Industrial Cybersecurity for Distributed Field Networks

Designing a secure WAN infrastructure to protect field assets and build cyber-resilient industrial operations

Contents

Introduction

Across all industries, organizations need advanced, agile, and secure WAN infrastructures to connect distributed Operational Technology (OT) assets to control centers and unlock the potential of digitization. Whether you are connecting roadways assets; first responder or public transport vehicles; water, oil, or gas infrastructures; renewable energy resources; power substations; Electric Vehicle (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 (NGFW) 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.

This solution brief is a subset of the Cisco Validated Design guide (CVD) on Industrial Security and is written to be a summary of the design guidance and capabilities contained within the guide. For more information on any of the technologies and best practices found in this solution brief, see the Cisco Industrial Security Design Guide.

Benefits

- Combining the best of OT networking with OT security to build a cyber-resilient industrial network
- Using defense-in-depth to distribute security across the architecture
- Controlling all traffic to and from remote sites with NGFW capabilities
- Gaining visibility into connected assets, vulnerabilities, and activities
- Segmenting and prioritizing the most critical traffic on the network across both LAN and WAN networks
- Providing zero-trust remote access to critical assets, empowering the line of business to manage and troubleshoot assets remotely

Cyber resiliency vs. cybersecurity

While we often talk about cybersecurity, 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 organization's ability to maintain its critical operations even in the face of cyberattacks.

Cybersecurity is, of course, part of a cyber-resilient architecture. Capabilities such as firewalls, segmentation, and the implementation of a zero-trust model mean 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, cybersecurity 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 maintains latency metrics when network paths go down, management plane security ensures that 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.

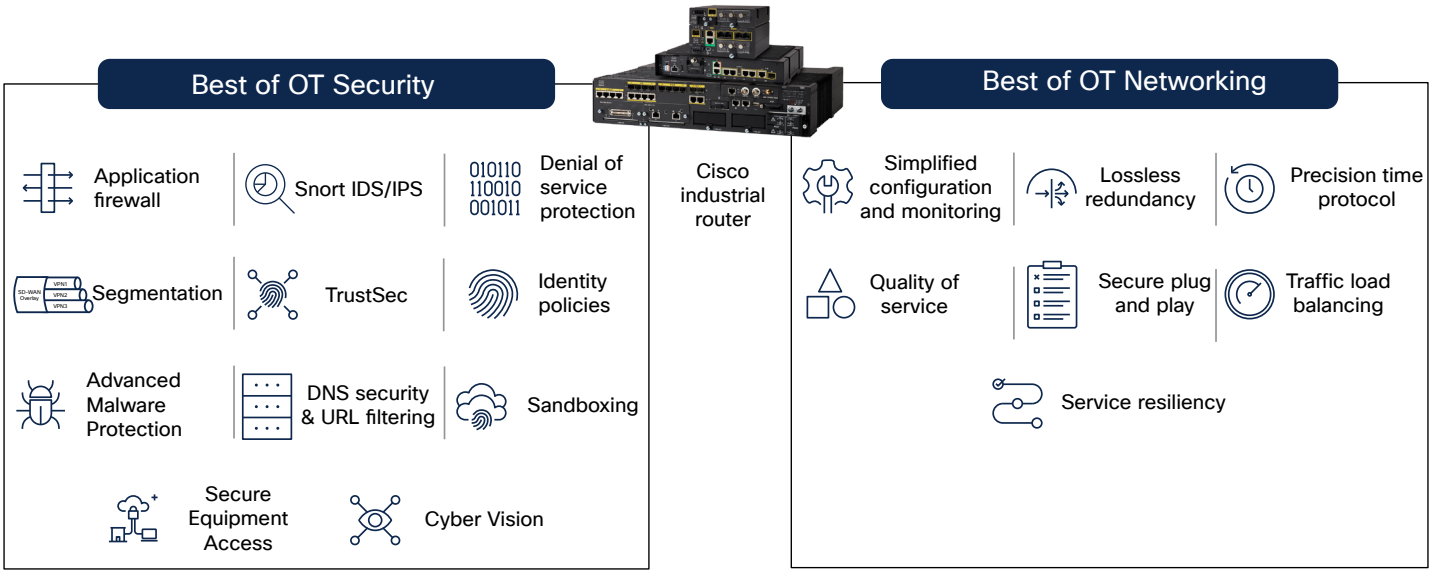

Figure 1. Best of OT networking and OT security with a Cisco industrial router

Understanding defense-in-depth

A major design philosophy in the Cisco Industrial Security Design Guide is the use of defense-in-depth to protect critical infrastructure. 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 defenses 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, they may also be able to bypass the next.

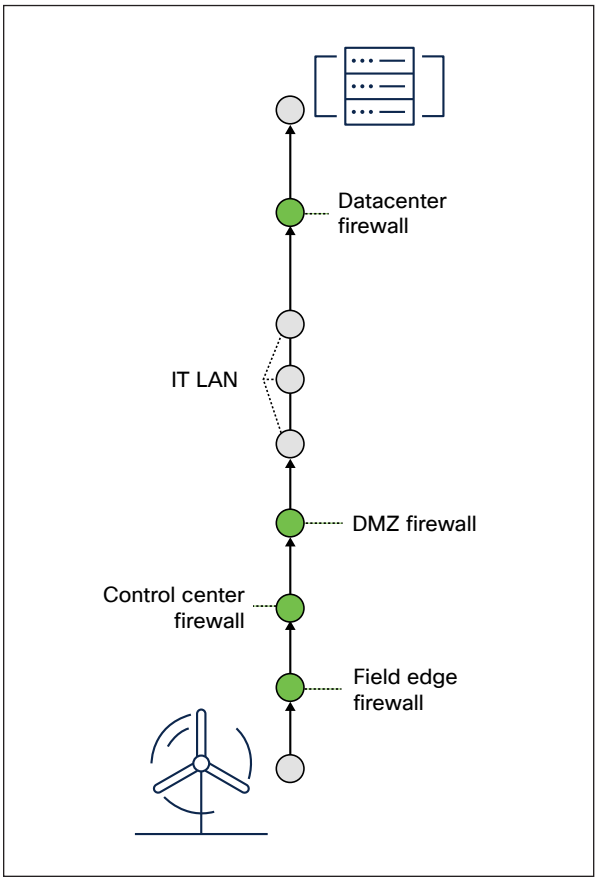

Figure 2. A typical data flow if using firewalls for defense-in-depth architectures

Organizations should implement appropriate technology for the use cases they need to protect. When implementing cybersecurity at the field edge, administrators should ask themselves:

- Is this traffic going directly to the internet, or should those policies instead be deployed at the internet breakout?
- Is Transport Layer Security (TLS) decryption required at every hop in the network, or should computationally expensive actions be reserved for strategic points in the network?
- Is an Intrusion Prevention System (IPS) required across all boundaries, subjecting the traffic to the same set of signatures that it has already been through?

Cisco's Industrial Security Design Guide does not have the scope to cover all possible use cases that an industry will face but rather acts as a guiding set of principles to build a cybersecurity 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 cybersecurity strategy, but there are often multiple enforcement points between field assets and their intended destinations. Use core security practices like Authentication, Authorization, and Accounting (AAA), port security, segmentation, and basic firewalling at the edge to ensure that only trusted devices cross the boundary using trusted ports, and supplement that with advanced inspection at the hub location. This approach 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 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 need to be done only once. It is important that 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) tools**. A SIEM tool such as Splunk® 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.

Using a Cisco industrial router to protect critical infrastructure

Cisco 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 the work of deploying and managing a large and complex WAN infrastructure from a central location.

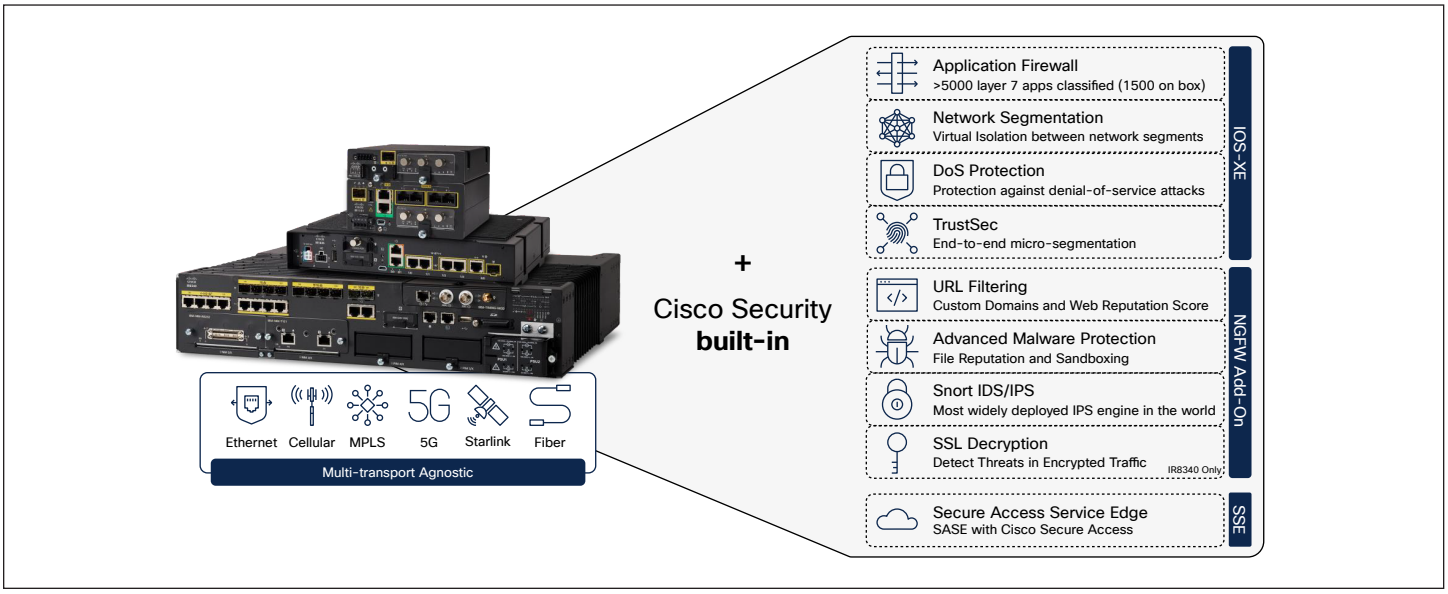

Figure 3. Cisco industrial router portfolio features

Cisco industrial routers also come with comprehensive 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 updated threat intelligence
- Malware protection and sandboxing
- URL filtering
- Integration with a Security Service Edge (SSE)

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

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 or IP address
- VLAN
- UDP or TCP port
- Objects
- User ID via Cisco Identity Services Engine (ISE)
- Fully Qualified Domain Name (FQDN)
- Security Group Tag (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: port 502
- DNP3: port 20000
- Ethernet/IP: port 44818
- OPC UA: port 4840
- MMS: port 102

In instances where NBAR2 may not recognize the protocol by name, security administrators can choose to open or close ports to achieve the same level of control. Alternatively, a custom application can be created to enable NBAR2 to recognize 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: In addition to creating firewall rules, NBAR2 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 became compromised, the others should not be affected.

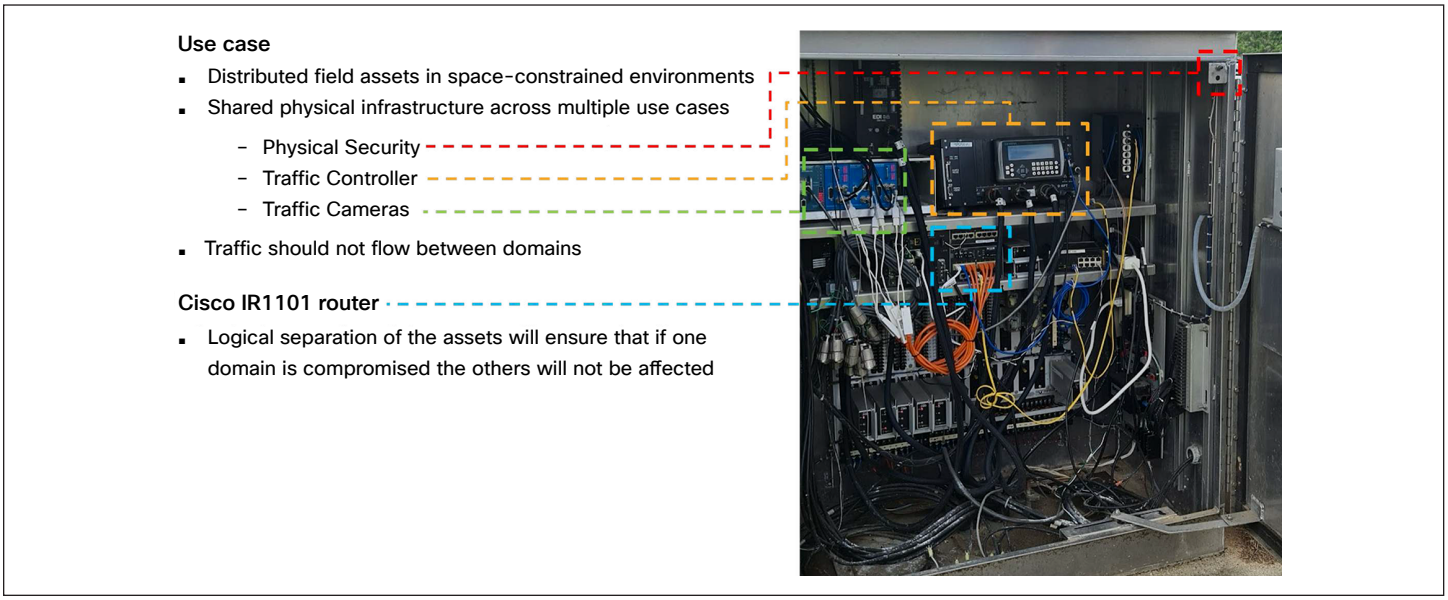

Figure 4. Example of segmentation requirements in a traffic cabinet

Note: In the context of Cisco Catalyst SD-WAN, the terms 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.

Authentication, authorization, and accounting

Most field networks currently have no port authorization enabled. Any device can be connected to an available Ethernet port and be given network service. This is problematic because distributed networks such as Intelligent Transportation Systems (ITS) and the utility grid are uniquely exposed due to weak physical security, leaving them readily accessible to the public—and bad actors. Thus, instead of having a "default open" posture, we move to the best practice of "default closed." By default, any connected devices have no network service and are given network service only once the network has established 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 ability to provide a username or password. Cisco Identity Services Engine (ISE) provides the ability to use 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. 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 Cisco IOS® XE enables administrators to restrict input to an interface by limiting and identifying the MAC addresses of the endpoints allowed to access a port. When port security is in use, 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 the port.

TrustSec

The TrustSec technology defines policies using logical device groupings known as Security Group Tags (SGTs). An 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, making them a key component of a defense-in-depth architecture intended to protect critical infrastructure.

Most use cases will be subject to multiple policy enforcement points when traversing the network, and as security architects, we need to coordinate 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 the IP address, which may change depending on the part of the network you're on and how many times you did Network Address Translation (NAT), an SGT provides a business context for a given device. For example, I may need to create policies for my employee-managed device (e.g., SGT 10) at a roadside cabinet, at the branch, and at the data center. Or I may need to create a policy for a traffic signal controller (e.g., SGT 20) at the cabinet, across the WAN, in the control center, and across the data center. Regardless of the IP address, policy for the device can be consistent across the architecture, and it becomes much easier to correlate log information across the entire network.

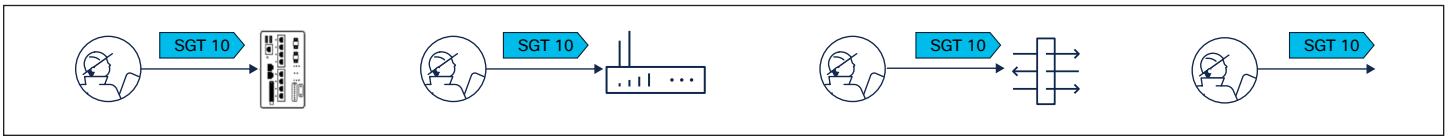

Figure 5. Importance of TrustSec and SGTs

TrustSec has three main functions—classification, propagation, and enforcement—which the Cisco Industrial Security Design Guide covers in detail. The level of support for each function with Cisco industrial routers can be found in the table below.

Table 1. TrustSec capabilities per Cisco industrial router platform

Router model	Classification		Propagation		Enforcement	
	Static	Dynamic	Inline	SXP	Zone-based firewall (ZBFW)	Security group ACL (SGACL)
IR1101	IP to SGT, subnet to SGT, Switch Virtual Interface (SVI) to SGT	Yes	Tunnel interface, Layer 3 ports	Speaker, listener	Yes	No
IR1800	IP to SGT, subnet to SGT, SVI to SGT	Yes	Tunnel interface, Layer 3 ports	Speaker, listener	Yes	No
IR8340	IP to SGT, subnet to SGT, SVI to SGT	Yes	Tunnel interface, Layer 3 ports	Speaker, listener	Yes	No*

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

TrustSec is an important technology when implementing a defense-in-depth security architecture. As we discussed previously, a misconception involving defense-in-depth is that it uses the same technology stack in multiple parts of the network. Figure 6 depicts a more realistic view of the security stack.

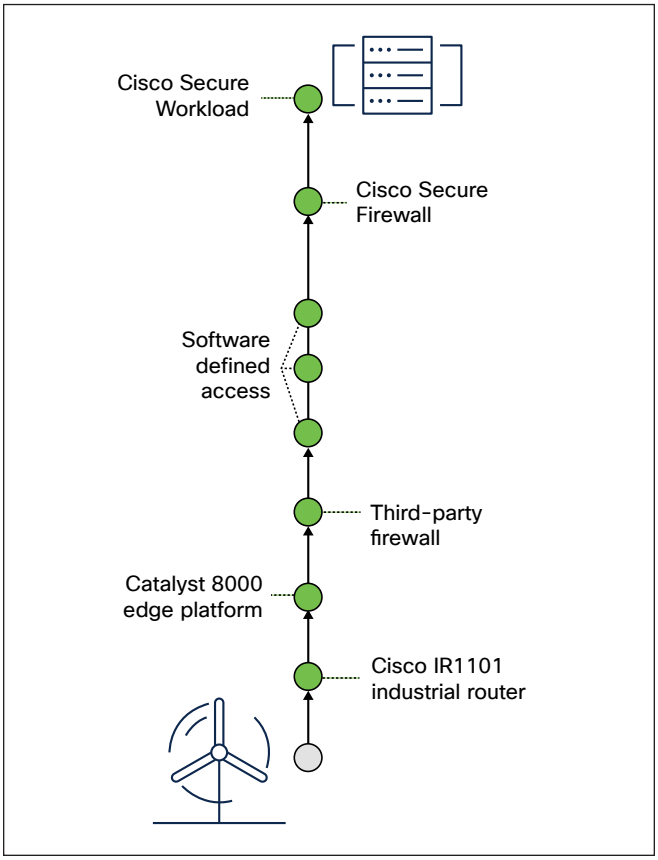

Figure 6. Example of a policy enforcement stack between IT and OT networks

Let's look at a flow originating from the OT network and destined for the data center:

- Traffic from an OT device must pass through a router/firewall/gateway at the field network edge. A policy decision must be made here to determine whether that device's traffic is permitted.
- The traffic traverses the WAN, whether it 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. This diagram depicts Cisco Software-Defined Access, which comes with its own policy engine.
- Upon reaching the data center, the traffic typically encounters another firewall.
- Data center security is also a hot topic, with products such as Cisco Secure Workload or Cisco Hypershield™ potentially being used here.

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

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 DoS (DDoS) attack that also uses ICMP but in a different way. The attacker sends ICMP echo request packets to the broadcast network using a spoofed source IP address set to the IP address of the target victim. Since the source IP address has been spoofed, the target now receives a large volume of ICMP echo replies from multiple devices and is overwhelmed by the flood of traffic, leading to network congestion.

To protect against DoS attacks, we recommended protecting both the data plane and the control plane packets.

Quality of service

QoS refers to a set of technologies and techniques used to manage and prioritize 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 are sensitive to latency and dropped packets.

Traditional QoS has its limitations because it can’t predict the changing bandwidth on the link. With adaptive QoS, the shapers at the edge can adapt to the available WAN bandwidth, including across 4G/LTE connectivity.

Control plane policing

Control Plane Policing (CoPP) improves security on Cisco industrial routers 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 are 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 the 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).

NGFW add-on

To expand upon the stateful firewall capabilities within the zone-based firewall, the Cisco Catalyst IR1835 and IR8340 can host an NGFW add-on 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 controls access to websites based on categories, reputation, and custom policies to prevent exposure to malicious sites.

Note: File detection and sandboxing rely on cloud connectivity.

Snort IDS/IPS

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 situations:

- A newly discovered vulnerability poses an immediate risk, and in this case, the fix or the patch is not available.
- The vulnerability is not highly critical, so it's not worth patching it outside the usual patch window.

In both cases, one must accept the interim risk and wait either 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 "virtual" fix to known vulnerabilities in the software until it has been patched or updated. Virtual patching is typically done by leveraging an IPS such as Snort. Hosting Snort in a Cisco industrial router enables the use of IPS policies that can prevent known vulnerabilities from being exploited over the network.

While Cisco industrial routers provide thousands of preconfigured Snort rules for access control, Snort also has the ability 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 offers "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**.

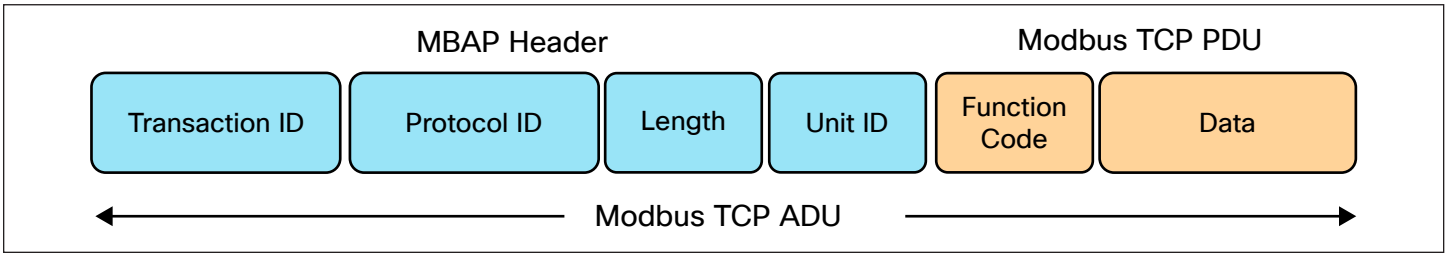

Figure 7. Modbus TCP packet structure

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)

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 over 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, URL filtering can give security administrators 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 cyberthreats. Detecting and removing malicious files before they enter your network is key to preventing breaches. Cisco 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 industry's largest collection of file reputation intelligence.

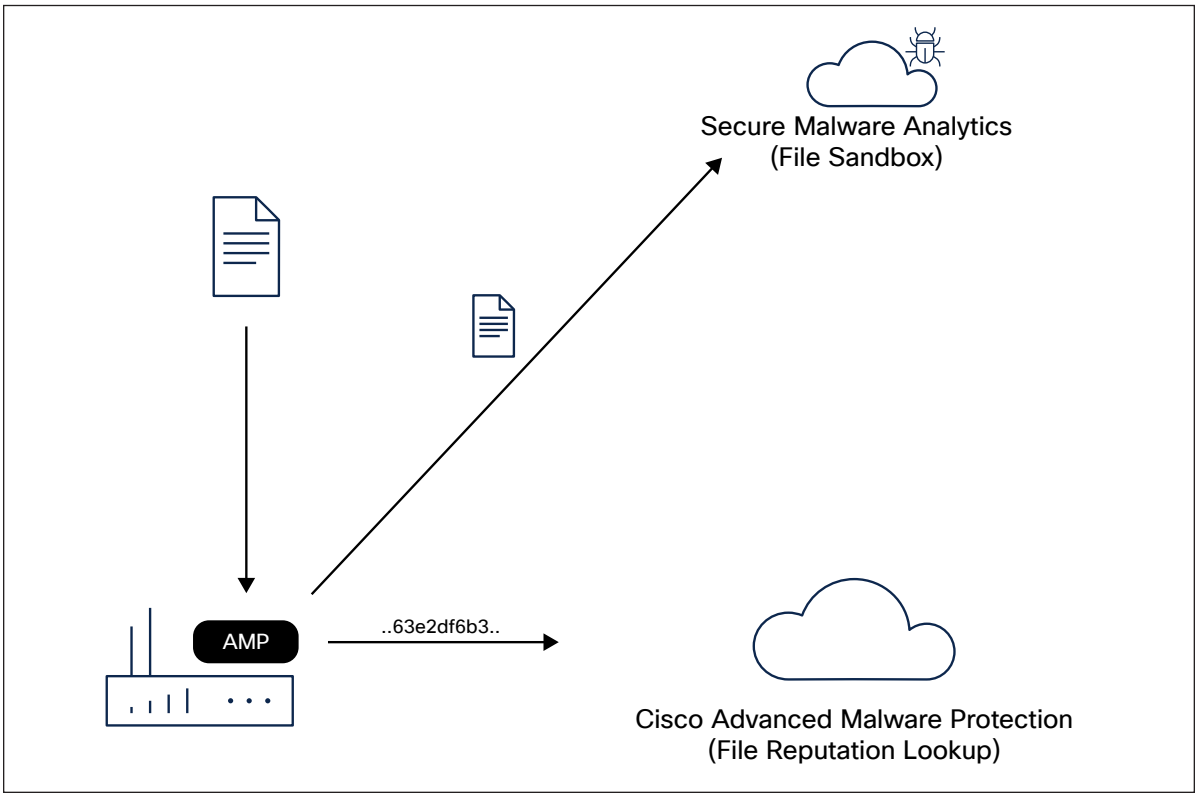

Figure 8. File inspection by the Cisco AMP and Secure Malware Analytics clouds

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.

Application hosting in IOx

Cisco IOx is an application enablement platform that enables 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 additional security applications such as Cisco Cyber Vision and Cisco Secure Equipment Access.

Asset visibility with Cisco Cyber Vision

As industrial networks can be quite old, be widely dispersed, and involve many contractors, operators often don't have an accurate inventory of what's on the network. Gaining comprehensive visibility into what's connected to your field network is critical both to efficiently manage resources and also to assess risks and prioritize what needs to be fixed to reduce your exposure to cyber risks.

Visibility also helps when implementing network segmentation by giving organizations a good understanding of the normal state of the OT network and distinguishing attacks from transient conditions or normal operations within the environment. Whether using a risk-based approach, a functional model, or other organizing principles, grouping components into levels, tiers, or zones is a precursor activity before organizations can consider applying policy to protect and monitor communication between zones. Implementing network monitoring in a passive mode and analyzing the information to differentiate between known and unknown communications may be a necessary first step in implementing security policies.

Cisco Cyber Vision is a software solution built into Cisco industrial routers and Cisco Industrial Ethernet switches to give comprehensive visibility and help implement security policies in field industrial networks. It consists of multiple sensors that perform Deep Packet Inspection (DPI), protocol analysis, and intrusion detection within your industrial network, and also includes 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.

The Cyber Vision sensor software can be hosted on a Cisco industrial router as an IOx application running in a dedicated CPU core to avoid impacting network performance. The sensor passively captures and decodes network traffic using DPI of industrial control protocols and can send safe active queries to assets using the semantics of the protocols at play. Because the Cyber Vision sensor is deployed directly in the router, there is no need to install dedicated appliances in field cabinets, where space can be an issue, or to build an out-of-band SPAN collection network. Because it decodes network traffic at the edge, it sends only lightweight metadata to Cyber Vision Center, adding a load of just 2% to 5% to your industrial network. This is critical where bandwidth is either limited, such as in distributed networks, or costly, such as over cellular interfaces.

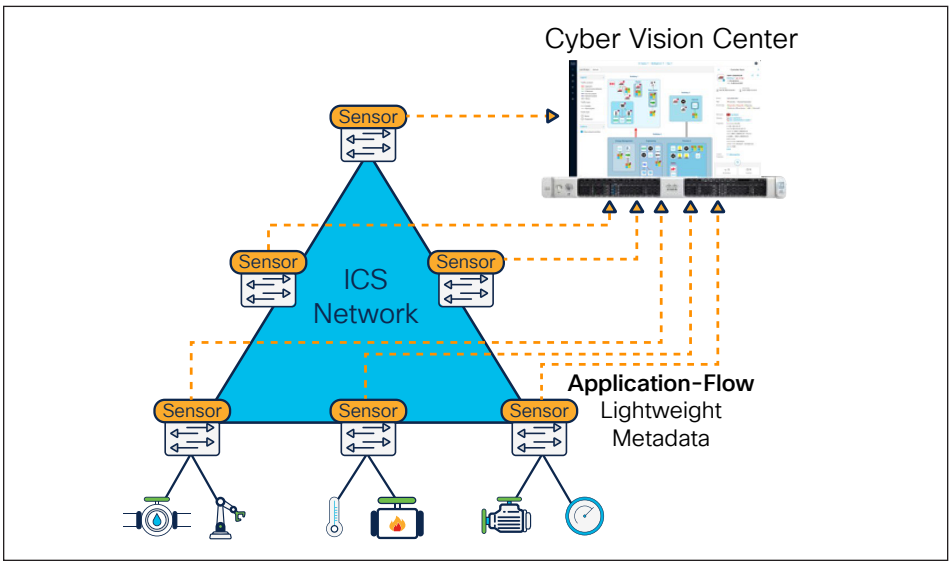

Figure 9. OT visibility using sensors embedded in the network infrastructure

Zero-trust remote access with Cisco Secure Equipment Access

Remote access is key to managing, maintaining, and troubleshooting field assets without time-consuming and costly site visits. Zero-Trust Network Access (ZTNA) solutions are gaining increased momentum for enforcing robust secure remote access in industrial networks. In general, when users log into a VPN, they are granted complete access to the entire network. ZTNA solutions verify users and connect them only to assets or applications they are authorized to access rather than to the network as a whole.

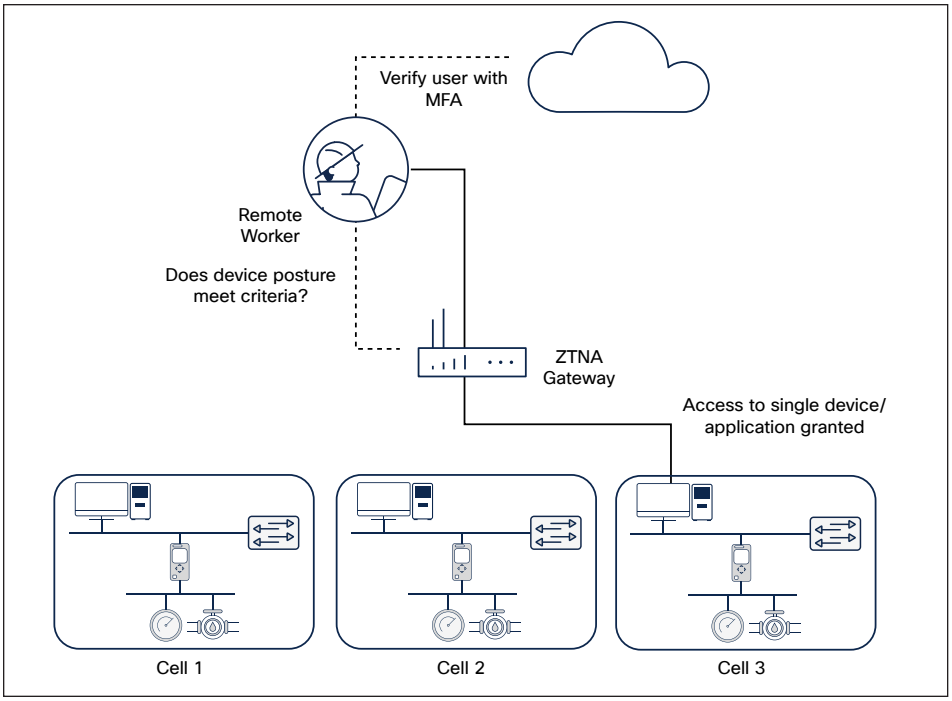

Figure 10. Using ZTNA to connect to a single device on the industrial control network

Because Multifactor Authentication (MFA), Single Sign-On (SSO), and security posture checks are built in, ZTNA can help address common security challenges in the workforce, such as phishing, malware, credential theft, remote access, and device security. This is done by securing the three primary factors that make up the workforce: users, their devices, and the applications they access.

Cisco Secure Equipment Access (SEA) is a ZTNA solution specifically designed for OT workflows. It is a cloud service that runs in Cisco's industrial network equipment, making it very simple to deploy and manage at scale. It empowers operations teams, vendors, and contractors to easily connect to remote OT assets such as traffic signal controllers, in-vehicle dispatch systems, cameras, and other systems deployed in the field, while enforcing least-privilege access controls based on identities and contexts: users can access only specific resources at specific times, and never the entire network.

With Cisco SEA, users just need a web browser to access remote assets. They connect to the SEA cloud portal, where they are authenticated and offered access only to the resources you have chosen, using only the protocols you specify, and only on the day and time you allow. 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 becomes cleartext only at the other end of the proxy (SEA Agent). Precautions still need to be taken 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 ICMPbased 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 those channels.

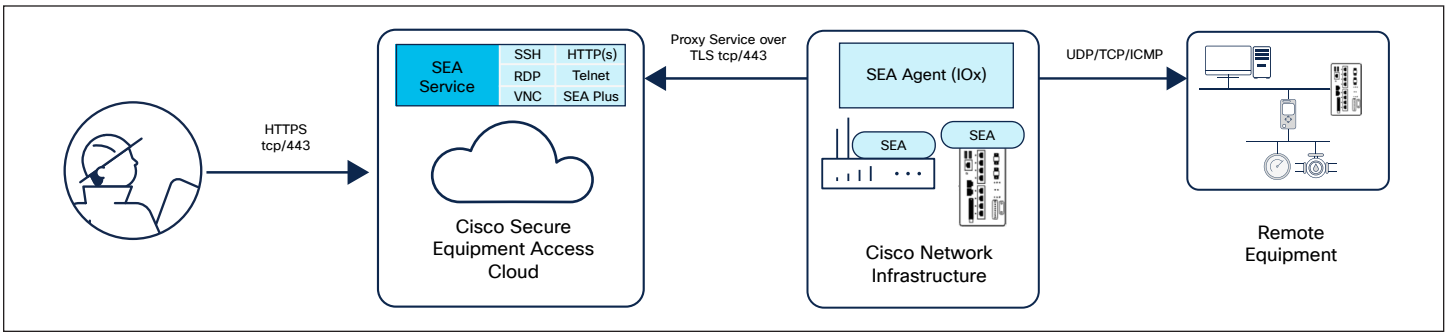

Figure 11. Components of Cisco SEA

The SEA cloud portal is a ZTNA trust broker where policy definition and enforcement are done. It gives security teams a single interface to manage users, assets, and policies for all sites. Session monitoring, session logs, and the ability to join and terminate a session are offered to admin users.

SEA's ZTNA gateway is a software feature in Cisco industrial routers and Cisco Industrial Ethernet switches. Not only does it eliminate the need for dedicated hardware, but it also simplifies secure remote access to OT assets sitting behind NAT boundaries. With SEA, Cisco industrial routers become comprehensive security appliances with unmatched cybersecurity capabilities in addition to advanced routing and SD-WAN features as well as full-featured NGFW.

Supported platforms

All Cisco 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 at least 8 GB of memory. The NGFW add-on for industrial routers provides Snort IDS/IPS, reputation-based URL filtering, and malware protection.

Table 2. Security features per Cisco industrial router platform

Feature	IR1101, IR18xx	IR1835	IR8340
Stateful packet inspection	Yes	Yes	Yes
Application visibility and control	Yes	Yes	Yes
VPN	Yes	Yes	Yes
Segmentation	Yes	Yes	Yes
DoS protection (QoS, CoPP, etc.)	Yes	Yes	Yes
AAA	Yes	Yes	Yes
Port security	Yes	Yes	Yes
TrustSec	Yes*	Yes*	Yes
IDS/IPS	No	Yes	Yes
Malware protection	No	Yes	Yes
File sandboxing	No	Yes	Yes
FQDN filtering	Yes	Yes	Yes
Reputation and category web filtering	No	Yes	Yes
TLS decryption	No	No	Yes
Cisco Cyber Vision sensor	Yes	Yes	Yes
Cisco SEA gateway	Yes	Yes	Yes

*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.

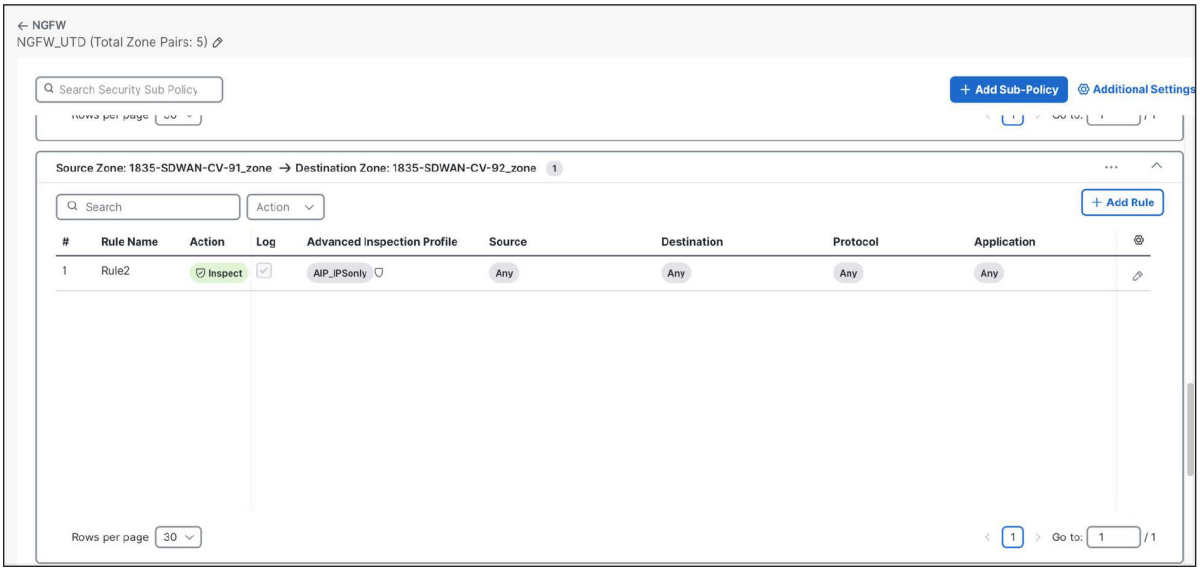

Figure 12. Firewall policy manager in Cisco Catalyst SD-WAN Manager

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 rules to allow, deny, or inspect the traffic. When using the inspect action, any traffic that matches this policy will be sent the NGFW add-on for advanced threat inspection.

The importance of 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 cyberthreats, 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. However, the IT team often 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.

With Cisco Catalyst SD-WAN, when a Cisco industrial router is provisioned or replaced in a field network, a device configuration containing all the security policies it needs to secure operations from day one is pulled into the router. The capabilities listed above, such as DoS protection, port security, and firewall rules do not need to be configured as an afterthought, reducing the risk of misconfiguration that could lead to a breach.

Summary

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 OT requirements plus a comprehensive networking and cybersecurity portfolio is a rare combination.

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.

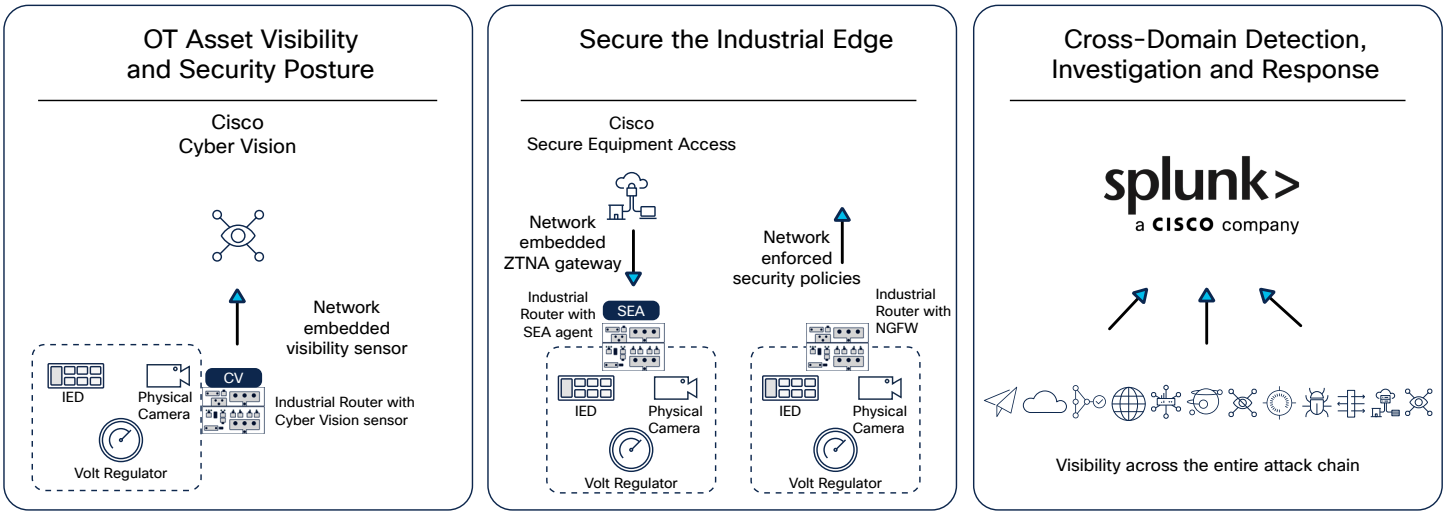

Figure 13. Cisco Industrial Threat Defense for distributed field networks

Talk to a Cisco sales representative or channel partner about how Cisco can help you secure your field industrial network. Visit cisco.com/go/iotsecurity or cisco.com/go/iotrouters to learn more.

C22-5284249-00 07/25