



Universal ZTNA with Zscaler Private Access™ Private Service Edge

Reference Architecture

Contents

About Zscaler Reference Architectures Guides	1
Who Is This Guide For?	1
A Note for Federal Cloud Customers	1
Conventions Used in This Guide	1
Finding Out More	1
Terms and Acronyms Used in This Guide	2
Icons Used in This Guide	3
Introduction	4
On-Premises and Universal ZTNA	5
Key Features and Benefits	7
New to ZTNA or ZPA?	7
Understanding ZPA Private Service Edge Deployments	8
Common Use Cases for a ZPA Private Service Edge	8
Deployment Requirements and Scalability	9
Software Updates and Shared Responsibility Model	10
Understanding IP Addressing and Public Accessibility	10
Confirm Your ZPA Private Service Edge Deployment	12
ZPA Private Service Edge Policy Design	13
Mapping ZPA Private Service Edges to Trusted Networks	13
Restricting User Traffic by Location	14
Client Forwarding Policies	15
Grouping App Segments or Segment Groups by Location	15
Understanding Control Plane Connections	16
Enrollment Overview – Private Service Edge	16
Enrollment Overview – App Connector	16
Enrollment Overview – Zscaler Client Connector	17

Understanding Traffic Forwarding	19
Traffic Forwarding – ZPA Private Service Edge	19
Traffic Forwarding – App Connector	19
Traffic Forwarding – Zscaler Client Connector	20
Recommendations for Gradual Migration to Your ZPA Private Service Edge	21
Migrating to Your ZPA Private Service Edge One App at a Time	22
Migrating to Your ZPA Private Service Edge One Group of Users at a Time	22
Combined Approach	23
Use Cases	24
On-Premises Users Accessing Applications on the Internal Network	24
Branch Users Accessing Applications Within the Branch Location	26
Remote Users Accessing Regional Internal Applications	28
Remote Users Accessing Regional Internal Applications in a Public Cloud	31
Remote and On-Premises Users Accessing Internal Applications	33
About Zscaler	36

About Zscaler Reference Architectures Guides

The Zscaler™ Reference Architecture series delivers best practices based on real-world deployments. The recommendations in this series were developed by Zscaler's transformation experts from across the company.

Each guide steers you through the architecture process and provides technical deep dives into specific platform functionality and integrations.

The Zscaler Reference Architecture series is designed to be modular. Each guide shows you how to configure a different aspect of the platform. You can use only the guides that you need to meet your specific policy goals.

Who Is This Guide For?

The Overview portion of this guide is suitable for all audiences. It provides a brief refresher on the platform features and integrations being covered. A summary of the design follows, along with a consolidated summary of recommendations.

The rest of the document is written with a technical reader in mind, covering detailed information on the recommendations and the architecture process. For configuration steps, we provide links to the appropriate Zscaler Help site articles or configuration steps on integration partner sites.

A Note for Federal Cloud Customers

This series assumes you are a Zscaler public cloud customer. If you are a Federal Cloud user, please check with your Zscaler account team on feature availability and configuration requirements.

Conventions Used in This Guide

The product name ZIA Service Edge is used as a reference to the following Zscaler products: ZIA Public Service Edge, ZIA Private Service Edge, and ZIA Virtual Service Edge. Any reference to ZIA Service Edge means that the features and functions being discussed are applicable to all three products. Similarly, ZPA Service Edge is used to represent ZPA Public Service Edge and ZPA Private Service Edge where the discussion applies to both products.



Notes call out important information that you need to complete your design and implementation.



Warnings indicate that a configuration could be risky. Read the warnings carefully and exercise caution before making your configuration changes.

Finding Out More

You can find our guides on the [Zscaler website](https://www.zscaler.com/resources/reference-architectures) (<https://www.zscaler.com/resources/reference-architectures>).
















You can join our user and partner community and get answers to your questions in the [Zenith Community](https://community.zscaler.com) (<https://community.zscaler.com>).

Terms and Acronyms Used in This Guide

Acronym	Definition
DC	Data Center
DMZ	Demilitarized Zone
DNAT	Destination Network Address Translation
IdP	Identity Provider
NOC	Network Operations Center
SDP	Software-Defined Perimeter
SSL	Secure Socket Layer (superseded by TLS)
TLS	Transport Layer Security
URL	Uniform Resource Locator
ZDX	Zscaler Digital Experience
ZIA	Zscaler Internet Access
ZPA	Zscaler Private Access
ZTE	Zero Trust Exchange
ZTNA	Zero Trust Network Access

Icons Used in This Guide

The following icons are used in the diagrams contained in this guide.

Zscaler Zero Trust Exchange		Router	
ZIA or ZPA Service Edge		Legacy Firewall	
Zscaler App Connector		Headquarters Location	
Laptop With Zscaler Client Connector Installed		Private Data Center	
Cell Phone With Zscaler Client Connector installed		Internet	
Laptop with VPN Agent Installed		Data Tunnel	
IPSec Concentrator		Database	
Generic Application or Workload			

Introduction

Zero trust network access (ZTNA), also known as the software-defined perimeter (SDP), is a set of technologies and functionalities that enable secure access to internal applications for users. It operates on an adaptive trust model, where trust is never implicit, and access is granted on a least privileged basis defined by granular policies. ZTNA gives remote users seamless, secure connectivity to private applications without ever placing them on the network or exposing apps to the internet.

Traditionally, ZTNA has been viewed as a way for an organization's remote employees to securely access internal applications. Previously, user VPNs were relied on for remote access, extending your organization's network perimeter to include the end user. When a user is on the network, they can move laterally. The same is true for viruses and ransomware, or malicious insiders.

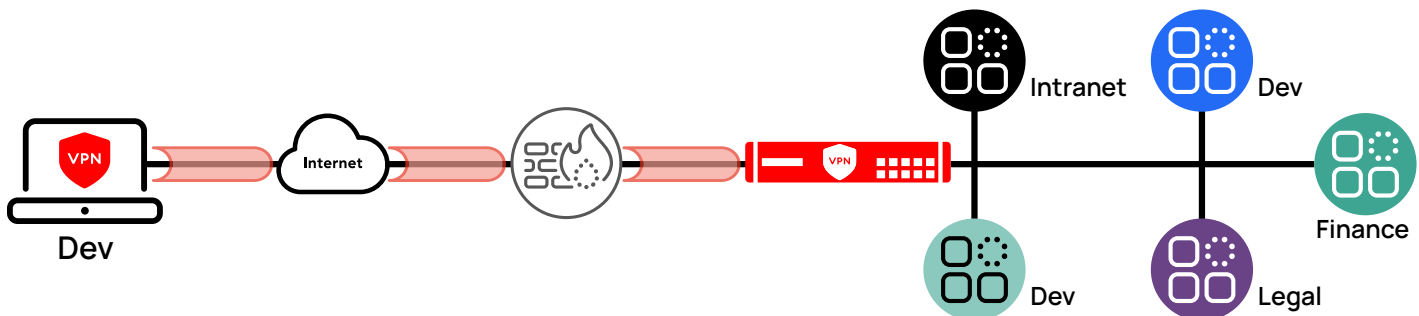


Figure 1. Traditional user VPNs put the user directly on your organization's network, allowing them to move laterally through the network

Placing controls around VPNs is also complicated. User issues from forgetting to turn on the VPN to turning it on unnecessarily lead to help desk complaints. Coordinating network access control (NAC) and firewall appliances, building quarantine VLANs, and segmenting your network by application requires constant monitoring and leaves little flexibility.

ZTNA works to counter this complexity through policy and visibility controls to applications. Zscaler Private Access (ZPA) delivers ZTNA to organizations that need to access internal applications in public clouds, private clouds, and in the organization's data centers. ZPA does this through policies made up of criteria you define.

When a user authenticates to your identity provider (IdP), they are evaluated for context in addition to their authentication response. This context can include things like location, device posture, user groups or department, and time of day. It can also include inputs from your endpoint security software running on user's devices.

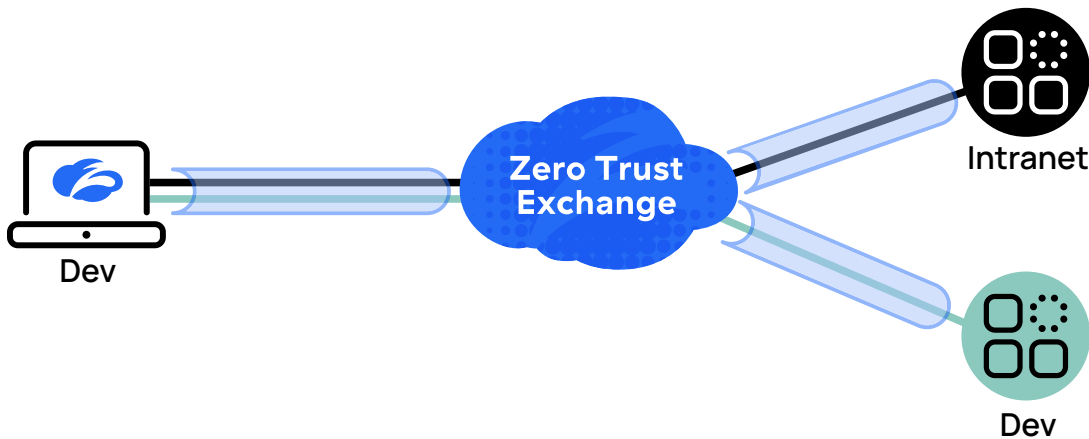


Figure 2. ZTNA restricts visibility to authorized users, hiding the resource completely from those not approved by policy

When the user attempts to access a resource, this context is used to match against a set of policies. If the user context matches the requirements you've set, the user is allowed to access the resource. If the requirements fail, it is as if the resource does not exist on the internet. It won't be reachable by hostname or IP address. Users have no way to find the application if they are not authorized to use it.

Several use cases exist for ZTNA, the first being remote access to internal resources. The second is using ZTNA to secure resources when accessed by users on the organization's premises. This in effect treats the organization's local network as if it were the internet. Many large organizations have taken this path, treating their internal network more like that of a coffee shop. There is no authentication to join, but you must authenticate and protect traffic to access any internal applications.

On-Premises and Universal ZTNA

Traditionally, the users at the organization's headquarters and branch sites have not been subject to these controls. In many organizations, the network at a site is deemed "safe" and the users "trusted" if they are within the network perimeter. These assumptions are being questioned as users transition to hybrid work, where the machine is used away from the organization's controlled work sites much more often.

Organizations are realizing this risk and are shifting to bringing ZTNA to their trusted sites. In this model, all resources are reachable only by using the ZPA service. By shifting to an on-premises ZTNA posture, you reduce your risk from machines entering the network in an infected state or from malicious insiders. Lateral movement through the network by viruses and ransomware is restricted, and your users have a consistent experience accessing resources. Universal ZTNA is the use of ZTNA for on-premises and remote users, with no distinction made as to the user's location.

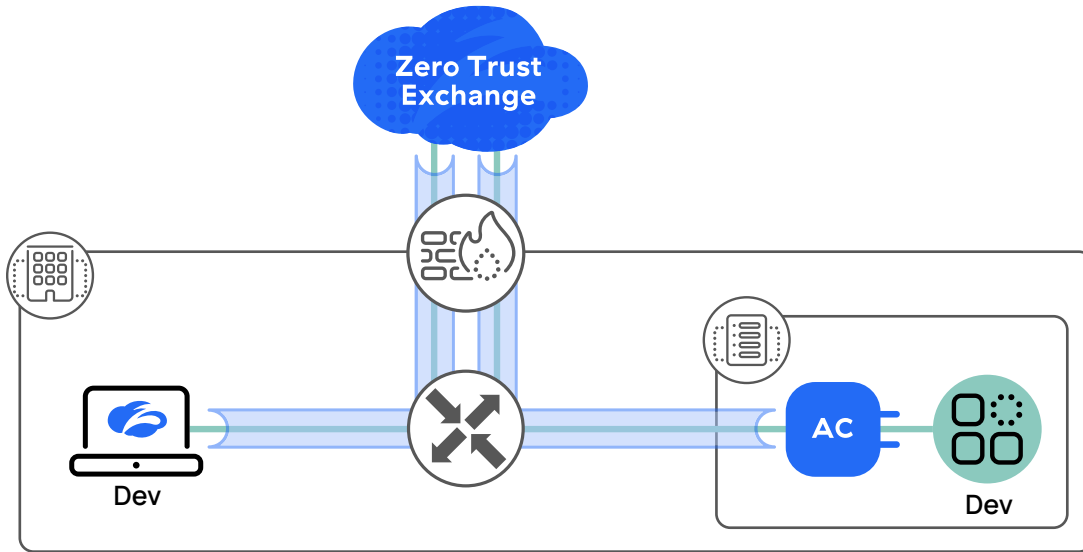


Figure 3. Cloud-based ZTNA can introduce network latency due to the trombone pattern of application access

Your network can be reduced to simply providing access, allowing you to move away from painful network segmentation and limitations of legacy firewalls. However, leveraging cloud-based ZTNA to access applications residing in your local data center can be inefficient. Sending your traffic to the cloud and back can introduce application latency.

Zscaler supports moving the Zscaler application broker services into your organization's demilitarized zone (DMZ) with a ZPA Private Service Edge. Your users attempting to reach local or remote resources will be brokered by your ZPA Private Service Edge, which connects your users to applications in your data center or in the cloud. There is no need to run two different security policies for local and remote users, keeping access controls consistent.

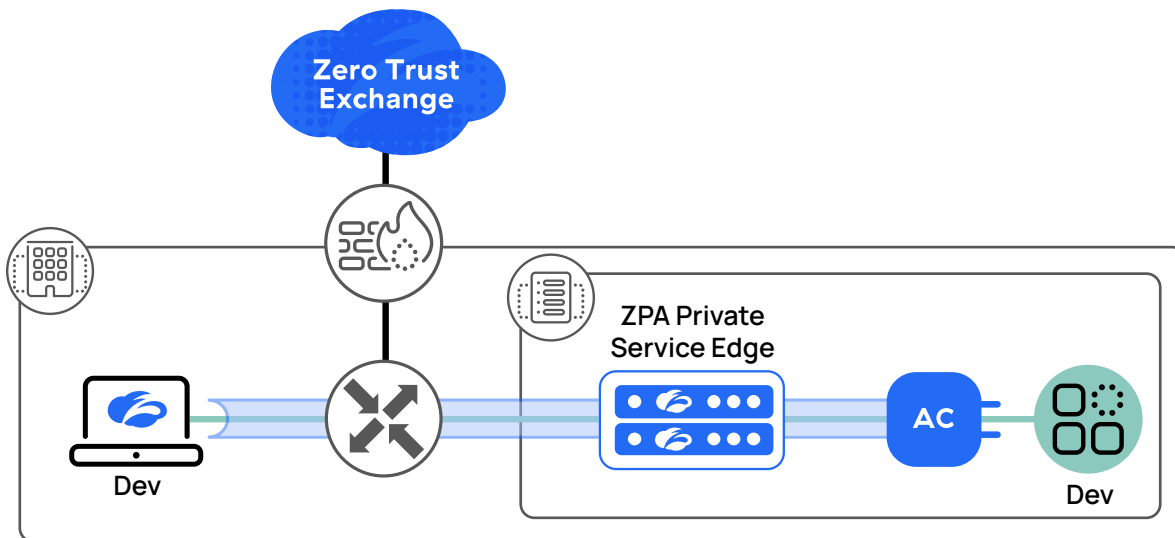


Figure 4. Leveraging your ZPA Private Service Edge allows local traffic to receive the same inspection and policy control without the latency

With your ZPA Private Service Edge deployed, your organization's local users access local apps without the need to first go to the internet. This appliance is deployed by you and managed by Zscaler as part of the Zero Trust Exchange (ZTE). When deployed, your ZPA Private Service Edge is dedicated to your organization, providing fast access to your organization's applications for your organization only. Your ZPA Private Service Edge can also serve your remote users when it is the geographically closest ZPA broker to them.

When is a ZPA Private Service Edge right for your organization? The most common deployment use cases are as follows:

- [On-Premises Users Accessing Applications on the Internal Network](#)
- [Branch Users Accessing Applications Within the Branch Location](#)
- [Remote Users Accessing Regional Internal Applications](#)
- [Remote Users Accessing Regional Internal Applications in a Public Cloud](#)
- [Remote and On-Premises Users Accessing Internal Applications](#)

In this guide, we cover the deployment considerations and best practices for deploying a ZPA Private Service Edge appliance pair. This will enable you to deploy ZTNA for campus data center applications. We also cover all three use cases: remote access, on-premises access, and universal ZTNA.

Key Features and Benefits

- Verify identity – Instead of trusting an IP address, establish the identity of the user and device using an identity provider (IdP) first.
- Set contextual policies – Define access policies based on user, device posture, location, and apps, and rely on a cloud service to enforce them.
- Improve visibility and adapt – Use logs to determine which users are accessing which apps, and automatically adapt based on any changes in context.
- Peerless security, beyond legacy VPNs and firewalls – Users connect directly to apps not the network, minimizing the attack surface and eliminating lateral movement.
- The end of private app compromise – First-of-its-kind app protection with inline prevention, deception, and threat isolation minimizes the risk of compromised users.
- Superior productivity for today's hybrid workforce – Lightning-fast access to private apps extends seamlessly across remote users, HQ, branch offices, and third-party partners.
- Unified ZTNA platform for users, workloads, and OT/IoT – Securely connect to private apps, services, and OT/IoT devices with the industry's most comprehensive ZTNA platform.

New to ZTNA or ZPA?

If you've just started on your journey to ZTNA or haven't learned about ZPA yet, we encourage you to explore the following links to make the most of this guide.

- If you are new to ZPA or the concept of zero trust network access, go to the Zscaler website for a video overview at [Cloud-Delivered Zero Trust Network Access](https://www.zscaler.com/capabilities/zero-trust-network-access) (<https://www.zscaler.com/capabilities/zero-trust-network-access>).
- If you are just learning about ZPA as an alternative to VPNs and network segmentation, go to the [ZPA product page](https://www.zscaler.com/products/zscaler-private-access) (<https://www.zscaler.com/products/zscaler-private-access>).
- View the [ZPA Private Service Edge data sheet](https://www.zscaler.com/resources/data-sheets/zpa-private-service-edge.pdf) (<https://www.zscaler.com/resources/data-sheets/zpa-private-service-edge.pdf>).

Understanding ZPA Private Service Edge Deployments

Your ZPA Private Service Edge extends the ZPA session broker functionality of our cloud into your data center locations hosting your private applications, whether on-premises or in the cloud. This software appliance is available as a lightweight virtual machine and as an RPM package. Your operations team will deploy and manage the software at operating system level. The software and updates to the ZPA Private Service Edge software and services are managed by Zscaler as a part of our global cloud. Unlike a ZPA Public Service Edge in the cloud, your ZPA Private Service Edge is only for use by your organization. By deploying your ZPA Private Service Edge device, your local and remote users can connect to your private data center applications.

Where we most often see a need for a ZPA Private Service Edge are in cases where private applications and users sit physically near one another. In these types of use cases, we don't want to send user traffic to the Zscaler Zero Trust Exchange (ZTE), only to have their data hairpin back to the same location it originated from. A ZPA Private Service Edge keeps the application local to the organization while providing external access to mobile users.

Regulations, data governance models, and software licensing can also cause applications to be less available than a traditional SaaS model allows. An application for patient health records, for example, might need to be stored and accessed locally due to the software license for the application. Regulations and the organization's data governance model might require that applications and data only reside in controlled facilities.

Common Use Cases for a ZPA Private Service Edge

For most organizations and applications, using ZPA Public Service Edge devices in the ZTE is the right choice and can satisfy most use cases. However, when users and internal applications are in close proximity, or when applications are extremely data intensive, a local ZPA Private Service Edge is often faster and provides a better user experience.

In this guide, we examine 5 use cases:

- [On-Premises Users Accessing Applications on the Internal Network](#)
- [Branch Users Accessing Applications Within the Branch Location](#)
- [Remote Users Accessing Regional Internal Applications](#)
- [Remote Users Accessing Regional Internal Applications in a Public Cloud](#)
- [Remote and On-Premises Users Accessing Internal Applications](#)

This guide covers best practice recommendations common to all deployment models and addresses specific recommendations for the use cases listed.



This guide is written for the production commercial ZPA cloud, and uses references to that cloud (e.g., *prod.zpath.net*) throughout. These references are different for tenants on other ZPA clouds, such as Preview or Government clouds (e.g., if an organization is using the ZPA Preview environment, the domain is *zpapreview.net* instead of *prod.zpath.net*). To learn which clouds are available, go to [Zscaler Config \(https://config.zscaler.com/private.zscaler.com/zscaler-app\)](https://config.zscaler.com/private.zscaler.com/zscaler-app).

ZPA Private Service Edges might be included in the ZPA edition you purchased, or you might need to purchase a license for the number of ZPA Private Service Edges you require. As of November 2021, every ZPA organization with a ZPA Business or higher subscription license receives one pair of ZPA Private Service Edges at no cost and additional pairs based on the quantity of ZPA seats purchased. ZPA Private Service Edges can be purchased as an add-on subscription to your service if you need more capacity.



ZPA Private Service Edge functionality is disabled by default in the ZPA tenant and must be enabled by Zscaler. Contact your Zscaler account team or Zscaler Support for assistance.

Deployment Requirements and Scalability

A ZPA Private Service Edge is a fully functional single-tenant instance that provides the functionality of a ZPA Public Service Edge inside your organization's environment. Your organization hosts the Private Service Edge either within your site or on a cloud service, and manages the operating system for that instance. Zscaler manages the Private Service Edge software packages.

ZPA Private Service Edges are deployed in pairs at a minimum, providing N+1 redundancy and resilience. As with a ZPA Public Service Edge, a ZPA Private Service Edge manages the connections between Zscaler Client Connector and ZPA App Connectors. Each ZPA Private Service Edge registers with the ZPA Cloud, which enables it to download relevant policies and configurations so it can enforce all ZPA policies for your ZPA tenant.

ZPA Private Service Edges can be deployed either as a lightweight virtual machine image or via RPM packages on Linux. Zscaler distributes images for deployment in enterprise data centers and local private cloud environments such as VMware, as well as public cloud providers such as AWS and Azure. The requirements for virtual machine or Linux server deployments are as follows:

Component	Size	Notes
Memory	8 GB	This is the recommendation for 2,000 users. Zscaler recommends an additional 1 GB of memory for every 1,000 users that connect to your ZPA Private Service Edge.
CPU	4 CPU cores for physical machines 4 CPU cores for virtual machines (VMs)	
Disk Space	16 GB	Thin provisioned for all deployment types
Network Interface	1 NIC	Minimum requirement
Bandwidth	500 Mbps max throughput	Deploy additional ZPA Private Service Edges to increase bandwidth

For a complete list of detailed deployment guides, go to [Private Service Edge Deployment Guides for Supported Platforms](https://help.zscaler.com/zpa/private-service-edge-management/private-service-edge-deployment-guides-supported-platforms) (<https://help.zscaler.com/zpa/private-service-edge-management/private-service-edge-deployment-guides-supported-platforms>).

Software Updates and Shared Responsibility Model

A deployed ZPA Private Service Edge uses a shared responsibility model for deployments. The host platforms are built on unmodified Linux images (currently CentOS 7.2), with default credentials and the minimum set of packages required for the system to function.

Since vulnerabilities are regularly found in core open-source components such as DNS resolvers and the Linux Kernel, Zscaler recommends either patching or using new Zscaler-distributed VM images on a regular basis, or protecting ZPA Private Service Edges using firewall policies. Additionally if you've installed a ZPA Private Service Edge as a package, Zscaler recommends that you take similar precautions.

Understanding IP Addressing and Public Accessibility

Your ZPA Private Service Edge uses multiple IP addresses for operation and management. As you develop your design, the number and types of IPs required depend on how you deploy your network.

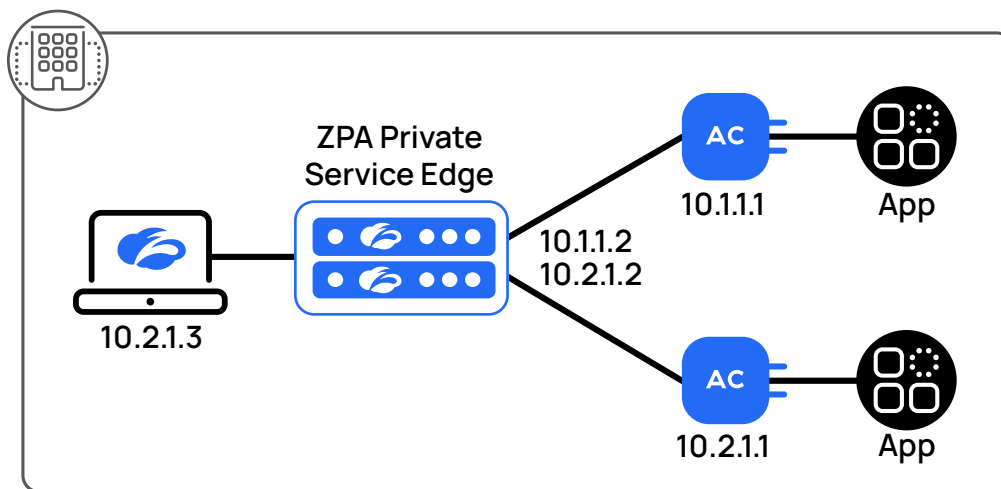


Figure 5. Your ZPA Private Service Edge leverages Publish and Listen IP addresses for Zscaler Client Connector to connect to

Listen IPs

By default, every available interface on your ZPA Private Service Edge is eligible for accepting connections from Zscaler Client Connector and App Connectors. If only a subset of ZPA Private Service Edge interfaces need to be used for accepting incoming connections, the IP addresses of such interfaces should be specified as Listen IPs. In most cases, there won't be a need to specify Listen IPs.

Publish IPs or Domains

Publish IP addresses and domains specify the addresses that your ZPA Private Service Edge listens on for incoming connections from Zscaler Client Connector or App Connectors. If these are not specified, then Zscaler Client Connector and the App Connectors try to connect on the Listen IPs.

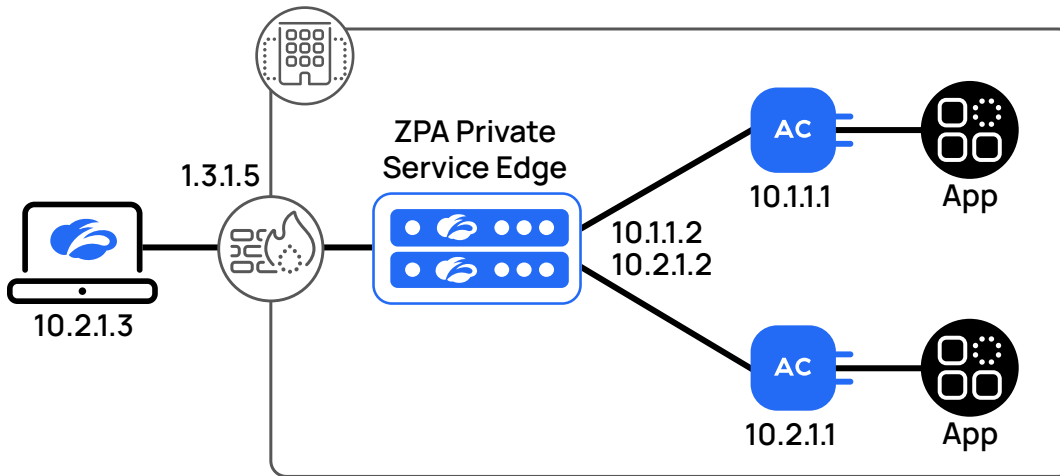


Figure 6. Private Service Edge Publish IPs

Firewalls in Front of ZPA Private Service Edge Deployment

In a scenario where a ZPA Private Service Edge is deployed behind a firewall, the firewall performs destination network address translation (DNAT) for the ZPA Private Service Edge's private IP. The firewall advertises a public IP on the internet. In this scenario, it is necessary to configure the public IP (1.3.1.5 in the previous image) advertised by the firewall as a Publish IP of the respective ZPA Private Service Edge.

Publicly Accessible

A Service Edge Group must be configured as Publicly Accessible for a deployment scenario where a user needs to access your ZPA Private Service Edge on its public IP address. The Publicly Accessible setting is under Service Edge Group and must be enabled.

Confirm Your ZPA Private Service Edge Deployment

Your operations team should set up monitoring for your instance to ensure that you are not exceeding the platform capacity. Zscaler provides both dashboard health monitoring and command line health checks for your ZPA Private Service Edge. Dashboards are ideal for your network operations team in a NOC. The dashboard provides the following information:

- Name – The name of your ZPA Private Service Edge.
- Last Updated – The timestamp showing the last time the ZPA cloud checked the health status.
- Public IP – The public IP address of your ZPA Private Service Edge.
- Private IP – The private IP address of your ZPA Private Service Edge.
- CPU Utilization – The CPU usage of your ZPA Private Service Edge.
- Memory Utilization – The memory usage of your ZPA Private Service Edge.
- Up Time – How long your ZPA Private Service Edge has been enrolled and running.

The command line instructions are similar in nature, but here your engineers interact with the servers using `systemd` on your ZPA Private Service Edge command line. This option is useful if your staff is more comfortable on the command line or is looking at automation to parse responses.

After you confirm your ZPA Private Service Edge is running, it functions exactly like a ZPA Public Service Edge. For on-premises users, or even remote users in countries where there is no ZPA cloud service, access to private applications is brokered through ZPA Private Service Edge for seamless, fast, and secure connectivity. Your operations team should monitor your ZPA Private Service Edge deployments to ensure they are properly sized and operating as expected.

- You can learn more about monitoring options at [Check ZPA Private Service Edge status](https://help.zscaler.com/zpa/managing-deployed-service-edges#Status) (<https://help.zscaler.com/zpa/managing-deployed-service-edges#Status>).
- Learn more about [Monitoring Private Service Edge Deployments](https://help.zscaler.com/zpa/monitoring-private-service-edge-performance) (<https://help.zscaler.com/zpa/monitoring-private-service-edge-performance>).
- If you are new to Linux, learn more about [systemd](https://en.wikipedia.org/wiki/Systemd) (<https://en.wikipedia.org/wiki/Systemd>).

ZPA Private Service Edge Policy Design

Policy configuration is primarily inherited from your ZPA policy without you needing to specify changes for cloud-based applications. What does change is defining your networks local to your ZPA Private Service Edge devices.

Mapping ZPA Private Service Edges to Trusted Networks

ZPA Private Service Edges support the concept of trusted networks. When specifying a trusted network, you tell the ZPA service that one or more networks segments should be directed to a local ZPA Private Service Edge on the same network. This feature requires Zscaler Client Connector 2.1 or later.

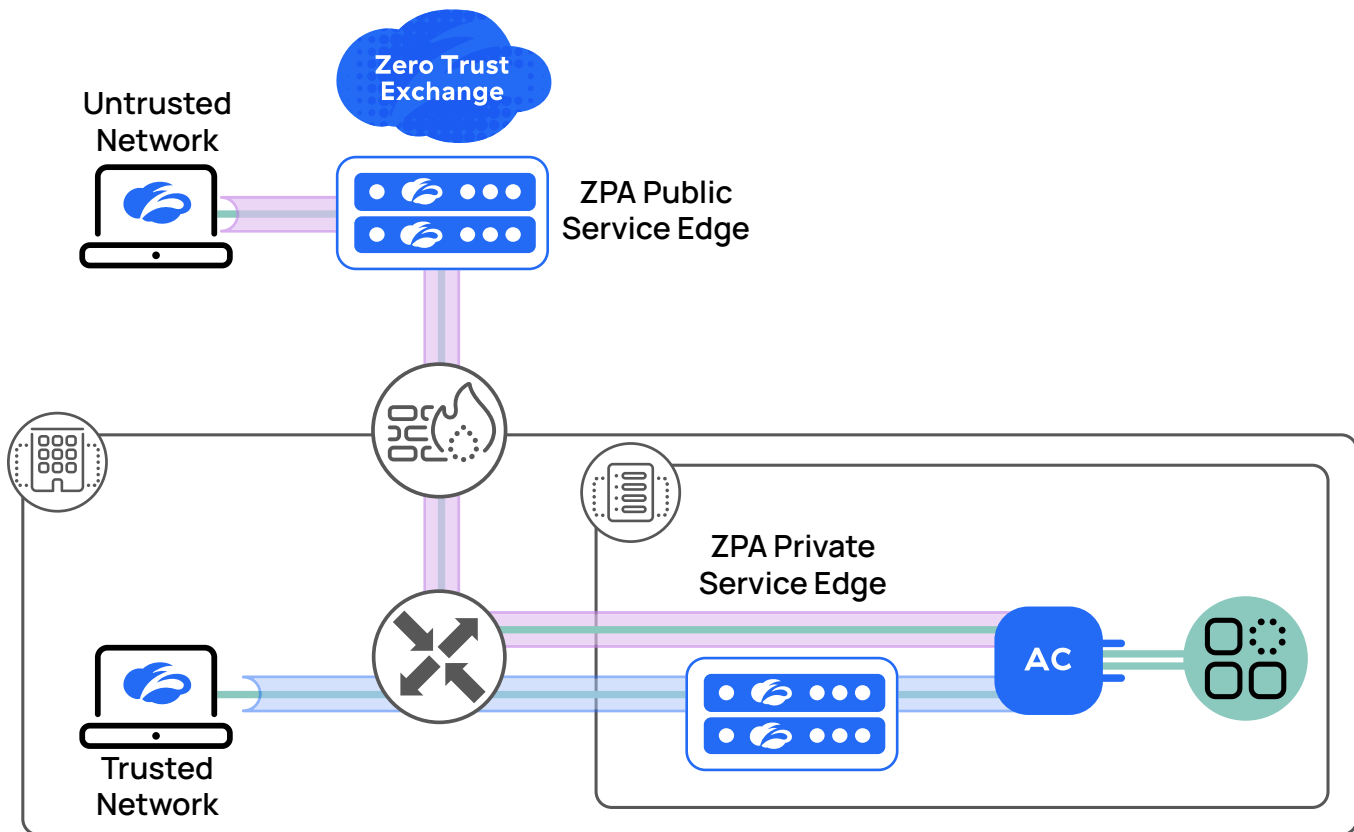


Figure 7. Map your trusted network segments to your local ZPA Private Service Edge

When the ZPA component of Zscaler Client Connector is started, it establishes a client tunnel to a ZPA Public Service Edge. The initial destination FQDN for this connection (*any.broker.prod.zpath.net*) resolves only to ZPA Public Service Edges; Zscaler Client Connector completes mutual TLS authentication with the initial ZPA Public Service Edge. That ZPA Public Service Edge associates Zscaler Client Connector with the appropriate ZPA tenant based on its client certificate, and performs a GeoIP lookup against the connection's source IP address. Zscaler Client Connector sends information about its local network, and the ZPA Public Service Edge determines whether the local network matches any of the defined trusted networks.

Based on this information, the initial ZPA Public Service Edge replies with a redirection message, in which it provides a prioritized list of more optimal Service Edges to connect to. If ZPA Private Service Edges are available, these appear first in the list. If any of those ZPA Private Service Edges match Zscaler Client Connector's current trusted networks, those ZPA Private Service Edges are prioritized; otherwise, ZPA Private Service Edges are prioritized solely by GeoIP. Zscaler Client Connector attempts connections to each Service Edge in the redirection response, in order, until a connection succeeds.

When “Publicly Accessible” is enabled and trusted networks aren’t involved, ZPA treats your ZPA Private Service Edge location just like ZPA Public Service Edge locations. If a ZPA Public Service Edge is closer to the user, the redirect is to that ZPA Public Service Edge; if a ZPA Private Service Edge is closer, the redirect is to that ZPA Private Service Edge. Logically, it behaves exactly as if Zscaler created a new data center with the coordinates of your ZPA Private Service Edge, except that the new data center is exclusive to your organization.



If you subscribe to both ZIA and ZPA, Zscaler Client Connector might reach ZPA by proxying through a ZIA Public Service Edge. This can affect the GeoIP phase of ZPA’s redirection decisions. The GeoIP lookup will be against the egress IP of the ZIA Service Edge, which is unlikely to be the actual location of the client. Bypassing the *any.broker.prod.zpath.net* domain out of ZIA can resolve this problem. Alternatively, trusted network associations can also alleviate this issue based on your network design.

Learn more about [trusted networks](https://help.zscaler.com/z-app/about-trusted-networks) (<https://help.zscaler.com/z-app/about-trusted-networks>).

Restricting User Traffic by Location

When you have multiple ZPA Private Service Edge deployments in different networks, it’s important to ensure that Zscaler Client Connector is connecting to the right ZPA Private Service Edge. This reduces connection delay times for your users.

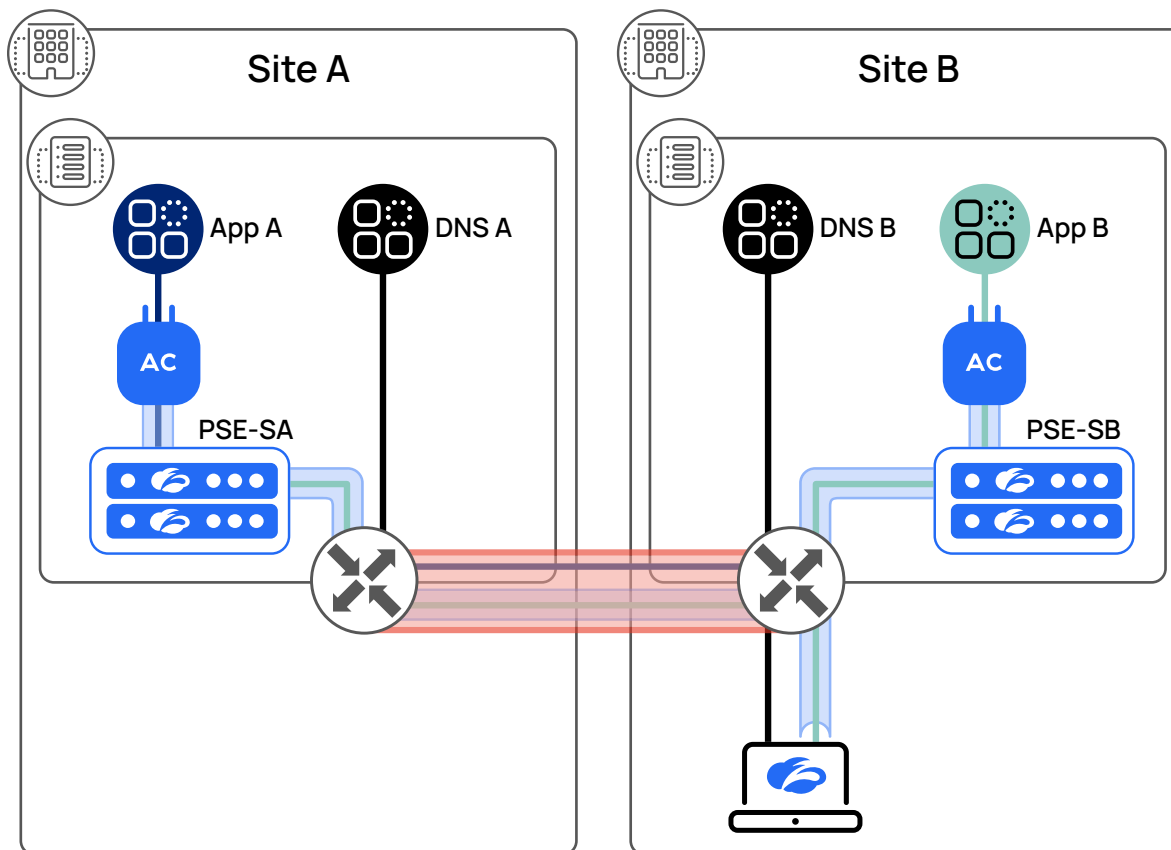


Figure 8. Ensure optimal ZPA Private Service Edge selection for your users with trusted networks

Consider an organization with ZPA Private Service Edges in two independent local networks at different locations, site A and site B as shown in the previous diagram. When a user establishes a connection from site B, ZPA geolocates the user and returns a list of ZPA Private Service Edges. Without specifying trusted networks, the user might be geolocated to a less optimal ZPA Private Service Edge instance nearby.

Defining trusted networks and ZPA Private Service Edge associates is recommended to optimize performance in these types of situations. The network's defined policy can determine which ZPA Private Service Edge is best suited to serve the request.

Here is a high-level configuration example to illustrate how to use trusted networks and Client Forwarding Policies to accommodate these scenarios. Continuing with our example, the organization has two independent sites A and B, each with its own DNS server, both serving the *safemarch.com* domain. The organization should configure three trusted network definitions:

1. Sitewide trusted network – DNS search suffix safemarch.com
2. Site A trusted network – DNS server A
3. Site B trusted network – DNS server B

When Zscaler Client Connector evaluates trusted networks, it has matching conditions for trusted networks 1 and 2 when the user is on site at site A, and for trusted networks 1 and 3 when the user is on site at site B. To associate users with the appropriate ZPA Private Service Edges, organizations should map their ZPA Private Service Edges at site A to trusted network 2, and map their ZPA Private Service Edges at site B to trusted network 3. No ZPA Private Service Edges should be mapped to trusted network 1. This also enables you to determine if your users are on or off your trusted network by checking trusted network 1 for other sites without a ZPA Private Service Edge present.

Client Forwarding Policies

On the client side, the forwarding profile tells Zscaler Client Connector how to treat traffic from your users' systems in different network environments. By [configuring trusted networks \(https://help.zscaler.com/z-app/configuring-trusted-networks-zscaler-app\)](https://help.zscaler.com/z-app/configuring-trusted-networks-zscaler-app), you can leverage the option in forwarding profiles to connect users to ZPA Private Service Edges when users are on private networks. Trusted network configuration can be accomplished by defining conditional criteria that verifies a client is connected to that specific network. For example, you can specify the specific DNS server as referenced previously and/or DNS Search Domains and/or Hostname and IP to resolve.

Grouping App Segments or Segment Groups by Location

Many organizations start their ZPA journey with wildcard app segments or large IP subnets associated with multiple App Connector groups in disparate locations, as well as broad connectivity policies aimed to replicate on-premises or open remote-access VPN access. As the solution matures, more granular definitions and policy are refined over time. The application discovery provides the visibility required to understand clearly which users connect to what applications and where those applications are located.

ZPA App Connectors form their data plane tunnels to the user's Private Service Edge. If you have users and applications close to your ZPA Private Service Edge, then the local App Connectors will provide the best experience. If you have file servers in multiple regions, creating specific file server app segments for each region and assigning those app segments to local or regional App Connectors will help to ensure that a user on a local network connects to the file servers in their location by the closest App Connectors. This happens automatically since the dynamic path selection process considers App Connector proximity to the requesting user when identifying the optimal path. However, if broad app segment definitions are associated with a wide variety of App Connector groups (as often happens during discovery), unnecessary health checks can be generated, or health checks won't occur if they exceed the capacity of a given App Connector group. This can impede the ability of the system to determine the optimal path.

As your zero trust policy matures, defining application segments and/or segment groups by physical and/or logical location reduces health-check overhead and makes the ZPA dynamic path selection more efficient at establishing the optimal routing for user traffic to back-end applications. This is generally true of ZPA overall and can be particularly important in ZPA Private Service Edge use cases, where a large number of users might be accessing a broad range of applications via an organization's internal network.

Understanding Control Plane Connections

In this section, we discuss how Zscaler Client Connector and App Connectors establish control plane connections to your ZPA Private Service Edge. This process, called enrollment, is how ZPA confirms that the various components belong to the same organization.

Enrollment Overview – Private Service Edge

Following initial startup, the ZPA Private Service Edge uses DNS to resolve *pb2br.prod.zpath.net*. The DNS response returns a set of ZPA Public Service Edges that are nearest to the ZPA Private Service Edge according to GeoIP lookup. Your ZPA Private Service Edge then connects to one of these ZPA Public Service Edge instances.



Your firewalls need to be configured to let the ZPA Private Service Edge make outbound connections to the ZPA Public Service Edge IP addresses. You can find a complete list of Zscaler IPs at [Cloud Enforcement Node Ranges](https://config.zscaler.com/zscaler.net/cenr) (<https://config.zscaler.com/zscaler.net/cenr>).

When your ZPA Private Service Edge connects to the ZPA service, it presents its provisioning key. The ZPA service validates the ZPA Private Service Edge. When accepted, the ZPA Private Service Edge generates a certificate signing request (CSR) that is signed and returned by the ZPA service. At the end of this process, the ZPA Private Service Edge is authorized to communicate with the ZPA Cloud.

Enrollment Overview – App Connector

The App Connector's control channel and config channel are directed to an optimal ZPA Public Service Edge during setup. After the connection is established to the ZPA Public Service Edge, the App Connector requests your ZPA Private Service Edge configuration for the tenant. The App Connector receives the full list of ZPA Private Service Edges and their published listening addresses. Next, the App Connector attempts to establish connections to all ZPA Private Service Edges. If the connection is successful, the App Connector identifies it as an available, accessible ZPA Private Service Edge. The App Connectors continue this process to build a matrix of available ZPA Private Service Edges.

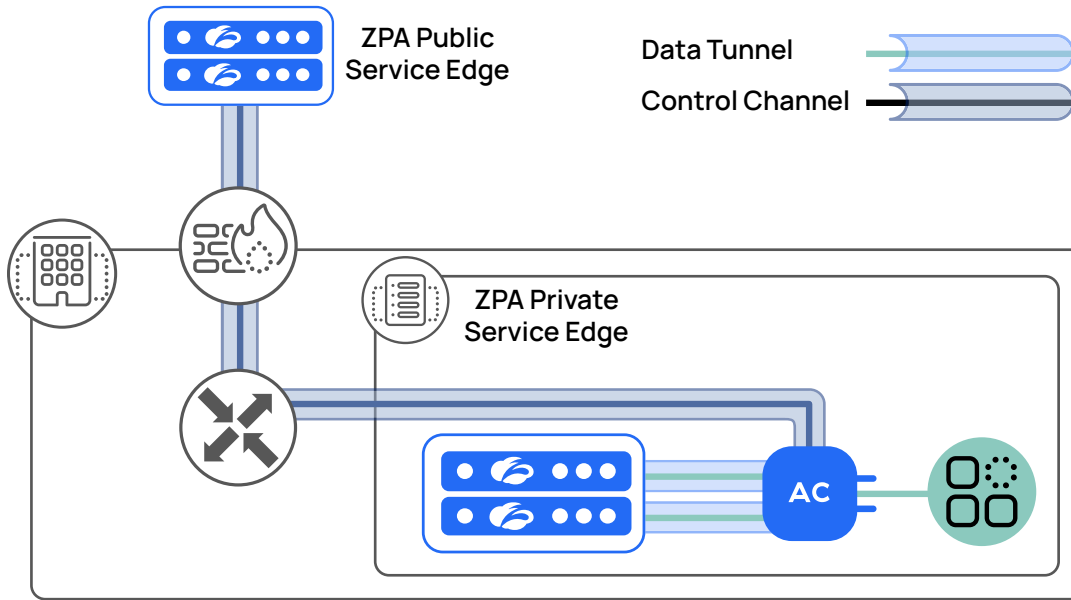


Figure 9. App Connector enrollment

When this process is complete, the App Connector maintains one control channel to a ZPA Public Service Edge for communication with the ZPA Cloud, known as the control channel, and additional control channels to all ZPA Private Service Edges it can reach, known as private control channels.

Learn more about [App Connectors](https://help.zscaler.com/zpa/about-connectors) (<https://help.zscaler.com/zpa/about-connectors>).

Enrollment Overview – Zscaler Client Connector

After enrollment, Zscaler Client Connector establishes one TLS connection for ZPA, called a client tunnel. When ZPA is started, a client tunnel is established to the nearest ZPA Public Service Edge. These tunnels are the only outbound connections established by Zscaler Client Connector for ZPA. Both control and data channels are established to the same Service Edge. Zscaler Client Connector is never simultaneously connected to multiple Services Edges.

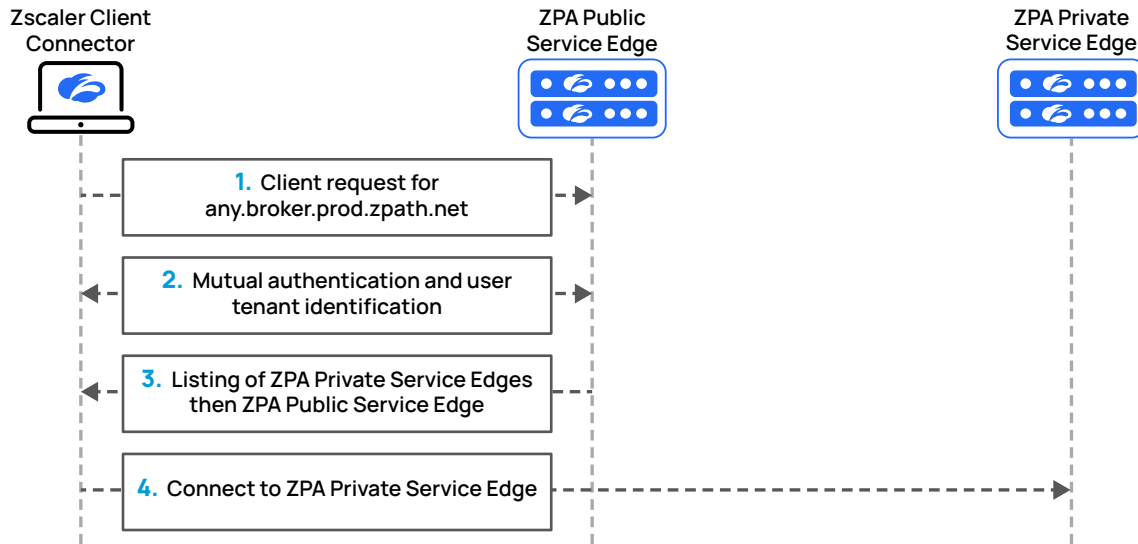


Figure 10. Zscaler Client Connector enrollment

1. The initial destination FQDN for this connection resolves to a ZPA Public Service Edge using the URL *any.broker.prod.zpath.net*.
2. During the mutual certificate validation by the ZPA Public Service Edge and Zscaler Client Connector, the Zscaler Client Connector certificate identifies the specific organization tenant, and thus the constellation for the organization, including ZPA Private and ZPA Public Service Edges. Zscaler Client Connector also sends information about its local network, and tells the ZPA Public Service Edge whether that network matches any defined trusted networks.
3. The ZPA Public Service Edge geolocates the Zscaler Client Connector location, and returns an ordered list of connection nodes. Two ZPA Private Service Edges matching the trusted network are set first, followed by the two Private Service Edges prioritized by GeoIP. Finally, the two nearest ZPA Public Service Edges are listed. This list is sent to Zscaler Client Connector via a redirection message.
4. Following this message, Zscaler Client Connector breaks the TCP connection in place with the ZPA Public Service Edge and begins to attempt connections on the entries in the redirection list (with ZPA Private Service Edges prioritized).

Learn more at [What Is Zscaler Client Connector?](https://help.zscaler.com/z-app/what-zscaler-app) (<https://help.zscaler.com/z-app/what-zscaler-app>).

Understanding Traffic Forwarding

After your ZPA Private Service Edge, Zscaler App Connectors, and Zscaler Client Connector are all enrolled, you are ready to begin forwarding traffic to the requested application. In this section, we discuss how your ZPA Private Service Edge, Zscaler Client Connector, and App Connector forward traffic.

Traffic Forwarding – ZPA Private Service Edge

Using the Certificate Authority configured within the ZPA Admin Portal, your ZPA Private Service Edge verifies that it shares the same tenant with Zscaler Client Connector that is attempting to connect. Zscaler Client Connector also has its certificate checked to ensure it is also a valid client of your organization.

Your ZPA Private Service Edge then establishes a separate TLS connection to the nearest ZPA Public Service Edge. This tunnel is used for traffic when the user requests an application that is not available by a directly connected App Connector. Your ZPA Private Service Edge uses this channel to request the application by the ZPA Public Service Edge network. These data connections are never shared among different clients. Even if multiple devices with Zscaler Client Connector have authenticated with the same user account, the data channels remain separate.

When a user requests an internal application, Zscaler Client Connector intercepts the application request and forwards it to your ZPA Private Service Edge. Your ZPA Private Service Edge matches the application request against the Access Policies, and if allowed, it continues to handle the request.

The next step is to determine if a local App Connector can broker the connection, or if it needs the help of ZPA Public Service Edges. ZPA Private Service Edges do not communicate with each other, only downstream to local App Connectors and Zscaler Client Connector, or upstream to the ZPA Public Service Edge.

Traffic Forwarding – App Connector

In addition to the control channels, the App Connector establishes a separate TLS-protected data connection to either the ZPA Public Service Edge or ZPA Private Service Edge to serve application data. An App Connector can establish multiple TLS connections to each ZPA Private Service Edge. This includes one private control channel, and one data channel to pass application traffic. In addition, the App Connectors maintain private control channels with every Private Service Edge that belongs to the same tenant and that it can reach over the network. They also establish and maintain its control channel with only one single Public Service Edge.

Upon startup, each App Connector downloads the set of applications it is responsible to serve. The App Connector then advertises its availability, responsibilities, and health check status to the public cloud via its control channel, as well as any ZPA Private Service Edges it can reach via direct connections. This data is tracked by the ZPA service. ZPA Private Service Edges track this data locally, so this information can be maintained without reliance on the public ZPA Cloud. App Connectors continually report application health status.

When your ZPA Private Service Edge begins handling a connection request, it reviews the most recent health check data that it received from all local App Connectors. This discovery data is passed through the private control channels.

If a directly connected App Connector can reach the requested application, one of them is selected to serve the request in the same way as it would through the ZPA Public Service Edge. If the app is available locally, a locally attached App Connector establishes a data connection with Zscaler Client Connector on the user's device.

As previously discussed, during initial authentication, your ZPA Private Service Edge establishes a TLS connection to a ZPA Public Service Edge, reserved for the specific client. Your ZPA Private Service Edge uses this connection to forward the client's application request to the ZPA Public Service Edge, if no directly connected App Connectors are available to serve the request. The ZPA Public Service Edge re-evaluates access policy to ensure the connection is authorized, even though it is already verified by your ZPA Private Service Edge. The ZPA Public Service Edge brokers the connection by associating the App Connector's data channel with your ZPA Private Service Edge's data channel. The App Connector establishes its data channel to the ZPA Public Service Edge. The ZPA Public Service Edge brokers the connection from your ZPA Private Service Edge, and your ZPA Private Service Edge can now broker the connection back to Zscaler Client Connector.

Learn more about [App Connectors](https://help.zscaler.com/zpa/about-connectors) (<https://help.zscaler.com/zpa/about-connectors>).

Traffic Forwarding – Zscaler Client Connector

Functionally, traffic handling with Zscaler Client Connector is identical to a deployment without ZPA Private Service Edges. In a public-only environment, the “client” would be Zscaler Client Connector. When the traffic must pass through your ZPA Private Service Edge to a ZPA Public Service Edge, the role of the “client” is played by your ZPA Private Service Edge.

Learn more at [What Is Zscaler Client Connector?](https://help.zscaler.com/z-app/what-zscaler-app) (<https://help.zscaler.com/z-app/what-zscaler-app>).

Recommendations for Gradual Migration to Your ZPA Private Service Edge

Transitioning from traditional network-centric security controls to on-premises zero trust can be a significant change for most organizations, and the best approach is to break it into phases. Typically, we see the most success using one of the following models:

1. [Migrating to Your ZPA Private Service Edge One App at a Time](#)
2. [Migrating to Your ZPA Private Service Edge One Group of Users at a Time](#)
3. [Combined Approach](#)

Each of these approaches can be applied using a Crawl-Walk-Run phased transition. For greatest control, these approaches can be combined starting with an initial pilot group of users, gradually migrating them to using all applications via ZPA, and then expanding the pilot to subsequent groups of users. We walk through each of these use cases in detail to discuss considerations and caveats.

Organizations that already use ZPA for remote access often enter the migration process with a Zscaler Client Connector Forwarding Profile that has ZPA set to Tunnel (enabled) when the user is off a trusted network and set to None (disabled) when the user is on a trusted network. This allows on-premises users to go directly to private application resources without going to a ZPA Public Service Edge first.

In this guide, it is expected that app segments have been defined for destination applications, and access policies have been created to allow access by authorized users prior to the introduction of ZPA Private Service Edge into your network. If you do not yet have those policies in place, we recommend reviewing our reference architecture [Zero Trust User-to-App Segmentation with ZPA](https://www.zscaler.com/resources/reference-architectures/zero-trust-user-to-app-segmentation-zpa.pdf) (<https://www.zscaler.com/resources/reference-architectures/zero-trust-user-to-app-segmentation-zpa.pdf>).

Migrating to Your ZPA Private Service Edge One App at a Time

In this phased transition, organizations often start with less sensitive applications to ensure connectivity and build confidence in the on-premises solution, then extend ZPA protection to the most sensitive applications.

Crawl

The first step in app-based migration to your ZPA Private Service Edge is to define a ZPA Client Forwarding Policy rule to bypass ZPA for all applications when the user is on a trusted network. This allows you to create a new Zscaler Client Connector Forwarding Profile that has ZPA set to Tunnel (enabled) when the user is on the trusted network, without changing the actual access path of the users. The user can still go directly to private applications when on-premises, but now ZPA is enabled, although it is not yet forwarding application traffic.

Walk

The next step is to start defining Client Forwarding Policy rules to forward specific applications to ZPA when the user is on a trusted network. These application-specific rules must be placed above the bypass rule. Generally, best practice is to start with simple and/or low-sensitivity applications, such as an internal website or file server. Test with users to ensure that access works as expected before incorporating more complex or critical applications into ZPA. Direct access is usually still permitted in this phase for all other apps. This allows you to easily remove a rule if you find it is impacting users and needs to be adjusted to ensure access.

Run

As access shifts to ZPA for each application or group of applications, direct access to those applications can be removed, so that the only available access path is through ZPA. When all application traffic is going through ZPA, the Client Forwarding Policy rule to bypass all applications can be modified to only forward allowed applications, since it will no longer be necessary to bypass any traffic. This also allows you to simplify your policy by removing explicit forwarding rules for individual applications, if desired. At this point, your ZPA Private Service Edge should broker all private application traffic for on-network users.

Migrating to Your ZPA Private Service Edge One Group of Users at a Time

In this phased transition, the gradual migration is by groups of users, rather than by groups of applications. Organizations often start with a pilot group, then extend access to additional groups after a comfort level is established. You might want to create a unique new group membership to identify users subject to on-network ZPA forwarding, so access can be expanded simply by adding more users to that group, instead of adding new groups to existing Forwarding Profiles/Client Forwarding Policy rules.

Crawl

The first step in group-based migration to your ZPA Private Service Edge is to identify or create a unique group membership for pilot users and ensure that it is available as a SAML attribute or SCIM group. Generally, best practice is to start with a small group whose access needs are relatively well understood. Define a ZPA Client Forwarding Policy rule to bypass ZPA for all applications, for that group only, when the user is on a trusted network. This allows you to create a new Zscaler Client Connector Forwarding Profile that has ZPA set to Tunnel (enabled) when the user is on the trusted network, without changing the actual access path of the users. The user can still go directly to private applications when on-premises, but now ZPA is enabled, although it is not yet forwarding application traffic.

Walk

The next step is to define a Client Forwarding Policy rule to only forward allowed applications to ZPA, for that user group only, when the user is on a trusted network. This forwarding rule must be placed above the bypass rule for the group. Now your ZPA Private Service Edge brokers all private application traffic for that group of users. Direct access is usually still permitted in this phase for all other apps. This allows you to easily remove a rule if you find it is impacting users and needs to be adjusted to ensure access. As each group of users successfully migrates to your ZPA Private Service Edge, ZPA forwarding can be introduced for another group of users.

Run

When access for all users is shifted to ZPA, direct access to private applications can be removed, so that the only available access path is through ZPA. When all private application traffic is going through ZPA, the Client Forwarding Policy rule to bypass all applications can also be removed, since it will no longer be necessary. At this point, your ZPA Private Service Edge should broker all private application traffic for on-network users.

Combined Approach

For environments that want an even more granular migration, these two approaches can be combined. This scenario follows the same top-level phases as the per-user migration, with an exception in the Walk phase: instead of using a Client Forwarding Policy rule to forward *all* authorized applications for the target group, you can create individual Client Forwarding Policies for specific applications as in the per-app migration.



If a ZPA Private Service Edge is reachable by Zscaler Client Connector on the same trusted network, the Zscaler Client Connector enrollment process prefers your ZPA Private Service Edge by default. Currently for on-premises users, it is not possible to have some of the users go through your ZPA Private Service Edge while other users go to a public ZPA Service Edge for the same local application.

Use Cases

In this section, we provide additional recommendations related to the most common use cases we encounter for a ZPA Private Service Edge. The most common deployment use cases we cover are:

- [On-Premises Users Accessing Applications on the Internal Network](#)
- [Branch Users Accessing Applications Within the Branch Location](#)
- [Remote Users Accessing Regional Internal Applications](#)
- [Remote Users Accessing Regional Internal Applications in a Public Cloud](#)
- [Remote and On-Premises Users Accessing Internal Applications](#)

On-Premises Users Accessing Applications on the Internal Network

This deployment scenario is relevant to organizations who are interested in implementing zero trust access for on-premises users. This scenario assumes that internal applications are hosted on-premises or in public cloud instances connected to the internal network via dedicated connections, such as MPLS or site-to-site tunnels. In addition, this scenario is relevant for organizations considering a ZPA deployment where only on-premises users access internal applications using ZPA Private Service Edges, while remote users access applications using ZPA Public Service Edges.

Deploying a ZPA Private Service Edge provides several advantages:

- Customers can implement zero trust access through a consistent policy framework for remote and on-premises users.
- Customers can achieve zero trust access for sensitive internal applications, without sending the application traffic through a ZPA Public Service Edge.
- Customers can reduce application access latency and bandwidth usage to the cloud, as the application traffic stays on the internal network.

For this scenario, Zscaler recommends that ZPA Private Service Edges are hosted in a data center. ZPA Private Service Edges can be set up in each data center or in a hub data center.

Networking Requirements

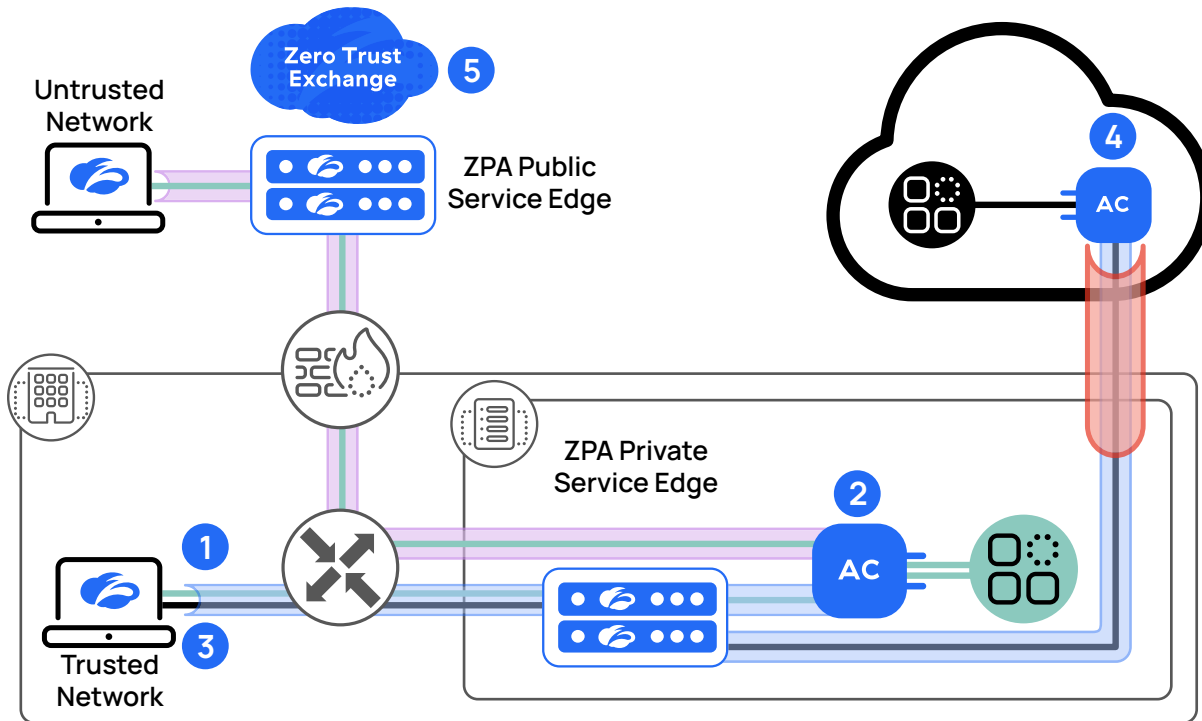


Figure 11. App Connectors with direct connectivity to your ZPA Private Service Edge

1. Each ZPA Private Service Edge must be able to accept incoming connections from internal sources on TCP port 443.
2. Each ZPA Private Service Edge must have a unique private IP address that can be reached over the internal network. This requirement does not apply to your ZPA Private Service Edge that is solely for remote users accessing on-premises applications.
3. For Zscaler Client Connector to connect to your ZPA Private Service Edge, Zscaler Client Connector must be able to reach your ZPA Private Service Edge IP address. If Zscaler Client Connector is unable to connect to any ZPA Private Service Edge, it connects to the closest ZPA Public Service Edge.
4. For Zscaler Client Connector to connect to a Private Service Edge, at least one of the following conditions must be met:
 - a. Your ZPA Private Service Edge is the closest ZPA Service Edge, given the geographic location of the user. The location of the user is determined based on the public IP address of the request from Zscaler Client Connector, as seen by ZPA.
 - b. ZPA Private Service Edge is configured to be on a specific trusted network, and Zscaler Client Connector is reporting itself being on the same trusted network.
5. Any traffic destined to `*.prod.zpath.net` from Zscaler Client Connector must be bypassed from ZIA or other proxies and go directly to the internet. Zscaler Client Connector always resolves `any.broker.prod.zpath.net` to initially connect to a ZPA Public Service Edge. ZPA Public Service Edges can redirect Zscaler Client Connector to your ZPA Private Service Edge based on the conditions in step 4. It is important that ZPA sees the true public IP address of the user for optimal ZPA Private Service Edge selection.

App Connectors should be able to reach your ZPA Private Service Edge's private IP address. When a user requests an application through your ZPA Private Service Edge, ZPA asks an App Connector to set up a data connection to this ZPA Private Service Edge. If Listen IPs are not set for your ZPA Private Service Edge, App Connectors in other data centers connect to your ZPA Private Service Edge's public IP (if available). Zscaler recommends the use of Listen IPs to impose preferred internal connectivity.



If no App Connector for the requested application can establish a data connection to the required ZPA Private Service Edge where the user is connected, that ZPA Private Service Edge relays Zscaler Client Connector's request to a ZPA Public Service Edge, and the App Connector connects to the ZPA Public Service Edge. So, there will be an additional hop through a ZPA Public Service Edge to complete the application request.

Configuration Requirements

In your Service Edge Group, disable the Public Access setting when only internal users will connect to your ZPA Private Service Edge. Specify a trusted network when Zscaler Client Connector is configured to report trusted networks. Specifying trusted networks assists ZPA with redirecting Zscaler Client Connector to your ZPA Private Service Edge on the same network as the user. If the user is closer to a ZPA Public Service Edge and you want the user to connect via your ZPA Private Service Edge instead, specifying a trusted network for your ZPA Private Service Edge is required.

Listen IPs and Publish IPs or Domains are optional but can add greater levels of control. Specify Listen IPs when only certain interfaces on your ZPA Private Service Edge are used for accepting incoming requests. If Listen IPs are specified, ZPA automatically advertises Listen IPs as Publish IPs. If Listen IPs are not specified, ZPA advertises the IPs of every available interface on your ZPA Private Service Edge.

Specify Publish IPs or Domains when the IP or domain that clients must use to reach your ZPA Private Service Edge differs from the available interface IPs. In most cases, Publish IPs or Domains don't need to be specified. If Publish IPs or Domains need to be specified, the resulting traffic should be routed to the private IPs of your ZPA Private Service Edge. Any one of the configured Publish IPs or Domains can be provided to Zscaler Client Connector and App Connectors for connecting to your ZPA Private Service Edge. If Publish Domains are specified, organizations must create corresponding DNS entries in the internal DNS resolvers.

Branch Users Accessing Applications Within the Branch Location

This deployment scenario is relevant to organizations who are interested in implementing zero trust access for on-premises users in a branch location. This deployment model assumes that users can access internal applications over the branch LAN. In addition, this scenario is relevant for organizations considering a ZPA deployment where only on-premises users access internal applications using ZPA Private Service Edges, while remote users access applications using ZPA Public Service Edges.

Deploying a ZPA Private Service Edge provides several advantages:

- Customers can implement zero trust access through a consistent policy framework for remote and on-premises users.
- Customers can achieve zero trust access for sensitive internal applications, without sending the application traffic through a ZPA Public Service Edge.
- Customers can reduce the application access latency, as the application traffic stays on the internal network.

For this scenario, Zscaler recommends that ZPA Private Service Edges are hosted in the Branch location.



If no App Connector for the requested application can establish a data connection to the required ZPA Private Service Edge where the user is connected, that ZPA Private Service Edge relays Zscaler Client Connector's request to a ZPA Public Service Edge, and the App Connector connects to the ZPA Public Service Edge. So, there will be an additional hop through a ZPA Public Service Edge to complete the application request.

Configuration Requirements

In your Service Edge Group, disable the Publicly Accessible setting when only internal users will connect to your ZPA Private Service Edge. Specify a trusted network if Zscaler Client Connector is configured to report trusted networks. Specifying trusted networks assists ZPA with redirecting Zscaler Client Connector to your ZPA Private Service Edge on the same network as the user. If the user is closer to a ZPA Public Service Edge and you want the user to connect via your ZPA Private Service Edge instead, specifying a trusted network for your ZPA Private Service Edge is required.

Listen IPs and Publish IPs or Domains are optional but can add greater levels of control. Specify Listen IPs when only certain interfaces on your ZPA Private Service Edge are used for accepting incoming requests. If Listen IPs are specified, ZPA automatically advertises Listen IPs as Publish IPs. If Listen IPs are not specified, ZPA advertises the IPs of every available interface on your ZPA Private Service Edge.

Specify Publish IPs or Domains when the IP or domain that clients must use to reach your ZPA Private Service Edge differs from the available interface IPs. In most cases, Publish IPs or Domains don't need to be specified. If Publish IPs or Domains need to be specified, the resulting traffic should be routed to the private IPs of your ZPA Private Service Edge. Any one of the configured Publish IPs or Domains can be provided to Zscaler Client Connector and App Connectors for connecting to your ZPA Private Service Edge. If Publish Domains are specified, organizations must create corresponding DNS entries in the internal DNS resolvers, as well as in external DNS if the ZPA Private Service Edge is publicly accessible.

Remote Users Accessing Regional Internal Applications

This deployment scenario is relevant to organizations who are interested in implementing zero trust access for remote users who are not near a ZPA Public Service Edge. This scenario assumes that internal applications are hosted on-premises or in public cloud instances connected to the internal network via dedicated connections, such as MPLS or site-to-site tunnels. This deployment model assumes that App Connectors in public cloud locations not connected to the internal network are able to reach your ZPA Private Service Edges over the public internet. In addition, this scenario is relevant for organizations considering a ZPA deployment where remote users access applications using ZPA Private Service Edges, while on-premises users might or might not send traffic through ZPA.

Deploying a ZPA Private Service Edge provides several advantages:

- Customers can achieve zero trust access for sensitive internal applications, without sending the application traffic through a ZPA Public Service Edge.
- Customers can reduce application access latency and enterprise bandwidth consumption, as the application traffic goes through a ZPA Private Service Edge that is geographically closer compared to a ZPA Public Service Edge.

For this scenario, Zscaler recommends that ZPA Private Service Edges are hosted in a data center. ZPA Private Service Edges can be set up in each data center or in a hub data center.

Networking Requirements

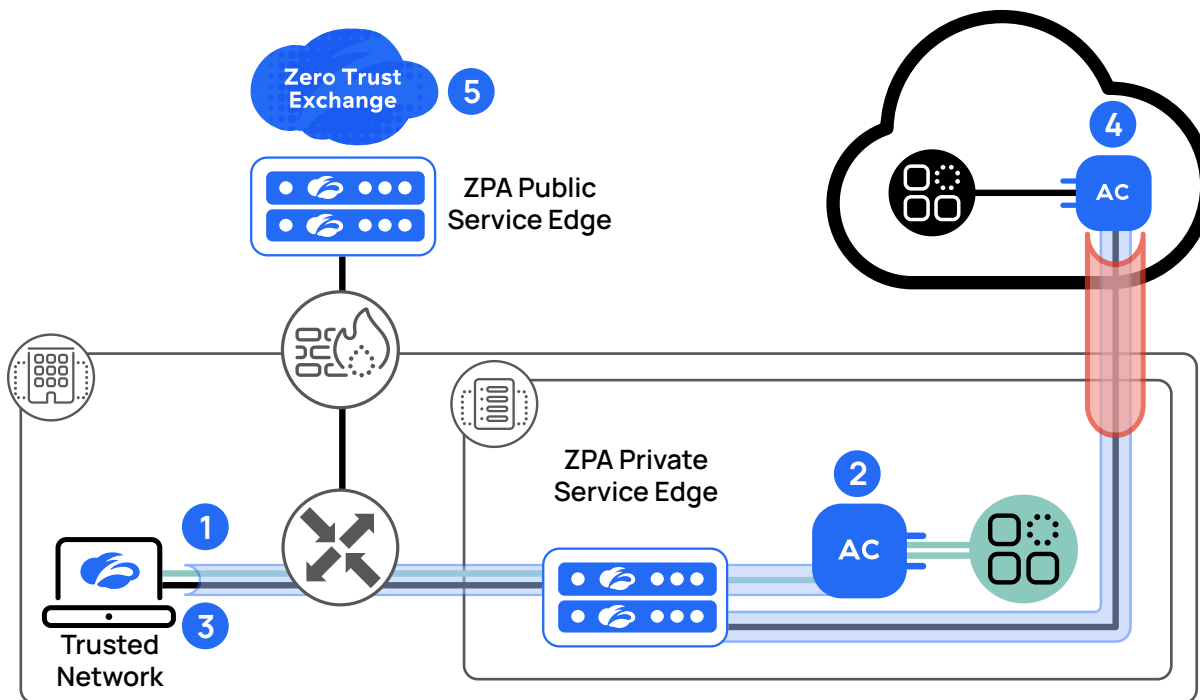


Figure 13. Remote users accessing regional internal applications

1. Each ZPA Private Service Edge must be able to accept incoming connections from external sources on TCP port 443.
2. Each ZPA Private Service Edge must have a unique public IP that can be reached by Zscaler Client Connector over the public internet. This requirement does not apply to your ZPA Private Service Edge that is solely internal facing.
3. For Zscaler Client Connector to connect to your ZPA Private Service Edge, Zscaler Client Connector must be able to reach your ZPA Private Service Edge's public IP address. If Zscaler Client Connector is unable to connect to your ZPA Private Service Edge, it connects to the closest ZPA Public Service Edge.
4. For Zscaler Client Connector to connect to a publicly accessible ZPA Private Service Edge, your ZPA Private Service Edge must be the closest ZPA Service Edge, given the geographic location of the user. The location of the user is determined based on the public IP address of the request from Zscaler Client Connector, as seen by ZPA.
5. Any traffic destined to `*.prod.zpath.net` from Zscaler Client Connector must be bypassed from ZIA or other proxies and go directly to the internet. Zscaler Client Connector always resolves `any.broker.prod.zpath.net` to initially connect to a ZPA Public Service Edge. ZPA Public Service Edges can redirect Zscaler Client Connector to your ZPA Private Service Edge based on the conditions in step 4. It is important that ZPA sees the true public IP address of the user for optimal ZPA Private Service Edge selection.

App Connectors should be able to reach your ZPA Private Service Edge's public or private IP address. When a user requests an application through your ZPA Private Service Edge, ZPA asks an App Connector to set up a data connection to this ZPA Private Service Edge. If Listen IPs are not set for your ZPA Private Service Edge, App Connectors in other data centers connect to your ZPA Private Service Edge's public IP. Zscaler recommends the use of Listen IPs to impose preferred internal connectivity.

If it is less optimal for App Connectors to reach your ZPA Private Service Edge on the public IP address, each ZPA Private Service Edge should have a unique private IP address that can be reached by App Connectors over the internal network (refer to scenario 2 under [Configuration Requirements](#)).



If no App Connector for the requested application can establish a data connection to the required ZPA Private Service Edge where the user is connected, that ZPA Private Service Edge relays Zscaler Client Connector's request to a ZPA Public Service Edge, and the App Connector connects to the ZPA Public Service Edge. So, there will be an additional hop through a ZPA Public Service Edge to complete the application request.

Configuration Requirements

For this use case, the Service Edge Group must have the Publicly Accessible setting enabled. Other settings depend on which scenario you deploy.

Scenario 1: All App Connectors can reach ZPA Private Service Edge's Public IP address.

Listen IPs and Publish IPs or Domains can be configured optionally. Specify Listen IPs when only certain interfaces on your ZPA Private Service Edge are used for accepting incoming requests. If Listen IPs are specified, ZPA automatically advertises Listen IPs as Publish IPs. In this case, the Listen IPs need to be public IPs. If Listen IPs are not specified, ZPA advertises the IPs of every available interface on your ZPA Private Service Edge. In this case, the IPs need to be public IPs.

Specify Publish IPs or Domains when the IP or domain that clients must use to reach your ZPA Private Service Edge differs from the available interface IPs. In most cases, Publish IPs or Domains don't need to be specified. If Publish IPs or Domains need to be specified, the resulting traffic should be routed to the private IPs of your ZPA Private Service Edge. For example, if your ZPA Private Service Edge is deployed behind a firewall, and the firewall is performing DNAT for your ZPA Private IP and advertising a public IP on the internet, the public IP advertised by the firewall must be configured as a Publish IP.

Any one of the configured Publish IPs or Domains can be provided to Zscaler Client Connector and App Connectors for connecting to your ZPA Private Service Edge. If Publish Domains are specified, organizations must create corresponding DNS entries in public DNS resolvers, as well as in internal DNS resolvers if your ZPA Private Service Edge is also accessed by internal users.

Scenario 2: Not all App Connectors can reach a Private Service Edge via a public IP address.

Each ZPA Private Service Edge requires a unique Publish Domain. Customers must specify a FQDN that can resolve on the internet to your ZPA Private Service Edge's public IP address and on the internal network to your ZPA Private Service Edge's private IP address.

If your organization has a ZPA app segment containing a wildcard (such as **.safemarch.com*) that includes the FQDN for your ZPA Private Service Edge, it will cause routing issues. A bypass config is needed, either in the ZPA Client Forwarding Policy or in the Zscaler Client Connector App Profile, so that traffic to your ZPA Private Service Edge is not intercepted as application traffic.

Zscaler Client Connector and App Connectors might resolve to your ZPA Private Service Edge's public or private IP address, depending on the location of the client. Clients can attempt to connect using the Publish Domain.

Remote Users Accessing Regional Internal Applications in a Public Cloud

This deployment scenario is relevant to organizations who are interested in implementing zero trust access for remote users to applications hosted only in public cloud instances. This deployment model assumes that users can access applications over the public internet, so no dedicated connections are required.

Deploying a ZPA Private Service Edge provides several advantages:

- Customers can achieve zero trust access for sensitive internal applications, without sending the application traffic through a ZPA Public Service Edge.
- Customers can reduce the application access latency, as the application traffic goes through a ZPA Private Service Edge that is geographically closer compared to a ZPA Public Service Edge.

For this scenario, Zscaler recommends that ZPA Private Service Edges are hosted in public cloud instances.

Networking Requirements

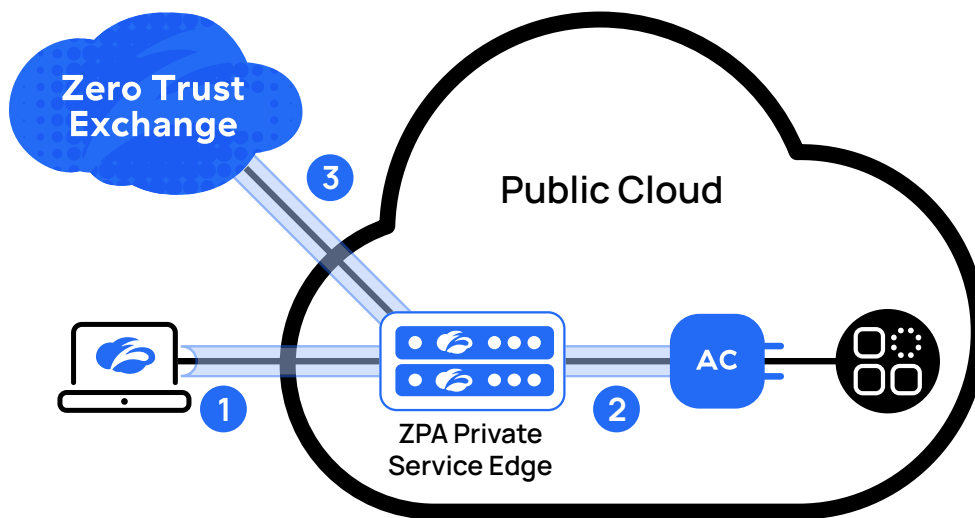


Figure 14. Remote users accessing internal applications via ZPA Private Service Edge in the public cloud

1. Each ZPA Private Service Edge must be able to accept incoming connections from external sources on TCP port 443.
2. Each ZPA Private Service Edge must have a unique public IP that can be reached by Zscaler Client Connector over the public internet.
3. For Zscaler Client Connector to connect to your ZPA Private Service Edge, Zscaler Client Connector must be able to reach your ZPA Private Service Edge's public IP address. If Zscaler Client Connector is unable to connect to your ZPA Private Service Edge, it connects to the closest ZPA Public Service Edge.

For Zscaler Client Connector to connect to a publicly accessible ZPA Private Service Edge, your ZPA Private Service Edge must be the closest ZPA Service Edge, given the geographic location of the user. The location of the user is determined based on the public IP address of the request from Zscaler Client Connector, as seen by ZPA.

Any traffic destined to `*.prod.zpath.net` from Zscaler Client Connector must be bypassed from ZIA or other proxies and go directly to the internet. Zscaler Client Connector always resolves `any.broker.prod.zpath.net` to initially connect to a ZPA Public Service Edge. ZPA Public Service Edges can redirect Zscaler Client Connector to your ZPA Private Service Edge based on the conditions in step 4. It is important that ZPA sees the true public IP address of the user for optimal ZPA Private Service Edge selection.

App Connectors should be able to reach your ZPA Private Service Edge's public or private IP address. When a user requests an application through your ZPA Private Service Edge, ZPA asks an App Connector to set up a data connection to this ZPA Private Service Edge. If Listen IPs are not set for your ZPA Private Service Edge, App Connectors in other data centers connect to your ZPA Private Service Edge's public IP (if available). Zscaler recommends the use of Listen IPs to impose preferred internal connectivity.

If it is less optimal for App Connectors to reach your ZPA Private Service Edge on the public IP address, each ZPA Private Service Edge should have a unique private IP address that can be reached by App Connectors over the internal network (refer to scenario 2 under [Configuration Requirements](#) for remote users accessing regional internal applications).



If no App Connector for the requested application can establish a data connection to the required ZPA Private Service Edge where the user is connected, that ZPA Private Service Edge relays Zscaler Client Connector's request to a ZPA Public Service Edge, and the App Connector connects to the ZPA Public Service Edge. So, there will be an additional hop through a ZPA Public Service Edge to complete the application request.

Configuration Requirements

For this use case, the Service Edge Group must have the Publicly Accessible setting enabled. Other settings depend on which scenario you deploy.

Scenario 1: All App Connectors can reach ZPA Private Service Edge's Public IP address.

Listen IPs and Publish IPs or Domains can be configured optionally. Specify Listen IPs when only certain interfaces on your ZPA Private Service Edge are used for accepting incoming requests. If Listen IPs are specified, ZPA automatically advertises Listen IPs as Publish IPs. In this case, the Listen IPs need to be public IPs. If Listen IPs are not specified, ZPA advertises the IPs of every available interface on your ZPA Private Service Edge. In this case, the IPs need to be public IPs.

Specify Publish IPs or Domains when the IP or domain that clients must use to reach your ZPA Private Service Edge differs from the available interface IPs. In most cases, Publish IPs or Domains don't need to be specified. If Publish IPs or Domains need to be specified, the resulting traffic should be routed to the private IPs of your ZPA Private Service Edge. For example, if your ZPA Private Service Edge is deployed behind a firewall, and the firewall is performing DNAT for your ZPA Private IP and advertising a public IP on the internet, the public IP advertised by the firewall must be configured as a Publish IP.

Any one of the configured Publish IPs or Domains can be provided to Zscaler Client Connector and App Connectors for connecting to your ZPA Private Service Edge. If Publish Domains are specified, organizations must create corresponding DNS entries in public DNS resolvers and in internal DNS resolvers, if your ZPA Private Service Edge is also accessed by internal users.

Scenario 2: Not all App Connectors can reach a Private Service Edge via a public IP address.

Each ZPA Private Service Edge requires a unique Publish Domain. Customers must specify a FQDN that can resolve on the internet to your ZPA Private Service Edge's public IP address and on the internal network to your ZPA Private Service Edge's private IP address.

If your organization has a ZPA app segment containing a wildcard (such as **.safemarch.com*) that includes the FQDN for your ZPA Private Service Edge, it can cause routing issues. A bypass config is needed, either in the ZPA Client Forwarding Policy or in the Zscaler Client Connector App Profile, so that traffic to your ZPA Private Service Edge is not intercepted as application traffic.

Zscaler Client Connector and App Connectors might resolve to your ZPA Private Service Edge's public or private IP address, depending on the location of the client. Clients can attempt to connect using the Publish Domain.

Remote and On-Premises Users Accessing Internal Applications

This deployment scenario is relevant to organizations who are interested in implementing zero trust access for both on-premises users and remote users when a ZPA Public Service Edge is not the closest ZPA Service Edge to the user. Both remote and on-premises users can access applications using your ZPA Private Service Edge, because your ZPA Private Service Edge is the closest ZPA Service Edge for users.

This deployment model assumes that branches and data centers are connected via dedicated connections such as MPLS and/or WAN solutions. On-premises Zscaler Client Connector and App Connectors can reach your ZPA Private Service Edge over the internal network. This deployment model also assumes that App Connectors in public cloud locations not connected to the internal network are able to reach your ZPA Private Service Edge over the public internet.

Deploying a ZPA Private Service Edge provides several advantages:

- Customers can achieve zero trust access for sensitive internal applications, without sending the application traffic through a ZPA Public Service Edge.
- Customers can reduce the application access latency, as the application traffic goes through a ZPA Private Service Edge that is geographically closer compared to a ZPA Public Service Edge.

For this scenario, Zscaler recommends that ZPA Private Service Edges are hosted in the on-premises data center and/or hub data centers.

Networking Requirements

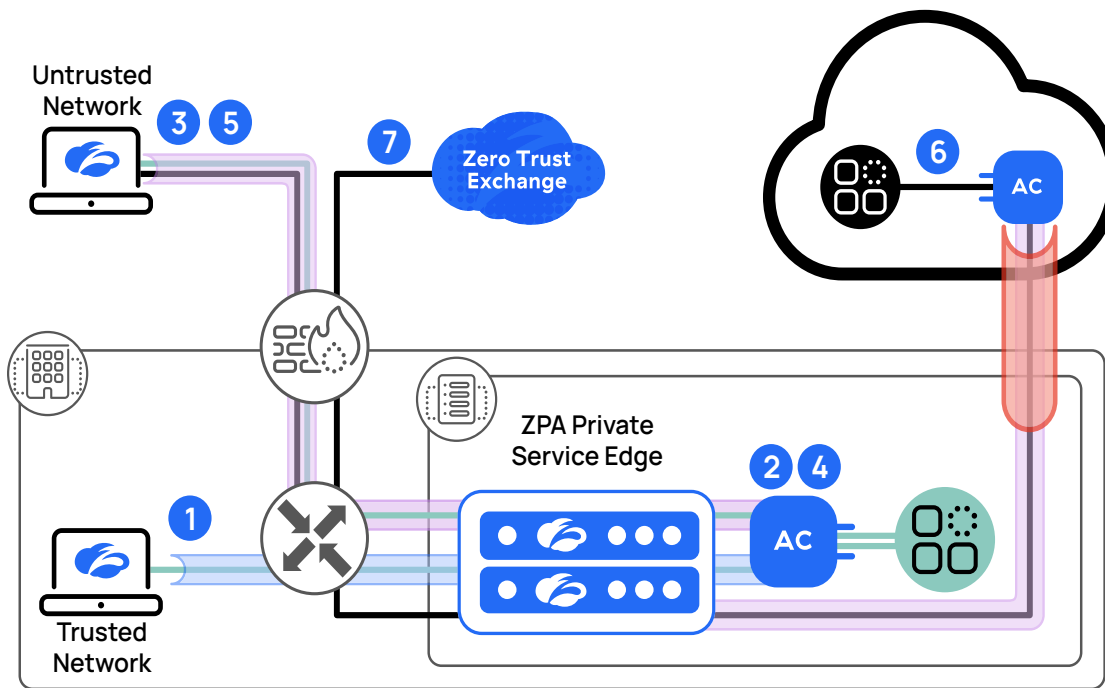


Figure 15. On-premises and remote users accessing internal applications

1. Each ZPA Private Service Edge must be able to accept incoming connections from both internal and external sources on TCP port 443.
2. Each ZPA Private Service Edge must have a unique public IP that can be reached by remote Zscaler Client Connector over the public internet.
3. Each ZPA Private Service Edge must also have a unique private IP address that can be reached by on-premises Zscaler Client Connector and App Connectors.
4. For Zscaler Client Connector to connect to your ZPA Private Service Edge, Zscaler Client Connector must be able to reach your ZPA Private Service Edge's public or private IP address. If Zscaler Client Connector is unable to connect to your ZPA Private Service Edge, it connects to the closest ZPA Public Service Edge.
5. ZPA only selects your ZPA Private Service Edge for a remote user if it is the closest ZPA Service Edge, given the location of the user. The location of the user is determined based on the public IP address of the user, as seen by ZPA.
 - a. For an on-premises user, ZPA selects your ZPA Private Service Edge if either your ZPA Private Service Edge is the closest ZPA Service Edge, or your ZPA Private Service Edge is configured to be on a specific trusted network when Zscaler Client Connector is reporting itself being on the same trusted network.
6. Any traffic destined to `*.prod.zpath.net` from Zscaler Client Connector must be bypassed from ZIA or other proxies and go directly to the internet. Zscaler Client Connector always resolves `any.broker.prod.zpath.net` to initially connect to a ZPA Public Service Edge. ZPA Public Service Edges can redirect Zscaler Client Connector to your ZPA Private Service Edge based on the conditions in step 4. It is important that ZPA sees the true public IP address of the user for optimal ZPA Private Service Edge selection.
7. App Connectors should be able to reach your ZPA Private Service Edge's public or private IP address. When a user requests an application through your ZPA Private Service Edge, ZPA asks an App Connector to set up a data connection to this ZPA Private Service Edge. If Listen IPs are not set for your ZPA Private Service Edge, App Connectors in other data centers connect to your ZPA Private Service Edge's public IP. Zscaler recommends the use of Listen IPs to impose preferred internal connectivity.

- a. If it is less optimal for App Connectors to reach your ZPA Private Service Edge on the public IP address, each ZPA Private Service Edge should have a unique private IP address that can be reached by App Connectors over the internal network (refer to scenario 2 under [Configuration Requirements](#) for remote users accessing regional internal applications).



If no App Connector for the requested application can establish a data connection to the required ZPA Private Service Edge where the user is connected, that ZPA Private Service Edge relays Zscaler Client Connector's request to a ZPA Public Service Edge, and the App Connector connects to the ZPA Public Service Edge. So, there will be an additional hop through a ZPA Public Service Edge to complete the application request.

Configuration Requirements

In this use case, the Service Edge Group must have the Publicly Accessible setting enabled. Specify a trusted network if Zscaler Client Connector is configured to report trusted networks. Specifying trusted networks assists ZPA with redirecting Zscaler Client Connector to your ZPA Private Service Edge on the same network as the user. If the user is closer to a ZPA Public Service Edge and you want the user to connect via your ZPA Private Service Edge instead, specifying a trusted network for your ZPA Private Service Edge is required.

Scenario 1: Zscaler Client Connector and all App Connectors can reach ZPA Private Service Edge's Public IP address.

Listen IPs and Publish IPs or Domains can be configured optionally. Specify Listen IPs when only certain interfaces on your ZPA Private Service Edge are used for accepting incoming requests. In this case, the Listen IPs need to be public IPs. If Listen IPs are specified, ZPA automatically advertises Listen IPs as Publish IPs. If Listen IPs are not specified, ZPA advertises the IPs of every available interface on your ZPA Private Service Edge. If Publish IPs or Domains need to be specified, the IPs should be the public IPs.

Your ZPA Private Service Edge is either configured with a public IP, or if your ZPA Private Service Edge is deployed behind a firewall, the firewall performs DNAT for your ZPA Private Service Edge's private IP and advertises a public IP on the internet. In such a scenario, the Public IP advertised by the FW needs to be configured as a Publish IP.

If a Domain is specified, it must be resolvable in public DNS to your ZPA Private Service Edge's public IP. Any one of the Published IPs or Domains will be provided to clients for connecting to your ZPA Private Service Edge.

Scenario 2: Not all App Connectors can reach a Private Service Edge via a public IP address

Each ZPA Private Service Edge requires a unique Publish domain. You must specify a FQDN that can resolve on the internet to your ZPA Private Service Edge's public IP address and on the internal network to your ZPA Private Service Edge's private IP address.

If your organization has a ZPA app segment containing a wildcard (such as **.safemarch.com*) that includes the FQDN for your ZPA Private Service Edge, it can cause issues. A bypass config is needed, either in the ZPA Client Forwarding Policy or in the Zscaler Client Connector App Profile, so that traffic to your ZPA Private Service Edge is not intercepted as application traffic.

Zscaler Client Connector and App Connectors might resolve to your ZPA Private Service Edge's public or private IP address, depending on the location of the client.

An alternate deployment model would be to deploy two separate ZPA Private Service Edge groups. The first ZPA Private Service Edge group is only reachable from the internal network for on-premises users. The second ZPA Private Service Edge group is reachable externally for remote users.

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so 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.

©2022 Zscaler, Inc. All rights reserved. Zscaler, Zero Trust Exchange, Zscaler Private Access, ZPA, Zscaler Internet Access, ZIA, Zscaler Digital Experience, and ZDX are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.