zentera™

Practical Considerations for Adopting Zero Trust Security

SEI Zero Trust Industry Day
August 31, 2022

Mike Ichiriu
VP of Marketing and Product

Outline of Today's Session

Addressing Zero Trust Confusion

- Defining Zero Trust (again)
- Check our definitions

About Zentera's Zero Trust Fabric

- How it works
- How it applies to the SEI example RFI

Avoiding Operational Pitfalls Created by Zero Trust Confusion

- Reviewing of some common pitfalls we've seen

Addressing Zero Trust Confusion

Zero Trust is…

- A security model where authenti
  authorization are used to appro
  network connection; and all un
  or unauthorized connection
  approved

Expanded requirements: n
Trust, but still must be add
Trust Architecture

- Required by M-22-09 or
- e.g. LAN encryption, encry
  specific attributes for AB

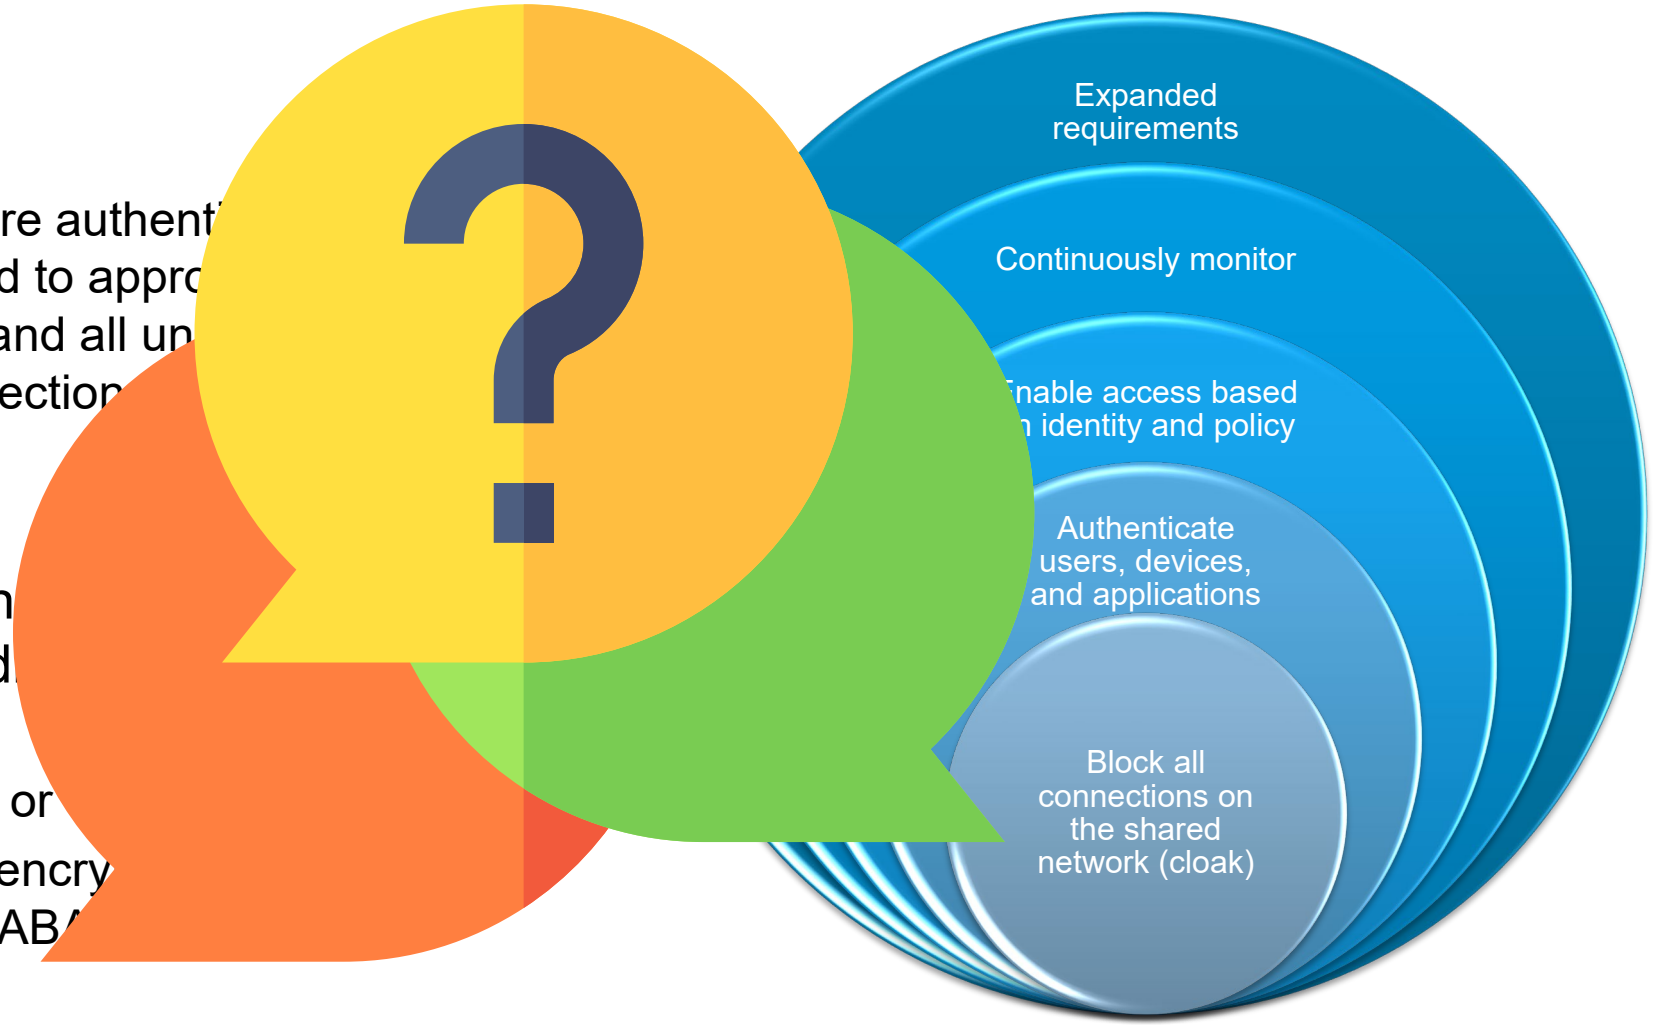

zentera™

let's play a game!

Which of the following completely enforces Zero Trust?

Easy: 10 points

a) VPN
b) Firewalls
c) VLANs
d) Cloud VPCs
e) None of the above

These legacy solutions either lack the concepts of user, device, and application identity, or lack identity-based controls

Which of the following completely enforces Zero Trust?

Medium: 20 points

a) IAM
b) PAM
c) SSO
d) None of the above

These functions may be *involved* in a Zero Trust solution, but don't *enforce* network security controls

Which of the following completely enforces Zero Trust?

Hard: 30 points

a) ZTNA
b) Micro-Segmentation
c) None of the above
d) It depends

Is the ZTNA a gateway that enforces Zero Trust only for remote access? Is the micro-segmentation Zero Trust Segmentation (based on identity)?

Which of the following completely enforces Zero Trust?

Master: 100 points

a) ZTNA Gateway + VLAN
b) Web application + IAM
c) VPN + IAM + MFA + NGFW
d) None of the above

Combinations of technologies may reduce attack surface, but don't get you all the way to Zero Trust with network security enforcement based on identity-based AuthN/AuthZ

Zentera's Zero Trust Fabric

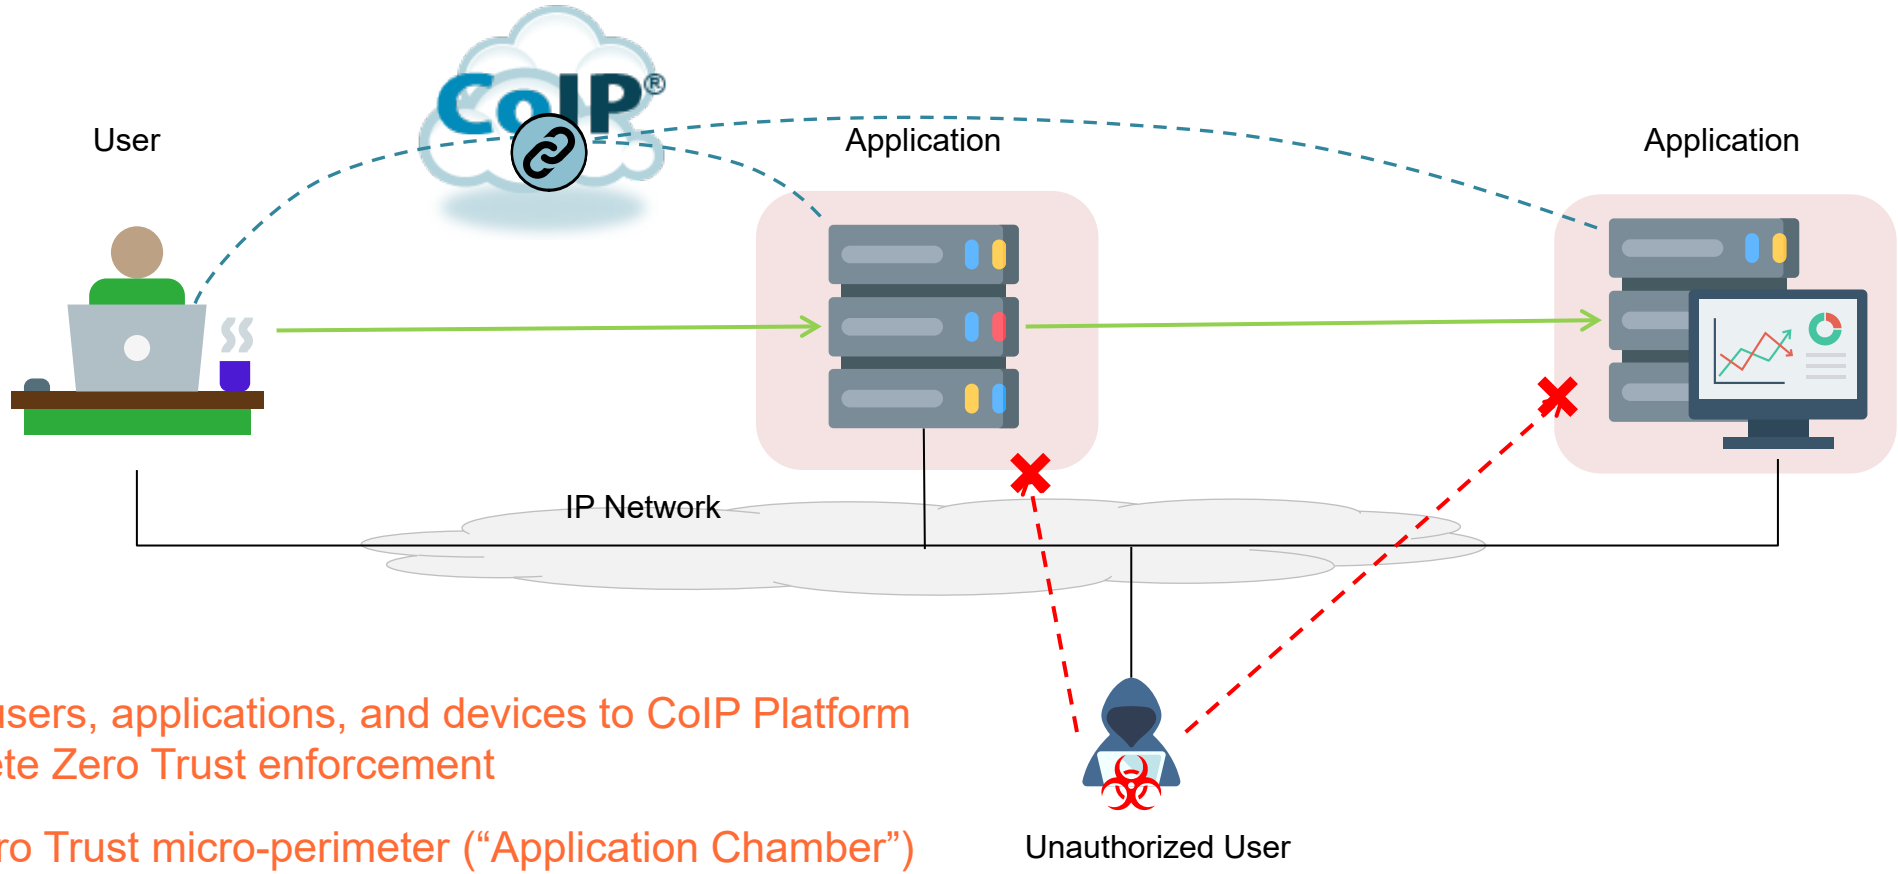

Onboard users, applications, and devices to CoIP Platform for complete Zero Trust enforcement

Create Zero Trust micro-perimeter ("Application Chamber")

Zero Trust authentication and authorization performed at the application – not in the network

Applications are protected from unauthorized remote or lateral access

Zentera's Zero Trust Fabric

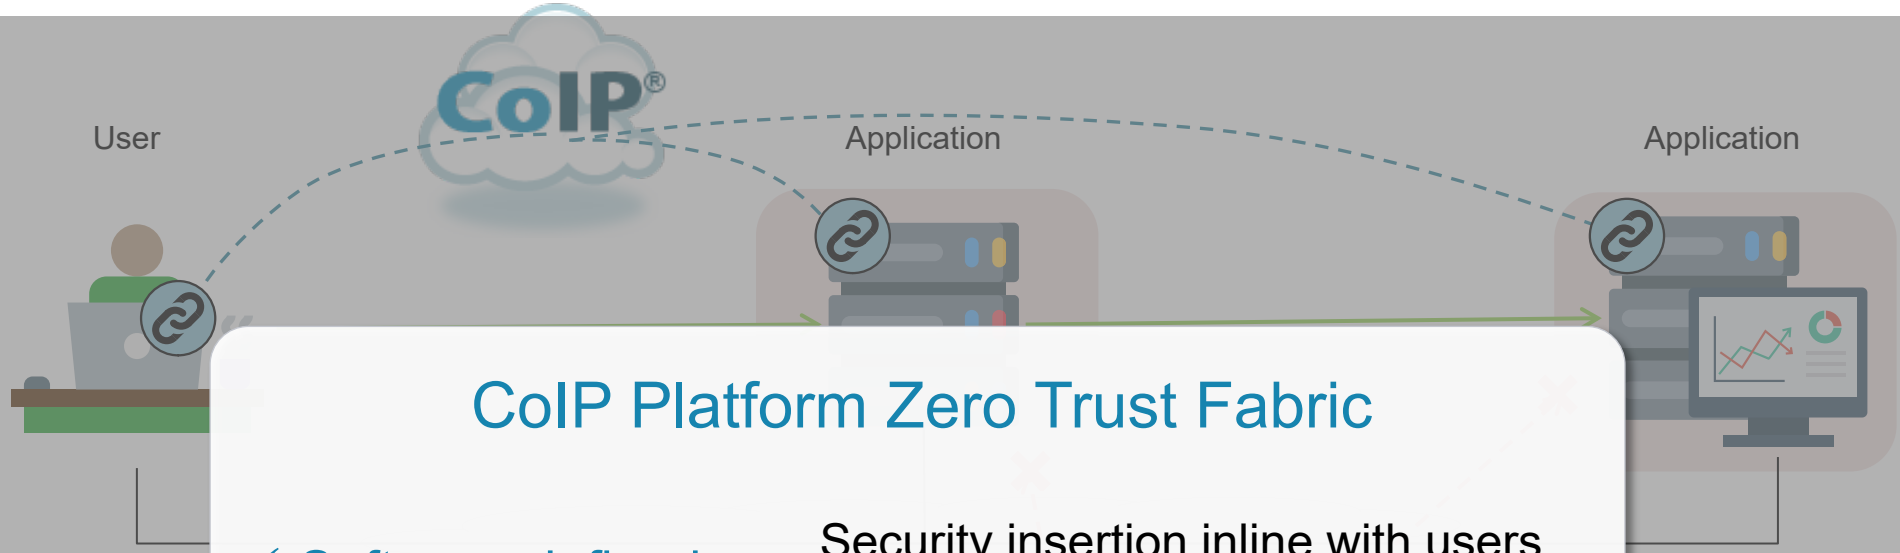

CoIP Platform Zero Trust Fabric

- ✓ Software-defined — Security insertion inline with users, applications, and devices
- ✓ Overlay — Deploys on top of existing networks and applications without changing their architecture
- ✓ Consistent — Controls work the same way in all hybrid environments

Onboard users, applications, and devices to CoIP Platform for complete Zero Trust enforcement

Create Zero Trust micro-perimeter ("Application Chamber") Multiple Users

Zero Trust authentication and authorization performed at the application – not in the network

Applications are protected from unauthorized remote or lateral access

zentera™

How It Works

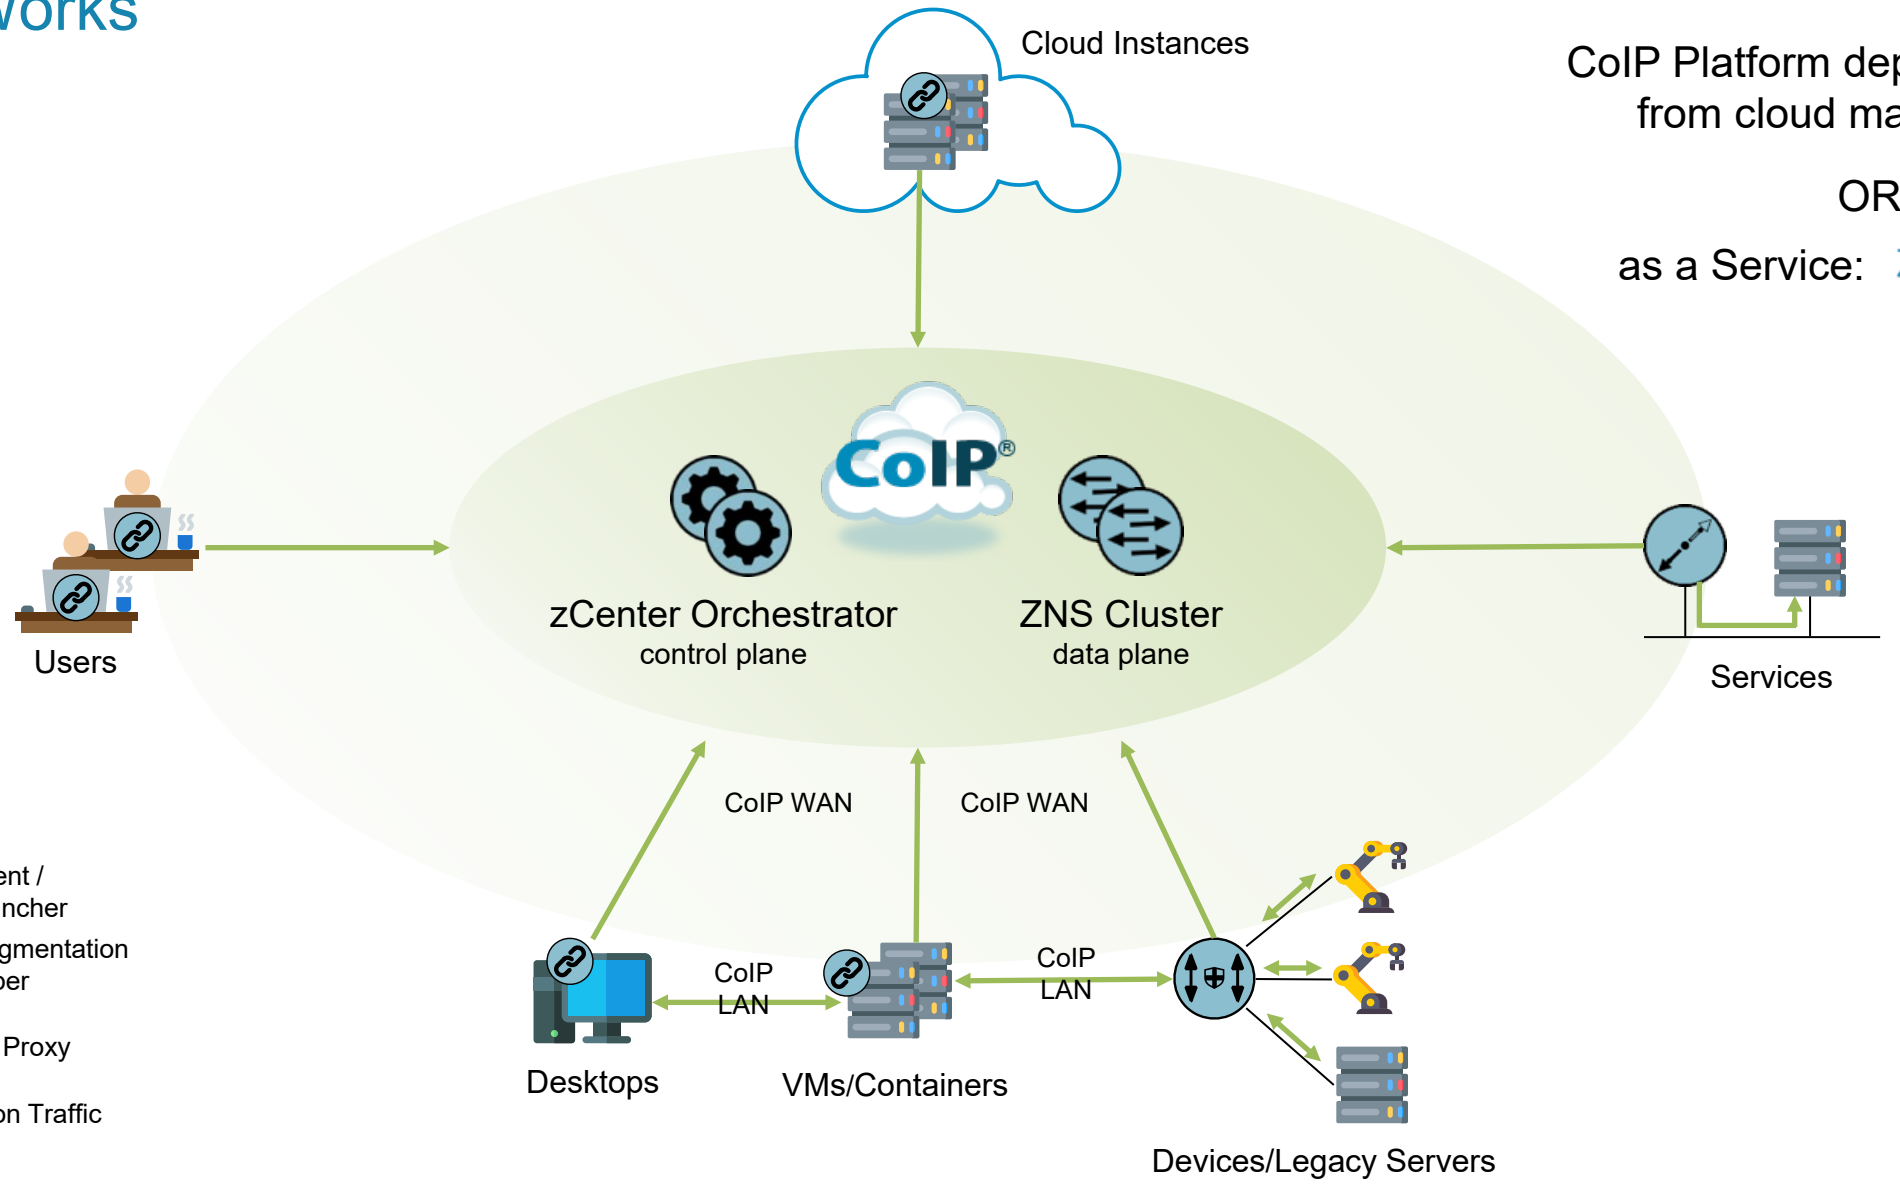

CoIP Platform deploys on-prem, from cloud marketplaces

OR

as a Service: zentera air

Zentera Zero Trust Architecture to Address SEI Example RFI

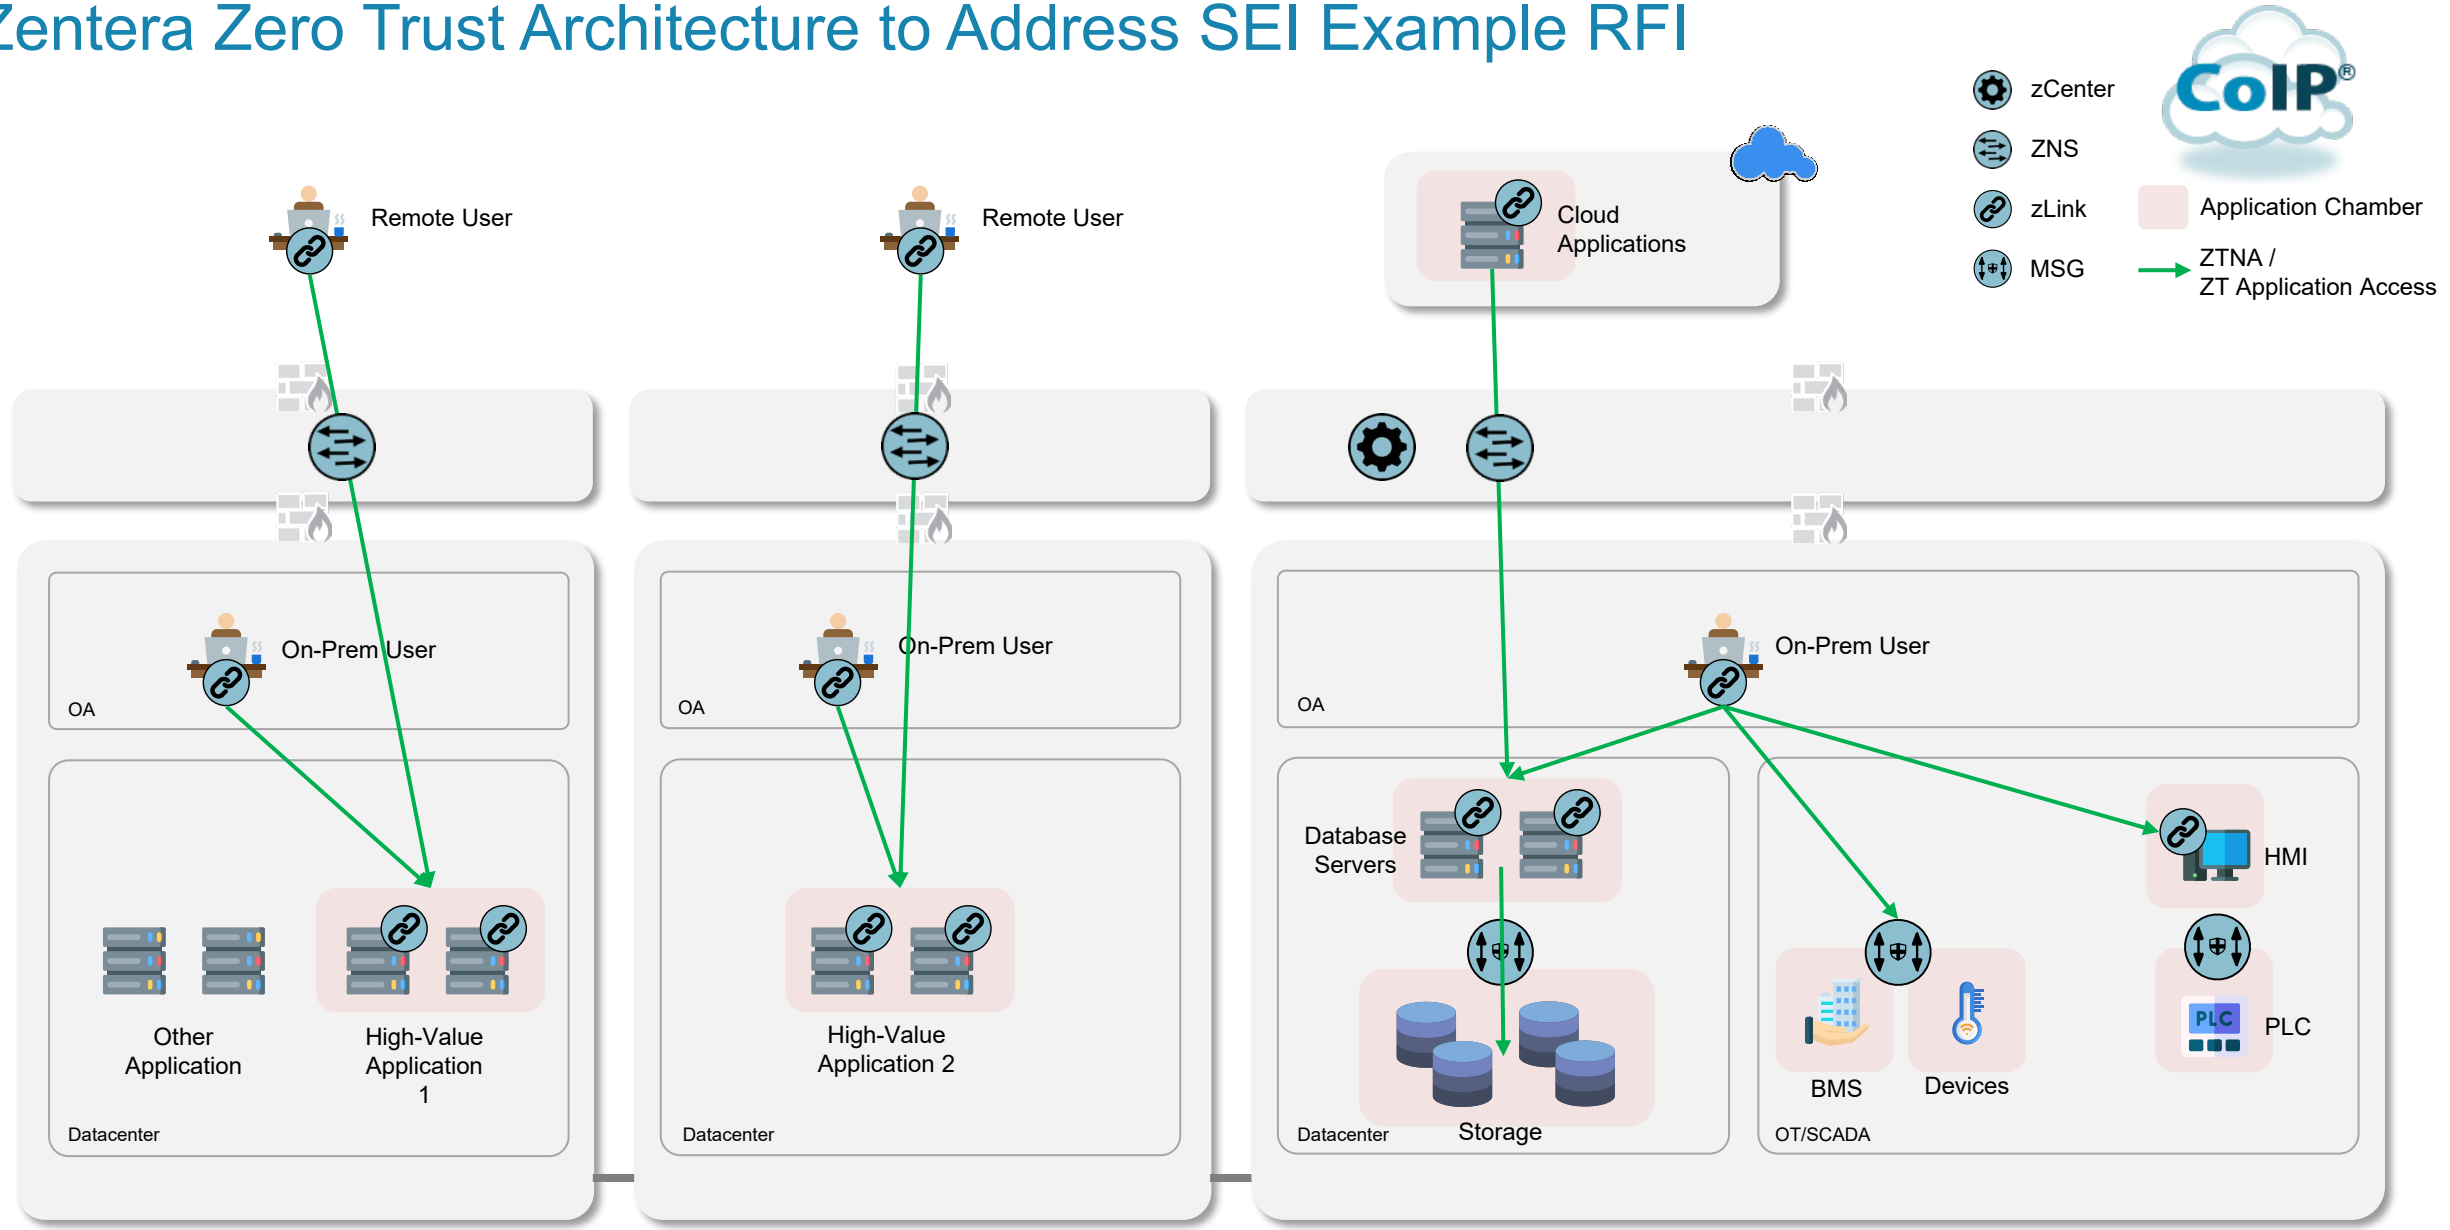

zentera™

Zero Trust Confusion Leads to Operational Pitfalls in Adoption

Ignoring the
Long Tail of Access

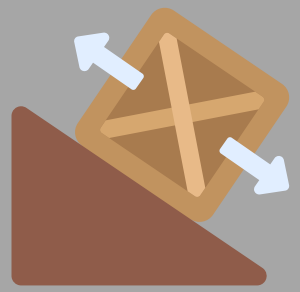

Impacting Production
Applications

Cementing New Infrastructure
Dependencies into Security

Thinking
"Outside-In"

These pitfalls are avoidable when you stay focused on the Zero Trust definition

zentera™

Pitfall #1: Boiling the Ocean

The Pitfall: Thinking that Zero Trust is an “all or nothing” proposition

The Problem: “Rip and replace” impacts existing security and compliance practices and creates organizational churn

Encourages corner-cutting (trading depth for breadth)

Examples: Aggressively decommissioning the VPN

The Reality: There *will* be a VPN until *all* applications are migrated / policies written (or your ZTNA has to *be* the VPN)

How to Avoid: Consider a gradual adoption model where each application is completely secured, even if temporary coexistence with VPN is necessary

An overnight transition is not reasonable; M-22-09 recommends starting with one FISMA Moderate application as an exercise

Pitfall #2: Ignoring the Long Tail of Access

The Pitfall: Thinking that Zero Trust only applies to the the main application data path

The Problem: Applications have lots of critical dependencies which leave attack surfaces open for lateral access

Examples: Application database backends, DNS, DHCP, NTP, NFS, LDAP, SIEM, performance monitors, ITSM tools, privileged IT admin access, CI/CD pipelines

The Reality: Even these services can become paths for lateral attack or data exfiltration (DNS tunneling)

How to Avoid: Choose a Zero Trust approach that enables controls and policies for *all* services on Day 1 – even if some are not yet fully Zero Trust (better to have policies than not)

As M-22-09 requires agencies to move toward encrypted DNS; having appropriate policies can help demonstrate compliance

Pitfall #3: Introducing Unintended User Friction

The Pitfall: Focusing Zero Trust adoption on solutions for one specific use case, creating inconsistent user experience

The Problem: Encourages users to explore workarounds and potentially creates security gaps

Examples: ZTNA that serves remote access only; users work differently when they come in to the office

The Reality: Users need a consistent access model to reduce cognitive load

How to Avoid: Avoid creating gaps in Zero Trust security

M-22-09 recommends user network access to be progressively de-emphasized;
don't forget this applies to on-prem!

zentera™

Pitfall #4: Impacting Production Applications

The Pitfall: Thinking that adding Zero Trust makes it OK to break a known working application

The Problem: Vendors/devops/network teams may not be available to support required changes, and management/users can't afford disruption. Things that used to work suddenly break. Compliance has to be re-done.

Examples:
Network segmentation requiring application re-IP
Adding native identity-checks into the application
Rewriting the application for LAN encryption
Security enforcement that forces application traffic to transit a SaaS cloud
Solutions that don't work with the existing WAF

The Reality: Significant changes in application behavior for Zero Trust can uncover hidden dependencies and architecture assumptions. Zero Trust adoption should not require firewall teams to open ports or shutting down existing security (e.g. WAF)

How to Avoid: Prioritize solutions that can decouple security from networking and minimize impact to existing applications

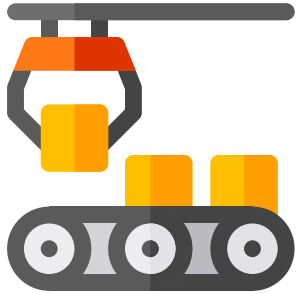

Consider how to implement application segmentation / isolation and full network encryption without breaking your application

Pitfall #5: Cementing New Infrastructure Dependencies Into Security

The Pitfall: Confusion that Zero Trust needs to be implemented by “programming” existing network infrastructure

The Problem: Increases lock-in; makes it hard to upgrade equipment, change technology stacks (multi-cloud), or even change security policies as applications move

Examples: Implementing infrastructure-specific segmentation (VLAN, ACL, SDN) that is limited to IT, but doesn’t work in all environments, OT, or cloud

The Reality: To be executable, scalable and even future-proof, security needs to *decouple from networking*

How to Avoid: Leverage a Zero Trust solution that can deploy as an inline overlay, delivering security enforcement in the OS or as close as possible to it

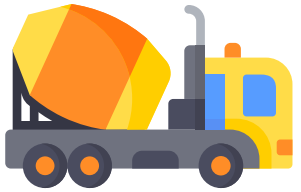

When done properly, Zero Trust should be an *enabler* for infrastructure innovation (e.g. multi-cloud)

Pitfall #6: Focusing on Zero Trust “Outside-In”

The Pitfall: Confusion about Zero Trust can lead to hyperfocus on protecting *users* instead of protecting *applications and data*

The Problem: You didn’t get on the news just because a user clicked on a phish… you’re on the news because your customer PII is on the dark web!

There are many other ways in… USB drives, SneakerNet, supply chain attacjs,,,

Examples: Urgent ”upgrades” that protect users - swapping EDR vendors, or deploying a SWG

The Reality: Zero Trust adoption needs to be focused on protecting things that matter – high value assets

How to Avoid: Start ”inside out” – FIRST answer the question: how should I protect these applications and data, and THEN decide how you will provide access

Establish and fund long-term efforts to achieve a complete Zero Trust makeover in parallel with day-to-day firefighting

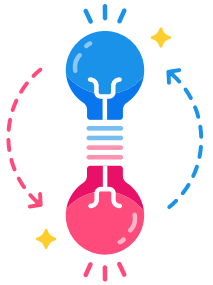

M-22-09 requires application segmentation / isolation and full network encryption

Conclusion

- There's a lot of confusion around Zero Trust today
  - Definition is straightforward, but applying it requires serious discernment!

- To succeed with Zero Trust
  - Be clear about what you are trying to achieve – avoid the pitfalls!
  - Consider security as an overlay to de-risk the project
  - Bite the bullet and start developing application policies as completely as possible
  - Even if you need to switch to another solution later, you know what your policies are!

Thank You!

Mike Ichiriu

VP of Marketing and Product

michiriu@zentera.net

www.zentera.net | info@zentera.net | 408-436-4811

Introduction

This document has been prepared by Zentera Systems in response to the Software Engineering Institute's Request for Information regarding Zero Trust, in support of SEI Zero Trust Industry Day and to provide information and guidance to help SEI analysts consider approaches and methodologies to recommend to Federal agencies and industry.

Zentera's responses herein should be considered for reference only, and should in no way be construed as providing commercial detail that can be used in a purchasing decision (e.g. information on quotation and delivery).

Please direct questions or correspondence regarding this document to Mike Ichiriu, VP of Marketing and Product (michiriu@zentera.net).

Key Challenges Posed by RFI, and Assumptions Made in this Response

The RFI asks the respondent to consider Zero Trust implementation in the position of a large agency, with a worldwide distributed workforce, multiple sites, and legacy systems. Memoranda from the Office of Management and Budget, such as M-22-09 and M-21-31, are key driving documents to help guide the agency in its effort to comply with the President's Executive Order 14028 on Improving the Nation's Cybersecurity.

The agency has 45,000 employees and 15,000 contractors. 75% of the agency's networks are IT, while the balance is made up by OT (building management, etc. – 15%) and SCADA (10%).

The three High Value Assets (HVA) to consider in the response are two legacy applications and one database that contains Protected Personal Information (PPI).

Assumptions Made by Zentera

To help guide the analysis, Zentera has made the following assumptions:

- The 45,000 employees and 15,000 contractors work out of 20 worldwide sites, although they may be working from other remote locations
- All 60,000 total users have a need to access the three HVAs
- At most 50% of the users would need to concurrently access the HVAs
- The two legacy application HVAs run on server farms, each of which has 150 servers and integrated storage
- The database HVA has a clustered database front-end consisting of 10 servers, with data stored across 12 storage appliances (filers)
- The agency has a total of 3,000 application servers that need to be protected with Zero Trust

Proposed Solution

Zentera's proposed solution using CoIP® Platform is outlined in the following diagram and explained in detail below.

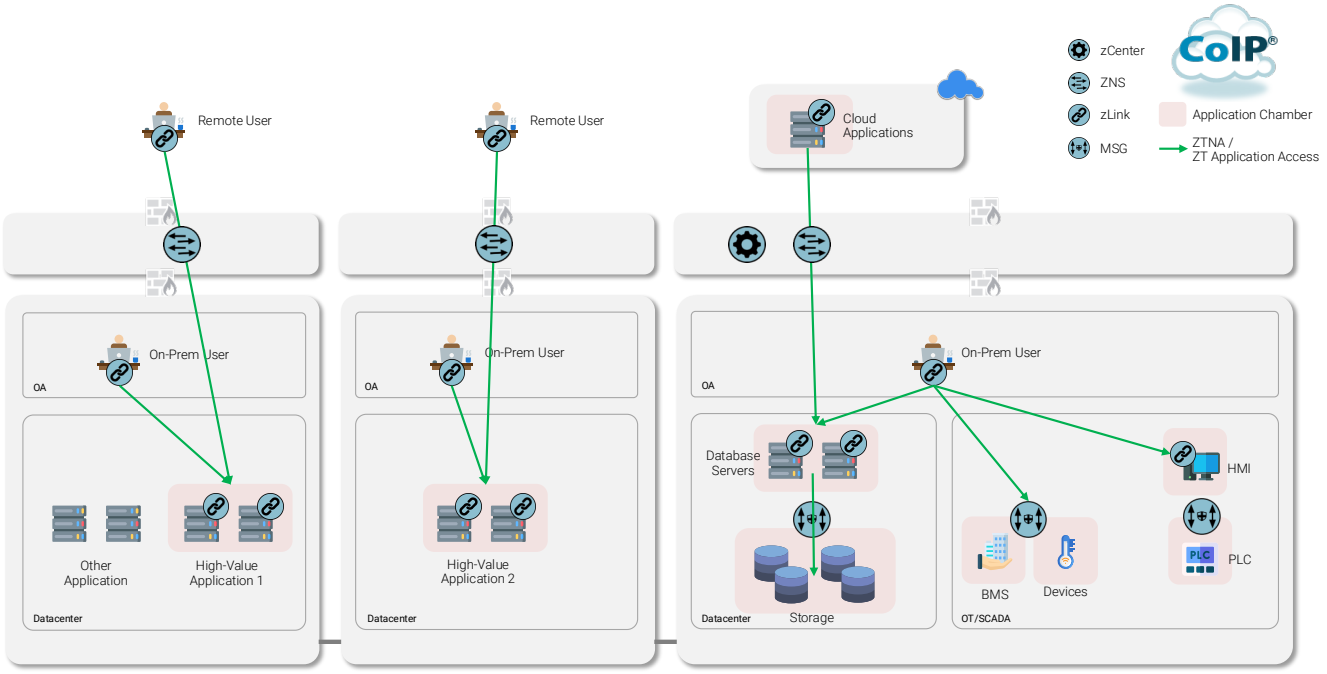

Overview of Solution Components

zCenter Orchestrator

The zCenter Orchestrator is the core of CoIP Platform; in NIST 800-207 parlance, it plays the role of the Policy Decision Point. It is aware of and authenticates all users, endpoints, and applications, and maintains the policies that define authorization.

User authentication is typically performed using an external Identity Provider (IdP, not shown). zCenter can work as a unified authentication enforcement system supporting multiple Identity Providers; supported authentication schemes include SAML 2.0, OAuth 2.0, OpenID Connect, or LDAP. Customers typically configure the IdP to perform multi-factor authentication (MFA), so the authentication flow is familiar to users.

Endpoint authentication is performed when endpoints register to zCenter. Endpoints running the zLink agent provide telemetry about the endpoint, including various hardware and software identifiers. These identity factors are pinned together with a unique identifier for the agent; subsequent changes in the identity factors can cause zCenter to deny the registration or even quarantine the endpoint for remediation. These identity factors may include the security posture, which can be included in the access policy to determine whether or not to grant access.

zCenter also collects information about the endpoint location; this can be used as part of the policy decision. The endpoint location is gathered from through IP geolocation based on the endpoint's observed public IP address. Endpoint location information is also gathered from cloud instance metadata that are provided by cloud providers.

Application authentication is configured in zCenter and enforced by zLink agents. Application authentication includes various identity factors such as checksum and path. Policies may incorporate the process tree (parent/child relationships between processes) to make context-based decisions about whether to allow communication for a particular application.

zCenter is natively multi-tenant and supports scoped access for service administrators, customer tenant administrators, project administrators, and read-only administrators. It also supports access through APIs, with key types and APIs that match administrator scoping.

The zCenter Orchestrator may be deployed on-premises as a virtual appliance, instantiated in a public cloud through cloud marketplaces and supports high availability and disaster recovery. Alternately, a zCenter Orchestrator may be purchased as-a-Service (Zentera Air). The above diagram shows one zCenter appliance, but in practice multiple would be used to address HA and DR requirements.

ZNS Network Switch

The ZNS Network Switch is a virtual appliance that provides data plane services for CoIP Platform. The ZNS bridges the two ends of a CoIP WAN connection and is typically used to bridge two disconnected network domains (e.g., serving remote access traffic). An individual ZNS node can support over 20Gbps of aggregated WAN traffic and can be clustered for redundancy and to scale out the performance to meet application requirements.

The ZNS does not expose any applications or network by default, and instead opens paths for application access selectively when directed by zCenter to do so. The ZNS also maintains separation for multi-tenant traffic.

ZNS appliances may be deployed on-premises as virtual appliances or instantiated in a public cloud through cloud marketplaces. ZNS works with Zentera Air SaaS, allowing customers to choose the simplicity of a SaaS control plane and work with a private on-prem data plane.

zLink Agent

The zLink Agent is the primary onboarding model for endpoints; in NIST 800-207 parlance, zLink serves as the Policy Enforcement Point.

zLink installs as user-mode software on practically any endpoint – Windows (Win XP and up), Linux (kernel 2.6.32 and up), or Mac (High Sierra and up). zLink supports the Application Chamber (discussed below) that enforces Zero Trust policies; it also intercepts network packets and compares them against a local subset of access policies that are relevant to the zLink. If additional checks or a connection setup is required, zLink submits the relevant information to zCenter for a decision and assistance in brokering the connection.

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

CoIP Launcher

CoIP Launcher is an agentless solution that provides ZTNA connectivity for users. From the user perspective, it operates much like a standard Web meeting product – browsing to a URL launches client software on the user's device. The user then interacts with the services through an in-browser User Portal. CoIP Launcher also

performs some enforcement functions, and can be categorized as a Policy Enforcement Point according to NIST 800-207.

CoIP Launcher supports all of the different access types. These include a basic "network mode", which enables the user to use standard TCP/IP clients (TCP, UDP, ICMP traffic types); several other built-in access types are supported (VNC, RDP, Secure Shell) with enhanced security protections (copy/paste controls, sftp/scp/ssh tunnel blocking, etc).

Presently, CoIP Launcher is supported on Windows, Mac, and Linux platforms.

Micro-Segmentation Gatekeeper

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 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.

Functional Overview

Application Chambers

An Application Chamber is a logical construct that defines a group of servers and enforces security policy controls. For example, applied to the standard 3-tier Web application model, the servers in the presentation tier (web servers) share identical access controls, as do the servers in the application and data tiers, and it is natural to create 3 Application Chambers – one for each tier.

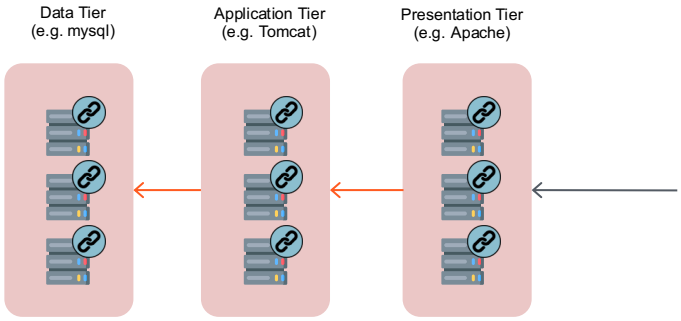

The Application Chamber is Zero Trust aware and can be used to filter access based the identity of users or endpoints. For example, an access policy may grant a database administrator access to the Data Tier; regardless of where the database administrator is located, the Application Chamber will grant access (subject, of course, to the other attributes specified in the policy). Similarly, endpoints in the Presentation Tier may make outbound connections to endpoints in the Application Tier; this policy will be applied consistently even as endpoints dynamically join or leave a tier (as service is scaled to meet a surge in traffic, for example).

When Application Chamber policy controls are turned on, the servers in the Chamber are cloaked from the shared local network, and access is allowed only based on Zero Trust policies that the administrator has

defined. For example, the admin might initially create an access policy that restricts DNS on TCP/UDP port 53 to a known set of DNS resolvers. When ready, the admin can change the policy to actively enforce the use of encrypted DNS over HTTPS or TLS – relevant policy changes are immediately pushed down to all PEPs.

Zero Trust Network Access and Application Access

CoIP Platform supports several modes of operation for ZTNA and application access.

CoIP WAN accesses cross network domains (remote access). All traffic is carried in TLS 1.3 tunnels, with multiple outbound tunnels bridged by the ZNS node under the direction of zCenter.

CoIP LAN accesses secure traffic between endpoints in a single network domain. Traffic is tunneled (and optionally encrypted) point-to-point, with the connection between the endpoints brokered by zCenter. CoIP LAN can be a good option for legacy unencrypted traffic.

Chamber-only accesses enforce security controls over existing application or network traffic, but do not otherwise tunnel or encrypt traffic between endpoints. Chamber-only can be a good option to avoid double encryption for traffic that is already encrypted.

How the Proposed Solution Addresses Zero Trust Requirements

Identity

Users

The proposal allows the agency to use its existing Identity Providers. CoIP Platform is capable of handling multiple Identity Providers, so consolidation is not a pre-requisite for beginning the Zero Trust journey.

Endpoints and Applications

CoIP Platform allows the agency to begin onboarding the HVA, based on pre-existing knowledge of which endpoints should be covered. CoIP Platform will collect identifying information about the endpoints which will be pinned for subsequent registrations.

The administrator may use CoIP Platform's Learn feature to identify any dependencies on network services, endpoints, or applications. The Learn feature presents the administrator with a record of network activity along with helpful information about the context of the access, including the application process owner, command and arguments which resulted in the network activity. This gives the admin the option to create a policy that allows or blocks that access. This helps the administrator develop a complete list of application policies or identify other endpoints that should be onboarded to CoIP Platform for Zero Trust enforcement.

Networks

The Application Chamber creates a new Zero Trust perimeter to enforce identity-based access. The access types enable the administrator to turn on Zero Trust policies for selected applications or to choose encrypted tunnels for some application traffic, satisfying the requirement in M-22-09 to encrypt local network traffic.

Data

CoIP Platform provides a variety of options to help secure access to the data in place.

Securing Data on Endpoints

For data stored on endpoints (data in application servers), the Application Chamber provides an effective protection against lateral access to the target, as well as preventing exfiltration. Several Zentera customers use CoIP Platform precisely for data protection; the Chamber is used to create a secure environment for users to access sensitive company intellectual property (IP). The Chamber is configured to prevent leakage (including blocking the use of Internet DNS servers to avoid DNS tunnel detection). User access is limited to an access method (e.g. Windows RDP) that is secured with Zero Trust and has additional controls (e.g. copy/paste blocking).

CoIP Platform also provides the capability to automatically mount NFS volumes by policy; when an authorized user initiates a ZTNA session to a designated endpoint, the zLink Agent will mount the appropriate folder(s) for security control. When the user session ends, the folder(s) are automatically unmounted.

Securing Data in the Network

The MSG provides a useful way to create Zero Trust protection for network-attached storage or filers. Deployed inline with the filer, the MSG enables the administrator to grant access to the filer to authorized users and endpoints based on Zero Trust policies. All other network access to the filer is blocked by default.

How the Proposed Solution Addresses Operational and Implementation Concerns

Zentera's unique model allows the agency to rapidly deploy Zero Trust controls throughout the existing network without disruption. This overlay technology and capability is critical to minimizing operational costs.

Managing the Transition to Zero Trust

One of the overarching goals driving Zentera's development of CoIP Platform has been to decouple security from the network. We have achieved that goal by presenting the Zero Trust security layer to applications *as a network stack* and presenting the Zero Trust security layer to networks *as an application stack*. Our onboarding models, the zLink Agent and MSG, deploy inline with the protected user, endpoint, or application.

This minimizes the need to change the network or application architecture to adopt Zero Trust. This is critical because it gives the agency agility in the Zero Trust implementation, and does not tie Zero Trust adoption to a network change or upgrade.

The agency may onboard the application to Zero Trust *in place*, without yet turning on Zero Trust enforcement. The administrator can run with the Application Chamber in detection mode to view policy violations without blocking them. Once satisfied with the results, the administrator can switch the Chamber into prevention mode to begin enforcing Zero Trust.

Mixed Infrastructure and Vendor Environments

CoIP Platform deploys as an overlay, completely independent of the infrastructure and vendor stack. CoIP Platform has been used to deliver Zero Trust security in every major cloud service provider environment, every major datacenter virtualization environment, on bare metal servers, and with many combinations of networking and network security tools from different vendors.

IT/OT

CoIP Platform's unique overlay deployment enables Zero Trust to be deployed to OT environments, as well as hybrid environments including IT, OT, and cloud. While not part of the RFI, we included Zero Trust protection for the BMS in OT and the SCADA network in the diagram above.

Roadmap and Implementation Plan

Near-Term and Long-Term Roadmaps

From the point a decision is made, deploying Zero Trust security for the 3 HVAs can be achieved in as little as 2-3 months from initial pilot to full production rollout. We have assisted some customers with small projects roll out Zero Trust for user access in a matter of days.

Based on our experience with rolling out Zero Trust protection, we recommend a phased approach that focuses on onboarding applications in order of priority. The highest value assets may be onboarded in the first phase,

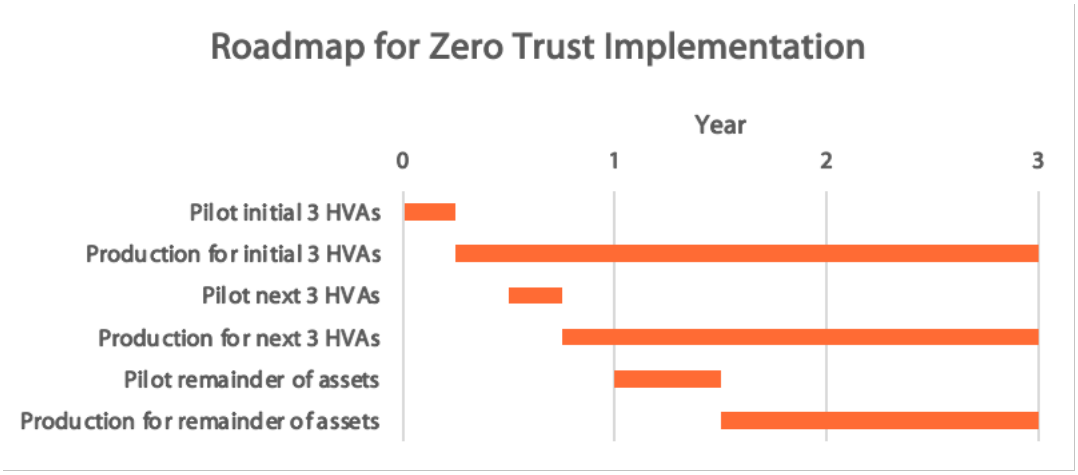

Implementation Plan

The steps below outline a potential Implementation Plan for Zero Trust Security for the three HVAs.

1. Deployment and Onboarding
    a. Deploy and configure zCenter and ZNS, including connections to existing Identity Providers
    b. Define Application Chamber for HVA #1
    c. Onboard HVA #1 to the Application Chamber with Zero Trust security disabled
2. Policy Creation and Refinement
    a. Configure HVA #1 access policy set to the best of the administrator's ability
    b. Run Learn to identify and eliminate gaps in the policy set
3. Pilot
    a. Set the Application Chamber for HVA #1 to detection mode
    b. Onboard test users chosen to cover a wide range of user scenarios and collect feedback
    c. Make policy adjustments as necessary based on detection results or user feedback
4. Production
    a. Roll out access to all users
    b. Set the Application Chamber for HVA #1 to protection mode
5. Monitor
    a. Log and monitor violation reports

This process (steps 1b – 5) may then be repeated for HVA #2 and #3.

The administrator should periodically review the policy set to ensure policies are still accurate and relevant. The administrator may choose to tighten policies over time or may need to adjust policies to enforce new requirements (e.g., enforcing the use of encrypted DNS).

User Impact

User Experience

Users typically find Zentera CoIP Launcher easier to use than VPN. Administrators distribute a URL to users, who log in with the corporate identity provider.

Once connected, they are presented with a User Portal that allows the user to see only the resources they have access to. The administrator can configure a launch action in the Portal that launches the appropriate client via a URI scheme. This streamlines access and makes it even easier for users to access internal resources; with a VPN, users are on the shared network and have to remember what action to take after signing in.

Training

In our experience onboarding customers with tens of thousands of users, minimal user training has been necessary. The user experience is natural and easier to use than VPN. Customers have adapted our user guides to create their own internal documentation and guides to support users.

IT Support

It is simple for existing helpdesk personnel to handle L1 support for Zentera Zero Trust. Our customers have typically assigned responsibility for L2 support to a small team of 1-2 who manage and direct the L1 team, and coordinate with Zentera factory technical support as L3.

Cost Considerations

Procurement

Implementation of Zentera CoIP Platform would require the agency to procure the core infrastructure to support onboarding applications and users; this includes the zCenter and ZNS.

Users and servers are both independently licensed. Zentera counts concurrent license usage, rather than named users. As described above,

Support and Maintenance

Our licensing model is based on annual subscription, and includes all product maintenance, updates, and upgrades. Standard support is included with the subscription, which includes web, email, and telephone support during business hours. Premium support packages are available, which provide access to dedicated technical support personnel for 24x7 response and issue escalation.

It is reasonable to use 10% of the ACV as a guideline for annual support cost.

Budgetary Costs

Budgetary costs for the rollout described above are included below, and fit within the agency's $3M implementation budget. These numbers are rough guidelines based on assumptions we have made, and do not constitute a binding quote.

	Scope	Solution Costs	User Licenses	Application Costs	Support	Annual Cost
Year 1	Initial HVAs	$ 250,000	$ 200,000	$ 40,000	$ 55,000	**$ 545,000**
Years 2 and 3	All agency applications	$ 250,000	$ 500,000	$ 200,000	$ 100,000	**$1,050,000**
					3-Year Cost	**$2,645,000**

Solution Costs refers to the software licenses required to implement the CoIP Platform base infrastructure. We recommend to implement the full infrastructure that will be needed for full production up front, to streamline onboarding applications to Zero Trust.

User Licenses refer to CoIP Launcher licenses needed to support user access. Zentera supports a concurrent licensing model, so we have used concurrency estimates as outlined in the Assumptions to minimize license cost.

Application Costs refers to licenses needed to support onboarding individual applications to Zero Trust – the zLink Agent and the MSG.

Support refers to premium support packages offered by Zentera, as described above.

Staffing Plan

Following the support model outlined above, staffing requirements are quite modest.

In the initial pilot phase, it is important to include security and network architects who can provide input about site- or application-specific idiosyncrasies to the team. As these projects ramp into production, they may transition to other efforts.

A dedicated project manager can help in planning and communication surrounding the agency's complete transition to Zero Trust, so we have modeled the project manager joining the team after the initial production launch of the first 3 HVAs.

Most important are the IT administrators who will be responsible for maintaining the system, pushing updates to endpoints (using automation tools such as Ansible as needed), and handling L2 support.

As described above, we anticipate that L1 support will be handled by the agency's existing IT helpdesk.

	Year 1				Year 2				Year 3			
	Q1	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**
Project Managers	-	-	1	1	1	1	1	1	-	-	-	-
Network Architects	1	1	1	-	-	-	-	-	-	-	-	-
Security Architects	1	1	1	1	1	1	1	1	-	-	-	-
IT Administrators	1	1	2	2	2	2	2	2	2	2	2	2
Project Headcount	**3**	**3**	**5**	**4**	**4**	**4**	**4**	**4**	**2**	**2**	**2**	**2**

Scalability Considerations

Impact of scaling down to small or medium-sized agency

Zentera's model scales linearly to small or medium-sized agencies. Core Zero Trust infrastructure requirements are unchanged; fewer ZNS nodes may be needed to support a less complex infrastructure (fewer sites and clouds). Licensing and support costs will scale with the reduced number of users, endpoints, and applications that are being covered.

Application Chambers

Cybersecurity Defense at Application-Scale

zentera™

"The Corporate Perimeter is Dead…"

The industry is on the verge what may be the most radical overhaul of enterprise network security tools and practices in history. That's because cloud and hybrid computing and remote work have exposed the fact that the traditional network perimeter security business model – delivered as part of the network infrastructure – is fundamentally broken.

Think about it. Firewalls are complex to program and maintain; they are intended to handle all corporate traffic at the network perimeter. Yet a firewall can only guarantee that packets flowing through it are secure. It can't secure packets it never sees.

Growing infrastructure complexity and cloud adoption are making it increasingly difficult to distributed traffic through specific security enforcement points in the network.

Many public companies have implemented network zones to protect financial data for years, driven by compliance requirements such as Sarbanes-Oxley (SOX). Typically implemented using VLANs with more firewalls, secure zones are reserved for the most sensitive applications and data. Could this model evolve to solve new security challenges?

But you can't just put firewalls everywhere… right? Clearly cost and effort just won't scale with traditional infrastructure, and agility targets won't be met either – today's applications and users are far too dynamic for infrastructure to keep up with. A more flexible solution is required to deploy alongside the applications and data and move elastically with them.

That future is here today, and it's called an *Application Chamber*.

Infrastructure solutions work well when requirements don't change.

They're not a great choice for things that are more dynamic – like next-gen applications and data.

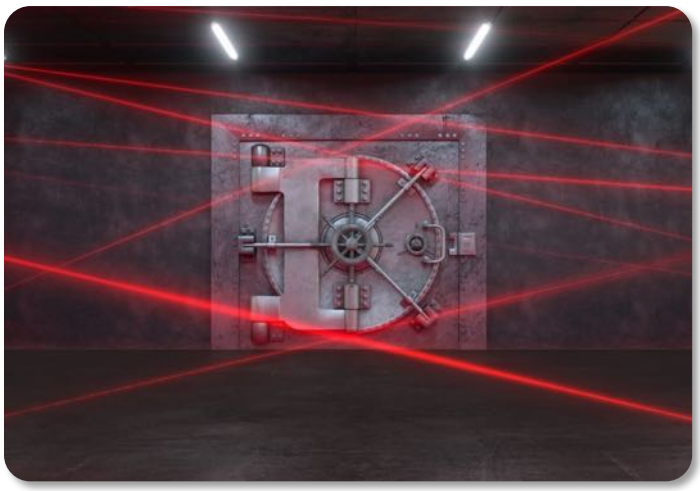

An Application Chamber wraps protections around sensitive applications and data. It is delivered as software, decoupled from underlay infrastructure, and provides multi-layered defenses against malware, insider threats, and ransomware.

"Long Live the Application Chamber (Perimeter)!"

An Application Chamber is a new "application-scale" perimeter that protects applications and data, right at the application edge, by ensuring that all network traffic, both inbound and outbound, undergoes a full security inspection. It deploys as software, enhancing the security of existing application servers without disrupting them or setting up new infrastructure.

Zentera's CoIP® Access Platform allows you to quickly configure Application Chambers to manage groups of servers by their role or function. Application Chambers support access policies that allow you to model tiered or interconnected applications, while automatic learning and easy policy templates allows you to rapidly secure existing applications.

Fully integrated with CoIP Access Platform's Zero Trust Network Access (ZTNA), you can provide on-prem or remote users instant, least-privilege access to specific chambered applications.

And best of all, Application Chambers work everywhere you need them, to protect bare metal and virtualized apps in any environment – on-prem, cloud, or even 3rd party-owned infrastructure.

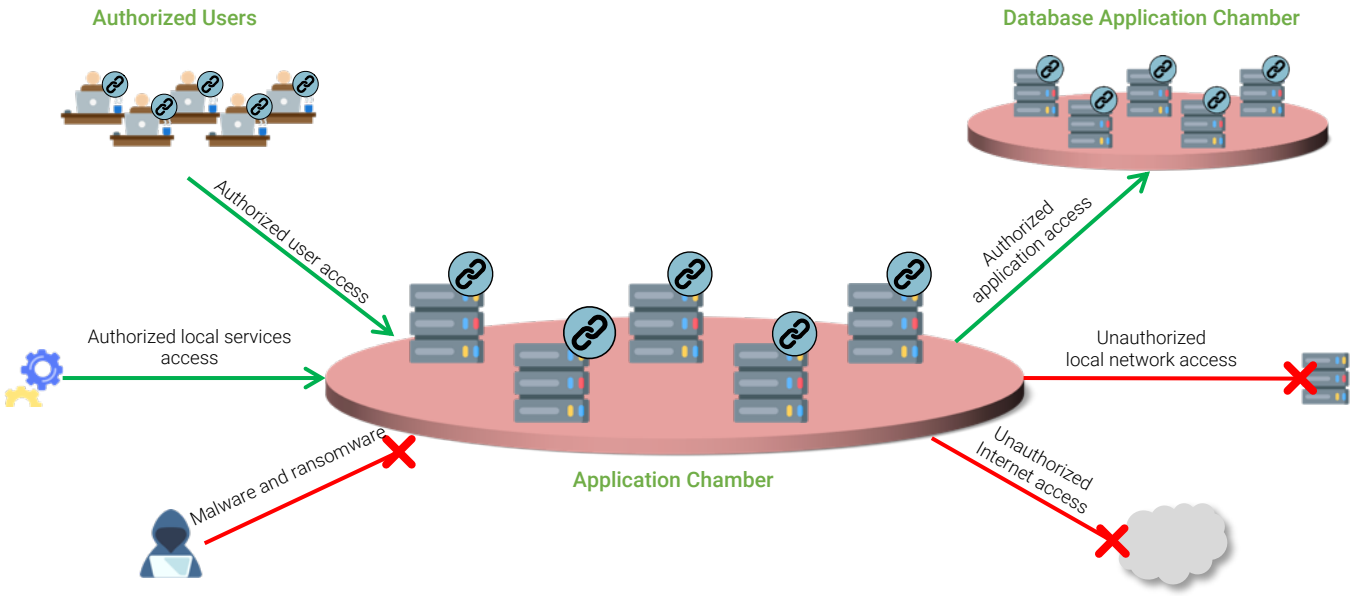

Application Chamber Features

Application Cloaking

Application-scale secure zone, inaccessible on the shared network

Ransomware Defense

Blocks ransomware from accessing vulnerable application servers

Data Leak Prevention

Blocks unapproved access methods, so data stays locked in the Chamber

Zero Trust

Natively integrated with ZTNA to authenticate and authorize each access

Template-Based Policies

Static policies that are enforced dynamically to react to changing workloads

Use Cases for Application Chambers

CoIP Access Platform's Application Chambers are incredibly powerful due to their tight integration with ZTNA. Together, they enable CoIP Access Platform to solve a variety of enterprise security and compliance challenges:

- **Cloaking sensitive applications (ERP, AI/Big Data, etc) to protect them from lateral attacks inside the enterprise**
- **Safeguarding critical data against exfiltration and insider threats**
- **Blocking ransomware and worm activity from reaching protected servers**
- **Enabling remote employees and 3rd parties to applications without exploring the network**
- **Maintaining a single consistent and auditable set of corporate policies across all hybrid environments**

How an Application Chamber Benefits Your Business

CoIP Access Platform enables enterprise to easily implement next-generation Zero Trust Security for application access security and control. The business benefits include:

- **Lower cyber-risk through dramatically reduced attack surface**
- **Promotes adoption of next-generation Zero Trust security, while avoiding "rip and replace" of the existing infrastructure**
- **Reduces employee stress by shifting the burden of maintaining cybersecurity to tools and processes**
- **Dramatically lower TCO and improved ROI compared to traditional infrastructure security, due to agility and consolidation of tools and methods across all hybrid environments**

More Resources

On the web: www.zentera.net

Email: sales@zentera.net

Phone: +1 (408) 436-4811

SOLUTION BRIEF

Secure IT/OT Convergence

Defending OT Operations with Zero Trust Security Segmentation and Controls

May 2022

zentera™

Hackers Are Increasingly Targeting OT Networks

Information technology (IT) and operational technology (OT) networks have traditionally been separate domains. IT networks, which support enterprise computing and applications, prioritize *confidentiality* of the data they contain – financial performance, customer lists, and so on. OT networks, on the other hand, drive revenue production for the business, and therefore prioritize *availability* – any downtime instantly results in lost revenue and may even present a safety hazard.

These details haven't escaped hackers, who have figured out how motivated OT companies are in keeping facilities operational. With ransomware and cryptocurrency making cyber extortion simpler than ever, the surge in reconnaissance activity against OT networks is bad news.

$4.24M
Average Cost of a Breach in Industrial Sector[2]

36%
Share of all OT Attacks from Ransomware[1]

2,204%
Increase in Reconnaissance Against OT, Jan-Sep 2021[1]

Potential Risks of OT Attacks

- Financial losses from production downtime and lost work in progress
- Costs to remediate and recover
- Loss of customer confidence and brand image
- Safety and environmental hazards
- Fines, lawsuits, and liability
- Bad publicity with shareholders and stakeholders
- Non-renewal of insurance coverage

[1]IBM Security, X-Force Threat Intelligence Index 2022
[2]IBM Security, Cost of a Data Breach Report 2021

The Divide Between IT and OT is Shrinking, and OT is More Exposed than Ever

The Purdue Reference Model for OT segmentation and security controls to protect mission-critical production assets was first proposed in1992. While it provides a useful logical model, it is often difficult to implement using traditional, inflexible networking technologies. As a result, many OT environments are flat and even mixed with IT today. This weak cyber defense for OT is showing signs of age in today's interconnected world: IBM Security's X-Force Threat Intelligence Index 2022 reports that manufacturing was the top attacked industry in 2021[1].

It's not that the production assets have changed – in fact, most OT assets and networks are designed to last for decades. But as companies embrace Digital Transformation – adopting work from home policies and migrating business intelligence into the cloud – the expanded attack surface brings this decades-old technology face-to-face with today's modern threats.

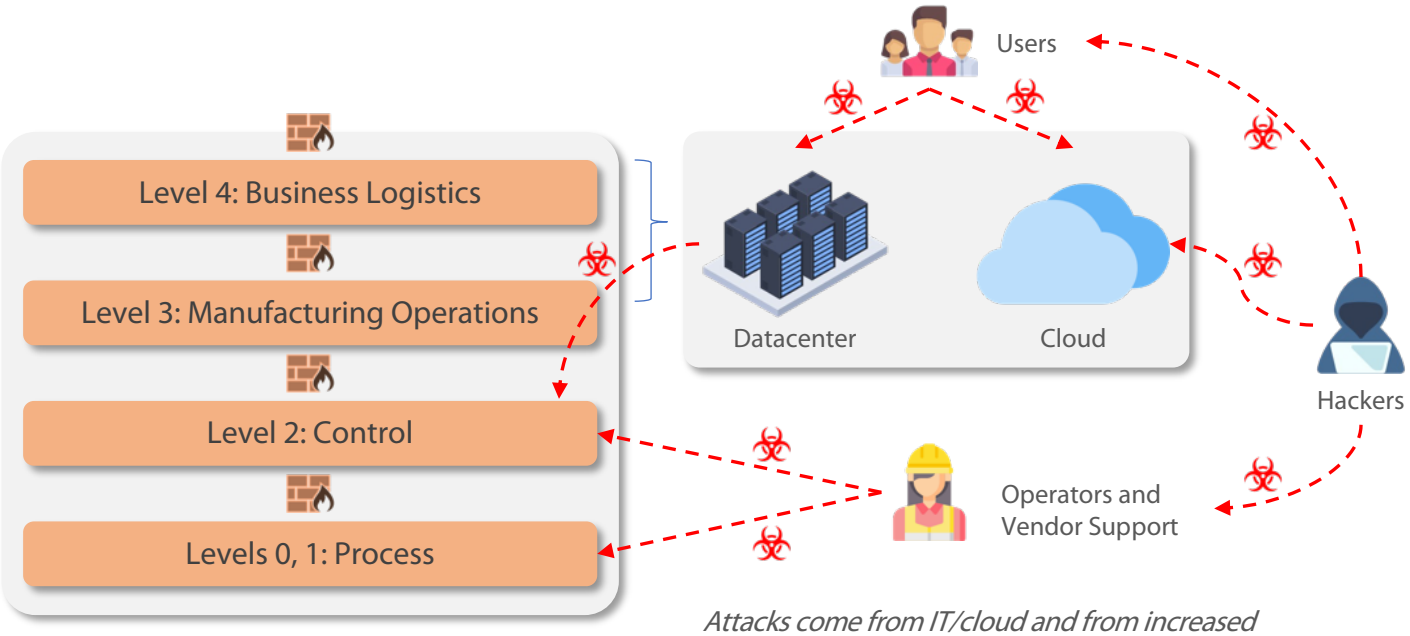

IT/OT with Purdue Reference Model Segmentation

Attacks come from IT/cloud and from increased application access. Digital Transformation gives hackers new options to attack critical production workloads through the expansion of Level 3/4 infrastructure, as well as through compromised remote users and 3rd parties

[1]https://www.ibm.com/downloads/cas/ADLMYLAZ
[2]https://www.ibm.com/downloads/cas/OJDVQGRY

Why Traditional OT Segmentation Has Fallen Short

The fundamental concept behind the Purdue Reference Model is sound. It makes complete sense to group OT assets into different levels based on application and criticality and control access to them. The problem is not in the concept – it's in the execution.

Network segmentation has been traditionally implemented by creating separate zones (e.g. VLAN) connected by a firewall or ACL as control points to control all cross-zone communications for devices. This approach creates many operational challenges:

Coarse Segmentation
Firewalls can be prohibitively expensive. Architects may choose to reduce cost and management overhead by implementing large zones, but this makes it difficult to specify access to specific devices or workloads in the OT environment. Detailed zones using VLAN/ACL create other challenges for policy implementation and management.

Lack of Identity-Based Controls
With traditional filtering, an unauthorized user can bypass the filters simply by launching the attack from any authorized server – potentially one in the cloud or inside the OT environment.

Inflexibility
With traditional segmentation, a device's zone is tied to the network topology. Policy updates mean network reconfiguration that can trigger downtime and so are discouraged, making it hard to keep security filtering up-to-date with the needs of the business.

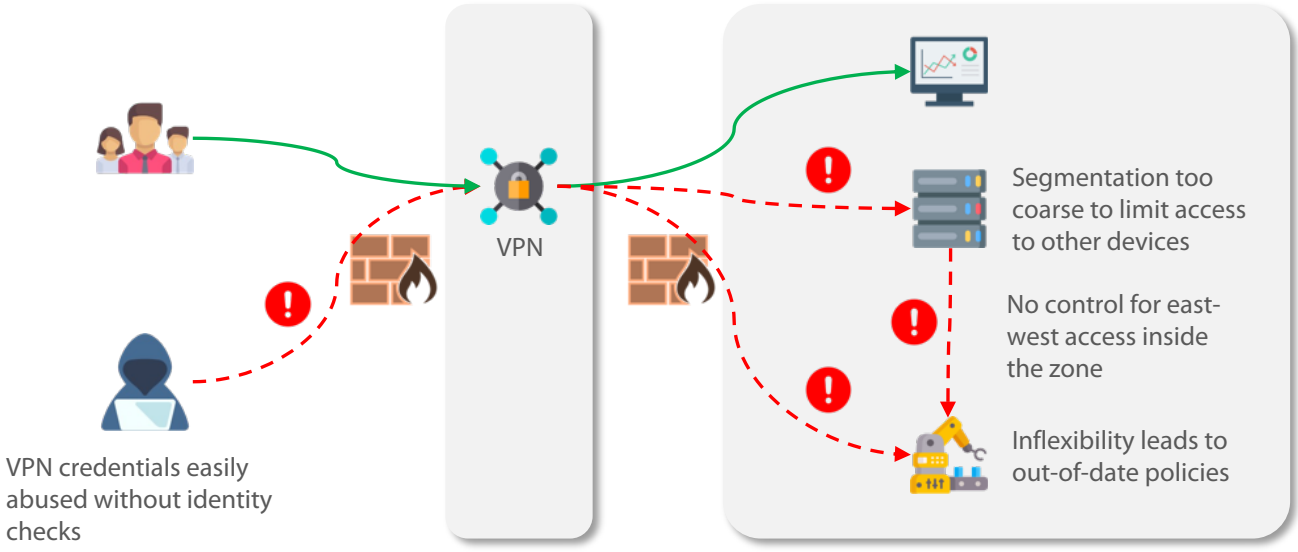

How Zentera Zero Trust Security Defends Converged IT/OT

Zentera Systems' CoIP® Access Platform is an advanced cybersecurity solution capable of layering application-scale segmentation and controls over existing IT and OT networks. Companies may use CoIP Access Platform as the basis of their Zero Trust plans in accordance with NIST 800-207 guidelines.

Once onboarded to CoIP Access Platform, IT and OT assets can be immediately assigned to logical zones called Application Chambers. Traffic flow into, out of, and between Chambers is controlled by identity-based policies that are easy to program and change. All policies are defined in a centralized orchestrator and enforced at users and applications for tight security control. The Zero Trust implementation completely overlays existing networking and firewall architectures without disruption.

Strong Zero Trust Security

Chamber and Access policies implement NIST 800-207 Zero Trust Security based on user, server, and application identity – not on IP address – and are applied properly even as users move and servers migrate across network environments

Workload Micro-Segmentation

Application Chambers provide granular segmentation, cloaking, and isolation of individual OT processes to cover the Detect and Protect categories of the NIST Cybersecurity Framework

Zero Trust Network Access

Least-privilege, application-aware, and VPN-free access with users, vendors, and contractors authenticated against your IdP

Overlay Application Network

Secure and application-aware direct access for cloud-based business applications and intelligence to Chambered OT assets without exposing the OT network to the cloud

Zero Touch Deployment

With onboarding options ranging from the Micro-Segmentation Gatekeeper for agentless protection of OT workloads to the powerful zLink agent for computers and servers, there's no need to disrupt or change existing applications or networks to deploy

Simple Operations and Change Management

Moving segmentation policies and flow control to CoIP Access Platform makes it simple and easy to onboard new users and devices or change policies without impacting already enforced policies

Apply Modern Segmentation and Controls to Any IT/OT Network

CoIP Access Platform enables OT administrators to create segmentation in any environment without VLANs, ACLs, or other network changes. Its combination of Application Chambers, ZTNA, and Overlay Application Network enable Zero Trust Security segmentation and controls for any converged IT/OT environment – even flat networks without any existing controls.

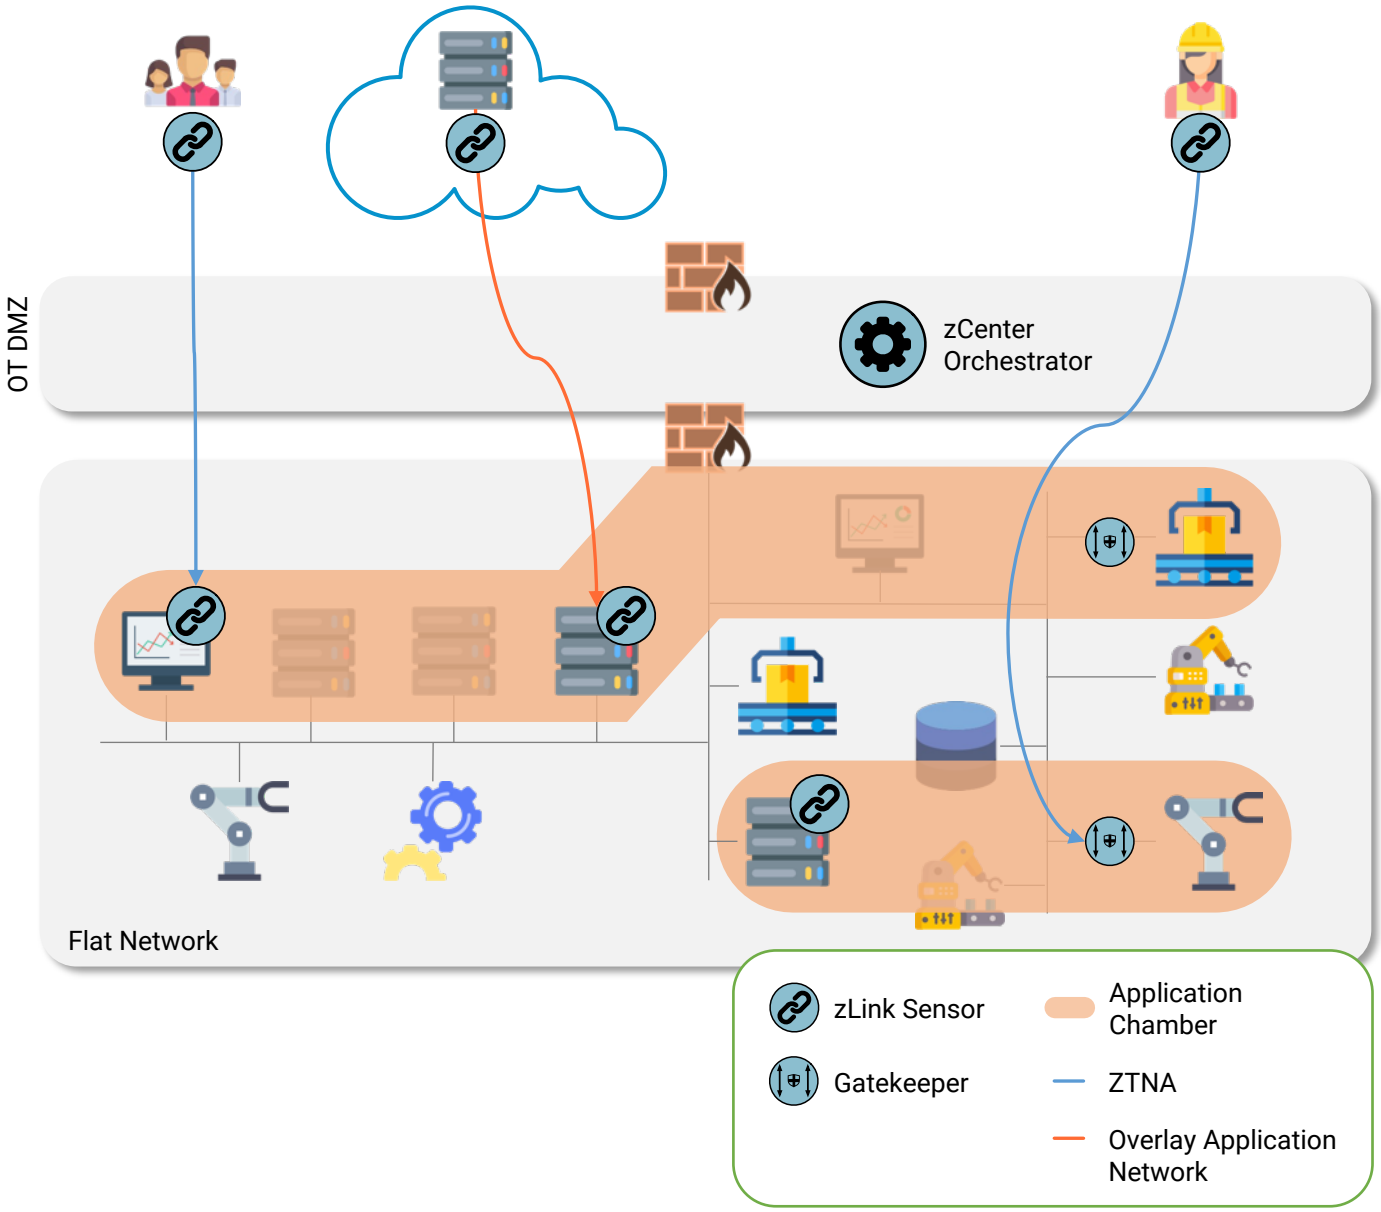

Advanced, Yet Simple to Adopt and Use

Adopting leading-edge network security solutions for secure developer access doesn't have to be complex.

All it takes to get up and running is a simple 3-step process. Once in place, CoIP Access Platform makes it easy to build powerful Zero Trust security policies to improve your security posture.

- Install CoIP Access Platform
- Onboard servers with zLink agents and devices with Micro-Segmentation Gatekeeper

- Create Chambers
- Configure Chamber and access policies
- Turn on policies in Detection mode
- Create user roles and onboard users

- Enforce policies in Prevention mode
- Monitor logs
- Optimize Chamber assignments and policies as needed

Secure IT/OT Convergence

Defending OT Operations with Zero Trust Security Segmentation and Controls

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™

The Future of SASE is Overlay

History Repeats Itself –
Why Next-Gen Security Will Be Over-the-Top

March 2021

zentera™

2020: The Year Everything Changed

While much has already been said about the impact of the global coronavirus pandemic, its impact on information technology (IT) cannot be understated. Businesses that had never had an online presence, such as restaurants, suddenly needed one just to stay alive, while cash businesses were forced to begin accepting electronic payments. At the other end of the spectrum, multibillion-dollar enterprises, faced with office closures lasting many months, had to digitize all business processes so that the enterprise could be piloted remotely and elastically. Digital transformation, previously a mere competitive advantage or an operational upgrade, was suddenly promoted to the level of existential need.

When it comes to benefiting from the pandemic, two IT trends stood out from all others: remote work and cloud migration. While both trends have been steadily growing for years, the coronavirus pandemic obliterated any last organizational resistance to change, realigning attitudes and behaviors overnight.

Two Major Trends, One Overarching Theme

On the fact of it, remote work and cloud migration appear quite distinct. One centers around employee productivity and workflows, to enable productive away from the office; the other is the battle to make the next-generation corporate datacenter elastic, replacing lumpy capital expenditures based on estimates of future computing and infrastructure capacity needs with operational expenses aligned with actual business consumption.

User, app, and data mobility is the corporate megatrend for the next decade

At first glance, it seems these trends have little to do with each other. However, upon closer inspection, it becomes clear that they share a common theme: they enable *mobility* of applications, data, and users in a new, agile world, where offices and datacenters are intentionally de-emphasized. This sea change in how companies view corporate-owned and controlled infrastructure has created huge stress on corporate IT and Infosec teams.

The writing is on the wall, but it will take time for this trend to play out, as delivery methods and security undergo successive waves of optimization.

The Mobility "Virtuous Cycle" Drives the Need for Elastic Security

While the idea of a fully-virtual office still seemed distant only a year or two ago, it has been realized in record time out of necessity. When the pandemic hit, simply companies responded by buying more licenses for virtual private networks (VPNs). But the corporate VPN and firewall were not designed to meet the security challenges introduced by many thousands of persistent remote connections to the corporate network infrastructure, accessing a wide range of corporate applications during the workday. VPNs and firewalls have become the bottleneck for not only user productivity, but also corporate security.

This need of supporting application access on a global scale has triggered many companies to re-think their traditional corporate datacenter models. Enterprises who had been experimenting with the cloud are suddenly pulling in cloud migration projects. As enterprise apps and data migrate to the cloud, users will come to expect a higher quality of service, in turn spurring apps to move again, into multicloud or edge computing. The natural result of this "virtuous cycle" is a world where users are mobile, and apps and data are elastically deployed to serve users.

These irreversible changes are affecting nearly every security sector: network security, endpoint security, identity, data protection, and threat analytics, to name a few. The old model of delivering security as functions embedded in the network cannot scale, and next generation security must also be elastic along with the applications it protects. Achieving this requires security to be completely decoupled from the network and delivered over-the-top, embedded along with the application and computing infrastructure.

The Historical Paradigm Shift – Mobility in Telephony

Over-the-Top delivery, or "OTT", is commonly associated with content. For example, Netflix and YouTube are delivered *on top of* IP networks, rather than through dedicated circuits like their cable TV predecessors. While it was Comcast's job to protect against cable theft, nobody expects them to be responsible for the security of Netflix's content – clearly, the security (user authentication and encryption, for example) are associated with the application owner, not the network provider.

> No one expects Comcast to provide Netflix's application security

But what does that have to do with today's corporate networks? How can they transition to an OTT model? We can look to another major historical technology transition as a guide: the move from fixed telephony to cellular.

Phase 1: Circuit-Switched Networks

It wasn't so long ago that desk phones were a fixture in corporate offices. These single-purpose devices assigned phone numbers based on which port on the PBX they connected to. An office move resulted in sending a technician out to the wiring closet and PBX to reprovision the network wiring. In other words, the *programming of the system is defined by the physical wiring topology*, and therefore, the identity of the phone is determined by which twisted pair it connects to. The physical security of the wiring closet helped ensure the security of the phone system.

Telephone networks were fixed function – which was fine, because there *were* no other applications at the time.

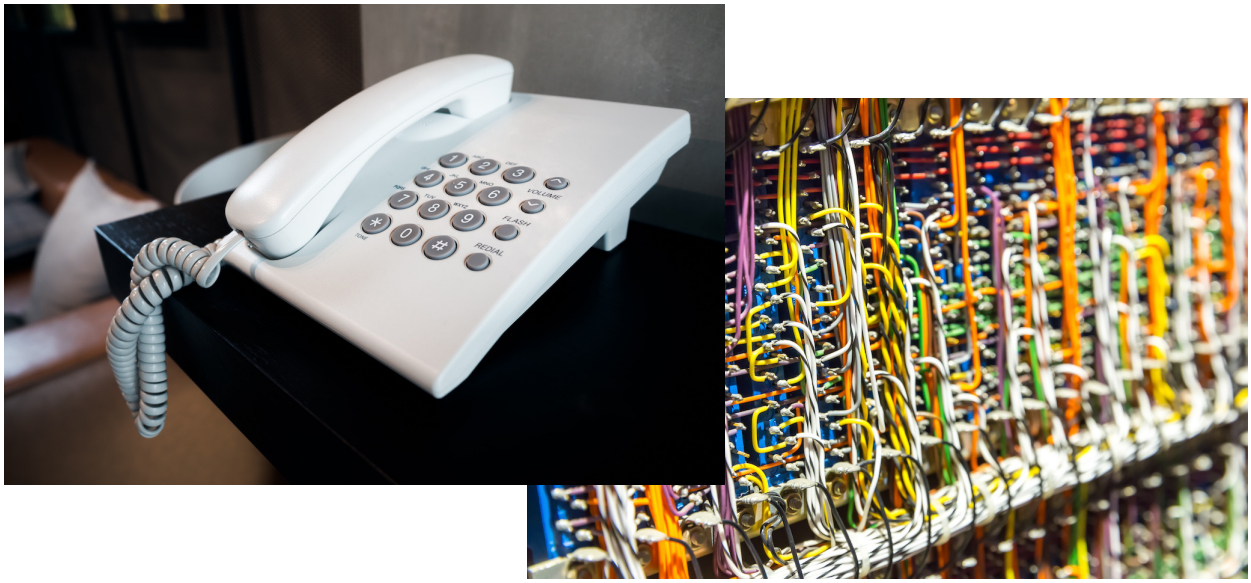

The circuit-switched network mobility paradigm.

Phase 2: Voice-over-IP

Voice-over-IP (VoIP) telephony exploded in the mid 1990s, with a radically different approach. Instead of relying on dedicated circuits, VoIP shares the same corporate data network infrastructure for the purpose of telephony. Identity moves out of the dedicated circuits and into the phone, so moving the employee extension to a new office becomes as simple as packing up the phone with the rest of the employee's belongings. Instead of being *unaware* of connectivity, the endpoint (the telephone) starts taking *responsibility* for it by using the shared network infrastructure.

At the same time, applications on a user's laptop allow them to make and receive calls even if they are not physically at their desk, providing a limited mechanism for user mobility.

With identity now determined by the phone, it became more important to secure the VoIP terminal. Manufacturers introduced certificates to help ensure the provenance of the device, and encrypted voice calls were introduced to help protect calls against snooping. In this transition, the security moved along with the application, out of the network and into the VoIP terminal, but still under the control of the corporate IP telephony team. In the meantime, the corporate date network continuously enhanced its performance and reliability on with a separate roadmap as a shared fabric for all applications.

A VoIP telephone uses a familiar session-based overlay, over standard Ethernet.

Phase 3: Cell Phones

Today, cell phones have become so compact and powerful that people are almost never without them. In fact, they are so ubiquitous that increasingly, companies even don't bother to provide desk phones for employees – it's far easier to call a user's cell phone than try to catch them at a desk phone. To provide good availability and quality of service to a mobile user, the telephony "application" had to become mobile as well. Along with the "smart phone" era, many more data applications have emerged on the single user device. In this trend the responsibility for services transitioned from companies to dedicated service providers. The shared infrastructure migrates from the private corporate data network to a new mixture including cellular networks and WiFi.

Cell phones have fulfilled the goal of mobility for telephony.

With cell phones, the implementation of security also transitioned to a fully over-the-top model, with Mobile Device Management (MDM) providing enterprises with the means to manage compliance and data through software installed on the end device.

While some cell phone applications may use a company WiFi for backhaul, they never travel east-west on the corporate network; they travel north-south, along with its application flow path, between cell phones and the service provider. As a result, there's no need for companies to manage dedicated telephony networks; they simply provide transport access, and manage security with OTT MDM.

Contrasting the Phases of Mobility in Telephony

	Wired Phone	VoIP Phone	Cell Phone
Application Location	Fixed (on desk)	Limited mobility (desk or laptop)	Full mobility (in user's hand)
Enterprise Network Required	Dedicated twisted pair network (RJ11)	Unified data communications network	Transport only (WiFi) for mobile apps
Security	**In the network wiring**	**In the endpoint, under corporate control**	**Fully software and elastic, OTT (MDM)**

The Telephony Transition to Mobility

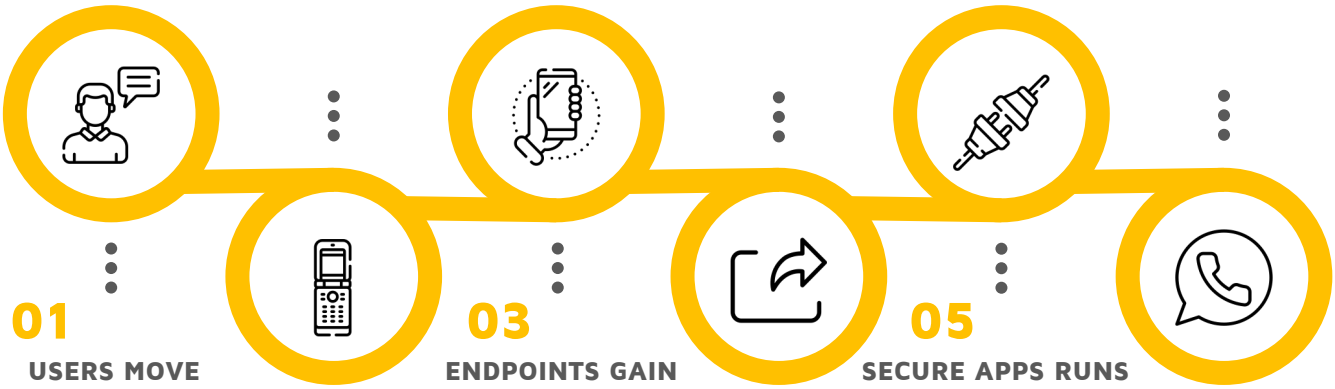

01 USERS MOVE

Users demanded telephony to support business needs

02 APPS FOLLOW

Telephony mobility fed user expectations, triggering the "virtuous cycle"

03 ENDPOINTS GAIN NETWORK AWARENESS

Responsibility for connectivity shifts from the network to the end device

04 SECURITY MOVES TO THE ENDPOINT

Identity and connection security responsibility moves to the end device

05 SECURE APPS RUNS OVER-THE-TOP

Carrier's voice calls can run over *any* data network, including corporate WiFi

06 DEDICATED NETWORKS RETIRE

Investment in corporate telephony networks shifts to more powerful data networks

Predictions for User, App, and Data Mobility

With history as a reference, we can make some predictions about how the current trend toward mobility will impact corporate apps and IT security over the next few years.

1. Users Move

Thanks to current events, this is a *fait accompli*; VPNs have become a vital lifeline for business continuity. And with more companies announcing that employees can work remotely from anywhere indefinitely, user mobility appears irreversible.

2. Apps (and Data) Follow

This phase has already begun. Enterprise IT has made cloud migration a priority and are actively working to move applications to cloud to meet the goal of business agility.

As employees get used to steadily improving application performance, the mobility virtuous cycle kicks in. However, data gravity and compliance issues will create friction, making it difficult for data to keep up with the applications. As a result, applications will increasingly be hybrid, with frontends deployed as close to the user as possible for low latency, with logic and data following behind.

As this phase matures, companies that have adopted one cloud provider will seek to adopt a second and third cloud provider for leverage, and may further migrate to edge computing for low latency.

Companies will need flexible deployment models that allow them to rapidly iterate and migrate applications to take advantage of shifting business priorities.

3. Endpoints Gain Network Awareness

Supporting agile application mobility in a complex hybrid environment means that the intelligence for application delivery will move into endpoints, and out of fixed-function network infrastructure.

We already see evidence of this in containerized applications, where mesh architectures such as Envoy allow DevOps to create and manage their own connectivity within a Kubernetes environment.

We predict that the mesh concept needs to extend beyond the confines of a Kubernetes deployment. All applications, new and legacy, will need a hybrid mesh that allows clients and servers to stay connected while they move flexibly within or across physical infrastructure silos.

4. Security Moves to the Endpoint

Just as in with mobile telephony, users, endpoints, and even applications need to be identified and authenticated. Security policies must be associated with an endpoint or application's *identity*, not with the physical network topology.

This is enabled by the concepts of Zero Trust security – although not all Zero Trust implementations support mobility.

Current network-based security point tools, such as firewalls and IDS/IPS gateways, will grow increasingly irrelevant – not because of their function, but because of their deployment model.

Protections virtually inserted at the endpoint or application edge will combine with the mesh approach in step 3 to create a *cybersecurity mesh* that provides consistent and centralized policies across all complex hybrid environments end to end, without the infrastructure and system integration effort that current point tools require.

5. Secure Apps Run Over-the-Top

Once security is baked into applications and the endpoints hosting them through a cybersecurity mesh, security and the application then begin to behave as one unit. Application workloads can then migrate between computing environments while retaining connectivity and security.

At this stage, the underlying data network plays a key role as a shared fabric facilitating physical connectivity, with high-performance and high-availability guarantees. However, application deployment itself and security policy creation will be orchestrated by the application owner, not the infrastructure owner.

Managed services providers will play a major role in this transformation, by helping implement and train enterprises in the new, flexible, OTT operational models.

6. Dedicated Networks Retire

As digital applications are rapidly reprovisioned across computing environments, fixed networks that are application-specific and require significant planning to build and investment to maintain will fall out of favor.

As companies increasingly make the remote work trend permanent, they will also deprioritize investment in office infrastructure, including the corporate LAN and WiFi. Corporate networks will still be important, but one of their primary purposes will be to provide a high-quality onramp transport for north-south application traffic over the Internet or high-speed pipes.

This trend will be supported by adoption of Zero Trust. Distrusting the internal network means treating all internal traffic as though it were going across the untrusted Internet – from a security perspective, there should no distinction between "internal" and "external" traffic.

Secure Access Service Edge (SASE) – or SASE Overlay?

The concept of the cybersecurity mesh outlined above aligns with Secure Access Service Edge (SASE), which was defined by Gartner in 2019. Recognizing that corporate users and resources have moved out of the traditional corporate perimeter, Gartner observed that enterprises will need new models that insert security from the cloud, thereby enabling the flexibility needed to handle mobility.

The benefits of SASE to security are clear; instead of having to manage disparate security appliances and enforcement points spread across many infrastructure edges, companies enjoy a centralized system with consistent security policy and services. Gartner's model promotes the convergence of security services with *networking services*, such as SD-WAN. Note that the efforts involved in converting existing static network infrastructure into dynamic networking services should not be underestimated. Connectivity at the network level traditionally involves tremendous infrastructure and operations (I&O) complexity; re-engineering of the shared network can impact many live applications and users that are depending on it.

We believe that coupling security to the *network*, rather than to the *application*, dramatically increases the scope of the solution and therefore the *solution complexity*, both for vendors as well as customers. The simple option is to let the network infrastructure continue to serve as a shared fabric doing what it's good at – high performance packet transport – and let security associate instead with the application. Elastic networking services will eventually be a requirement. However, it may not be ideal to couple that with the security services as part of the implementation stack.

We call this alternate model a SASE *overlay*, as it seeks to provide the same security benefits of SASE, using the physical network to create a cybersecurity mesh, but not coupling to it for implementation of security policy enforcement and operation.

A SASE overlay simplifies enterprise security by creating a unified stack of security services that can be inserted between users and resources, wherever they are.

Decoupled from physical networks, the SASE overlay connects and protects users, endpoints, and applications, without requiring stitching together network silos at Layer 3.

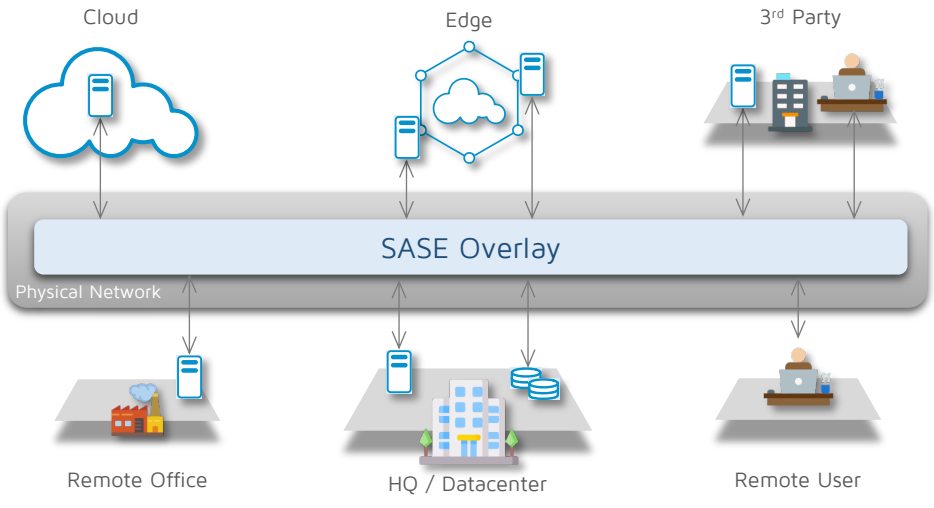

Zentera Systems' SASE Overlay: CoIP® Access Platform

CoIP Access Platform is an advanced SASE overlay providing end-to-end security for users, endpoints, and applications in complex and distributed computing environments. Unlike legacy VPNs and firewalls, it builds connectivity and enforces security for authorized accesses on-demand. Following the Zero Trust model, every access is authorized against centralized policies that are based on user, endpoint, and application identity, so they don't need to be rewritten when the application or user moves within or across network silos.

At the heart of CoIP Access Platform is the CoIP SASE Overlay, a session-based overlay technology that decouples applications from the underlying IP network silos and relieves IT and DevOps from the operational drudgery of building and maintaining network connectivity. With the CoIP SASE Overlay, distributed users and hybrid applications can be set up to connect across sites in minutes, without tickets to open brownfield firewalls or configure VPNs to connect subnets, and even to cloud VPCs.

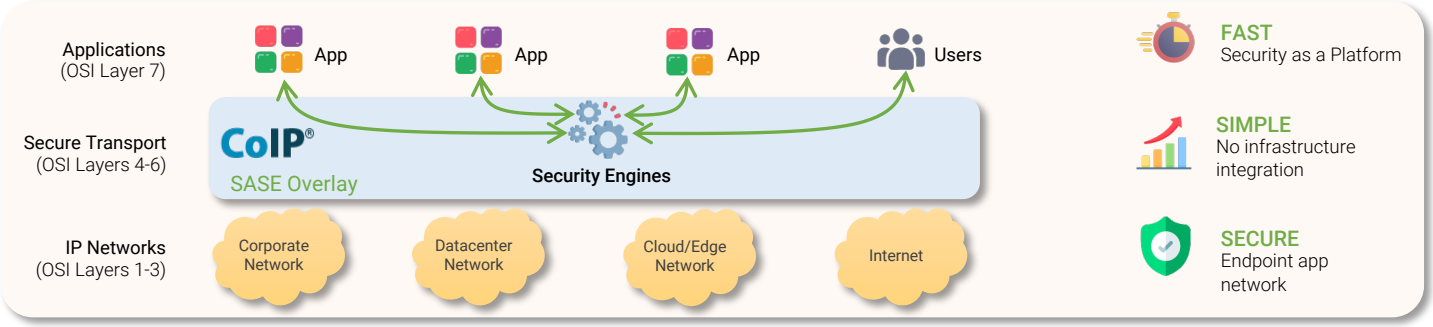

CoIP Access Platform's SASE overlay delivers secure connectivity across disconnected IP network silos.

Using CoIP Access Platform, customers can instantly connect applications and users running in complex hybrid environments, inserting security services such as:

- Zero Trust Network Access (ZTNA) for least-privilege remote user access to corporate resources and applications
- Application and micro-segmentation for cloaking distributed applications in the cloud or on-prem
- Layer 7 firewalling and threat prevention
- TLS 1.3-based link encryption
- Threat detection and prevention

CoIP Access Platform supports a variety of onboarding models that enable nearly any application workload to be onboarded to the CoIP SASE Overlay easily.

Agent-based

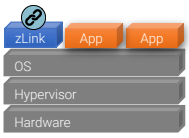

The zLink agent acts as a security sensor which builds an identity profile of the endpoint and applications running on it. It creates a virtual interface that connects to the CoIP SASE Overlay, allowing applications to communicate with each other through the cybersecurity mesh. It enforces endpoint security policies through periodic checking of endpoint trust factors, as well as end-to-end connectivity policies through an application-aware firewall. As a security enforcement point, it associates with the *application* rather than the network, allowing workload migration to be decoupled from the underlying physical network.

Agentless

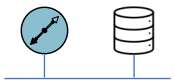

The Gateway Proxy is a virtual appliance that deploys near a service, acting as an onramp/offramp for the CoIP SASE Overlay. Remote applications can access the service through the Gateway Proxy, with firewall and access control policies applied automatically.

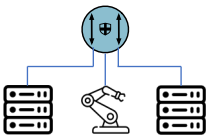

The Micro-Segmentation Gatekeeper is a hardware appliance, deployed inline with assets that cannot support an agent, but require higher security at an individual device level – for example, east-west micro-segmentation, access control, or inline threat prevention. These capabilities and agentless deployment model make it ideal for protecting critical databases, IoT devices, or OT equipment.

As an increasing number of existing applications are adapted to a mobile, digital future, this breadth of onboarding options ensures that nearly any application can be protected through the CoIP SASE Overlay.

Mobility in Action: Cloud Migration with a SASE Overlay

Let's examine how the CoIP SASE Overlay can enable security to migrate along with the application, in support of improved service quality for users.

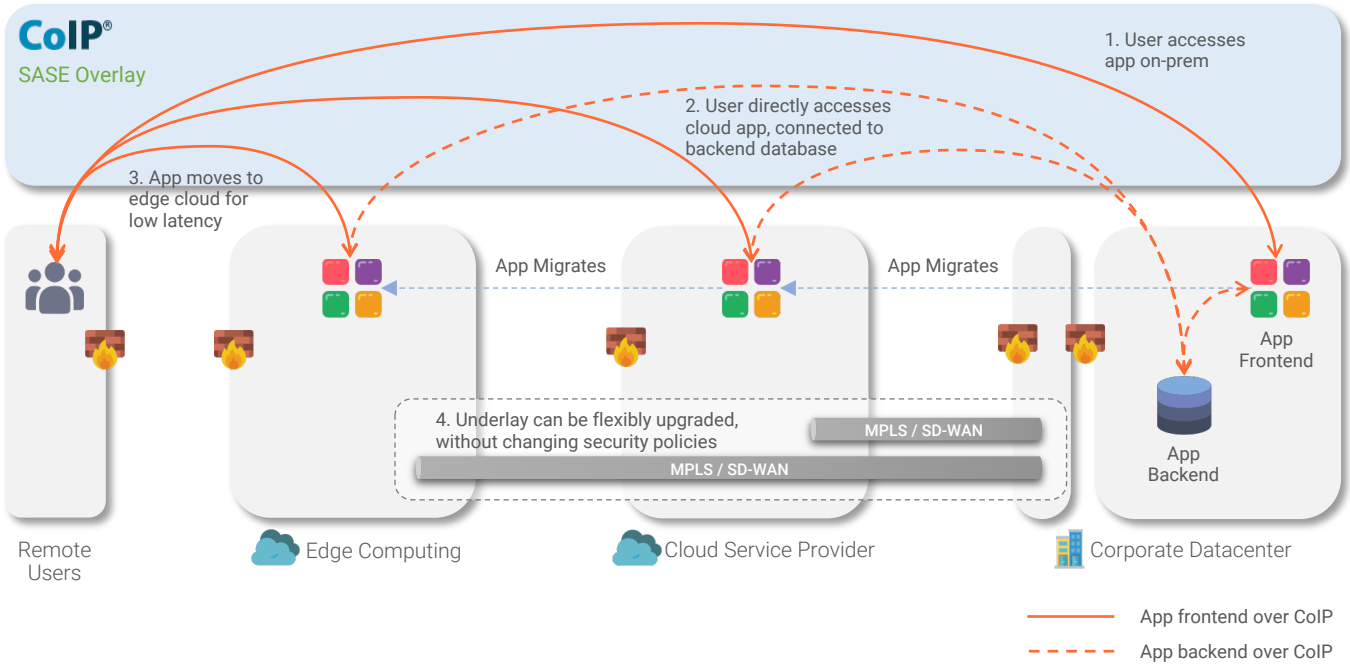

Deploying SASE as an overlay enables connectivity and security policies to be written specifying user access to applications. As these policies are decoupled from the IP addresses and topology in each network, the application components can move freely without impacting security policy.

As the application frontend migrates closer and closer to the user to optimize performance or enabling new business opportunities, the user access method and experience do not change, and regardless of whether the user is at home, in a coffee shop, or using guest WiFi at a customer site, the same Zero Trust security checks are enforced to ensure minimal access to the application.

The backend application tiers and database can also move or optionally remain on premises (for example, for compliance reasons). CoIP Access Platform provides minimal access to the app backend, so that only authenticated and authorized app frontend servers can access the app backend servers in the north-south direction through overlay.

Additionally, CoIP Access Platform enables all connectivity to be flexibly delivered over the Internet or any other private network transport. As the SASE overlay decouples security policies from the network topology, the team can optimize the access path by deploying high-speed links over time replacing Internet access, without changing any of the security policies or firewall rules in the overlay.

Conclusion

Zentera's overlay approach to SASE has significant operational as well as security benefits to the cybersecurity and applications teams, by decoupling their efforts from the requirements of an infrastructure build. In this way, CoIP Access Platform can directly address many of the technical and operational challenges from scaling the use of cloud, multi-cloud, and hybrid applications, to provide high performance, consistent and end-to-end security, and streamlined access controls across all environments.

For more information on CoIP Access Platform, please visit our website at https://www.zentera.net. To speak with an architect, please contact us at support@zentera.net.

zentera™

CoIP® Access Platform

Securing Application Access for the Next-Gen Enterprise

Remote Work has Created New Cyber Risks

It wasn't that long ago that employees could be expected to come to the office to access files, applications, or machines needed to get work done. And then, the COVID-19 pandemic hit.

How things have changed.

Business leaders and employees alike changed their expectations overnight. Now, all but the most critical workloads must be accessible from anywhere. Most companies solved that problem with tried and trusted VPN to connect remote workers to the corporate network.

But decades-old VPN and firewall network infrastructure concepts are just not up to the task of managing granular access in today's rapidly changing and hyperconnected business environments. It's no wonder attacks against the VPN have skyrocketed since the pandemic began.

It turns out that hackers have paid attention to the remote work trend, and have stepped up their efforts to penetrate VPNs.

My company uses a VPN. Why isn't that enough?

VPNs provide connectivity, by connecting a user into the corporate network. While it encrypts communications in transit, it doesn't control *who* can access *what*. On the other hand, it's not practical to expect the corporate firewall to block all network attacks, including zero-day exploits.

This means that VPN vulnerabilities, stolen credentials, or a compromised laptop can give malicious actors instant access to your most sensitive applications and data on-prem.

[1]https://www.helpnetsecurity.com/2021/06/15/vpn-attacks-up/

Identifying Key Remote Access Concerns: Think Like a Hacker

Before we can decide whether tools and policies for remote access are secure, we must think about the nature of remote access – how it should be used, and how it can be abused.

The former is straightforward; authenticated users must be able to access apps and data.

So how can remote access be abused? Here, it helps to think like a hacker. Let's assume we have gained VPN credentials for XYZ Corporation and intend to do some damage. What options do we have?

Exploitation, in which we take over control of XYZ servers

Reconnaissance, which includes exploring the XYZ network to find other exploitable targets and valuable targets (e.g. applications and databases).

Inject Command and Control (C2) malware, which creates another connection to XYZ Corp, to be used in case we ever lose access to the VPN.

Inject ransomware with which we can snarl XYZ's systems and bring XYZ's operations to a screeching halt.

Exfiltrate sensitive data, either to monetize ourselves, or to sell to another party.

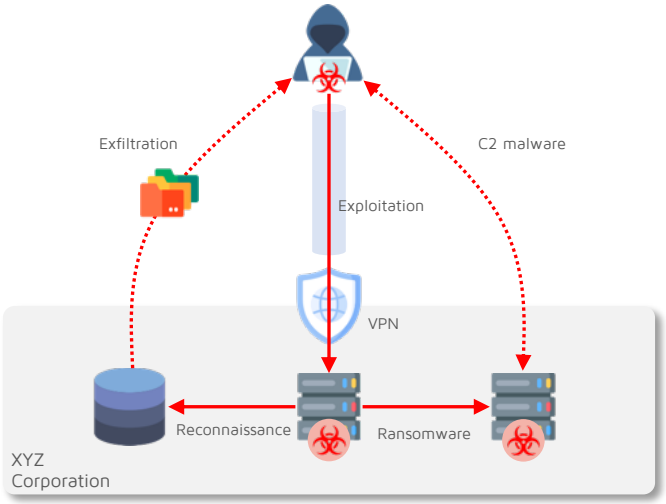

To protect XYZ Corp, we'll need to pay attention to attacks in multiple *directions*.

North-South – between the XYZ network and a remote user or site

East-West – laterally between machines on the XYZ network

The Fundamental Principle for Cyber Protection of Remote Access

From our thought experiment, we can observe the following:

- VPN-borne attempts to exploit servers or inject malware and ransomware are typically *inbound* to the organization from a remote hacker, and then move laterally, closer to targets

- Attempts to exfiltrate data are *outbound* from the organization – but need not be carried over the VPN. It's just as easy for a hacker to upload critical data to an online storage service like Dropbox

But valid users still need access to applications and data. Logically, we can preserve user access while blocking abuses by:

1. **Identifying the remote access user using multiple factors**
   This makes it much harder for a hacker to use stolen credentials

2. **Giving the user access only to specific applications and data, not networks**
   This prevents hackers from exploiting other targets, performing reconnaissance, and roaming freely within the corporate network

3. **Restricting the access methods available to the user**
   Limiting access to specific application ports, and even specific software binaries where possible, we can make it much harder for a hacker to exploit a server; we can also prevent the remote user from exfiltrating corporate data

4. **Segmenting target application servers away from the rest of the corporate network**
   This protects the application server from malware and insider attacks that may be present on the corporate network

5. **Controlling application server access to external services, such as on-prem services and the Internet**
   Control over this access helps block exfiltration and C2 activity, while still allowing access for approved uses such as software updates.

These principles can be summed up by one fundamental principle: **give users and application the least-privilege they need to do their jobs.**

How CoIP Access Platform Enforces Least-Privilege for Cyber Attack and Data Leak Prevention

CoIP Access Platform is a next-generation Zero Trust Security platform that enables you to quickly and easily control access between users and servers to achieve least-privilege control. It does this by merging *connecting* and *blocking*: connecting things based on whitelisted application and operations behavior and blocking the rest.

Zero Trust Network Access (ZTNA)

CoIP Access Platform implements ZTNA, authenticating users and servers based on multiple identity factors. Policies allow you to grant access on an individual server or application basis.

By connecting the remote user to an *application*, rather than connecting them to the network, there's no network path to be abused. In this way, CoIP Access Platform mitigates inbound cyber attacks, such as exploitation, reconnaissance, and even ransomware.

Application Chamber

CoIP Access Platform also implements an Application Chamber, which apply software-defined network packet filtering and application controls around a server or group of servers – similar to the concept of an application firewall. An Application Chamber provides hardened security controls for groups of application servers, and effectively blocks lateral attacks of the application.

As shown below, the combination of ZTNA and Application Chamber provides powerful and holistically integrated protection to prevent cyber attacks as well as data leaks.

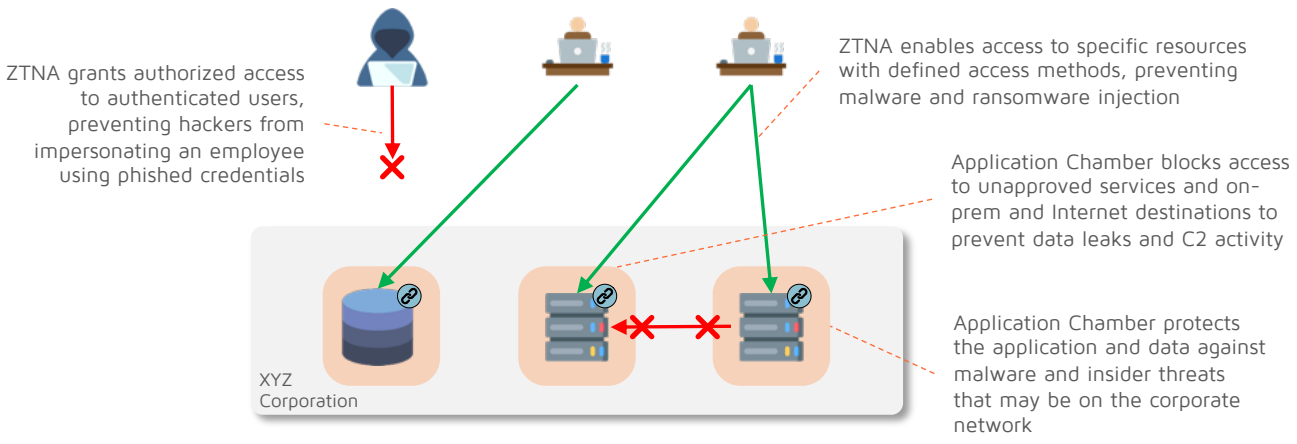

Use Cases for Secure Application Access

CoIP Access Platform elegantly enables a wide range of use cases for secure application access and protection, including:

- **Application Chamber preventing advanced cyber attacks**
- **Virtual Chamber for collaboration with ecosystem partners**
- **Virtual Chamber for secure software and hardware development**
- **Virtual Chamber for vendor secure access and support**
- **Remote Desktop for Work-From-Home initiatives**
- **Secure Remote Access for customer support**
- **Private Application Access during mergers and acquisitions**

Business Benefits

CoIP Access Platform enables enterprise to easily implement next-generation Zero Trust Security for application access security and control. The business benefits include:

- Protects applications from cyber attacks and data leaks inside the network or over the VPN
- Authenticates and authorizes users for least-privilege, Zero Trust access to applications
- Avoids "rip and replace" with software-defined overlay technology that deploys easily on existing infrastructure
- Dramatically lower TCO and improved ROI, compared to traditional infrastructure methods

More Resources

On the web:
www.zentera.net

Email:
sales@zentera.net

Phone:
+1 (408) 436-4811