Zero Trust User-to-App Segmentation with ZPA

Reference Architecture

Contents

©2022 Zscaler, Inc. All rights reserved.

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**).

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

©2022 Zscaler, Inc. All rights reserved.

Terms and acronyms used in this guide

Acronym	Definition
CNAME	Canonical Name record
DNS	Domain Name Server
FQDN	Fully Qualified Domain Name
IdP	Identity Provider
RA	Reference Architecture
SAML	Security Assertion Markup Language
SCIM	System for Cross-Domain Identity Management
SIEM	Security Information and Event Management
VPN	Virtual Private Network
ZPA	Zscaler Private Access™

Icons used in this guide

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

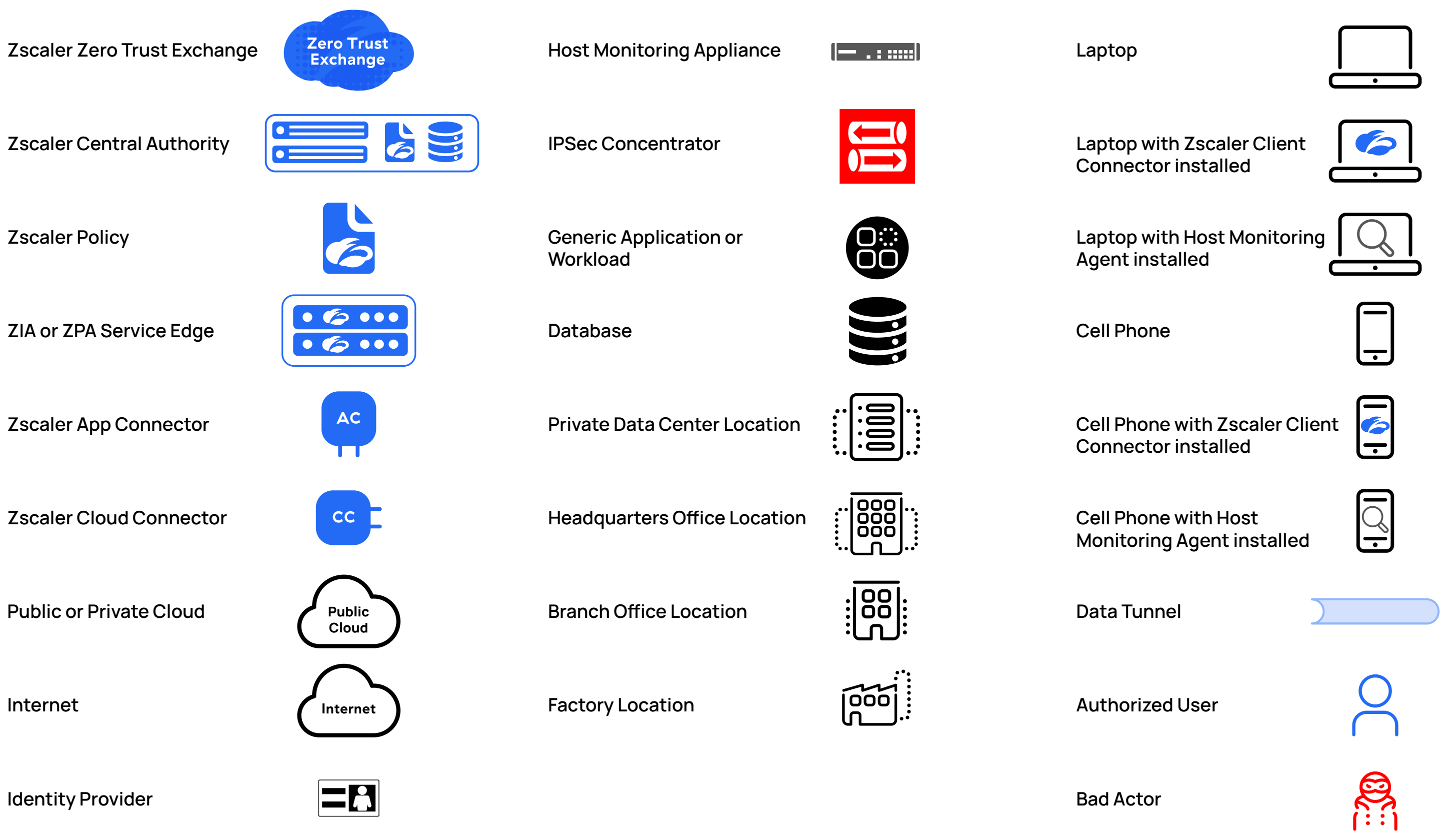
Zscaler Zero Trust Exchange
Zero Trust Exchange
Zscaler Central Authority
Zscaler Policy
ZIA or ZPA Service Edge
Zscaler App Connector
AC
Zscaler Cloud Connector
CC
Public or Private Cloud
Public Cloud
Internet
Internet
Identity Provider
Host Monitoring Appliance
IPSec Concentrator
Generic Application or Workload
Database
Private Data Center Location
Headquarters Office Location
Branch Office Location
Factory Location
Laptop
Laptop with Zscaler Client Connector installed
Laptop with Host Monitoring Agent installed
Cell Phone
Cell Phone with Zscaler Client Connector installed
Cell Phone with Host Monitoring Agent installed
Data Tunnel
Authorized User
Bad Actor

©2022 Zscaler, Inc. All rights reserved.

Introduction

Zero trust has become a popular model for secure user access to applications and resources. Moving away from the traditional VPN style of network access, zero trust is an approach that focuses on granular user-to-app segmentation. Decoupling the user from network-based application access is the first step towards zero trust security, making applications invisible unless the user is authorized to access them.

Zscaler Private Access (ZPA) gives users access to applications without requiring users to share a network context with the applications. Where VPNs assign a user an IP address and place the user on the network, ZPA enables users to connect only to allowed applications, with no access to adjacent applications or systems. The user can be anywhere, and the application can be hosted in any location. Granular context-based policy can control application visibility and access by end users.

In this guide, we review several reference architectures that illustrate how you can leverage the ZPA solution to limit users to specific applications. These models can be designed and deployed to deliver zero trust access that meets your organization's goals for application and workload security.

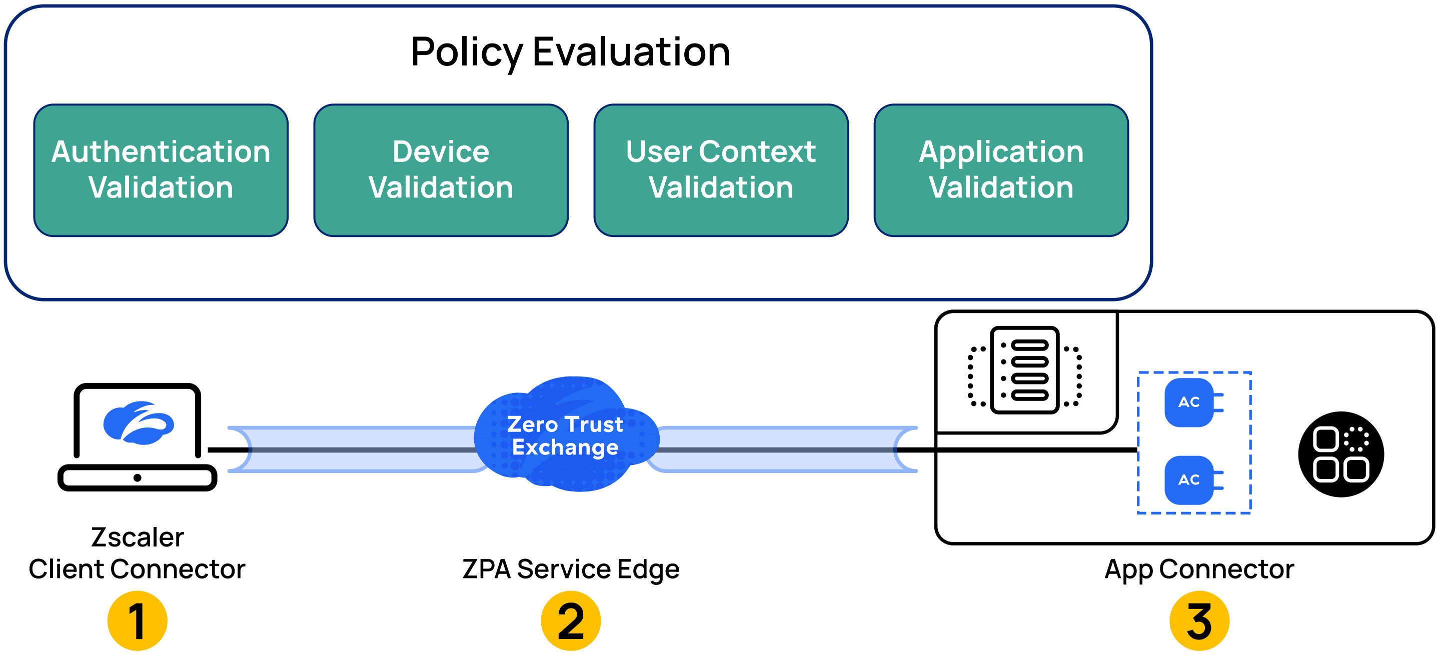

Figure 1. Zero Trust Access

Flexible policy selection is more than a simple identity- and role-based structure in ZPA. The user policies can be defined by examining the entire user context available to the system, including things like the user ID, device, location, and compliance. Based on those values, a set of policies can be enforced. This means you can provide one level of access to a user on an IT-issued laptop at a branch office, and another level of access when that same user is on their personal phone at a coffee shop.

©2022 Zscaler, Inc. All rights reserved.

Key Features and Benefits

- 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.
- 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.
- Unified platform for users, workloads and OT/IoT access: Private apps, services, and OT devices stay in easy, secure reach with the industry's most comprehensive ZTNA platform.

New to ZPA and Zero Trust?

If this is your first time reading about ZPA, we encourage you to watch this **three-minute overview** (**https://youtu.be/1KLbE243dLY**) covering the benefits of ZPA over traditional VPN solutions.

You can find additional information, case studies, demo videos, and a free test drive of ZPA on the **Zscaler website** (**https://www.zscaler.com/products/zscaler-private-access**).

To learn more about zero trust, see our **zero trust microsite** (**https://www.zscaler.com/it-starts-with-zero**).

To learn more about the zero trust architecture, we recommend the National Institute of Standards and Technology paper (NIST 800-207) on **zero trust architecture** (**https://www.nist.gov/publications/zero-trust-architecture**).

©2022 Zscaler, Inc. All rights reserved.

Getting Started with ZPA

ZPA allows you to define a set of policies, authentication, and application definitions that combine to provide user access to applications. The ZPA service ensures that applications are available and that users have the correct policy applied. The following ZPA components and configuration elements compose the foundation for implementing zero trust principles via user segmentation.

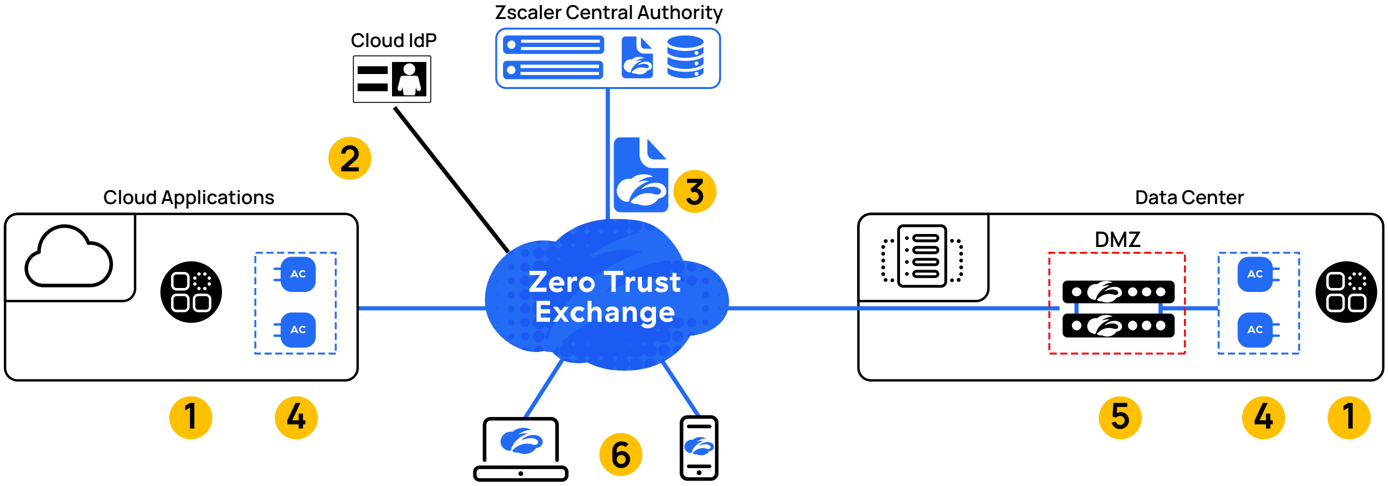

Figure 2. Component overview of the ZPA solution

1. Applications, Application Segments, and Segment Groups – An application is a Fully Qualified Domain Name (FQDN), local domain name, or IP address that is defined on a standard set of ports. Applications must be defined within an application segment. Zscaler recommends using FQDN whenever possible. An application segment is a set of defined applications on shared ports across one or more back-end servers. A segment group is a set of application segments combined for policy purposes. Applications can be grouped into application segments and segment groups based on access type, authorized users, etc.
2. SAML and SCIM Attributes – SAML and SCIM attributes such as group membership, role, etc. are used in access policy rules to provide least-privilege access to applications. These attributes may originate in existing authentication/authorization repositories, such as Active Directory, binding users to relevant groups that reflect onboarding, movement to different departments, changes such as termination, etc.
3. Access Policy – Access policy rules enable context-based access control. To configure an access policy rule, you must first define which applications or segment groups the rule controls, and then define the context required for access. Additional context for access policies may include device posture, access type, network location, and other context provided by the SAML Identity Provider (IdP).
4. App Connectors – App Connectors provide a secure, encrypted, authenticated interface between a customer's servers and applications and the ZPA cloud for delivering user traffic to back-end applications.
5. Private Service Edge – A ZPA Private Service Edge is a single-tenant instance that provides complete broker functionality of a ZPA Public Service Edge in an organization's environment. For an on-premises user connecting via a local Private Service Edge and local App Connector, all control-plane and data-plane traffic stays within the network; the Private Service Edge communicates to the Zscaler cloud for management plane (configuration, logging, etc.) and delivery of user traffic to remote resources.
6. Zscaler Client Connector – Zscaler Client Connector is an agent that resides on your mobile or desktop devices. Supported on popular operating systems such as Windows, macOS, Android, iOS, and CentOS, this agent connects your devices to Zscaler's Zero Trust Exchange. At the ZPA Service Edge, requests are evaluated and approved users are connected to appropriate applications.

©2022 Zscaler, Inc. All rights reserved.

User and Application Discovery for Policy Development

In many cases, you won't have sufficient information to generate fully granular policies for user access. You may not have a list of every user that needs access to applications, or you may not have a full list of applications in use within your organization. There are three primary methods to define the context required to set up policies:

- Asking the application owners about application access by users and groups
- Gathering context from existing logs and data sources
- Leveraging ZPA's own application discovery capabilities

These are not mutually exclusive paths, and it can be beneficial to use all three methods. The reference architecture examples described later in this guide address each of these methods in practice. The goal is to figure out which applications need to be accessed via ZPA and which users and user groups need access to each application.

Zscaler recommends wherever possible that you leverage application discovery. In this mode, ZPA is initially set up to allow a set of users to access a wide range of destinations. Based on the resulting user traffic, ZPA generates a discovered applications list that can be combined with user activity logging to gather granular context. This list is a non-intrusive collection with no scanning or active mechanisms; it is simply a passive observation of user requests and destination applications.

To enable ZPA's application discovery, start by defining application segments containing wildcard domains or IP subnets, and creating an access policy to allow specific users to access those application segments. After applications are discovered, you can collaborate with the application owners to create granular application segment definitions and gradually phase out wildcard application segments. Your configuration policy will initially contain:

1. Wildcard application segments with a wide range of ports
2. Access policies with limited restrictions, allowing users to access many applications

Application discovery does not have to be applied to all users or all applications simultaneously. This level of visibility can be gathered by one use case, one user community, one resource, or one group of resources at a time. Instead of taking a monolithic approach to zero trust, you can take a phased approach: start with applications common across your organization and add more complex use cases over time as discovery provides necessary context. You can apply the lessons you learn in initial phases to refine subsequent phases.

©2022 Zscaler, Inc. All rights reserved.

Transitioning to a Zero Trust ZPA Deployment

One common driver for ZPA deployment is to replace a remote access IPSec VPN. In the VPN model, endpoints are connected to a network, and access security must be handled by network segmentation or application controls. In your initial deployment of ZPA, you can do something similar with wildcard discovery. You can deploy ZPA in a policy configuration that mirrors the open-access VPN model with wildcard app segments and broad access policies. This allows you to make one conceptual change at a time: first change the transport to the applications, then work on granularity of control. This also reduces risk by ensuring that no user applications break due to policy restrictions.

At a high level, the starting point will look something like the following illustration, where users have access to various applications.

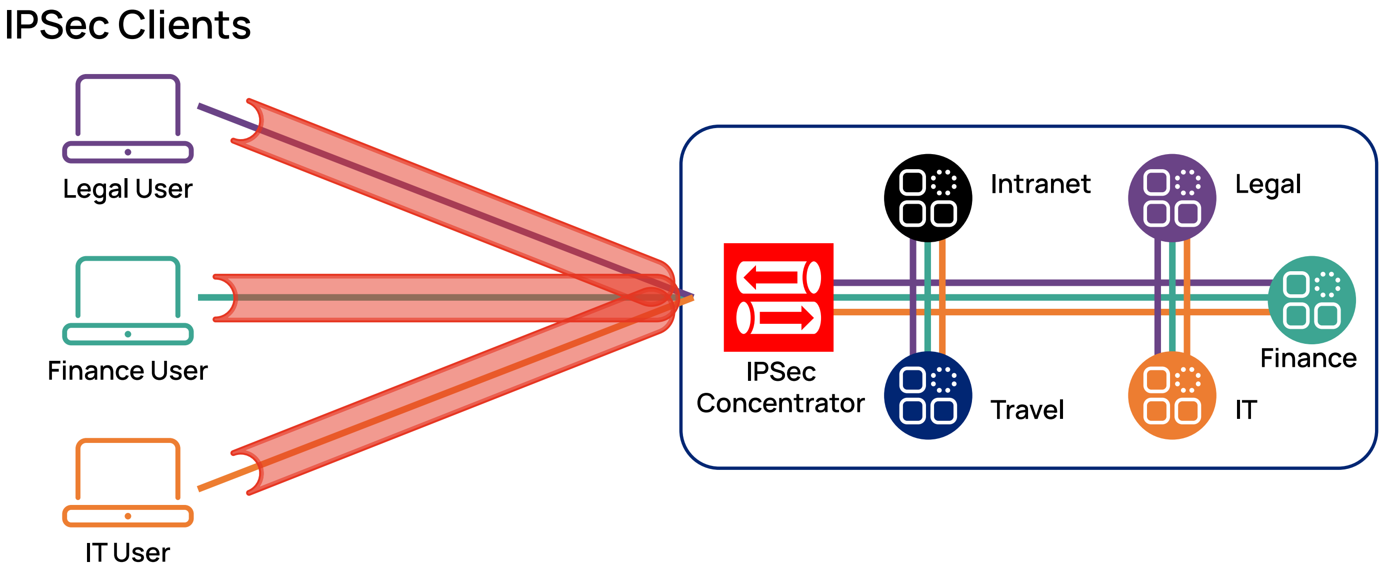

Figure 3. Remote access IPSec VPN puts users directly on the network

When transitioning to ZPA for access, it should be as transparent as possible to the end user. By setting up ZPA to use wildcard discovery, you are replicating the user's existing controls. After migration to Zscaler, initial setup will look something like the following illustration, mirroring your IPSec VPN access policy so users have a similar experience and access to the applications.

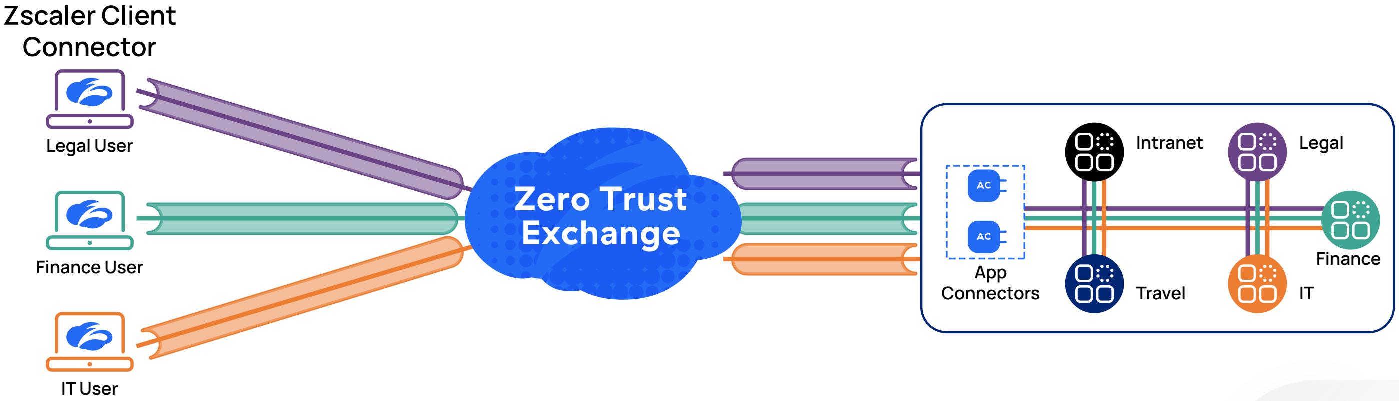

Figure 4. ZPA in wildcard discovery mode initially mirrors the access of an IPSec VPN

©2022 Zscaler, Inc. All rights reserved.

After your initial group is migrated to ZPA, you need to begin building more granular policy. The goal is to refine the app segment definitions, user context, and access policies based on what is learned in discovery, and to provide access to required applications only to users and user groups who should have access to these resources. This can be achieved in multiple ways depending on the requirements. The same process of discovery and refinement allows you to extend your ZPA protection beyond the remote-access use case and also encompass on-premises user access.

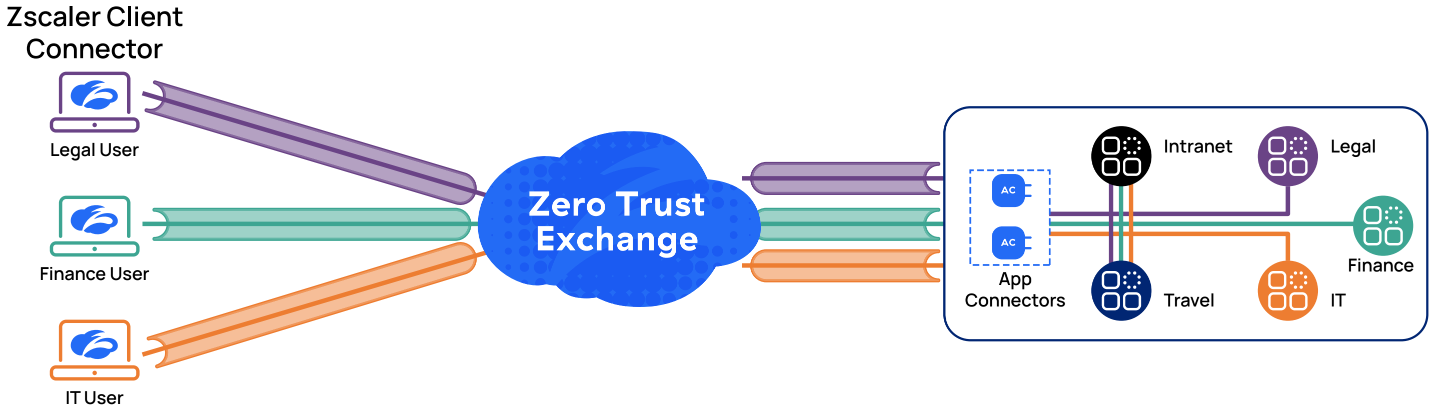

Figure 5. After policy refinement, users can only access the appropriate resources, no matter their location

How your organization progresses toward a zero trust transformation will be defined by the kinds of access you want to grant and the relationship of the user, device, or location to your organization. Some use cases require using fully granular zero trust policies. Examples include third-party access to internal tools, or during an acquisition where the user community is easily identified and the target resources are clearly scoped.

For other use cases, such as employee roles where the traditional approach has been full network access, a different approach may be required. A targeted access policy can allow only those users to get to a broad range of resources (e.g., *.internal.safemarch.com), enabling initial connectivity as well as dynamic application discovery. The resulting user activity logs offer visibility to inform development of more granular policies, so you can evolve towards more granular zero trust controls for those users as context becomes available.

It's important to remember that "success" at zero trust does not necessarily mean eliminating all wildcard access. There may be use cases in your organization where it's appropriate to provide a certain user community access to a full range of applications within a certain domain or subdomain. Granularity then might mean ensuring that only authorized users on appropriate devices, possibly with additional context controls, are permitted access via that wildcard policy.

©2022 Zscaler, Inc. All rights reserved.

Reference Architectures (RA) Overview

The following reference architectures are based on real-world ZPA deployments by a variety of enterprises and municipalities, representing a range of approaches to achieving zero trust security with ZPA. Broadly speaking, the transition to zero trust can be accomplished with varying degrees of speed and control along the way. The first five architectures represent three main categories:

- Fast/Open: RA1 approach is to wildcard everything and carve out application segments over time to gradually restrict access as needed. You can use this model when replacing a legacy user-access VPN, or if you need to rapidly respond to users moving from organization sites to work from home.
- Slow/Controlled: RA2 and RA3 both avoid wildcard use and require the customer to apply a range of mechanisms to identify and define the user groups and network details needed to access various applications. RA2 describes using data feeds from other elements of the environment, such as end-user devices, while RA3 describes using data directly provided by application owners. Both approaches offer a very planned and controlled approach. However, they can require a longer timeline to work through transformation and can also miss applications that are not well documented or understood.
- Variable speed/Risk-based control: RA4 and RA5 both describe the use of well-defined application segments for some applications while allowing open access (via wildcards or open on-premises access) for lower risk situations. RA4 delineates based on criticality of the data, resource, or application. RA5 delineates based on user groupings.

Reference Architecture 1: Wildcard discovery

This approach starts with wildcard applications and monitors the interactions between users and applications. You then rely on these discovered applications and patterns to configure granular application segments. The steps are highlighted in the following flowchart.

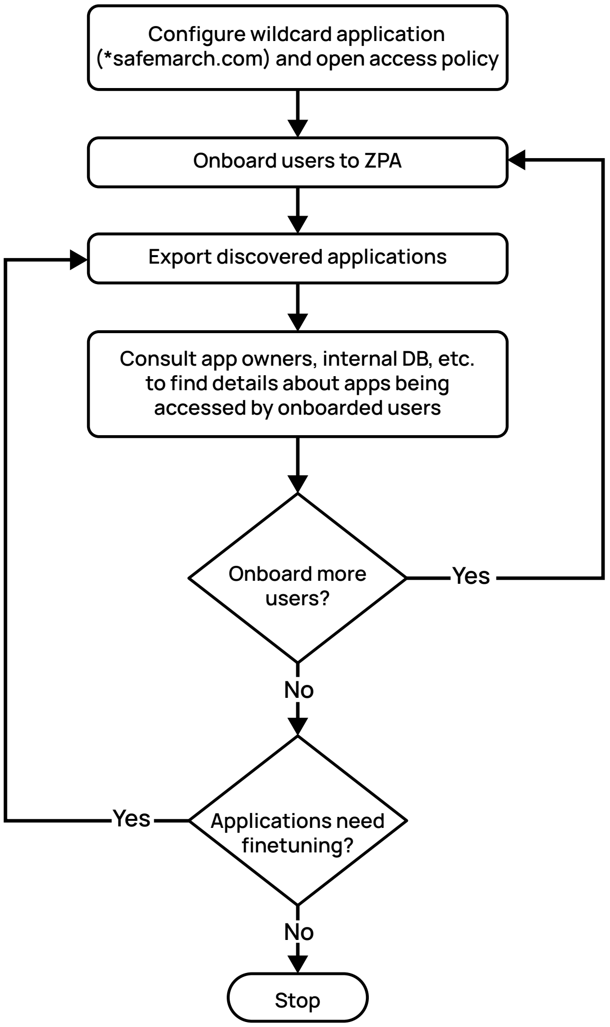
Configure wildcard application (*safemarch.com) and open access policy
Onboard users to ZPA
Export discovered applications
Consult app owners, internal DB, etc. to find details about apps being accessed by onboarded users
Onboard more users?
Yes
No
Applications need finetuning?
Yes
No
Stop

Figure 6. Policy refinement from wildcard discovery to granular policy

You need to build a well-defined process for evaluating and creating policy around discovered applications. It's very likely that your initial discovery will identify many applications. Some will be easily identified as business applications and grouped appropriately. Others will need to be more closely evaluated. The following steps outline the process to export discovered applications from the ZPA Admin Portal:

1. Determine what each discovered application represents. This may involve engaging with application owners and/or consulting internal databases to map FQDNs or IPs to relevant applications.
2. Determine whether the discovered application needs a new application segment or needs to be added to an existing application segment. Also determine if this requires creating a new access policy or updating an existing access policy.
3. Proceed to make changes in the ZPA Admin Portal. APIs can be leveraged, particularly for application segment configuration or updates, as there can be many FQDNs, IPs, or ports that need to be updated. You can learn more about ZPA APIs in the **Zscaler Help Portal** (**https://help.zscaler.com/zpa/zpa-api**).

The iterative process of configuring granular applications can continue even after all users are onboarded, as more applications are discovered based on user access patterns. For example, a finance application may be discovered only during financial yearly closing, as it is used only at the end of the year. Alternately, certain application segments may remain as wildcards, such as application domains (e.g., *.dev.safemarch.com) used in a dev environment where developers frequently spin up new servers and decommission servers after testing is complete.

Reference Architecture 2: Out-of-band discovery

This model takes the opposite approach and does not start your ZPA deployment with any wildcard applications. In this approach, you set up monitoring prior to onboarding users to ZPA. This model assumes that application owners do not have granular details of the applications, such as FQDN, IP, protocol, and ports. Gathering this information is divided into two steps:

1. Perform user behavior analysis by pulling user access data from systems in the environment, such as firewalls and user endpoints. You should aim to gather as much information as possible. This may include username, endpoint IP address, FQDN, IP, protocol, ports accessed, user vs. machine generated requests, etc.
2. Discovery performed in step 1 results in understanding different applications that are accessed by users. These applications can now be defined and added to app segment definitions.

User group mapping is managed at the SAML IdP level. If a user moves from one department to another, gets promoted, is terminated, etc., these changes are handled at the IdP level by adding or removing users from groups.

Step 1 requires capturing data. In this example, we consider data from user endpoints running a host monitoring tool such as CrowdStrike. Traffic pattern data from CrowdStrike is pulled into a database. The database should be connected to the DNS server to resolve IP addresses to hostnames and gather context on the applications, and connected to the IdP to understand which users and groups are accessing these applications.

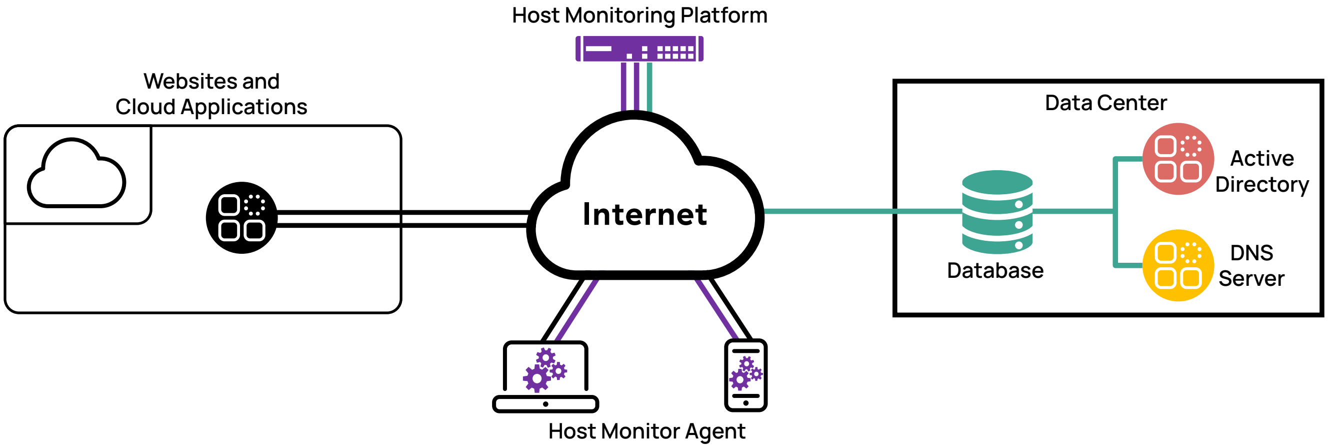

Figure 7. Monitoring user application access via host monitoring agent

The following table is an example of a traffic pattern captured in the database.

IP	Port	Protocol	User
10.1.1.1	443	TCP	user1@safemarch.com
10.1.1.1	443	TCP	user2@safemarch.com
192.168.10.15	22	TCP	user1@safemarch.com
192.168.10.100	443	TCP	user3@safemarch.com

©2022 Zscaler, Inc. All rights reserved.

Using the connection to the DNS server, the FQDN for the 10.1.1.1 IP can be determined. In our example, it is bitbucket.safemarch.com. Connection to the IdP helps you understand the roles for users that are accessing bitbucket. In our previous example, both user1@safemarch.com and user2@safemarch.com are from the Engineering group. You can determine how critical a resource is to a user or group by monitoring the number of times the applications are being accessed. Using this method, the traffic database can be enhanced with useful information about the applications and users.

IP	FQDN	Port	Protocol	User	User Group
10.1.1.1	bitbucket.safemarch.com	443	TCP	user1@safemarch.com	Engineering
10.1.1.1	bitbucket.safemarch.com	443	TCP	user2@safemarch.com	Engineering
192.168.10.15	jump-us.safemarch.com	22	TCP	user1@safemarch.com	Engineering
192.168.10.100	tableau.safemarch.com	443	TCP	user3@safemarch.com	Finance

Based on this data, it can be determined that users from the Engineering group need access to bitbucket.safemarch.com on TCP port 443. Now this can be translated into an app segment and access policy in the ZPA Admin Portal. After the configuration is complete, these users can be migrated to the ZPA platform.

Application Segment Definition				
App Segment Name	**Segment Group Name**	**FQDN**	**Port**	**Protocol**
Eng_App_Seg	Eng_Seg_Grp	bitbucket.safemarch.com	443	TCP
Eng_App_Seg	Eng_Seg_Grp	bitbucket.safemarch.com	443	TCP
Eng_App_Seg	Eng_Seg_Grp	jump-us.safemarch.com	22	TCP
Fin_App_Seg	Fin_Seg_Grp	tableau.safemarch.com	443	TCP

Access Policy Definition			
Policy Name	**Segment Group/App Segment**	**SAML Attribute**	**Action**
Policy to Allow ENG Applications	Eng_Seg_Grp	memberOf=Engineering	Allow
Policy to Allow FIN Applications	Fin_Seg_Grp	memberOf=Finance	Allow

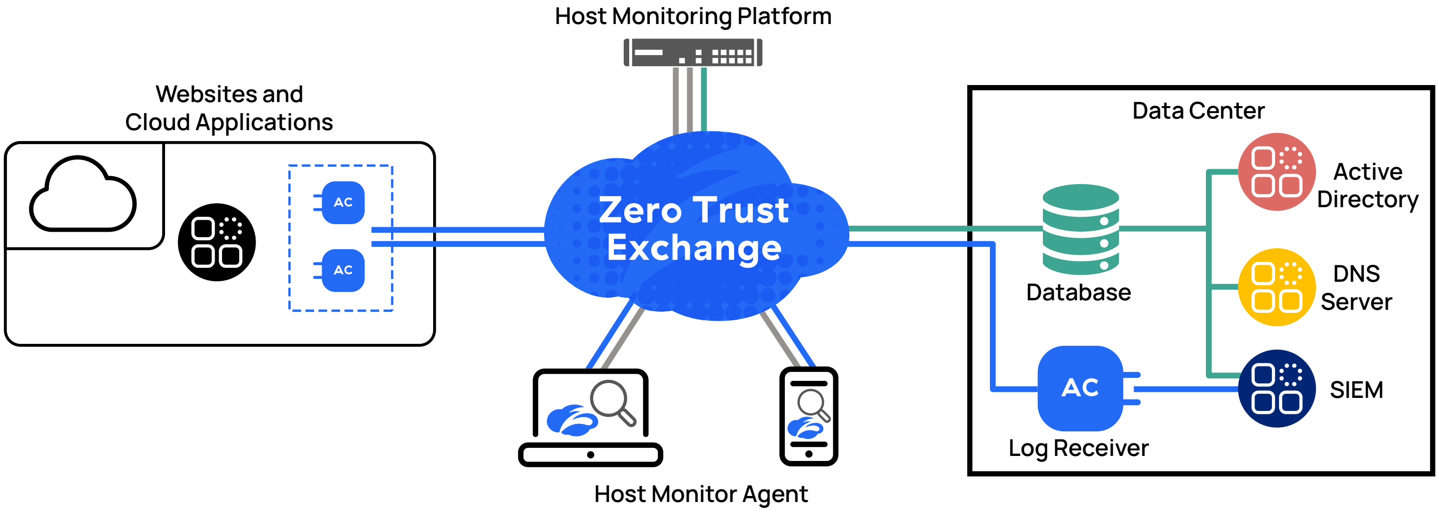

Figure 8. Fully integrated discovery and logging

After starting migration of users to ZPA, you can connect the database to your Security Information and Event Management (SIEM) server in addition to your IdP and DNS servers. SIEM logs can be referenced by the database to determine if an application is already migrated to ZPA or not.

Over time, the migration process follows a traffic light approach.

- Green: Users can be migrated if there are app segments and access policies defined in the ZPA Admin Portal for all resources being accessed by the user group.
- Amber: Some applications or policy are not yet defined for this user group.
- Red: No applications or policy are defined for this user group.

To build the database, it is helpful to engage a database analyst and an API expert. You can fetch data from systems such as CrowdStrike and correlate this information by further connecting to your IdP, DNS, and SIEM servers. You can then configure the resulting app segments and access policies in the ZPA Admin Portal. One customer who followed this approach had a 6-member team that was able to perform initial discovery and successfully start onboarding the first group of users to ZPA in approximately 7 days.

Reference Architecture 3: App owner discovery

As in our previous model, this approach does not start with any wildcard application discovery. Application owners are expected to supply information about the applications, and then users can be onboarded. This is divided into two steps:

1. Application owners provide details about the applications, users, and user groups who need access to applications.
2. Using the information supplied in step 1, your ZPA administrator can complete ZPA configuration.

This approach requires close collaboration with application owners within the organization before deploying the ZPA solution. Application owners help supply detailed information about applications and who should have access to their applications. This provides input for both:

1. Application segment configuration
2. Access policy definition

©2022 Zscaler, Inc. All rights reserved.

The following table shows application owners supplying definitions of applications and the users and user groups who need access to them.

Application Name	Application FQDN/IP	Protocol	Port	User/Group
Automation Tool	automation1.eu-west.amazonaws.com automation2.eu-west.amazonaws.com dev-test.cd2.safemarch.com	TCP+UDP	1433-1434	Group_ENG
AutoCad	autocad1.cd1.safemarch.com autocad2.cd2.safemarch.com 10.163.1.10 10.161.4.110	TCP	27000-27001 2080-2080 28000-28000	Group_ARCH
RDP Server	rdp.ad.safemarch.com	TCP	3389-3389	Group_ADMIN

You can take the above information and perform relevant configuration in the ZPA Admin Portal. This includes configuring the application segment and access policy so that relevant users and groups have access to their applications. This can be done manually via the ZPA Admin Portal or via API.

Reference Architecture 4: Application-focused granularity

This architecture focuses on the applications in use in your organization, and optionally uses DNS subdomains to facilitate access granularity. Applications are placed into one of three categories:

1. High-value assets that are defined as granular applications. There are no wildcard app segments for high-value assets, and limited users have access to these applications. Common examples include infrastructure, security, or finance applications.
2. General applications that may be defined in wildcard app segments, and access may be granted to a broader set of users. These are applications that most users in an organization should be able to access, such as internal web resources, timesheet apps, etc.
3. Third-party facing applications that are sets of resources or tools that an outside vendor, contractor, or consultant may access. This may include an HVAC system managed by a vendor, or a registration system where a contract manufacturer enters serial numbers.

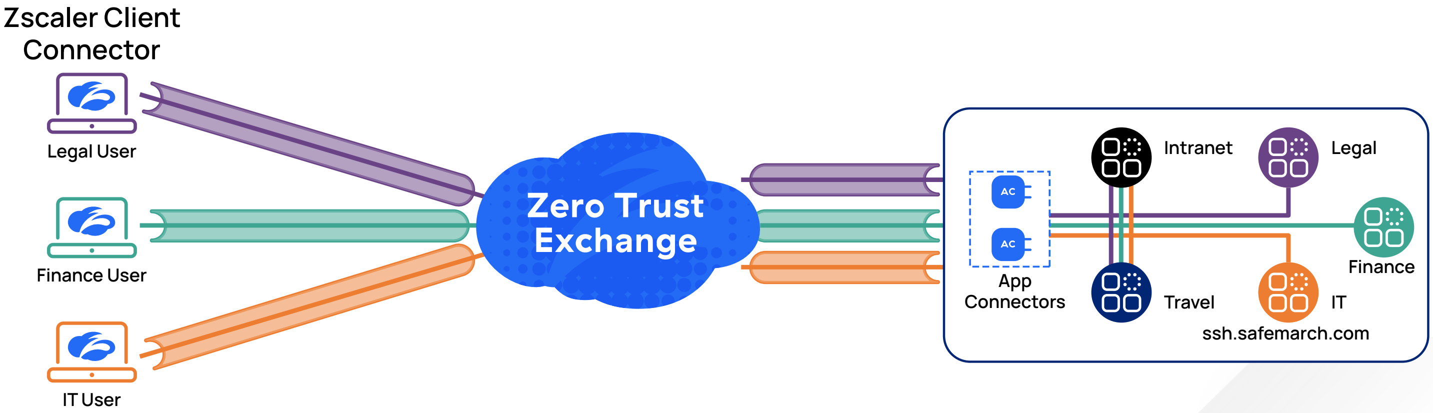

Figure 9. The host resides in the same common domain name space as other applications

Infrastructure applications and other highly restricted applications can be placed in a new namespace created to better handle the setup. DNS CNAMEs can be created for these applications during the transition. This creates an alias for each entry in the new namespace without requiring the original host to change its name. For example:

<CNAME> ssh.admin.safemarch.com

<A record> ssh.safemarch.com

Now when a user needs to access the ssh.safemarch.com server, the user can go to ssh.admin.safemarch.com, which will be picked by ZPA. After the request reaches the relevant app connector, DNS resolution ensures traffic is directed to the correct server.

There is one minor caveat to this approach. Users accessing the infrastructure resources need to be informed that the new FQDN must be used for accessing these resources.

Old FQDN = ssh.safemarch.com

New FQDN = ssh.admin.safemarch.com

If the old FQDN remains reachable, some users may fail to switch to the new FQDN. This can be resolved by establishing and communicating a transition timeline for when the old FQDN will be removed, requiring use of the new FQDN and app access via ZPA.

Now you can group infrastructure applications into app segments, for example:

FQDN = *.admin.safemarch.com

Ports = TCP: 22-22, 3389-3389

Application Name	Application FQDN/IP	Protocol	Port	User/Group
Server	*.admin.safemarch.com	TCP	22-22 3389-3389	Group_IT

Access policy can be defined to allow access to these applications for only the IT user group. This ensures that only specific users or groups have access to the high-value asset applications. General and third-party facing applications may be a combination of wildcard and FQDN-specified applications, depending on requirements.

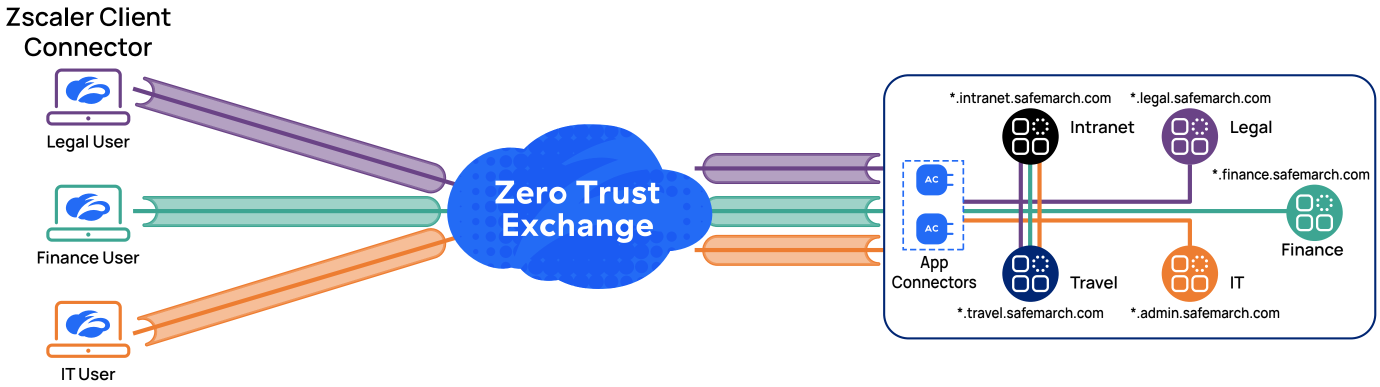

Figure 10. Applications are behind granular access policies, and only approved users can access each application

This architecture heavily leverages DNS subdomains to simplify the setup. If there are servers being deployed in AWS, then a subdomain should be dedicated for that environment. For example, the organization may designate edp.aws.safemarch.com, so the application segment can be defined as *.edp.aws.safemarch.com with relevant TCP/UDP ports. An access policy can be created for the users who need access to this environment.

Reference Architecture 5: User-focused granularity

This approach is most appropriate if your organization currently has a flat network and wants to build boundaries between applications, but not via network segmentation. Instead, you implement controls at the server level. You first identify applications, then leverage IP firewalls on hosts to allow only traffic from authorized sources.

Initially, your traffic first comes from jump hosts. These are hardened systems that your users log into for access to your secure applications. This is an intermediate step to using ZPA App Connectors as your rollout progresses. This enables a 3-phase approach to dividing applications based on the user communities who need to access them.

1. Prepare by transitioning from direct user access to using jump hosts.
2. Onboard your early adopters to ZPA.
3. Complete transition to ZPA and remove jump host access.

IT professional groups and early adopters in other groups are engaged to identify application usage by each group. Zscaler Client Connector is installed and configured for access both on- and off-premises. Applications used by third-party external users, and by some subsets of internal users, are defined as granular applications with limited user access to these applications. Applications used by multiple internal user groups—such as file servers, Active Directory infrastructure services, etc.—can be defined in enterprise-wide app segments, so that access can be granted to broader sets of users.

User context is captured in an IdP, preferably cloud-based. As application groupings are identified, corresponding user groupings are defined. Device compliance requirements—such as domain membership, disk encryption, etc.—are captured for protection of sensitive applications and information.

Phase 1: Preparation

In the first phase, IP firewalls on the server-side hosts are used to limit access to RDP hosts, or virtual machines used as jump hosts for application access. Instead of open access, users connect through these jump hosts to access resources.

1. Ensure user identity is linked to employee roles.
2. Install and configure application connectors.
3. Configure ZPA to allow user access to internal jump hosts based on groupings and optional posture checks.
4. Configure IP firewalls on destination servers to only accept access from jump hosts.

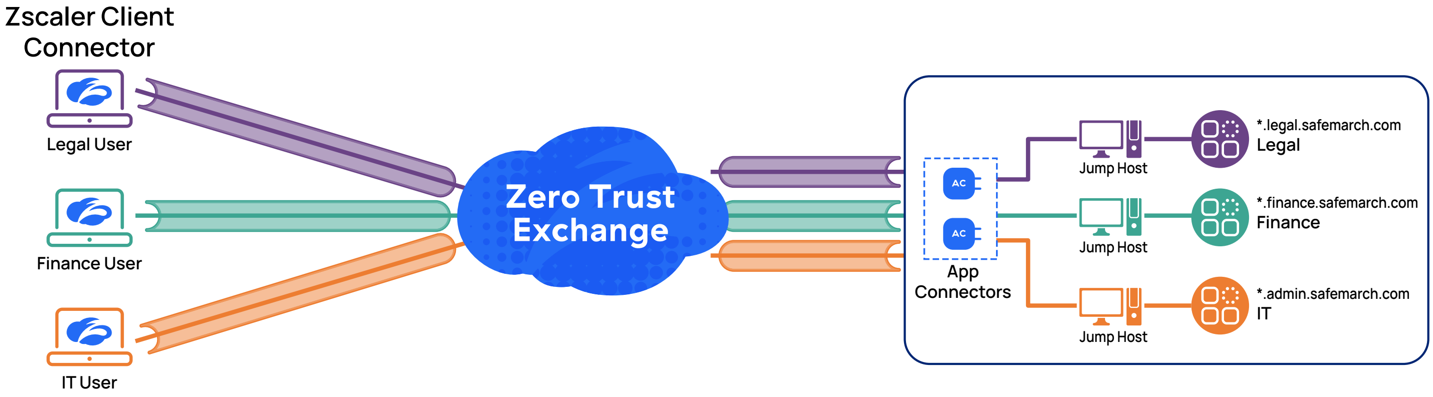

Figure 11. In Phase 1, jump hosts are used to allow access to specific applications

Phase 2: Onboard early adopters

In the second phase, server-side IP firewall rules are updated to also allow traffic from App Connectors. This allows authorized users to connect directly to the required applications without needing to go through the jump hosts, while jump host access is maintained as a fallback. More applications are identified and grouped, and further posture checking may be added to reflect device compliance requirements.

1. Update IP firewall rules to allow access from both jump hosts and ZPA App Connectors.
2. Define known application segments and groups for user access to high-risk applications.
3. Onboard early adopters from each employee role using skilled "champions" who in turn can help others.
4. Identify and add applications that are not accessible while off-premises.
5. Apply posture checks (https://help.zscaler.com/client-connector/configuring-device-posture-profiles-zpa) as needed for sensitive applications.

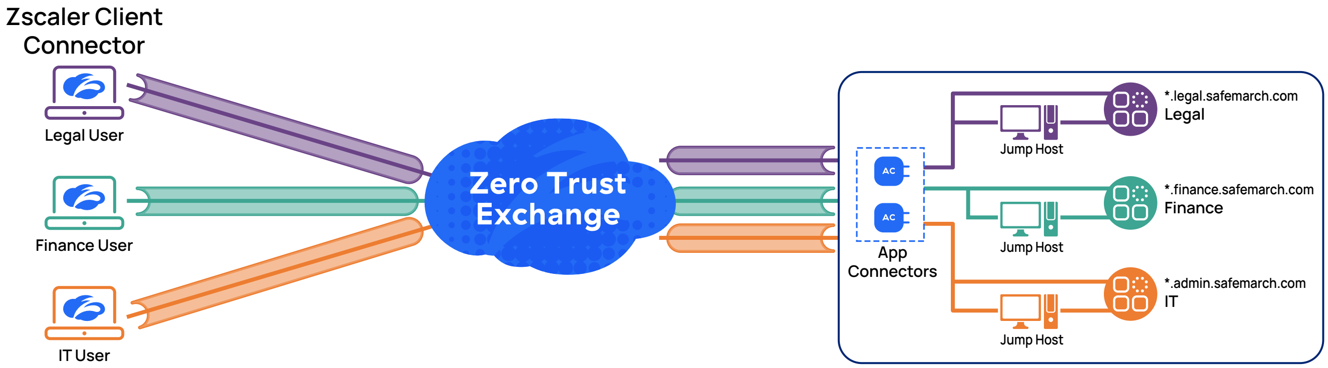

Figure 12. Users shift primarily to ZPA for access, with jump hosts acting as a fallback access mechanism

Phase 3: Complete transition and remove jump host access

In this final phase, you remove the jump hosts from the network. Access policy should now be well defined, and all users should be accessing their applications via ZPA.

1. Validate that all necessary applications are accessible from on- or off-premises, only by authorized users, with no jump host required.
2. Ensure all users are running Zscaler Client Connector on their devices to gain access.
3. Identify and validate any residual traffic between applications that is not flowing through an App Connector.
4. Remove all unnecessary server-side IP firewall access rules, including jump host access.

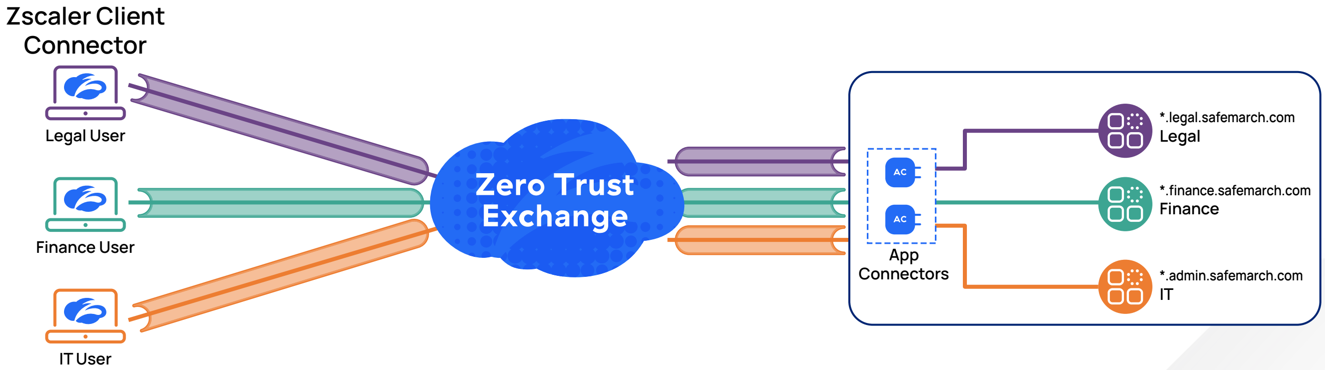

Figure 13. All access via ZPA, jump hosts removed from the network

©2022 Zscaler, Inc. All rights reserved.

Considerations for Multi-Entity Organizations

For a large multi-entity organization, such as a municipality with multiple agencies or an enterprise composed of many subsidiaries, environmental elements can be leveraged to ease deployment. Domain membership can be a useful way to differentiate between users in disparate agencies or departments. If each sub-entity has its own domain, access policies can incorporate a posture check to ensure that the device is in the appropriate organization for the resource the user is requesting to access. This also prevents access from personal devices. This can be particularly useful for jump host or remote desktop access scenarios.

Similarly, strong policies and structure around DNS zones and entries can help ensure that traffic is routed to appropriate App Connectors. Because user application access is provisioned ideally based on hostname and not on IP address, dynamic path selection can ensure optimal performance as well as appropriate delivery. Zscaler best practice is to configure FQDNs, rather than IP addresses, in app segments whenever possible. Two underlying controls can help to enable these benefits:

1. Configure App Connectors to use the most appropriate DNS servers to ensure that they can resolve DNS for applications they are eligible to serve.
2. Configure App Segments to be mapped only to App Connectors that are desirable for delivering that application traffic.

In organizations with multiple independent DNS environments, where cross-environment application access is required, you may want to have all App Connectors resolve to a centralized DNS server with trust relationships to each sub-entity's DNS servers. For other use cases like an M&A use case, where acquirer and acquired company are running on overlapping IP address space, it may be critical to ensure that App Connectors resolve DNS only for the application domains in their specific environments. Taking the time to evaluate your use cases and think through the finer points of name resolution across your entire landscape can streamline the implementation of your zero trust controls.

Conclusion

In this guide, we described a few approaches that can be used to facilitate your transformation from legacy, network-based application access to a user-based approach using application segments and access policies applied through ZPA. We discussed various techniques your organization may use to identify, define, and characterize the applications and application segments that support your transformation journey. We recognize that transformation can be daunting; ZPA eases the transformative process by enabling a flexible, phased approach that adapts to any organization's timeline and risk appetite.

©2022 Zscaler, Inc. All rights reserved.

Additional Resources

Description	URL
Customer Clinic: Going Beyond App Discovery	https://community.zscaler.com/t/going-beyond-app-discovery-what-to-know-about-zpa-access-policies/13299
About ZPA Private Service Edges	https://www.zscaler.com/resources/data-sheets/zpa-private-service-edge.pdf
Master article for ZPA APIs	https://help.zscaler.com/zpa/zpa-api
API for Application Segment	https://help.zscaler.com/zpa/application-segment-use-cases
API for Segment Group	https://help.zscaler.com/zpa/application-segment-group-use-cases
API for Access Policy	https://help.zscaler.com/zpa/access-policy-use-cases
Zero Trust for Ransomware Recovery	https://www.vanillaplus.com/2021/10/15/64465-zero-trust-helps-to-regain-control-after-a-ransomware-attack/
Customer case study	https://www.zscaler.com/resources/case-studies/man-energy-solutions.pdf

©2022 Zscaler, Inc. All rights reserved.