CoIP® Platform
Zero Trust Architecture

Implementing NIST SP 800-207 with a Zero Trust Overlay

zentera™

About CoIP Platform

Zentera's CoIP Platform provides a comprehensive suite of tools to help organizations rapidly transition to a Zero Trust security model. Onboarding a resource to CoIP Platform allows an administrator to define and enforce security controls limiting access to the resource to authenticated and authorized users, devices, and applications. It monitors network activity to expose implicit trust relationships and provides the necessary tools to convert them to explicit policies for security enforcement.

CoIP Platform is uniquely positioned as a security overlay to the existing networks. This allows resources to be onboarded and brought into a Zero Trust posture without modifying the resource's network infrastructure. The overlay approach enables Zero Trust controls to be quickly deployed to secure resources in bare metal, hyperconverged infrastructure, cloud, and even OT environments.

About this Document

This document covers CoIP Platform's theory of operation and demonstrates how CoIP Platform maps to the Zero Trust Architecture described in NIST SP 800-207.

For more detailed technical and operational documentation, please refer to our support site, at support.zentera.net (registration may be required).

CoIP Vision and Philosophy

The core design objective of CoIP Platform is to support the dynamic creation and management of security control points that can enforce granular access policies for network traffic in a scalable way.

About Scalability

By "scalable," we do not refer only to the number of devices, resources, or users being managed, but also the complexity of the network and application architectures, hosting environments, and device types that may be encountered. An enterprise might have many applications hosted in a datacenter; providing Zero Trust access to these resources is a clear design goal for any vendor. Peeling back the onion, we quickly encounter hidden complexity – resources and services have dependencies on details that may not be visible from a white-boarded block diagram. For example:

- Load balancers deployed to protect service availability
- IP-based power distribution units that monitor and control the power on each datacenter rack
- Management access to each server blade's baseboard management controller (BMC)

The ultimate goal of any Zero Trust security is to protect the resource or service from attacks – it should not matter whether an attack is lateral (already "in the network") or is coming from the north-south remote access. Clearly, therefore, it is *not sufficient* to focus only on the "front-door" access to the resource. Zentera recognized this requirement years ago and has developed a wide range of onboarding methods and controls to enable a path to holistic protection for the entire resource.

About Infrastructure and Operations

Zentera also recognized that it is not realistic for organizations to embark on a massive infrastructure upgrade in order to enforce Zero Trust controls. Some of the challenges and problems with a "network rip and replace" approach include:

- It triggers significant planning for procurement and implementation;
- It creates downtime for the resource during the infrastructure changeover; and
- It potentially impacts other applications and resources which are dependent on the existing infrastructure, creating the opportunity for an extended period of instability.

Zentera's approach addresses this concern by deploying security as a software overlay, totally decoupled from existing network infrastructure. Resources may be onboarded to CoIP Platform by installing software without a restart, minimizing any potential for service disruption; policies can be developed and applied to step up the security level gradually to minimize impact to existing applications and resources. This also allows for "break glass" scenarios for OT environments, where security may need to be rolled back temporarily to support availability.

The overlay approach also enables the onboarding of a resource to CoIP Platform to be completely decoupled from all other resources in the same facility or datacenter. In contrast, any infrastructure upgrade or change can immediately and irreversibly impact adjacent resources that happen to share the underlying networks.

Our design and architecture decisions that have been made to enable Zero Trust controls as close as possible to the user, application, or resource that is being accessed.

Foundational Concepts of CoIP Platform Security

Applications

An *Application* in CoIP Platform is a logical group of resources with a shared role. This construct allows the resources to be managed as a group, both from the perspective of managing Application Chambers and defining Access Policies.

The first diagram below illustrates a standard 3-tier web services architecture. The constituent endpoint servers have been split into 3 Applications ("Presentation," "Logic," and "Data") which allows full control of the computing flow into, out of, and between the tiers.

The second diagram illustrates an alternate model, where all resources are grouped into one large Application. This flexibility allows administrators to trade off security granularity for ease of management.

Application Chambers

CoIP Platform defines an *Application Chamber* for each logical Application A Chamber is like an application firewall that protects a group of endpoints; Chamber Policies apply to all endpoints within the Chamber.

Application Chambers are typically used to filter the endpoint's physical network to cloak them from the shared network and to enforce access policies, such as allowing network services like DNS, while blocking other unauthorized access.

In the 3-tier web services example below, Chamber Policies may be set to deny unauthorized inbound traffic by default, or to block outbound-first access to the Internet to block data leaks.

Access Policies

An Access Policy defines allowed access to resources based on specified attributes, which may include user and endpoint identity, user role, identity of the client or server application, location, and other factors.

Access Policies are used to enable access from users to a Chamber, or between Chambers. Notably, the IP address of the resource is *not* configured in the policy; it is automatically included and enforced by the system. This means that assigning an endpoint to a Chamber automatically applies all relevant Access Policies to that endpoint.

An Access Policy also defines whether the CoIP Overlay or physical network are used for the access.

Granular Model of a 3-Tier Web Service

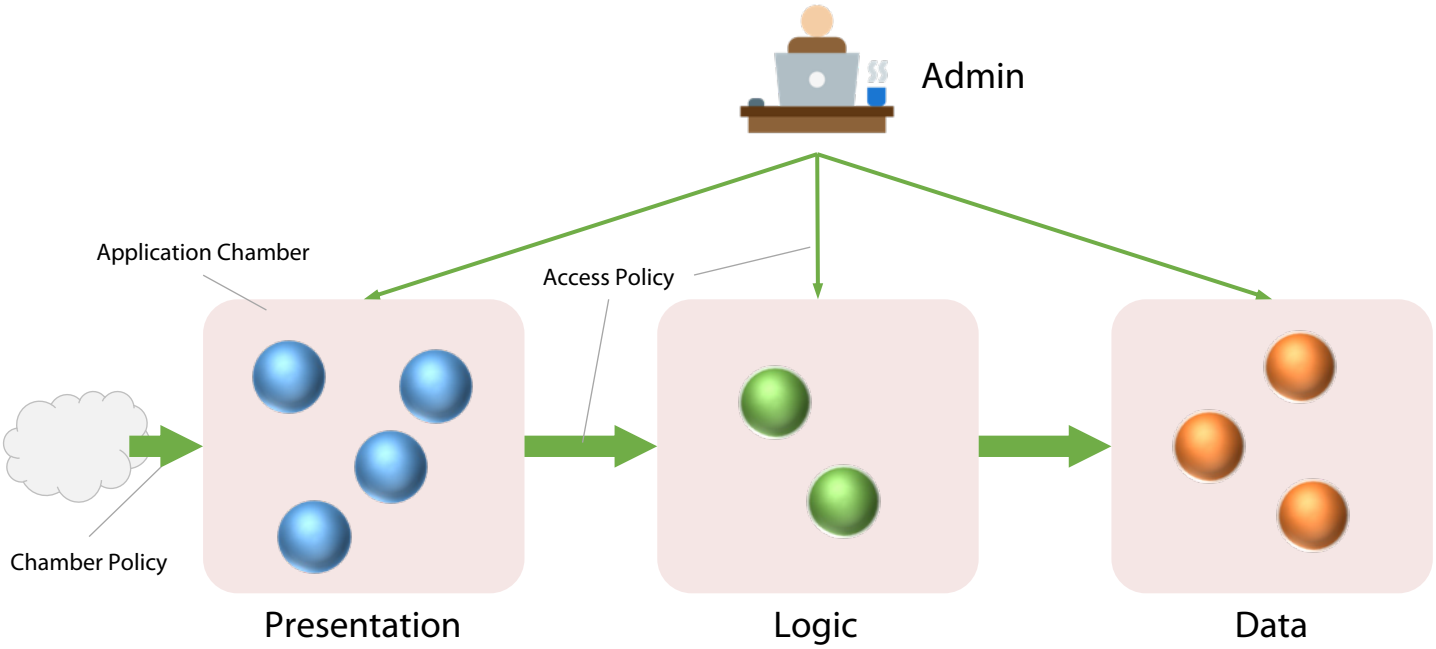

Coarse Model of a 3-Tier Web Service

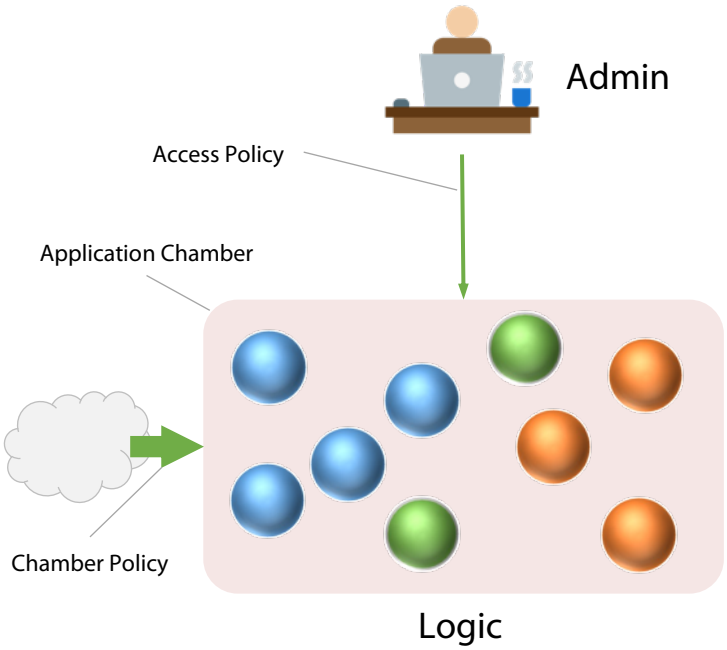

Resource Onboarding

CoIP Platform supports both agent-based and agentless architectures for resource onboarding. In either case, the CoIP architecture is designed to enable security enforcement as close as possible to the user, application, or resource that is being accessed. This provides a host of benefits which simplify the setup and maintenance.

Benefits of Enforcing Security Controls at the Application Edge

- ✓ A segmentation boundary in the endpoint OS provides ultimate granularity for policy controls
- ✓ Simpler to define policies for individual resources and assets, instead of coupled policies for aggregated applications
- ✓ Security controls deployed to the endpoint can travel with the workload for migration and mobility
- ✓ Deploying inline with endpoints eliminates the need to use IP routing to steer traffic to the security checkpoint
- ✓ Makes security controls environment-independent, enabling consistent controls to be applied on any compute stack or cloud

As noted earlier, CoIP Platform enables resources to be onboarded without immediately changing security policies or routing.

Overlay Networking

In addition to enforcing security at the endpoint, CoIP Platform is also able to tunnel traffic between endpoints in overlay tunnels; this capability can be used to connect applications across network domains or to secure traffic in the LAN environment.

The CoIP Overlay is typically encrypted with TLS 1.3, although tunneling without encryption is also supported to avoid double encryption.

In the 3-tier web service example of the previous section, the CoIP Overlay may be used to enable administrator ssh access to an endpoint, even though 22/tcp access from the physical network is blocked by the Chamber.

CoIP Platform Components

zCenter

Orchestrator for Centralized Policy Management

At the heart of CoIP Platform is a centralized orchestration service, zCenter. zCenter is running the control plane and responsible for authenticating all subjects and resources (users, devices, and application) and maintaining all policies that authorize access. zCenter may be configured to connect to external systems to provide authentication services (for example, identity providers) or information needed to make a policy decision (for example, IP geolocation providers).

When a network session is created, zCenter is responsible for determining whether the session is allowed by the current set of security policies. If allowed, zCenter orchestrates a data connection; otherwise, the connection is denied.

ZNS

ZNS Network Switch

The ZNS is a switch used to run the data plane and connect overlay connections across disconnected network domains. In the Zero Trust architecture model, ZNS can be viewed as the Zero Trust edge gateway with private resources behind the distributed ZNS gateways.

ZNS nodes can be clustered to provide redundancy and optimize traffic routing.

zLink

zLink Agents

The zLink agent is used to onboard compute endpoints, such as bare metal machines, virtualized instances, or cloud instances. zLink is run in th OS user space, and can install to a wide range targets, including Windows, Linux, and Mac OS, and includes end-of-life operating systems such as Win XP, Windows Server 2003, and RedHat 5. It can also be deployed to Raspberry Pi systems.

The zLink agent enforces the Application Chamber and initiates Access Policies. It is responsible for gathering telemetry about the endpoint as well as the applications using the network.

zLink agents can be managed and upgraded by IT directly through the zCenter portal, or through standard IT tools.

Micro-Segmentation Gatekeeper (MSG)

MSG

The Micro-Segmentation Gatekeeper (MSG) is a hardware appliance that sits inline between network and application devices and enforces security policies. It enables the onboarding of devices which cannot support the zLink Agent. The MSG enforces the Application Chamber for protected downstream devices.

The MSG typically inserts inline between the device and the access switch and supports Layer 2 traffic, with a hardware bypass to maintain availability in the event of power loss or system failure. The MSG may alternately be placed between the access switch and the aggregation switch, providing options to balance cost against the granularity of enforcement, with interface options that range from Gigabit Ethernet to 100G.

Gateway Proxy

Gateway Proxy

The Gateway Proxy is a virtual appliance that terminates remote overlay traffic, retransmitting the traffic on the local network.

Deployed at the same level of hierarchy as endpoints, it can filter remote access traffic from or to a remote client or server. However, as a proxy it cannot create or enforce an Application Chamber.

Mapping CoIP Platform to NIST SP 800-207

NIST SP 800-207 illustrates several possible ways to achieve a Zero Trust Architecture. CoIP Platform takes a blended approach to support all architecture models listed in the document and achieve the benefits previously listed – rapid implementation of Zero Trust, on any technology stack and architecture, without requiring infrastructure changes.

The components listed in the previous section map directly to the logical components of a Zero Trust Architecture, listed in SP 800-207 §3. The diagram shown below illustrates the basic CoIP deployment in the style of SP 800-207, and illustrates the following mappings:

- The zCenter orchestration implements the Policy Engine and Policy Administrator, together the *Policy Decision Point*
- The zLink Agent implements the Policy Enforcement Point

Basic CoIP "Type 1" LAN Mapped to NIST SP 800-207

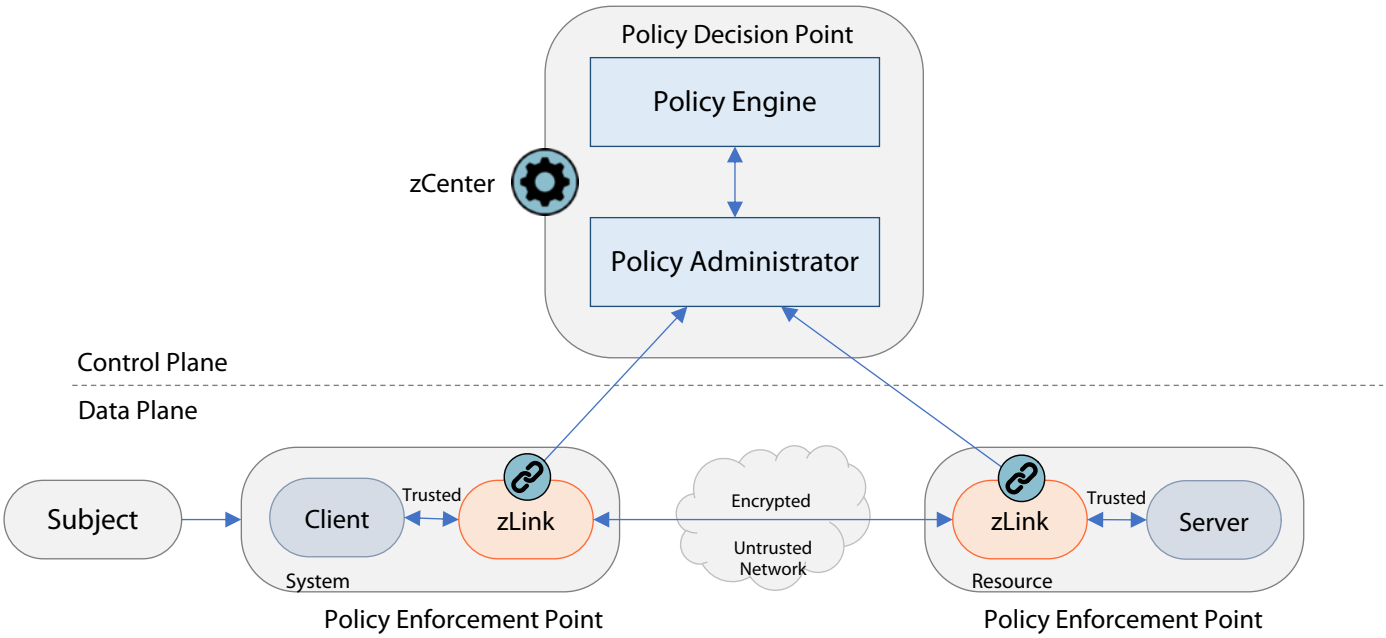

Some enhancements of this CoIP Zero Trust Architecture over the base SP 800-207 architecture worth noting include:

Mutual Zero Trust

As shown in the diagram, zLink is deployed on *both* the source and destination. This creates a Zero Trust enforcement point ensuring that both source and destination are authenticated and authorized by zCenter.

No Implicit Trust Zone

Section 3.2.1 in NIST SP 800-207 shows a gateway and resource adjacent to each other, as an implementation of the CSA Software Defined Perimeter. There is an implicit trust zone between the gateway and the resource that must be secured in some other way (for example, by implementing a Resource Enclave as shown in §3.2.2 / Figure 4).

The CoIP LAN Type 1 model shown above assumes the entire network is untrusted and secures the access (including optionally encrypting the session) all the way to the OS of the resource without creating an implicit trust zone.

Bidirectional Communications

The CoIP model also makes the Zero Trust flow truly symmetric, so that bidirectional communications become possible. With CoIP, it is possible to configure a policy that enables mutual communications between two servers – for example, to secure an IT management workflow using Zero Trust between an application server and a management server.

Micro-Segmentation

The zLink Agent can enforce an Application Chamber around the source system or the destination resource to cloak them on the shared network.

Other CoIP LAN/WAN Types

The previous example illustrated the CoIP LAN "Type 1" deployment model. CoIP supports various deployment models, which can be mixed to model different cases as outlined below.

Type	Source	Destination
1	zLink Agent	zLink Agent
2	zLink Agent	Gateway Proxy
3	Gateway Proxy	Gateway Proxy
4	User	zLink Agent
5	User	Gateway Proxy

Types 1/4

Redrawn using Zentera's typical architecture schema, the CoIP LAN Type 1 (server to server) and Type 4 (user to server) connections are shown below.

Note that for simplicity, the zCenter and ZNS are shown outside of any particular network; they can be placed anywhere, as long as there is a route to them from all system components.

CoIP LAN/WAN, Types 1 and 4

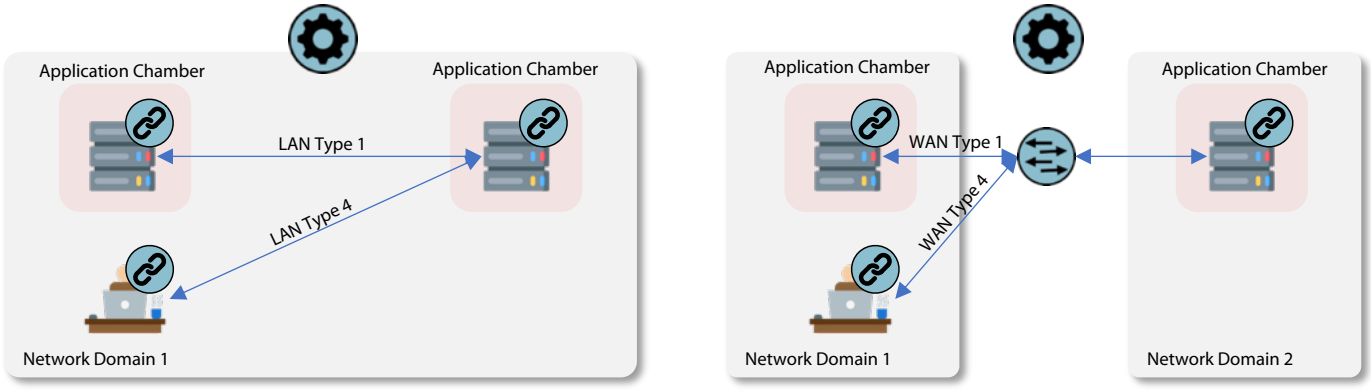

Types 2, 3, and 5

The other types involve a Gateway Proxy, and in doing so implement variations on the Device Agent/Gateway model of NIST SP 800-207 §3.2.1.

Although a Gateway Proxy is capable of filtering traffic that flows through it, a Gateway Proxy does not perform full Zero Trust identity checks on the resources behind it, nor can it implement an Application Chamber. The zLink Agent or Micro-Segmentation Gatekeeper are a better fit for use cases that require full segmentation of the resource.

CoIP LAN /WAN, Types 2, 3, and 5

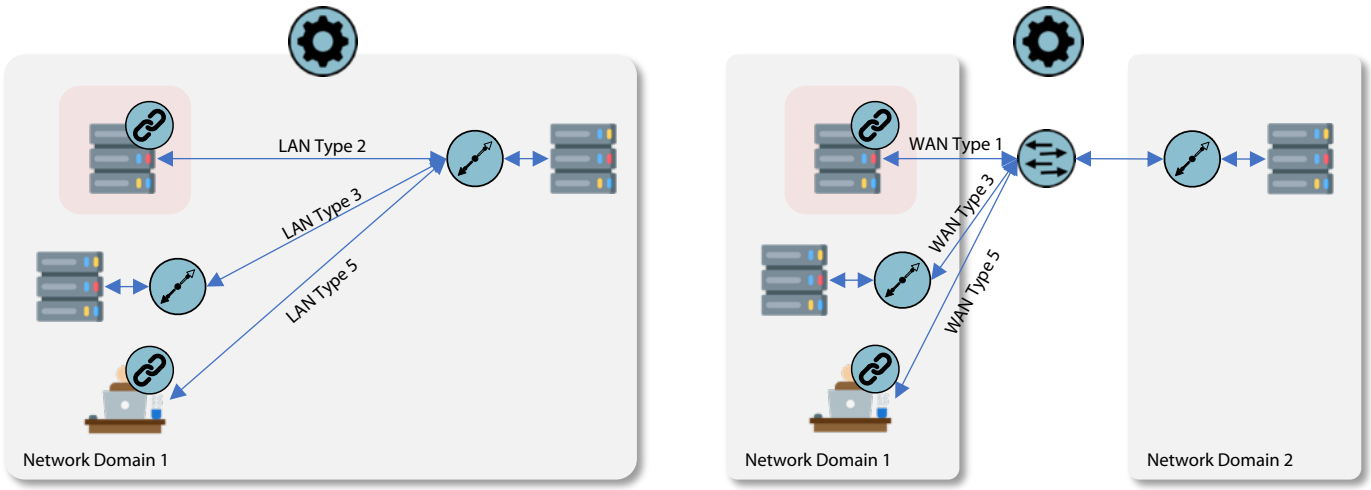

CoIP Platform Trust Algorithm

NIST SP 800-207 §3.3.1 outlines various trust models that can be used to control access. CoIP Platform uses a criteria-based and singular trust algorithm. The design intent is to keep system behavior repeatable, so that it is easy to understand and debug access control allow/deny decisions. At the same time, we provide flexible criteria to make it easy to specify the desired system behavior.

Implementing the NIST SP 800-207 Use Cases with CoIP Platform

NIST SP 800-207 outlines several use cases that can be secured with Zero Trust. It's simple to implement all cases with CoIP Platform.

Enterprise with Satellite Facilities

This use case, described in §4.1, covers the use of Internet backhaul to connect remote workers to the corporate resources they need to access. CoIP can be used to eliminate the need for a VPN, and also enforces a consistent Zero Trust experience whether the employee is remote or is in the Enterprise HQ.

This is simple to implement with CoIP Platform, either deployed by the enterprise, or with Zentera's SaaS version of CoIP Platform, Zentera Air.

As shown below, resources can be onboarded with the zLink Agent, or with a Gateway Proxy. If there is a requirement to segment the resource, the Micro-Segmentation Gatekeeper can be used instead.

Enterprise with Satellite Facilities

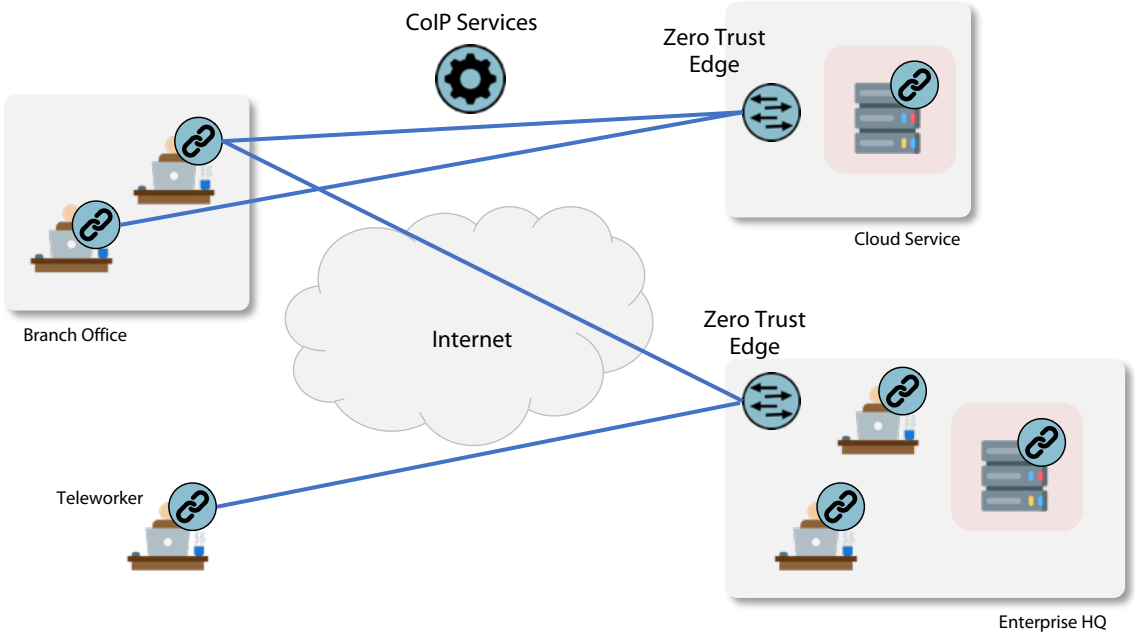

Multi-cloud / Cloud-to-Cloud Enterprise

This use case, described in §4.2, uses a Zero Trust connection to connect a web front-end in one cloud provider to a database backend in a separate cloud provider.

This use case is also simple to implement with CoIP Platform. The enterprise may also deploy ZNS along private routes to ensure that the enterprise-to-cloud and cross-cloud traffic uses dedicated backhaul, for example, MPLS.

As the zCenter only implements the control plane, its placement is not critical, so long as all endpoints have a route to it.

The zLink Agent is used to onboard cloud instances and can apply the Application Chamber to protect the instance from unauthorized access. This allows the enterprise to apply a Zero Trust segmentation model as part of the security stack to meet the cloud concept of shared responsibility.

Multi-cloud / Cloud-to-Cloud Enterprise

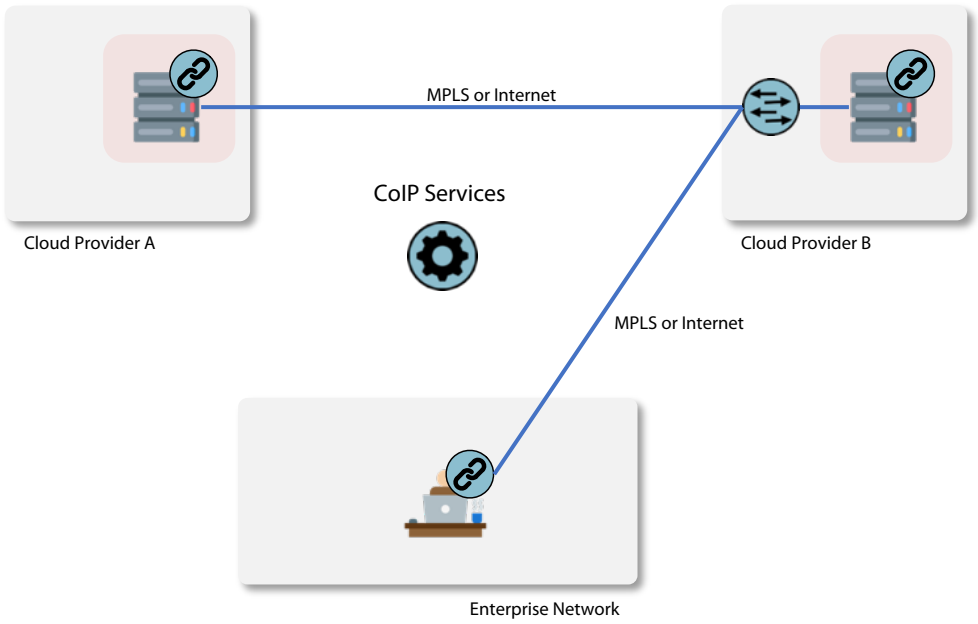

Enterprise with Nonemployee Access

This use case, described in §4.3, uses Zero Trust to provide differentiated levels of access for users of different roles.

The scenario described in SP 800-207 includes a Conference Center, where visitors are allowed to access the Internet but not enterprise resources. This is already a common feature in enterprise conference rooms; guest Wi-Fi access typically does not allow visitors access to the enterprise network.

The scenario also depicts a contractor, who has access to the local network.

There are therefore actually two critical requirements here:

1. Enable access for the employee, who is also on the conference center Wi-Fi

2. Make it possible for visitors/contractors to be able to use the corporate network to get their job done in a secure manner

The first requirement is straightforward – it is equivalent to enabling remote access for the employee.

The second, however, requires more thought. Simply directing the contractor service tech's traffic to the Internet as a solution simply preserves the enterprise network as a large implicit trust zone. A Zero Trust enterprise should be able to segment critical resources, such as the Internal App Server, Database, and even the Contracted HVAC, to remove any implicit trust. This model is shown with CoIP Platform below; even though the contractor service tech has access to the corporate network, he is not authorized to access resources protected by Zero Trust other than the Contracted HVAC.

A third potential requirement, which is not shown in the scenario, is to enable remote access for the contractor service tech, so he is able to access the Contracted HVAC without being physically on-site. This is included in the CoIP version of this use case, shown below. If this is done, the enterprise may also choose to grant the contractor only guest Wi-Fi access when on-site.

In this case, the zCenter and ZNS services are shown in the corporate DMZ.

Enterprise with Nonemployee Access

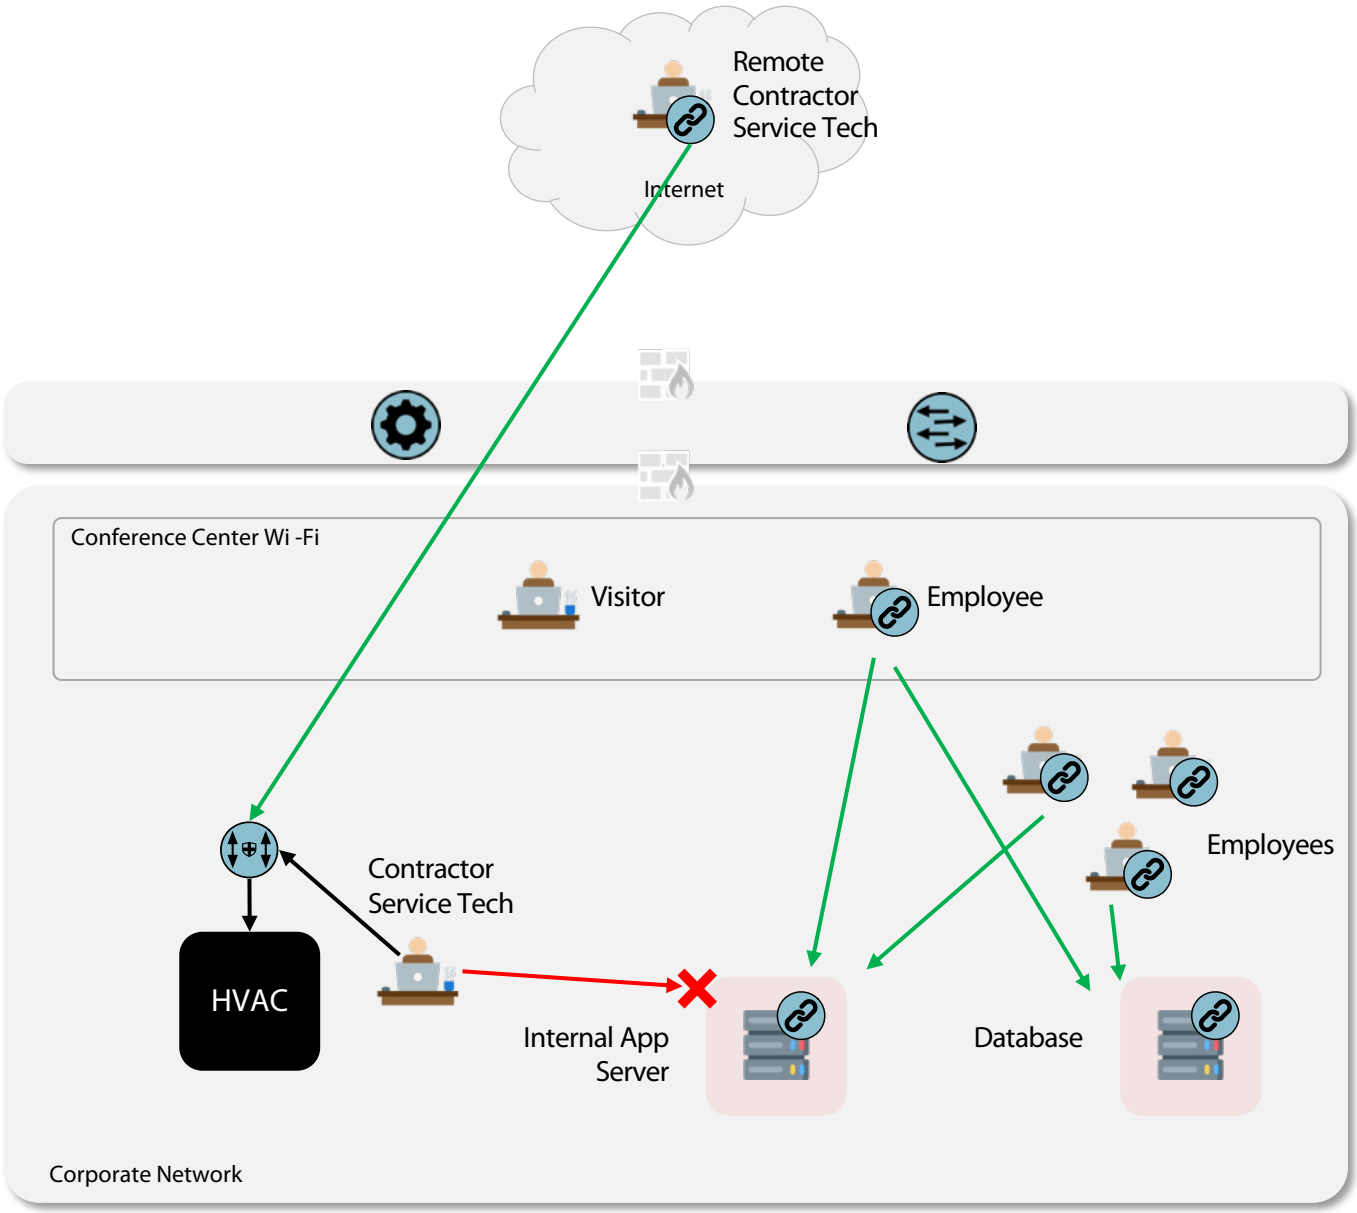

Collaboration Across Enterprise Boundaries

This use case, described in §4.4, uses Zero Trust to provide least-privilege access to resources in a network controlled by a third party.

In this case, we show the zCenter and ZNS are shown in Enterprise A's corporate DMZ. SP 800-207 suggests that federated ID may be desirable. It is our experience that for truly critical resources, enterprises may not prefer for management of identities and policy to be the responsibility of the partner enterprise IT.

Figure 11 in SP 800-207 suggests that the resources are databases, but it is also possible that the users in Enterprise B may have OS level login access to the two database machines. Therefore, it is important to consider applying segmentation controls, such as Application Chambers, to both machines, so that Enterprise B's users cannot abuse their access rights to move laterally within Enterprise A.

Collaboration Across Enterprise Boundaries

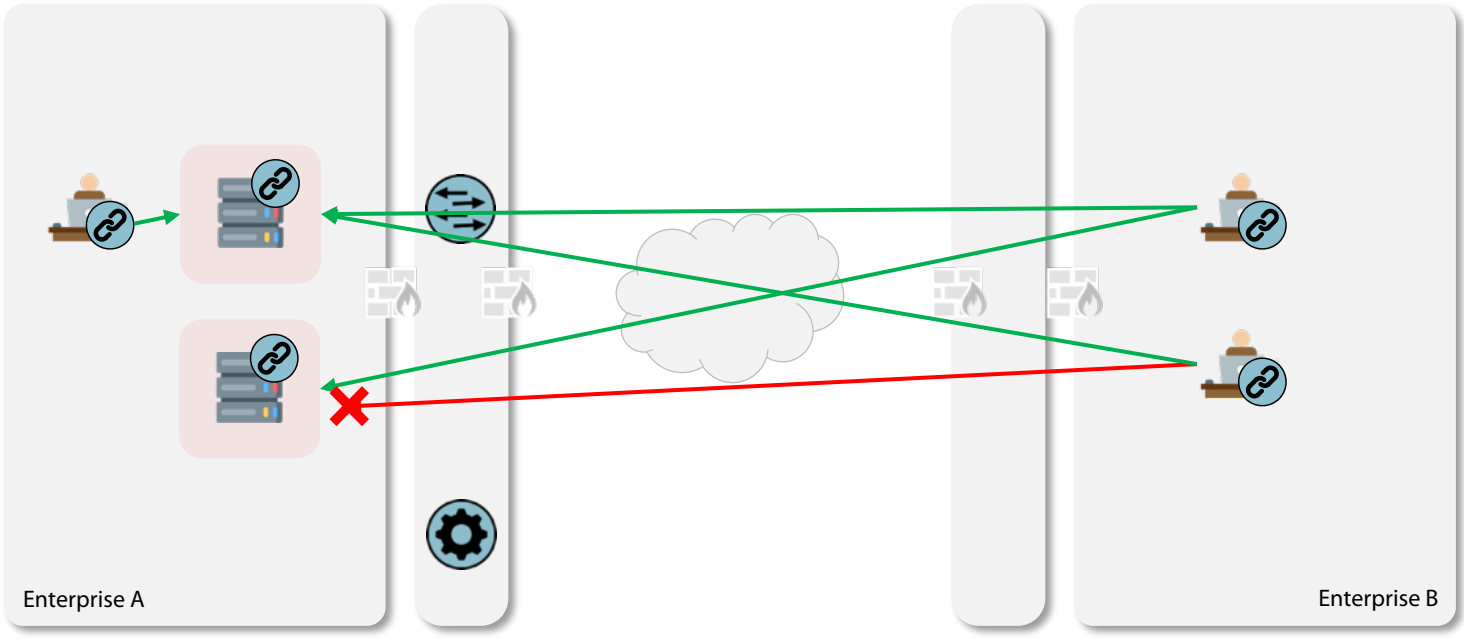

CoIP® Platform Zero Trust Architecture
Implementing NIST SP 800-207 with a Zero Trust Overlay

About Zentera

Zentera Systems is the leader in Zero Trust Security solutions for the digitally-transformed enterprise. Founded by experts in networking, security, and remote access, we offer award-winning Zero Trust networking, security, and multi-cloud connectivity that overlays any fragmented infrastructure and deploys rapidly on premises or as a service.

Our global enterprise customers and network of partners use our products to secure employee and third-party access, protect against data leaks, and instantly defend applications in complex hybrid and IT/OT environments. Based in Silicon Valley, we have received numerous recognitions, including Cool Vendor for Cloud Security by Gartner.

More Resources

On the web:
www.zentera.net

Email:
sales@zentera.net

Phone:
+1 (408) 436-4811

zentera™