



# Traffic Forwarding in Zscaler Internet Access™

Reference Architecture

# Contents

<b>About Zscaler Reference Architectures Guides</b>	<b>2</b>
Who Is This Guide For?	2
A Note for Federal Cloud Customers	2
Conventions Used in This Guide	2
Finding Out More	2
Terms and Acronyms Used in This Guide	3
Icons Used in This Guide	4
<b>Introduction</b>	<b>5</b>
New to ZIA and Traffic Forwarding?	6
<b>Solution Overview</b>	<b>6</b>
<b>Fixed Site - Transparent Forwarding</b>	<b>10</b>
GRE Tunnel (Recommended)	10
IPSec VPN	18
Dedicated Proxy Ports	24
Surrogate IP for Fixed Site Deployments (Recommended)	25
<b>Mobile Users - Explicit Forwarding</b>	<b>27</b>
Zscaler Client Connector (Recommended)	28
PAC Files	32
<b>Specialized Forwarding Cases</b>	<b>34</b>
Sending Traffic from a Non-Zscaler Source IP	34
Load Balancing across Multiple WAN Links (Bonded DSL, etc.)	37
No Default Route Networks	37
Proxy Chaining	39
<b>About Zscaler</b>	<b>40</b>

## 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>).

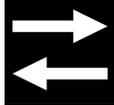
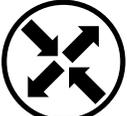
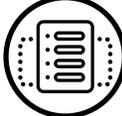
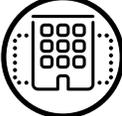
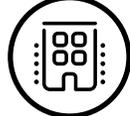
You can join our user and partner community and get answers to your questions in the [Zenith Community](https://community.zscaler.com) (<https://community.zscaler.com>).

## Terms and Acronyms Used in This Guide

Acronym	Definition
CSV	Comma Separated Value file
DC	Data Center
DPD	Dead Peer Detection
GRE	Generic Routing Encapsulation
IaaS	Infrastructure as a Service
IKE	Internet Key Exchange
IP	Internet Protocol
IPSec	Internet Protocol Security
MDM	Mobile Device Management
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NAT-T	NAT Transversal
PAC	Proxy Auto-Configuration
PFS	Perfect Forward Secrecy
PSK	Pre-Shared Keys
SA	Security Associations
SaaS	Software as a Service
SAML	Security Assertion Markup Language
SSL	Secure Socket Layer (superseded by TLS)
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
VCS	Version Control System
ZDX	Zscaler Digital Experience
ZIA	Zscaler Internet Access
ZPA	Zscaler Private Access
ZTE	Zero Trust Exchange

## Icons Used in This Guide

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

					
Zscaler Zero Trust Exchange	ZIA or ZPA Service Edge	Zscaler App Connector	Branch Connector	Zscaler Load Balancer	Virtual IP
					
Zscaler Client Connector on Devices	Laptop and Phone	Desktop Workstation	IoT Device	NAT Device	Router
					
PAC File Server	SD-WAN	Private Data Center Location	Headquarters Office Location	Branch Office Location	Factory Location
					
Airport	Private Residence	Legacy Firewall	Legacy Security Appliance	Internet	Data Tunnel
					
Authorized User	Bad Actor	Dropped Traffic			

## Introduction

The Zscaler Internet Access™ (ZIA™) service is a secure internet and web gateway delivered in the cloud as a part of the Zero Trust Exchange™, the world's largest security cloud. ZIA provides a full security stack with ZIA Service Edges enforcing your policies and protecting your users wherever they are, from the head office to their home offices. The same authentication and policy follow the user, and that policy can be adjusted based on location, device in use, and more. User traffic is inspected and forwarded on or blocked according to your defined policy.

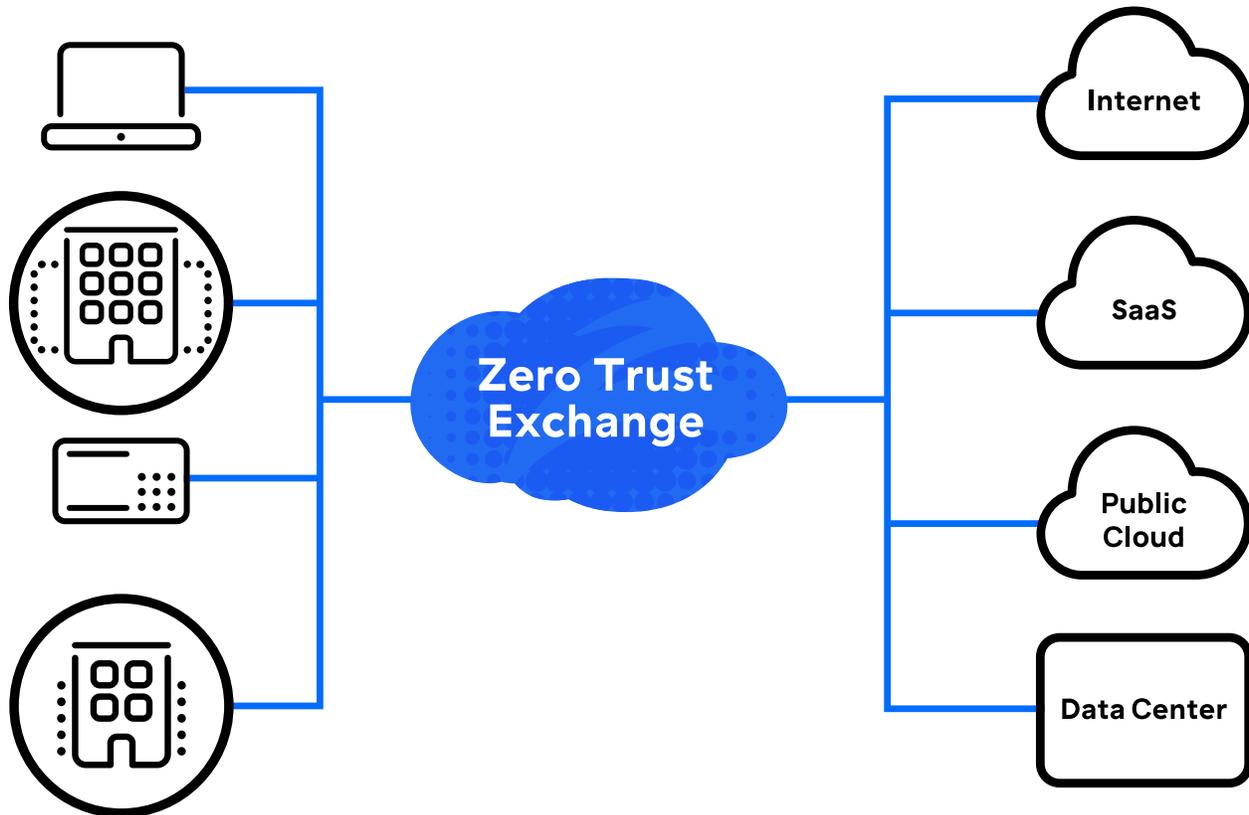


Figure 1. Zscaler ZIA cloud-based security

To provide this protection, you first must get your traffic to the nearest ZIA Service Edge. Unlike previous centralized models that backhauled traffic to a central location, ZIA is available in 150+ data centers around the world. Many of these data centers are in internet exchange peering points with other cloud infrastructure and application providers, such as Microsoft, Amazon, and Google.

In this guide, we discuss best practice for forwarding traffic to the ZIA Service Edge, where your users are authenticated and your policy is enforced.

This guide contains three primary sections:

1. Forwarding traffic from known site locations, such as a campus location or branch office. These locations have fixed infrastructure but might use commercial or commodity internet connections.
2. Forwarding traffic from mobile users connecting directly to the internet, including remote workers, field teams, and traveling executives.
3. Special cases in traffic forwarding, including private source IP and bonded internet links.

We cover when to use each solution for each access type, allowing you to layer them to provide seamless protection from the office to the home.

## New to ZIA and Traffic Forwarding?

- Watch a short video and download the ZIA e-book from the [Zscaler website \(https://www.zscaler.com/products/zscaler-internet-access\)](https://www.zscaler.com/products/zscaler-internet-access).
- Learn about Zscaler partners in the [SD-WAN space \(https://www.zscaler.com/partners/technology/sd-wan\)](https://www.zscaler.com/partners/technology/sd-wan).
- Learn about [Zscaler Client Connector \(https://www.zscaler.com/platform/zscaler-client-connector\)](https://www.zscaler.com/platform/zscaler-client-connector).

## Solution Overview

ZIA is an advanced web proxy solution designed and built for the cloud. The strength of the service comes from its distributed nature. No matter where your users are located, they are always directed to the nearest ZIA Service Edge for traffic processing. Each ZIA Service Edge delivers full TLS/SSL interception for all of your traffic, web proxying, firewall, anti-virus, anti-malware, and security services. The same authentication and policy are applied at each location. Logging is comprehensive and gives you visibility into your entire organization.

Zscaler has deployed ZIA Service Edge devices in 150+ data centers around the world. This includes peering at internet exchanges with other service providers. This wide distribution means your users can connect to the nearest ZIA Service Edge, keeping latency low for your web-based applications by eliminating backhauling. You can view Zscaler's extensive [data center map \(https://trust.zscaler.com/data-center-map\)](https://trust.zscaler.com/data-center-map).

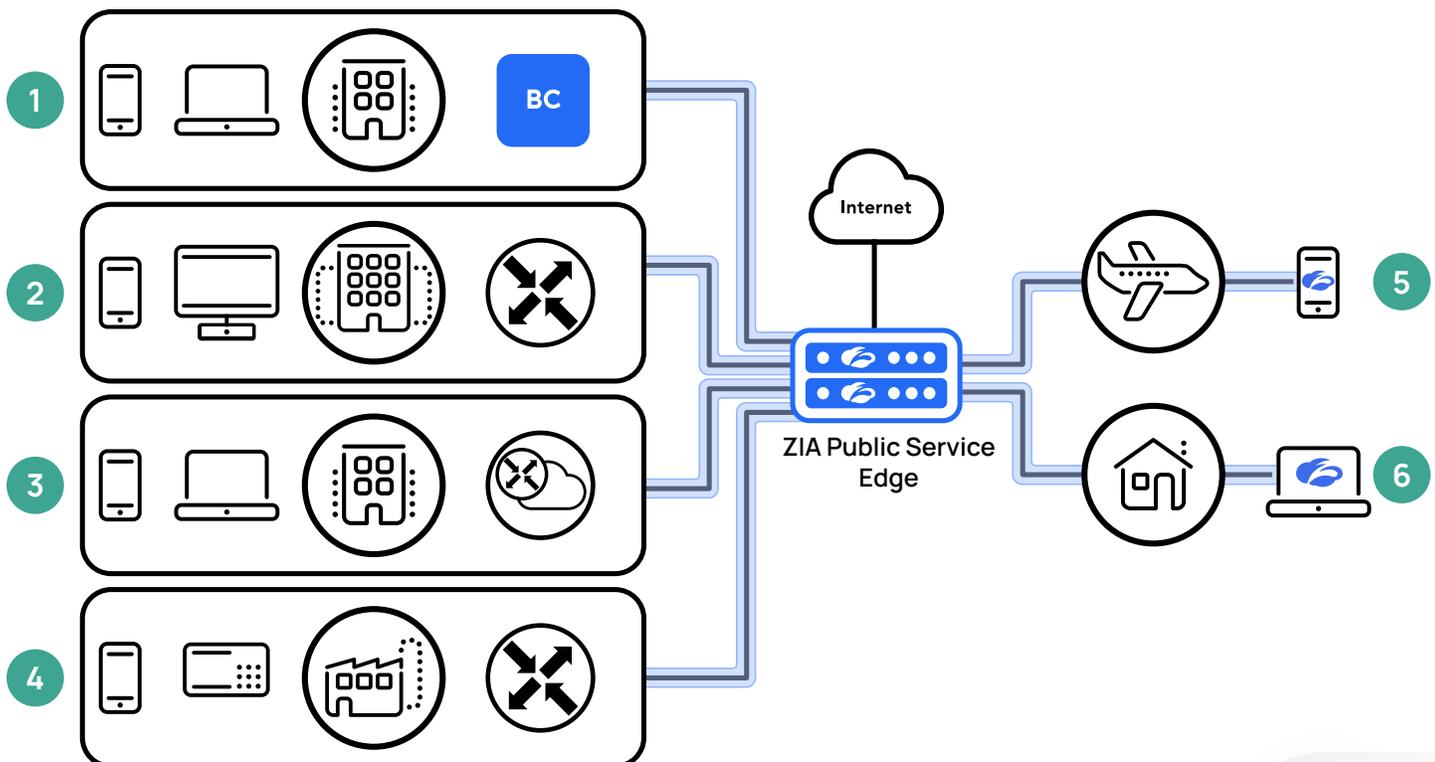


Figure 2. Traffic forwarding from fixed or mobile locations

In the previous graphic, we highlight the five forwarding methods covered in this guide. These methods fall into two categories: transparent proxy used at known locations, and explicit proxy used in forwarding mobile traffic. In most use cases, you'll use more than one solution to meet connectivity needs across your users.

## Transparent Proxy – Fixed Sites

1. Zscaler Branch Connector – Simplify traffic forwarding and centrally manage your remote networks with Zscaler Branch Connector. Available in hardware formats for small to medium sites, and as a virtual machine (VM) for larger sites and data centers. Branch Connector establishes tunnels to the Zero Trust Exchange and can act as your internet gateway.
2. Generic Routing Encapsulation (GRE) (recommended) – GRE wraps a simple header around your traffic destined for the nearest ZIA Service Edge. GRE is available on many enterprise routers and SD-WAN devices. Note that GRE requires each of your organizations' locations to have a static IP address. Learn more about [GRE \(https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre\)](https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre).
3. Internet Protocol Security (IPSec) – IPSec forms an encrypted tunnel between your device router or SD-WAN device and the ZIA Service Edge. Typically, IPSec VPN is only used when the gateway device doesn't support GRE or have a static IP address. Note that IPSec VPNs have bandwidth constraints. Learn more about [IPSec \(https://help.zscaler.com/zia/about-ipsec-vpns\)](https://help.zscaler.com/zia/about-ipsec-vpns).
4. Dedicated Proxy Ports – This subscription service provides you with dedicated ports on the ZIA Service Edge infrastructure, where you can forward traffic to these ports from your gateway device. To prevent abuse of proxy ports, authentication must be enabled for all users. Learn more about [dedicated proxy ports \(https://help.zscaler.com/zia/configuring-dedicated-proxy-ports\)](https://help.zscaler.com/zia/configuring-dedicated-proxy-ports).



For information about Zscaler Cloud Connector as a forwarding mode for your cloud applications, see the following reference architectures focused on Zscaler Cloud Connector in Amazon Web Services and Microsoft Azure:

- [Zero Trust Security for AWS Workloads with Zscaler Cloud Connector \(https://help.zscaler.com/cloud-connector/zero-trust-security-aws-workloads-zscaler-cloud-connector\)](https://help.zscaler.com/cloud-connector/zero-trust-security-aws-workloads-zscaler-cloud-connector)
- [Zero Trust Security for Azure Workloads with Zscaler Cloud Connector \(https://help.zscaler.com/cloud-connector/zero-trust-security-azure-workloads-zscaler-cloud-connector\)](https://help.zscaler.com/cloud-connector/zero-trust-security-azure-workloads-zscaler-cloud-connector)

## Explicit Proxy – Mobile Users

5. Zscaler Client Connector (recommended) – This lightweight agent is included in your ZIA subscription. It is installed on user devices and is the mobile gateway to the ZIA service. Zscaler Client Connector detects if you are on a trusted network at one of your configured locations and builds a tunnel to the nearest ZIA Service Edge. Zscaler Client Connector on the end device can also be leveraged for use with other Zscaler services, including Zscaler Private Access™ (ZPA™) or Zscaler Digital Experience™ (ZDX™). This software supports the most common operating systems, including Windows, macOS, Android, iOS, and Linux. Learn more about [Zscaler Client Connector \(https://help.zscaler.com/z-app/what-zscaler-app\)](https://help.zscaler.com/z-app/what-zscaler-app).
6. Proxy Auto-Configuration (PAC) Files – PAC files provide a mechanism for setting up forwarding between a browser and ZIA. This is a JavaScript file that sets the proxy server address and optionally additional forwarding rules. This older technology is typically only used on devices without general purpose operating systems where Zscaler Client Connector cannot be installed. Learn more about [PAC files \(https://help.zscaler.com/zia/about-pac-file\)](https://help.zscaler.com/zia/about-pac-file).

## Forwarding Decision Tree

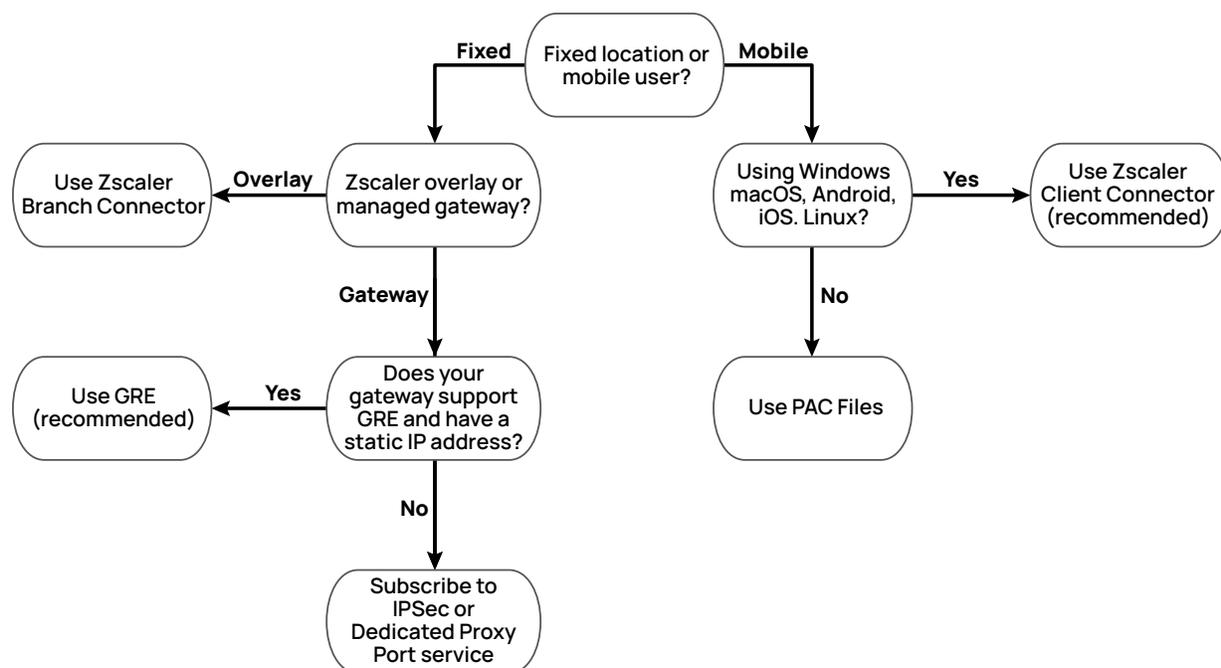


Figure 3. Forwarding mode decision tree

The method you use to connect to ZIA depends on where the forwarding is happening. You will likely use a mix of technologies across your organization. Use the previous diagram to quickly select a solution for each of your use cases.

Zscaler Branch Connector sets up tunnels for both ZIA and ZPA services to the Zscaler Zero Trust Exchange (ZTE), providing a simple network overlay. You can leverage the hardware versions for your branch gateways and the VM version in a campus data center. This extends the secure access service edge (SASE) protections of Zscaler to your network edge.

If you plan to leverage your existing network routers and equipment, you need to configure locations and connections back to the ZTE. For your large locations, you can use GRE from your gateway router, and you might use a mix of GRE and IPsec at your branch locations via an SD-WAN device. Proxy ports are used only when no other forwarding option is available. Zscaler recommends that you use GRE for your locations as a best practice from your own router or SD-WAN device.

Your mobile users can install Zscaler Client Connector to get them connected when they are off site. PAC files can be used as a last resort if Zscaler Client Connector can't be installed.



No matter which transparent forwarding option you choose, Zscaler recommends installing Zscaler Client Connector on all devices. This ensures your users are protected with the same security in and out of the office. Zscaler Client Connector is included in your ZIA subscription for all of your users.

It is also possible to move to an internet-only organization, combining ZIA with Zscaler Private Access (ZPA). In this model, there is no "inside" the network. All of your devices use Zscaler Client Connector to connect to cloud applications, either your private applications or SaaS-based offerings. Your organization's network would be simplified to provide access to the internet, and all applications would be hosted in the cloud.

## Special Forwarding Cases

While most forwarding choices are determined by location and client, there are some use cases that require additional consideration. These include no default route networks or bonding multiple uplinks from branch sites. We also cover the case where you must use a fixed set of IP addresses to access a particular application.

Geopolitical considerations can also exist. For example, deployments in China require that your traffic is inspected in the country. For more information, see [Deploying Zscaler Internet Access in China](https://help.zscaler.com/zia/deploying-zscaler-internet-access-china) (<https://help.zscaler.com/zia/deploying-zscaler-internet-access-china>).

## Fixed Site - Transparent Forwarding

In this chapter, we focus on traffic forwarding at fixed sites. This includes any location where you have deployed an enterprise-grade internet gateway, such as on a campus or branch location. Your choice of connections depends on two factors: equipment capabilities at the site, and if the site has a static IP address.

On the hardware side, we are looking at your internet gateway device. This can be a gateway router on a campus, a Zscaler Branch Connector, or an SD-WAN box at a branch site. These devices typically support GRE, IPSec, or both. You need to determine which tunneling mechanisms are supported based on your equipment and software licenses.

The second requirement is a static IP address. This is required for GRE, but not for an IPSec VPN or Dedicated Proxy Port connections. If you cannot obtain a static IP address, you need to leverage one of these services or have your users connect via Zscaler Client Connector. Zscaler recommends the use of GRE whenever possible from fixed locations. This protocol ensures all traffic is forwarded to the ZIA Service Edge with the least amount of overhead.



Typically, consumer-grade devices such as a combination modem, router, and access point are not capable of either GRE or IPSec. If your branch is leveraging consumer grade DSL or cable, you might not be able to use that device to transparently connect to ZIA. Instead, the users in that location should be considered mobile users. See [Mobile Users - Explicit Forwarding](#) for more information.

While not a forwarding mechanism on its own, Zscaler recommends enabling Surrogate IP for all transparent forwarding types. This feature enables the mapping of private IP addresses to users so that the users' policies are applied to all traffic. We discuss this further in [Surrogate IP for Fixed Site Deployments \(Recommended\)](#).



Surrogate IP requires visibility into private IP addresses with user authentication. User traffic must enter the tunnel without NAT translation, and your locations must have authentication enabled.

### GRE Tunnel (Recommended)

Zscaler recommends the use of GRE whenever possible. This low-overhead protocol is ideal for ensuring all your internet-bound traffic is inspected by the ZIA Service Edge. This service is transparent to users as the traffic is tunneled from the gateway device to the ZIA Service Edge.

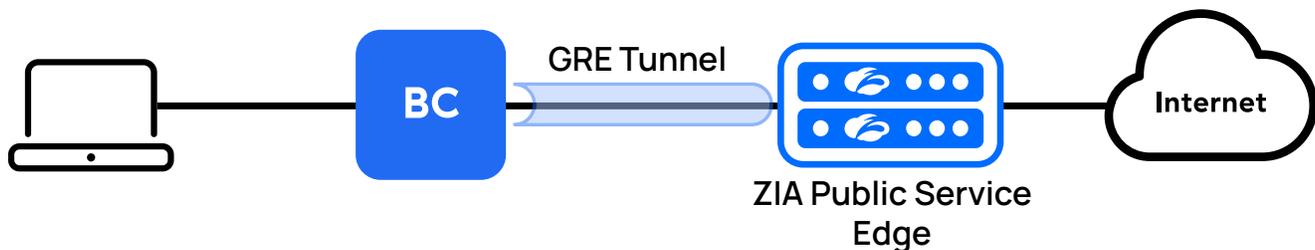


Figure 4. Traffic tunneled over GRE

This protocol supports internet-bound traffic by creating one or more logical interfaces in your gateway device. As traffic bound for the internet is received on the gateway device, a new set of headers is added to the front of your original IP packet. The first is a new IP header with the destination address of the ZIA Service Edge currently in use. The second is a GRE header with information about the protocol. Your entire packet including your original IP header is preserved within the payload, allowing for source-IP logging and policy actions based on source address.

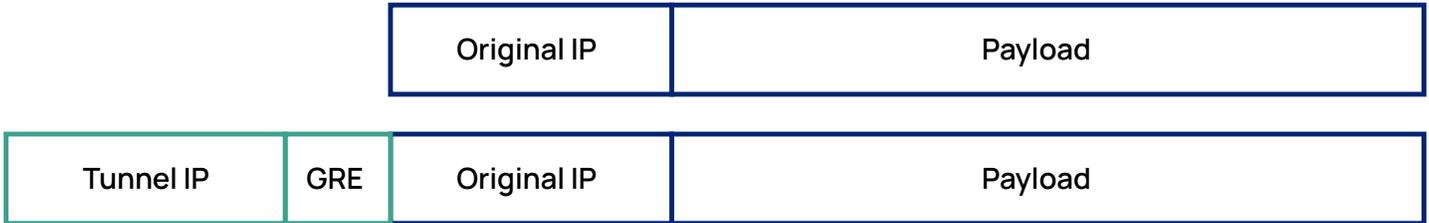


Figure 5. IP packet before and after GRE encapsulation

This is similar to how a traditional VPN operates. The difference is that the traffic is not encrypted; it is simply wrapped in an additional header. This also means that end-to-end security is the job of the client and internet application being accessed. As an example, a request sent over HTTP instead of HTTPS is still visible inside a GRE tunnel if intercepted. This is the same as if no tunneling occurred. When the traffic exits the GRE tunnel and is processed by the ZIA Service Edge, it exits with the same level of encryption as it entered the tunnel.

### GRE Bandwidth Considerations

Each GRE tunnel can support up to 1 Gbps if the users' internal IP addresses are not behind a NAT device. This allows the ZIA load balancers to shift traffic to different ZIA enforcement nodes in the same data center.

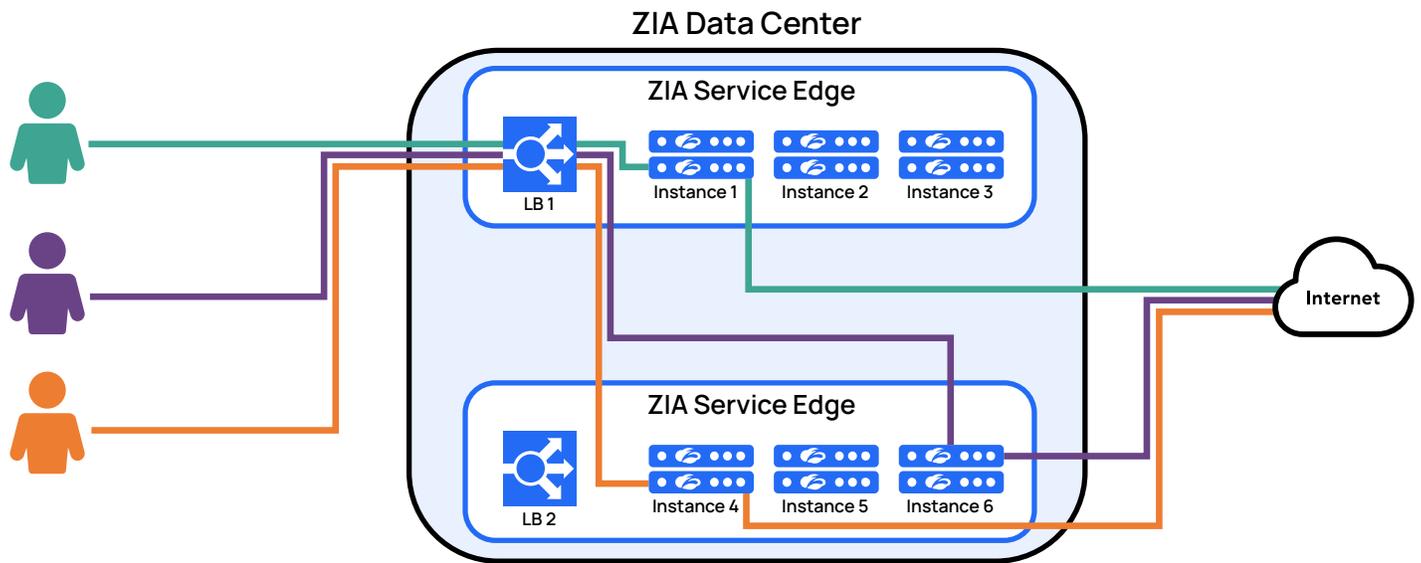


Figure 6. Load balancing between a pair of ZIA Service Edge devices

Users are load balanced to an appropriate ZIA Service Edge instance for inspection and enforcement. This balancing is done using the source and destination IP addresses, and the user is directed to the same Service Edge instance.

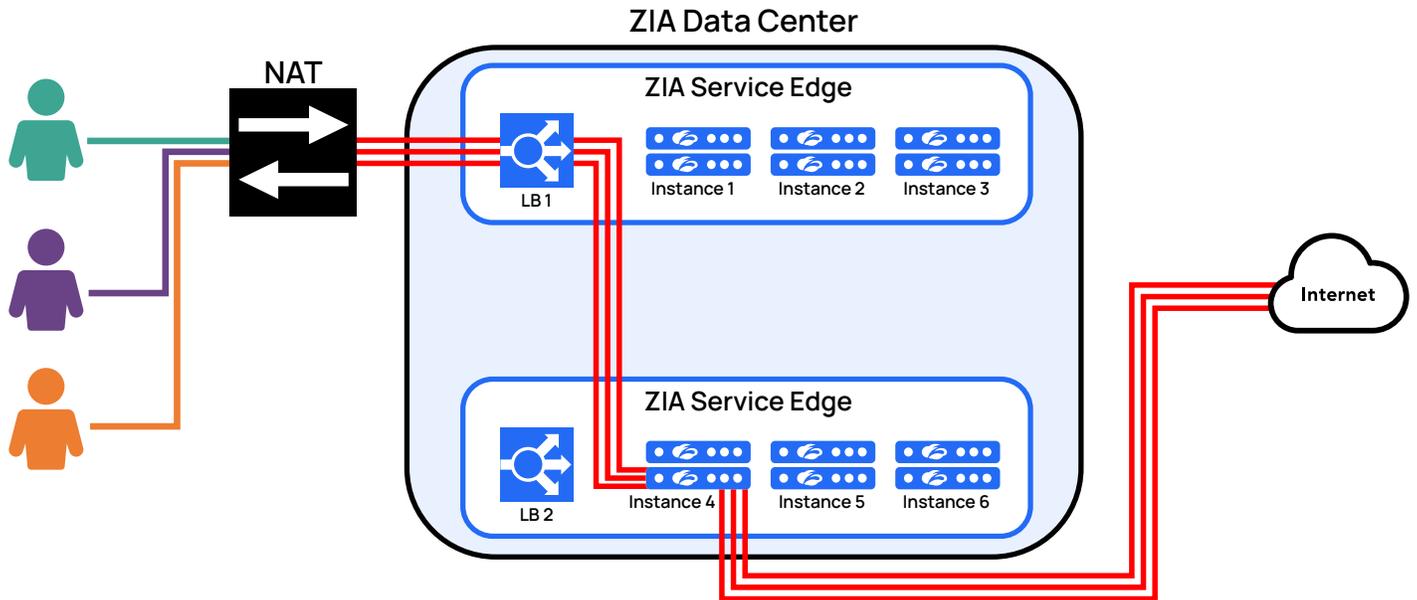


Figure 7. No load balancing when a NAT device exists between users and the ZIA Service Edge

If your users are behind a NAT device before being tunneled to the ZIA Service Edge, all of your user traffic will appear to be coming from a single user when it reaches the ZIA Service Edge. Your bandwidth is constrained to 250 Mbps in this configuration. Zscaler best practice is to tunnel traffic to the ZIA Service Edge without NAT translation.

Traffic more than 1 Gbps from a single location can be handled by adding more GRE tunnels. Each tunnel must originate from a unique public IP address. In these scenarios, Zscaler recommends that you construct a load balancing strategy that ensures that the same IP address travels across the same tunnel. Failure to do so can cause users to continuously authenticate and they might not be able to access their applications.

Learn more about [GRE bandwidth limits](https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre#supported-bandwidth) (<https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre#supported-bandwidth>).

## GRE Redundancy

GRE configuration supports building redundant tunnels between your network gateway and two different ZIA Service Edge data centers. One tunnel operates in active mode, and another in standby mode. In the event of a loss of connectivity to the primary data center, your gateway router, or SD-WAN device will fail over to the backup data center.

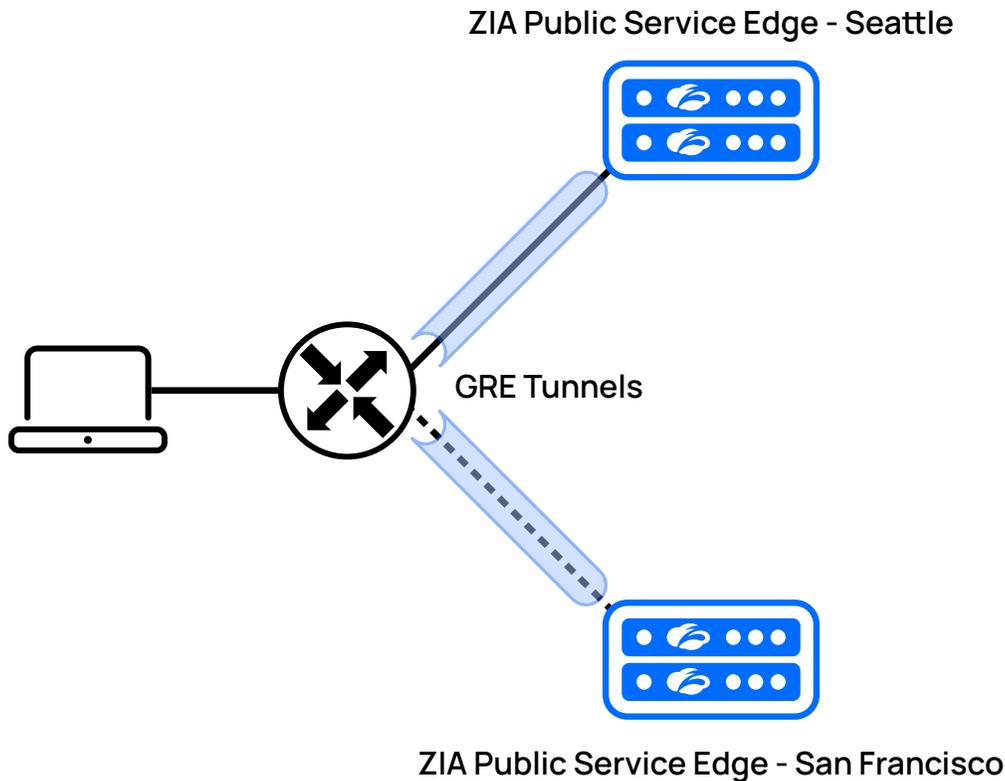


Figure 8. Redundant GRE tunnels from the network gateways, both from routers and SD-WAN devices

In the previous image, the network gateway has two GRE tunnels configured. The primary tunnel goes out to our Seattle ZIA Public Service Edge, with a backup tunnel connecting to San Francisco. Zscaler requires that you build primary and secondary GRE tunnels from each internet egress point. Additionally, if you have multiple ISPs at a single location, primary and secondary tunnels must be built from each ISP. This ensures your users have high availability access to the service.

## GRE Setup

GRE tunnels are built as logical interfaces on your gateway routers, Zscaler Branch Connector, or SD-WAN devices. You can work with Zscaler Support to determine where tunnels will terminate. The device that you build your tunnel on depends on your network design and the types of devices. At larger campus sites, there are often internal and external routers on either side of a firewall. At a branch site, there might only be a single Zscaler Branch Connector or SD-WAN device. Zscaler always recommends redundant links and redundant edge devices where possible.



If you need specific data centers to be used due to language or legal requirements, Zscaler Support can work with you to ensure the proper data centers are selected.

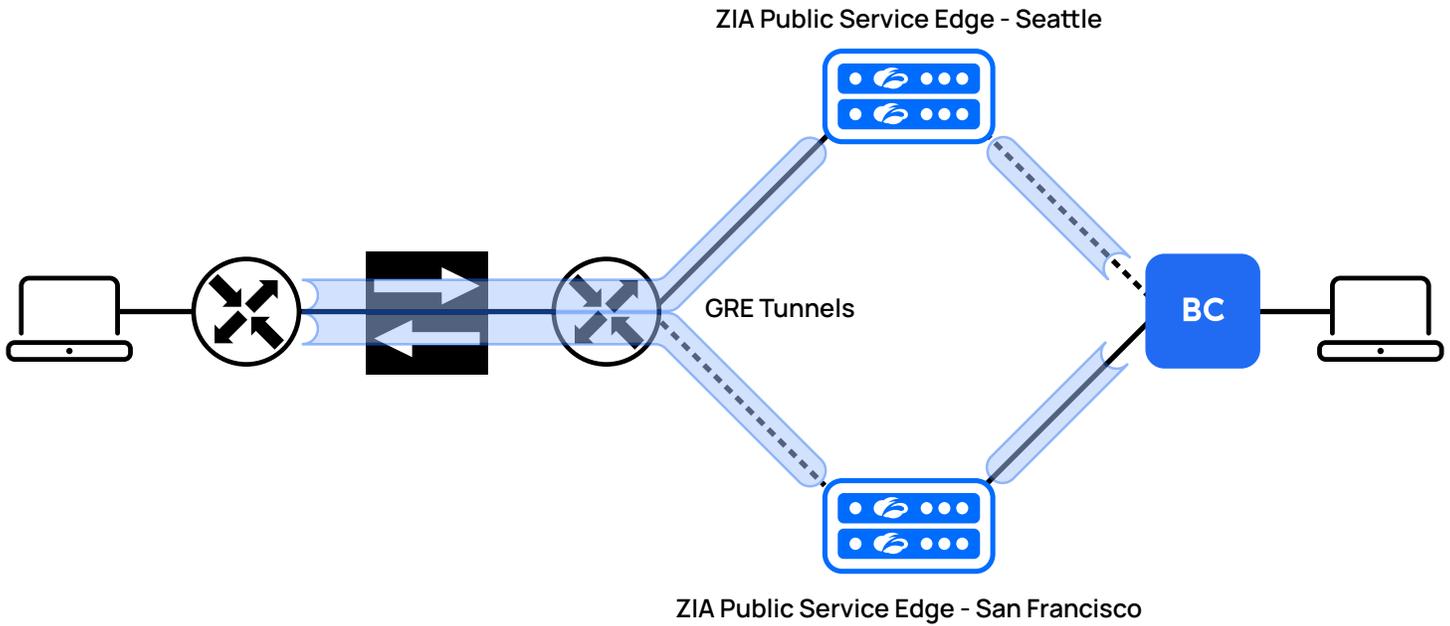


Figure 9. GRE setup on a campus and in a branch

When deploying on a campus with multiple routers, Zscaler recommends using the internal gateway router before the firewall as your tunnel source. Firewalls are typically responsible for performing NAT. As mentioned previously, NAT interferes with the ZIA Service Edge’s ability to distinguish users from one another due to all users sharing the same IP address.

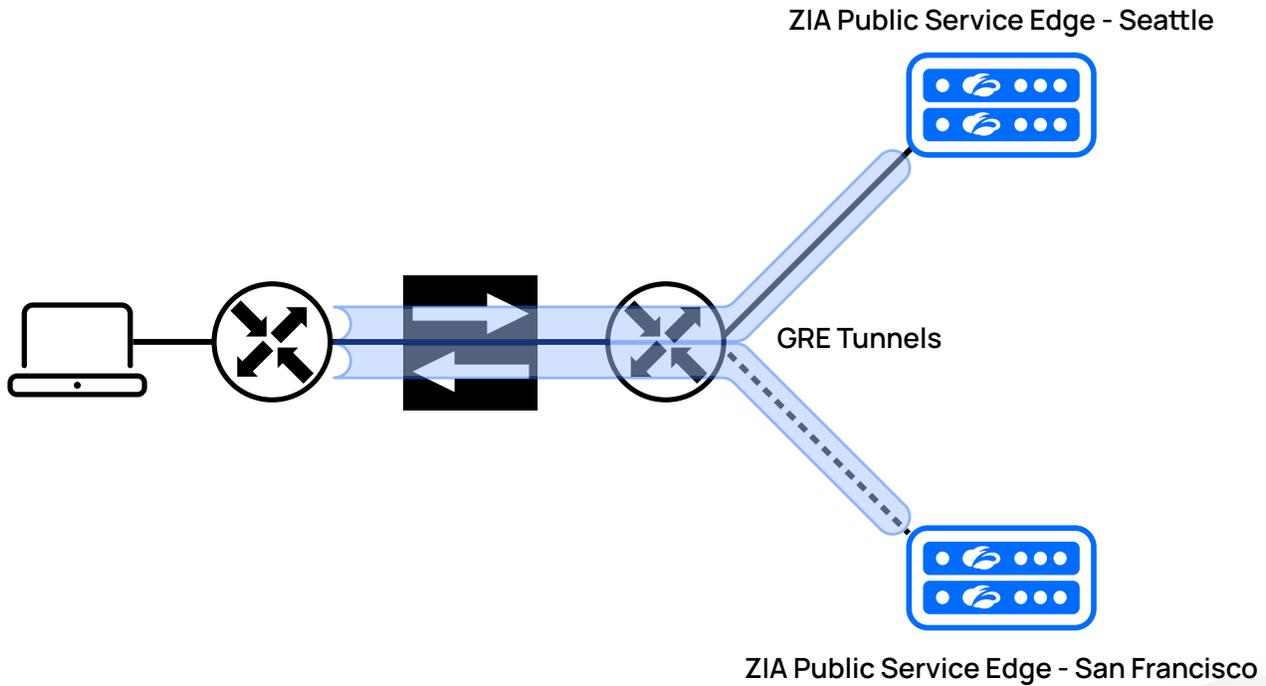


Figure 10. Campus network with GRE tunnels extending from the interior gateway through the firewall

If it is not possible to deploy in front of the firewall, you must bypass the NAT feature for traffic destined to the internet. You also need to ensure your internet-facing router is aware of your internal network and has a route back. This allows the ZIA Service Edge to see your internal address space.

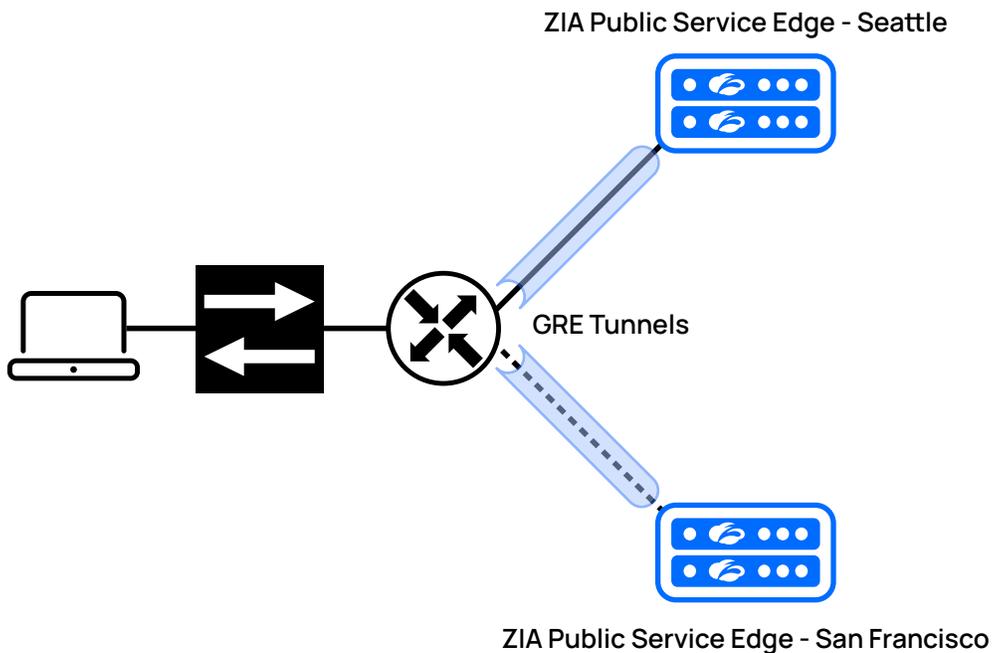


Figure 11. If you cannot tunnel from your inner gateway, you should bypass your firewall NAT

For smaller sites such as a branch office, you might have a simplified network with security services collapsed into a single device such as a Zscaler Branch Connector or an SD-WAN device. This device typically acts as a multi-function firewall, gateway router, and policy-based forwarding engine. When building out a network from a single device, ensure that internet-bound traffic is forwarded via the GRE tunnel and that NAT is disabled.

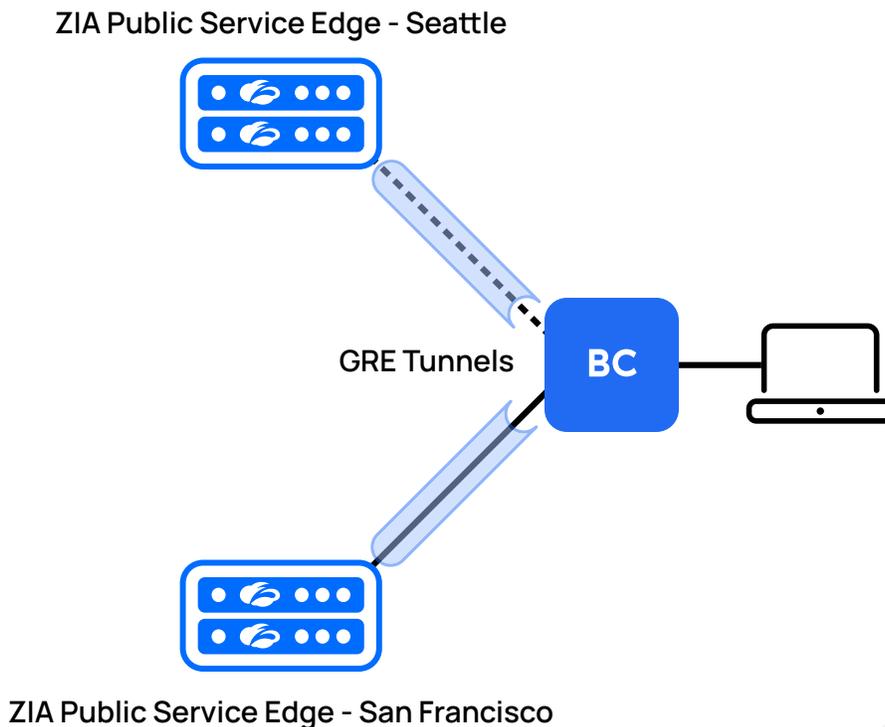


Figure 12. GRE tunnels from your Branch Connector or SD-WAN edge

The last step in GRE tunnel setup is ensuring that your maximum transmission unit (MTU) and maximum segment size (MSS) values for the GRE tunnel are set correctly. These values are based on the MTU and MSS configuration of the WAN interface of your router, Branch Connector, or SD-WAN device. It is Zscaler best practice to ensure that these values match. By configuring the proper values, you reduce fragmentation and see higher throughput.

To calculate the MTU and MSS, you need to start with your WAN link's configured MTU in bytes. In this example, we use 1500 bytes as MTU. To find our MSS, we subtract the IP (20 bytes) and TCP (also 20 bytes) headers from the MTU.

```
WAN Link MTU = 1500
```

```
WAN Link MSS = 1460 [MTU(1500) - IP(20) - TCP(20)]
```

Using these numbers, we can now determine our GRE tunnel MTU and MSS. Recall from earlier in the guide that a GRE tunnel adds a new IP header (20 bytes) and a GRE header (4 bytes). So, we remove the IP and GRE headers from the WAN link MTU, and calculate the MSS as we did before.

```
GRE MTU = 1476 [WAN Link MTU(1500) - IP(20) - GRE(4)]
```

```
GRE MSS = 1436 [GRE MTU(1476) - IP(20) - TCP(20)]
```

## GRE Monitoring and Failover

The GRE protocol does not contain a detection mechanism for tunnel failures. This means you need to deploy other mechanisms to detect tunnel failure or service interruptions. Using these tools, you can enable automatic failover from your primary to your backup tunnel.

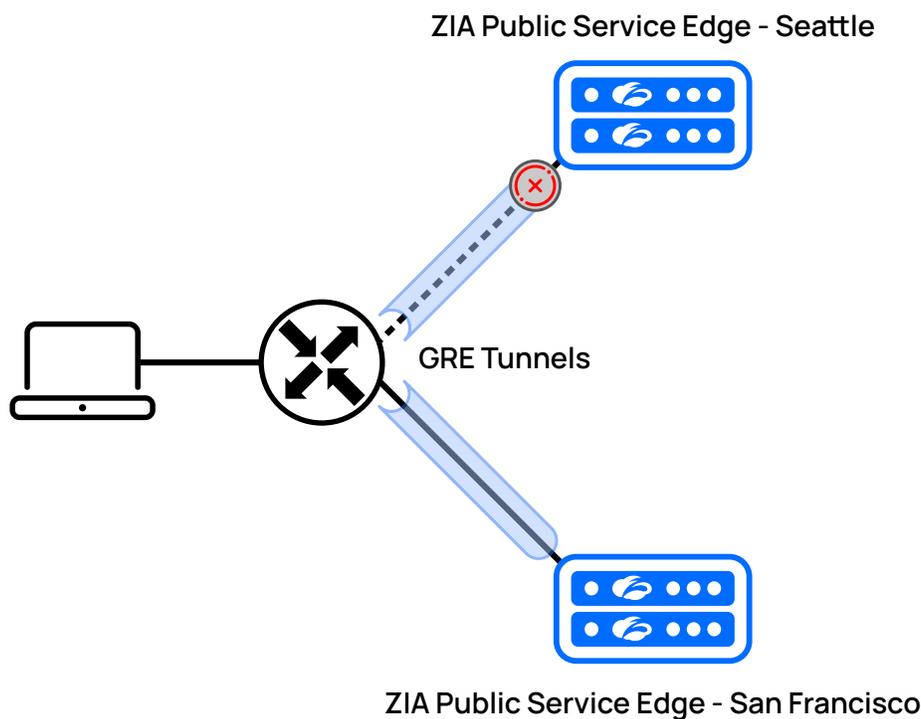


Figure 13. GRE failover in the event a data center becomes unreachable

Zscaler recommends combining multiple tools to check for tunnel and service availability. While the GRE standard does not have built-in protection mechanisms, many vendors have built tools to monitor GRE, which can be used to trigger a switch from the primary tunnel to the backup tunnel.



Many of these mechanisms require that you use network tools to check if a service or host is reachable. Zscaler requires that this be a host you control. Using hosts such as Google trigger protection mechanisms against the Zscaler IP address space. This impacts all Zscaler users, as Google triggers a Captcha for all user transactions.

Many router vendors have implemented GRE keepalive mechanisms to detect tunnel failure. Essentially, the routers send a network ping to a remote station. If a configured number of replies are missed, the service triggers a failover. This tool confirms that a path between the two devices is working. Ideally you pick a target site, like a branch location, to use as your connection test. Note that your vendor implementation can differ slightly, so be sure to check their documentation on implementation of GRE keepalives.

Because the GRE keepalive is trying to reach an interface on another host, it can only tell us that there is a path to the host, and that the host is responding. It cannot check for application availability. Layer 7 health checks are supported by some vendors as well, which can be used to trigger a failover of tunnels when a path is up but a service is unreachable. Zscaler offers a Layer 7 health check target for each cloud, using the target `http://gateway.[your-zscaler-cloud-name].com/vpntest`.

Because each implementation is proprietary, it is beyond the scope of this guide to give detailed advice on configuration. Zscaler recommends you check with your router or SD-WAN vendor to determine what support is available for service checks and failover.

## GRE Summary

Zscaler recommends that you tunnel traffic via GRE without a NAT device between your users and the ZIA Service Edge. Surrogate IP should be enabled, as well as user authentication at each location. This gives you the most visibility and flexibility in your policy and logging, as well as the highest levels of performance. Zscaler requires two GRE tunnels be configured to two different data centers in an active/standby configuration. As a best practice, Zscaler recommends implementing monitoring and automated tunnel failover. Check with your gateway router or SD-WAN vendor for availability and configuration settings.

NAT	Disable or place after the tunnel entry point
Surrogate IP	Enable
Authentication at location	Enable
Number of tunnels	N+1

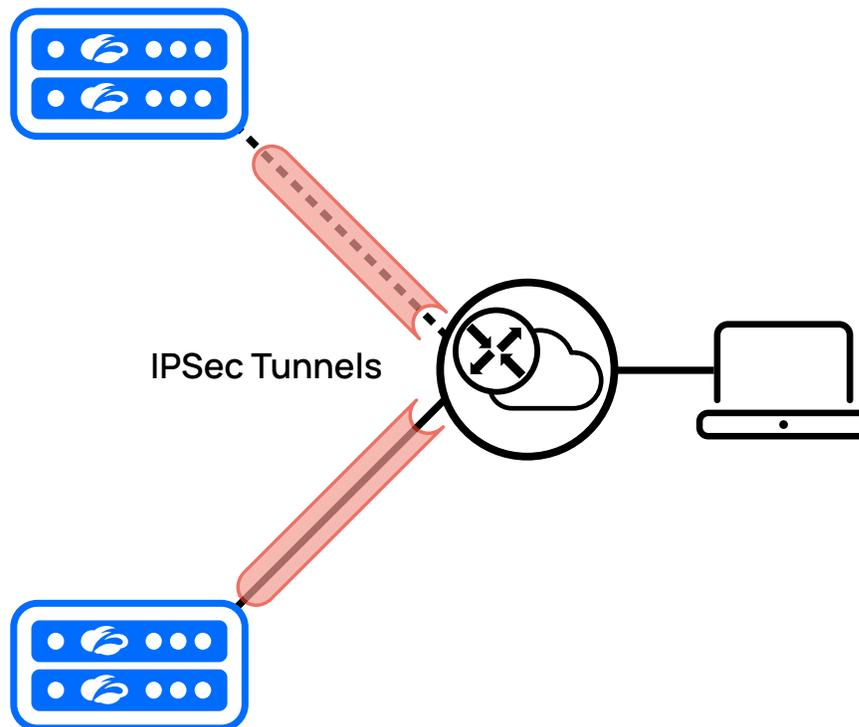
- Learn more about Zscaler's [GRE implementation and configuration](https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre) (<https://help.zscaler.com/zia/about-generic-routing-encapsulation-gre>), including guides for Cisco and Juniper routers.
- Read about SD-WAN vendor configuration files on our [help site](https://help.zscaler.com/zscaler-technology-partners/network) (<https://help.zscaler.com/zscaler-technology-partners/network>).
- View details of the [GRE protocol](https://tools.ietf.org/html/rfc2784) (<https://tools.ietf.org/html/rfc2784>).

## IPSec VPN

IPSec VPNs are traditional virtual private network tools. They create a tunnel between endpoints (or peers) and encrypt that traffic with shared keys before it crosses the tunnel. In traditional remote networking, this was often done to protect an organization's internal data as it transited the public network between sites in a hub-and-spoke design.

Today almost all internet applications and services have moved to being encrypted with Transport Layer Security (TLS), or its predecessor Secure Socket Layer (SSL), for secure communication. This means your traffic is already encrypted from your user's device to the destination, and the additional overhead of IPSec adds little value. The only benefit is that unsecured sites using HTTP are encrypted between the Service Edge and Zscaler. However, that same connection goes out as HTTP and is still unencrypted after it leaves the ZIA Service Edge.

### ZIA Public Service Edge - Seattle



### ZIA Public Service Edge - San Francisco

Figure 14. IPSec tunnels extend from your SD-WAN device like GRE tunnels

This additional layer of security comes at a cost of reduced performance. Traffic must first be decrypted from the IPSec tunnel, and then decrypted again to process the TLS/SSL traffic at the ZIA Service Edge. ZIA Service Edge devices use hardware acceleration to achieve high rates of throughput, but still don't match the low overhead of GRE tunnels.



Increased security and trust requirements mean additional configuration by both your organization and Zscaler. For this reason, IPSec requires an additional subscription.

Zscaler only recommends using IPSec when mandated by policy or regulation, or because your internet gateway does not support GRE tunnels. View a [list of internet gateways \(https://help.zscaler.com/zia/about-ipsec-vpns#zscaler-interoperability-list\)](https://help.zscaler.com/zia/about-ipsec-vpns#zscaler-interoperability-list) known to be interoperable with Zscaler.

## Confidentiality, Integrity, and Authentication

Central to IPSec, VPNs are the concepts of confidentiality, integrity, and authentication. Each of these plays a part in setting up a secure IPSec tunnel between peers and protecting data in transit. Make decisions about encryption ciphers and hash digests based on your policy that relates to these functions.

First, let's define what these terms mean in the context of IPSec and what protections they provide:

- Confidentiality – Ensures that data cannot be read by unauthorized parties.
- Integrity – Verifies that data was not modified during transit.
- Authentication – Verifies the identity of peers.

The negotiation for confidentiality and integrity is handled by the Internet Key Exchange (IKE) protocol, either version 2 (IKEv2) or the older version 1 (IKEv1). Depending on which version you select, you will have different options for confidentiality, integrity, and authentication configuration. Zscaler recommends using IKEv2 wherever possible. IKEv2 has better performance and fixes vulnerabilities that have been found in IKEv1. Check with your gateway router or SD-WAN vendor on their IKE version support.

IKE performs two phases when setting up an IPSec connection:

- Phase 1 – Authenticate peers and set up a secure channel for exchange of Phase 2 messages. This phase protects the messages in Phase 2.
- Phase 2 – Negotiate parameters and set up security associations (SA). This is where tunnel mappings are established between peers.

The selections you need to make around integrity and confidentiality will have different parameters for each phase.

Authentication is the one parameter that is not negotiated between peers. This is because it has already been configured on each side prior to the connection being set up beforehand. Zscaler recommends using pre-shared keys (PSK) with long keys. These are configured in the ZIA Admin Portal and can be managed via an uploaded CSV. The pre-shared keys are used in a Diffie-Hellman (Public Key Encryption) exchange to protect the rest of the key exchange. After authentication and negotiation, peers switch to a shared key for faster encryption and decryption.

Integrity's job is to ensure that the message has not been changed in transit between peers. This step runs a hash function over the message to ensure the message received is the same one that was sent. The following table outlines Zscaler's recommendation for integrity settings:

Integrity	Phase 1	Phase 2
IKEv2	SHA-256 or SHA-1	MD5
IKEv1	SHA-1	MD5

Confidentiality is where you would normally encrypt your traffic in transit. In a traditional IPSec VPN scenario between two remote sites, you want to use the strongest encryption possible. This case, however, is different in that we're using IPSec as a forwarding tunnel and not for data security. Its only job is to get your traffic from your edge device to the ZIA Service Edge. The table below outlines Zscaler's recommended settings for Phase 1 and Phase 2.

Confidentiality	Phase 1	Phase 2
IKEv2	AES-256	NULL
IKEv1	AES-128	NULL

You might find it surprising that Zscaler recommends NULL or no encryption in Phase 2. Remember that all the traffic you are placing in this IPsec tunnel eventually goes out to the internet the same way it entered the tunnel. No additional security protections are applied to it. As mentioned previously, almost all of this traffic is already encrypted with TLS/SSL. For this reason, Zscaler recommends you set confidentiality to NULL in Phase 2.

## Perfect Forward Secrecy

A feature that is often used in IPsec VPNs is Perfect Forward Secrecy (PFS). This is a mechanism where the SA keys negotiated in Phase 2 are automatically renewed. This renewal takes place after an amount of traffic or time has passed. The idea behind PFS is that should an attacker break your key, they can only read the material protected by that key. Zscaler recommends disabling PFS as we are using NULL encryption, so there is no reason to refresh the key.

## NAT Transversal

Zscaler supports IPsec when the originating endpoint is placed behind a NAT device, called NAT Transversal or NAT-T. The purpose of NAT-T is to allow IPsec sessions to be established when one of the peers is behind a NAT device.

When enabled, peers detect if a NAT is present in the path during Phase 1 negotiations. They then verify that both peers support NAT-T. If so, they encapsulate the packet inside a new set of NAT-T headers.

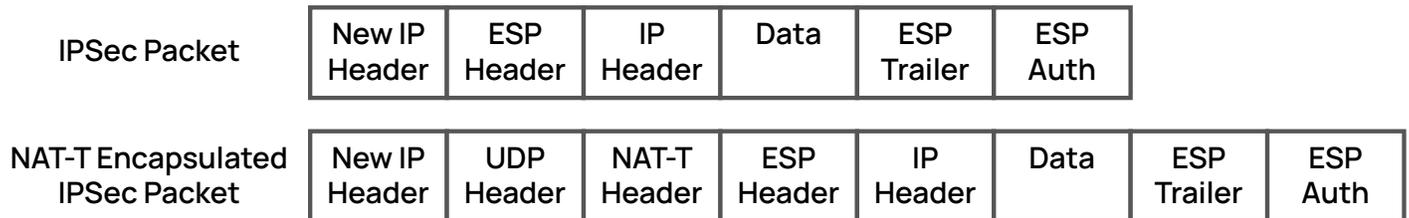


Figure 15. IPsec and NAT-T encapsulation

The first is a new IP header. This retains the same tunnel destination but changes the protocol to the User Datagram Protocol (UDP). The UDP header sets its source port as 500, and retains the IPsec peer's destination port. Finally, a NAT-T header is inserted behind the UDP header. The packet transits the NAT device, and the new headers are stripped off by the peer. Zscaler recommends enabling NAT-T.

## Bandwidth

Each IPSec tunnel to a ZIA Public Service Edge or ZIA Private Service Edge is formed from a single IP address. Each tunnel configured has a maximum throughput of 400 Mbps. If you need additional throughput from your site, you can add additional IP addresses and tunnels to your configuration.

### ZIA Public Service Edge - Seattle

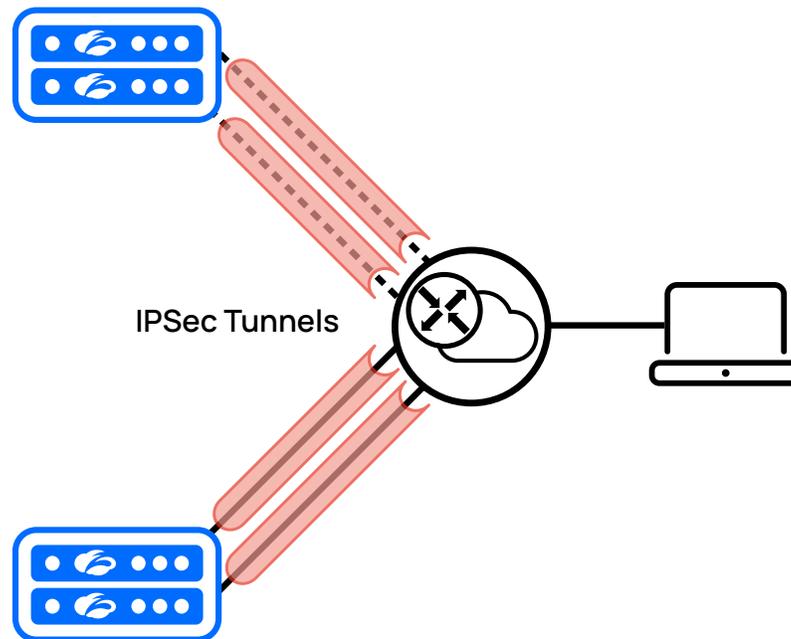


Figure 16. Bandwidth for IPSec can be increased by adding additional links

In general with IPSec VPNs, you are connecting to your own network from a remote location, so there is no need to translate the IP addresses of stations. Ensuring that NAT is not in use allows for more granularity with policy and enables the use of [Surrogate IP](#) discussed later in this guide.



The ZIA Virtual Service Edge does not support termination of IPSec connections. If your organization is using a ZIA Virtual Service Edge, you need to build tunnels using GRE or the [Zscaler Client Connector](#) software discussed later in this guide.

## Redundancy

As with GRE, Zscaler requires that you configure redundant tunnels to two different data centers. One tunnel will operate in active mode, and another in standby mode. In the event of a loss of connectivity to the primary data center, your gateway router or SD-WAN device will fail over to the backup data center.

### ZIA Public Service Edge - Seattle

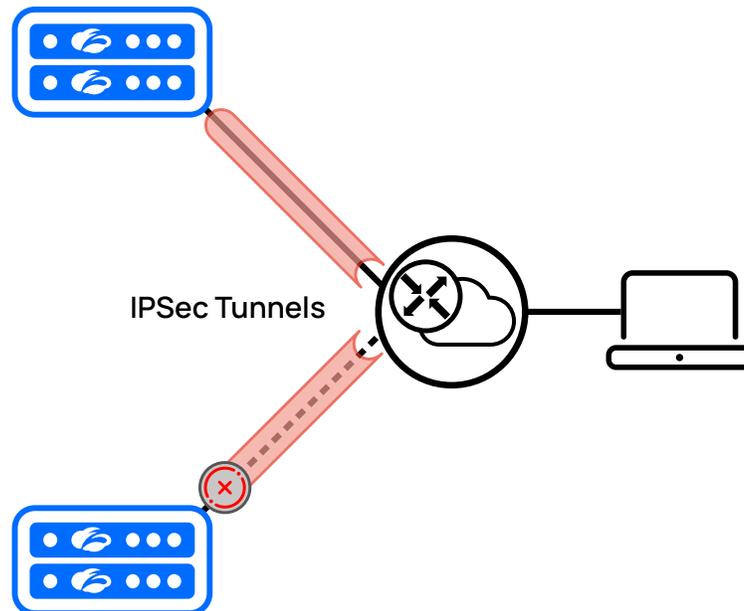


Figure 17. IPsec failover when a data center becomes unreachable

Zscaler requires that you build primary and secondary IPsec tunnels from each internet egress point. Additionally if you have multiple ISPs at a single location, primary and secondary tunnels must be built from each ISP.

## IPsec Setup

Configuration of IPsec requires that you configure credentials for each peer, as well as the parameters they use to negotiate encryption ciphers and hash digests. You need the following pieces of information to complete this setup:

- IP address or hostname of your internet gateway or SD-WAN device
- Shared secret
- IP addresses or hostnames of the ZIA Public Service Edge or ZIA Private Service Edge. You can find the ZIA Public Service Edge IP addresses and hostnames on the [Zscaler Help Portal \(https://help.zscaler.com/zia/locating-hostnames-and-ip-addresses-zens\)](https://help.zscaler.com/zia/locating-hostnames-and-ip-addresses-zens).
- Your decisions on IPsec parameters

You must first add VPN credentials and select an authentication type. Zscaler supports adding up to 16,000 credentials on the platform. You can also add and remove credentials by importing credentials via CSV. Zscaler supports two identification methods: FQDN and IP address. You assign credentials to a location for use in authenticating the peer gateway.

The two credential types have slightly different requirements: FQDN has a user ID field, whereas IP simply uses the public IP address of the peer. FQDN and IP use a pre-shared key (up to 255 characters). Check with your internet gateway vendor for supported modes of operation. You can view a list of [deployment guides \(https://help.zscaler.com/zia/configuring-ipsec-vpn-tunnel\)](https://help.zscaler.com/zia/configuring-ipsec-vpn-tunnel) for internet gateways using IPsec.

## Monitoring and Failover

IPSec supports a protocol called Dead Peer Detection (DPD) to assist with failover. The purpose of DPD is to alert the system when its peer has become unresponsive. Instead of waiting for the SA to expire, the systems instead check periodically for station responses using a HELLO message. This is a more robust protocol than simply pinging a peer to detect presence, as the remote peer must respond correctly to the HELLO message.

DPD goes one step further by trying to reduce the overhead of detection by using traffic as an indicator of the peer existing and continuing to function. Traditional VPN concentrators typically terminate very large numbers of stations when they are used for user access. In this case, sending heartbeats and maintaining timers for each station introduces a large amount of traffic and computational overhead.

Instead, DPD is used only when a peer needs to send traffic after the tunnel has been inactive for some time. When the systems are exchanging valid IPSec traffic, there is no need to continue to check that the peer is active. Traffic passing in a bidirectional manner is enough to prove the station is available.

Zscaler recommends enabling DPD on your connections so that any downtime due to a lost connection is minimal.

## IPSec Summary

Traditional IPSec VPNs have been used for decades to protect information between peers. Zscaler supports IPSec as a tunnel mechanism for forwarding, but the need for full security is reduced since all traffic is headed for the internet instead of a corporate data center. The continued growth of a TLS/SSL secured internet ensures that your data is protected from casual interception without the additional overhead of IPSec encryption.

Zscaler recommends using IPSec when:

- It is mandated by policy or regulation.
- Your internet gateway does not support GRE tunnels.
- Your internet gateway does not have a static IP address.

Zscaler recommends the following settings for IPSec:

Integrity	Phase 1	Phase 2
IKEv2	SHA-256 or SHA-1	MD5
IKEv1	SHA-1	MD5
Confidentiality		
IKEv2	AES-256	NULL
IKEv1	AES-128	NULL
Perfect Forward Secrecy	Disable	
NAT Transversal (NAT-T)	Enable	
Dead Peer Detection (DPD)	Enable	
NAT	Disable or place after the tunnel entry point	
Surrogate IP	Enable	
Authentication at location	Enable	
Number of tunnels	N+1	

- View a list of [deployment guides \(https://help.zscaler.com/zia/configuring-ipsec-vpn-tunnel\)](https://help.zscaler.com/zia/configuring-ipsec-vpn-tunnel) for internet gateways using IPSec.
- Learn more about [Dead Peer Detection \(https://datatracker.ietf.org/doc/html/rfc3706\)](https://datatracker.ietf.org/doc/html/rfc3706).

## Dedicated Proxy Ports

As an alternative to tunneling, Zscaler supports dedicated proxy ports. In this model, your internet gateway or SD-WAN device forwards user traffic to a specific port number assigned to your organization. This is an additional subscription service available from Zscaler, and is only recommended when your gateway device is not capable of providing GRE or IPsec tunnels.

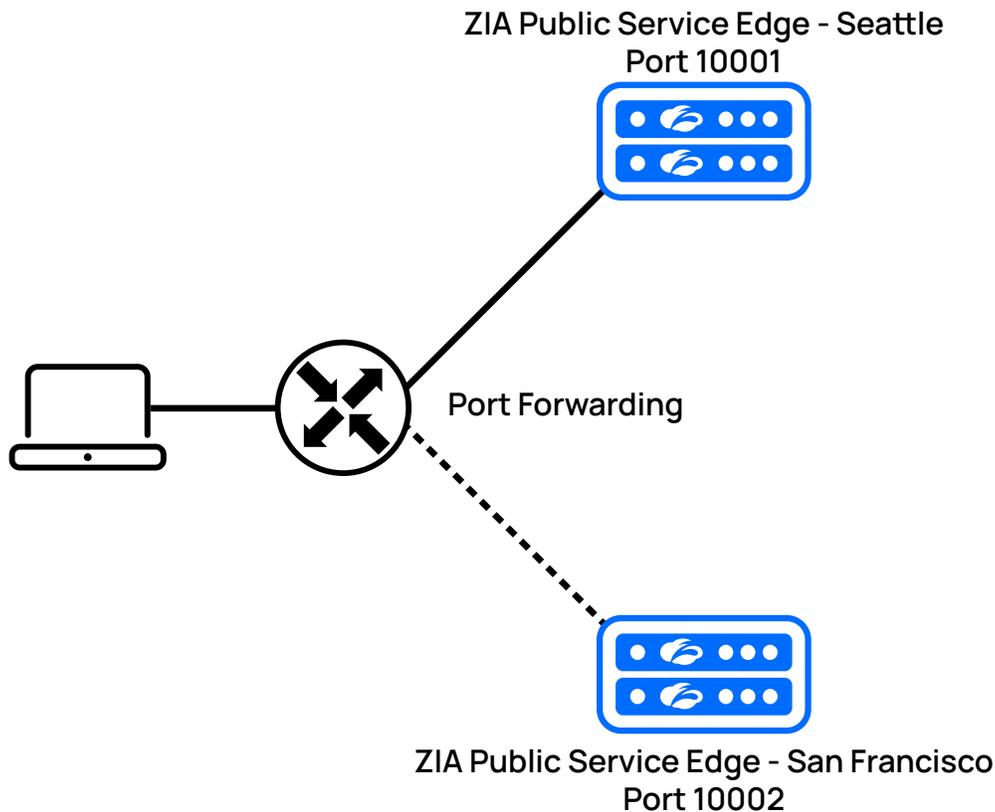


Figure 18. Dedicated proxy ports at two data centers

The ports are tied to locations you have designated. Because anyone can send traffic to a dedicated proxy port, Zscaler requires that your users authenticate to prevent misuse of the system. For remote users, they need to use a PAC file to access proxy ports. A more robust method for remote users is to deploy Zscaler Client Connector to your remote stations.

Configuration of dedicated proxy ports is handled on the Location Management page of the ZIA Admin Portal. You add your assigned port(s) and select the Enforce authentication checkbox. Zscaler recommends enabling Surrogate IP on proxy ports and for remote users.

NAT	Disable
Surrogate IP	Enable
Authentication at location	Enable

Learn more about [configuring proxy ports](https://help.zscaler.com/zia/configuring-dedicated-proxy-ports) (<https://help.zscaler.com/zia/configuring-dedicated-proxy-ports>).

Learn more about [configuring locations in ZIA](https://help.zscaler.com/zia/configuring-locations) (<https://help.zscaler.com/zia/configuring-locations>).

## Surrogate IP for Fixed Site Deployments (Recommended)

Surrogate IP is a feature within ZIA that maps private IP addresses to authenticated users. This feature allows you to apply policies assigned to a user when that user’s traffic cannot be authenticated. Without this service enabled, the location policies apply to unauthenticated user traffic. Some examples of traffic that would fall into this category include:

- Web-based applications that do not support cookies like Google Earth.
- HTTPS traffic that you have chosen not to decrypt.
- Transactions from unknown user agents, such as applications with a custom user agent or that are not web-based.

When a user authenticates to ZIA, the service makes note of the private IP address in use. Any other traffic that ZIA sees coming from the same private IP address is mapped to the user and their applied policies. Users can have multiple private IP addresses mapped to them at the same time, such as a laptop and smartphone. In the ZIA logs, you can see the user mapped to all of these IP addresses so that you don’t have to correlate across them.

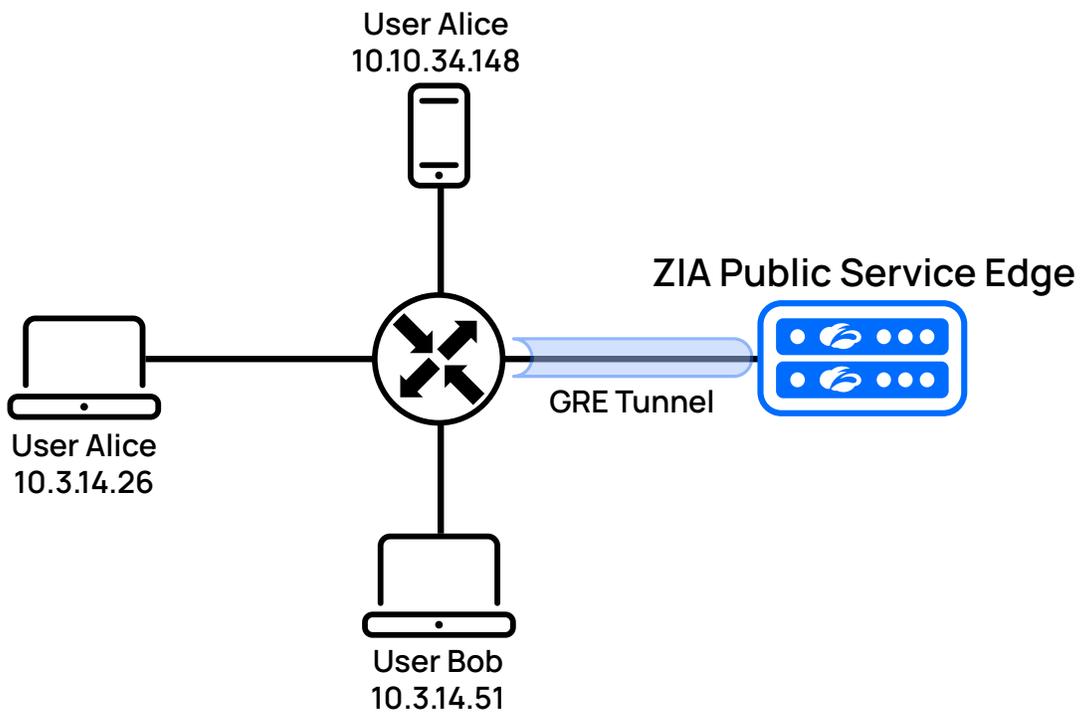


Figure 19. User credential and station IP addresses

In the previous image, the policy for the user and logging from ZIA’s perspective would look like the following table:

Source	Policy and Logging
Authenticated as Alice	User Alice
Authenticated as Bob	User Bob
10.3.14.26	User Alice
10.3.14.51	User Bob
10.10.34.148	User Alice

ZIA continues to map users to these private IP addresses to the user until:

- The user logs out.
- The configured idle time is reached.
- A new user authenticates using the same private IP address.

To successfully use Surrogate IP, the following requirements must be met:

- A GRE or IPSec tunnel without NAT, or a dedicated proxy port subscription.
- The location the user is coming from must have authentication enabled.

If ZIA sees different users log in and send traffic from the same IP address within one minute, Surrogate IP is disabled for 5 minutes. All traffic is assigned that location's policies instead of their more specific policies that would be assigned to the user. This only applies to traffic that cannot be authenticated. Therefore, NAT cannot be used with Surrogate IP, as all users look like they are using the same IP address. As mentioned previously, the system must be able to map single IP addresses to single users.

Zscaler recommends enabling Surrogate IP for all fixed site locations.

Learn more about [Surrogate IP](https://help.zscaler.com/zia/about-surrogate-ip) (<https://help.zscaler.com/zia/about-surrogate-ip>).

## Mobile Users - Explicit Forwarding

Today, part- and full-time work-from-home employees are rapidly becoming the norm in many organizations. Users connecting while outside the enterprise is now commonplace at almost every level of the organization.

At the same time, the shift to cloud services and infrastructure over a corporate data center continues to gain momentum. These applications are either being outsourced to Software as a Service (SaaS) vendors such as Salesforce, or moving to virtual data centers in Infrastructure as a Service (IaaS). In both cases, the fastest route from your user to the application is direct internet access.

What has lost popularity are the remote access VPNs backhauling traffic to a corporate data center for inspection and policy enforcement using hardware appliances. This hub-and-spoke network design adds latency to web applications. Using ZIA eliminates this backhaul by having the client connect to the nearest ZIA Service Edge able to service their traffic.

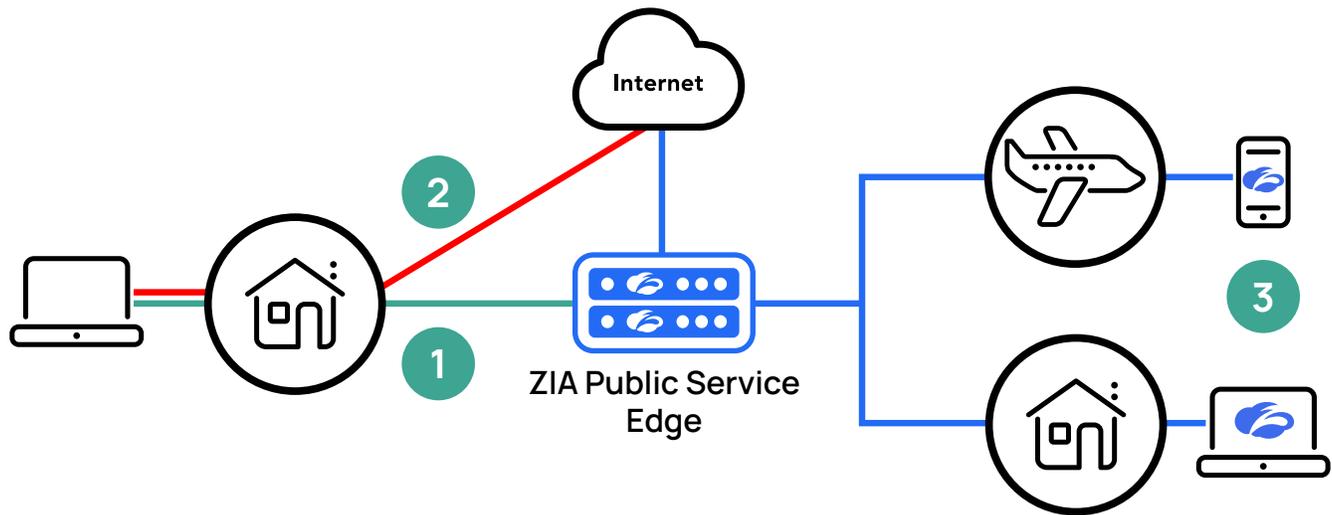


Figure 20. Zscaler Client Connector or PAC files for mobile and remote workers

Zscaler offers two solutions to connect mobile users to the ZIA Service Edge: Zscaler Client Connector or PAC files. Zscaler recommends Zscaler Client Connector over PAC files. This software package is included in your Zscaler subscription, and is required for other Zscaler services such as ZPA and ZDX.

## Zscaler Client Connector (Recommended)

Zscaler Client Connector is a lightweight agent that sits on your user's endpoint, supporting ZIA, ZPA, and ZDX services. Depending on which Zscaler services you subscribe to and what applications the end user is accessing, the traffic is automatically forwarded to the correct service. Unlike VPNs, there is no need to backhaul traffic to your data center first. Users go straight to the cloud by connecting directly to a ZIA Service Edge.

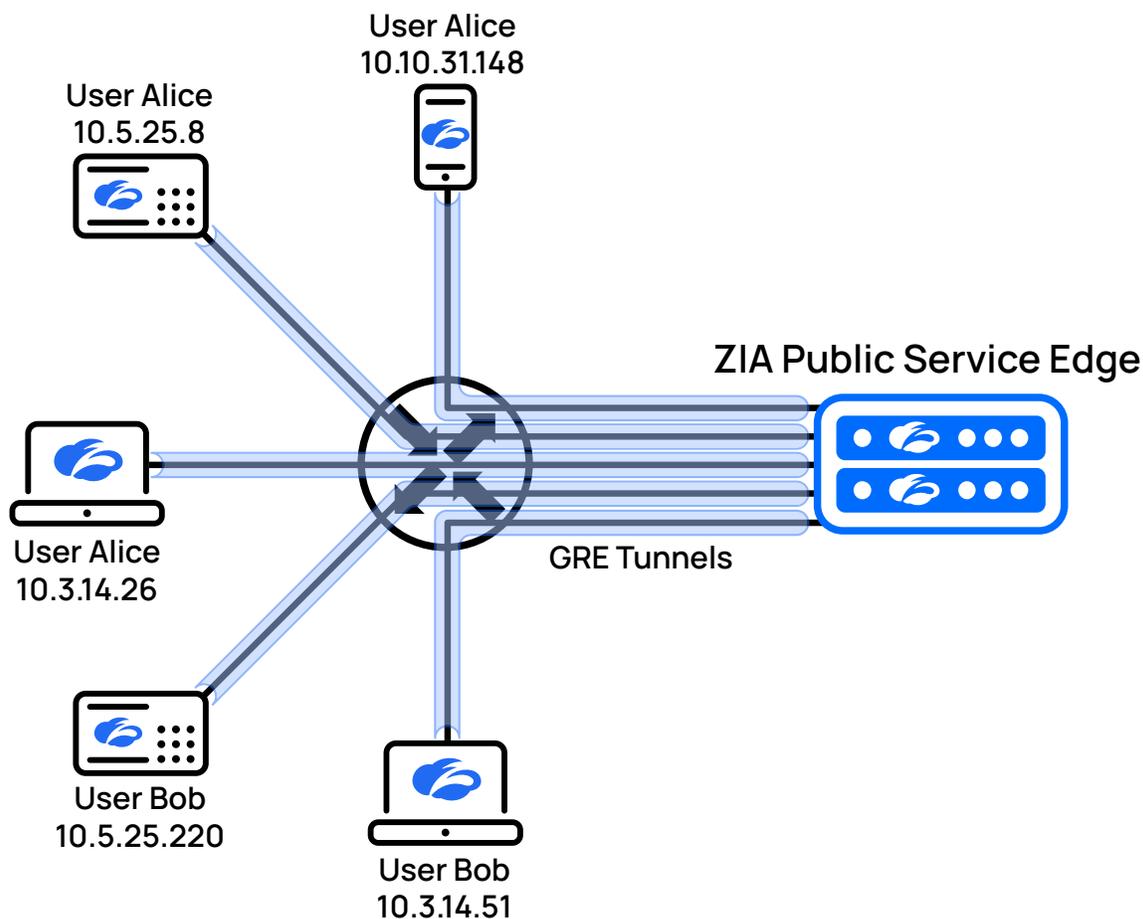


Figure 21. Zscaler Client Connector works with many operating systems and devices

When installed on a user's device, Zscaler Client Connector creates a virtual network adapter. For the ZIA use case, by default your user traffic is tunneled directly to the nearest ZIA Service Edge for inspection.

When your users attempt to use the internet, this virtual adapter captures that traffic flow. Zscaler Client Connector then uses geolocation to determine the closest ZIA Service Edge node, and builds a lightweight tunnel called a Z-Tunnel to that node. The user traffic is then placed inside the tunnel and forwarded to the ZIA Service Edge for inspection and policy enforcement.

Zscaler Client Connector can also be set to disable itself temporarily when it detects that it is on a trusted network, or if a captive portal is blocking access to the internet.

## Operating System Support and Installation

Zscaler Client Connector supports a wide variety of the most common operating systems, including Windows, macOS, CentOS, Ubuntu (20.04+), iOS, and Android. Zscaler Client Connector can be deployed via direct download or Mobile Device Management (MDM) for silent installs.

The desktop versions of the Zscaler Client Connector agent provide administrators the ability to preconfigure the client. This eliminates user error and help desk calls; administrators only need the ability to install software. You can host the software on your intranet, a publicly available location, or through MDM push. When installed, your users are required to log in to the client to access resources. After authentication successfully completes, the client device is available for inspection in the Zscaler Client Connector Admin Portal.

You can prevent users from disabling Zscaler Client Connector or uninstalling the software by using an admin-provided password. This ensures that your users are always protected when accessing internet resources.

## Authentication Options

Zscaler Client Connector supports several authentication options. Except for Kerberos, all authentication methods supported by ZIA are also supported by Zscaler Client Connector.

While ZIA supports many other options, not all of these are suitable for large scale deployments. The complete list of supported authentication options includes Identity Federation using SAML, directory server, Zscaler Authentication Bridge, one-time link, one-time token, and passwords.

Zscaler recommends the use of SAML where possible. SAML is a modern, flexible, and robust authentication system and is used widely by application vendors. SAML and Zscaler Client Connector support two-factor authentication as well for added security.

Learn more about [authentication options](https://help.zscaler.com/zia/choosing-provisioning-authentication-methods) (<https://help.zscaler.com/zia/choosing-provisioning-authentication-methods>).

View a list of Zscaler's leading [authentication provider partners](https://www.zscaler.com/partners/technology/identity-access-management) (<https://www.zscaler.com/partners/technology/identity-access-management>).

## Trusted Network Detection

Anyone who has supported traditional user access VPNs knows the difficulties associated with users not understanding when to launch their VPN client. Users at home forget to log in, generating help desk calls about access. Other users log in while on the organization's campus, slowing their connection and using up licenses that remote users need to connect.

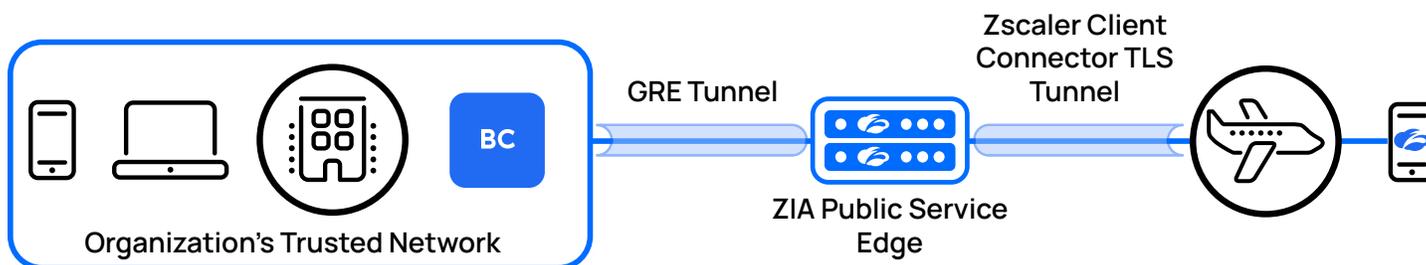


Figure 22. Trusted network connection builds a Zscaler Client Connector tunnel only when on an untrusted network

Zscaler Client Connector contains a feature called trusted network detection. You configure the software to understand when it is on one of your organization's network segments and disable itself. This allows your traffic to be tunneled by your internet gateway using GRE or IPsec. When the software detects that it is not connected to a trusted network segment, it automatically tries to connect to the nearest ZIA Service Edge.

Zscaler Client Connector uses the following options to determine if a client is on a trusted network. Except for Pre-Defined Trusted Networks, you can specify a match on ANY or ALL configured criteria for a more robust detection:

- DNS Server (recommended) – Configure a list of IP addresses which are your internal corporate DNS servers. Zscaler Client Connector verifies at least one DNS server can be reached.
- DNS Search Domains (recommended) – Enter a list of search domains. Zscaler Client Connector compares this list to the search domain of the active network adapter.
- Hostname and IP – A hostname and the IP addresses where the hostname resolves when users are on the corporate network.
- Pre-Defined Trusted Networks – Configure all your trusted networks, then add them to your configuration. This detection option cannot be combined with any of the other three options in the policy, and is the only criteria used if configured.

After your criteria is defined, Zscaler Client Connector automatically disables itself when it determines it is on a trusted network. Any time a network adapter connects to a network, it reassesses the trusted/untrusted criteria and determines if Zscaler Client Connector should be active or not.

Zscaler recommends using DNS Server and/or DNS Search Domains for detecting trusted networks. These are the most static checks possible. Combining them gives you the most assurance of being on a trusted network.

## Fail Open Options

There are times where you might choose to have Zscaler Client Connector fail open. The supported options are:

- Captive portal detected – A captive portal is blocking access to the internet.
- ZIA Service Edge is not reachable – Something on the network is preventing the client from connecting to the nearest ZIA Service Edge.
- Z-Tunnel setup issues – Zscaler Client Connector is unable to establish a Z-Tunnel, the lightweight tunnel established from the client to the ZIA Service Edge used to forward traffic.

The use of captive portals is widespread in hospitality Wi-Fi networks, such as those found in airports and hotels. Captive portal detection allows the client to disable itself for a short time while the user authenticates to the portal, and then re-enables itself after the timer expires. You can set a value from 1 to 60 minutes. Zscaler recommends enabling captive portal detection. The disable time setting should be set as low as you feel is reasonable for users to interact and pass a captive portal registration system.

The ZIA Service Edge unreachable use case occurs when the client is unable to establish a connection due to network configuration. This could be a completely isolated network, such as in a lab setting. In this case, you can choose to either fail open, sending traffic directly to the internet/local network, or fail close and disable internet access. The use of this feature should be governed by your organization's policy, and how you want to treat user traffic when it cannot be secured on the internet. Zscaler Client Connector continues to try to establish a connection in the background and re-enables when connected.

The Z-Tunnel setup issue occurs when Zscaler Client Connector can locate a ZIA Service Edge but cannot establish a tunnel. This could be a case where the ZIA Service Edge resolves with DNS, but the tunnel setup is blocked by a network firewall. In this case, you can choose to either fail open, sending traffic directly to the internet/local network, or fail close and disable internet access. The use of this feature should be governed by your organization's policy, and how you want to treat user traffic when it cannot be secured on the internet. Zscaler Client Connector continues to try to establish a connection in the background and re-enables when connected.

Learn more about [configuring fail-open settings \(https://help.zscaler.com/z-app/configuring-fail-open-settings-zscaler-app\)](https://help.zscaler.com/z-app/configuring-fail-open-settings-zscaler-app).

## TLS/SSL Inspection with Zscaler Client Connector

Zscaler recommends inspecting all traffic your users are sending, including that encrypted with TLS/SSL. As mentioned previously, almost all traffic today is encrypted with some version of TLS or the older and more vulnerable SSL. To do this, your clients need to trust the certificates issued by the ZIA Service Edge for the services they want to reach. This allows the ZIA Service Edge to decrypt and inspect the traffic before re-encrypting with the service's actual certificate.

Certificate installation for ZIA can be handled automatically with Zscaler Client Connector. The root certificate for ZIA is added to your client's trust store, enabling the station to trust the short-term certificates issued by the ZIA Service Edge. For mobile devices, you might need to manually add the root certificate to your device. You need to enable inspection for mobile traffic in your policy.

Learn more about [TLS/SSL Inspection with Zscaler Internet Access \(https://help.zscaler.com/zia/tls-ssl-inspection-zscaler-internet-access\)](https://help.zscaler.com/zia/tls-ssl-inspection-zscaler-internet-access).

Learn more about [configuring TLS/SSL inspection \(https://help.zscaler.com/zia/configuring-ssl-inspection-policy\)](https://help.zscaler.com/zia/configuring-ssl-inspection-policy).

## Zscaler Client Connector Summary

Zscaler Client Connector is a lightweight agent that allows you to directly forward your users' traffic to a ZIA Service Edge when users are away from the office. This agent is included as a part of your ZIA subscription.

Unlike traditional VPNs, the client can auto-detect when it needs to enable or disable itself. This ensures that your user traffic goes to the closest ZIA Service Edge for inspection and policy enforcement without the need to backhaul traffic to the organization's offices. Available for a wide variety of operating systems, Zscaler recommends installing Zscaler Client Connector on all your end user devices.

## PAC Files

A Proxy Auto-Configuration (PAC) file is a set of instructions written in JavaScript that tell a browser how to reach a web proxy. The original design for PAC files was created in 1996 by the Netscape corporation and released in Netscape Navigator 2.0. Today, PAC files are supported by every major browser in the market.

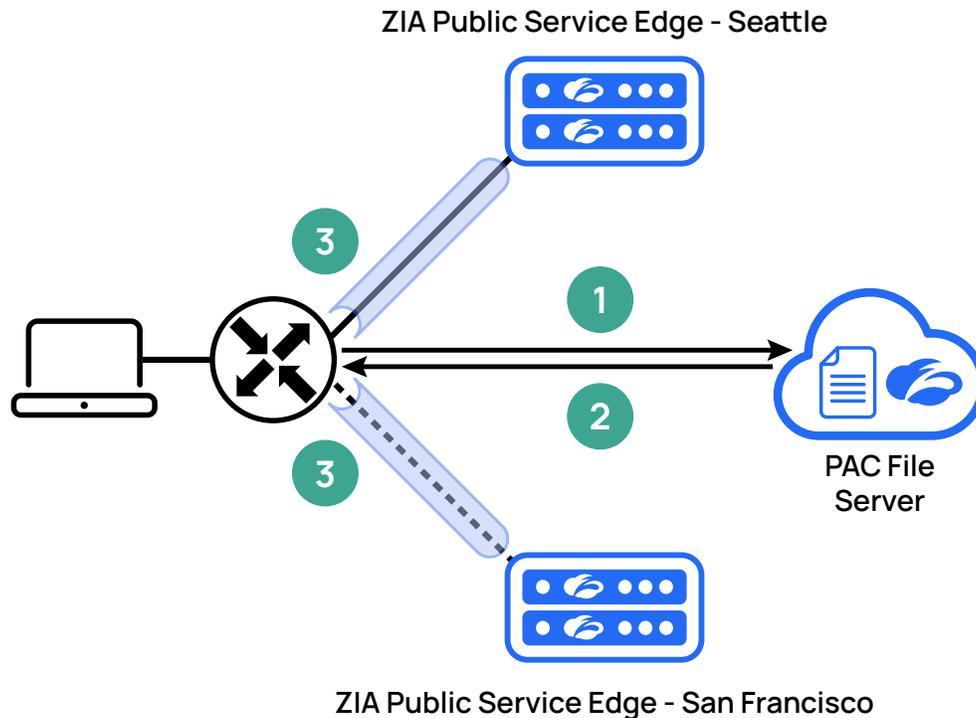


Figure 23. PAC file transaction for web traffic forwarding

To use a PAC file, a browser is preconfigured with the URL where it can retrieve the ZIA PAC file. When launched:

1. The browser first retrieves the PAC file by reaching out to the PAC file server.
2. Based on the geolocation of the station requesting the PAC file, one is returned with the two nearest ZIA Public Service Edge locations. The Zscaler PAC file uses variables that are replaced with the nearest node when they are returned.
3. The web browser then forwards traffic through the nearest ZIA Public Service Edge.

The JavaScript within a PAC file is parsed in top-down and first match wins. Commands are matched against hostnames or IP addresses. There are two options for PAC file matches:

- **DIRECT** – The browser should send the traffic directly and bypass the proxy.
- **PROXY** – Proxy the traffic. Note that the proxy statement also requires the IP address of the proxy and the port number.

Zscaler supplies a default PAC file as part of your subscription. This file contains the minimum command set needed to reach the nearest ZIA Service Edge. Both a primary and secondary gateway are included in this file. You can also upload a custom PAC file to the ZIA Admin Portal with customized forwarding commands. See [Custom PAC Files](#) for more information.

PAC files as an older technology come with several limitations. This is a tool designed before most stations were mobile and when computer systems generally only had a single network interface. So, Zscaler can only recommend the use of PAC files where the traffic is all web-based, and no other forwarding option (Zscaler Client Connector, GRE, IPSec, or dedicated proxy port) is feasible for the device.



Not all applications support PAC files. If you are using native non-web applications that access internet hosts, you should check with your vendor on support for PAC files.

## Custom PAC Files

Starting with the default PAC file found in the ZIA Admin Portal, Zscaler supports the upload of customized PAC files. However, the creation of PAC files being a legacy technology requires that you spend time testing changes before pushing them out to user. Additionally, there are tradeoffs to be made between size and speed of execution. The maximum file size ZIA supports is 256 KB. A long PAC file takes time to parse, and a complex regex to save space can slow down execution as well.

If you plan to customize your PAC file, Zscaler strongly encourages your JavaScript developers to follow our development best practices. Your developers should be familiar with these practices before they begin writing code for the PAC file. The following help documents should be shared with your development team:

- [PAC file best practices](https://help.zscaler.com/zia/best-practices-writing-pac-files) (<https://help.zscaler.com/zia/best-practices-writing-pac-files>)
- [Writing a PAC file](https://help.zscaler.com/zia/writing-pac-file) (<https://help.zscaler.com/zia/writing-pac-file>)

In the ZIA Admin Portal, you can and should test the PAC file for syntax issues before saving and deploying. After you have fully tested your PAC file, you can deploy it to your users. When the file is published, it takes effect for all users in your organization. Zscaler recommends updating your PAC files during maintenance windows when possible.



Always test your PAC file before deploying it to your users, and make sure you have a backup of your previous file. Using a version control system (VCS) such as Git can help you ensure that you can revert to a known good version if a mistake is made. You must use a plain text editing tool to edit the files. You should not use word processors such as Microsoft Word or Google Docs, as these introduce additional incompatible formatting in their file output.

The final piece of deploying a custom PAC file is obfuscating the URL of the PAC file. Because this file must be made public for your users' machines to access, this means anyone who has the URL can access the content. This can contain information you don't want to be public. The ZIA Admin Portal generates a URL for your use when you enable Obfuscate URL. Zscaler strongly recommends enabling Obfuscate URL.



When you enable obfuscation, you need to update your client machines to the new URL location. Until that occurs, your users use the default system PAC file. This can lead to issues with accessing resources until the client device is updated with the new URL location.

Learn more about [uploading a new PAC file](https://help.zscaler.com/zia/using-custom-pac-file-forward-traffic-zia) (<https://help.zscaler.com/zia/using-custom-pac-file-forward-traffic-zia>).

## PAC File Summary

PAC files are a well supported but ultimately limited technology for forwarding traffic from web browsers to the Zscaler proxy. It is an older technology that has limitations and complexities that you should understand before deploying them. Zscaler recommends that PAC files only be used when another forwarding option is not usable by the client device.

## Specialized Forwarding Cases

Like all complex technologies, there are often special cases that require special consideration. In this section, we cover several items that can affect your organization and its ability to forward traffic to the ZIA service. If your organization needs to use one of these forwarding cases, please contact your Zscaler Account team, technology partners, and local ISPs. Together these groups can explain your options and help you plot a path forward with ZIA.

### Sending Traffic from a Non-Zscaler Source IP

When your traffic is processed by the ZIA service, its source IP address is modified for the ZIA Service Edge that is processing the traffic. This is necessary for the traffic to flow bidirectionally through the ZIA Service Edge.

In some cases, this presents a problem with some applications and vendors that require the source IP address for traffic comes from a preset list of IP addresses. Any traffic from a non-registered IP address is automatically dropped. For example, patient health records can only be accessible from a particular hospital's IP address to help protect patient privacy.

There can also be geographic restrictions based on IP address. If Zscaler does not have a ZIA Public Service Edge in the IP address range in question, the traffic is denied. Examples include government services for a country, where only residents of that country are allowed to access those services.

In these situations, your traffic needs to be handled differently to comply with these IP-based requirements. Zscaler supports three methods for traffic forwarding that can comply with this requirement:

- Use Zscaler IP addresses as source with your vendor – If your vendor accepts Zscaler IP addresses, you simply need to update that application with the Zscaler IP address range. Find a list of [Zscaler IP address ranges by cloud \(https://config.zscaler.com/zscaler.net/cenr\)](https://config.zscaler.com/zscaler.net/cenr).
- Source IP Anchoring – By using ZPA App Connectors in a public cloud such as Amazon AWS, Microsoft Azure, Google GCP, or your on-premises data center, your traffic is routed through that App Connector which then becomes the source IP address for the traffic.
- ZIA Virtual Service Edge or ZIA Private Service Edge – A virtual or physical device is deployed in your data center and managed by Zscaler. This extends the Zscaler cloud into your local environment and allows you to use NAT as traffic leaves the ZIA Service Edge.

Each of these methods requires additional work and subscriptions. The next section briefly describes each option.

## Using Zscaler IP Addresses

With this solution, you update your application with Zscaler's ZIA Public Service Edge addresses. This can also be a subset of the addresses if your access only comes from fixed site locations. This could be as small as your primary and secondary data centers, or as large as the entire ZIA Public Service Edge IP address range. Contact your application vendor to understand the specifics of their requirements.

## Source IP Anchoring

Combining ZIA for traffic inspection and ZPA App Connectors for final forwarding, Source IP Anchoring allows you to control the egress IP for your traffic. In this deployment model, you deploy redundant pairs of ZPA App Connectors in a public cloud with a static IP address. This is the IP address that your applications or websites see as the source for your traffic. On the ZIA side, you configure your policy to identify traffic that is required to come from your IP address, and forward that traffic post-inspection to the ZPA App Connector.

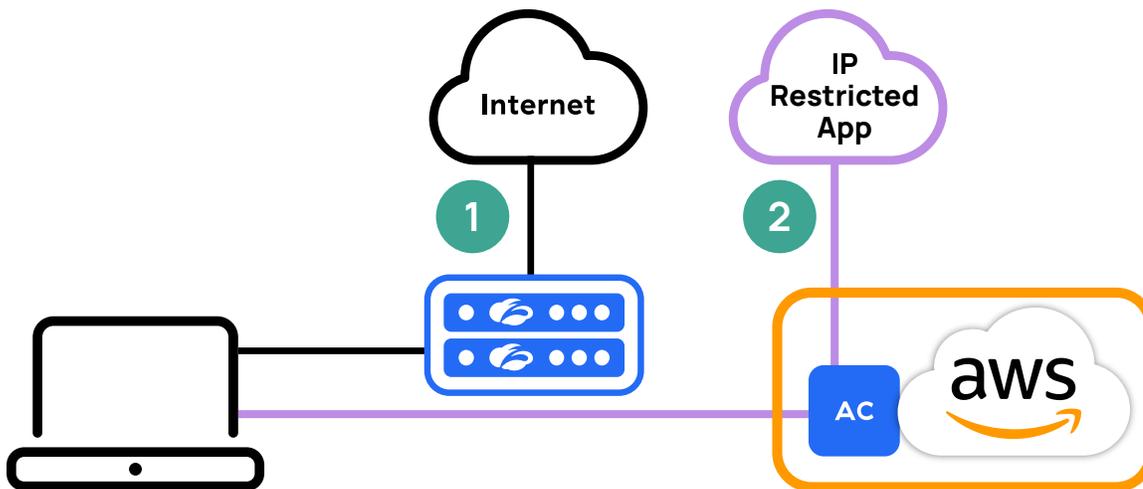


Figure 24. Zscaler Source IP Anchoring via a Zscaler App Connector and a public cloud host

Zscaler recommends only forwarding traffic that requires a set source IP address to be sent to the cloud via policy. You select those applications, URLs, or IP addresses that need to see a static source IP address. Traffic that does not require a set source IP address can continue to be forwarded by the ZIA Service Edge directly as normal.

Forwarding all your traffic to a public cloud can quickly become cost prohibitive in terms of bandwidth and the number of App Connector VMs required. If you need all your traffic sourced from static IP addresses, Zscaler recommends deploying ZIA Virtual Service Edges or ZIA Private Service Edges at your campus locations.



This model requires you to have both ZIA and ZPA subscriptions.

Learn more about [Source IP Anchoring \(https://help.zscaler.com/zia/about-source-ip-anchoring\)](https://help.zscaler.com/zia/about-source-ip-anchoring).

## ZIA Virtual Service Edge and ZIA Private Service Edge

ZIA Private Service Edge and ZIA Virtual Service Edge devices extend the Zscaler cloud into your data center. In this deployment model, a virtual or physical set of servers are deployed in your data center and managed by Zscaler. These are not standalone appliances you manage, but are instead private extensions of the Zscaler public cloud into your organization.

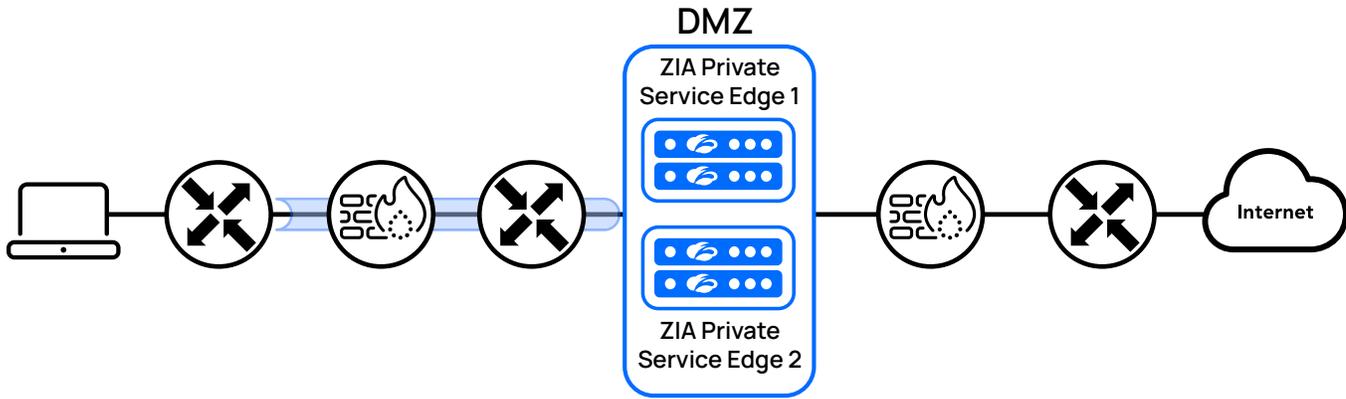


Figure 25. A pair of ZIA Private Service Edge devices deployed in your organization's DMZ

This deployment model consists of two services: a ZIA Load Balancer and a ZIA Service Edge. Both are deployed in pairs, for a total of four devices per location. This allows for active and standby load balancers, and use of all Service Edge instances. These load balancers are internal to the ZIA Private Service Edge. If your organization routinely exceeds 2 Gbps of traffic, physical load balancers are recommended. When deployed, Zscaler fully manages the systems including updating software and signatures.

For most organizations, a ZIA Private Service Edge or ZIA Virtual Service Edge is not required or recommended. Deploying this infrastructure requires additional approvals and subscriptions. The primary reasons for deploying a ZIA Private Service Edge are:

- To have data processed in a location where geopolitical factors make tunneling traffic out of the country impossible. This allows you to enforce policy at your organizational edge and still allow for lawful interception by government agencies.
- Your use of internet applications require that your static IP address must be preserved, instead of being able to use one of Zscaler's public IP addresses. This is a model used by some software vendors to authenticate traffic as being from a legitimate customer, or for applications licensed by location.
- Your campus bandwidth requirements are greater than 1 Gbps for upload or 2 Gbps for download, and you don't want to configure multiple GRE tunnels to multiple data centers.
- To ensure that your users see localized content when using web applications. Some applications use IP addresses to determine a user's location and language. In this case, tunneling your user's traffic to another destination can result in incorrect language selection based on the location of the ZIA Public Service Edge.
- Your physical location encounters high latency when trying to reach the nearest ZIA Public Service Edge.

If these conditions apply to your organization, Zscaler recommends engaging with your local team and technology partners. These experts can help to determine if a ZIA Private Service Edge or ZIA Virtual Service Edge is the right solution. If approved, you need to decide between a physical or virtual deployment model.



The ZIA Private Service Edge and ZIA Virtual Service Edge require additional software subscriptions, and the ZIA Private Service Edge has an additional cost for the device hardware. Redundant pairs are deployed in both instances.

Learn more about the [ZIA Private Service Edge](https://help.zscaler.com/zia/about-service-edge) (<https://help.zscaler.com/zia/about-service-edge>).

Learn more about the [ZIA Virtual Service Edge](https://help.zscaler.com/zia/about-virtual-service-edges) (<https://help.zscaler.com/zia/about-virtual-service-edges>).

## Load Balancing across Multiple WAN Links (Bonded DSL, etc.)

In some locations, network access is not available at the speeds required for your organization to operate effectively. This lack of infrastructure can lead to aggregating multiple WAN links to gain the necessary bandwidth. WAN link aggregation hardware is set up to load balance traffic across the links.

When configuring your load balancing, ensure that users are placed onto the same connection WAN link. You want to ensure that users are going to the same ZIA Service Edge for their traffic.

When routing user traffic over multiple links, there is always a chance that the users will be routed to a different ZIA Service Edge due to the ISP's backend. If this occurs, the user must keep reauthenticating, potentially causing issues for your users. Applications can have issues servicing traffic if users arrive from different Zscaler IP addresses for each transaction to the same application.

It is Zscaler best practice that each user communicates consistently with the same ZIA Service Edge while on the same network.

## No Default Route Networks

In a no default route network, there is no default route from your network gateway to an upstream ISP router. Instead, your gateway routers are configured with static routes to proxy servers, usually in the organization's demilitarized zone (DMZ). Traffic from remote sites is tunneled back to network gateways, typically over MPLS connections. Local users download PAC files from a local resource inside the network, such as a web server or the proxy itself. Remote users are considered uncommon, and are forced to use remote-access VPNs and PAC files to connect back to the corporate network. All network access is provided via static routes from gateway routers. These gateway routers push traffic through stacks of application proxies and appliances that are too expensive and complex to deploy to remote sites.

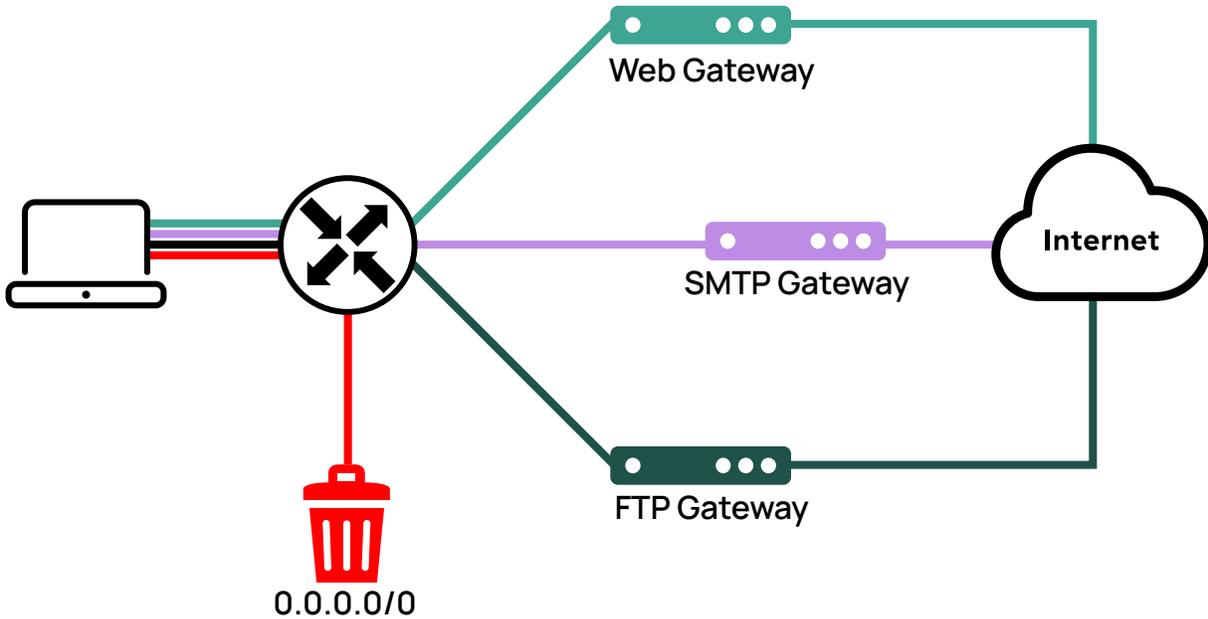


Figure 26. No default route network

No default gateway networks are an older model, and the world has changed significantly since its introduction. In the modern workplace where cloud-based applications, commodity ISP connections, and work-from-anywhere are common, this model encounters serious limitations. The organization's data center is moving to SaaS and IaaS providers in the cloud. MPLS is being replaced by commodity internet connections that provide faster access to cloud-based applications. Users are working remotely and now experience latency and content issues when they are forced back to common gateways.

While Zscaler can support no default route networks, you should be aware the Zscaler's SLA does not apply. Due to the complexity of serving different PAC files inside and outside of the network, Zscaler cannot guarantee access in the same manner as a traditionally designed network.

 No default route network designs void Zscaler's SLA. If you must use a no default route network, contact your Zscaler Account team to discuss your options.

Zscaler strongly recommends the use of a standard gateway approach and tunneling for your organization's network. For remote users, Zscaler recommends the use of Zscaler Client Connector.

Learn more about PAC file options with a [no default route network](https://help.zscaler.com/zia/implementing-zscaler-no-default-route-environments) (<https://help.zscaler.com/zia/implementing-zscaler-no-default-route-environments>).

## Proxy Chaining

Proxy chaining occurs when one proxy server forwards traffic to another proxy server. In this case, your legacy on-premises server is configured to forward traffic to the ZIA Service Edge. This is a transitional network design until you can configure another forwarding method in your network.

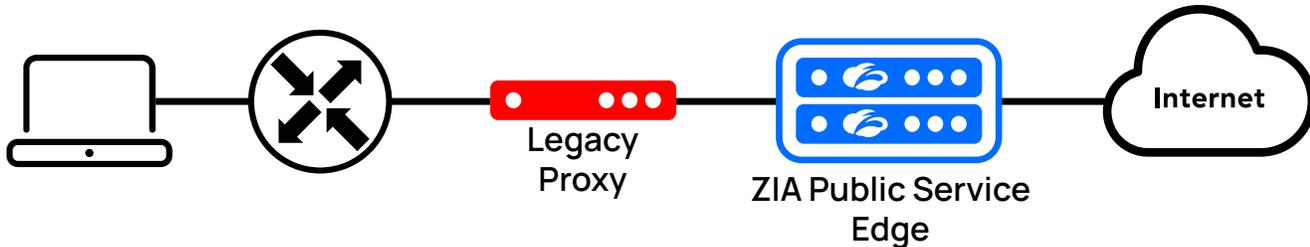


Figure 27. Proxy chaining with your existing proxy forwarding traffic to a ZIA Public Service Edge

Typically, proxy servers only support manual failover in the event of a network outage. So, Zscaler only recommends proxy chaining as a short-term solution until a more robust forwarding solution can be implemented.

Learn more about [configuring proxy chaining](https://help.zscaler.com/zia/configuring-proxy-chaining) (<https://help.zscaler.com/zia/configuring-proxy-chaining>), including from Microsoft ISA or Squid Proxy servers.

## 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.

©2024 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.