CHAPTER 2

CPwE Cloud Connectivity Design Considerations

This chapter provides an overview of design considerations—using diverse technologies for threat detection and prevention—to integrate an end-to-end IACS solution from the skid/equipment/machine to the enterprise to the cloud within a CPwE architecture. Although this design guide focuses on end-to-end FactoryTalk solution uses cases, the holistic and diverse industrial security technology best practices apply generically to IACS device to cloud use cases. This chapter includes the following major topics:

It is important to mention that this design guide does not discuss options for the method used by the various cloud connected solutions to connect to the cloud application; these design decisions have already been made by the developers. In this design guide, the various cloud connected solutions use the TLS protocol to establish secure communication with the cloud applications. This design guide does not suggest options for gateway to cloud communication. This design guide does provide infrastructure recommendations to monitor the TLS traffic between the Industrial Zone and the cloud application for anomalous traffic, as well as recommendations for firewall configurations to limit the traffic from the Industrial Zone to the cloud.

Security Policy Considerations

Most companies understand the need for security and, therefore, have policies governing enterprise systems such as the internet or email use. It is, however, common for security policies to often be insufficient or, in some cases, nonexistent in Industrial Zone systems.

With the recent popularity of providing computing resources, analytics, application software, and data storage in the cloud, there is a natural desire to extend this capability to industrial operations. Vendors are providing industrial operators with cloud-based Software as a Service (SaaS) not only to leverage their expertise in areas such as analytics, but to deliver on the promise of cloud-based technologies to also allow their customers to view data anywhere and from any device.

Enterprise IT has adopted the use of cloud-based resources, which has positioned that department to be aware of and responsible for:

- Cloud security policy considerations and ultimately the generation of cloud security policies
- Cloud infrastructure security best practices
- Cloud application security best practices

The opportunity for OT and IACS assets to send data securely to a trusted cloud to provide information on product quality, skid/equipment/machine operations, and performance has been sufficiently advantageous to realizing the development of cloud connected devices. Cloud connected devices are typically dual role where they support communication to IACS devices and to a trusted cloud-hosted application.

Risk Assessment and Risk Management

Until recently most industrial operations have been segmented with no connectivity to the internet unless there was a specific business need for highly restricted remote support, site-to-site communications through a VPN tunnel, or other less common purposes. With today's desire to send information to cloud-based services, one must consider the risk versus reward proposition of sending industrial operational data to trusted cloud applications. According to organizations such as Rockwell Automation, Cisco Systems, and Open Web Application Security Project (OWASP), the following risks should be considered:

- Data Ownership and Protection

  The traditional approach to business software applications is to run them in-house on an infrastructure built and maintained by the organization using the applications. Therefore, all data resides within the organization and it has complete control over the data and how it is protected. An organization using cloud services must understand to what extent the cloud provider's personnel have access to their data. When an organization uses these cloud-enabled services, they are outsourcing business processes to the cloud provider which requires access to the organization's data. The industrial operator should confirm that both the cloud provider and the Internet Service Provider (ISP) are trusted and provide the necessary network and security services to help protect connectivity and data as required by their business and security policies.

- User Identity Management and Federation

  Organizations must understand how cloud providers identify users and manage their accounts for accessing data in the cloud. In addition, they must understand the risks associated with logon accounts and how the cloud provider mitigates these risks. These risks include password guessing, password theft, password reset, hijacking of user login sessions, and revocation of access. As an alternative to creating a separate island of user names and passwords, some cloud providers may offer integration with an organization's in-house authentication systems. Through integration, existing in-house logon accounts managed by the organization can be used to access data in the cloud.

- Regulatory Compliance

  Organizations using cloud providers face different challenges regarding regulatory compliance for data stored in the cloud. They must consider whether data entrusted to a cloud provider carries legal/regulatory protection and breach notification requirements, such as protected health information (PHI) governed by HIPAA and HITECH, personally-identifiable information (PII) governed by state privacy laws, and payment card information regulated by the Payment Card Industry's (PCI's) Data Security Standard (DSS).

- Business Continuity and Resiliency

  Business continuity and resiliency refer to the ability of an organization to conduct business operations in adverse situations. Adverse situations include disruptions not only to the information technology infrastructure, but also any disruptions affecting the ability of the cloud service provider to deliver its services at defined service levels, including, for example, the loss of key personnel or the loss of access to business offices.

  When an organization uses a cloud provider, the organization cedes control of business continuity planning for the data and services entrusted to the provider. As a result, the organization must consider carefully how the cloud provider would maintain continuity of services if affected by adverse situations.

- User Privacy and Secondary Uses of Data

  Organizations must understand how a cloud provider helps protect and use information about different types of users. They should consider to what extent a cloud provider can disclose information about its employees, its customers, or its business. This information includes specific information or aggregate statistics. It includes information collected from an individual's use of the cloud provider's information systems, such as characteristics of user behavior (for example, links clicked, options selected, etc.) and productivity measurements.

- Service and Data Integration

  Organizations must understand how their users will access the data and services of a cloud provider. Typically, this access will be over the internet or a virtual private network (VPN) using a web browser or a software application downloaded from the cloud provider. If the organization will be interfacing any of its systems with the systems of the cloud provider, they must understand the technical aspects of how the interface will work. An example of this may be if the organization wants to implement "back-end" or batch processing of Health Level 7 (HL7) or Electronic Data Interchange (EDI) transactions. In both cases (user access and system interfaces), organizations must understand the risks associated with electronic communication across the internet or wide area networks (WANs). This includes interception of data in transit, falsification or corruption of data, and verification of client and server endpoints.

- Multi-tenancy

  In a cloud computing environment, multi-tenancy refers to the sharing of information technology infrastructure among multiple clients (different customers of a single cloud service provider). This infrastructure includes telecommunications circuits, network equipment, servers, storage, and application software. Multi-tenancy allows cloud providers to achieve economies of scale, which would be impossible for an individual organization to attain, enabling organizations to obtain higher levels of service at lower costs.

  Risks with multi-tenancy include one client accessing the data of another client, unintentional mixing of one client's data with another client's data, one client affecting the quality of service provided to another client, and cloud provider application software upgrades affecting client business operations. While cloud providers can be expected to have adequately mitigated these risks given that multi-tenancy is core to the cloud business model, an organization should understand how the cloud provider achieves isolation between clients. Isolation approaches include use of virtualization technologies such as virtual machines, application-level isolation through processes, threads, or application-managed contexts, and database-level isolation using separate database instances, tablespaces, or record identifiers

- Incident Response and Forensic Analysis

  Incident response and forensic analysis refer to activities conducted by an organization when there is a security incident requiring immediate response and subsequent investigation. These incidents include malicious acts or mistakes by the employees or former employees of the organization, resulting in data breaches. When an organization uses a cloud provider, it does not have access to the underlying log files and other low-level system-level information typically used for forensic examination.

- Infrastructure and Application Security

When an organization uses a cloud service provider, it trusts the provider to properly secure its applications and infrastructure. This is a highly complex activity requiring an extensive array of personnel with advanced technical skill sets and threat knowledge.

- Non-production Environment Exposure

  A cloud provider typically operates multiple environments where cloud data and services exist. These environments include what is normally referred to as a production environment, which is where cloud subscribers have the primary copy of their data and where they conduct their business operations.

  Cloud providers also typically operate other environments for purposes such as software development, testing, training, and demonstrations to potential customers. These other environments may be populated with copies of data from the production environment. In other words, an organization's data may be copied into several places to support the necessary business operations of the cloud provider. The data contained in these copies may or may not be de-identified, a process whereby individual customer information is rendered untraceable to a specific customer and individual business information is made untraceable to an organization.

Architectural Considerations

Network and security technology architecture decisions are made with several key factors in mind to help determine how IACS is implemented. Some of the key architectural decision factors are:

- Align the architecture to support security policies.
- Consider best practices for safety and security provided by IACS vendors.
- Technical security controls such as zoning, firewalls (protection and detection), end host protection, application security, and proxy services to support an organization's security policy.
- Technical control deployments will align with an organization's risk profile while balancing business aspects and budget constraints.

In the following sections four different security architectures are reviewed, which have been classified as Platinum through Bronze to denote, respectively, higher-layered through lower-layered security architectures. Cisco, Rockwell Automation, and Panduit realize there are many different possible combinations that could be covered, so four security architectures were selected to help provide a scalable solution to cover the typical scenarios found within plant environments.

Regardless of which multi-layered security architecture is deployed, the manufacturer must confirm that the cloud provider and ISP are trusted and provide the necessary network and security services to help protect connectivity and data as required by the manufacturer's business and security policies.

Cisco, Rockwell Automation, and Panduit recommend that a risk assessment be conducted to survey IACS assets and identify any potential threat vectors before the selection and deployment of any security architecture.

Architectures at a Glance

Four architectures are discussed in this design guide, each with varying key security technologies to help monitor and limit the traffic between the cloud connected devices and the cloud application. Table 2-1 summarizes the architectures.

Table 2-1 Architectures at a Glance

Architecture	TLS Encrypted Communications between the Industrial Zone and Cloud Applications	IDMZ	Industrial Firewall at the Cell/Area Zone	TLS Monitoring	DNS Protection
Platinum	X	X	X	X	
Platinum Architecture with Cisco Umbrella	X	X	X	X	X
Gold	X	X	X		
Silver	X	X			
Bronze	X				

The IDMZ offers a buffer between the Enterprise and the Industrial Zone networks and is a widely accepted method for providing a layered security model. In this design guide, one option is to place the TLS proxy in the IDMZ.

The Industrial Firewall placed within the Cell/Area Zone can be used to limit the traffic of the cloud connected devices to the DNS server and the cloud. The industrial firewall device is used to define the boundary of the security zone and is used if the cloud connected device is compromised as a means to limit the traffic reaching other IACS applications within the Industrial Zone.

TLS URL filtering and monitoring can be accomplished by the Cisco WSA or through the Cisco Umbrella service. TLS monitoring is used to spot unusual traffic patterns or payload sizes in an attempt to monitor and alert the end user to possible malicious traffic. In addition, decryption policies within the Cisco WSA can allow for advanced malware protection by examining the payloads of traffic.

Cisco Umbrella is a security solution that helps protect against threats by helping ensure that any IACS traffic to the internet or cloud is going to safe websites or domains. DNS requests are sent to the Umbrella service, which is constantly updated with information received from Cisco's Talos security intelligence team. This helps ensure that web traffic from end users or the cloud connected device is legitimate, but it can also block malware, phishing, and command and control callbacks over any port or protocol. If the cloud connected device was in any way compromised, Cisco Umbrella would keep the malware from reaching websites that could jeopardize the network.

See the Cisco Umbrella page at: https://umbrella.cisco.com.

Platinum Security Architecture

This security architecture provides multiple layers of defense-in-depth using diverse technologies for threat detection and prevention. This helps to prevent a single vulnerability from taking down IACS operations within the Industrial Zone. The cloud connected device typically is configured to communicate with the IACS devices using the CIP protocol and send and retrieve data to and from the cloud based on IACS device, such as lifecycle status information. The Platinum architecture contains an IDMZ for brokered services such as a TLS proxy and is the Cisco, Rockwell Automation, and Panduit security architecture recommended for manufacturers that require cloud-based connectivity yet have a lower tolerance to risk.

There are several key security technologies for threat detection and prevention that help to create this multi-layered security architecture, including:

- A TLS proxy that can inspect the outbound and inbound encrypted TLS traffic between the cloud connected device and the cloud-based software.
- Zoning—An IDMZ with Firewall(s) and replicated services to buffer the Industrial Zone from the Enterprise Zone.
- Zoning—An Industrial Firewall which creates a security zone, to help enforce security policies within a Cell/Area Zone, that helps protect the Industrial Zone from the ingress/egress point of the TLS traffic.

A TLS proxy, such as the Cisco WSA could be in the IDMZ to monitor and inspect the traffic between the cloud connected device and the cloud-based application. See Figure 2-1.

Figure 2-1 Platinum Security CPwE Cloud Connectivity Use Case with TLS Proxy in the IDMZ

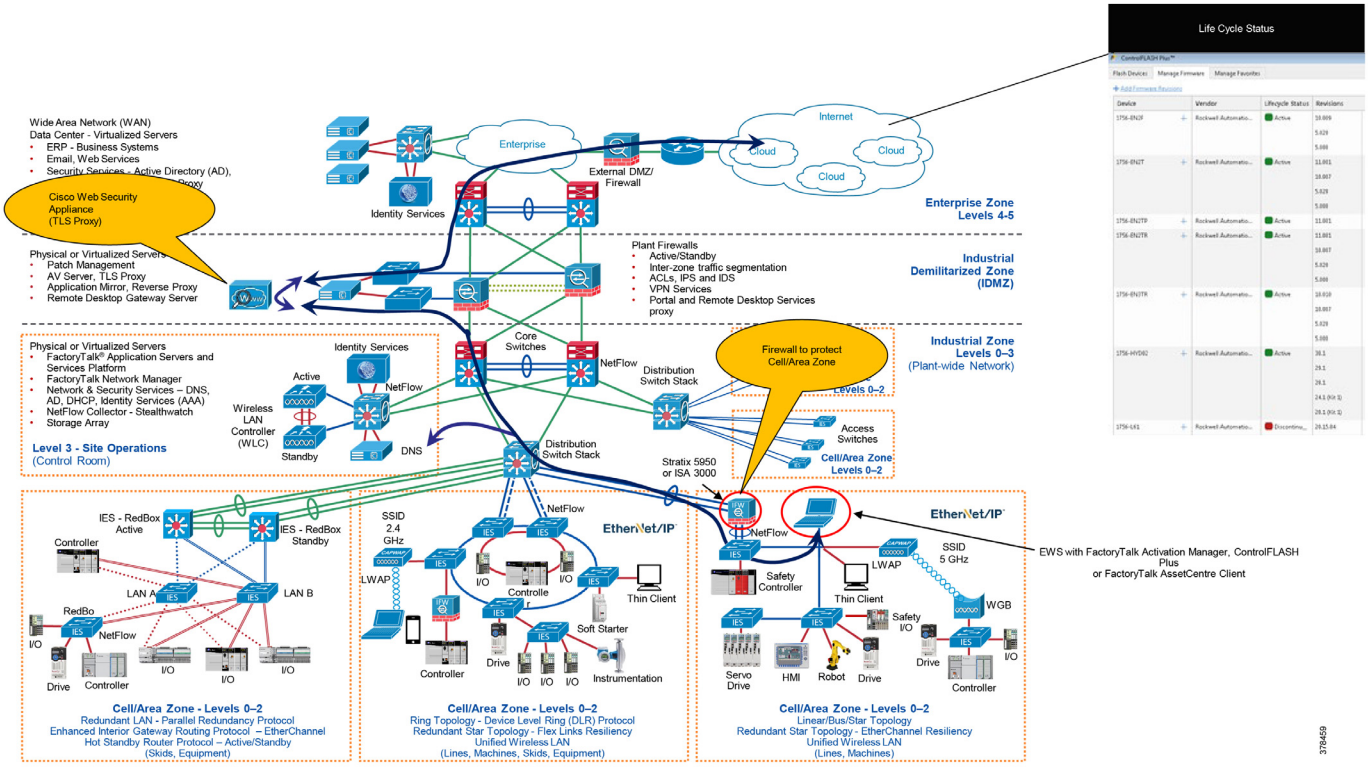

Some organizations may have already implemented a TLS proxy in the Enterprise Zone and want to use these same appliances for Industrial Zone cloud connectivity as shown in Figure 2-2. This does not follow one of the key tenets of the IDMZ, which is brokered services but still provides TLS proxy services to the IACS. In addition, an IDMZ TLS proxy can be used to broker inspected traffic to an upstream proxy device that may exist in the Enterprise Zone.

Figure 2-2 Platinum Security CPwE Cloud Connectivity Use Case with TLS Proxy in the Enterprise Zone

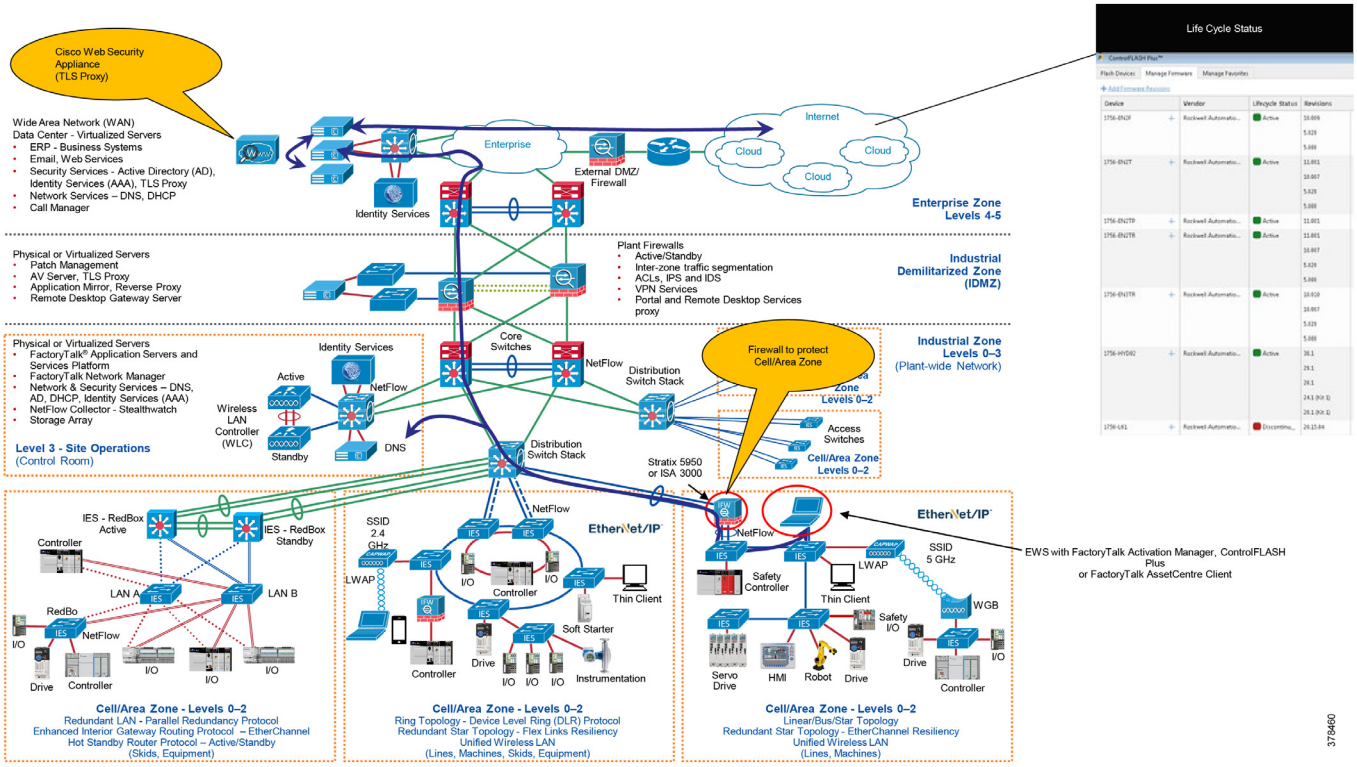

Another architecture option for providing TLS proxy services is to use a cloud-based solution such as the Cisco Umbrella TLS proxy and DNS services. See Figure 2-3.

Figure 2-3 Platinum Security CPwE Cloud Connectivity Use Case with TLS Proxy Provided by Cisco Umbrella Cloud Solution

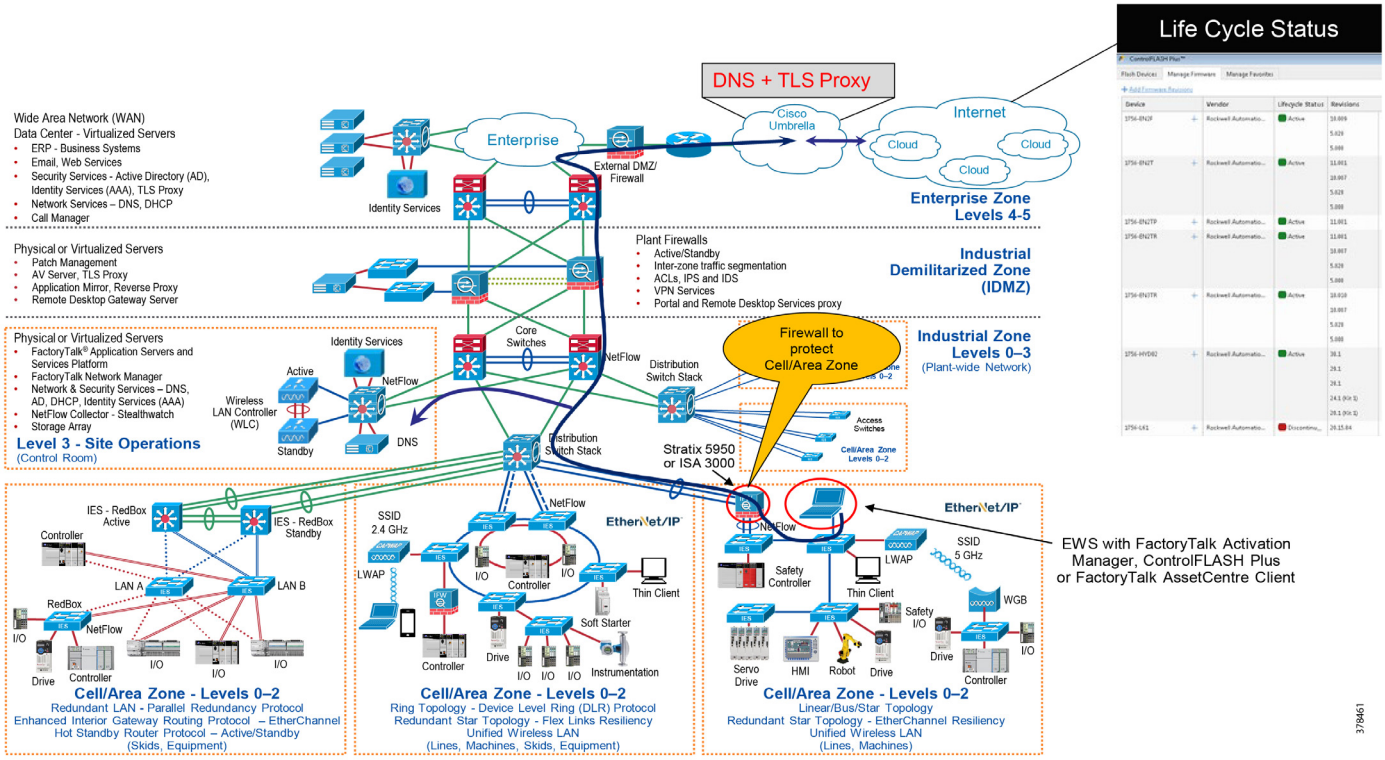

Gold Security Architecture

The Gold security architecture is depicted without a TLS proxy but still incorporates zoning with an IDMZ and Industrial Firewall(s). See Figure 2-4.

In this architecture, the TLS traffic between the cloud connected device and the cloud-based application is no longer inspected through the TLS proxy. However, firewalls can provide URL filtering and whitelisting/blacklisting to lower overall risk of the architecture. This presents the possibility of undetected malware traveling between the cloud-based application and the cloud connected device.

Figure 2-4 Gold Security CPwE Cloud Connectivity Use Case

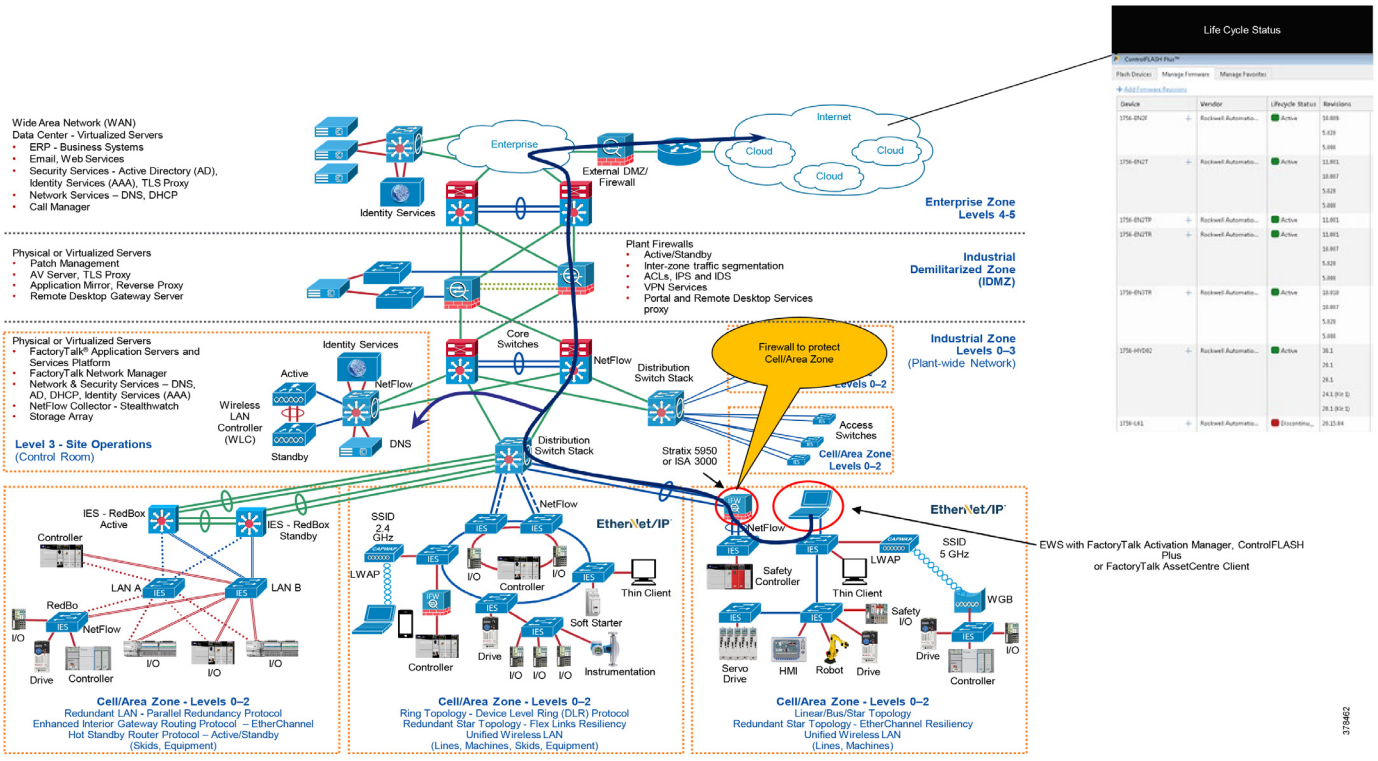

Silver Security Architecture

This architecture keeps the IDMZ concept with firewall(s) and replicated services to buffer (zoning) the Industrial Zone from the Enterprise Zone. The biggest difference between this architecture and the Gold architecture is there are no Industrial Firewall(s) acting as a security boundary (zoning) for the Cell/Area Zones. Industrial firewalls define smaller, more granular security zones that help enforce security policies to help protect the Industrial Zone from the ingress/egress point of the TLS traffic. See Figure 2-5.

Figure 2-5 Silver Security CPwE Cloud Connectivity Use Case

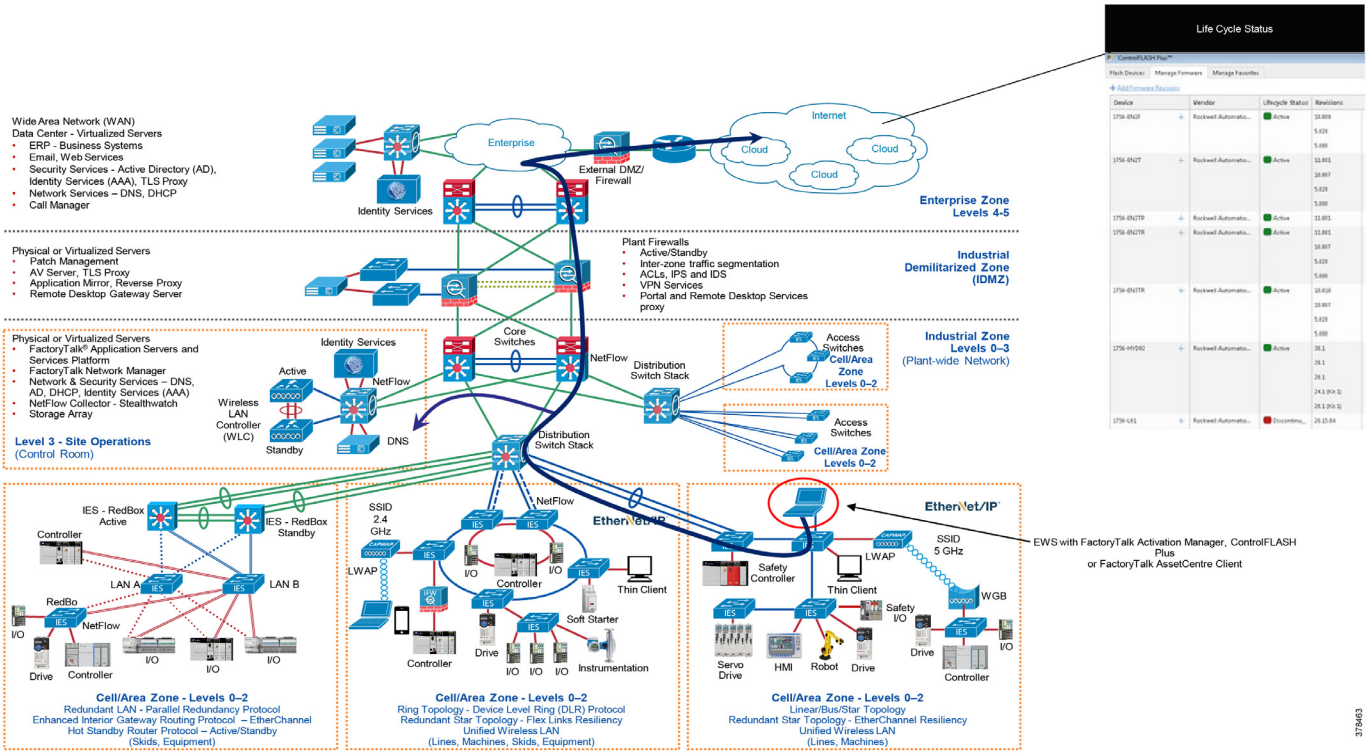

Bronze Security Architecture

This security architecture has the fewest defense layers of diverse industrial security best practices and is only recommended for manufacturers with a higher tolerance to risk. See Figure 2-6.

Figure 2-6 Bronze Security CPwE Cloud Connectivity Use Case

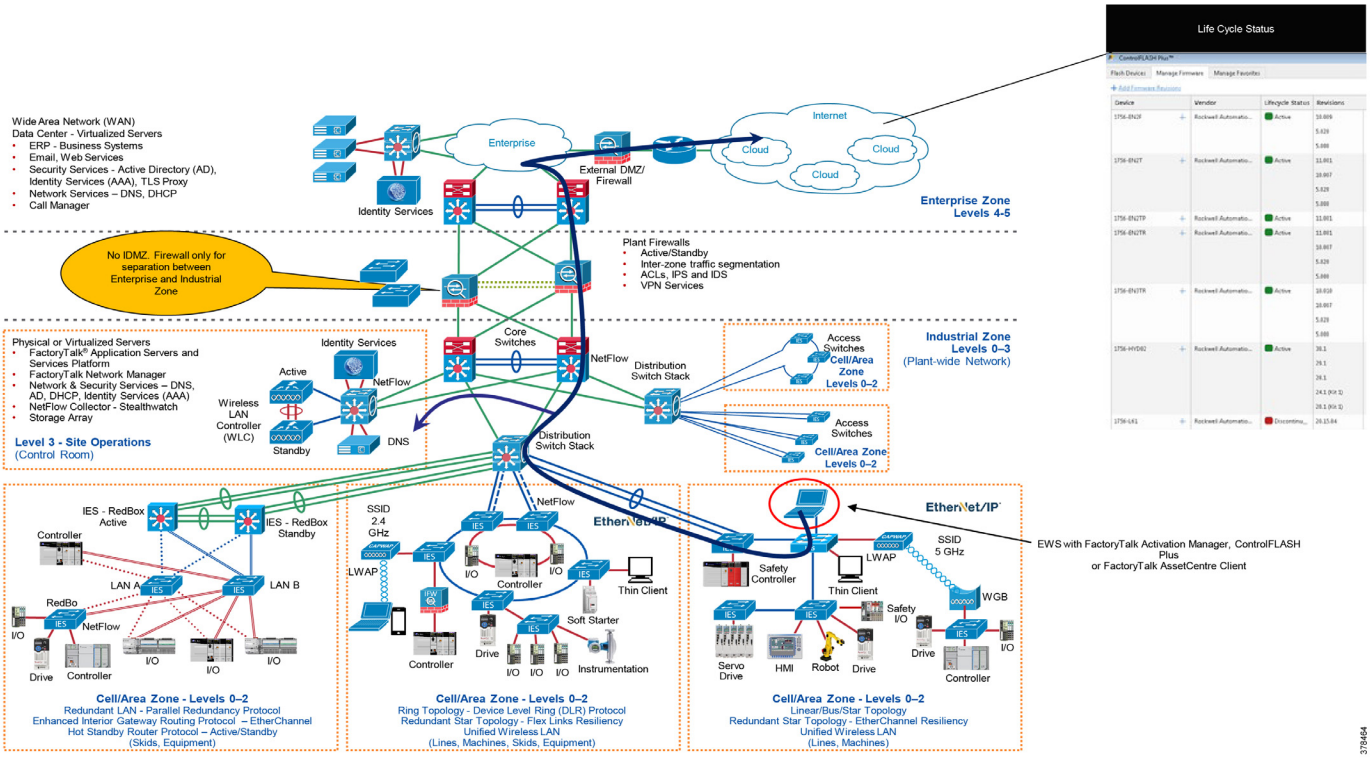

While a firewall is used to define security boundaries, many security standards organizations, as well as Rockwell Automation and Cisco Systems, maintain that a firewall alone is not an acceptable layered security architecture for most risk averse manufacturers. When using a firewall, simple “permit” and “deny” rules define the traffic that is allowed through the firewall device. If the firewall permits traffic from a host in the Enterprise Zone to the Industrial Zone and either host is compromised, it is unlikely the firewall will deny or detect malicious traffic.

For further information about deploying firewalls, see:

- NIST 800-41 Guidelines on Firewalls and Firewall Policy: http://csrc.nist.gov/publications/nistpubs/800-41-Rev1/sp800-41-rev1.pdf
- Deploying Industrial Firewalls within a Converged Plantwide Ethernet Architecture Design and Implementation Guide (document number ENET-TD002A-EN-P):
  - Cisco site: https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/5-0/Firewalls/DIG/CPwE-5-IFS-DIG.html
  - Rockwell Automation site: http://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td002_-en-p.pdf

Trusted versus Untrusted Security Zones

Security zones are established by grouping assets with common and related security requirements. These zones allow for technical and non-technical controls to be implemented to mitigate risk. Assets within the security zone are considered “trusted” while any assets outside the security zone are considered “untrusted”.

This concept aligns with the IEC 62443-3-2 standard that discusses the zones and conduits model where the security zones represent assets with common security requirements and the conduits represent the communication channels that exist between the security zones.

When implementing a cloud connected device within the Industrial Zone, one must consider if the cloud connected device should be placed within its own security zone or within an existing security zone. For instance, if the cloud connected device is placed within a new security zone, communication to each IACS asset in another security zone must be considered “untrusted” and therefore a technical security control such as an industrial firewall should be implemented to limit and inspect the traffic to other security zones. If the cloud connected device is placed within an existing security zone, then the communications to IACS devices within the same security zone is considered “trusted”.

Most open industrial protocols do not support device authentication and authorization, so the best method of risk mitigation is to create security zones and secure the communications between the security zones. With the inclusion of CIP Security to Rockwell Automations portfolio, the cloud connected device could be limited to the number of devices to which it has access. This can be used to enforce conduits within a security zone with devices that may have varying risk tolerance. Understanding device-to-device IACS communication is paramount when implementing technical controls because limiting traffic to known communication paths and protocols is the foundation for successfully implementing these types of controls.

Each designer and implementer must determine their security profile (through a risk assessment process) to determine the placement of the cloud connected device within the Industrial Zone.

Technology Considerations

There are different security technologies used within CPwE Cloud Connectivity to provide a scalable and multi-layered security architecture.

The following section briefly discusses the technologies described in the previous security architectures.

Transport Layer Security (TLS)

The first common technology used for communications between the FactoryTalk applications and the cloud applications is the IETF TLS. TLS provides encrypted communications between the two participating endpoints, which in CPwE Cloud Connectivity is the FactoryTalk application and the cloud hosted destination software. A consideration of the participating endpoints in TLS is ensuring a trust relationship between the devices accomplished using digital certificates. TLS traffic presents a challenge to monitor or inspect because the traffic is encrypted and therefore two commonly used methods are utilized to provide TLS proxy functionality:

- The first method is to provide a TLS proxy that is capable of decrypting, inspecting, and re-encrypting the traffic. This method can cause latency and the proxy could create a network bottleneck.
- An alternative method of providing the TLS monitoring is Cisco Encrypted Traffic Analytics (ETA), which monitors the encrypted traffic looking for malware without opening the packet. ETA inspects the initial data packet of the connection and monitors the sequence of packet lengths and times thereafter, which offers clues about traffic contents beyond the beginning of the encrypted flow. In general, ETA does not provide the same security features and benefits of using the Cisco Web Security Appliance.

For more information on the Cisco Web Security Appliance, Cisco Umbrella, and Cisco ETA see:

- Cisco Web Security Appliance:
https://www.cisco.com/c/en/us/products/security/web-security-appliance/index.html
- Cisco Umbrella:
https://umbrella.cisco.com
- Cisco Encrypted Traffic Analytics:
  - https://www.cisco.com/c/en/us/solutions/enterprise-networks/enterprise-network-security/eta.html
  - https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Campus/CVD-Encrypted-Traffic-Analytics-Deployment-Guide-2019JUL.pdf

Industrial Demilitarized Zone (IDMZ) Proxy Technologies

A key tenant in the Platinum, Gold, and Silver architectures is the use of the IDMZ. This security zone provides a buffer between the Enterprise and Industrial Zone. The IDMZ is an additional layer of defense-in-depth to securely share IACS data and network services among the Industrial Zone, the Enterprise Zones, and the cloud. The demilitarized zone concept is commonplace in traditional IT networks, but is still in early adoption for IACS applications.

For secure IACS data sharing, the IDMZ contains assets that act as brokers between the zones. Multiple methods to broker IACS data across the IDMZ exist:

- Use an application mirror, such as a PI-to-PI interface for FactoryTalk Historian.
- Use Microsoft® Remote Desktop Gateway (RDG) services.
- Use a Web Security Appliance with TLS proxy server capabilities.

These broker methods, which help to hide and protect the existence and characteristics of the Industrial Zone servers from clients and servers in the Enterprise Zone, are covered in *Securely Traversing IACS Data across the Industrial Demilitarized Zone Design and Implementation Guide* (see URLs below).

In Platinum architectures (Figure 2-1) the TLS proxy is placed in the IDMZ to monitor TLS traffic between the cloud connected device and the cloud.

For a complete understanding of an IDMZ, see the *Securely Traversing IACS Data across the Industrial Demilitarized Zone Design and Implementation Guide*:

- Rockwell Automation site:
http://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td009_-en-p.pdf
- Cisco site:
https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/3-5-1/IDMZ/DIG/CPwE_IDMZ_CVD.html

Proxy Servers and Capabilities

A proxy server resides between a trusted zone and untrusted zone, typically in the IDMZ. It provides additional protection by helping prevent direct communication between the clients within the trusted zone to servers within the untrusted zone. The goal of a proxy is to be the middle man (proxy) between clients and servers.

In essence, it works as an intermediate device that intercepts a request from an originator or client device then proceeds to make the connection on behalf of the client to the target or server device. Once the connection has been established by the target or server device, the proxy server continues to "proxy" the data stream until the connection is closed.

A Forward Proxy is configured to handle requests for a group of clients to an unknown, untrusted or arbitrary group of resources that are outside of their control.

A Reverse Proxy is where the proxy is intended to be on the same network as the HTTP(S) servers and its purpose is to serve up content for these HTTP(S) servers. In this scenario, the reverse proxy is helping protect the servers that are providing the content to the clients.

A TLS proxy is a forward or reverse proxy that uses a TLS connection to provide secure communications between an originator or client device and a target or server device. It can be dual homed to allow for the usage of the Web Cache Communication Protocol (WCCP) to redirect traffic to the TLS proxy from the IDMZ ASA firewalls.

Cisco Web Security Appliance (WSA)

The Cisco WSA is an all-in-one highly secure web gateway that brings strong protection, complete control, and investment value. It offers an array of competitive web security deployment options, each of which includes the market-leading global threat intelligence infrastructure from Cisco. The Cisco WSA correlates threats collected from their network to produce a behavior score, known as a web-reputation score, on which to act. It applies and enforces web-reputation scores on parent sites and subsites. The Cisco WSA defends against malware and advanced persistent threats using multiple layers of anti-malware technologies and intelligence from Cisco Talos updated every three to five minutes. Every piece of web content accessed from HTML to images to Flash files is analyzed using security and context-aware scanning engines. Cisco WSA analyzes traffic in real time, breaks it into functional elements, and pushes elements to best-designed malware engines for inspection while maintaining high processing speeds.

Cisco Advanced Malware Protection (AMP)

Cisco AMP is an additionally licensed feature available to all Cisco WSA devices. AMP is a comprehensive malware-defeating solution that provides malware detection and blocking, continuous analysis, and retrospective alerting. AMP augments the malware detection and blocking capabilities already offered in the Cisco WSA with enhanced file reputation capabilities, detailed file-behavior reporting, continuous file analysis, and retrospective verdict alerting. Cisco AMP provides the ability to sandbox PDF, Microsoft Office software, archived/compressed files, and Windows portable executable file. The Cisco AMP feature was not tested as part of this CRD.

WSA Proxy Modes—Transparent and Explicit Proxies

Since the goal of a proxy is to be the middle man between HTTP(S) clients and HTTP(S) servers the Cisco Web Security Appliance (WSA), as a web proxy, will have two sets of TCP sockets per client request:

- Client->WSA
- WSA->Origin server

The WSA can be configured for “transparent” or “forward” from its web user interface. This is slightly misleading, as this is really “transparent” or “explicit” mode, both of which are forward proxy deployments. Reverse proxy is where the proxy is intended to be on the same network as the HTTP(S) servers and its purpose is to serve up content for these HTTP(S) servers.

The main difference between transparent and forward mode on the WSA is that in transparent mode, the WSA will respond to both transparent and explicit HTTP(S) requests. Whereas in explicit, the WSA ONLY responds to explicit HTTP(S) requests. For transparent redirection, the Web Cache Communication Protocol (WCCP) is used to intercept and forward traffic.

The WSA will always send its upstream request as a transparent style request, since the WSA is acting as its own client, UNLESS the WSA is configured to specifically use an explicit upstream proxy such as a proxy in the Enterprise Zone.

The ability of the WSA HTTP proxy to obtain the request from the client can be defined as one of two ways: Transparent or Explicit.

Each of these deployments have several specific configuration requirements:

1. Explicit proxies require that each workstation and web browser have its proxy settings pointing to the proxy server. All web traffic unless otherwise specified is sent to the proxy.
    a. For HTTPS traffic, the proxy servers certificate must be installed in the workstations trust store
2. Transparent proxies require another technology to redirect traffic to the proxy such as Policy Based Routing (PBR) or Web Cache Communication Protocol (WCCP). Only traffic that passes through the redirecting device and is configured to be redirected is sent to the proxy.
    a. For HTTPS traffic, the proxy servers' certificate must be installed in the workstations trust store

There are a few differences between explicit and transparent HTTP(S) requests:

1. An explicit request has a destination IP address of the configured proxy. A transparent request has a destination IP address of the intended web server (DNS resolved by the client).
    a. Depending on which deployment is being used, this item should be considered when configuring firewalls along the communication path
2. When requests are being redirected to the WSA transparently, the WSA must pretend to be the OCS (origin content server), since the client is unaware of the existence of a proxy. On the contrary, if a request is explicitly sent to the WSA, the WSA will respond with its own IP information back to the client.
3. The URI for a transparent request does not contain the protocol with the host:
    - Transparent—GET / HTTP/1.1
    - Explicit—GET http://www.google.com/ HTTP/1.1

Both will contain an HTTP Host header that specifies the DNS host.

Web Cache Communication Protocol (WCCP)

The WCCP specifies interactions between one or more routers, layer 3 switches or firewalls and one or more web proxies. Regarding the Platinum Architecture, WCCP is enabled on the Ingress interface of the IDMZ Firewall but can also be utilized on most Layer 3 switches that provide full routing services. The purpose of the interaction is to establish and maintain the transparent redirection of selected types of traffic that flow through a group of routers. The selected traffic is redirected to a group of web proxies in order to optimize resource usage and lower response times.

The flow of work for redirection has these steps:

1. The user enters a URL into a browser.
2. The URL is forwarded to domain name system (DNS) for address resolution.
3. The URL is resolved to the IP address of the web server.
4. The client initiates a connection to the server with a SYN request.
5. On the device, the WCCP web proxy service intercepts the HTTP(S) request (TCP port 80 or 443) and redirects the request to proxies via GRE:

- If using HTTPS the proxy requires a certificate exchange between the proxy and the client to transparently proxy HTTPS. This requires that the proxy's certificate be imported into the trusted store of each web browser on each workstation. Self-signed certificates were evaluated as part of the CPwE testing using the Cisco WSA.
- If there is a cache hit, the proxy responds to the original GET with the requested content and uses the source IP address of the origin server in the response pack.
- If the requested content is not already stored on the proxy, there is a cache miss:
  - The proxy establishes a connection to the origin server, uses its own IP address as the source, and sends the HTTP GET.
  - The server responds to the proxy with content.
  - The proxy writes a copy of the cacheable content to the disk.

WCCP Service Groups

Once connectivity is established, the device and web proxies form service groups to handle the redirection of traffic whose characteristics are part of the service group definition.

A web proxy transmits a WCCP2_HERE_I_AM message to each WCCP enabled device in the group at HERE_I_AM_T (10) second intervals to join and maintain its membership in a service group. The message may be by unicast to each device or by multicast to the configured service group multicast address. Multicast functionality was not evaluated as a part of the CPwE testing.

- The Web-Cache Identity Info component in the WCCP2_HERE_I_AM message identifies the web proxy by IP address.
- The Service Info component of the WCCP2_HERE_I_AM message identifies and describes the service group in which the web proxy wishes to participate.
  - For the testing completed as a part of the CPwE a user-configurable Service Group was created.

Table 2-2 Service Groups

Service Group	Type	Description
Service 0	Web-cache	Web caching service that permits the device to redirect HTTP traffic to the proxy.
Service 53	DNS	DNS caching service that permits the device to redirect DNS client requests transparently to the proxy.
Service 60	FTP-native	Caching service that permits the device to redirect FTP native requests transparently to a single port on the proxy.
Service 70	https-cache	Caching service that permits the device to intercept port 443 TCP traffic and redirect this HTTPS traffic to the proxy.
Service 80	rtsp	Media streaming service that permits the device to redirect Real Time Streaming Protocol (RTSP) client requests to a single port on the proxy.
Service 81	mmst	Media caching service that permits the device to use TCP-based Microsoft Media Server (MMST) redirection to route Windows Media Technology (WMT) client requests to TCP port 1755 on the proxy.
Service 82	mmsu	Media caching service that permits the device to use User Datagram Protocol (UDP)-based Microsoft Media Server (MMSU) redirection to route WMT client requests to UDP port 1755 on the proxy.
Service 83	wmt-rtsp	Media streaming service that allows the device to redirect RTSP requests from Windows Media Service 9 clients to UDP port 5005 on the the CE.
Service 90-97	user configurable	User-defined WCCP services that support up to eight ports for each WCCP service. When you configure these user-defined services, you must specify whether to redirect the traffic to the HTTP caching application, to the HTTPS application, or to the streaming application on the proxy.

Table 2-2 Service Groups

Service Group	Type	Description
Service 98	custom-web-cache	Caching service that permits the device to transparently redirect HTTP traffic to the proxy on multiple ports other than port 80.
Service 99	reverse-proxy	Caching service that permits the device to redirect HTTP reverse proxy traffic to the proxy on port 80.

A service group is identified by Service Type and Service ID. There are two types of service groups:

- Well-known services
- Dynamic services

Well-known services are known by both the redirecting device and web caches and do not require a description other than a Service ID.

In contrast, dynamic services must be described to a redirecting device. The device may be configured to participate in a particular dynamic service group, identified by Service ID, without any knowledge of the characteristics of the traffic associated with that service group. The traffic description is communicated to the device in the WCCP2_HERE_I_AM message of the first web proxy to join the service group. A web proxy uses the Protocol, Service Flags, and Port fields of the Service Info component to describe a dynamic service. Once a dynamic service has been defined, the device discards any subsequent WCCP2_HERE_I_AM message that contains a conflicting description. The device also discards a WCCP2_HERE_I_AM message that describes a service group for which it has not been configured.

The numbers 0 to 254 are dynamic services, and the web cache service is a standard, or well-known, service. What this means is that when the web cache service is specified, the WCCP v2 protocol has predefined that TCP destination port 80 traffic is to be redirected. For the numbers 0 to 254, each number represents a dynamic service group. The WCCP proxies are to define a set of protocols and ports that are to be redirected for each service group. Then, when the device is configured with that same service group number (wccp 0 ... or wccp 1 ...), the device performs redirection on the specified protocols and ports as directed by the device.

Figure 2-7 is an example that shows Web-Cache Identity Info.

Figure 2-7 Web-Cache Identity Info

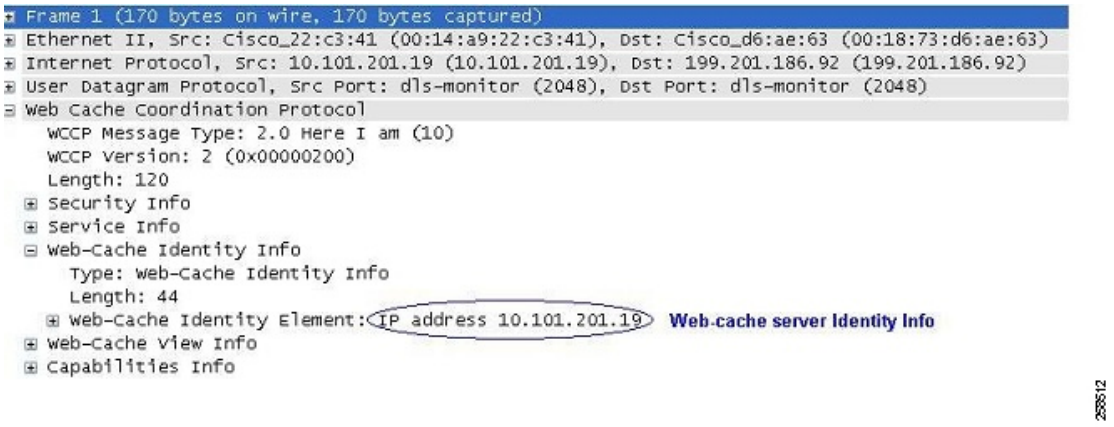

Figure 2-8 is an example that shows that the Web-Cache is part of service group 0.

Figure 2-8 Web-Cache Part of Service Group 0

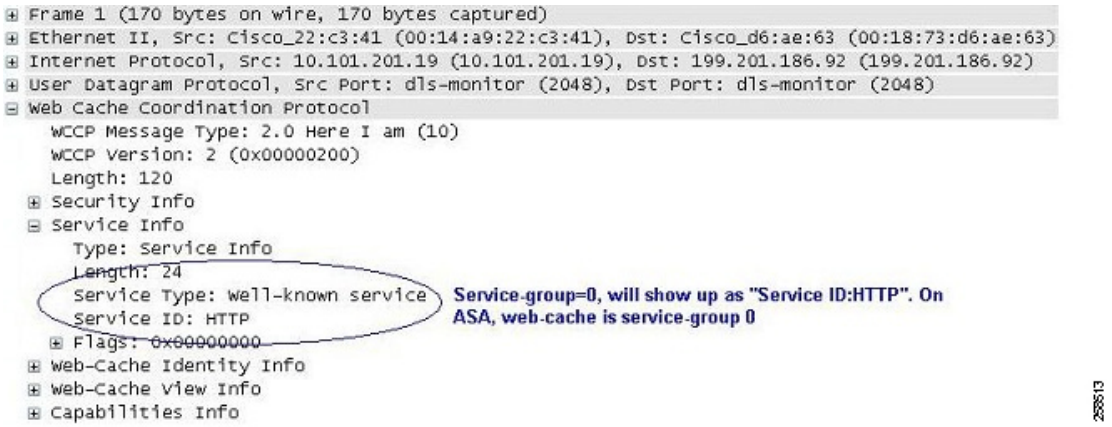

Figure 2-9 is an example that shows a Web-Cache server as part of custom service group 91 and the ports whose traffic is redirected to the server.

Figure 2-9 Web-Cache Server Part of Service Group 91 and Redirected Ports

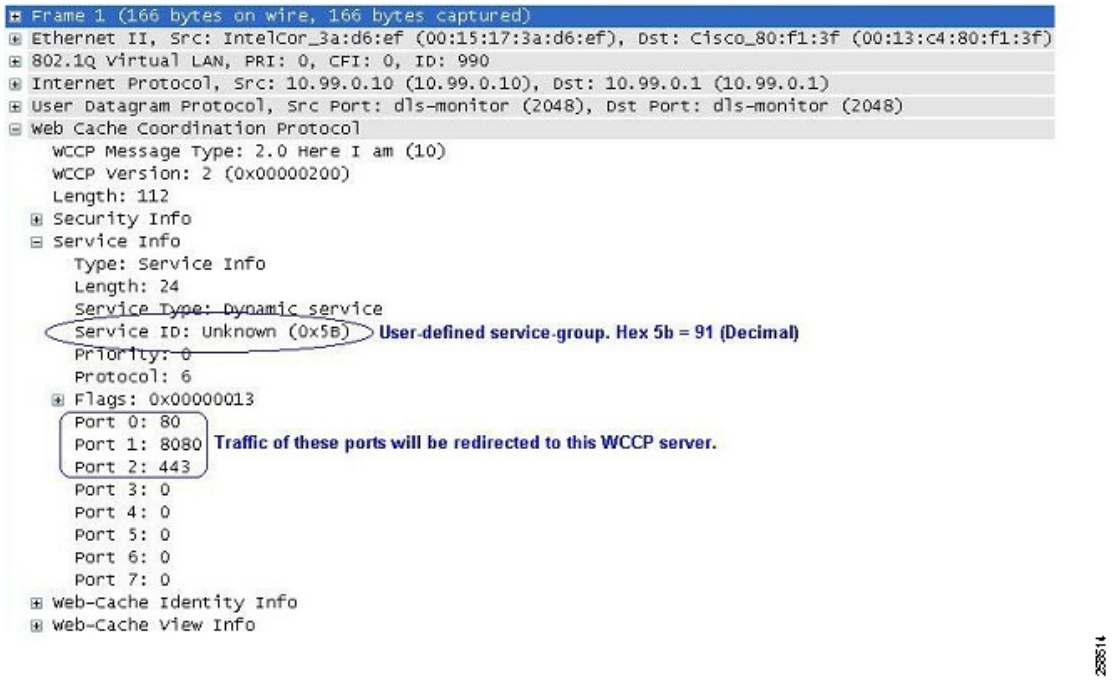

The device responds to a WCCP2_HERE_I_AM message with a WCCP2_I_SEE_YOU message.

- If the WCCP2_HERE_I_AM message was unicast, the router responds immediately with a unicast WCCP2_I_SEE_YOU message.
- If the WCCP2_HERE_I_AM message was multicast, the router responds with the scheduled multicast WCCP2_I_SEE_YOU message for the service group.

Figure 2-10 is an example of the device's "I See You" message, which shows that the router joins service group 91 (a custom service group) and redirects ports 80, 8080, and 443 to the web proxy server.

Figure 2-10 "I See You" Message

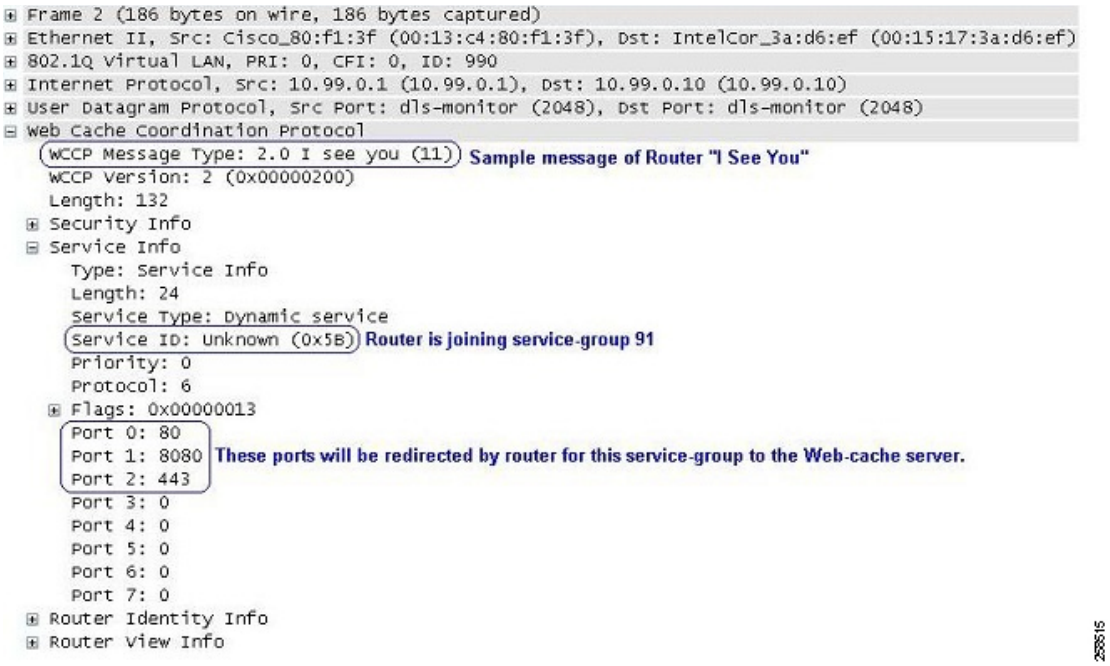

When you redirect traffic using WCCP, keep the following behavior in mind:

- The device selects the highest IP address configured on any interface as the WCCP router ID. This address is used to establish a GRE tunnel with the device. When the device redirects packets to the WCCP-enabled device, the device sources the redirect from the router ID IP address (even if it is sourced out another interface) and encapsulates the packet in a GRE header. For WCCP to work, the interface whose IP address is chosen as the router ID must be in the UP state and there must be a route to the device.

Figure 2-11 is an example of a GRE packet.

Figure 2-11 GRE Packet

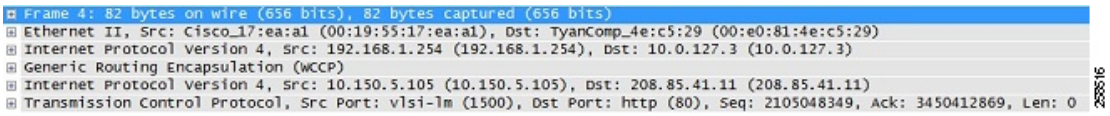

- An inbound access rule always takes higher priority over WCCP. For example, if an interface ACL does not permit a client to communicate with a server, then the matching traffic is simply dropped, it is not redirected.
- TCP intercept, authorization, URL filtering, inspection engines, and IPS features are not applied to a redirected flow of traffic.
- When a device cannot service a request and returns a packet to the redirecting device, then the contents of the traffic flow is subject to all the other configured features of the device.
- If you have two WCCP services and they use two different redirection ACLs that overlap and match the same packets (with a deny or a permit action), the packets behave according to the first service group found and installed rules. The packets are not passed through all service groups.

Limitations

Note These limitations may be based on the device. The limitations below are from a Cisco ASA.

- WCCP redirection is supported only on the ingress of an interface. The only topology that the ASA supports is when client and proxy are behind the same interface of the ASA and the proxy can directly communicate with the client, without going through the device.
- Multiple routers in a service group.
- Multicast WCCP.
- The Layer 2 redirect method.
- WCCP source address spoofing.
- Wide Area Application Services (WAAS) devices.
- AAA for network access does not work in combination with WCCP.
- Does not support IPv6 traffic for redirection.
- When the ASA determines that a packet needs redirection, it ignores TCP state tracking, TCP sequence number randomization, and NAT on these traffic flows.
- WCCP does not support ACLs that include a user, user group, service group, or a fully qualified domain name object.
- The maximum number of services, including those specified with a dynamic service identifier is 256.

Industrial Firewall (IFW)

Another key tenant in the Platinum and Gold architectures is the use of Industrial Firewalls to enforce security policies within the Cell/Area Zone. In some cases, the cloud connected device will communicate with Industrial Zone IACS assets to gather information such as device inventory and verify lifecycle status with the cloud. The Industrial Firewall can be configured to constrain the communication of the cloud connected device to appropriate IACS devices. The Industrial Firewall with CIP deep packet inspection (DPI) can also be configured to limit the types of CIP commands that can be sent to a device. The placement of the Industrial Firewall will dictate the amount of security granularity one can achieve with this device.

For a complete understanding of an IFW, see the *Deploying Industrial Firewalls within a Converged Plantwide Ethernet Architecture Design and Implementation Guide*:

- Rockwell Automation site:
  http://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td002_-en-p.pdf
- Cisco site:
  http://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/5-0/Firewalls/DIG/CPwE-5-IFS-DIG.html

Cloud Connected Devices to Cloud Applications Test Cases

CPwE Cloud Connectivity tested the following items:

- Cloud connected devices using FactoryTalk AssetCentre and ControlFLASH Plus connecting to the Rockwell Automation cloud API.

  - The Rockwell Automation Cloud API is a cloud residing application providing lifecycle status information.
- FactoryTalk Activation Manager's ability to reach out to the cloud for activation validity purposes such as Get New Activation, Rehost and Renew.
- IDMZ and Industrial firewall configurations as well as the Cisco WSA configurations while positioned within the IDMZ.

Platinum Security Architecture-Cloud Connected Devices Use Case

The Platinum security reference architecture that was tested contained the following key components that were configured and functionally verified:

- Cisco Web Security Appliance
- FactoryTalk AssetCentre Server and Client
- ControlFLASH Plus connecting to various Rockwell Automation assets
- FactoryTalk Activation Manager
- Centralized Activation Server
- Standalone FactoryTalk Activation Manager client
- DNS server required to resolve the various cloud IP Addresses.
- Allen-Bradley Stratix 5950 Industrial Firewall—This firewall was configured to allow DNS requests and TLS traffic to the Cloud.
- IDMZ ASA firewalls—This firewall was configured to allow DNS requests and TLS traffic to the Cloud.

Gold Security Architecture—Cloud Connected Devices Use Case

The gold security reference architecture that was tested contained the following key components that were configured and functionally verified:

- FactoryTalk AssetCentre Server and Client
- ControlFLASH Plus connecting to various Rockwell Automation assets
- FactoryTalk Activation Manager
- Centralized Activation Server
- Standalone FactoryTalk Activation Manager client
- DNS serverr equired to resolve the various cloud IP Addresses.
- Allen-Bradley Stratix 5950 Industrial Firewall—This firewall was configured to allow DNS requests and TLS traffic to the Cloud.
- IDMZ ASA firewalls—This firewall was configured to allow DNS requests and TLS traffic to the Cloud.