deployment server is not a required component. For more information on the deployment server see [About deployment server and forwarder management](#).

License Manager

A license manager is required by other Splunk components to enable licensed features and track daily data ingest volume. The license manager role has minimal capacity and availability requirements and can be collocated with other management functions.

On the topic of licensing, customers can purchase a commercial end-user license for Splunk Enterprise based on either data volume or infrastructure. They provide access to the full set of Splunk Enterprise features within a defined limit of indexed data per day (volume-based license), or vCPU count (infrastructure license).

Splunk does have a free license tier, which gives very limited access to Splunk Enterprise features. The free license does not expire, and is for a standalone, single instance use only installation. The free license allows indexing of 500MB of data per day, and if this is exceeded a violation warning will be presented to the user. Search functionality will be disabled if there are several license violation warnings.

For more information in licensing, see [How Splunk Enterprise licensing works](#).

Monitoring Console

The monitoring console provides dashboards around usage and health monitoring of a Splunk environment. It contains several prepackaged platform alerts that can be customized to provide notifications for operational issues.

For more information on the monitoring console, see [About the Monitoring Console](#).

Cluster Manager

The cluster manager is the required coordinator for all activity in a clustered deployment. For example, an indexer cluster is a group of indexers configured to replicate each other's data, so that the system keeps multiple copies of all data. This process is known as index

replication. By maintaining multiple, identical copies of data, indexer clusters prevent data loss while promoting data availability for searching.

Splunk Enterprise clusters feature automatic failover from one indexer to the next. This means that, if one or more indexers fail, incoming data continues to get indexed and indexed data continues to be searchable.

For more information see [The basics of indexer cluster architecture](#).

Search Head Cluster Deployer

The search head cluster deployer (SHC-D) is needed to bootstrap a search head cluster and manage Splunk configurations deployed to the cluster. Each search head cluster requires its own SHC-D function. The SHC-D is not a runtime component and has minimal system requirements. The SHC-D is typically collocated with other management roles.

For more information see [Deploy a search head cluster](#).

Splunk Enterprise Concepts and Features

Search Processing Language

A Splunk search is a series of commands and arguments. Commands are chained together with a pipe “|” character to indicate that the output of one command feeds into the next command on the right.

```
search | command1 arguments1 | command2 arguments2 | ...
```

At the start of the search pipeline is an implied search command to retrieve events from the index. Search requests are written with keywords, quoted phrases, Boolean expressions, wildcards, field name/value pairs, and comparison expressions. The AND operator is implied between search terms. For example:

```
sourcetype=access_combined error | top 5 uri
```

This search retrieves indexed web activity events that contain the term “error.” For those events, it returns the top 5 most common URI values. Search commands are used to filter unwanted events, extract more information, calculate values, transform, and statistically analyse the indexed data. Think of the search results retrieved from the index as a dynamically created table. Each indexed event is a row. The field values are columns. Each search command redefines the shape of that table. For example, search commands that filter events will remove rows, search commands that extract fields will add columns.

Figure 109 Splunk search result, highlighting events as rows and the fields as columns

i	Time	Event
>	06/06/2025 06:35:39.871	<pre> { [-] CVSS: 7.8 CVSSVersion: 3 cve: CVE-2019-6858 cvssVectorString: CVSS:3.0/AV:L/AC:L/PR:N/UI:R/S:U/C:H/I:H/A:H fullDescription: id: b589d5fc-6888-4aa1-be53-4f80575d1ed6 links: [[+]] publishTime: -6795364515871 solution: This vulnerability is fixed in version 11.06.08 and is available for download below: https://schneider- electric.box.com/s/ilobctct4e4fgwwuk2l0kuvkn05mzd04v summary: A CWE-120:Buffer Copy without Checking Size of Input vulnerability exists which could cause a stack overflow when an attacker uses malicious input with excessive recursion. title: Vulnerability in Saitel DP (866e) & Saitel DR (HJc) vendorId: } Show as raw text host = 192.168.200.80 source = Vulnerabilities sourcetype = cisco:cybervision:vulnerabilities </pre>
>	06/06/2025 06:35:39.871	<pre> { [-] CVSS: 5.8 CVSSVersion: 2 cve: CVE-2014-2249 cvssVectorString: AV:N/AC:M/Au:N/C:N/I:P/A:P fullDescription: The web server of the affected PLCs (port 80/tcp and port 443/tcp) might allow CSRF(Cross-Site Request Forgery) attacks, compromising integrity and availability of the affected device. id: 8b0da72f-09aa-4ffa-85ef-7664ec5c9a3c links: [[+]] publishTime: -6795364494871 solution: Siemens provides firmware update V1.5.0 which fixes the potential vulnerabilities. As a general security measure Siemens strongly recommends to protect network access to S7-1500 CPUs with appropriate mechanisms. It is advised to follow recommended security practices and to configure the environment according to operational guidelines in order to run the devices in a protected IT environment. summary: In the SIMATIC S7-1500 CPU firmware multiple vulnerabilities were discovered. The vulnerabilities allowed attackers to perform Denial of Service attacks with specially crafted HTTP(S), ISO-TSAP, or Profinet network packets. The integrated web server was also vulnerable to Cross-Site Request Forgery, Cross-Site Scripting, header injection, and open redirect attacks as well as privilege escalation. The vulnerabilities might have been exploited over the network without authentication. title: Vulnerabilities in SIMATIC S7-1500 CPU 1 & SIMATIC S7-1200 CPU 1 vendorId: SSA-456423 } Show as raw text host = 192.168.200.80 source = Vulnerabilities sourcetype = cisco:cybervision:vulnerabilities </pre>

Fields

Fields appear in event data as searchable name-value pairing such as `user_name=alice` or `ip_address=10.0.0.1`. Fields are the building blocks of Splunk searches, reports, and data models. When a search is run against event data, Splunk software looks for fields in that data. Look at the following example search.

```
Status=404
```

The search finds events with status fields that have a value of 404. When you run this search, Splunk Enterprise does not look for events with any other status value. It also does not look for events containing other fields that share 404 as a value. As a result, this search returns a set of results that are more focused than you get if you used 404 in the search string.

Fields often appear in events as key=value pairs such as `user_name=alice`. But in many events, field values appear in fixed, delimited positions without modifying keys. For example, events might have the `user_name` value always appear by itself after the timestamp and the `user_id` value.

```
Nov 15 09:32:22 00224 johnz
```

Nov 15 09:39:12 01671 dmehta

Nov 15 09:45:23 00043 sting

Nov 15 10:02:54 00676 lscott

Splunk Enterprise can identify these fields using custom field extraction. Splunk automatically extracts host, source, and sourcetype values, timestamps, and several other default fields when it indexes incoming events. It also extracts fields that appear in event data as key=value pairs. Custom field extraction is available in Splunk Enterprise.

For more information on fields and field extractions, see the [Knowledge Manager Manual](#).

Subsearches

A subsearch runs its own search and returns the results to the parent command as the argument value. The subsearch is run first and is contained in square brackets. For example, the following search uses a subsearch to find all syslog events from the user that had the last login error:

```
sourcetype=syslog [ search login error | return 1 user ]
```

Optimizing Searches

The key to fast searching is to limit the data that needs to be pulled off disk to an absolute minimum. Then filter that data as early as possible in the search so that processing is done on the minimum data necessary.

- Partition data into separate indexes, if you will rarely perform searches across multiple types of data. For example, put web data in one index, and firewall data in another.
- Limit the time range to only what is needed. For example -1h not -1w, or earliest=-1d.
- Search as specifically as you can. For example, fatal_error not *error*
- Use post-processing searches in dashboards.
- Use summary indexing, and report and data model acceleration features.

Common Search Commands

Command	Description
chart/timechart	Returns results in a tabular output
dedup	Removes subsequent results that match a specified criterion
eval	Calculates an expression
fields	Removes fields from search results
head/tail	Returns the first/last N results
lookup	Adds field values from an external source
rename	Renames a field. Use wildcards to specify multiple fields
rex	Specifies regular expression named groups to extract fields
search	Filters results to those that match the search expression
sort	Sorts the search results by the specified fields

stats	Provides statistics, grouped optionally by fields
mstats	Similar to stats but used on metrics instead of events
table	Specified fields to keep in the result set. Retains data in tabular format
top/rare	Displays the most/least common values of a field
transaction	Groups search results into transactions
where	Filters search results using eval expressions. Used to compare two different fields

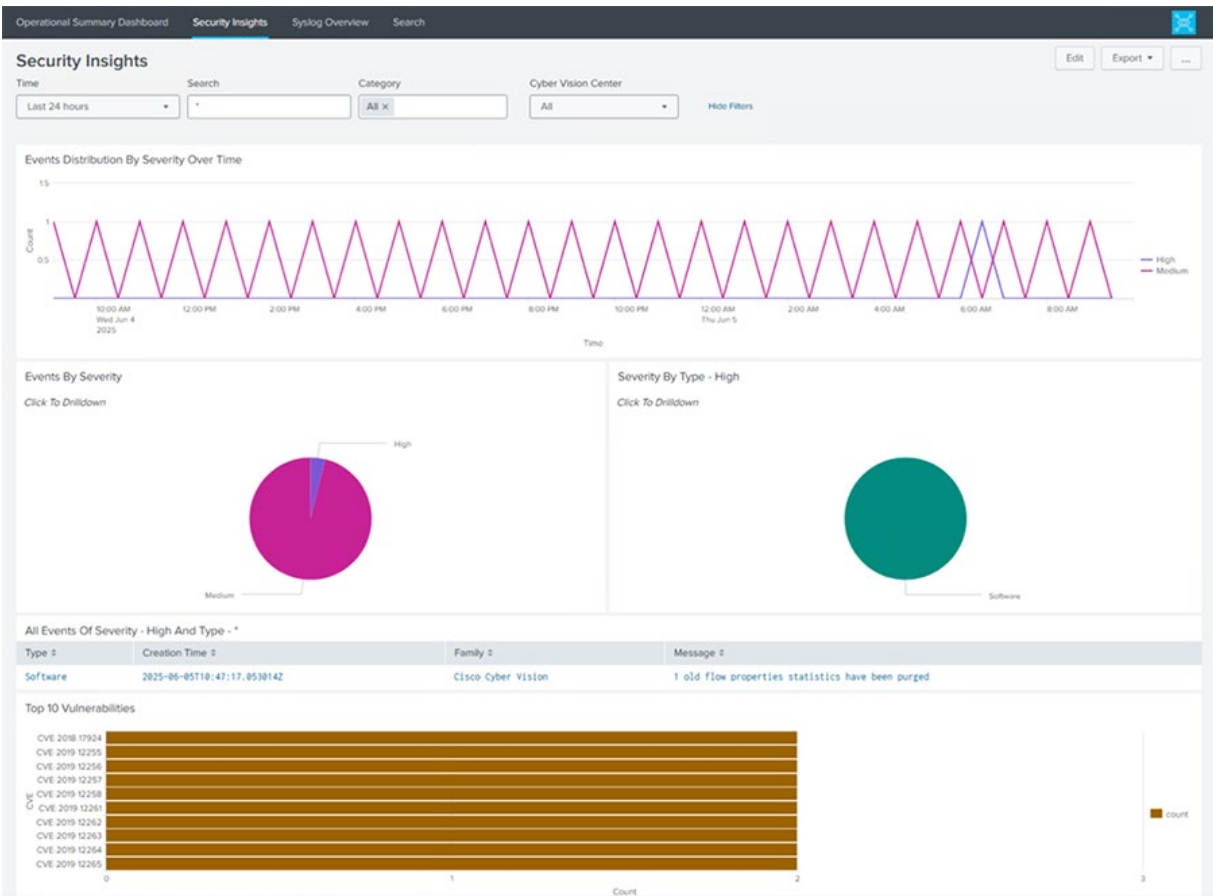
Note: it is not the purpose of this design guide to act as a training guide for creating SPL queries. Example queries will be used later in the document when providing design guidance for the deployment of Cisco specific apps, but for a more complete overview, including a tutorial, see Splunk's [Search Manual](#).

Apps

Splunk Apps

Splunk Applications are comprehensive solutions that enhance the Splunk platform by providing pre-built dashboards, reports, workflows, and visualizations designed for specific use cases or industries. These apps cater to a wide range of scenarios, such as IT operations, security monitoring, and business analytics, enabling users to derive actionable insights quickly without the need to create custom configurations. By offering ready-to-use functionality, Splunk Apps help organizations streamline their workflows and maximize the value of their data.

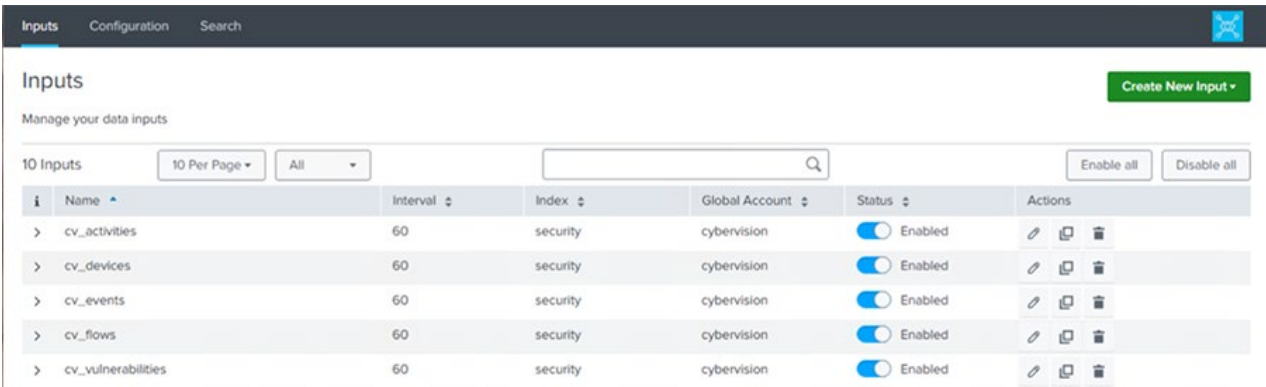
Figure 110 Security insights dashboard within the Cisco Cyber Vision Splunk app



Splunk Add-ons

Splunk Add-ons are lightweight extensions that focus on data integration and enrichment. They enable the collection, parsing, and normalization of data from various sources, such as servers, devices, cloud services, or third-party applications. Add-ons ensure that raw data is properly formatted and enriched, making it ready for analysis within the Splunk platform. While they do not include visualizations or dashboards, add-ons are essential for ensuring seamless data ingestion and compatibility with Splunk's analytics capabilities.

Figure 111 Cyber Vision Add-on for Splunk input configuration menu



Splunk Add-ons are required for getting data into the Splunk platform from many data sources, while the Splunk apps are optional, and provide default visualisations utilising the data extracted by the add-on.

Alerts, Dashboard & Reports

Alerts provide notifications when search results for both historical and real-time searches meet configured conditions. Alerts can be configured to trigger actions like sending alert information to designated email addresses, posting alert information to an RSS feed, and running a custom script, such as one that posts an alert event to syslog.

Figure 112 Splunk Alerts and Report creation dashboard

Searches, Reports, and Alerts

Searches, reports, and alerts are saved searches created from pivot or the search page. [Learn more](#)

4 Searches, Reports, and Alerts

Type: All

App: Search & Reporting (search)

Owner: Administrator (cisco)

filter

10 per page

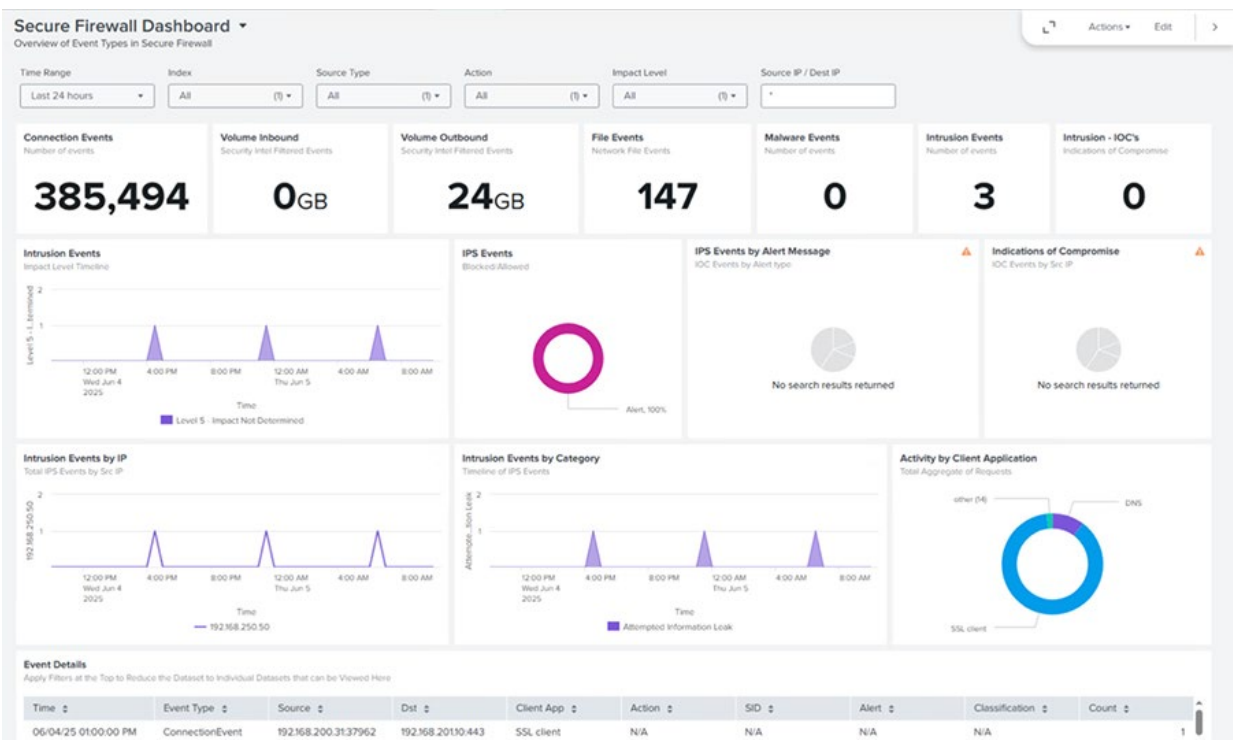
Name	Actions	Type	Next Scheduled Time	Display View	Owner	App	Alerts	Sharing	Status
<div>FTD File Event</div> <div>FTD File Event</div>	<div>Edit</div> <div>Run</div> <div>View Recent</div>	Alert	2025-06-05 10:15:00 EDT	none	cisco	search	69	Private	Enabled
<div>Unwanted Vendors Alert</div> <div>Unwanted Vendors Alert</div>	<div>Edit</div> <div>Run</div> <div>View Recent</div>	Alert	2025-06-06 00:00:00 EDT	none	cisco	search	23	Global	Enabled
<div>Vulnerability Alert</div> <div>Vulnerability Alert</div>	<div>Edit</div> <div>Run</div> <div>View Recent</div>	Alert	2025-06-09 06:00:00 EDT	none	cisco	search	0	Global	Enabled
<div>Vulnerable Device with Intrusion Attempts</div> <div>Vulnerable device with Intrusion Attempts</div>	<div>Edit</div> <div>Run</div> <div>View Recent</div>	Alert	2025-06-05 10:15:00 EDT	none	cisco	search	0	Global	Enabled

Figure 113 Triggered alerts on the Splunk alerts dashboard

24 Triggered Alerts			
Enable or Disable Manage triggered alerts Filter by Last: 24 Hours 5 per page			
Alert Name	Instance	Time Triggered	Description
Unwanted Vendors Alert		Jun. 05, 2025 9:00 AM	Unwanted Vendors Alert
Unwanted Vendors Alert		Jun. 05, 2025 8:00 AM	Unwanted Vendors Alert


Dashboards contain panels of modules like search boxes, fields, charts, and so on. Dashboard panels are usually connected to saved searches or pivots. They display the results of completed searches and data from real-time searches that run in the background.

Figure 114 Cisco Secure Firewall dashboard within the Cisco Security Cloud Splunk App



Splunk Enterprise allows you to save searches and pivots as reports and then add reports to dashboards as dashboard panels. Run reports on an ad hoc basis, schedule them to run on a regular interval, or set a scheduled report to generate alerts when the result meets conditions.

Figure 115 Sample report highlighting privileged account usage over time



Alerts, Dashboard and Reports will be explored in depth later in the design guide.

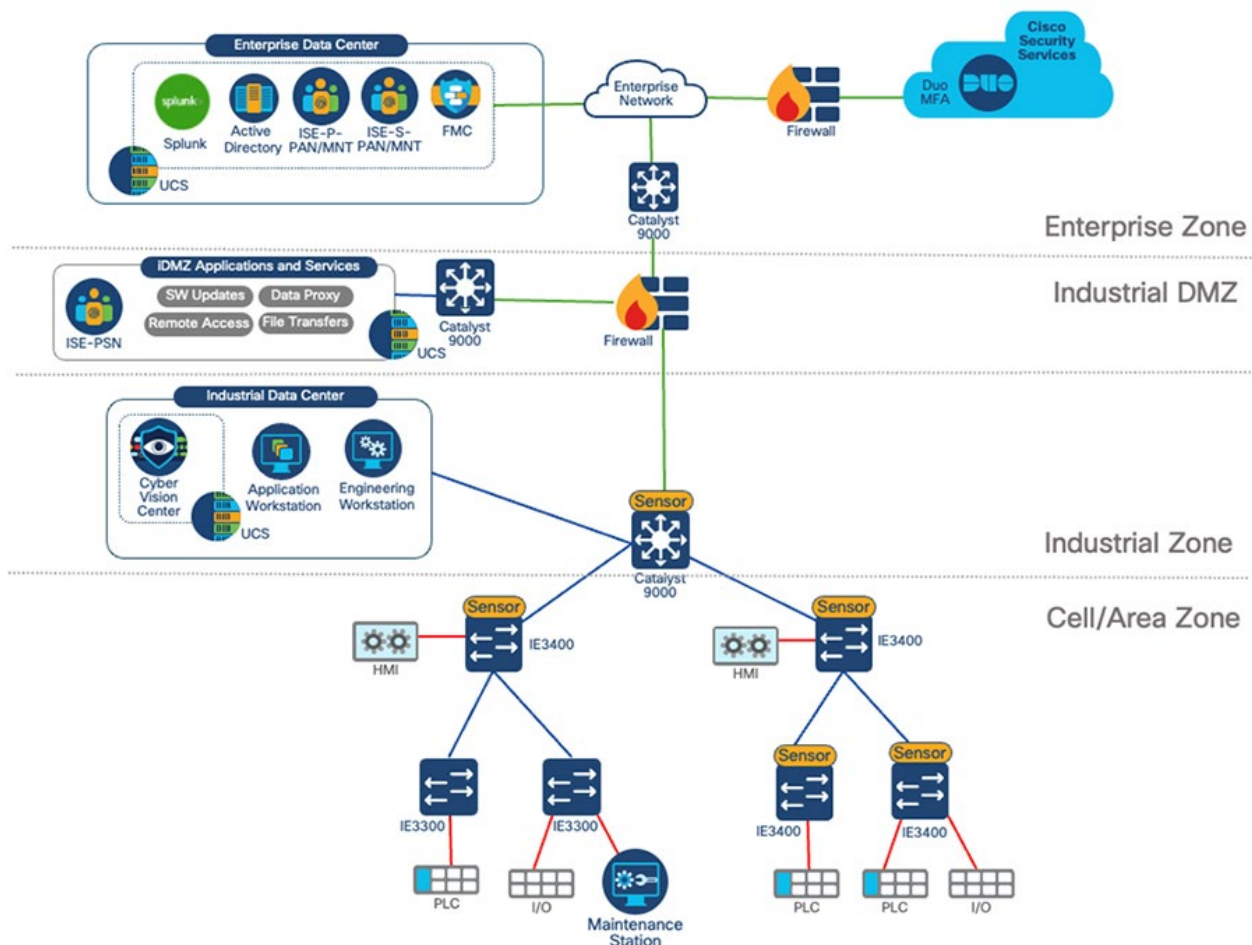
Splunk Design Considerations

Install Considerations

The simplest deployment is the one you get by default when you first install Splunk Enterprise on a machine: a standalone instance that handles both indexing and searching. You log into Splunk Web or the CLI on the instance and configure data inputs to collect machine data. You then use the same instance to search, monitor, alert, and report on the incoming data.

However, a single instance is often not the recommended approach for load balancing and resilience across a production environment and is therefore recommended to install specialized instances of Splunk Enterprise on multiple machines. The architecture chosen for this design guide was the single server deployment as the focus of this design guide is data onboarding and extracting value from that data using Splunk.

Figure 116 Architecture used during validation testing for Splunk in OT environments



To understand which deployment model is best suited for your environment, Splunk has curated a [Topology selection guidance](#) questionnaire. The design details of the chosen deployment can also be found in the same documentation.

Getting Data In

Splunk can ingest almost any type of data, by either through a pull model, or listening for devices and applications to push data into it. The design guide will detail how data was ingested for the following services:

- Cisco Cyber Vision
- Cisco IOS-XE
- Cisco ISE
- Cisco Firewall Management Center
- Cisco Duo

Cisco Cyber Vision Integration

Note: For installation instructions see [Appendix F](#).

As mentioned previously in the guide, the Cisco Cyber Vision integration has two components – the *Cyber Vision Splunk Add On* and the *Cyber Vision Splunk App*. The Add On is the component responsible for ingesting data into Splunk, the App is a Cisco created visualisation that uses data ingested from the Add On. If the intention is to create custom visualisation from Cyber Vision data, only the Add On is required.

The Cyber Vision Add On ingests data in two ways – a pull model, using the Cyber Vision APIs, and a push model, listening to Syslog events sent directly by Cyber Vision.

Required Ports:

- **API:** TCP 443
- **Syslog:** UDP/TCP 514 (customizable)

The intention for Syslog generated events, is to be notified in real time when an event does happen. For example, if Cyber Vision was to trigger an event from a Snort rule, a Syslog message would be created and a SOC analyst could be immediately notified of its presence. However, by default the Cyber Vision Add On for Splunk polls data from Cyber Vision every 60 seconds, making Syslog unnecessary as events are also retrieved via API. If API ingestion time was to be adjusted, Syslog would be recommended to complement it.

When connecting Cyber Vision to Splunk, SOC analysts have access to the following data sources:

- Devices
- Activities
- Flows
- Vulnerabilities
- Events

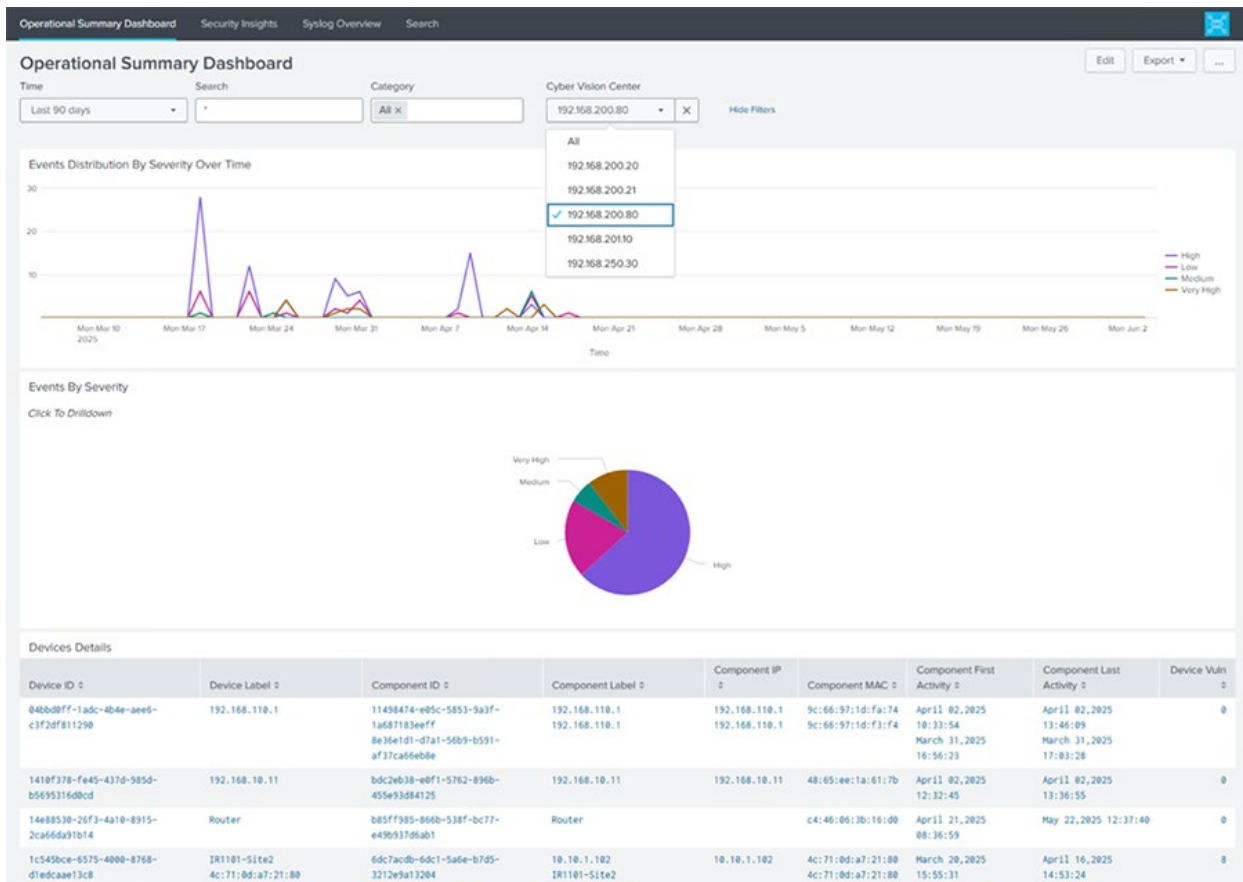
Using Splunk to aggregate data across the entire Cyber Vision deployment

Cyber Vision is recommended to be installed locally at an OT facility. Therefore, for organizations who are managing many facilities, it is common for multiple Cyber Vision Centers to be deployed. This architecture results in scattered event logs, as each instance would need to be looked at individually to get security insights across a distributed deployment. When

integrating Cyber Vision into Splunk, all of the Cyber Vision Centers can be aggregated into a single app, meaning a single pane of glass can be used to visualise and query data.

The Cyber Vision Splunk App provides a default visualisation for this aggregated view, where data can be looked at individually, focusing on one site at a time, or collectively, where an organizations security posture can be viewed globally.

Figure 117 Cyber Vision data can either be aggregated across all Cyber Vision Centers, or data can be filtered on a per site basis



When deploying a distributed Cyber Vision architecture, it is recommended to take advantage of Splunk as the global collector for OT visibility data.

Cisco Networks App for IOS-XE data collection

Note: For installation instructions see there is a help page in the application after download.

The Cisco Networks App for Splunk Enterprise includes dashboards, data models and logic for analyzing data from Cisco Switches & Routers (Cisco IOS, IOS XE, IOS XR and NX-OS devices), WLAN Controllers and Access Points.

Required Ports:

- **Syslog:** UDP/TCP 514 (customizable)

The Cisco Networks app is designed to give insights to both network and security operators across switching, routing and wireless deployments. With IOS-XE being a shared operating

system across both enterprise and rugged deployments, this app is not limited to OT environments, however the scope of what it was used for in this design guide is. The app provides visibility into events such as port flapping, performance issues, and time drift of network equipment. However, from a security perspective, the app gives analysts insights into access control lists and authentication requests.

When using TrustSec in the OT network, syslog messages are generated when there is an SGACL match on the switch. Splunk provides users the ability to take these SGACL events and create their own custom dashboards for monitoring the security events across the OT network.

Cisco Identity Services Engine App

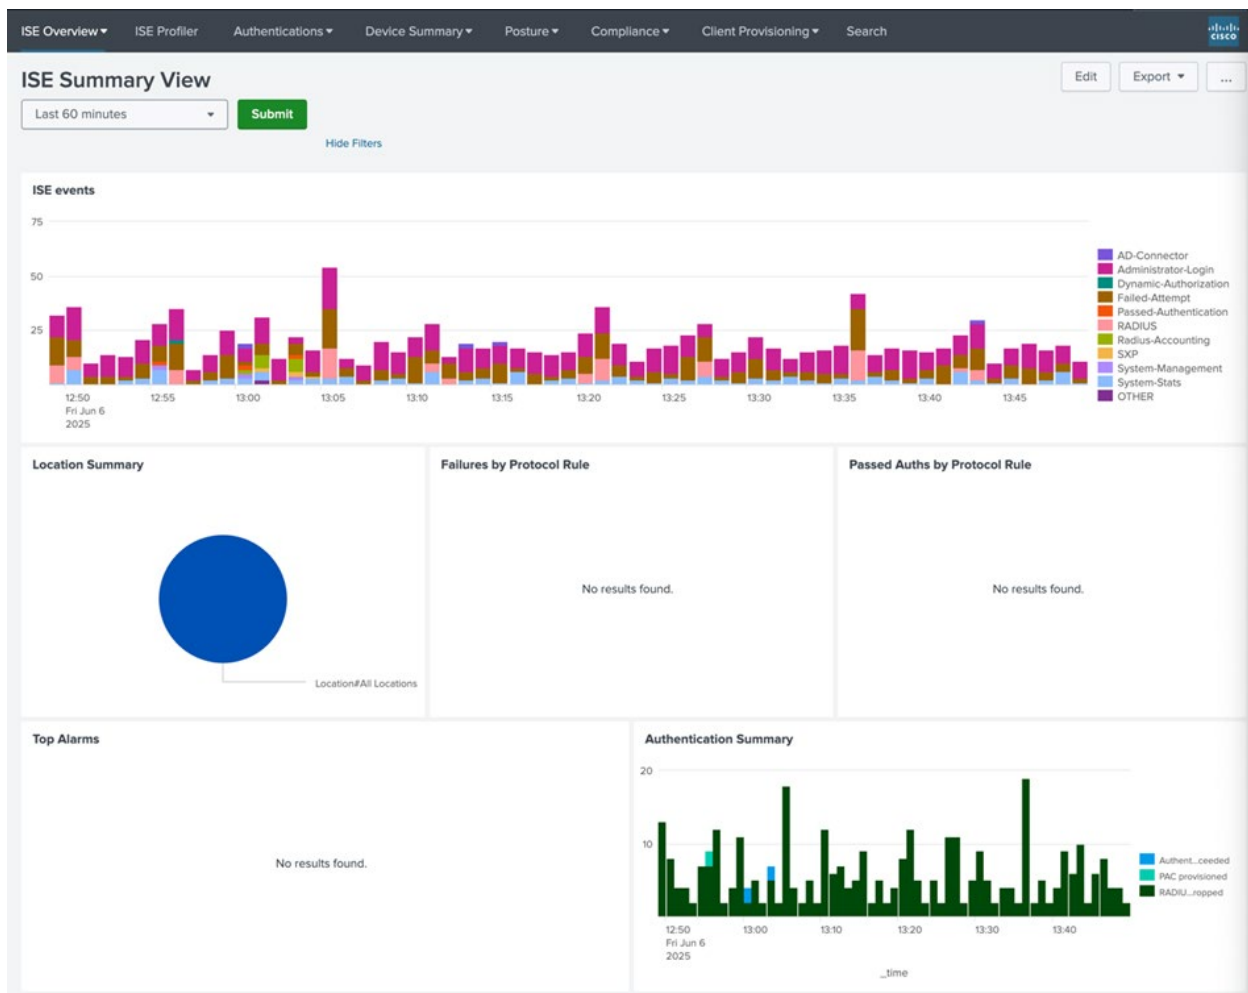
Note: For installation instructions see [Appendix G](#).

The Splunk for Cisco ISE add-on allows for the extraction and indexing of the ISE AAA Audit, Accounting, Posture, Client Provisioning Audit and Profiler events. This integration allows any Splunk user to correlate ISE data with other data sources (e.g. with firewall events or application data) to get deeper operational and security visibility. It also includes sample dashboards and reports for profiling, authentication, system statistics, alarms, and location awareness.

Required Ports:

- **API:** TCP 443
- **Syslog:** UDP/TCP 1468 (customizable)

Figure 118 Summary view in the ISE app for Splunk



Cisco Security Cloud App

Cisco Security Cloud App is a web application that offers a centralized platform to integrate Cisco security products with Splunk. At the time of writing this guide, the following applications have been included:

- Cisco Duo
- Cisco Secure Malware Analytics
- Cisco Secure Firewall Management Center
- Cisco Multicloud Defense
- Cisco XDR
- Cisco Secure Email Threat Defense
- Cisco Secure Network Analytics
- Cisco Secure Endpoint
- Cisco Vulnerability Intelligence

Security Cloud App combines the concept of Splunk Technical Add-ons and Splunk app into a single offering, providing a comprehensive solution for monitoring and analysis. Cisco Security Cloud App has built-in health checks, and constant monitoring to ensure operational integrity.

The rich user interface renders a cohesive user experience that provides detailed instructions for each Cisco product to help facilitate the setup process. For installation instructions, see the [Cisco Security Cloud App for Splunk User Guide](#).

*Note: For this version of the design guide, **Cisco Secure Firewall Management Center (FMC)** and **Cisco Duo** were used as other apps on this list were not used earlier on in this design guidance. If other apps in the list are introduced later in guidance for security industrial networks they will be added to this section too.*

Cisco Secure Firewall Management Center

When integrating Cisco FMC and Cisco Secure Firewall with Splunk for log ingestion, it is crucial to understand the capabilities and limitations of the available integration methods: eStreamer, Syslog, and API. Below is a detailed comparison and design guidance to help you decide which method(s) to use based on your requirements.

eStreamer

- The Secure Firewall System Event Streamer (eStreamer) uses a message-oriented protocol to stream events and host profile information to Splunk
- Data provided by eStreamer includes:
 - Intrusion events
 - Connection events
 - File and malware events
 - Security intelligence events
 - User activity (if identity policies are configured)
 - Host profile data
- eStreamer gets updated on an interval, so will not handle real time log ingestion
- Required Port: TCP 8302

Syslog

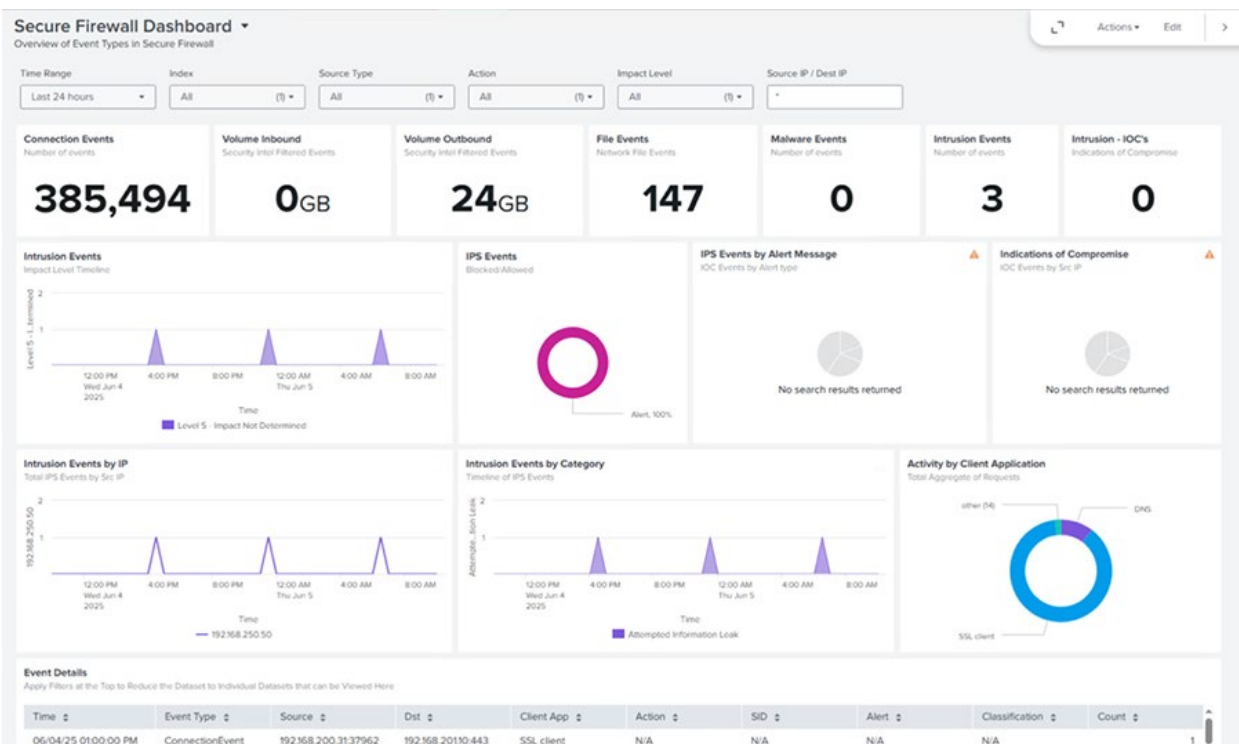
- Syslog is a standard protocol that the firewall uses to forward logs in real time
- The following data is sent over syslog as they occur:
 - Connections events
 - Intrusion events
 - Security intelligence events
 - Audit logs
- Syslog is easy to configure, and enables Splunk to react to firewall logs in real time. It is recommended to pair syslog with eStreamer for complete coverage.
- High-volume environments may require careful running to avoid performance issues.
- Required port: UCP/TCP 1025 (customizable)

API

- Splunk can also use the RESTful API provided by Cisco FMC for querying specific data.
- The following data is available over API:
 - Configuration data
 - Policy and rule definition
 - Event data
- Port required: TCP 443

The recommendation is to install eStreamer as a minimum for its depth in data. Install syslog if real-time monitoring is required, and use the API access method to configuration data or rules are required.

Figure 119 Cisco Secure Firewall Dashboard within the Cisco Security Cloud App for Splunk



Cisco Duo

In this design guide, Cisco Duo was integrated into Cisco Secure Equipment Access for an added layer of protection for remote access into OT networks. There is only one method for integrating Duo into Splunk, and that is through an API connector.

Required Ports:

- **API:** TCP 443

When connecting Duo to Splunk, SOC analysts have access to the following events:

- Account logs
- Activity logs

- Administrator logs
- Authentication logs
- Telephony logs
- Trust monitor logs
- Endpoint logs

The Duo logs enable analysts to get insights into which countries users are attempting to log in from, what device they were using, or if they were failing their MFA prompt. Singular Duo logs may not mean anything in isolation, but when a burst of events come in, it could signal a problem. For example, a legitimate user may accidentally fail an MFA prompt, or maybe they are logging in from a new PC. However, if the same user is logging in from multiple countries in the space of minutes, or they have repeated failed authentication attempts, it may be a sign that the account has been compromised. Creating alerts and repetitive data will be covered later in the guide.

Note: When doing validation testing, the Duo functionality within the Cisco Security Cloud app was not working. To test Duo, the standalone app was used. This standalone app will be deprecated soon, and the combined app should be fixed. Reach out to your Cisco representative for a date of completion for the Duo fix if it still broken at the time of reading this.

Index Design and Life-Cycle Management Recommendations

Splunk Enterprise stores indexed data in buckets, which are directories containing both the data and index files into the data. An index typically consists of many buckets, organized by age of the data. An index typically consists of many buckets, and the number of buckets grows as the index grows. As data continues to enter the system, the indexer creates new buckets to accommodate the increase in data. As a bucket ages, it rolls through stages.

Bucket State	Description	Searchable?
Hot	New data is written to hot buckets. Each index has one or more hot buckets.	Yes
Warm	Buckets rolled from hot. New data is not written to warm buckets. An index has many warm buckets.	Yes
Cold	Buckets rolled from warm and moved to a different location. An index has many cold buckets.	Yes
Frozen	Buckets rolled from cold. The indexer deletes frozen buckets, but you can choose to archive them first. Archived buckets can later be thawed.	No
Thawed	Buckets restored from an archive. If you archive frozen buckets, you can later return them to the index by thawing them.	Yes

The default scheme for bucket handling should be fine for most users, but if you are indexing large amounts of data, have specific data retention requirements, or otherwise need to carefully

plan your aging policy, it is recommended to explore the detailed information at [How the indexer stores indexes](#).

Deployment Considerations for Splunk Indexes

1. Number of Indexes

It is best practise to create separate indexes for separate kinds of data sources. Logs with different retention periods should be in separate indexes (e.g., compliance logs vs. ephemeral app logs, security logs vs. generic networking logs). This helps with the performance of the queries by keeping the dataset smaller.

2. Index Size

To determine the size of the index you need to calculate the amount of data that you are going to ingest from that particular data source and the time period for which the data needs to be stored. Plan storage based on: Daily ingestion volume, Retention policies and Data compression ratios (typically 50-80% savings in Splunk).

3. Maximum Index Size

To set the maximum index size on a per-index basis, use the *maxTotalDataSizeMB* attribute in the *indexes.conf* file or on the Splunk Index creation UI. When this limit is reached, buckets begin rolling to frozen.

4. Retention

The size and age of event data is enforced on an index-by-index basis. One can granularly control the retention policy for the different kinds of data based on compliance regulations of how long the data needs to be stored if we store the data in different indexes. Define retention policies for each index based on business or compliance needs.

Use the *frozenTimePeriodInSecs* setting in the *indexes.conf* file to define how long data should be retained before being moved to the frozen state (archived or deleted).

Example: *frozenTimePeriodInSecs* = 2592000 (30 days).

Create separate indexes for data with different retention requirements (e.g., logs with a 30-day retention vs. compliance data with a 7-year retention).

5. Set Index Size Limits

Use the *maxTotalDataSizeMB* setting in *indexes.conf* to limit the total size of each index. This ensures indexes do not grow uncontrollably and consume all available disk space.

Example: *maxTotalDataSizeMB* = 50000 (50 GB).

Splunk will automatically delete the oldest data from cold buckets when the size limit is reached.

6. Efficiency/Speed

The more isolated the data is, the fewer false positives and less data needs to be sorted through so the more efficient and faster searches will be.

7. Data Access Needs

Place high-frequency search data in dedicated indexes to improve search performance.

8. Data Volume Management

Use index settings *maxTotalDataSizeMB* to limit the size of an index to prevent disk over-utilization. Monitor the growth of indexes and adjust their configurations as needed.

9. Use Multiple Indexers

For large-scale deployments, distribute indexes across multiple indexers for better load balancing, storage management, and search performance. Use indexer clustering for high availability and disaster recovery.

10. Index Replication and Clustering

For high availability, configure indexer clustering with replication policies to prevent data loss. Use the *replicationFactor* setting to ensure multiple copies of index data exist.

11. Access Control

Leverage Splunk's Role-Based Access Control (RBAC) to restrict access to indexes based on user roles. Assign specific roles to sensitive data indexes (e.g., only security teams can access `security_events`).

12. Use Metrics Indexes for Metric Data

Store metrics data (e.g., CPU, memory, application performance) in metrics indexes instead of event indexes for better performance and reduced storage requirements.

13. Monitoring and Maintenance

Regularly monitor index health and usage using Splunk's Monitoring Console. Identify and address potential issues such as bucket corruption, indexer disk space shortages, or excessive index growth. Perform periodic index cleanup to remove old or unnecessary data.

Use the Splunk Monitoring Console or custom dashboards to monitor:

- Daily data ingestion volume.
- Index size trends over time.
- Disk space usage on indexers.

Set up alerts to notify administrators if disk usage exceeds a threshold (e.g., 80% utilization).

14. Archive or Delete Frozen Data

Configure the *coldToFrozenDir* setting to archive frozen data to cheaper storage (e.g., AWS S3, HDFS, or on-premises storage).

If archival is not required, let Splunk delete frozen data automatically using *coldToFrozenScript* or default deletion behavior.

By following these best practices and considerations, you can ensure that your Splunk indexes are well-optimized for performance, scalability, and manageability while meeting business and compliance requirements.

Splunk Indexes used in this design guide

Data Input From	Index Name
Cybervision REST-API	cisco_cybervision
Cybervision Syslog	
Cisco ISE Syslog	cisco_ise_syslog
Cisco FMC/FTD	cisco_secure_fw
Cisco FTD Syslog	cisco_sfw_ftd_syslog
Cisco FMC API	cisco_sfw_api
IOS-XE Switches	cisco_catalyst
IOS-XE Routers	cisco_catalyst
Cisco Duo	cisco_duo

Using the Search & Reporting app

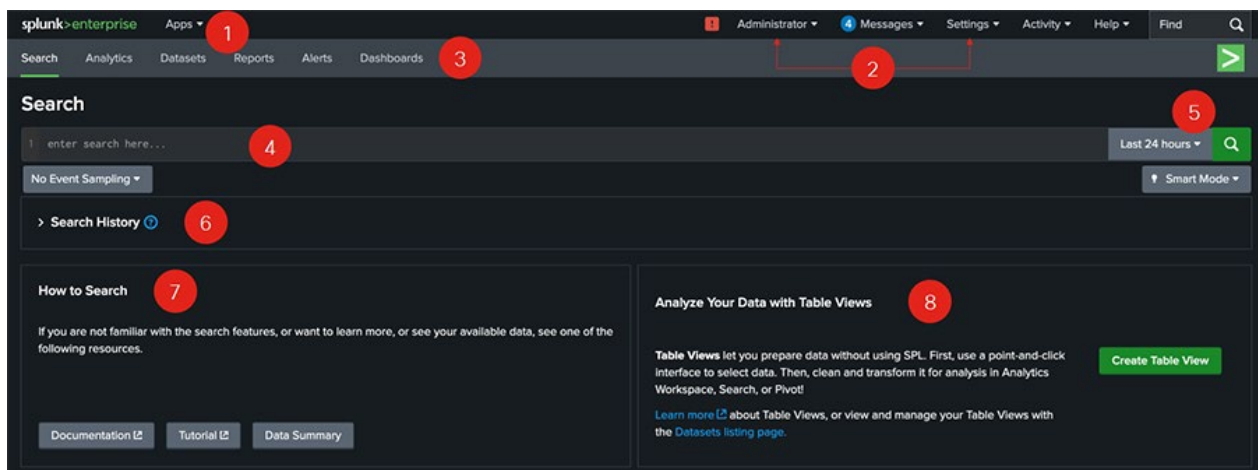
The Search app, the short name for the Search & Reporting app, is the primary way you navigate the data in your Splunk deployment. The Search app consists of a web-based interface (Splunk Web), a command line interface (CLI), and the Splunk SPL.

To open the Search app, from Splunk Home click **Search & Reporting** in the **Apps** panel. This opens the Search Summary view in the Search & Reporting app.

Before you run a search, the Search summary view displays the following elements: App bar, Search bar, Time range picker, **How to Search** panel, **Search History** panel, and the **Analyze Your Data with Table Views** panel.

Some of these are common elements that you see on other views. Elements that are unique to the Search Summary view are the panels below the Search bar.

Figure 120 Splunk Search page



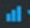
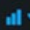
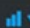



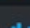




1	Applications menu	Switch between Splunk applications that you have installed. The current application, Search & Reporting app, is listed. This menu is on the Splunk bar.
2	Splunk bar	Edit your Splunk configuration, view system-level messages, and get help on using the product.
3	Apps bar	Navigate between the different views in the application you are in. For the Search & Reporting app the views are: Search, Datasets, Reports, Alerts, and Dashboards.
4	Search bar	Specify your search criteria.
5	Time range picker	Specify the time period for the search, such as the last 30 minutes or yesterday. The default is Last 24 hours .
6	Search history	View a list of the searches that you have run. The search history appears after you run your first search.
7	How to Search	Use the links to learn more about how to start searching your data, as well a summary of the data that you have access to.
8	Analyze Your Data with Table Views	Create curated collections of event data into datasets that you design for a specific business purpose.

The Data Summary dialog box shows three tabs: Hosts, Sources, Sourcetypes. These tabs represent searchable fields in your data. Selecting a host, source, or source type from the Data Summary dialog box is a great way to see how your data is turned into events.

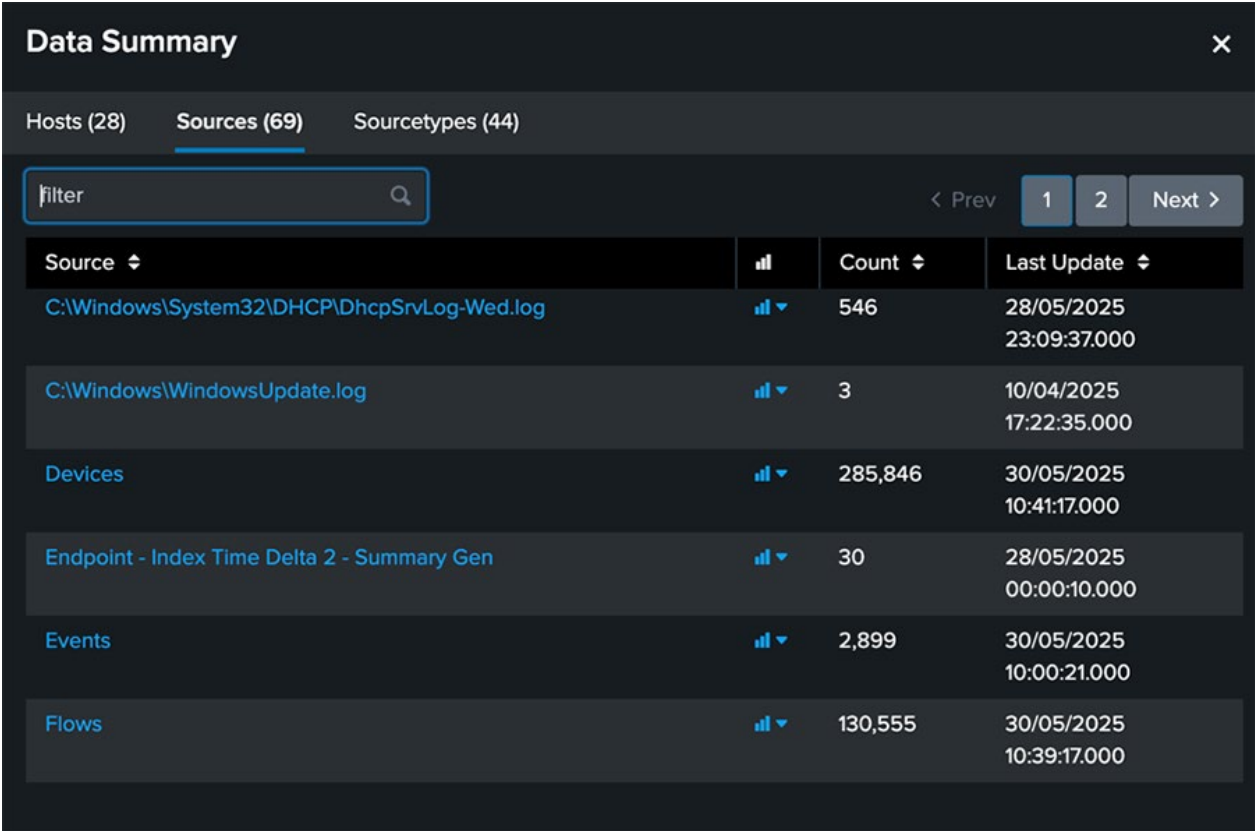
The **host** of an event is the host name, IP address, or fully qualified domain name of the network machine from which the event originated. In a distributed environment, you can use the host field to search data from specific machines.

Figure 121 Example data summary page in the Search & Reporting app showing the hosts where data originates

Data Summary ×			
Hosts (28) Sources (69) Sourcetypes (44)			
<div>filter 🔍</div>			
Host ⬇		Count ⬇	Last Update ⬇
192.168.201.2		419,882	30/05/2025 08:27:38.000
192.168.201.3		84,131	30/05/2025 10:41:33.000
192.168.201.4		127	12/05/2025 07:22:43.000
192.168.201.5		17	09/05/2025 11:38:34.000
192.168.201.6		13,930	30/05/2025 10:41:48.000
192.168.201.7		280	30/05/2025 10:17:26.000
192.168.204.5		25,817	20/10/2024 16:26:02.000
192.168.250.30		53,079	30/05/2025 10:40:24.000
IS-AD		5,970,661	30/05/2025 10:42:17.000
duo_api		24,306	30/05/2025 10:41:15.000

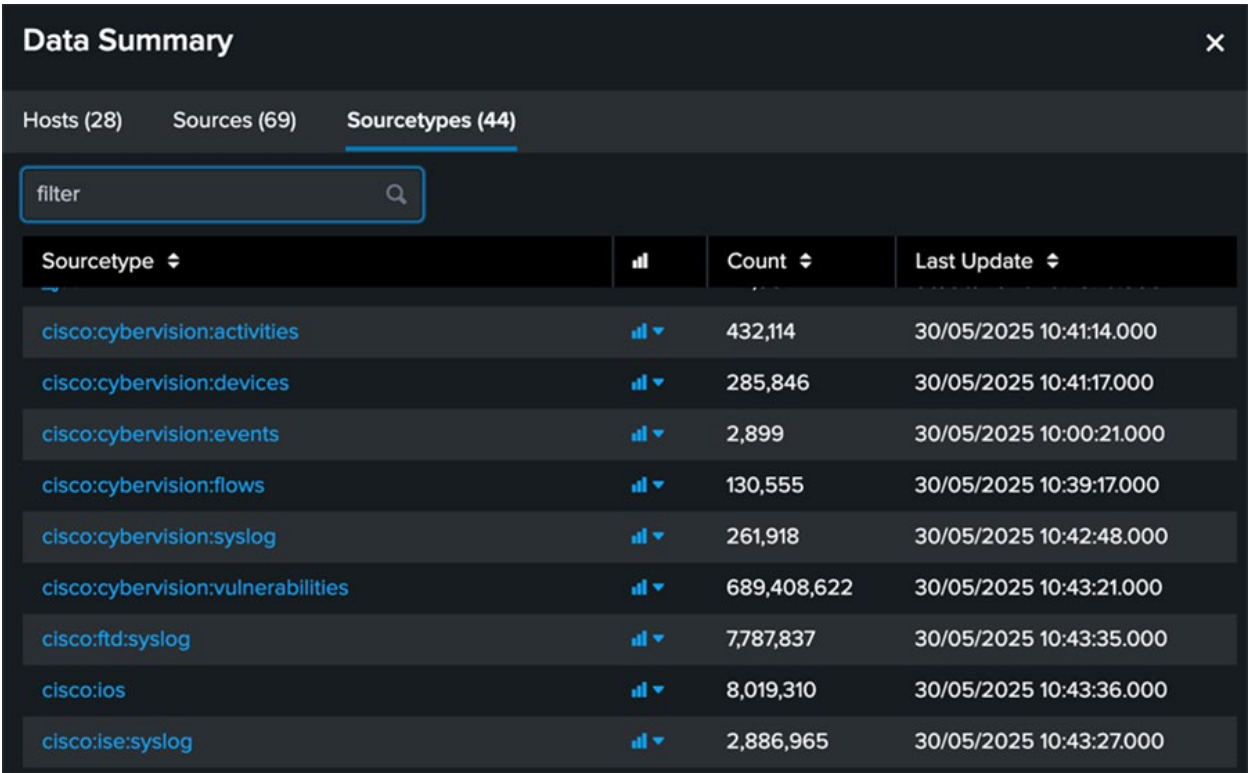
The **source** of an event is the file or directory path, network port, or script from which the event originated.

Figure 122 Example data summary page in the Search & Reporting app showing the data sources available for creating queries



The **source type** of an event tells you what kind of data it is, usually based on how the data is formatted. This classification lets you search for the same type of data across multiple sources and hosts. For example, searching for `cisco:cybervision:events` would provide Cyber Vision events across all of the Cyber Vision nodes in the environment.

Figure 123 Example data summary page in the Search & Reporting app showing the sourcetypes available

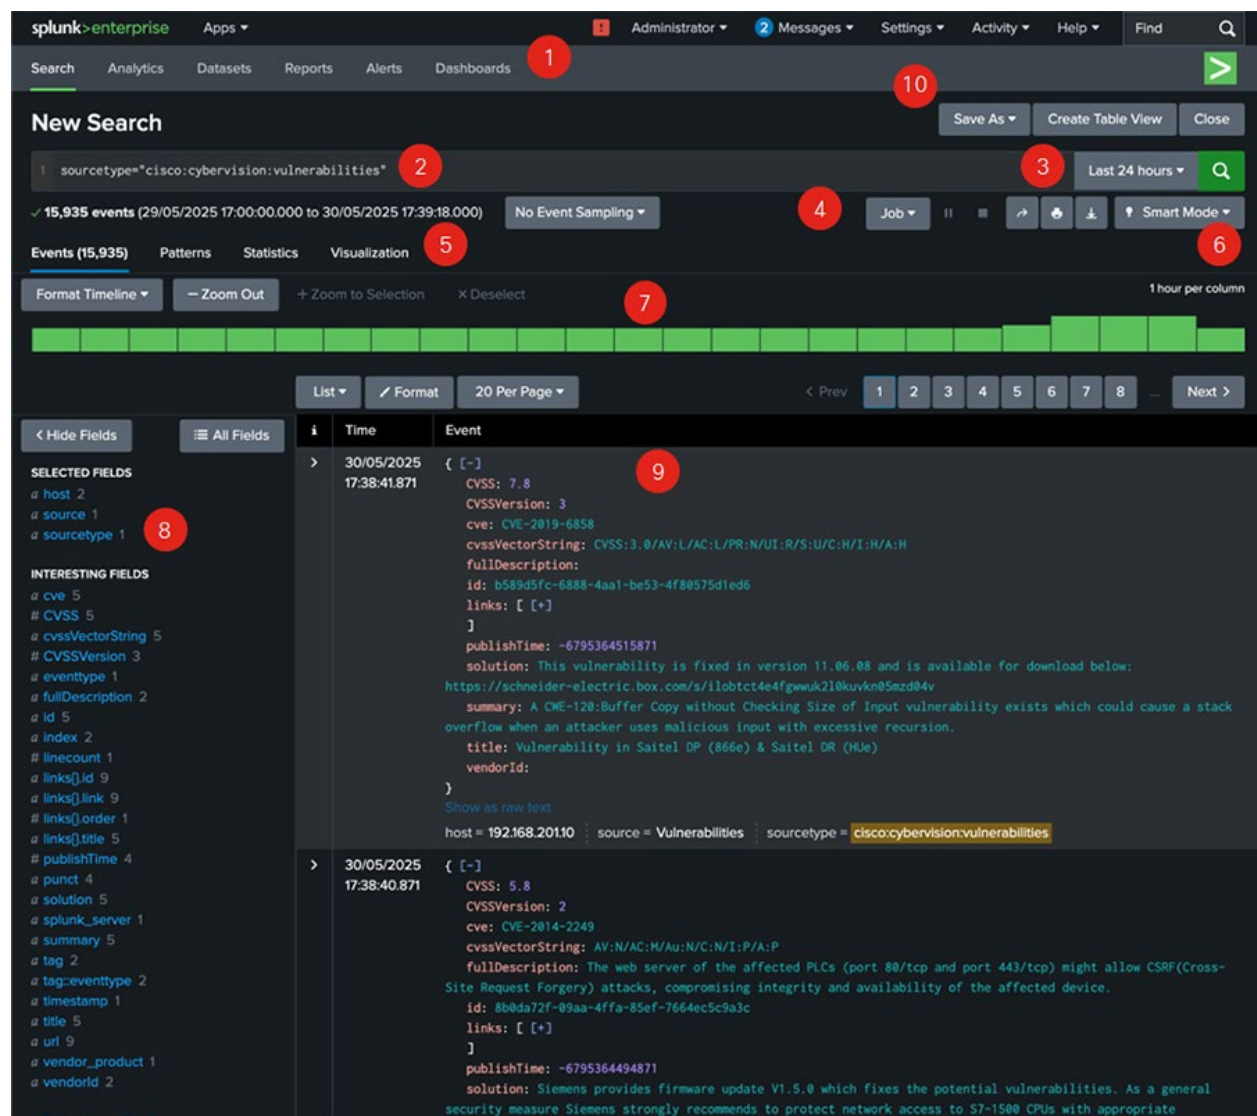


New Search View

The **New Search** view opens after you run a search or when you click the **Search** tab to start a new search. The App bar, Search bar, and Time range picker are still available in this view. Additionally, this view contains many more elements: search action buttons and search mode selector; counts of events; job status bar; and tabs for Events, Patterns, Statistics, and Visualizations.

For users who are following this guide, but do not yet have access to their own dataset, or are new to the platform, Splunk does offer a [Search Tutorial](#). The tutorial includes a dataset, sample search queries, enriching events with lookups, creating reports and creating dashboards. This design guide will focus on sample queries using the data ingestion options from the previous section.

Figure 124 Example search query to highlight the data available to a user after an initial search has run



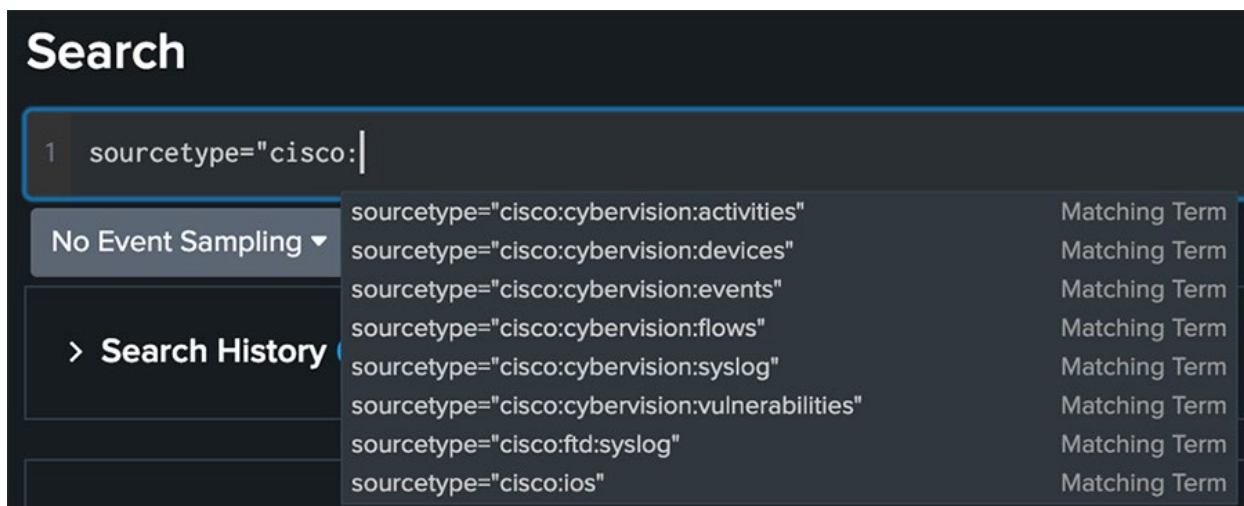
- 1 **Apps bar** Navigate between the different views in the Search & Reporting app: Search, Metrics, Datasets, Reports, Alerts, and Dashboards.
- 2 **Search bar** Specify your search criteria.
- 3 **Time range picker** Specify the time period for the search.

4	Search action buttons	Actions that you can perform, including working with your search Job, sharing, printing, and exporting your search results.
5	Search results tabs	The tab that your search results appear on depends on your search. Some searches produce a set of events, which appear on the Events tab. Other searches transform the data in events to produce search results, which appear on the Statistics tab.
6	Search mode menu	Use the search mode selector to provide a search experience that fits your needs. The modes are Smart (default), Fast, and Verbose.
7	Timeline	A visual representation of the number of events that occur at each point in time. Peaks or valleys in the timeline can indicate spikes in activity or server downtime. The timeline options are located above the timeline. You can format the timescale, zoom out, or zoom to a selected set of events.
8	Fields sidebar	Displays a list of the fields discovered in the events. The fields are grouped into Selected Fields and Interesting Fields .
9	Events viewer	Displays the events that match your search. By default, the most recent event is listed first. In each event, the matching search terms are highlighted. To change the event view, use the List , Format , and Per Page options.
10	Save As menu	Use the Save As menu to save your search results as a Report, Dashboard Panel, Alert, or Event Type.

Using the Search Assistant

The Search Assistant is a feature in the Search app that appears as you type your search criteria. The Search Assistant is like autocomplete for SPL query creation. Let's go back to the previous image, where the search query `sourcetype="cisco:cybervision:vulnerabilities"` was used to search for all the vulnerability data across the deployment.

Figure 125 Splunk search assistant feature



As letters are typed into the Search bar, the Search Assistant shows match terms to help auto complete the query, enabling users to navigate the data without the need to remember every index name or source type available to them. The recommendation is to learn the fundamentals of SPL query language, get familiar with the most common terms, and then Splunk will provide guidance from there.

Fields sidebar

When you add data to the Splunk platform the data is indexed. As part of the index process, information is extracted from your data and formatted as name and value pairs, called **fields**. When you run a search, the fields are identified and listed in the Fields sidebar next to your search results. The fields are divided into two categories.

- **Selected fields** are visible in your search results. By default, host, source, and sourcetype appear. You can select other fields to show in your events.
- **Interesting fields** are other fields that have been extracted from the events in your search results.

Figure 126 Highlighting the CVE field shows all unique values available for further queries

The screenshot shows the 'New Search' interface with the query `sourcetype="cisco:cybervision:vulnerabilities"`. It displays 15,920 events. The 'INTERESTING FIELDS' section on the left lists various fields, with 'cve' highlighted. A modal for the 'cve' field is open, showing 5 values (100% of events). The modal includes a 'Selected' toggle set to 'Yes' and a table of top values.

Values	Count	%
CVE-2014-2249	3,184	20%
CVE-2017-12090	3,184	20%
CVE-2019-6858	3,184	20%
CVE-2020-35198	3,184	20%
CVE-2021-33012	3,184	20%

Fields provide a mechanism to help explore the data available to you. For example, after running the search query for vulnerabilities detected in the OT network, the CVE values were extracted as an interesting field from the returned data.

Figure 127 Follow on search query after clicking on the CVE of interest

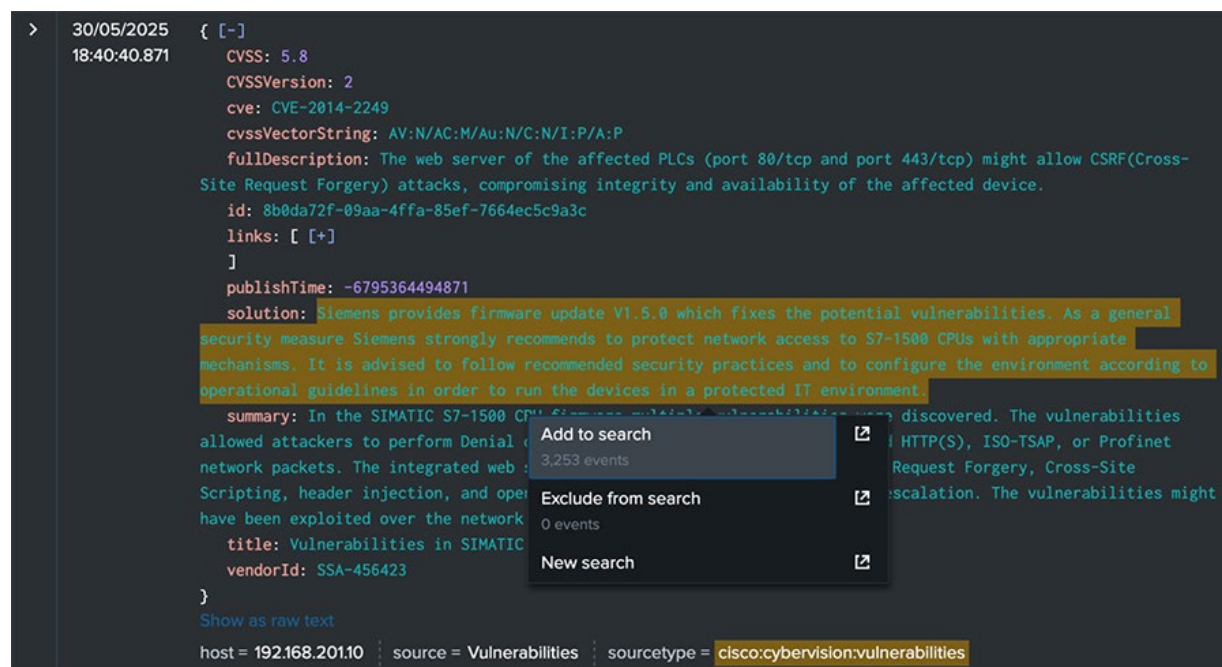
The screenshot shows the 'New Search' interface with the follow-up query `sourcetype="cisco:cybervision:vulnerabilities" cve="CVE-2014-2249"`. It displays 3,253 events. The interface includes buttons for 'Save As', 'Create Table View', and 'Close'. The 'Last 24 hours' filter is selected, and the 'Smart Mode' is enabled.

Click on one of the extracted CVE's simply created a new search, focusing on the CVE that had been chosen.

Interacting with the events


Another method to navigate through the data and help build search queries is by interacting with the events directly.

Figure 128 Adding data directly from the event into a new search



By clicking on the field that is of interest, there is an option to Add to search. In this scenario, we have noticed that the vulnerability has a solution related to a firmware update. This has created an interest to understand all of the vulnerabilities that can be cleared via a firmware patch.


Figure 129 Example query that searches for all vulnerabilities in the database with the term firmware update in the solution



While clicking on Add to search adds the full body of text to the search, a simple modification using an asterisk either side of “firmware update” and removing the original CVE provides all vulnerabilities in the database that contain the word “firmware update” in the solution field.

There is still one problem with this search. The previous query resulted in 8803 events with the firmware update event. Due to the nature of the Cyber Vision integration, vulnerability information is kept up to date with each query, so if a vulnerability continues to exist, Splunk will create a new event for it every time it polls the system. To combat this, we use the dedup command which removes duplicates based on the field provided.

Figure 130 Using the dedup search command to focus on unique CVE entries



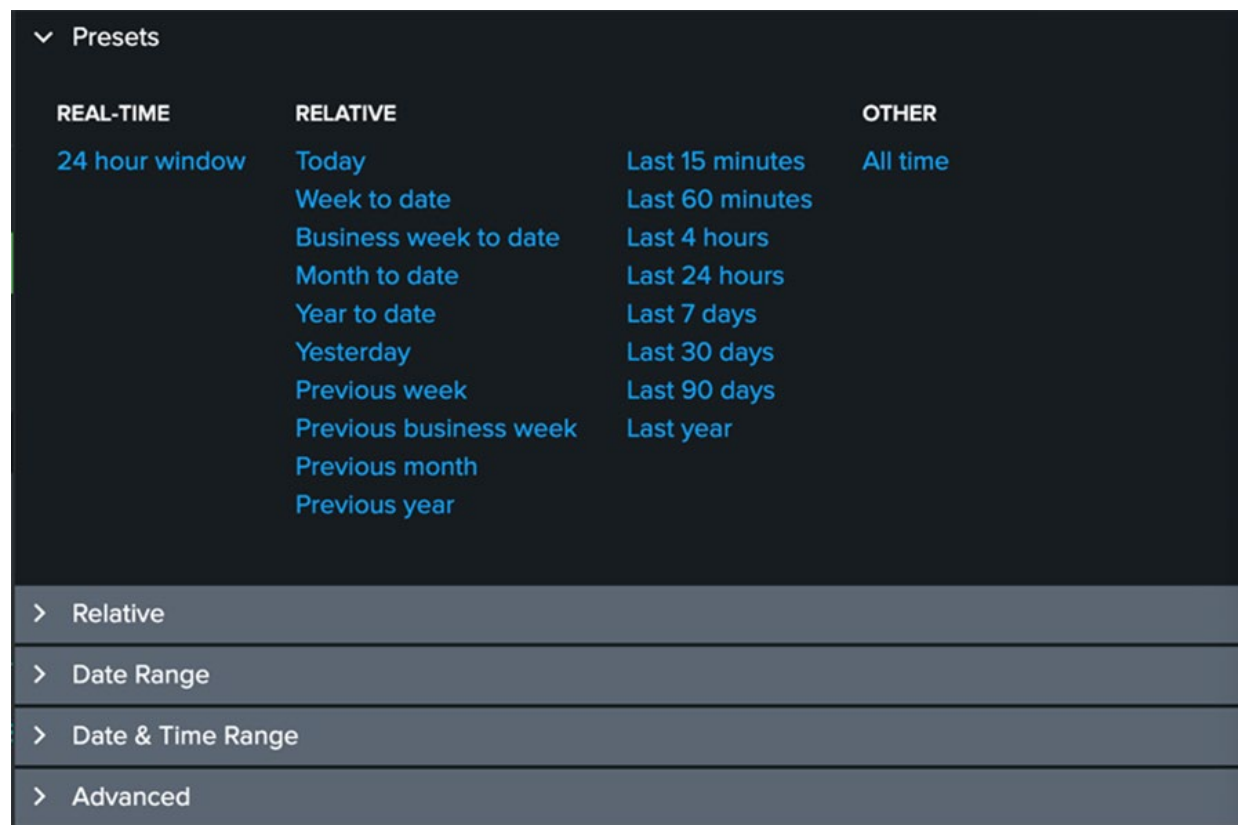
In this example, the CVE was used to only show unique vulnerabilities in the network that can be removed via a patch. The de-duplication of data reduced the results from 8803 events, to only 2.

Note: more sample queries will be included later in the design guide when discussing custom dashboards.

Time range picker

Time is the single most important search parameter that you specify.

Figure 131 Choosing a time range for the search query



Use the time range picker to retrieve events over a specific time period. Real-time searches displays a live and continuous view of events as they stream, into the Splunk platform. Real-time searches are the default when the system is first used, capturing all events that have occurred in the last 24 hours. Historical searches a those with a distinct time range, such as the past hour, the previous day, or “between 2 and 4 last Tuesday”. The time range picker has many preset time ranges that you can select from, but you can also type a custom time range. When searching for data, and more importantly when building reports and dashboards, make sure to pick an appropriate time range. Monthly reports should analyse the previous 30 days, weekly reports should analyse the previous 7. Incident investigation may require historical searching to narrow in to a specific indicator of compromise.

For more information, see [About searching with time](#).

Search actions

There are a wide range of search actions you can perform, including working with your search Jobs, saving, sharing, exporting, and printing your search results.

Search actions were not considered in this design guide, however, for more information, see:

- [Perform actions on running searches](#)
- [About jobs and job management](#)
- [Export search results](#)

Splunk's Alert Framework

Alerts provide an analyst a way to monitor for and respond to specific events by using a saved search to look for events in real time or on a schedule. Alerts trigger when search results meet specific conditions, and then alert actions can be used to respond to those triggers.

There are two types of alerts, scheduled and real-time. Real-time alerts searches continuously for events and are useful in situations where immediate monitoring and responses are important.

Example scenarios:

- The presence of malicious traffic should be investigated immediately, such as when a Snort signature has triggered indicating there may have been a compromise to the network.
- Key operational events may be restricted unless operating within a maintenance window. If a firmware update or program download was detected in the OT network, and it is not during a known maintenance window, it may need to be investigated with urgency.

Schedule alerts search for events on a regular basis and monitor whether they meet specific conditions, which is useful if immediate or real-time monitoring is not a priority.

- An organization may have an approved list of vendors that should be used in the OT environment. Once a day, Splunk can search the Cyber Vision index for all of the vendors seen in the network and generate an alert if a device deviates from the approved list.
- A new CVE may be prevalent in the news and an initial search was created to look for the presence for this CVE in the network. The initial search came up empty, but this query may run on a schedule to alert a user if this vulnerability does appear as it needs to be prioritised for remediation.

Alert Trigger Conditions

An alert can search for events on a schedule or in real time, but it does not have to trigger every time search results appear. Trigger conditions help monitor patterns in event data or prioritize certain events. The trigger condition works as a secondary search to evaluate the alert's initial search results. If the secondary search does not return results, the alert does not trigger.

For example

```
index="security" sourcetype="cisco:ise:syslog" MESSAGE_TEXT="RADIUS Request dropped"
```

checks for RADIUS requests that have been dropped by ISE, which can happen from time to time across a large deployment especially with the presence of employee laptops that may be temporary out of compliance and need an update. However, a large number of them may indicate there is something wrong in the system.

The first step is to modify the search to count the number of events returned in the search:

```
index="security" sourcetype="cisco:ise:syslog" MESSAGE_TEXT="RADIUS Request dropped" | stats count
```

The following custom triggering condition is then added to trigger the alert only if the count exceeds 100:

search count > 100

Alert Actions

Alert actions help respond to triggered alerts. Multiple alert actions can be taken, and the following options are available:

- Email notification
- Webhook
- Output results to a CSV lookup file
- Monitor triggered alerts
- Sending an alert to Splunk mobile users

For more information, see [Set up alert actions](#).

Creating Alerts

To create a schedule alert:

- Navigate to the **Search** page in the **Search and Reporting app**
- Create a search
- Select **Save As > Alert**
- Follow the prompts to build either scheduled or real-time alerts.

The following search queries were used during alert validation testing:

- Return the number of assets in the Cyber Vision database that belong to an unwanted vendor list.

```
index="security" sourcetype="cisco:cybervision:devices"
| dedup id
| search asset_vendor IN ("Belkin International Inc.", "Siemens")
| stats count by asset_vendor, id
```

- Return the number of assets in the Cyber Vision database that do not belong to a known vendor list.

```
index="security" sourcetype="cisco:cybervision:devices"
| stats dc(id) as "Device Count" by asset_vendor
| eval is_unwanted=if(asset_vendor IN ("Siemens AG", "Siemens", "Cisco Systems, Inc"), "No", "Yes")
| where is_unwanted="Yes"
| rename asset_vendor as Vendor
| table Vendor, "Device Count"
```

- Find devices that have known vulnerabilities in Cyber Vision and have also triggered an intrusion event in the firewall

```
(index=security sourcetype="cisco:cybervision:devices"
vulnerabilitiesCount > 0) OR (index="cisco_secure_fw"
EventType="IntrusionEvent")
| eval ResponderIP=if(index=="cisco_secure_fw", ResponderIP, dest_ip)
| stats values(label) as "Label", values(dest_ip) as "Device IP",
values(vulnerabilitiesCount) as "Vulnerabilities Count",
count(eval(index=="cisco_secure_fw")) as IntrusionEventCount by
ResponderIP
```

```
| where IntrusionEventCount > 0 AND 'Vulnerabilities Count' > 0
| table "Device IP", "Label", "Vulnerabilities Count",
"IntrusionEventCount"
```

For more information on alerts see the Splunk [alerting manual](#).

Building Custom Dashboards and Visualizations

Splunk platform visualisations are used to organize and communicate data insights. Visualisations and dashboards help users monitor or learn about important metrics and trends. Dashboards can be built using Classic Dashboards, the original tool based on Simple XML, or Dashboard Studio, a UI-driven experience with customizable layouts and advanced visualisation tools.

This best way to get started with building customised dashboards is by following the Splunk tutorials:


- [Splunk Dashboard Studio tutorial](#)
- [Splunk XML Dashboard tutorial](#)

The recommendation in this design guide is to use Dashboard Studio as some of the features of the XML dashboard have been deprecated, such as the scheduled reporting that will be discussed in the next section, and Dashboard studio offers greater customisation.

A common use case for building custom dashboards is to aggregate data across multiple individual apps into a single pane of glass. To help get you started on that journey, this design guide will show the process of how to pull existing charts from other Splunkbase applications into your own dashboard as a starting point.

Starting in the Cyber Vision app, there is a pie chart that represents the number of events in the system by their severity rating. By hovering over the widget, a menu appears in the bottom right which enables users to **Open in Search**.

Figure 132 Cyber Vision App widget for displaying events by severity



Getting the search used to generate the chart gives us the first thing needed for creating widgets – the data source. Copy the search, this will be used in the next step:

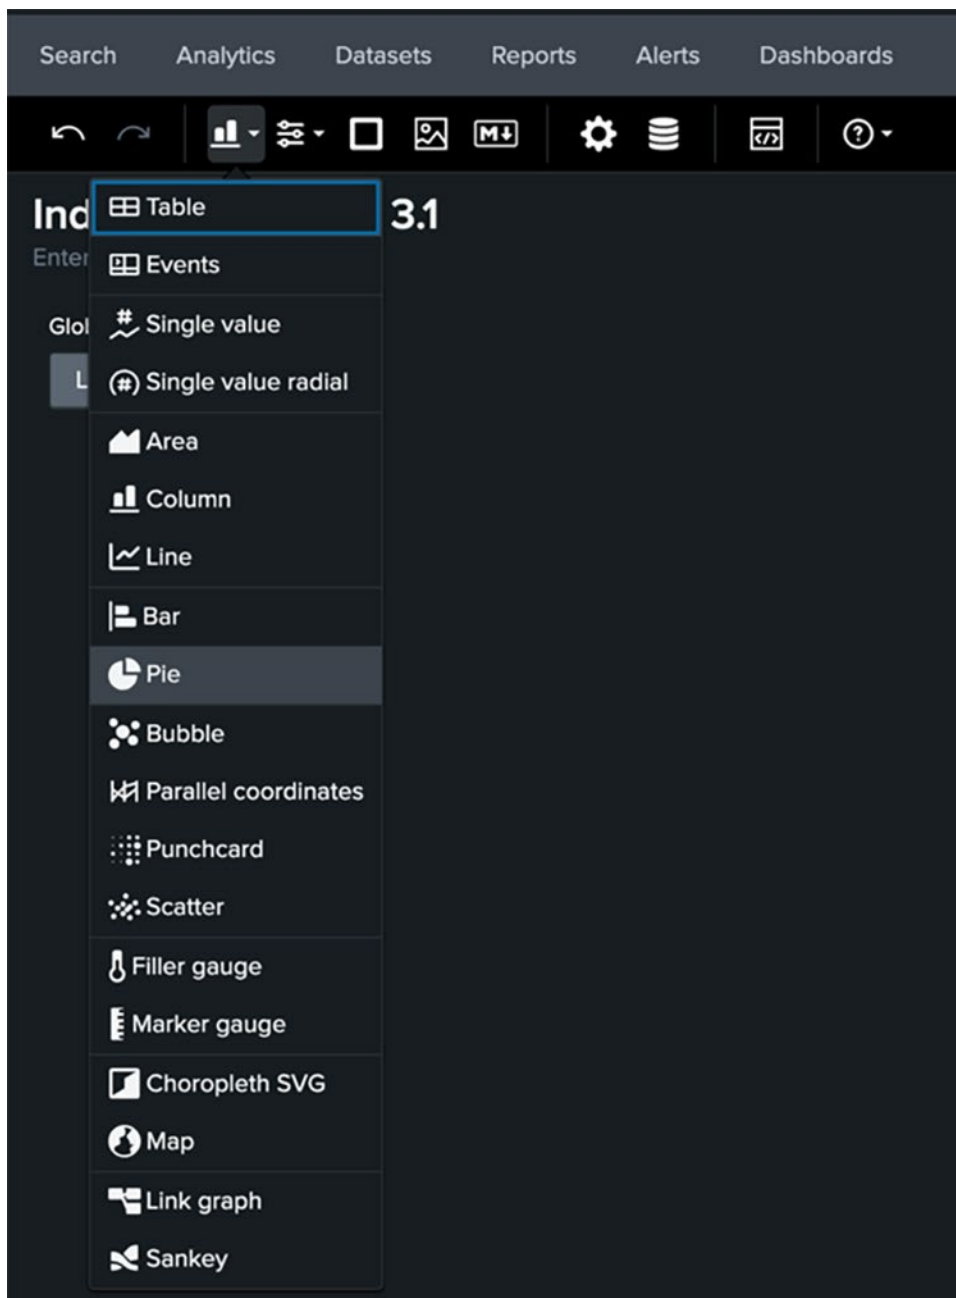
```
sourcetype="cisco:cybervision:events" category IN ("*") "*" | dedup id | stats count by severity
```

Navigate to the **Search & Reporting** app and open the **Dashboards** tab. Click **Create New Dashboard** and choose **Dashboard Studio**. You are now ready to build the dashboard. There are many out of the box visualisation options available to choose from:

- Area and line charts
- Bar and column charts
- Bubble charts
- Choropleth SVG
- Events Viewer
- Filler and maker gauges
- Icons
- Link graphs
- Maps
- Parallel coordinates
- Pie charts
- Punchcard charts
- Sankey diagram
- Scatter charts
- Shapes
- Single value visualisations
- Table
- Trellis layout

The original widget that the data has been copied from was a pie chart, so let's start there. In the top left corner of the dashboard editor, click on the chart icon and choose **Pie**.

Figure 133 Splunk Dashboard Studio chart selector



Click + **Create search** and copy the search from earlier in the process.

Figure 134 Adding the data source to the chart

New data source

Data source name

CV events by severity


☐ Access search results or metadata ?

SPL query [Open in search](#)

```
sourcetype="cisco:cybervision:events" category IN ("
  *") "*" | dedup id| stats count by severity
```

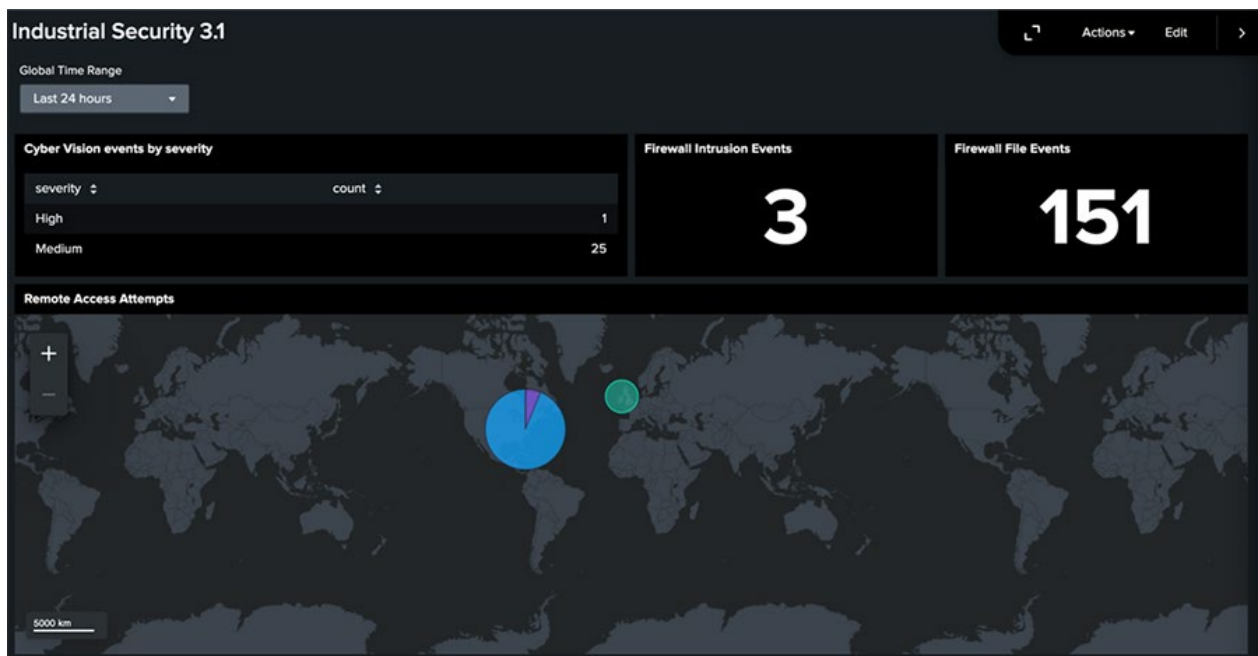
Click Apply and close and the widget will be copied to the dashboard. However, just because the original widget uses a pie chart does not mean this has to be the case. It is recommended to experiment with the visualization types to find the one that best fits the data. The search query does not need to change, users can experiment by changing the options in the Visualization type dropdown. For example, the image below shows the data in a bar chart instead.

Figure 135 Changing the original visualization type to a bar chart



Continue this process until all of the visualisation are complete. Experiment with color and sizing to meet visual needs when the dashboard gets busy. The following image represents a custom dashboard with various widgets across the apps that were used in validation testing.

Figure 136 Sample Splunk Studio Dashboard using widgets from existing Splunkbase apps



To recreate the dashboard above, use the following prompts:

Cyber Vision events by severity - Table

```
sourcetype="cisco:cybervision:events" category IN ("*") "*"
| dedup id
| stats count by severity
```

Firewall Intrusion Events – Single Value

```
| tstats count from datamodel=Cisco_Security.Secure_Firewall_Dataset where
nodename=Secure_Firewall_Dataset.Intrusion_Events
Secure_Firewall_Dataset.index IN (activedirectory, audit_summary, ba_test,
cim_modactions, cisco_catalyst, cisco_duo, cisco_sdwan_firewall,
cisco_secure_fw, cisco_sfw_api, cisco_sfw_ftd_syslog, cms_main, duo,
endpoint_summary, gia_summary, history, ioc, main, mc_artifacts,
mc_aux_incidents, mc_events, mc_incidents_backup, mc_investigations,
networking, notable, notable_summary, risk, security, sequenced_events,
summary, summary_cv_vulnerabilities, threat_activity, ubaroute, ueba, whois,
win_ad, windows) sourcetype IN (*) (Secure_Firewall_Dataset.initiator_ip=* OR
Secure_Firewall_Dataset.responder_ip=*)
Secure_Firewall_Dataset.Intrusion_Events.impact IN (*)
Secure_Firewall_Dataset.inline_result="*"
```

Firewall File Events – Single Value

```
| tstats count from datamodel=Cisco_Security.Secure_Firewall_Dataset where
nodename=Secure_Firewall_Dataset.File_Events Secure_Firewall_Dataset.index IN
(activedirectory, audit_summary, ba_test, cim_modactions, cisco_catalyst,
cisco_duo, cisco_sdwan_firewall, cisco_secure_fw, cisco_sfw_api,
cisco_sfw_ftd_syslog, cms_main, duo, endpoint_summary, gia_summary, history,
ioc, main, mc_artifacts, mc_aux_incidents, mc_events, mc_incidents_backup,
mc_investigations, networking, notable, notable_summary, risk, security,
```

```
sequenced_events, summary, summary_cv_vulnerabilities, threat_activity,  
ubaroute, ueba, whois, win_ad, windows) sourcetype IN (*)  
(Secure_Firewall_Dataset.initiator_ip=* OR  
Secure_Firewall_Dataset.responder_ip=*)
```

Remote Access Attempts (Duo) - Map

```
index="duo" extracted_eventtype=authentication  
| iplocation ip  
| geostats count by City
```

For more information see Splunk [dashboards and visualizations](#).

Using Dashboards to build Reports

Reports are created in Splunk when a search is saved for later use. A report can be as simple as a search that is run on a regular interval, and the results of that search is sent over an email. However, the recommendation in this design guide is to use the dashboards as they create more meaningful reports.

Dashboards in Splunk can be saved as a PDF simply by clicking on **Actions > Download PDF** from the dashboard. Alternatively, users can create a schedule for email delivery.

Figure 137 Scheduling a weekly report using the custom dashboard from the previous section

Edit export schedule [X]

Dashboard: **Industrial Security 3.1**

Schedule: ☒

Run every week on Monday at 6:00

The schedule is set to the timezone of this Splunk instance

Email recipients: abc123@yourdomain.com

Comma separated list of email addresses.

☐ Show CC and BCC

Priority: ☐ Critical ☐ High ☒ Normal ☐ Low ☐ Very low

Subject: Industrial Security 3.1

The email subject, recipients, and messages can include tokens that insert text based on the search results. [Learn more](#)

Message: Hi VIP,
Here is your weekly report for Industrial Security 3.1.
Regards,
Splunk

☐ Include link to dashboard
Use with discretion when sharing Splunk dashboard links with any email recipient.

File type: ☒ PDF ☐ PNG

[Send test email] [Cancel] [Save]

Note: Exporting dashboard PDFs, scheduling PDF delivery, and printing PDFs with Classic Simple XML dashboards is deprecated in Splunk Enterprise as of version 9.4.0 and Splunk Cloud Platform as of 9.3.2408. Although this feature continues to function, it might be removed in a future version, hence why this design guide recommends to use Dashboard Studio only.

For more information on reports see the Splunk [reporting manual](#).

Splunk Enterprise Security

Splunk Enterprise Security (ES) is a premium app which is used in conjunction with Splunk Enterprise or the Splunk Cloud Platform. While Splunk Enterprise as a standalone product offers a huge amount of capabilities, the out of the box capabilities are mainly limited to the applications that have been developed on Splunkbase. With Splunk ES, the Splunk Threat Research Team delves deep into detection engineering, providing 1800+ out-of-the-box detections that align to industry frameworks like MITRE ATT&CK, so that threats can be found and remediated faster.

Splunk ES provides a comprehensive threat detection, investigation, and response (TDIR) solution, which is key to the security monitoring strategy of today's enterprise infrastructure. Splunk Enterprise Security combines the best features and functionalities of Splunk's SIEM, SOAR, and threat intelligence management capabilities to identify security threats and effectively respond to them.

Splunk Enterprise Security offers the following benefits:

- Unified user experience and a seamless integrated workflow for case management, alert triage, investigation, and response
- Aligned taxonomy with Open Cybersecurity Schema framework (OCSF) and industry standards
- Enhanced detection and turnkey capabilities to implement risk-based alerting that creates high confidence alerts for investigations
- Alert aggregation capabilities using finding groups that map to pre-determined rules based on common security frameworks and techniques
- Automation with Splunk SOAR and full access to actions and playbooks.

Note: Splunk ES is out of scope for this version of the design guide. More details will be added at a later date. In the meantime, for more information, see the a list of demo videos at [Splunk Enterprise Security Features](#).

OT Security Add-On for Splunk

The Splunk for OT Security add-on expands the capabilities of Splunk's platform to monitor for threats and attacks, compliance, incident investigation, forensics, and incident response across the broad spectrum of assets and topologies that define modern manufacturing, energy, and public sector organizations.

The solution, comprised of an app and related documentation, provides the following features:

- **Expanded Asset Framework and Asset Center:** Ability to store and analyze additional asset attributes including facility/site id, asset criticality, asset types, classification, vlan, zone, and other data alongside traditional IT asset elements. Assets can be segmented by site or into multiple entity zones when attributes like IP Addresses and Host Names may be reused among different sites.
- **Integration with leading OT Security partner technologies:** Ingest asset inventory, vulnerabilities, and alerts from leading OT-ready systems, such as Cisco Cyber Vision.

- **Making using OT Data easier:** Prebuilt dashboards, reports, and other content related to perimeter monitoring, infrastructure monitoring, and centralized monitoring of multiple OT Security solutions. This content is in direct response to customers wanting quicker time to value from their OT data.
- **Prioritized vulnerability matching:** Evaluate, filter, and score matching vulnerabilities using iteratively executing correlation queries and dynamically calculated Asset Risk scores.
- **Integrated OT Asset Behavior Profiling:** Monitor asset behavior profiles to detect activity changes on critical assets that may represent increased threat risk.
- **Constructing and evaluating asset baselines:** Create baseline groups and baselines to verify assets follow a consistent hardened setup. Baselines can be created from data and extended to customer baseline types.
- **OT ready Correlation Searches:** Extend the deep bench of existing Enterprise Security correlation searches that monitor identity, endpoint, network and access in Splunk with OT-specific searches including mapping to common security frameworks including the MITRE ATT&CK for ICS.
- **Support for key elements of NERC CIP:** Dashboards and associated reports reviewed by trusted practitioners and NERC CIP auditors to help clients focus on NERC CIP requirements where Splunk can be assistive in compliance monitoring and audit support.

Note: the OT security add-on for Splunk is out of scope for this version of the design guide. More details will be added at a later date. In the meantime, for more information, see the [OT Security Add-on for Splunk: Technical Guide and Documentation](#).

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/Center-Appliance/Release-4-1-2/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-Cisco-Cyber-Vision-with-Cisco-Identity-Services-Engine-via-pxGrid_3_1_1.pdf

XDR

- Cisco Cyber Vision XDR Integration -
https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/GUI-Administration-Guide/Release-5-0-0/b_cisco-cyber-vision-GUI-administration-guide/m_integrations.html?bookSearch=true#xdr

Appendix B – Example TrustSec Configuration in Plant Networks

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

Table 6: User Interface for TrustSec Configuration

Configuration Item	Configuration Target	Configuration Tool*
Define and create SGTs and Policies	ISE	Cisco Catalyst Center or ISE
Define ISE as AAA server on network settings	Industrial switches and ISE	Cisco Catalyst Center or Industrial Switch and ISE
Enable device tracking on access ports	Industrial switches	Cisco Catalyst Center or Industrial Switch
Port-based Authentication	Industrial switches	Cisco Catalyst Center (templates) or Industrial Switch
Fall back policy and static entries	Enforcement switches (Distribution switch or Industrial switch)	Cisco Catalyst Center (templates) or Switch
Propagation (SXP or inline tagging)	Industrial switches, Distribution switch and ISE	Cisco Catalyst Center (templates) and ISE or Switch and ISE
Enable enforcement	Distribution switch or Industrial switch	Cisco Catalyst 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 Catalyst Center (templates) or switch

***Bold** represents method used for the CVD

Define and Create SGTs and Policies Using Cisco Catalyst Center

1. From the Cisco Catalyst 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.

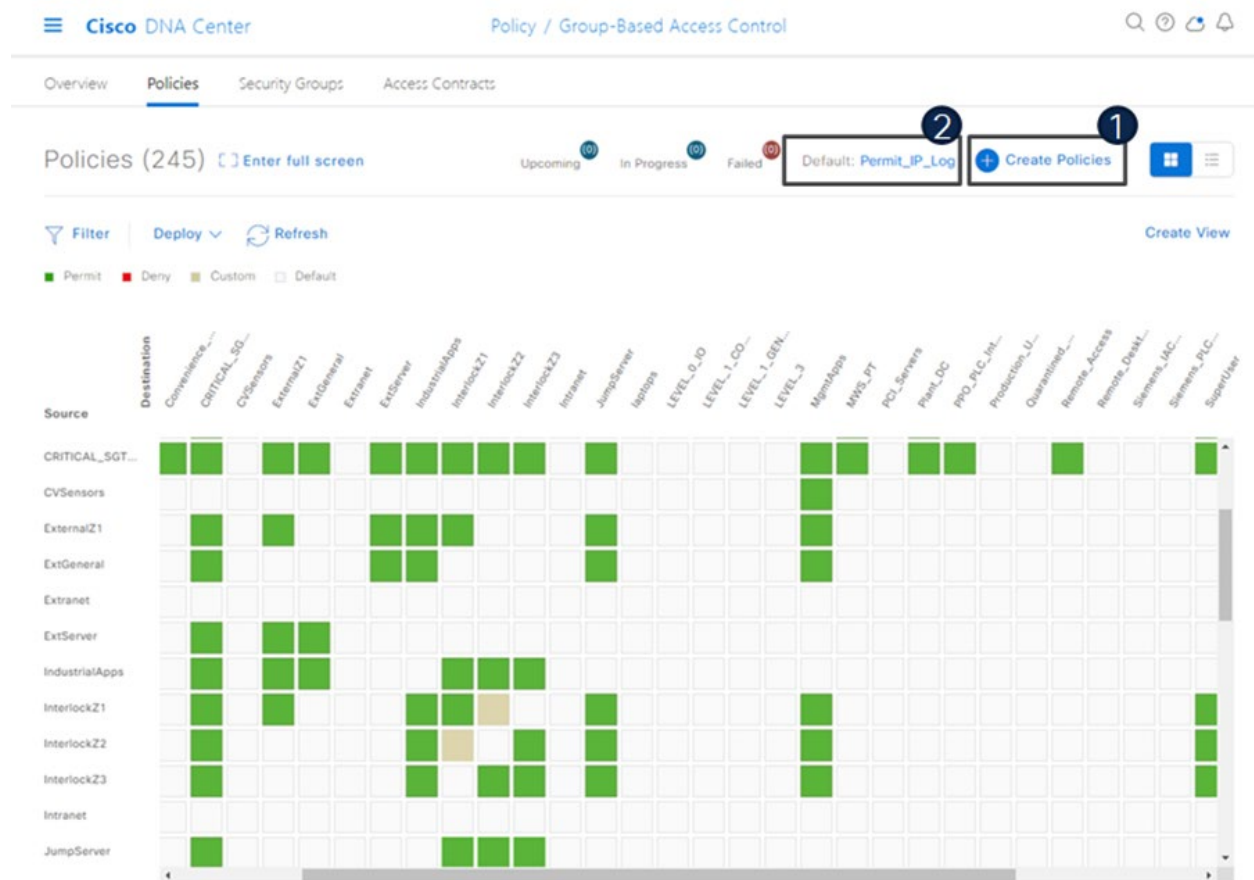
After creating the SGTs in Cisco Catalyst Center, the policy matrix can be updated to suit the enforcement intent. To make changes to the TrustSec policy matrix in Catalyst Center, do the following:

1. From the Cisco Catalyst 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.

Figure 108 shows the TrustSec policy matrix in Cisco Catalyst 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: Do not 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 138 TrustSec Policy Matrix in Cisco DNAC



Define ISE as the AAA Server using Cisco Catalyst Center

When a device is provisioned in the inventory, Cisco Catalyst Center configures AAA server information, CTS authorization commands, and RADIUS server groups. In addition, Cisco Catalyst 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 Catalyst 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 Catalyst Center.

Enable Device Tracking on Access Ports using Cisco Catalyst Center

Cisco Catalyst 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 Catalyst 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 switch-initiated 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, CATALYST CENTER, 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 Catalyst 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 **Create**.
6. Navigate to **Work Centers > TrustSec > SXP > SXP Devices**.
7. Click **Add**.
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 **Save**.

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 **Add**.
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 **Save**.

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. Check 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).
 - i. From the Expression drop-down list, choose Custom Attribute > assetGroup.
 - ii. From the logic drop-down list, choose **Contains**.
 - iii. 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 139 ISE Profiling Policy using Cyber Vision Group Data

Profiler Policy List > CVC_group_Interlock2

Profiler Policy

* Name	CVC_group_Interlock2	Description	<input type="text"/>
Policy Enabled	<input checked="" type="checkbox"/>		
* Minimum Certainty Factor	40	(Valid Range 1 to 65535)	
* Exception Action	NONE	▼	
* Network Scan (NMAP) Action	NONE	▼	
Create an Identity Group for the policy	<input checked="" type="radio"/> Yes, create matching Identity Group <input type="radio"/> No, use existing Identity Group hierarchy		
* Parent Policy	NONE	▼	
* Associated CoA Type	Global Settings	▼	
System Type	Administrator Created		

Rules

If	Condition	CUSTOMATTRIBUTE_assetGroup_CONT...	
----	-----------	------------------------------------	--

Conditions Details

Expression: CUSTOMATTRIBUTE:assetGroup CONTAINS Interlock2

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**.

Figure 110 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 140: Example Authorization Policies in ISE

	Status	Rule Name	Conditions	Profiles	Security Groups	Hits	Actions
			Search				
1	✓	SuperUser	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS Wired802_1x InternalUser-IdentityGroup EQUALS User Identity Groups:WS-user 	PermitAccess x	SuperUser	0	⚙️
	✓	Contractor	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS Wired802_1x InternalUser-IdentityGroup EQUALS User Identity Groups:Contractor 	PermitAccess x	Select from list	0	⚙️
2	✓	Interlock1	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled:CVC_group_Interlock1 	PermitAccess x	InterlockZ1	0	⚙️
	✓	Interlock2	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled:CVC_group_Interlock2 	PermitAccess x	InterlockZ2	0	⚙️
	✓	Interlock3	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled:CVC_group_Interlock3 	PermitAccess x	InterlockZ3	0	⚙️
	✓	External	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled:CVC_group_External 	PermitAccess x	ExtGeneral	0	⚙️
	✓	External1	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled:CVC_group_External1 	PermitAccess x	ExternalZ1	0	⚙️
3	✓	MABDefault	Normalised Radius-RadiusFlowType EQUALS WiredMAB	PermitAccess x	Select from list	23	⚙️

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
```


Appendix C – Installing custom Snort rules in Cisco Catalyst SD-WAN

The following steps show how to upload custom OT Snort signatures to Cisco Catalyst SD-WAN Manager.

Create the rules file

Note: The [official Snort documentation](#) describes the syntax in detail.

In general, each line is a different rule. In the example ruleset above, each rule starts with an action (alert, block, drop, log, pass, react, reject, rewrite), followed by some matching criteria. Snort can match on header level fields like well-known TCP/UDP port number, as well as payload data, such as a specific operation from a SCADA protocol, including [DNP3](#), [MODBUS](#), [CIP](#), [IEC 104](#), [MMS](#), and [S7CommPlus](#). Each protocol will have different fields, such as function codes for DNP3, that are defined in the Snort .cc files, such as [dnp3_map.cc](#), all of which are linked to in the previously mentioned supported SCADA protocol list. Looking at this file we can see, for example that DNP3 function code 2 is associated with the “write” function.

```
static dnp3_map_t func_map[] =
{
{ "confirm", 0 },
{ "read", 1 },
{ "write", 2 },
{ "select", 3 },
{ "operate", 4 },
{ "direct_operate", 5 },
{ "direct_operate_nr", 6 },
...
}
```

After defining the match criteria, a “msg: “ field is added for each rule – this message will be shown in the Catalyst SDWAN Manager’s Monitoring > Device > Intrusion Prevention page, whenever a rule is matched.

Next, in each rule, a “classtype” field is required, as defined in the [Snort documentation](#). This class type is also associated with a priority / severity level which is reflected in the IPS dashboard.

Finally, a signature ID (sid) field is included in each rule to help with searching/correlation of events. Each rule will have a unique signature ID value.

The rules used for validation follow:

```
alert tcp any any -> any 20000 (msg:"DNP3 confirm";
dnp3_func: 0; classtype:misc-activity; sid:1000000;)
```

```

block tcp any any -> any 20000 (msg:"drop DNP3 read";
  dnp3_func: 1; classtype:misc-activity; sid:1000001;)
block tcp any any -> any 20000 (msg:"block DNP3 write";
  dnp3_func: 2; classtype:misc-activity; sid:1000002;)
alert tcp any any -> any 20000 (msg:"DNP3 select";
  dnp3_func: 3; classtype:misc-activity; sid:1000003;)
alert tcp any any -> any 20000 (msg:"DNP3 operate";
  dnp3_func: 4; classtype:misc-activity; sid:1000004;)
alert tcp any any -> any 20000 (msg:"DNP3
  direct_operate"; dnp3_func: 5; classtype:misc-activity;
  sid:1000005;)
alert tcp any any -> any 20000 (msg:"DNP3
  direct_operate_nr"; dnp3_func: 6; classtype:misc-
  activity; sid:1000006;)
alert tcp any any -> any 20000 (msg:"DNP3 immed_freeze";
  dnp3_func: 7; classtype:misc-activity; sid:1000007;)
alert tcp any any -> any 20000 (msg:"DNP3
  immed_freeze_nr"; dnp3_func: 8; classtype:misc-activity;
  sid:1000008;)
alert tcp any any -> any 20000 (msg:"DNP3 freeze_clear";
  dnp3_func: 9; classtype:misc-activity; sid:1000009;)
alert tcp any any -> any 20000 (msg:"DNP3
  freeze_clear_nr"; dnp3_func: 10; classtype:misc-activity;
  sid:1000010;)
alert tcp any any -> any 20000 (msg:"DNP3
  freeze_at_time"; dnp3_func: 11; classtype:misc-activity;
  sid:1000011;)
alert tcp any any -> any 502 (msg:"Modbus Read Coils
  request"; modbus_func: 1; classtype:misc-activity;
  sid:1000133;)
alert tcp any any -> any 502 (msg:"Modbus Read Discrete
  Inputs"; modbus_func: 2; classtype:misc-activity;
  sid:1000134;)
alert tcp any any -> any 502 (msg:"Modbus Read Holding
  Registers"; modbus_func: 3; classtype:misc-activity;
  sid:1000135;)
alert tcp any any -> any 502 (msg:"Modbus Read Input
  Registers"; modbus_func: 4; classtype:misc-activity;
  sid:1000136;)
drop tcp any any -> any 502 (msg:"Modbus Write Single Coil";
  modbus_func: 5; classtype:misc-activity; sid:1000137;)

```

```
alert tcp any any -> any 502 (msg:"Modbus Write Single
Register"; modbus_func: 6; classtype:misc-activity;
sid:1000138;)

alert tcp any any -> any 502 (msg:"Modbus Read Exception
Status"; modbus_func: 7; classtype:misc-activity;
sid:1000139;)

alert tcp any any -> any 502 (msg:"Modbus Diagnostics";
modbus_func: 8; classtype:misc-activity; sid:1000140;)

alert tcp any any -> any 502 (msg:"Modbus Get Comm Event
Counter"; modbus_func: 11; classtype:misc-activity;
sid:1000141;)
```

Upload and use file in Cisco Catalyst SD-WAN Manager

1. After the custom rules file is created, upload it to Catalyst SD-WAN Manager through the Administration > UTD Snort Subscriber Signature page.
2. In the Policy Group page, associate the Custom Signature Set with an Intrusion Prevention Policy, from the Group of Interest link.
3. Associate the Intrusion Prevention Policy to an Advanced Inspection Profile that is used in the NGFW policy.

Appendix D – Cisco Cyber Vision vs. Cisco Secure Network Analytics

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.

Appendix E – Cisco SecureX

Note: Cisco SecureX is end of life and has been replaced with Cisco XDR. Most of this section still applies, but now exists under the Cisco XDR logo. Some of the screens may be different, but the content is much the same. A future version of this guide will update the content accordingly.

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:

- **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 shouldn't 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 app.

Figure 141 SecureX Ribbon in Cyber Vision

Device	Group	First activity	Last activity	IP	MAC	Risk score	Tags	Activities	Vuln	Var
1769-L18ER/B LOGIX5318ER (Port1-Link00)	Cell1	Mar 29, 2022 9:44:40 AM	Nov 15, 2022 2:11:43 PM	192.168.3.40	14:54:33:9b:77:76	82	Controller, SNMP Agent, Rockwell Automation	8	11	3
1769-L16ER/B LOGIX5316ER	Paint	Mar 29, 2022 9:44:48 AM	Nov 15, 2022 2:11:43 PM	192.168.3.50	14:54:33:91:cb:ee (+ 1 other)	87	Controller, SNMP Agent, Rockwell Automation	16	11	12
PWS	Paint	Jul 7, 2022 9:39:10 AM	Nov 15, 2022 2:11:41 PM	192.168.3.30	00:0c:29:c7:c8:76	54	HTTP Client, HTTPS Client, Windows	2153	0	0
BehrWorkWin10	Paint	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:39 PM	fe80:c22:d6cad:4ec:1299 (+ 1 other)	00:50:56:92:46:07	45	Engineering Station, IP-v6 Link Local, Windows	24	0	0
plcxb1d0ed	Cell3	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:29 PM	192.168.3.55	28:63:36:94:9e:7e (+ 1 other)	85	IO Module, Controller, SNMP Agent	11	29	0
kp8xb1d2c	Cell3	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:29 PM	192.168.3.56	28:63:36:44:93:bb (+ 1 other)	65	IO Module, Operator Panel, SNMP Agent	9	0	0
1734-AENTR/B Ethernet Adapter	Cell2				1a:4:28:13	82	IO Module, Rockwell Automation	6	8	0
Stratix 5800-MMS10EAR	Paint				1e:5:99:2b:er	59	IO Module, Network Switch, SNMP Agent, Cisco, Rockwell Automation	4	0	0
1734-AENTR/B Ethernet Adapter	Cell1				1a:4:25:16	75	Rockwell Automation	5	8	0

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 142 SecureX Threat Response Example

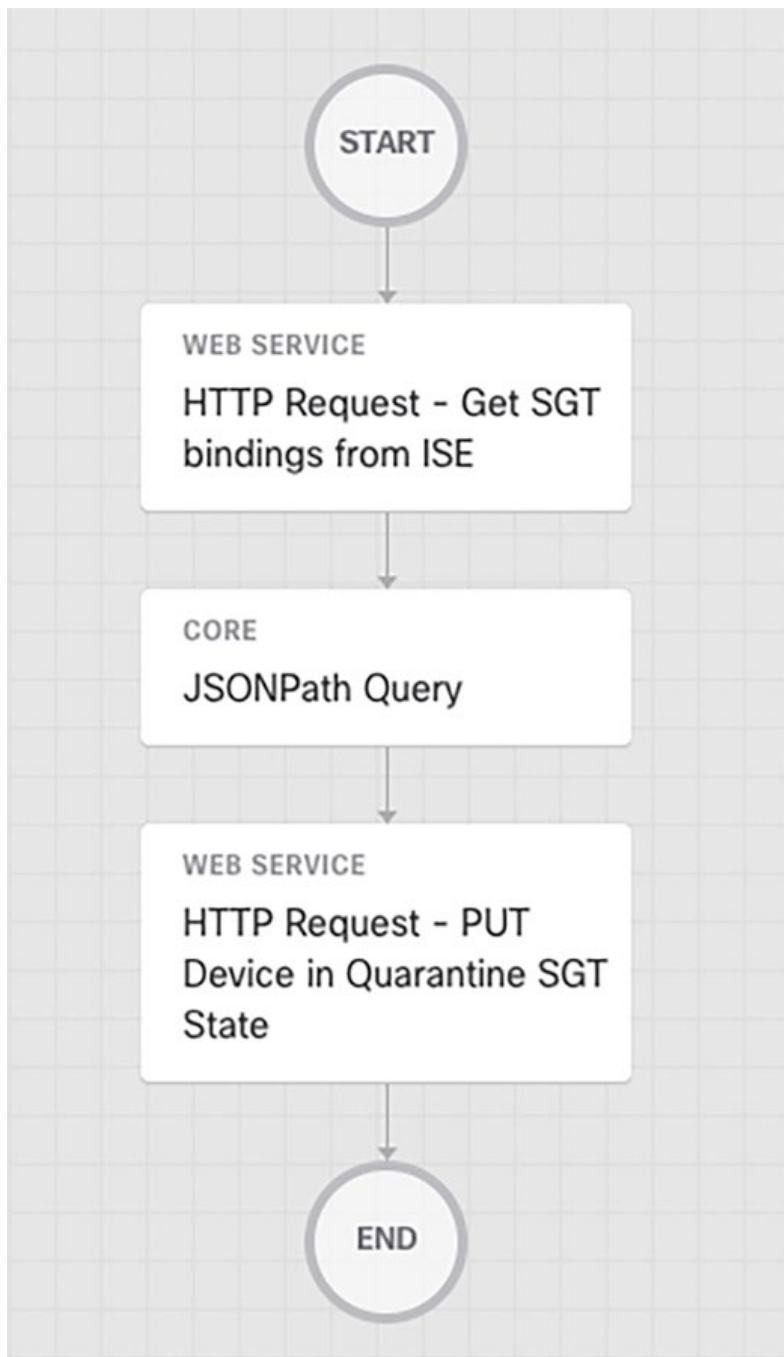
	Status	Rule Name	Conditions	Profiles	Security Groups	Hits	Actions
				Search			
1	✓	SuperUser	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS Wired802_1x InternalUser-IdentityGroup EQUALS User Identity Groups:WS-user 	PermitAccess x	SuperUser	0	⚙️
	✓	Contractor	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS Wired802_1x InternalUser-IdentityGroup EQUALS User Identity Groups:Contractor 	PermitAccess x	Select from list	0	⚙️
2	✓	Interlock1	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled.CVC_group_Interlock1 	PermitAccess x	InterlockZ1	0	⚙️
	✓	Interlock2	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled.CVC_group_Interlock2 	PermitAccess x	InterlockZ2	0	⚙️
	✓	Interlock3	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled.CVC_group_Interlock3 	PermitAccess x	InterlockZ3	0	⚙️
	✓	External	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled.CVC_group_External 	PermitAccess x	ExtGeneral	0	⚙️
	✓	External1	AND <ul style="list-style-type: none"> Normalised Radius-RadiusFlowType EQUALS WiredMAB IdentityGroup-Name EQUALS Endpoint Identity Groups:Profiled.CVC_group_External1 	PermitAccess x	ExternalZ1	0	⚙️
3	✓	MABDefault	Normalised Radius-RadiusFlowType EQUALS WiredMAB	PermitAccess x	Select from list	23	⚙️

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. 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 143 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 third-party 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.

Figure 144 Invoking SecureX Orchestration Workflow from the Ribbon in Cyber Vision

The screenshot displays the Cyber Vision interface with a table of 9 devices. The table columns include Device, Group, First activity, Last activity, IP, MAC, Risk score, Tags, Activities, Vuln, and Var. A context menu is open over the table, showing options for investigation and orchestration. The 'SecureX Orchestration' section is expanded, showing various actions like 'Quarantine Device', 'Shutdown Interface POC', and 'Take Orbital Forensic Snapshot'.

Device	Group	First activity	Last activity	IP	MAC	Risk score	Tags	Activities	Vuln	Var
1769-L18ER/B LOGIXS318ER (Port1-Link00)	Cell1	Mar 29, 2022 9:44:40 AM	Nov 15, 2022 2:11:43 PM	192.168.3.40	14:54:33:9b:77:76	82	Controller, SNMP Agent, Rockwell Automation	8	11	3
1769-L16ER/B LOGIXS316ER	Paint	Mar 29, 2022 9:44:48 AM	Nov 15, 2022 2:11:43 PM	192.168.3.50	14:54:33:91:cb:ee (+ 1 other)	87	Controller, SNMP Agent, Rockwell Automation	16	11	12
PWS	Paint	Jul 7, 2022 9:39:10 AM	Nov 15, 2022 2:11:41 PM	192.168.3.30	00:0c:29:c7:c8:76	54	HTTP Client, HTTPS Client, Windows	2153	0	0
BehrWorkWin10	Paint	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:39 PM	fe80::c22:d6cad:4ec:1299 (+ 1 other)	00:50:56:92:46:07	65	Engineering Station, IPv6 Link Local, Windows	24	0	0
plcxb1d0ed	Cell3	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:29 PM	192.168.3.55	28:63:19:2f:3d:50 (+ 1 other)		Controller	11	29	0
kp8xb1d62c	Cell3	Mar 29, 2022 9:44:18 AM	Nov 15, 2022 2:11:29 PM	192.168.3.56	28:63:19:2f:3d:50 (+ 1 other)			9	0	0
1734-AENTR/B Ethernet Adapter	Cell2							6	8	0
Stratix 5800-MMS10EAR	Paint							4	0	0
1734-AENTR/B Ethernet Adapter	Cell1							5	8	0

The context menu shows the following options:

- Investigate in Threat Response
- Create Judgement
- AMP for Endpoints
- Search for this IP
- Secure Endpoint - BehrHome
- Search for this IP
- SecureX Orchestration
 - Quarantine Device
 - Shutdown Interface POC
 - Take Orbital Forensic Snapshot
 - Take Forensic Snapshot and Isolate
 - Submit URL to Threat Grid
 - Move Computer to AMP Triage Group
 - AMP Host Isolation with Tier 2 Approval

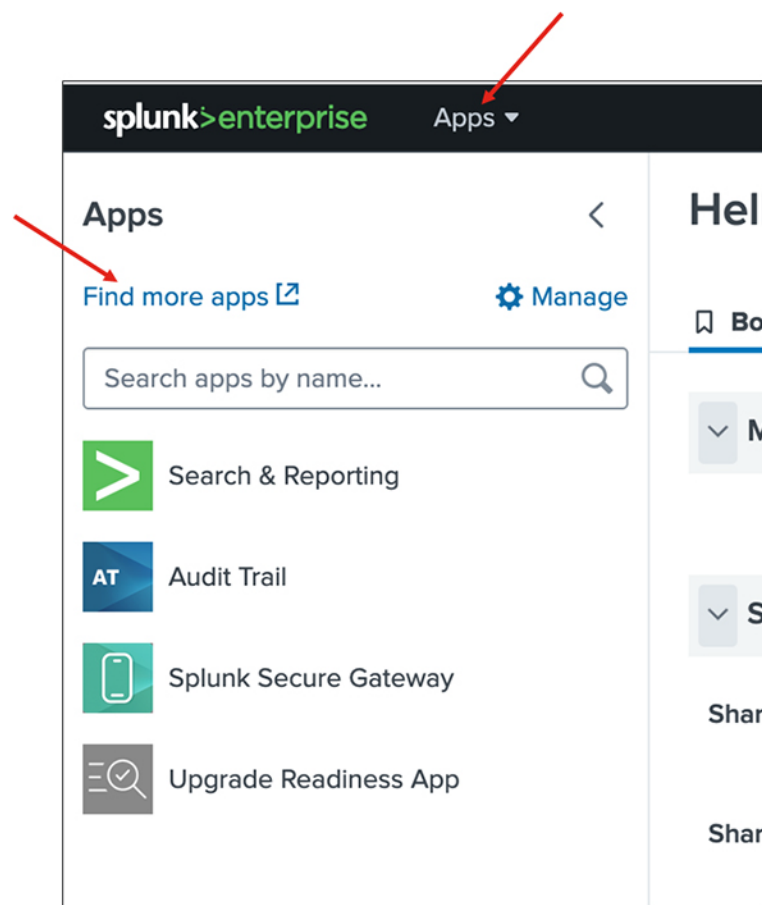
Appendix F – Cisco Cyber Vision and Splunk

Quick Deployment Guide

Note: before proceeding with the install, make sure time is synchronized between Splunk and Cyber Vision.

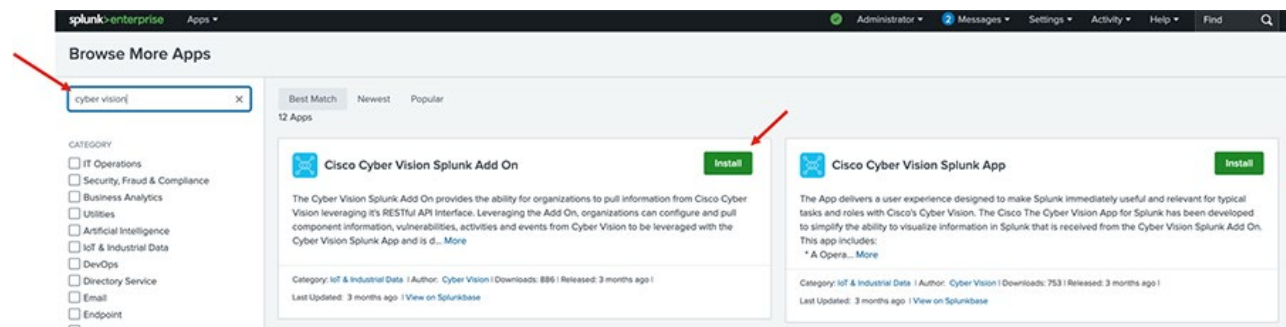
- In Splunk Enterprise, click **Apps** > **Find more apps** to navigate to Splunkbase

Figure 145 Splunk Apps menu



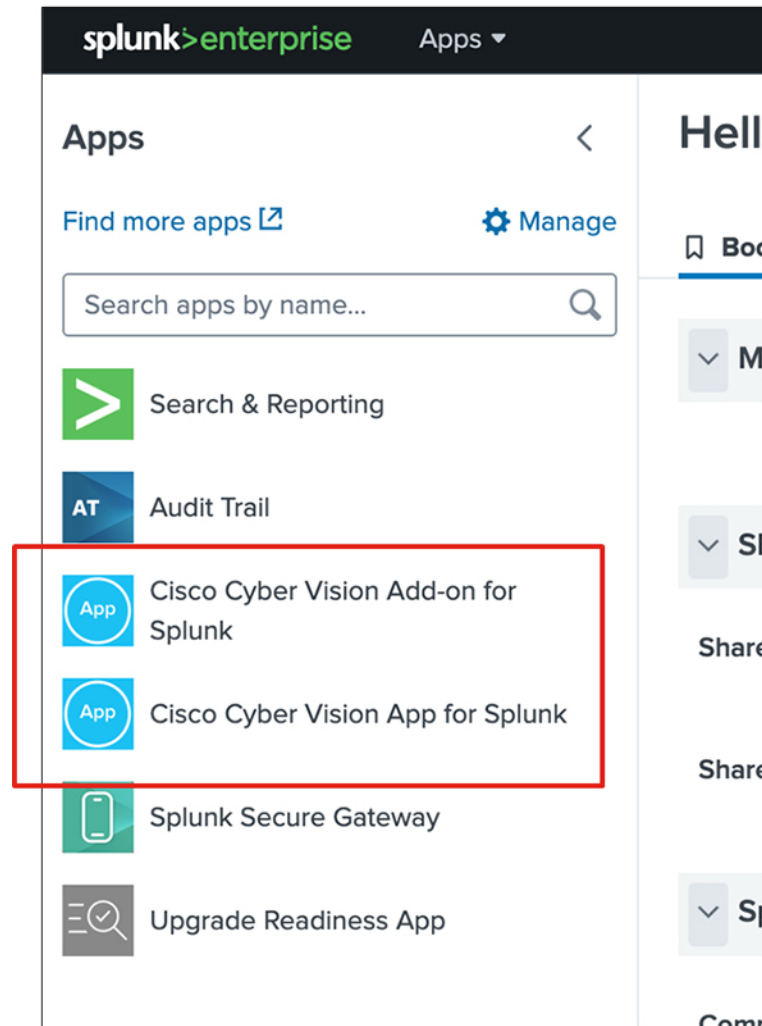
- Type “Cyber Vision” into the search filter box
- Click **Install** on **Cisco Cyber Vision Splunk Add On**

Figure 146 Splunk Cyber Vision Add on and App in Splunkbase



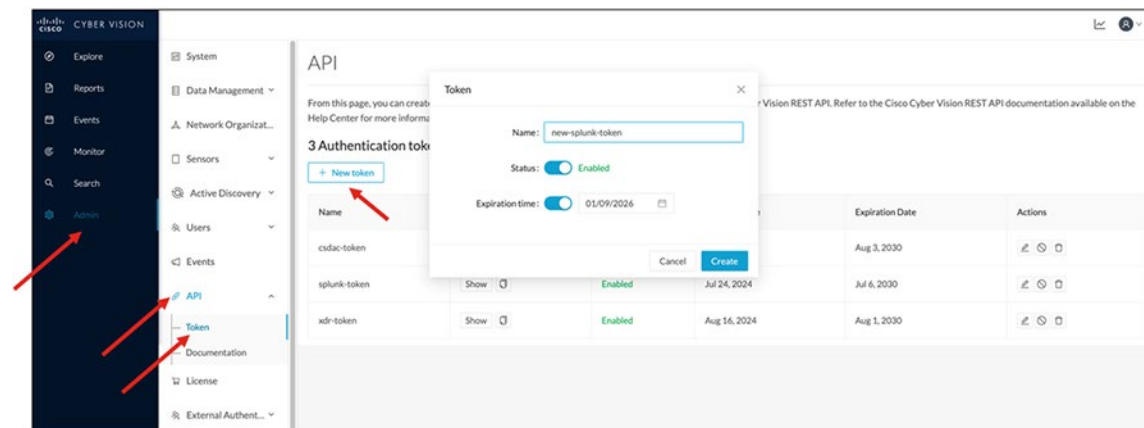
- Enter your Splunk account credentials and click **Agree and Install**.
- Repeat the same process for the **Cisco Cyber Vision Splunk App**.
- In Splunk Enterprise, click on Apps and verify that both apps are successfully installed.

Figure 147 Splunk Apps menu after the Cyber Vision apps have been installed



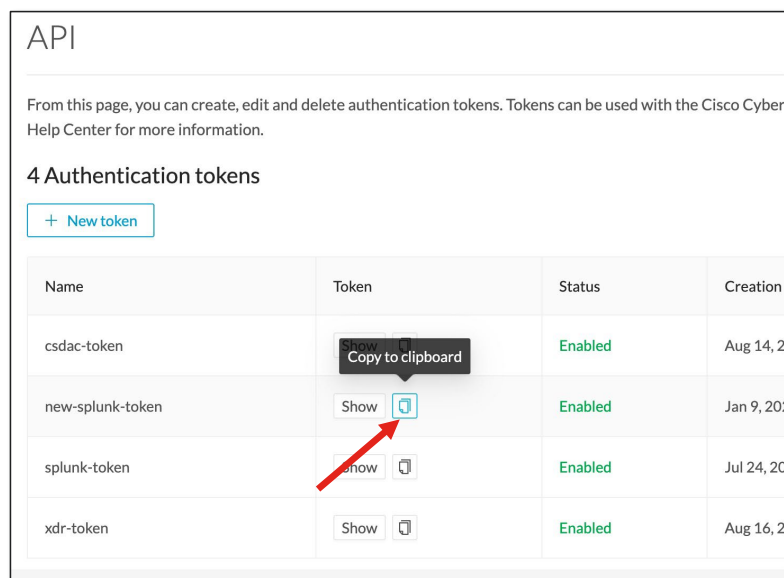
- In Cyber Vision, navigate to **Admin > API > Token** and click **+ New Token**

Figure 148 Generating an API token in Cyber Vision



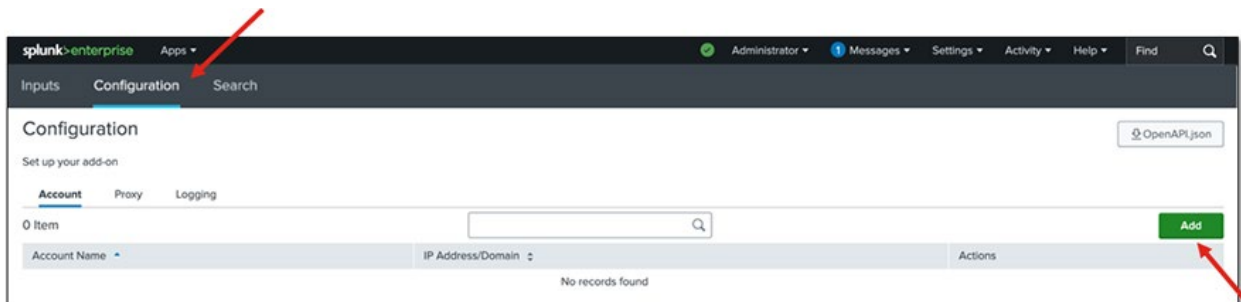
- Copy the token to the clipboard to be used in a subsequent step

Figure 149 Copying the API token to the clipboard



- Navigate to the Splunk Cyber Vision Add On and select **Configuration > Add**

Figure 150 Cyber Vision add on configuration page



Note: before processing, we need to decide which SSL certificate verification method to use. There are 3 possibilities

- No SSL certification verification (ok for lab/PoC use)
- SSL certificate verification, self-signed certificate
- SSL certificate verification, certificate authority

Proceed to the appropriate section depending on the decision.

No SSL certification verification

- In the Splunk CLI, change VERIFY_SSL (line 10) in the file TA_cisco_cybervision_utils.py from True to False

Figure 151 Disabling certificate verification in Splunk Enterprise



- In the Splunk Cyber Vision Add On, use the API token from the earlier steps and complete the fields to connect to Cyber Vision

Figure 152 Adding Cyber Vision to Splunk Enterprise with no certificate

The image shows the 'Add Account' dialog box in the Splunk Cyber Vision Add On. The fields are filled with the following values:

- Account Name:** cv_instance1
- IP Address/Domain:** https://172.26.136.10
- API Token:** [Masked]
- Use Custom CA Certificate:** ☐

The 'Add' button is highlighted in green.

SSL certificate verification, self-signed certificate

- When using certificates, the Cyber Vision Center name resolution must work on the Splunk side. To check the Cyber Vision Center FQDN, use the CLI command `/opt/sbs/bin/sbs-syem-fqdn`
- In Splunk, add a DNS server or use the local 'hosts' file
- When adding the account to Splunk, Cyber Vision Center FQDN must be used to define the center account
- Download the ca.pem file from Cyber Vision by navigating to <https://center-ip/ca.pem> and add it to the appropriate field

Figure 153 Adding Cyber Vision to Splunk Enterprise with a self-signed certificate

Update Account

Center FQDN → * Account Name:
Enter a unique name for this account.

* IP Address/Domain:
Enter the IP Address of the Cisco Cyber Vision in format https://<ip address> or https://<domain-name>

Center API token → * API Token:
Enter API Token generated from Cyber Vision for above account.

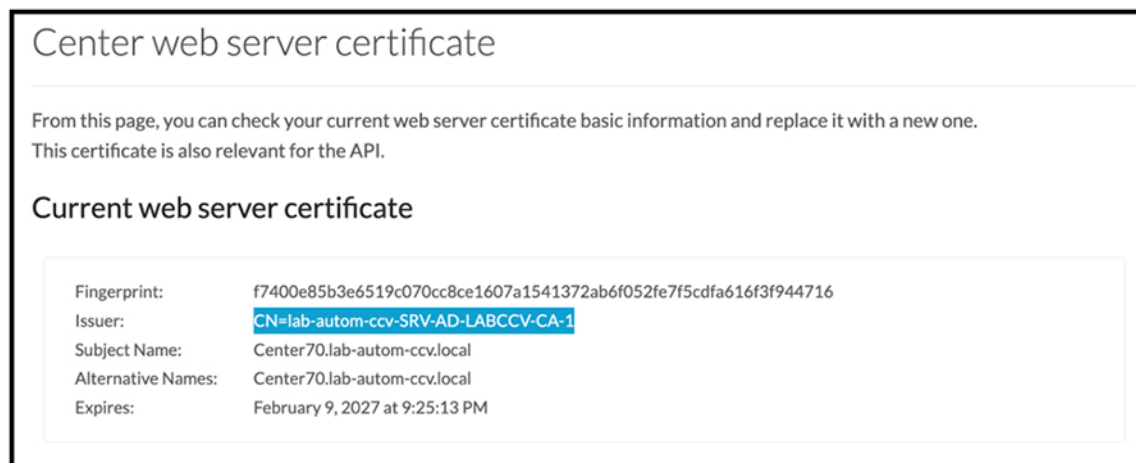
Use Custom CA Certificate: ☒

Center ca.pem → Custom CA Certificate:
Enter the Custom CA Certificate for this account.

SSL certificate verification, certificate authority

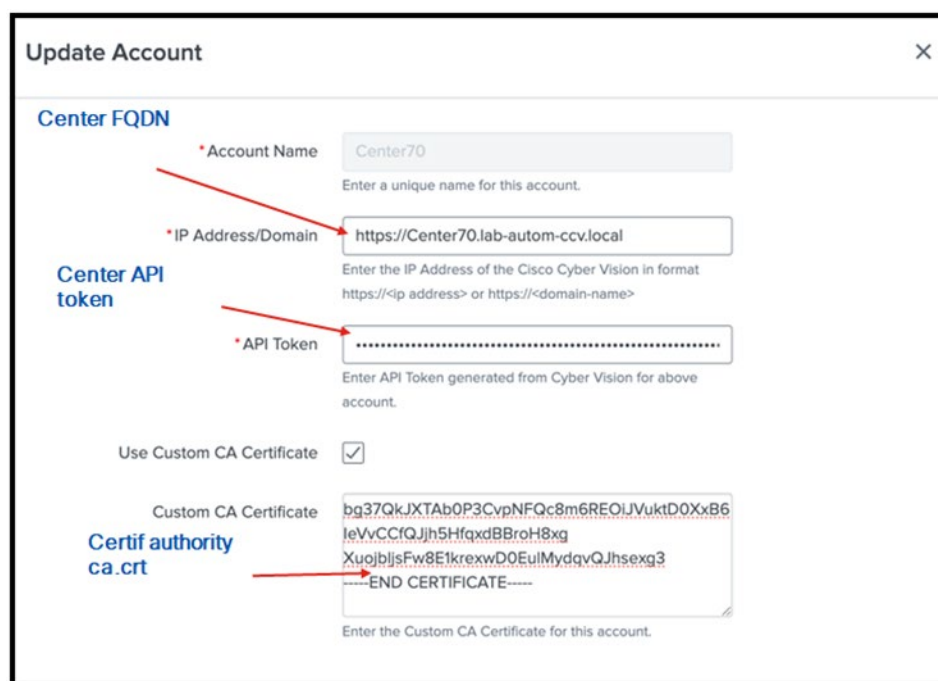
- When deploying in a production environment, it is recommended to issue a web server certificate from a certificate authority.

Figure 154 Adding a signed certificate to Cyber Vision



- Cyber Vision Center name resolution must work from Splunk.
- Repeat the same steps as defined previously, but using the certificate authority cert instead.

Figure 155 Adding Cyber Vision to Splunk Enterprise with a signed certificate



Adding Input(s)

- In the Cyber Vision Add On, navigate to **Inputs**.
- Click Create New **Input** > **Cyber Vision Events**
- Complete the fields as shown in the image.

Figure 156 Adding Cyber Vision events data input to Splunk

Add Cyber Vision Events [X]

*Name
Enter a unique name for the data input

*Interval
Time interval of input in seconds.

*Index

*Global Account

*Start Date
Enter the Start Date in format of YYYY-MM-DDTHH:MM:SSZ

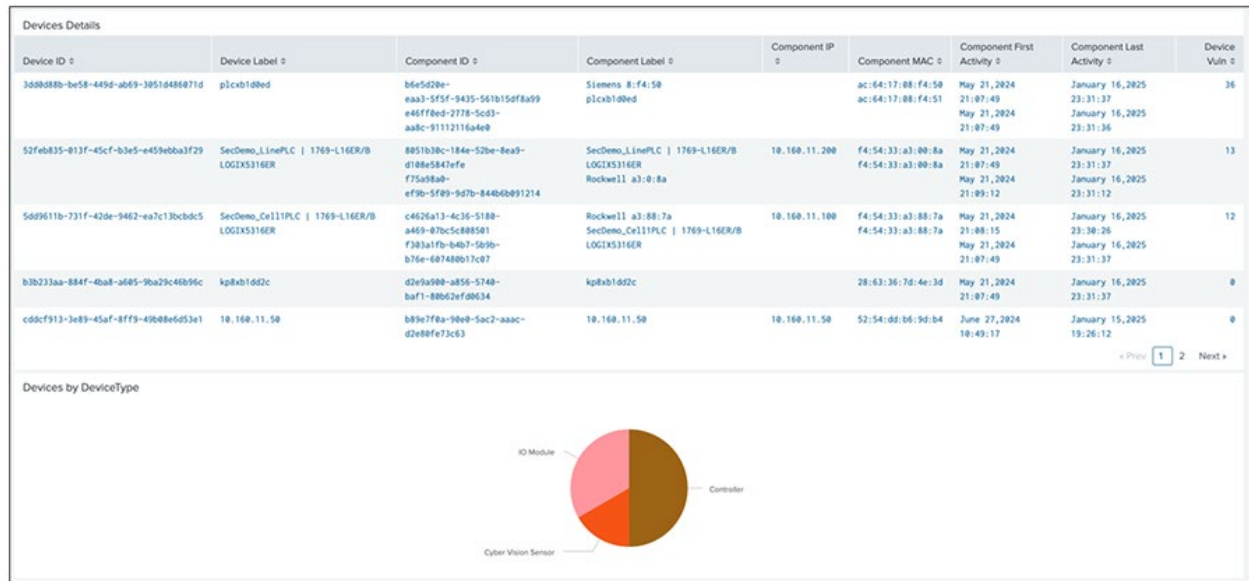
Cancel Add

Note: It is recommended to start with a long interval between updates. After a few days, check the data usage required for your deployment and then modify the interval to an appropriate shorter timespan (if desired). The “Start Date” field only applies the first time the add on pulls data from Cyber Vision. After that, the add on asks for any new data that has occurred since the last data request.

- Repeat the process for **Devices**, **Flows**, **Activities** and **Vulnerabilities**.

Data can now be visualised in the Cyber Vision App for Splunk. An example dashboard can be seen below.

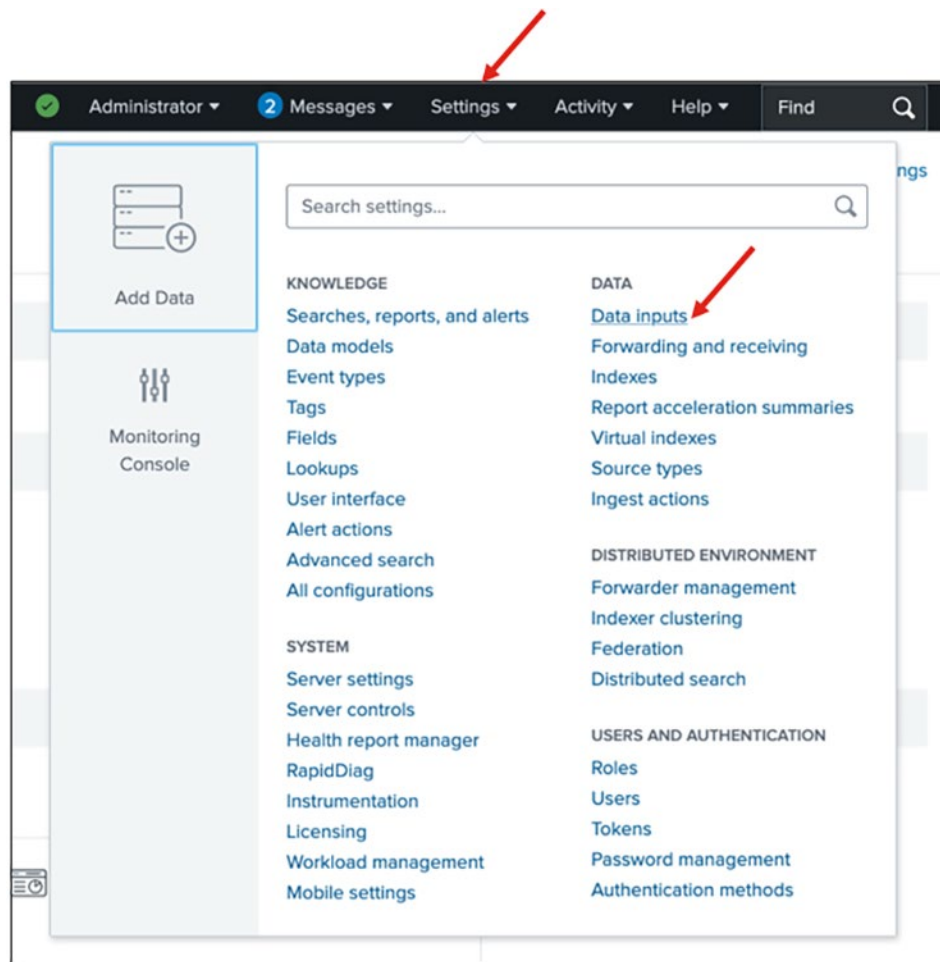
Figure 157 Sample dashboard from the Cyber Vision app for Splunk



The other option for data ingestion is through Syslog, which provides more real time data ingestion. This is a recommended option if the add on has been configured for long update intervals.

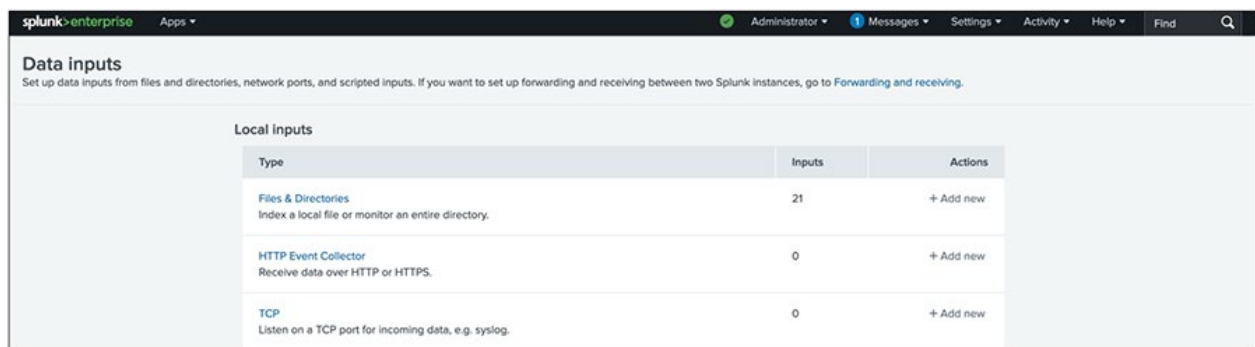
- In Splunk Enterprise, navigate to **Settings > Data inputs**

Figure 158 Splunk Enterprise settings menu for data inputs navigation



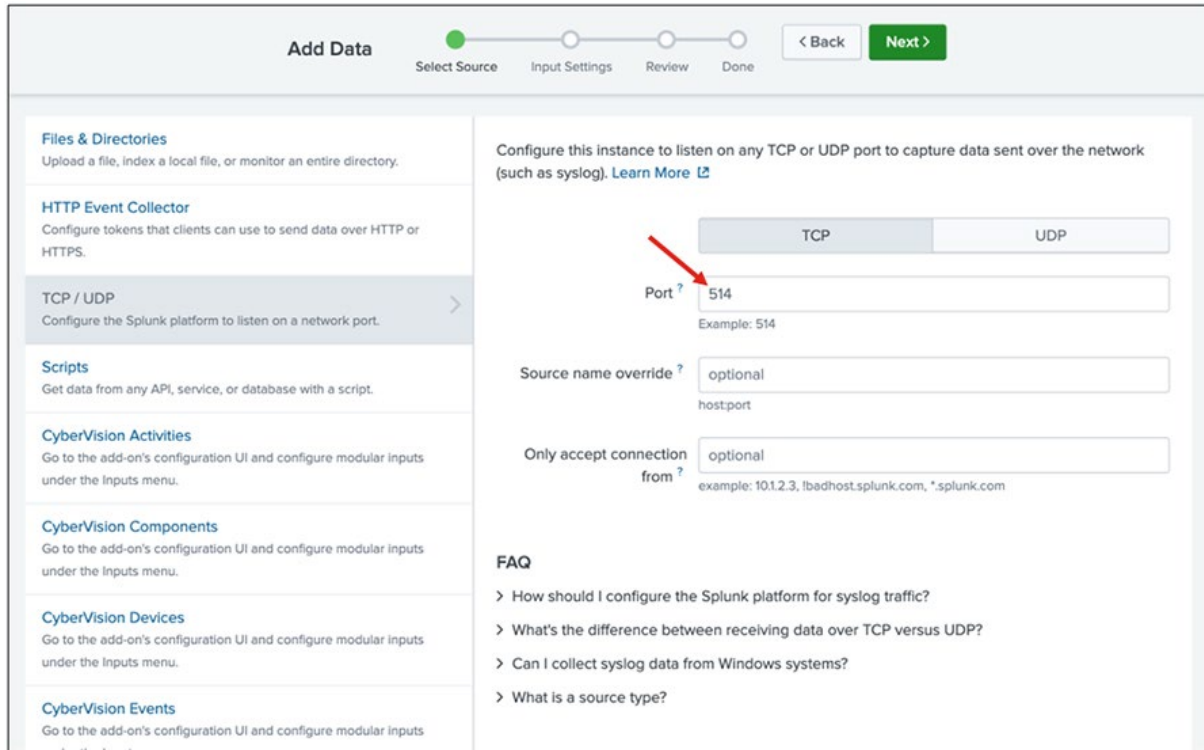
- Click on **TCP > + Add New**

Figure 159 Adding TCP data input to Splunk



- Specify **port 514** and click **Next**

Figure 160 Configuring the TCP input with port 514 which is the standard syslog port



Add Data

Select Source Input Settings Review Done

< Back Next >

Files & Directories
Upload a file, index a local file, or monitor an entire directory.

HTTP Event Collector
Configure tokens that clients can use to send data over HTTP or HTTPS.

TCP / UDP >
Configure the Splunk platform to listen on a network port.

Scripts
Get data from any API, service, or database with a script.

CyberVision Activities
Go to the add-on's configuration UI and configure modular inputs under the Inputs menu.

CyberVision Components
Go to the add-on's configuration UI and configure modular inputs under the Inputs menu.

CyberVision Devices
Go to the add-on's configuration UI and configure modular inputs under the Inputs menu.

CyberVision Events
Go to the add-on's configuration UI and configure modular inputs under the Inputs menu.

Configure this instance to listen on any TCP or UDP port to capture data sent over the network (such as syslog). [Learn More](#)

TCP **UDP**

Port ?
Example: 514

Source name override ?
host:port

Only accept connection from ?
example: 10.1.2.3, lbadhost.splunk.com, *.splunk.com

FAQ

- > How should I configure the Splunk platform for syslog traffic?
- > What's the difference between receiving data over TCP versus UDP?
- > Can I collect syslog data from Windows systems?
- > What is a source type?

- Select **sourcetype cisco:cybervision:syslog** and select the Cyber Vision Add on as the App content

Figure 161 Input setting ensures data is indexed into the correct format upon arrival

Add Data | Select Source | **Input Settings** | Review | Done | < Back | Review >

Input Settings

Optionally set additional input parameters for this data input as follows:

Source type
The source type is one of the default fields that the Splunk platform assigns to all incoming data. It tells the Splunk platform what kind of data you've got, so that the Splunk platform can format the data intelligently during indexing. And it's a way to categorize your data, so that you can search it easily.

Select | New
cisco:cybervision:syslog

App context
Application contexts are folders within a Splunk platform instance that contain configurations for a specific use case or domain of data. App contexts improve manageability of input and source type definitions. The Splunk platform loads all app contexts based on precedence rules. [Learn More](#)

App Context: Cisco Cyber Vision Add-on for Splunk (TA-cisco_c_...)

Host
When the Splunk platform indexes data, each event receives a "host" value. The host value should be the name of the machine from which the event originates. The type of input you choose determines the available configuration options. [Learn More](#)

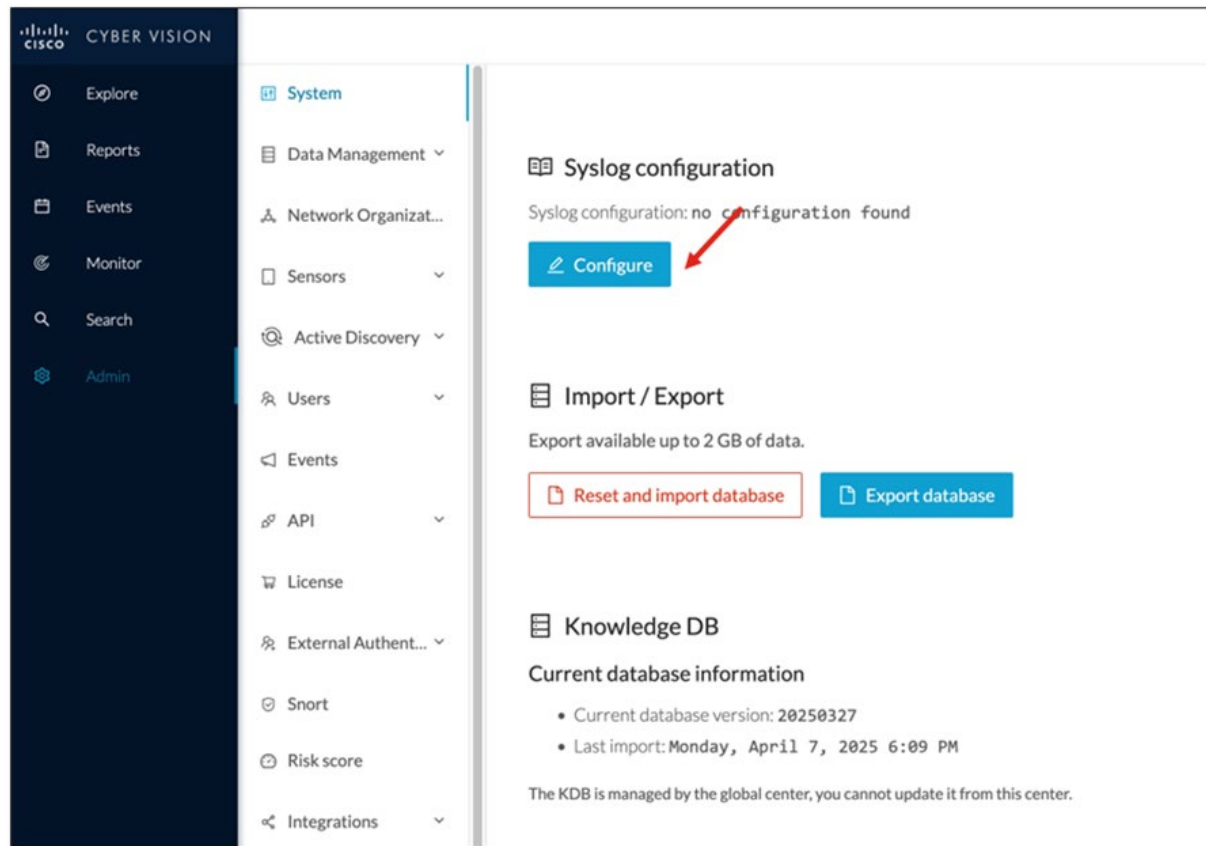
Method ? | IP | DNS | Custom

Index
The Splunk platform stores incoming data as events in the selected index. Consider using a "sandbox" index as a destination if you have problems determining a source type for your data. A sandbox index lets you troubleshoot your configuration without impacting production indexes. You can always change this setting later. [Learn More](#)

Index: Default | Create a new index

- Review information and click **Submit**.
- In Cyber Vision, navigate to Admin > System and click Configure under Sylog configuration.

Figure 162 Syslog configuration menu in Cyber Vision



- Add the Splunk IP address and specify port 514. *Note: The use of port 514 is configurable, just make sure both sides of the configuration match.*

Figure 163 Configuring syslog to be sent to Splunk

The screenshot shows the 'Syslog configuration' form. It has four main fields: 'Protocol' set to 'TCP', 'Host' set to '172.26.154.120', 'Port' set to '514', and 'Format' set to 'RFC3164/CEF'. At the bottom are 'Save configuration' and 'Cancel' buttons.

Nothing more is needed from this point, but data may not appear immediately. Users need to wait for the next Cyber Vision event to occur so a syslog message is generated.

Appendix G – Cisco ISE and Splunk Quick Deployment Guide

- In Splunk Enterprise, navigate to **Apps > Find more Apps**
- Search for “Identity Services Engine” and install both the **Splunk Add-On for Cisco Identity Services** and the **Splunk for Cisco Identity Services (ISE)** apps.

Figure 164 Cisco ISE on Splunkbase

The screenshot displays two application cards on the Splunkbase interface. The top card is for 'Splunk for Cisco Identity Services (ISE)', featuring a green 'Install' button. The bottom card is for 'Splunk Add-on for Cisco Identity Services', featuring a green 'Install' button and a small 'AddOn+' icon. Both cards provide a description of the app's functionality, installation instructions, and metadata such as category, author, download count, and release date.

Splunk for Cisco Identity Services (ISE) Install

Cisco Identity Services Engine (ISE) is a security policy management and control platform. It automates and simplifies access control and security compliance for wired, wireless, and VPN connectivity.

The Splunk for Cisco ISE add-on allows for the extraction and indexing of the ISE AAA Audit, Accounting, Posture, Client Provisioning Audit and Profiler events. This integration allows any Splunk user to correlate ISE data with other data sources (e.g. with firewall events or application data) to get deeper operational and security visibility. It also includes sample dashboards and reports for profiling, authentication, system statistics, alarms, and location awareness.

For installation instructions, see the README file within the default directory.
[Less](#)

Category: IT Operations, Security, Fraud & Compliance | Author: Jason Conger | Downloads: 17617 | Released: 2 years ago |
Last Updated: 2 years ago | [View on Splunkbase](#)

Splunk Add-on for Cisco Identity Services Install

The Splunk Add-on for Cisco ISE allows a Splunk software administrator to collect Cisco Identity Service Engine (ISE) syslog data. You can use the Splunk platform to analyze these logs directly or use them as a contextual data source to correlate with other communication and authentication data in the Splunk platform. This add-on provides the inputs and CIM-compatible knowledge to use with other Splunk apps, such as Splunk Enterprise Security and the Splunk App for PCI Compliance.

[Less](#)

Category: IT Operations, Security, Fraud & Compliance | Author: Splunk Inc. | Downloads: 30020 | Released: 2 years ago |
Last Updated: 2 years ago | [View on Splunkbase](#)

- In Splunk, navigate to **Settings > Data > Indexes**

- Click **New Index** to create a new index to store all the data received in ISE
- In Splunk, navigate to **Settings > Data > Data inputs**
- Click + **Add** new in the **TCP** row to create a syslog listener for ISE events
- Choose port 1468, and click Next
- Select **Source Type cisco:ise**, **App Context Cisco ISE (Splunk_CiscoISE)** and **Index** as the newly created index from the previous step.

Figure 165 Splunk settings for consuming syslog from Cisco ISE

Add Data

Progress: Select Source (Completed) → Input Settings (Completed) → **Review** (Active) → Done (Pending)

Review

Input Type TCP Port
 Port Number 1468
 Source name override N/A
 Restrict to Host N/A
 Source Type cisco:ise
 App Context Splunk_CiscoISE
 Host (IP address of the remote server)
 Index security

< Back Submit >

- In ISE, navigate to **Administration > System > Logging > Remote Logging Targets**
- Click +**Add** to add Splunk as a new syslog target
- Enter the connection details for the Splunk server, change the **Target Type** to **TCP SysLog** and increase the **Maximum Length** for messages to **8192**.

Figure 166 Cisco ISE remote logging target configuration

Log Settings
Remote Logging Targets
Logging Categories
Message Catalog
Collection Filters

Remote Logging Targets List > New Logging Target

Logging Target

* Name	Splunk Enterprise	Target Type	TCP SysLog
Description		Status	Enabled
* Host / IP Address	192.168.200.31		
* Port	1468	(Valid Range 1 to 65535)	
Facility Code	LOCAL6		
* Maximum Length	8192	(Valid Range 200 to 8192)	
Include Alarms For this Target	<input type="checkbox"/>		
Comply to RFC 3164	<input type="checkbox"/> ⓘ		
Buffer Messages When Server Down	<input type="checkbox"/>		
Enable Server Identity Check	<input type="checkbox"/>		
Buffer Size (MB)	100	(Valid Range 10 to 100)	
Reconnect Timeout (Sec)	30	(Valid Range 30 to 120)	

Submit Cancel

- Navigate to **Logging Categories** and for each category you are interested in sending to Splunk, click **Edit** and add the new logging target to the selected list.

Figure 167 Cisco ISE logging categories

Deployment

Licensing

Certificates

Logging

Maintenance

Upgrade

Health Checks

Backup & Restore

Admin Access

Settings

Log Settings

Remote Logging Targets

Logging Categories

Message Catalog

Collection Filters

Logging Categories

Selected 0 Total 29

Edit

	Parent Category	Category	Targets	Severity	Local Log ...
<input type="radio"/>	AAA Audit	AAA Audit	LogCollector,Splunk_AllinOn	INFO	enable
<input type="radio"/>		Failed Attempts	LogCollector,ProfilerRadius...	INFO	enable
<input type="radio"/>		Passed Authentications	LogCollector,ProfilerRadius...	INFO	enable
<input type="radio"/>	AAA Diagnostics	AAA Diagnostics	LogCollector,Splunk_AllinOn	WARN	enable
<input type="radio"/>		Administrator Authentication and Auth...		WARN	enable
<input type="radio"/>		Authentication Flow Diagnostics		WARN	enable
<input type="radio"/>		Identity Stores Diagnostics	Splunk_AllinOn	WARN	enable
<input type="radio"/>		Policy Diagnostics	Splunk_AllinOn	WARN	enable
<input type="radio"/>		RADIUS Diagnostics	LogCollector,Splunk_AllinOn	WARN	enable
<input type="radio"/>		Guest	LogCollector	INFO	enable
<input type="radio"/>		MyDevices	LogCollector	INFO	enable
<input type="radio"/>		AD Connector	LogCollector,Splunk_AllinOn	INFO	enable
<input type="radio"/>		TACACS Diagnostics	LogCollector	WARN	enable

- At this point all the connections have been successfully setup, however, we may need to provide the Splunk user with enough permission to read from the newly created security Index (this step will be non-applicable if the default Index was used).
- In Splunk, navigate to **Settings > user and Authentication > Roles**
- Select the role of “user” and enable the “Default” checkbox for the newly created index for ISE.
- Click **Save**.

Appendix H – References

- MITRE ATT&CK for ICS - <https://attack.mitre.org/matrices/ics/>
- Networking and Security in Industrial Automation Environments Design and Implementation Guide - https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/Industrial_Automation/IA_Horizontal/DG/Industrial-AutomationDG/Industrial-AutomationDG.html
- Cisco Cyber Vision Home Page - <https://www.cisco.com/site/us/en/products/security/industrial-security/cyber-vision/index.html>
- Cisco Cyber Vision Data Sheet - <https://www.cisco.com/c/en/us/products/collateral/se/internet-of-things/datasheet-c78-743222.html>
- Cisco Cyber Vision Architecture Guide - https://www.cisco.com/c/dam/en/us/td/docs/security/cyber_vision/Cisco_Cyber_Vision_Architecture_Guide_V2-00.pdf
- Cisco Cyber Vision GUI Administration Guide - https://www.cisco.com/c/en/us/td/docs/security/cyber_vision/publications/GUI-Administration-Guide/b_cisco-cyber-vision-GUI-administration-guide.html
- Cisco Identity Service Engine Home Page - <https://www.cisco.com/site/us/en/products/security/identity-services-engine/index.html>
- Cisco Identity Services Engine Administrator Guide - https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2.html
- Cisco SecureX Home Page - <https://www.cisco.com/site/us/en/products/security/securex-platform/index.html>
- Splunk Enterprise docs - <https://docs.splunk.com/Documentation/Splunk>
- Verizon 2025 Data Breach Investigations Report - <https://www.verizon.com/business/resources/reports/dbir/>
- Waterfall 2025 OT Cyber Threat Report - <https://waterfall-security.com/ot-insights-center/ot-cybersecurity-insights-center/2025-threat-report-ot-cyberattacks-with-physical-consequences/>

Appendix I – Acronyms and Initialisms

Table 7 lists the acronyms and initialisms used in this design guide.

Table 7: 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
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 and 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

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)