White Paper

Implementing Segmentation in Phases

Stopping the lateral movement of threats and preventing breaches

Introduction

Segmentation is an effective security strategy to stop the lateral movement of threats inside a cloud or data center environment—whether across virtual private clouds (inter-cloud) or inside them (intra-cloud). Segmentation is not defined by a single implementation or security control; a broader and more effective view of segmentation considers multiple types of controls, all working together to prevent threats from spreading across the network—across users and workloads.

After working with numerous customers who have purchased the full suite of products, Zscaler has built a complete segmentation strategy that focuses on the biggest risk areas first, and then layers on subsequent controls in a way that makes microsegmentation attainable. This strategy is implemented in three phases and follows zero trust principles. Zscaler believes that segmenting and securing application connectivity is a journey that requires a thoughtful and measured approach—one that balances both security requirements and operational considerations. The three phases of segmentation are:

Phase one: user-to-application segmentation

Phase one starts by eliminating the largest threat to your infrastructure: user connectivity. Users are often the weakest point in an organization's security posture; therefore, an effective segmentation solution must protect against threats from compromised users spreading across the network and compromising workloads. Zscaler for Users segments and secures user-to-app communication and prevents any threats from infecting users communicating with the internet before they go on to access private applications.

Phase two: application-to-application segmentation across VPCs, regions, and clouds

Phase two ensures that all ingress and egress points are reviewed and remediated by evaluating connectivity and reducing the attack surface wherever possible. This prevents threats from infected applications or workloads spreading from one environment to another. Environments include virtual private clouds (VPCs), virtual networks (VNets) in public clouds, and on-premises data centers. Zscaler for Workloads segments and secures app-to-app communications and prevents any threats from infecting workloads communicating to the internet.

Phase three: application-to-application segmentation within VPCs

Phase three layers on granular process-level control to further lock down connections routing through the Zscaler Zero Trust Exchange™, as well as block the potential for server-to-server lateral movement. Zscaler for Workloads uses identity to segment and secure application-to-application communication down to the process level; only verified applications can communicate within the VPC to deliver zero trust security.

In summary, the phased approach to segmentation first controls access into the environment and enforces least-privileged access by defining and filtering all ingress and egress points. This consolidation (and filtering) of all inbound connectivity drastically simplifies the microsegment policy set to make microsegmentation more achievable in a realistic time frame.

For segmentation to be effective, it must follow zero trust principles—where the network is assumed to be untrusted, and the identity and context of users and workloads must be verified before any communication is allowed. Let's look at a three-phased approach to segmentation based on zero trust.

Zero trust segmentation, phase one: user-to-app

Infected users are the number one risk to an organization. Legacy architectures admit the user to the corporate network and allow them to access any application across the enterprise, which is extremely risky. For example, a salesperson logged in via a VPN, who should have access only to the CRM application, could see all the other applications in the environment, and could even get access to a sensitive HR application. To reduce this risk, user access to applications should be restricted to only the necessary applications (i.e., user-to-app segmentation).

Unlike legacy architectures, Zscaler never places users on the network. The Zero Trust Exchange architecture directly connects users to applications for true zero trust security. Because all communications go through the Zero Trust Exchange, the network is no longer routable. This reduces security risks and greatly simplifies segmentation policies without the network complexity.

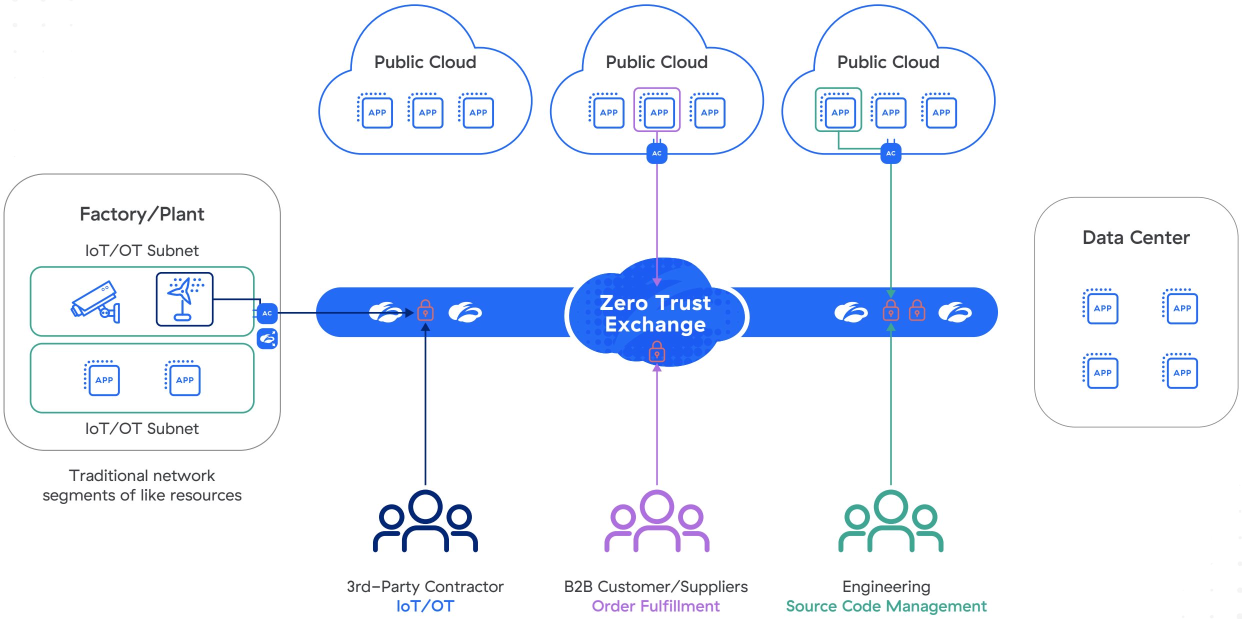

Figure 1: Phase 1, user-to-app segmentation. Business policies connect users to apps, not networks. See the colored lines connecting specific users to specific applications and nothing else. This eliminates the risk of lateral threat movement.

Zscaler makes it simple to achieve user-to-application segmentation. First, an application segment (e.g., a CRM application) needs to be defined. Zscaler provides the ability to define an application segment as specifically or granularly as needed. Then, the application segment needs to be added to a policy where specific users, based on their identity and context (e.g., SAML attributes, device status), are allowed to access the application segment—that's it. To simplify policies further, application segments can be added to segment groups.

Implementing phase one is very simple, especially for existing Zscaler customers who may have started with a more open connectivity-oriented approach, because nothing new needs to be deployed on end user devices, in the cloud, or in data centers. Zscaler Client Connectors and App Connectors are used to enforce segmentation. All that needs to be done is defining segmentation policies in the Zero Trust Exchange. User-to-application segmentation is a foundational first step that is often tackled before organizations move to phase two.

Zero trust segmentation, phase two: app-to-app across environments

The second phase of segmentation is to ensure that application workloads in public clouds can securely communicate with other applications across VPCs, cloud regions/providers, and back to on-premises data center applications. This approach secures the remaining ingress and egress points—i.e., traffic leaving a VPC or re-evaluating VPC connectivity to on-premises data centers—and ensures there is no open, persistent network connectivity presenting an attack surface and increased risk of lateral threat movement. Additionally, outbound/egress protection for workloads ensures that workloads can securely communicate to internet destinations.

Zscaler uses the same Zero Trust Exchange (explained in Phase 1) to directly connect application segments to other application segments without placing them on the network, yielding true zero trust security and eliminating lateral threat movement across VPCs, regions, and cloud environments.

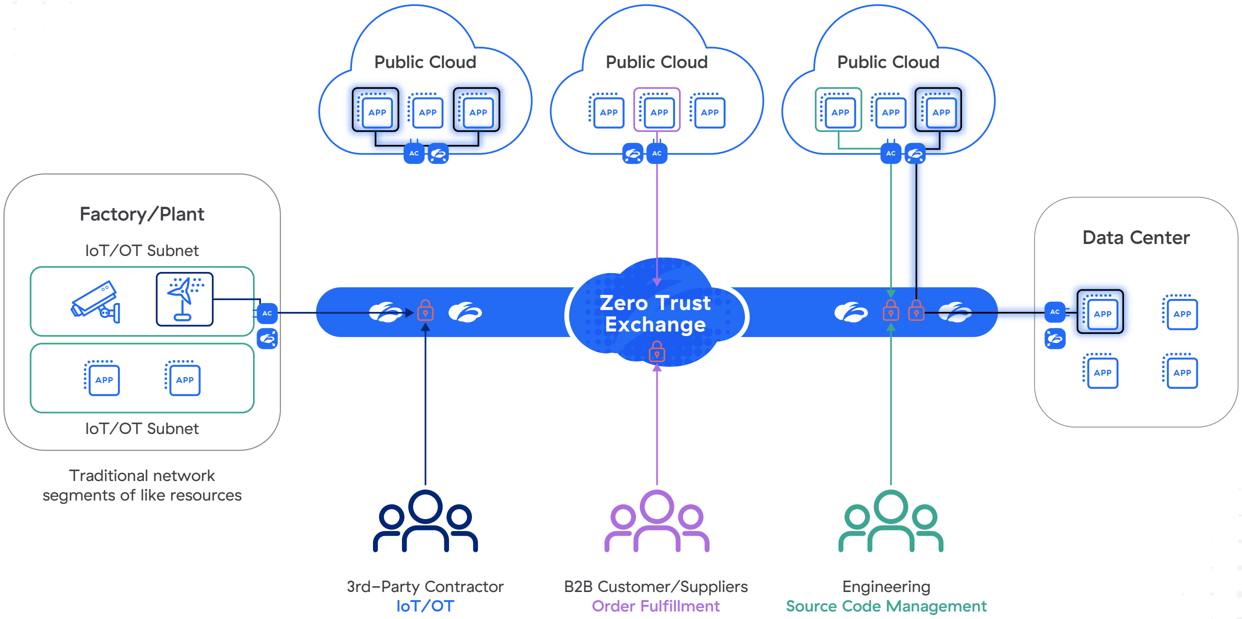

Figure 2: App-to-app segmentation—lines highlighted in a blue glow show how only specific application workloads securely communicate across two regions of a public cloud, across clouds, and from the public cloud to data centers.

Legacy approaches to workload connectivity and segmentation have relied on extending the network to the cloud. This approach uses network-based VPNs and firewalls to create full mesh networks and separate network segments. This network-based approach requires potentially thousands of firewall rules, which are complex to create and manage—and it offers little security benefit, since the network remains routable, and applications can still be discovered and exploited.

In contrast, Zscaler Private Access for Workloads uses the Zero Trust Exchange architecture to directly connect applications to other applications across any environment. The zero trust architecture that secures user-to-application access is also used for securing application-to-application access.

This approach further simplifies segmentation while extending zero trust security to application-to-application communication. Similar to phase one, with the Zero Trust Exchange, there is no IP-connected routable network, so lateral movement across environments (e.g., VPCs, regions, or clouds) is prevented while only approved applications can communicate. Because applications are not on the network, segmentation policies can be much simpler.

Implementing phase two is simple because only a Cloud Connector virtual appliance needs to be installed at the VPC/VNet gateways. Zscaler provides automated scripts and templates (e.g., Terraform and CloudFormation) for installing Cloud Connectors in the public cloud environments—these can be downloaded from the AWS and Azure marketplaces. Moreover, policies are easily defined in the familiar ZPA console.

With phases one and two implemented, you will have made significant progress with your segmentation initiatives to stop the lateral movement of threats from users to applications as well as between applications across environments. To further secure app-to-app communications within a VPC, you can use microsegmentation.

Zero trust segmentation, phase three: app-to-app within VPCs

This phase of segmentation adds another layer of security to applications that protects communications within a VPC environment by getting rid of the flat network and stopping lateral movement. This is also commonly referred to as microsegmentation.

Legacy firewall-based approaches to microsegmentation result in thousands of rules that are impossible to manage. This approach is not secure because attackers, once inside an environment, can piggyback on approved firewall rules. These address-based controls have no visibility into the identity of the communicating applications, which makes them exploitable.

Zscaler Workload Segmentation is the only zero trust solution that microsegments applications using their cryptographic identity fingerprint, rather than address-based controls, to protect process-to-process level communication within a VPC or data center environment. This is the highest level of protection for the most sensitive applications.

Zscaler Workload Segmentation provides zero trust protection by verifying the identity of the communicating software, including processes, before allowing any communication. Using identity this way enables you to reduce policies by up to 90% vs. firewall rules, dramatically simplifying microsegmentation. Zscaler Workload Segmentation provides visibility into application communications and tests effectiveness before enforcement using simulated policies, which adapt to allow for the introduction of new applications or application upgrades in your environment. Workload Segmentation uses an agent-based approach.

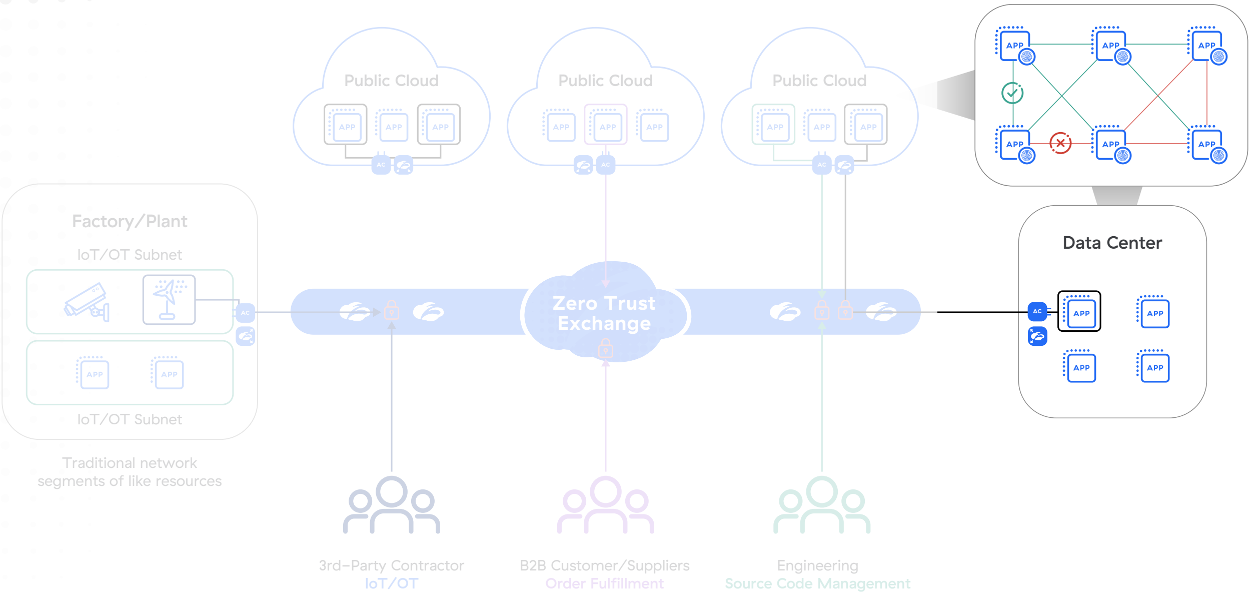

Figure 3: Identity-based segmentation at the process level. See the top-right corner—workload communications inside a VPC or a data center are protected by identity verification. Workloads are fingerprinted to establish identity.

Before beginning this phase, all workload and user access should ideally be secured by Zscaler Private Access for Workloads in phases one and to avoid “local unmanaged” policies, which occur when agents cannot be installed on certain workloads. If phases one and two are not done, you may need to accommodate numerous IP ranges representing groups of clients on the network who should have access to certain applications—which creates lateral movement risk because you cannot verify the source.

Stronger security with the phased approach to segmentation

In summary, accomplishing the three phases of segmentation will offer the highest level of protection from lateral threat movement while advancing your organization’s zero trust security initiatives. Zscaler delivers a holistic segmentation solution that is easy to implement and delivers measurable value at every phase.

1. User-to-application segmentation	**2. Application-to-application segmentation across environments**	**3. Application-to-application segmentation within a VPC**
Phase one eliminates the greatest source of risk—compromised users infecting application workloads in the cloud/data center	Phase two provides the second layer of protection to prevent threats from entering an environment (e.g., when work-loads connect to the internet) and infected workloads moving laterally across the entire cloud/data center environment	Phase three layers on true microsegmentation, at the application process level, to prevent lateral movement of threats within a VPC

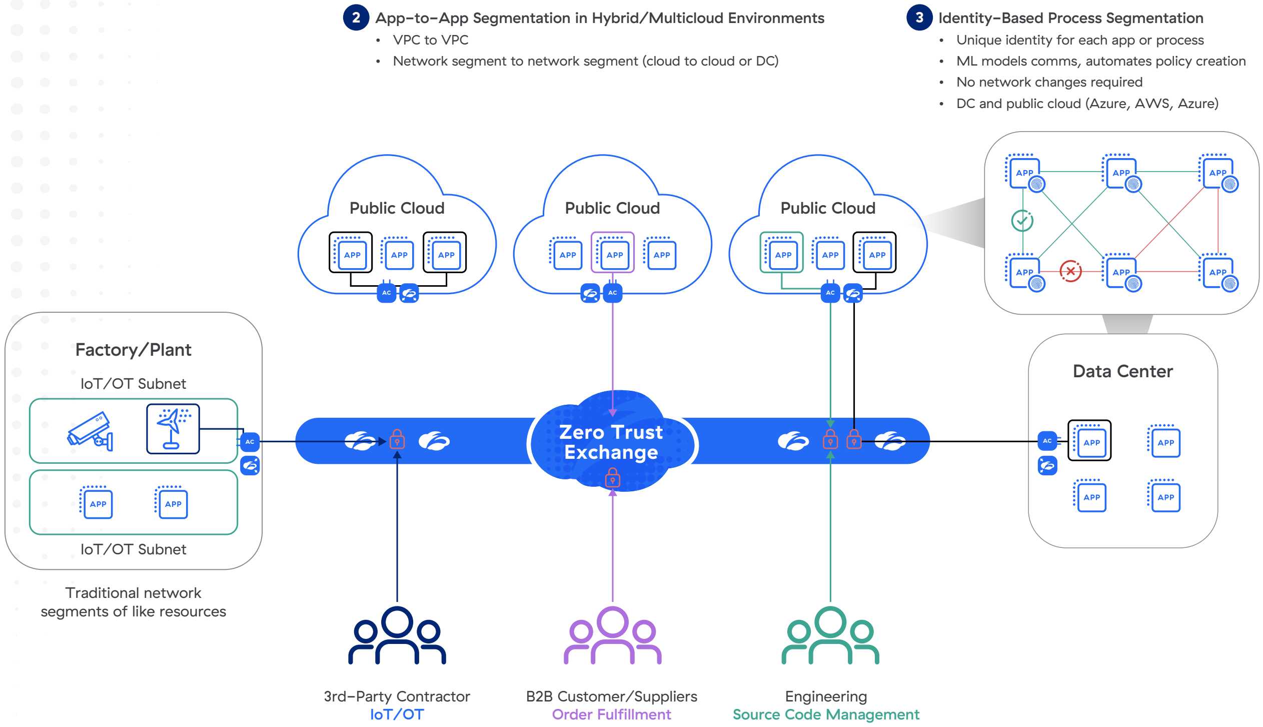

2 App-to-App Segmentation in Hybrid/Multicloud Environments

- VPC to VPC
- Network segment to network segment (cloud to cloud or DC)

3 Identity-Based Process Segmentation

- Unique identity for each app or process
- ML models comms, automates policy creation
- No network changes required
- DC and public cloud (Azure, AWS, Azure)

1 User-to-App Segmentation

- Business policies connect users to apps, not networks
- No lateral threat movement

Figure 4: The three phases of segmentation work together as a single solution.

Talk to your Zscaler representative today to see a demo. You can also schedule an architecture workshop to start building your segmentation plan.

zscaler™ | **Experience your world, secured.™**

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so that customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SASE-based Zero Trust Exchange is the world's largest inline cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

+1 408.533.0288 Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134 zscaler.com