paloalto® NETWORKS

STRATA™ BY PALO ALTO NETWORKS

Reference Architecture for OT and IoT Device Security in Industrial Control Systems

This document provides a high-level reference architecture for OT and IoT device security in ICS using Palo Alto Networks Next-Generation Firewalls, Cortex Data Lake, and the IoT Security service.

Content

Introduction

Digital transformation of operational technology (OT) environments has led to a rapid growth in the number of IP-connected devices in environments such as ICS, SCADA, and DCS. These devices include those used for core industrial processes such as HMIs, PLCs, RTUs, and IEDs. As industrial IoT comes online, there will also be some devices such as IoT gateways that communicate outside of the traditional OT perimeters to reach public clouds or even multi-clouds. These cyber-physical systems are also blended with more general IP-connected devices such as CCTV cameras, thermostats, badge readers, multi-function printers, and IP phones intended to be used on-premises, yet having the capabilities to be connected to the internet. Given this growth and the associated security risks that come with these devices, securing OT and IoT devices in industrial control systems (ICS) has become a high priority for many industrial companies.

Such companies are actively looking at more effective ways not only for taking stock of all of their cyber-physical devices but also for gaining more intelligence about the devices themselves and how they are communicating over the network. With this knowledge, they are better able to assess baseline device applications, associated risk, and what kind of network access policies or remediative measures should be set for these devices. This paper presents a high-level reference architecture for OT and IoT device security in ICS with such priorities in mind. It explains how to harness the powerful capabilities of the Next-Generation Firewall, Cortex Data Lake, and the IoT Security service to deliver comprehensive device detection and threat prevention capabilities across the entire OT and IT domains.

Security Design Objectives

Before presenting the reference architecture, we will first look at the baseline OT/ICS environment, and then we will cover the high-level design objectives.

Baseline OT/ICS Network

While ICS architectures are known to vary potentially quite dramatically from plant to plant, it is possible to create a general baseline network onto which we can overlay our reference security architecture. Figure 1 shows this generic ICS topology within a plant comprising networking equipment and network security devices as well as some of the common devices connected at the access layer.

Starting at the top, we have a pair of high availability (HA) firewalls that secure traffic between IT (Level 4), the DMZ (Level 3.5), and OT (Levels 0-3). Within OT, we have a three-tier switching architecture with an HA pair of core switches interconnected with distribution switches below and followed by access switches beneath those. The access switches provide connectivity for various IP-connected devices within the subcategories of OT, IoT/IIoT, and IT. Examples of these devices include:

- **OT devices**—HMIs, PLCs, IEDs, plant historians, distributed I/O, physical access control systems (PACS), safety systems
- **IoT devices (usually non-internet-facing)**—CCTV cameras, HVAC equipment, PACS
- **IIoT devices—IIoT gateways, smart sensors, unmanned** vehicles
- **IT devices**—workstations, servers, printers, tablets, mobile phones

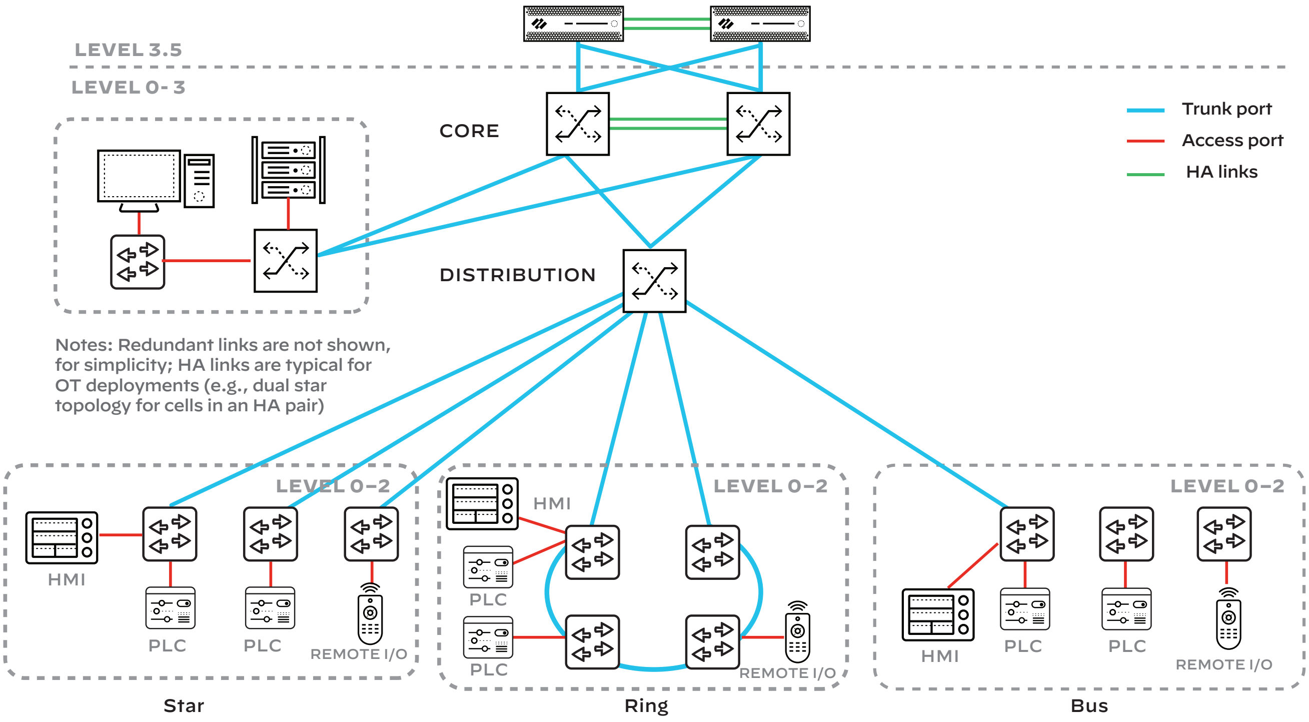

Notes: Redundant links are not shown, for simplicity; HA links are typical for OT deployments (e.g., dual star topology for cells in an HA pair)

Figure 1: Generic ICS network

While different network topologies could be found at the access layer, including star/dual-star, ring, and bus, they all connect back up to the distribution switch or to each other via Layer 2 802.1q trunk ports represented in blue. These trunk ports carry the different VLANs, which provide coarse-grained segmentation of the various subsystems or security zones. The approaches for creating zones could be, for example, those prescribed in the ISA 62443 standard. The red lines represent the access ports for the individual devices. Management of VLANs and inter-VLAN routing is handled at the distribution switch in a router-on-a-stick configuration.

NOTE: To reduce the complexity of the drawing, we have intentionally omitted some of the redundant links commonly found in OT environments. For example, each access switch in a dual-star configuration will typically have a second connection back to a second line card on a distribution switch or even a second distribution switch (HA pair).

Security Design Objectives

Having defined our baseline ICS/OT network, we can now discuss the main security objectives. The security architecture we implemented should do all of the following.

Enumerate and profile all IP-connected devices in the ICS. We want to identify all devices with an IP address. These include not only those communicating north-south but also devices communicating east-west.

Support Zero Trust architecture and approach, down to device-level policy. We want to be able to realize the best practices prescribed by Zero Trust, including the creation of a segmentation gateway that supports granular, business-driven policy. We want to be able to create flexible policies based on application, user, content, and, of course, device. For device policy, we want to have control over parameters such as category, type, vendor, model, OS family, and OS version.

Provide 24/7 real-time monitoring and automatic detection of anomalous/malicious behavior. It is humanly impossible to monitor every individual device in OT and respond to malicious/anomalous incidents in a timely manner. A key design objective is to implement an architecture that utilizes machine learning to automate the detection and response over malicious device traffic.

Minimize the number of devices needed to implement security controls. OT security teams need to be mindful of the risk of "appliance sprawl" where multiple point products are stitched together to try to achieve the desired controls. This comes with many costs and operational overhead as well as higher risk of gaps in security controls. We want a single solution to cover all OT/IIoT/IoT devices rather than having multiple point solutions. Furthermore, we want the sensing, enforcement, and threat prevention functionalities to be unified in the same device as much as possible.

Support integration with minimal disruption to existing OT networks. The approach should be easy to overlay onto existing OT topologies, with minimal changes to the existing network and disruption of service. Furthermore, it should support the high availability, performance, and form factor requirements of OT.

Scale easily. The solution must be able to scale easily from small control systems up to global, multi-plant OT infrastructure.

Implement an approach that supports IT-OT convergence. Ideally, the approach taken should be one that could be applied in both IT and OT to facilitate consolidation of the security architecture and collaboration between IT and OT. It should be able to identify IT, OT, IoT, and IIoT devices.

With the objectives in place, we can now look at the reference architecture and how it addresses them.

Security Reference Architecture

Architecture Overview

Shown in figure 2 is the high-level security architecture overlaid onto our baseline ICS design. It comprises three salient components: Palo Alto Networks Next-Generation Firewalls (NGFW) placed at different strategic locations, Cortex Data Lake service for aggregating device traffic from the NGFWs, and the IoT Security service. We will now cover these components and the roles they play in the overall architecture.

Next-Generation Firewall

Single-Pass Architecture

The single-pass parallel processing (SP3) architecture of the NGFW enables Zero Trust by providing granular visibility and control at the level of applications (App-ID), users (User-ID), and contents (Content-ID). With the advent of PAN-OS® 10.0, the NGFW also has the capability of identifying and controlling devices (Device-ID). It allows fine-grained and least privilege segmentation of OT that aligns with the zones and conduits model of ISA 62443, which is a parallel to the security zones and segmentation gateway model of the Zero Trust architecture.

With App-ID™, industrial protocols and applications such as Modbus, CIP EtherNet/IP, and OPC can be identified and controlled even down to the function-code level. The latest list of supported application identifiers for ICS/OT is covered in our App-ID for ICS/SCADA tech brief.[1] User-ID™ enables role-based access control while Content-ID™ allows visibility and firewall policy based on content. Device-ID™ works with our IoT Security service to detect devices and implement device-level policy based on several parameters, including device category, type, vendor, model, operating system (OS) family, and OS version.

Subscription Services

While network segmentation and least privilege controls with strong authentication dramatically increase the protection surface, there is still a potential for malicious traffic to traverse these allowed communication channels through the NGFW. To address this challenge, various threat services are offered as subscriptions on the firewall. These services work

[1] App-IDs for ICS/SCADA Technical Brief, https://www.paloaltonetworks.com/resources/whitepapers/app-ids-industrial-control-systems-scada-networks.

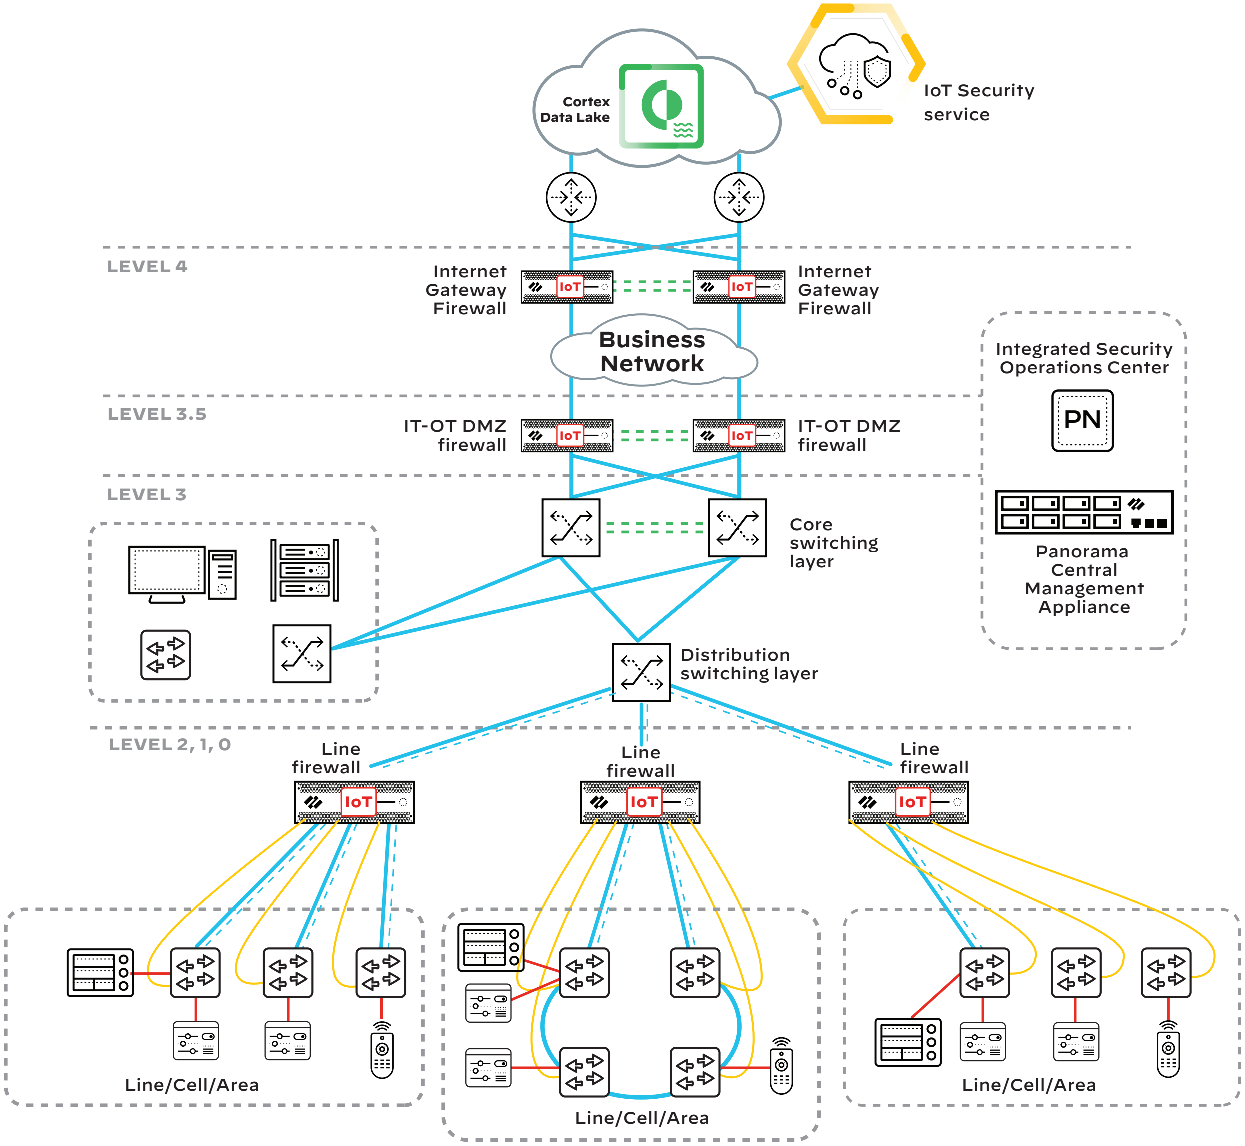

Figure 2: Reference architecture for device security in ICS/OT

in parallel in the single-pass engine to improve performance and eliminate the challenges associated with having multiple point solutions, such as increased operational overhead and reduced security effectiveness. Here are short descriptions of these services.

IoT Security

The IoT Security service provides IoT device identification, risk assessment, enforcement, and incident response. We will discuss this in more detail shortly.

Threat Prevention

Threat Prevention provides an IPS/IDS service to protect against exploits, malware, and command-and-control (C2) traffic. It helps address the ever-present challenge of protecting unpatched and unpatchable OT systems, such as legacy HMI, Historian, PLC, and engineering workstations.

WildFire

WildFire® next-generation malware analysis and sandboxing technology uses community-sourced threat intelligence and advanced malware analysis to automatically and quickly detect and prevent unknown threats. This is great for stopping the next zero-day variants of advanced threats, such as OT-specific malware, remote access trojans, and ransomware. WildFire could, for example, help with detecting and stopping the next zero-day variation of malware like Crash-Override, Triton, BlackEnergy and Petya.

URL Filtering

URL Filtering enables safe internet access from OT by automatically preventing attacks that leverage the web as an attack vector, including phishing links in emails, phishing sites, HTTP-based attacks, malicious sites, and pages that carry exploit kits.

DNS Security

The DNS Security service applies predictive analytics for automated protections to thwart attacks that use DNS.

GlobalProtect

GlobalProtect™ network security for endpoints extends the prevention capabilities of the NGFW to mobile workers, regardless of their location. This is ideal for ensuring consistent access control and threat services when using mobile devices for OT, such as maintenance laptops, tablet-based HMIs, and smartphones used by operators, engineers, plant managers and third-party workers roaming the plant floor.

Data Loss Prevention

Data Loss Prevention discovers, monitors, and protects an organization's sensitive data, such as PII and intellectual property, minimizing the risk of data breaches and enhancing both data privacy and compliance.

SD-WAN

SD-WAN provides secure and reliable software-defined WAN that delivers an optimal user experience for cloud applications from your branches and retail locations, without compromising security.

For this paper, we will be focusing on the IoT Security service. Learn more about the different subscriptions on the NGFW by accessing the "Subscriptions" section of the PAN-OS 9.1 Administrator's Guide.[2]

Interface Deployment Modes

The NGFW supports multiple physical and logical interface deployment modes. Users can mix and match these modes within a single appliance. This provides the high level of flexibility required when trying to address different OT use cases. We will see shortly how this flexibility enables the hybrid interface deployment mode discussed in this paper. Table 1 shows the different physical and logical interface deployment modes supported by the NGFW.

Table 1: NGFW-Supported Physical and Logical Interface Modes

Physical (Ethernet) Interface Modes	Logical Interface Modes
Tap mode	VLAN
Virtual wire	Loopback
Layer 2	Tunnel
Layer 3	Decrypt mirror
Aggregate interfaces	
HA	

For our core OT reference architecture, we will be focusing on the Layer 2 physical interface—with VLAN logical interfaces to secure north-south traffic—that is in line with the firewall and tap mode to passively monitor east-west traffic. Learn more about the different interface modes of the NGFW by accessing the "Interface Configuration" section of the PAN-OS 9.1 Administrator's Guide.[3]

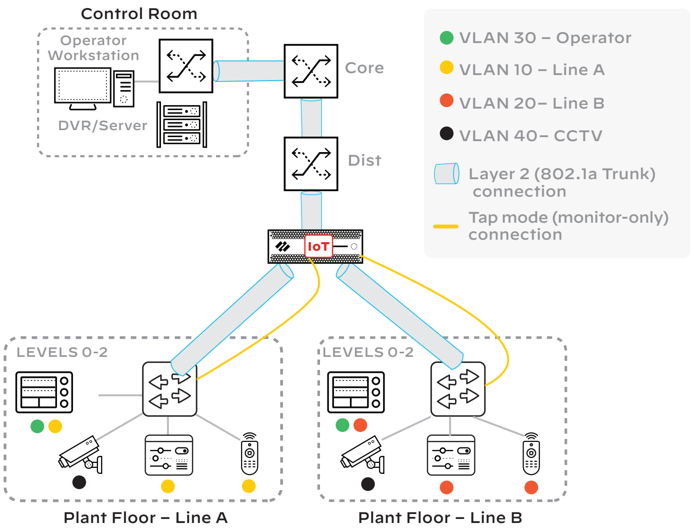

Figure 3: ICS firewalls monitoring and protecting levels 3, 2, and 1

Interface and Services Configuration

Figure 3 shows a zoomed-in view of a simpler version of our baseline reference model, including a single plant firewall securing traffic between levels 3, 2, and 1. As previously mentioned, redundant devices and HA links are intentionally removed to reduce the complexity of diagrams and facilitate the communication of key concepts. It should be noted that the NGFWs have robust HA capabilities to support the high availability requirements of OT. More details can be found in the "High Availability" section of the PAN-OS Administrator's Guide.[4] The network implements VLANs for two manufacturing cells, an operator VLAN and CCTV VLAN. The NGFW is at the center, providing Layer 7 visibility and segmentation as well as threat services for the ICS environment. The key configuration aspects of the system follow.

Layer 2 for North-South Traffic

The NGFW Layer 2 physical interfaces aggregate 802.1q trunk ports from the distribution and access switches. Within each trunk port are the pre-existing logical VLANs determined by the asset owner per segmentation frameworks such as ISA 62443 and possibly using risk-based approaches such as hazards and operability studies (HAZOP). The NGFW isolates each of the VLANs within the trunk port as logical Layer 2 subinterfaces. By isolating the VLANs as subinterfaces, it is now possible to treat each VLAN as a separate security zone and gain Layer 7 visibility as well as apply Layer 7 controls and threat services. With this approach, security teams can

[2] PAN-OS Administrator's Guide – Subscription Services, https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/subscriptions/all-subscriptions.html.

[3] PAN-OS Administrator's Guide – Interface Configuration, https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/networking/configure-interfaces.html.

[4] PAN-OS Administrator's Guide – High Availability, https://docs.paloaltonetworks.com/pan-os/8-1/pan-os-admin/high-availability.html.

leverage existing VLAN segmentations and avoid having to reconfigure the IP addresses for the existing devices.

Tap Mode for East-West Traffic

There are many instances in OT where automation and security teams deem it best not to pass industrial automation and control traffic through a firewall. This could be due to pre-existing internal architectural standards, for example, or requirements by automation vendors. Whatever the reason, organizations still value the ability to identify OT devices and quarantine them if they exhibit unusual or malicious behavior. The tap mode interface configuration can be utilized here to provide the required east-west visibility. As mentioned previously the NGFW can support multiple interface configuration types simultaneously. Some ports can be configured as Layer 2 for north-south traffic and others as tap mode to monitor the SPAN ports on the access switches for east-west.

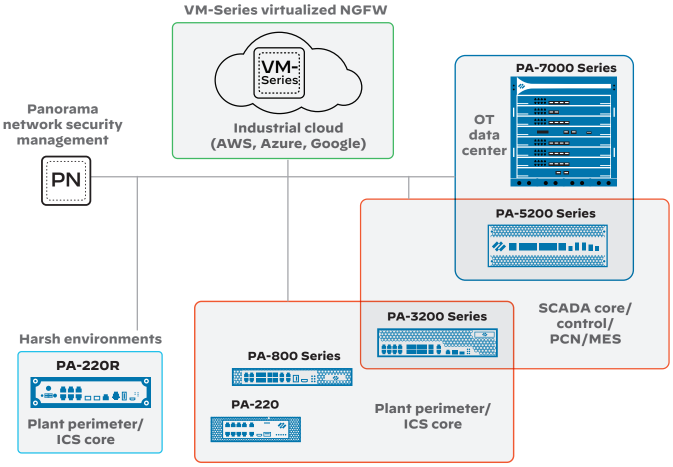
VM-Series virtualized NGFW
VM-Series
Industrial cloud (AWS, Azure, Google)
Panorama network security management
PN
OT data center
PA-7000 Series
PA-5200 Series
SCADA core/ control/ PCN/MES
PA-3200 Series
Harsh environments
PA-220R
Plant perimeter/ ICS core
PA-800 Series
PA-220
Plant perimeter/ ICS core

Figure 4: Mapping of NGFW portfolio to OT network security applications

Figure 4 shows the different physical and virtualized NGFWs offered. In addition to physical appliances for environmentally controlled environments, a ruggedized PA-220R and VM-Series virtual firewalls for public cloud usage are also offered. All NGFWs implement the SP3 architecture and support the aforementioned security services, ensuring consistency in approach across your enterprise. Furthermore, all appliances can be managed centrally using Panorama™, our central management platform. Security teams can pick from any of these different NGFW models to ensure an optimized deployment.

With the basic segmentation and monitoring schemes described previously, we can now discuss the next step, which is the data aggregation.

Cortex Architecture and Cortex Data Lake

High-Level Overview of the Cortex Platform

Aggregation of the security telemetry from the NGFWs is a critical first step before analyzing the data. Cortex™ Data Lake is the component of the architecture that performs this function. It is part of the Cortex platform—the industry's only open and integrated AI-based continuous security platform. Figure 5 shows a high-level diagram of the Cortex platform architecture.

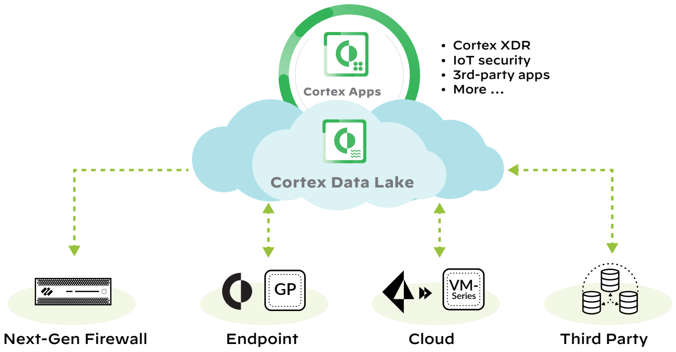

Figure 5: Mapping of NGFW portfolio to OT network security applications

In the Cortex architecture, Palo Alto Networks products send rich network, endpoint, and cloud data to Cortex Data Lake. The physical NGFW appliances and virtual NGFWs (VM-Series) are respectively responsible for sending network and cloud security telemetry to the data lake. Endpoint security telemetry is provided to it by Cortex XDR™ endpoint agent, which is a lightweight agent that you deploy on your hosts. The agent includes endpoint protection, making it easy for

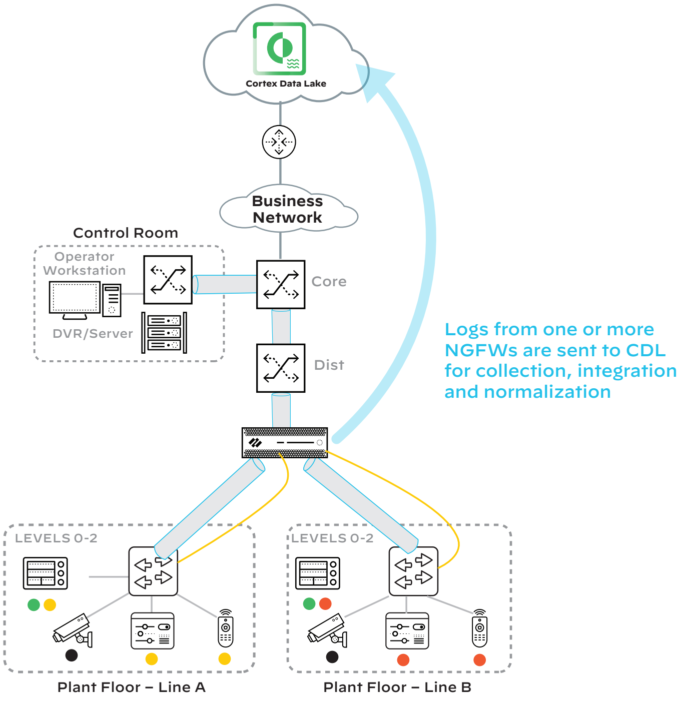

Figure 6: Collection of security telemetry from plant firewall to Cortex Data Lake

5. Cortex Microsite at PaloAltoNetworks.com, https://www.paloaltonetworks.com/cortex.

you to block malware, exploits, and fileless attacks while also collecting all the data you need for detection and response.

Various Cortex applications, such as Cortex XDR and the IoT Security service from Palo Alto Networks as well as third-party applications, are available to make use of the rich, correlated security information in Cortex Data Lake. Cortex applications can apply machine learning to automatically detect threats and communicate back to the enforcement points, NGFWs, and endpoints to adjust policy and block threats. Learn more about the Cortex platform at the microsite on the Palo Alto Networks website.[5]

Cortex Data Lake for OT

In our reference architecture, Cortex Data Lake is used to collect the telemetry data from the plant firewalls. This is depicted in figure 6.

As a cloud-based data lake, it is very easy to aggregate telemetry across multiple geographically distributed plant firewalls and scale the storage requirements on a pay-as-you-grow basis.

IoT Security Service

With the security telemetry from the OT firewalls aggregated in a central data lake, we can now discuss the functionality provided by the IoT Security service.

Service Overview

The IoT Security service is a Cortex application that applies machine learning technology on the rich NGFW security telemetry in the data lake, specifically the extended application logs (EAL), to identify and secure devices. NGFWs simply need to be enabled with an IoT Security service subscription to access the automated device security functions that follow.

OT and IoT Device Visibility

Unlike other ICS network security monitoring solutions, which are dependent on pre-existing signatures to identify devices, the IoT Security service uses machine learning to discover devices. In other words, there is no need to wait for a vendor-provided signature database before a device can be identified. Furthermore, the approach does not probe actively, thereby eliminating the risk of such methods accidentally causing OT/IoT devices to crash. As a cloud-based service, the IoT Security service is also able to utilize the device-level intelligence across the Palo Alto Networks customer base to make the device personalities even more precise.

Device Risk Assessment

The service utilizes several Palo Alto Networks and third-party intelligence sources to assess the risk associated with devices. This helps organizations prioritize the patching or isolation of high-risk assets. Customers can use the Threat Prevention service in conjunction with the IoT Security service to provide virtual patching of these identified devices by stopping their associated exploitation vectors.

Native Policy Enforcement

While other ICS network security monitoring solutions only detect devices and are dependent on other security devices for enforcement, Palo Alto Networks provides both detection and policy enforcement natively. This eliminates the need for additional sensors and the associated operational costs. Here the IoT Security service applies machine learning to analyze Cortex Data Lake data and make recommendations for policy, which the users can automatically or manually implement. Using Device-ID on the NGFW, users can create intuitive device-level policies that are based on device objects configured using six device attributes: category, type, vendor, model, operating system (OS) family, and OS version. This approach is more effective than creating device policies based only on IP address, and networking constructs such as security zones, as the Device-ID policies are applicable to devices irrespective of IP address or location in the network.

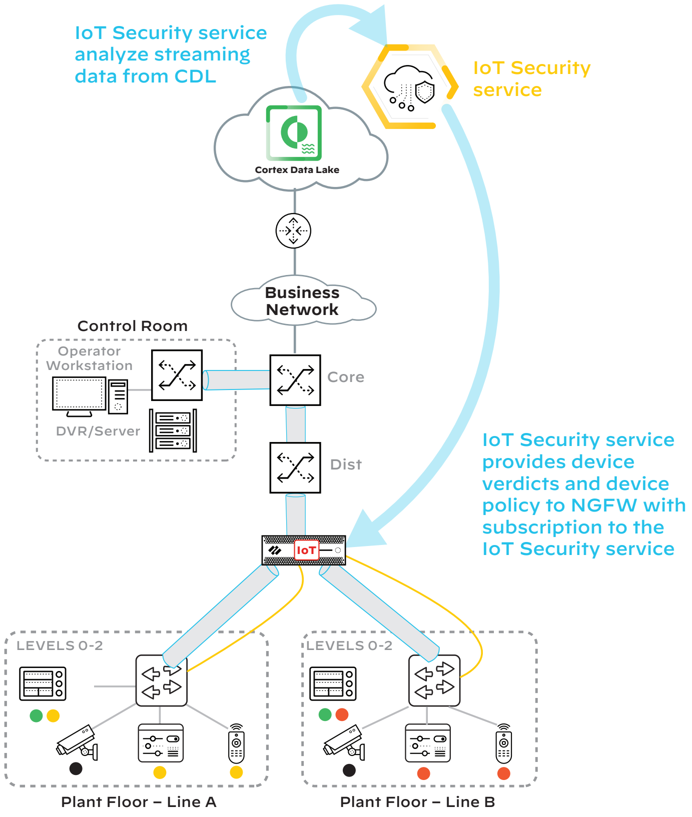
IoT Security service analyze streaming data from CDL
IoT Security service
Cortex Data Lake
Business Network
Control Room
Operator Workstation
DVR/Server
Core
Dist
IoT Security service provides device verdicts and device policy to NGFW with subscription to the IoT Security service
IoT
LEVELS 0-2
LEVELS 0-2
Plant Floor – Line A
Plant Floor – Line B

Figure 7: IoT Security service providing ML/AI-based security services

IoT Security Service Deployment in OT

Figure 7 shows the addition of the IoT Security service to our reference architecture. Here the service analyzes the streaming data from the CDL using machine learning technology. Device verdicts and device policy are sent back to NGFWs that subscribe to the IoT Security service.

OT Use Cases for the IoT Security Service

Several OT security use cases can be discussed with all of our key architecture elements covered.

Asset Inventorying, CMDB Integration, Supply Chain Management

Visibility over the OT and IoT assets in your ICS environment is an important first step. The IoT Security service, used in conjunction with existing or newly deployed NGFWs acting as sensors, is able to identify the IP-connected OT devices, such as HMIs, PLCs, RTUs, and IIoT gateways, along with the IT and IoT devices deployed in OT. This device information can be shared with central configuration management databases (CMDB) via the RESTful API on the NGFW.

Having a thorough understanding of your OT and IoT assets is also critical from a supply chain management standpoint. This helps you assess to what degree your environment is using products from approved vendors and manage risk better for when there may be violations or exceptions. Asset inventorying and supply chain management challenges could get exacerbated as the number of OT and IoT devices grows. The IoT Security service uses ML/AI to automate this process and is available 24/7. Furthermore, the same approach allows it to identify cyber-physical assets and general IoT/OT devices. Users benefit from the cloud-based approach as it is easy to gain visibility of all devices across an often highly geographically distributed OT infrastructure.

Zero Trust Policy for Inter-Zone Traffic

After establishing a baseline for OT traffic, the machine learning technology of the service can make recommendations for device-level policies consistent with Zero Trust architecture. The user can decide to implement these recommendations as is, make adjustments to better suit their environment and preferences, or just use them for informational purposes. Device-level policy using Device-ID can be applied in conjunction with App-ID, User-ID, and Content-ID to achieve Zero Trust architectures. Threat Prevention services can be applied natively in the same device, the NGFW, to make sure unpatched and unpatchable devices and systems are protected from exploits and malware, while preventing outbound command-and-control traffic.

Quarantining Devices Exhibiting Anomalous Behavior

The IoT Security service not only identifies the OT and IoT devices but also applies machine learning on the ICS protocol and application traffic to establish baselines and detect anomalies. The NGFWs support over 500 App-IDs for ICS, including some for functional-level App-IDs (e.g., Read, Write, Restart). For example, the IoT service can detect when a human machine interface, which is rarely used for more than reading PLCs, is suddenly used to program coils and registers. This may not be indicative of malicious use but is certainly an anomaly that would be worth getting informed about for possible incident response.

The IoT Security service reduces the burden on SOC teams by automatically detecting such anomalies and alerting the security analyst. Devices exhibiting anomalous/malicious traffic that goes through the firewall can be quarantined using firewall policy. Anomalous/malicious devices not going through the firewall can be placed on separate VLANs that do traverse the firewall for monitoring and possible access limitation or outright blacklisting.

Other Considerations

Alternate Architectures

While the hybrid mode comprising inline Layer 2 interfaces and tap mode interface configurations is our baseline, other configurations are certainly possible.

VWIRE

Users who want to deploy the NGFW inline with the least amount of changes to the network can use this bump-in-the-wire mode. It is completely transparent and requires no change to the network. Similar to the Layer 2 interface option, VWIRE subinterfaces can be mapped to VLANs. There is a tradeoff in the number of physical interfaces needed (more are required for VWIRE vs. Layer 2), and the policy implementation is more involved. However, if inline controls are required with minimal disruption to the network, VWIRE may be optimal.

Tap Mode

Users who are not ready to deploy inline firewalls between the distribution switches and access switches can deploy the NGFWs purely in tap mode to monitor SPAN ports. The IoT Security service still works in this completely passive model, providing visibility to devices, alerting, and policy recommendations.

Using the IoT Security Service with Older PAN-OS Versions

PAN-OS 10.0 is the optimal OS version for the IoT Security service with Device-ID functionality for policy enforcement. However, users of older PAN-OS versions, namely PAN-OS 8.1 through 9.1, can still use the IoT Security service where policy enforcement is achieved through the use of Dynamic Access Groups (DAG). The IoT Security service sends device policy recommendations to the Panorama management appliance, and from there, users can push policies to the individual firewalls. Policies are associated with individual devices by assigning their IP into the associated DAG. For more details on using the IoT Security service with older versions of PAN-OS, please see the PAN-OS Administrator's Guide.

Integrated IT and OT Architectures

A major benefit of the machine learning-based approach of the IoT Security service is that it doesn't really care what kind of IP-connected device it is in order to apply its detection capabilities. This makes it ideal for accommodating both IT and OT environments. Other point solutions using signature-based approaches are either only addressing one side of the enterprise or, if they try to address both, are constantly trying to keep up in terms of creating all the OT and IoT device signatures. This approach is simply not scalable, whereas the machine learning approach used by the IoT Security service is scalable and ideal for a converged IT-OT security architecture as shown in figure 8.

Here we see the plant firewalls, along with the IT-OT perimeter firewall, internet gateway firewalls protecting Level 4, and even VM-Series firewalls protecting the public cloud. All of these firewalls, once enabled with the IoT Security service, can send data to Cortex Data Lake and get information about the devices, risks, and recommended policies to protect the devices in their respective domains.

Summary

In summary, both OT and IoT device security are critical for safe enablement of Industry 4.0. The challenge will only increase as the number of IP-connected devices in OT grows and OT becomes more converged with business networks and third-party infrastructure, such as vendor networks and public cloud infrastructure.

Palo Alto Networks offers a unique IoT security solution that provides the needed device security capabilities via our Next-Generation Firewall, coupled with Cortex Data Lake and the IoT Security service. The IoT Security service differentiates itself from other IoT security products in that it uses a cloud-based machine learning architecture to provide fast and accurate detection of devices, flexibility to accommodate IT and OT devices, and scalability to support large and small OT infrastructure. Moreover, because the sensor is simultaneously the enforcement point, our approach provides native protection of devices. This is in contrast to other OT sensors, which provide only visibility and are dependent on other products for protection.

Request more information or a demonstration on how you can better secure your OT and IoT devices in your ICS by contacting your local sales representative or sending us email at **secure_ics@paloaltonetworks.com**.

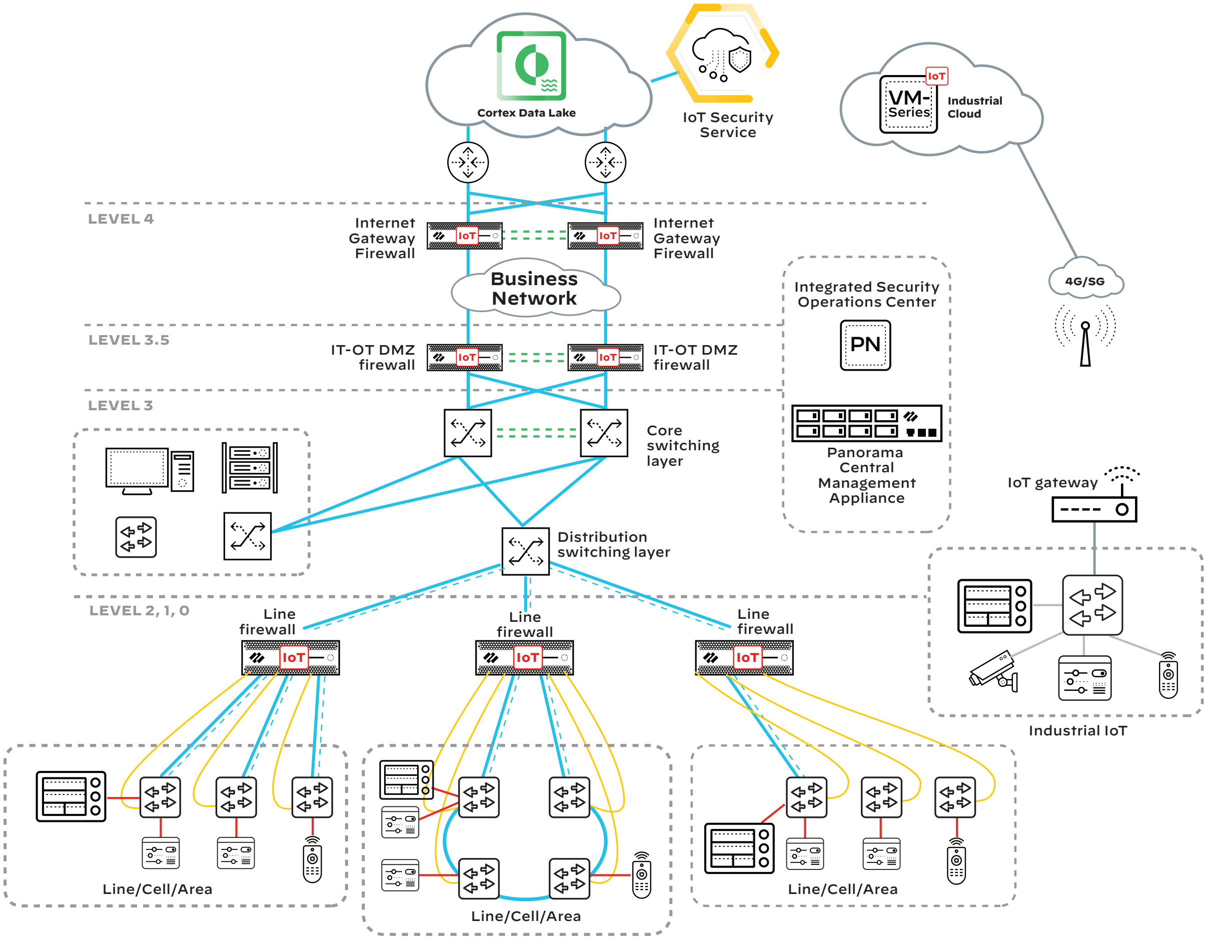

Figure 8: Converged IT-OT security architecture

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com