Zero Trust Security for Azure Workloads with Zscaler Cloud Connector

Reference Architecture

Contents

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

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.

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.

Finding out more

You can find our guides on the Zscaler web site at **https://www.zscaler.com/resources/reference-architectures**.

You can join our user and partner community and get answers to your questions at **https://community.zscaler.com**.

Terms and acronyms used in this guide

Acronym	Definition
ZIA	Zscaler Internet Access
ZPA	Zscaler Private Access
ZTE	Zero Trust Exchange
ACL	Access Control Lists
AWS	Amazon Web Services
AZ	Availability Zone
CA	Certificate Authority
DLP	Data Loss Prevention
DTLS	Datagram Transport Layer Security
IaaS	Infrastructure as a Service
IPS	Intrusion Prevention System
LSS	Log Streaming Service
MITM	Man-in-the-Middle
NIST	National Institute of Standards and Technology
NSS	Nanolog Streaming Service
PaaS	Platform as a Service
SaaS	Software as a Service
SIEM	Security Information and Event Management
SSL	Secure Sockets Layer
TLS	Transport Layer Security
VM	Virtual Machine
VNets	Virtual Networks
VPN	Virtual Private Network

Icons used in this guide

Zscaler Zero Trust Exchange 

ZIA or ZPA Service Edge 

Zscaler App Connector 

Zscaler Cloud Connector 

Azure Load Balancer 

Azure Application Gateway 

Azure Virtual Machine 

Azure Application or Workload 

AWS Application or Workload 

Generic Application or Workload 

Private Data Center Location 

Headquarters Office Location 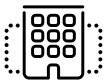

Branch Office Location 

Factory Location 

Authorized Use 

Bad Actor 

Data Tunnel 

Introduction

The shift to cloud services has rebuilt the enterprise data center off-premises and outside of traditional security boundaries. Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) enable organizations to quickly build out and scale their platforms and services. Securing services across multiple clouds, vendors, and support features requires a different approach than that of the traditional data center.

Securing this communication through the layering of legacy Access Control Lists (ACL), on-premises firewalls, and service-chaining has always been both complicated to build and difficult to maintain. Private applications were accessed via Virtual Private Networks (VPNs) that extended the network to locations in an any-to-any access model. This large, flat network gave users a single location to connect to for access to private applications.

Leveraging the cloud breaks these models. You now have multiple vendors across different clouds, products, and services. Your policy must be interpreted at each cloud and application to determine how best to implement it with the tools available. This risk goes up given a mistake, potentially exposing your organization to a host of network-born attack vectors.

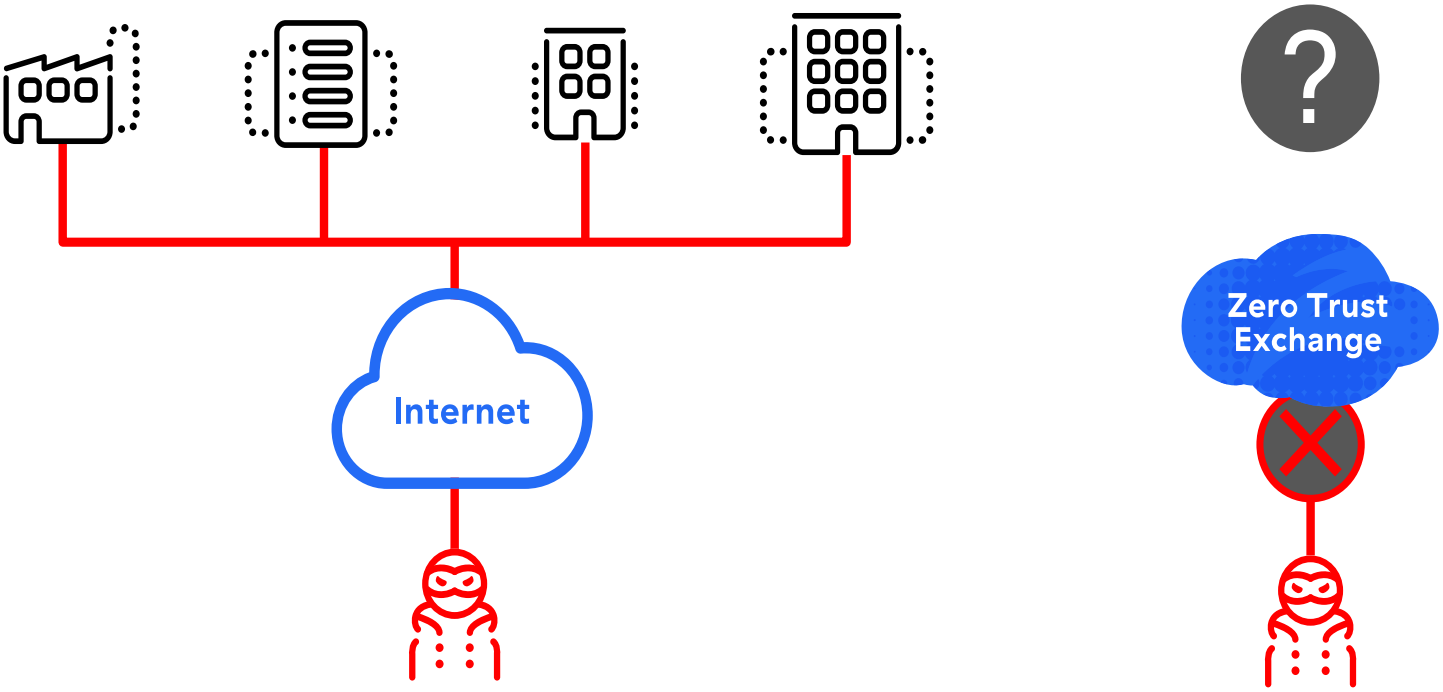

Figure 1. Zero trust moves from exposed network security to user-centric policy enforcement

Ideally, an organization's security policy should be at the foundation of its network design. Connectivity to and from devices happens as a product of the security policy and not the other way around. This is the heart of the Zscaler Zero Trust Exchange (ZTE) model.

Users must be authorized before they can connect to that service. Even knowing the application's hostname and the services it provides won't give the attacker any information, as that service won't resolve until the user authenticates. Your applications are effectively hidden from the internet and each other until you define policy to allow access.

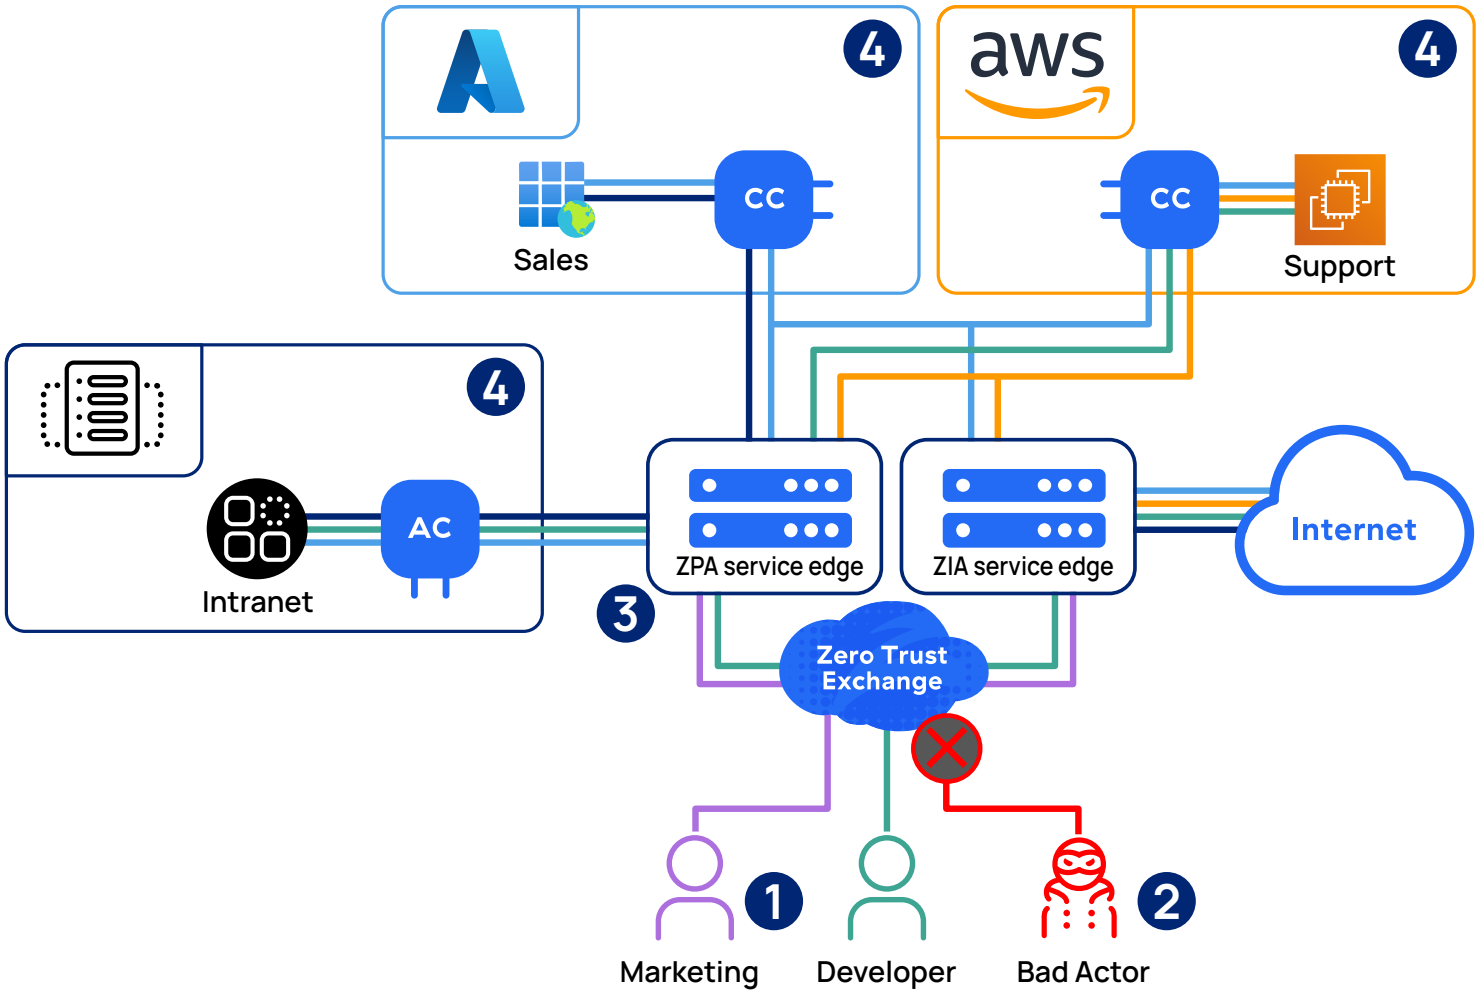

Figure 2. Zero trust principles applied to users and workloads

1. **Authentication** – All users must first authenticate to Zscaler. Based on multiple criteria such as user group membership, device posture, and location, the user is assigned a set of policies. These include the ability to see internal applications.
2. **ZIA Service Edge** – When traffic from users or workloads needs to be routed to the internet, a ZIA Service Edge inspects the traffic. If your policy allows the traffic out, the return traffic is also scanned for malicious content on its way back to the user.
3. **ZPA Service Edge** – Traffic from users or workloads bound for other internal applications is handled by ZPA. Based on the user's authentication and assigned policy, only approved resources are resolved. All other resources are hidden from unauthorized users as if the services do not exist.
4. **Cloud Connector** – Deployed in front of your internal applications, Cloud Connector creates a set of outbound tunnels to ZIA and ZPA. They decide where the tunnel connects based on policy.

In the previous image, your users in marketing have workloads running in Microsoft Azure, and your developers have workloads in Amazon Web Services (AWS). All users have access to internal applications in the data center, as well as general internet access. In this case, each user and workload are limited to which applications they can resolve and access, based on the policy applied to them.

Your marketing user (purple) can access their workloads in Azure, the data center, and the internet based on policy. Your developer (green) can access their workloads in AWS, the data center, and the internet. Finally, your workloads in Azure, AWS, or your data center can all reach one another and the internet via the ZTE, without the need to set up additional VPN links. All these connections are subject to the policy you set.

Cloud Connector ensures that cloud workloads adhere to organizational security policy when accessing both public and private endpoints. This is achieved by intelligently forwarding traffic to the Zscaler Internet Access (ZIA) and Zscaler Private Access (ZPA) platforms. Cloud Connector also enables multi-cloud connectivity and enforces a security policy for cloud-to-cloud traffic.

Key Features and Benefits

- **Security** – Secures all inbound and outbound traffic to the internet. The security capabilities available through the ZIA platform for server internet access are Transport Layer Security (TLS)/Secure Sockets Layer (SSL), Intrusion Prevention System (IPS), Firewall, Data Loss Prevention (DLP), and more.
  To learn more about ZIA, visit **https://www.zscaler.com/products/zscaler-internet-access**.
- **Connectivity** – Provides seamless connectivity from private or public cloud applications to the internet.
- **Performance** – Ensures better end-user experience and application performance by direct peering relationships with SaaS providers (e.g., Microsoft Office 365, Amazon Web Services, and Microsoft Azure).
- **Reduces Cost** – Consolidates multiple products (e.g., Squid proxies, firewalls, third-party NAT appliances, URL filtering, etc.) into a single solution. Additionally, the same policy applied to user traffic can be applied across the cloud infrastructure.
- **Highly Scalable** – Ease of implementation across 1K service accounts in public clouds and single solution scales to connect 10K+ server environments in public clouds (e.g., AWS, Azure, etc.).
- **Ease of Deployment** – Fully orchestrated deployment for Azure using Terraform scripts or ARM templates.
- **Real-Time Visibility** – Dashboards and Insights provide unparalleled visibility into your users and applications, and the health of your organization's applications and servers.

New to Zscaler Cloud Connector?

If this is your first time reading about Zscaler Cloud Connector, we encourage you to watch the following video to see a demo and hear real examples of how Cloud Connector solves today's security challenges at **https://community.zscaler.com/t/innovative-approach-to-secure-server-access-to-the-internet/12804**.

If you are new to Azure, an Azure Fundamentals course is available at **https://docs.microsoft.com/en-us/learn/paths/az-900-describe-cloud-concepts/**.

Zscaler Internet Access (ZIA) provides outbound internet protection for users. Learn more at **https://www.zscaler.com/products/zscaler-internet-access**.

Zscaler Private Access (ZPA) provides private access to applications, not networks. Learn more at **https://www.zscaler.com/products/zscaler-private-access**.

To learn more about zero trust, visit our zero trust microsite at **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 (NIST) paper at **https://www.nist.gov/publications/zero-trust-architecture**.

Cloud Infrastructure Protection using Cloud Connector

When most organizations began to experiment with cloud infrastructure, the first concern was how to protect communication to that virtual machine. As cloud knowledge matured, so did the complexity of the workloads that were shifted to the cloud. With this complexity, it becomes critical to secure traffic:

- As it moves within the cloud
- As it enters or exits the cloud
- While in transit between clouds

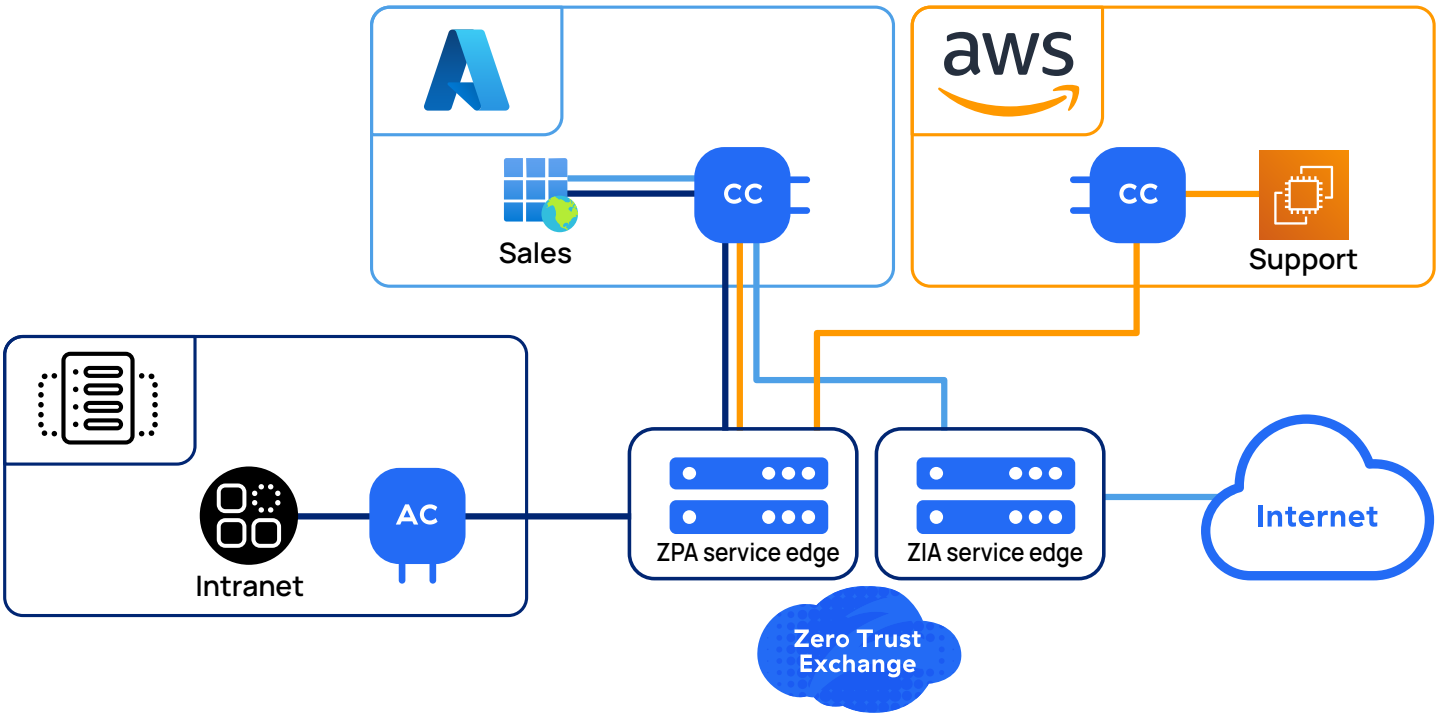

Figure 3. Workload communication between private and public applications

The communication may be from private workloads (IaaS or physical DC) to public workloads (SaaS internet application) or between private workloads (IaaS to IaaS, or physical DC to IaaS). Securing these communications channels with physical or virtual appliances is cumbersome and can lead to inconsistent configuration.

In the example above, our application sales.azure.internal.safemarch.com sits behind a cloud connector with access to both ZPA and ZIA platforms. In this model, the workload can reach out to the support workloads in AWS, allowing the sales team to file support and product requests without logging into the support portal. The sales portal is accessed by our intranet workload in our data center to pull deals and rankings for the company dashboard. Finally, our sales workload can reach the internet to update our cloud CRM, which in turn only accepts connections from Zscaler IPs for our tenant.

Zscaler Cloud Connector virtual machines extend the security of ZIA and ZPA to cloud native workloads. ZIA protects your workload traffic communicating with a public application. ZPA protects your communications between private workloads. This allows organizations to secure all workload communications over any network. The Zscaler Zero Trust Exchange allows workloads to communicate with each other with a granular security policy applied.

- Applications-to-Internet Communications may need to access any internet or SaaS destination, such as third-party APIs, software updates, etc. with a scalable, reliable security solution that inspects all transactions and applies advanced threat prevention and data loss protection controls.
- Application-to-Application Communication to other public clouds and corporate data centers for multi/hybrid cloud connectivity, delivered with better security and a dramatically simplified operational model as compared with traditional solutions like proxies, virtual firewalls, and IDS/IPS.
- Application-to-Application Communications within a Virtual Private Cloud by securing process-to-process communications. This achieves microsegmentation of traffic with no changes to the application or the network.

Cloud Connector is delivered in several form factors. It is available as a virtual appliance in both Amazon Web Services and Microsoft Azure, as well as VMs for on-premises deployment.

If you are deploying on Microsoft Azure:

- Zscaler recommends the Standard D2s v3 instance size to support Cloud Connector as it offers the best performance of 300 Mbps (unidirectional).
- The appliance is available on the Azure Marketplace at: **https://azuremarketplace.microsoft.com/en-us/marketplace/apps/zscaler1579058425289.zia_cloud_connector_app**

For on-premises deployments, the image requires:

- VMware ESXi and CentOS/Linux (KVM) images
- 2 virtual CPUs
- 4 GB of RAM

Deploying Cloud Connector VMs via Scripts

Cloud Connector VMs can be deployed in either two ways: ARM templates or Terraform scripts. ARM templates allow a user to deploy the appliance directly from the Azure Marketplace via guided workflow, are user-friendly, and work well with brownfield deployments. Terraform is the most flexible option. Its goal is to be as "hands-off" as possible by automatically configuring items without user intervention. However, Terraform is more complex in its initial setup. Both options allow you to automate your deployment and achieve the same results. For a detailed look at these deployments, see Deploying Cloud Connector via Terraform and Deploying Cloud Connector via ARM Template in this guide.

High Availability

Cloud Connector leverages Azure Load Balancer functionality to achieve high availability and horizontal scalability. In this model, inbound traffic from workload Virtual Networks (VNets) are directed to the front-end IP address of the load balancer. When traffic returns from the internet, the Cloud Connector appliance strips off Datagram Transport Layer Security (DTLS) encapsulation and forwards the traffic back to the originating workload VNet.

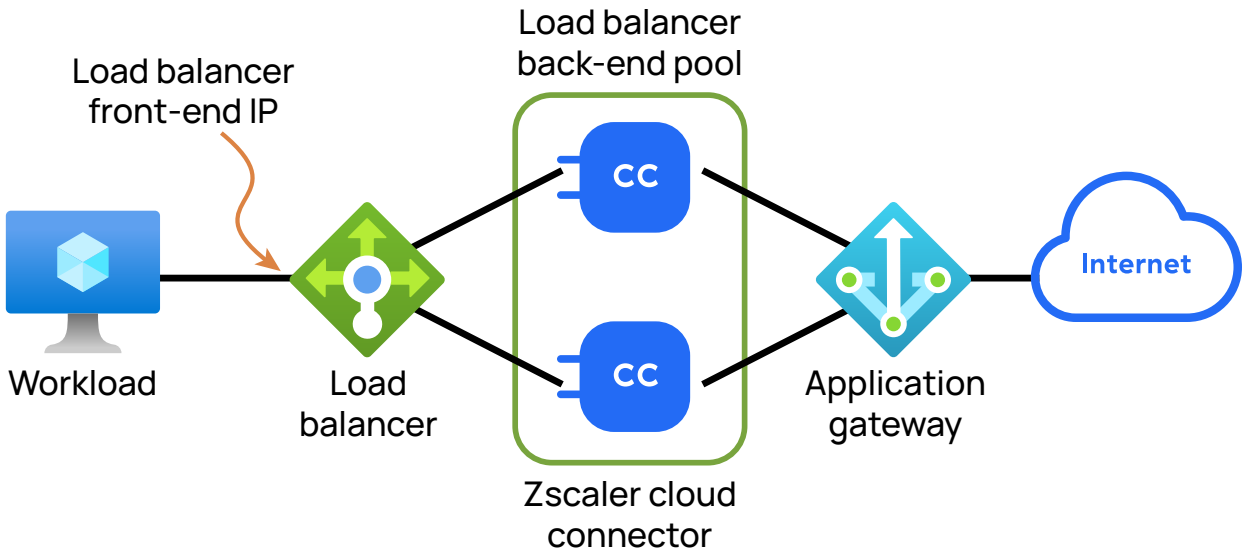

Figure 4. Cloud Connectors receive outbound traffic from the load balancer

Zscaler recommends a minimum of two Cloud Connector appliances, each in a different Availability Zone (AZ). Workloads within those same availability zones should then leverage their respective Cloud Connector appliances. If a Cloud Connector appliance fails, load balancer functionality automatically redirects traffic to the active appliance in the adjacent AZ.

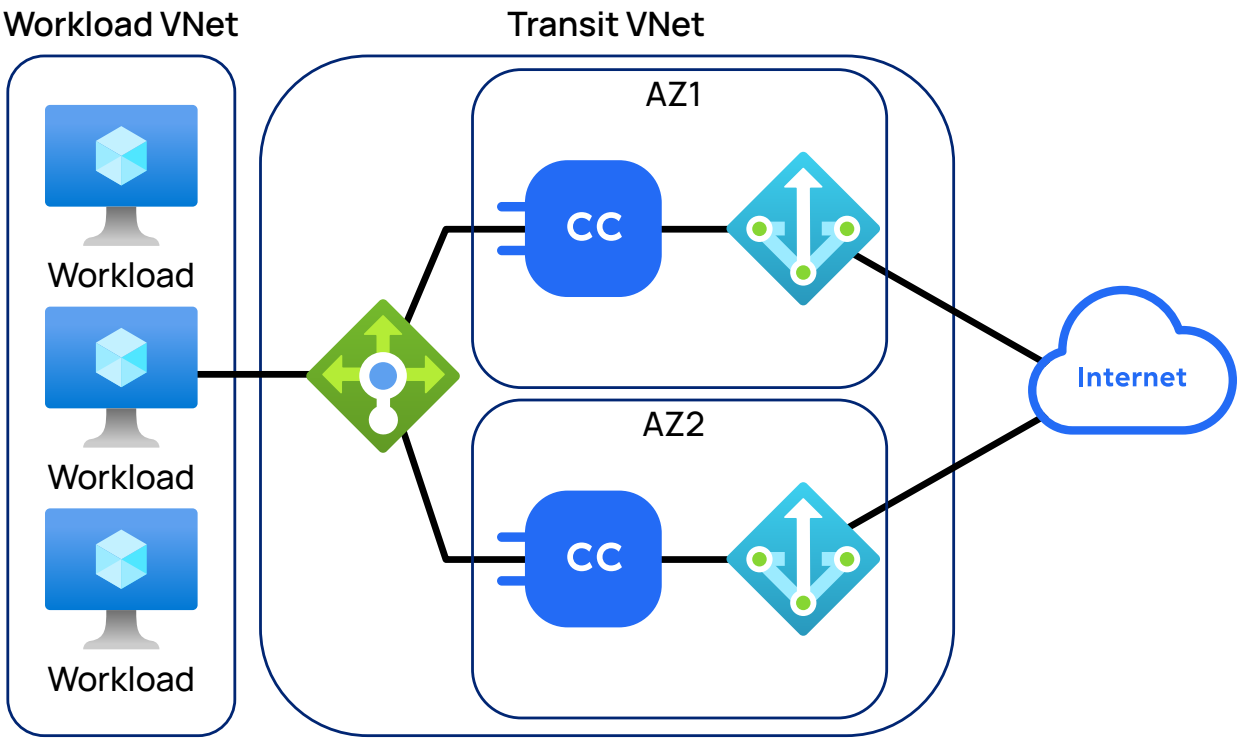

Figure 5. Cloud Connectors deployed in redundant pairs across availability zones

By default, Azure Load Balancer uses a 2-tuple hash (source and destination IP address) to balance traffic. Azure Load Balancer, by default, probes the HTTP Probe Port every 15 seconds. Two probes must fail before considering an appliance down (for a total outage of 30 seconds). However, Microsoft Azure allows tuning down to 5 seconds with two failed attempts as necessary (for a total outage of 10 seconds).

In addition to the appliance-level redundancy, Cloud Connector also maintains redundant DTLS tunnels to the Zscaler cloud. Primary and secondary/backup nodes can be set within the Cloud Connector portal for ZIA, or left as automatic (as is the case with ZPA), wherein the Cloud Connector chooses geographically proximate brokers to connect to.

Terraform scripts can be used to instantiate this functionality, or it can be built manually. Zscaler provides a Terraform template for your use at **https://help.zscaler.com/cloud-connector/about-cloud-automation-scripts**.

Learn more about Azure Load Balancer at **https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview**.

Scalability

Cloud Connector supports two methods of scaling: vertical and horizontal. With vertical scaling, the Cloud Connector can be instantiated with a higher footprint of vCPU and RAM. However, throughput and connection capacity scales linearly with additional resources, and most cloud providers have a bandwidth cap per workload.

Cloud Connector can also be scaled horizontally, wherein multiple appliances are instantiated within multiple availability zones around a region. Inbound traffic to the Cloud Connector appliance can then be load-balanced across all available paths. Zscaler recommends horizontal scaling by deploying additional Cloud Connectors.

Logging

Cloud Connector can utilize built-in logging functionality through the Insights page of the portal. In this dashboard, you can review Session Insights, DNS Insights, and ZIA Tunnel Insights. All three facilities allow you to review traffic that passes through the Cloud Connector appliance from a different perspective.

You can learn more about the insights page at **https://help.zscaler.com/cloud-connector/cloud-connector/analyzing-traffic-using-insights**.

For additional information, particularly related to the disposition of this traffic, further insights can be found in the ZIA and ZPA pages:

- ZIA analytics: **https://help.zscaler.com/zia/about-dashboards**
- ZPA dashboard and diagnostics: **https://help.zscaler.com/zpa/dashboard-diagnostics**

Cloud Connector supports both the Nanolog Streaming Service (NSS) for ZIA use cases and Log Streaming Service (LSS) for ZPA use cases. NSS uses a Virtual Machine (VM) to stream traffic logs in real time to your Security Information and Event Management (SIEM) system, such as Splunk or ArcSight. LSS operates in a similar way, with the deployment of a ZPA App Connector VM that receives the log stream and then forwards it to the log receiver.

Both services enable real-time alerting and correlation of logs with your other devices. NSS and LSS can be configured from the Cloud Connector portal.

NSS and LSS require separate subscriptions.

- To learn more about NSS, visit **https://help.zscaler.com/zia/about-nanolog-streaming-service**.
- To learn more about LSS, visit **https://help.zscaler.com/zpa/about-log-streaming-service**.

Upgrading

Cloud Connector is based on Zscaler OS, so software updates and OS updates are provided by Zscaler. When a Cloud Connector is deployed, the software is automatically upgraded to the latest version. A Cloud Connector checks for new software daily and upgrades itself automatically at midnight (local time, based on the deployed cloud region).

As mentioned throughout this guide, Zscaler recommends that Cloud Connector appliances be deployed as redundant, high-availability appliances. Specific to software upgrades performed by Zscaler, this ensures that the customer incurs no downtime. Once an appliance is rebooted to accept a new update, Azure Load Balancer automatically moves traffic over to the redundant, active appliance.

Although cloud IaaS providers such as Azure are responsible for ensuring the security and availability of their infrastructure, organizations are ultimately still responsible for the security of their workloads, applications, and data. To learn more about the shared responsibility model, visit **https://docs.microsoft.com/en-us/azure/security/fundamentals/shared-responsibility**.

Deployment and Design Options

The following section outlines your options when deploying Cloud Connector. You can design your network using the tools that best match your cloud deployment. We recommend that you review each use case to familiarize yourself with the various options, which can be combined to meet your organization's deployment needs. For example, in many production environments Cloud Connector would be deployed in a Transit/Egress VNet to perform outbound Internet and Zscaler Private Access. Although the use cases do not depend on one another, the concepts and logic depicted within them progressively build on one another.

Pre-Deployment Considerations

Cloud Connectors and Availability Zones

Zscaler recommends that Cloud Connector appliances be installed in pairs for high availability. When building high-availability pairs of Cloud Connector appliances, Zscaler recommends that each appliance be instantiated within different availability zones. This ensures that individual Cloud Connector appliances exist on physically separate pieces of underlying hardware from one another.

To learn more about Azure Availability Zones, visit **https://docs.microsoft.com/en-us/azure/availability-zones/az-overview**.

Network Connectivity

Zscaler recommends employing NAT Gateways for internet access. Just like Cloud Connector, a NAT Gateway is also deployed in an availability zone. NAT Gateways can be shared across availability zones. Zscaler recommends that NAT Gateways also be instantiated as a pair, with one gateway in each of the Cloud Connector availability zones.

To learn more about Azure NAT Gateway, visit **https://docs.microsoft.com/en-us/azure/virtual-network/nat-gateway/nat-gateway-resource**.

Deploying Cloud Connector via Scripts

Cloud Connector can be deployed directly from the Azure Marketplace, through ARM Template scripts or Terraform scripts. While supported, deploying Cloud Connector manually is not best practice outside of lab environments. Zscaler recommends using scripting to provide consistent deployments. ARM Template scripts in Azure are more "native" to their respective platforms, however the preferred method for deploying the appliances is via Terraform.

Deploying Cloud Connector via Terraform

Zscaler Terraform scripts provide complete end-to-end automation to not only instantiate the Cloud Connector appliances, but all the secondary and tertiary components as well (in a best practice configuration). Terraform scripts can be downloaded from the Cloud Connector portal in two versions:

- **Starter Deployment Template** – Instantiate a Resource Group containing Cloud Connector appliance, VNet, Route Tables, Subnet, NAT Gateway, and Network Security Groups for use cases where only ZIA is required. In addition, Terraform also creates a VM instance for use as a Management/Bastion host in the VNet that Cloud Connector is deployed in. This host is not a requirement long term, but is recommended for easier troubleshooting and testing.
- **Starter Deployment Template with Load Balancer** – Instantiate Cloud Connector in high-availability mode using Azure Load Balancer, along with required cloud constructs mentioned previously for use cases where high availability is a requirement. In addition, Terraform also creates a VM instance for use as a Management/Bastion host in the VNet that Cloud Connector is deployed in. This host is not a requirement long term, but is recommended for easier troubleshooting and testing. To learn more about Azure Load Balancer, visit **https://azure.microsoft.com/en-us/services/load-balancer/**.

It is important to note that Terraform does not modify brownfield deployments. When executing Terraform scripts, new VNets, Route Tables, Subnets, and VM instances are spawned to support the current workflow. It is the customer's responsibility to integrate the new deployment into their existing environment. This may mean that the new Cloud Connector VNet is peered with existing VNets, or that new workloads are installed within the Cloud Connector VNet. Bear this in mind when considering whether Terraform is the correct option to use when integrating with a brownfield environment.

For detailed deployment instructions and to find the templates listed above, visit **https://help.zscaler.com/cloud-connector/about-cloud-automation-scripts**.

Deploying Cloud Connector via ARM Template

For customers seeking a more native automation option for deploying Cloud Connector, Zscaler offers ARM Templates. Though ARM Templates can be used in greenfield situations, their value shines when a customer is seeking brownfield integration, since many of the prerequisites are already satisfied if a customer has an existing Azure buildout.

It should be noted that although ARM Templates scripts work well with brownfield deployments, it is still your responsibility to integrate them into the environment.

For detailed deployment instructions and to find the templates listed above, visit **https://help.zscaler.com/cloud-connector/about-cloud-automation-scripts**.

Directing Traffic to Cloud Connector

Cloud Connector acts as a gateway to cloud workloads. Directing traffic through the Cloud Connector is as simple as modifying the default gateway route of the workload route table to point to the appliance, or to the Azure Load Balancer IP. In most circumstances, this ensures that both internet-bound traffic destined for ZIA, and DNS traffic that requires modification for ZPA Use Cases where redirection to an App Connector is necessary, are appropriately handled.

For example, with a single instance of Cloud Connector the workload route table can be updated with a default route using the IP address of the service interface of the Cloud Connector appliance as the target. The Cloud Connector appliance uses the service subnet and route table created during the deployment process. The default route for this route table should point towards the NAT Gateway, also created in the deployment process. A public subnet and route table should have also been created during the deployment process and reference the corresponding internet gateway with its default route.

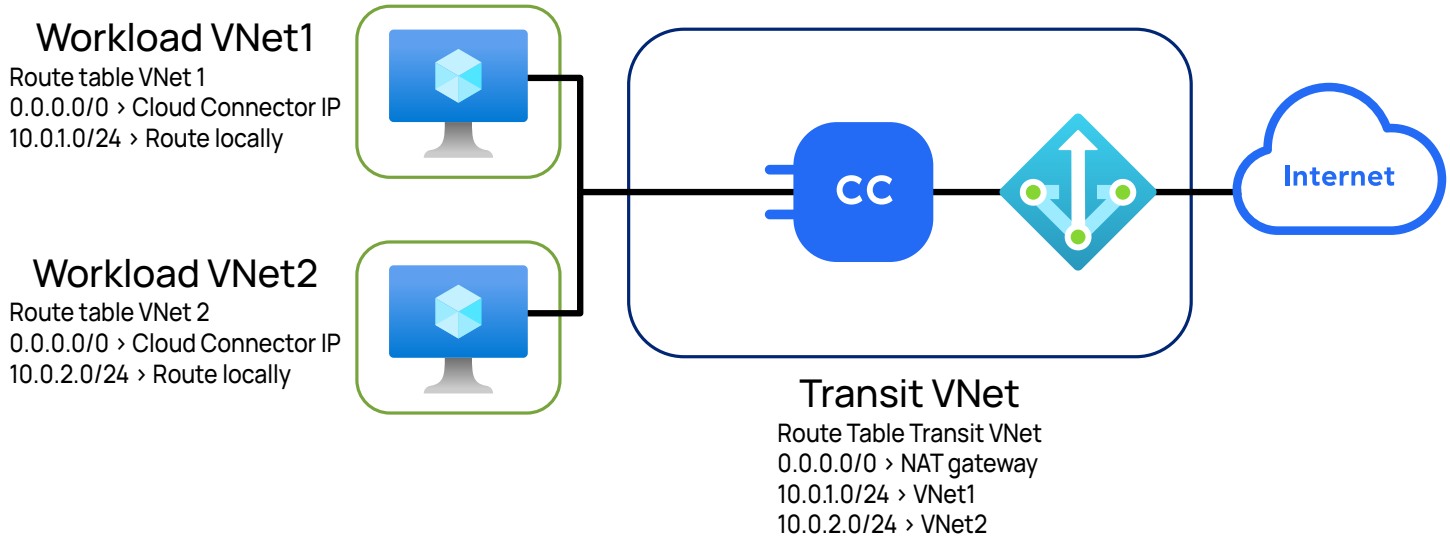

Figure 6. Default routes for workloads go through the Cloud Connector

In the case of hub and spoke, wherein the Transit/Egress VNet is the "hub" and workload VNets are the "spokes," the Transit VNet should be peered with all workload VNets. A default route in the service subnet route table of the Cloud Connector appliance directs traffic towards the NAT Gateway by default. In the workload VNet, a default route is present to direct traffic across the VNet peering towards the Cloud Connector appliance. As with all network traffic, ensure you have routing set up as well so that returning traffic from the internet is correctly directed back towards the initiating host.

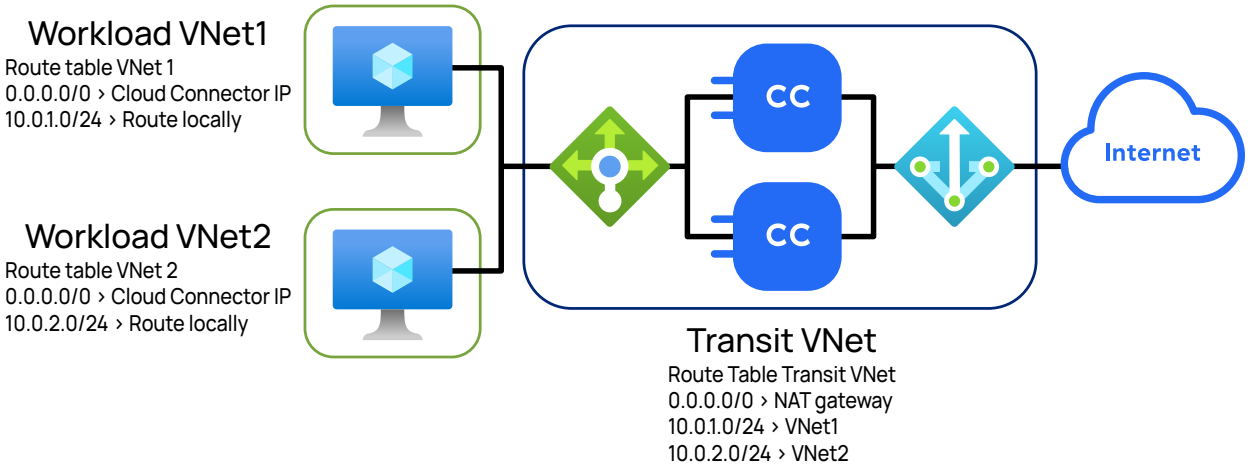

Figure 7. Transit VNets provide access to one or more private subnets

When the Azure Load Balancer is in use in high-availability use cases, Transit/Egress VNet route tables do not change. However, workload route tables are adjusted to use the load balancer front-end IP address as their default gateway.

Forwarding Options

After traffic has reached the Cloud Connector, there are three Traffic Forwarding options available to direct traffic out of the Azure cloud:

- **Direct** – Traffic matching the criteria defined bypasses the Cloud Connector and is routed out of the service interface, where it follows AWS route tables towards the destination.
- **Zscaler Internet Access (ZIA)** – Traffic matching the criteria defined is forwarded to the Zscaler Internet Access cloud for inspection.
- **Zscaler Private Access (ZPA** – Traffic matching the criteria defined is forwarded to the Zscaler Private Access cloud for inspection.

Each of the three options permits the administrator to define a range of match criteria. In general, macro forwarding logic can be defined within the Cloud Connector portal, whereas ZIA or ZPA can perform more granular inspection.

Traffic Forwarding policy is in the Policy Management section of the Cloud Connector portal. Rule creation and assessment models ZIA and ZPA workflows. More specific rules should be ordered near the top, while more broad rules ordered towards the bottom. Match criteria is as follows:

General

- **Location** – Locations identify the various VNets from which your workloads send traffic. As Cloud Connector appliances are brought online, the VNet they are installed within automatically populates this menu. It should be noted, however, that in a Transit/Egress VNet scenario, downstream VNets do not automatically populate. In such a case, you must use Source or Destination IP/FQDN as match criteria. In ZIA, if the traffic is from a known location, the service processes the traffic based on the location settings. For example, the service checks whether the location has authentication enabled and proceeds accordingly. It also applies any location policies that you configure and logs internet activity by location.
- **Location Group** – If necessary, location groups can be created to organize various cloud Vnets, such as a “Dev VNets” location group, “Prod VNets” location group, etc. If there are many locations and associated sub-locations within your organization, consider using location groups.
- **Branch and Cloud Connector Groups** – Branch and Cloud Connector groups allow you to match traffic transiting specific Cloud Connector appliances.

Source

- **Source IP Groups** – When multiple source IP addresses must be matched across multiple policy rules, it is operationally more efficient to create source IP groups. These groups allow you to organize IP addresses for easier rule creation and visualization.
- **Source IP Addresses** – This match criteria allows you to specify the source IP address of the workload.

Destination

- **Destination IP Address / FQDN** – For individual IP address/FQDN matching, enter the value you want to be matched in this field.
- **Destination IP / FQDN Group** – You can group together destination IP addresses and FQDNs that you want to control in a Forwarding Policy rule by specifying IP addresses, countries where servers are located, and URL categories.

Wildcard domain identifiers ("*") are not currently supported.

- **Destination Country** – This match criteria allows you to specify the destination country of the remote machine.

Destination criteria is not supported when Zscaler Private Access is selected as the Forwarding Method.

After configuring a Forwarding Method and match criteria, you must choose an action. By default, for ZIA use cases, the Cloud Connector appliance uses geolocation to locate a ZIA Enforcement Node in geographic proximity to the appliance. Alternatively, you can manually specify which Enforcement Node to use by configuring a gateway under the Forwarding Methods section of the Administration menu.

Gateway selection criteria is not supported when Zscaler Private Access is selected as the Forwarding Method. Cloud Connector automatically selects a broker.

Lastly, specifically for ZPA use cases, Cloud Connector also allows for the filtering of DNS requests/responses. In the Administration menu within DNS Control, administrators can add additional rules to permit or deny specific DNS requests from workload segments. More importantly, this functionality can be used to determine which traffic gets consumed by ZPA, and therefore which synthetic IP Pool is used to address traffic within Microtunnels.

Choosing the Correct Design Model

Cloud Connector is extremely flexible in the ways in which it can be deployed: directly adjacent to the workloads it services, or in a dedicated island by itself where traffic can be directed through it via Azure networking constructs like VNet Peering. There is no single design model that fits every environment. Many organizations pull elements from all design models to suit their goals. There are three main questions to ask when determining how best to get started:

- *Is ZPA a requirement?* Zscaler Private Access requires workload DNS queries to transit the Cloud Connector so a synthetic IP Address can be assigned to the connection. Consider how DNS is employed within the cloud. If using cloud-hosted DNS servers, it is possible that DNS resolution requests are never directed across the Cloud Connector which would break ZPA. For this reason, consider how DNS resolution requests transit Cloud Connector, such as if a public DNS server outside of the cloud is used (or if a custom DNS server is used that forwards these requests). Additionally, if this cloud implementation also services inbound requests from remote clouds, consider pointing App Connectors towards real DNS servers in this scenario.
- *Is high availability a requirement?* Zscaler recommends that high availability be employed in all use cases. However, when deploying directly into the workload VNet, compute costs can quickly spiral out of control, particularly if there are many VNets requiring appliance(s). For this reason, you may consider using a dedicated Transit/Egress VNet with VNet Peering. This allows you to maintain high availability without a large compute footprint.
- *Will Cloud Connector be deployed within the workload VNet, or in a dedicated VNet?* For small environments, only a handful of VNets Cloud Connector can be deployed directly within the workload VNet. However, the number of VNets and VM instances tend to increase as an organization grows larger and invests further in the cloud. As new VNets are added, they will require new appliances. As you consider where the Cloud Connector appliances will be installed, ensure you plan for adequate growth in the number of workloads and VNets that Cloud Connector will protect. If the future state of the environment becomes operationally cumbersome, or if the environment already contains several VNets, it may be best to consider a Transit/Egress VNet approach for Cloud Connector.

Use Case: Direct to Internet using Zscaler Internet Access

Implementing Cloud Connector to provide outbound Internet access through ZIA is one of the first steps to cloud workload protection. The following deployment model represents a recommended option that can be leveraged to satisfy this business requirement and offer a foundation to build on when looking to implement services like ZPA.

In this model, Cloud Connectors can be installed directly into the workload VNet adjacent to the individual workloads they will service. As with all deployment models, Zscaler highly recommends deploying Cloud Connector in high availability. The following image assumes redundant appliances are being deployed with an Azure Load Balancer:

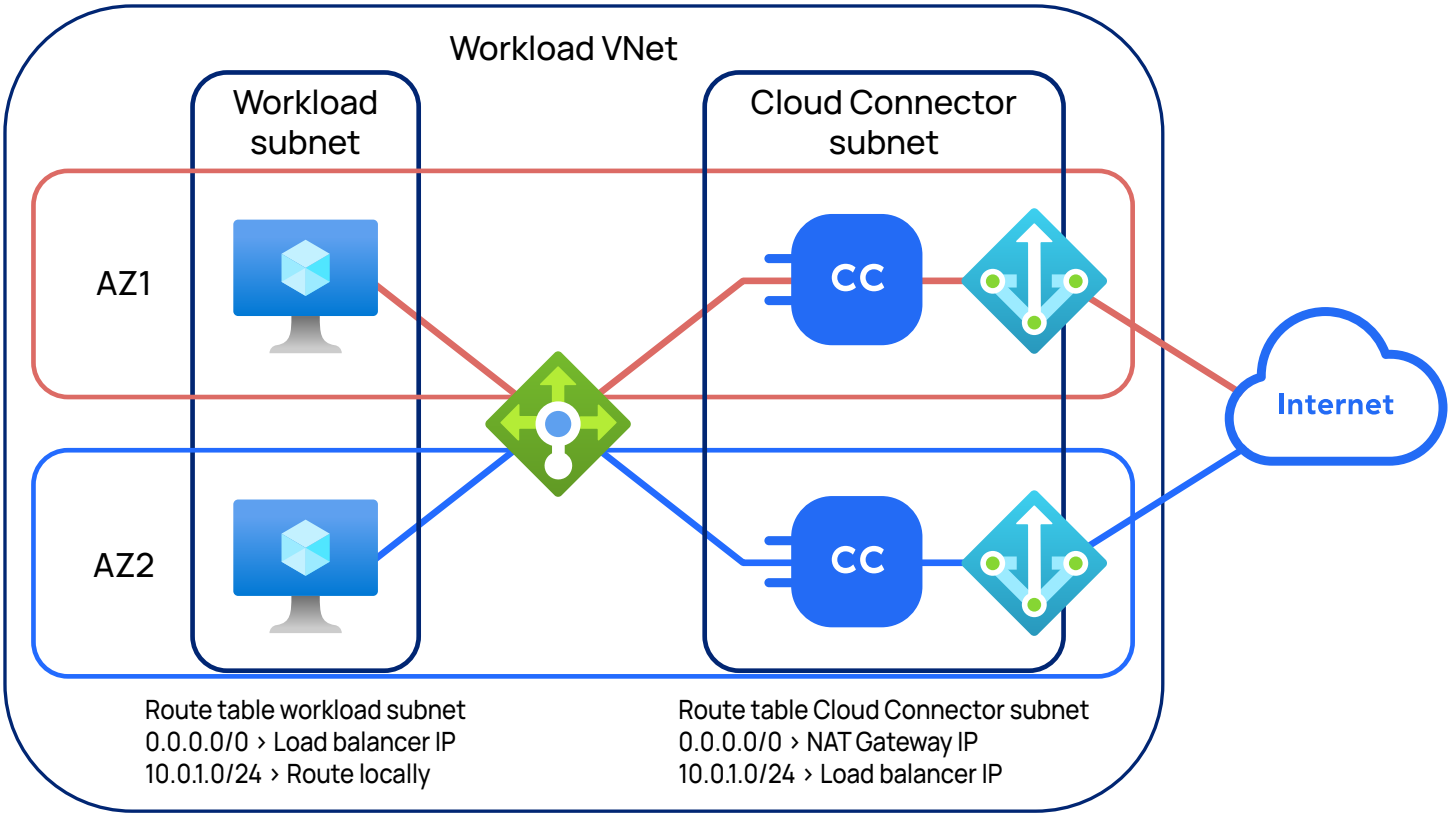

Figure 8. Azure Load Balancer providing redundancy between Cloud Connectors

The primary benefit to this design option is its simplicity and time to implement. Since each Cloud Connector instance is spawned within the workload VNet that it services, routing is made simple. Likewise, whether via Terraform or ARM Template, Zscaler automation can implement this model in a matter of minutes. From a cost perspective, customers are only paying for egressing data fees one time (as the workload traffic leaves the Cloud Connector), as opposed to the double billing that can occur when using a Transit/Egress VNet.

If you have many workload VNets, however, this design option can be cumbersome. Any cost savings associated with egress fees may be eliminated by the increased compute footprint, since separate Cloud Connector VM instances are required per workload VNet. Additionally, this option requires the modification of many route tables to direct traffic accordingly, which is further complicated when high availability enters the picture.

When implementing this design option, the first step is to consider which automation technique to employ. This has been discussed in the **Terraform** and **ARM Template** sections.

- If using ARM Templates, consider deploying the HA Cloud Connector Application within a separate availability zone to instantiate the Azure Load Balancer.
- If using ARM Templates scripts, it is recommended to deploy NAT Gateway. Since NAT Gateway(s) operate within a single availability zone, Zscaler recommends creating a second NAT Gateway in a different AZ so that an infrastructure failure of one AZ does not affect both NAT Gateways. Ensure that the Cloud Connector in the first AZ is in a different subnet than the Cloud Connector in the second AZ. Then, associate each NAT Gateway to each respective subnet.

For brownfield implementations, ARM Templates may provide more seamless integration. For greenfield implementations, consider using Terraform.

Use Case: Leveraging VNet Peering

Cloud Connector can also be placed in a dedicated VNet wherein outbound workload traffic is directed through a centralized hub, such as a Transit VNet. Transit VNets with VNet Peerings is a design option that is growing in adoption as organizations seek to address scalability concerns and operational deficiencies imposed by deploying services directly within workload VNets.

This model closely resembles a traditional hub-and-spoke network since the hub Transit/Egress VNet, where Cloud Connector operates, receives traffic from many workload spoke VNets. As with all deployment models, Zscaler highly recommends deploying Cloud Connector in high availability. The following image assumes redundant appliances are being deployed:

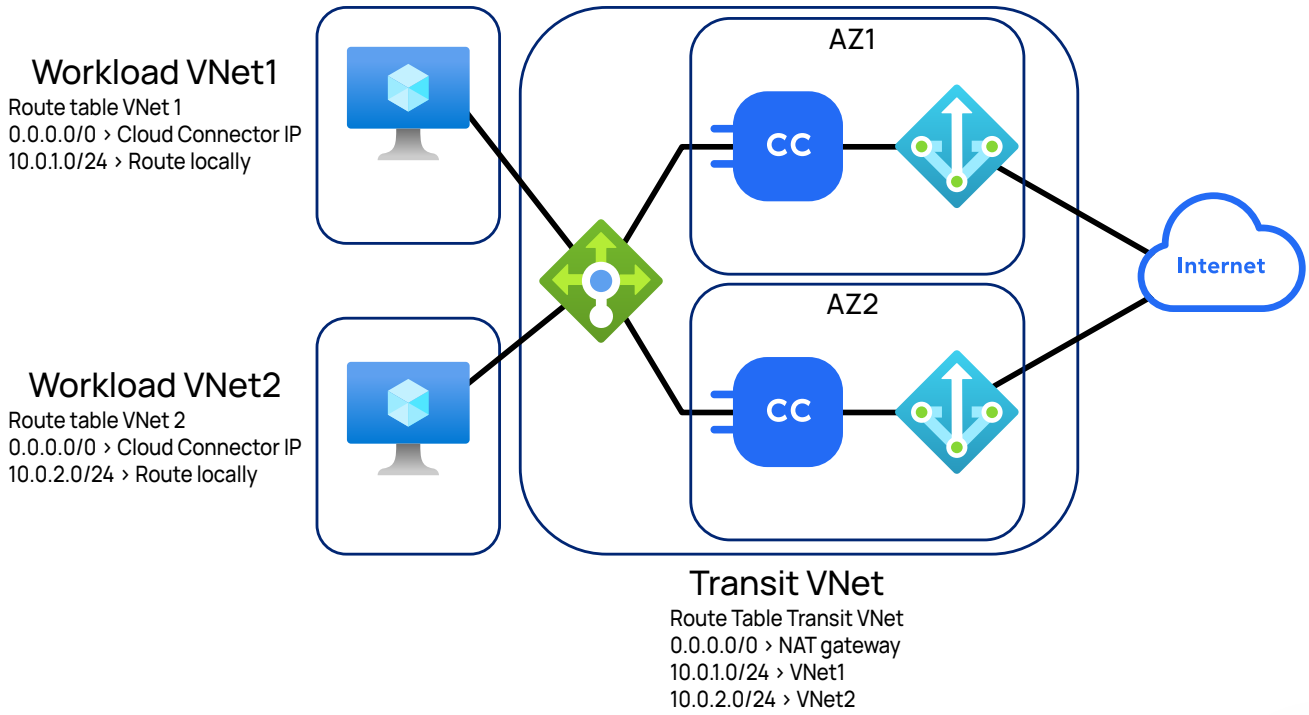

Figure 9. Cloud Connectors deployed in a highly available configuration within the transit VNet

Deploying Cloud Connector using a Transit/Egress VNet allows the organization to simplify cloud routing and, more importantly, reduce the compute footprint required when deploying directly to the workload VNets. In this option, only a single pair of Cloud Connector appliances is necessary for a Transit/Egress VNet. Spoke VNet workloads requiring internet or private access are simply directed towards the front-end load balancer IP address using a simple default route, where they can then be directed towards the Cloud Connector appliances for outbound routing.

By default, Terraform installs Cloud Connector using a Transit/Egress VNet model, though it can be customized to suit an organization's deployment needs.

You should be aware that the potential exists for double billing in this model. Microsoft Azure bills its customers based on egressing traffic out of a VNet. Specifically in this model, the same traffic egresses a VNet twice (once as it travels from the workload VNet to the Transit/Egress VNet, then again as it leaves the Transit/Egress VNet to the internet).

When implementing this design option, the first step is to consider which automation technique to employ. This has been discussed at length in the **Terraform** and **ARM Template** sections previously.

- If using ARM Templates, consider deploying a second Cloud Connector appliance within a separate availability zone. This can be done by simply re-running the ARM Template workflow again. ARM Templates are not capable of instantiating Microsoft Azure Load Balancer. Azure Load Balancer needs to be set up manually or leverage Terraform scripts downloaded from the Cloud Connector portal.
- If using ARM Templates scripts, it is recommended to deploy NAT Gateway. Since NAT Gateway(s) operate within a single availability zone, Zscaler recommends creating a second NAT Gateway in a different AZ so that an infrastructure failure of one AZ does not affect both NAT Gateways. Ensure that the Cloud Connector in the first AZ is in a different subnet than the Cloud Connector in the second AZ. Then, associate each NAT Gateway to each respective subnet.

Use Case: Integrating Zscaler Private Access for Private Cloud Application Access

Assuming that Cloud Connector has been instantiated and traffic directed through it, we can now add support for ZPA. This use case is growing in popularity as organizations seek to depart from legacy VPN technologies to interconnect cloud and on-premises workloads. An important consideration with Cloud Connector is that it is designed to facilitate outbound workload traffic towards a remote destination. When the destination is in a customer-controlled location, we must consider how this traffic ingresses into the remote facility. We do this using the Zscaler App Connector appliance, where App Connector VMs sit adjacent to the workloads they provide access to.

This model builds on the foundation provided in the direct-to-internet and VNet Peering use cases discussed previously. Cloud Connector provides outbound connectivity for cloud workloads to an on-premises data center, which uses App Connector appliances (VMs) sitting in an application server segment to provide inbound connectivity. Both appliances build DTLS tunnels to the ZPA Broker and establish a Microtunnel between the source (cloud) workload and the destination data center workload. The traffic within the Microtunnel targets synthetic proxy IP addresses inside the Cloud Connector and App Connector, respectively.

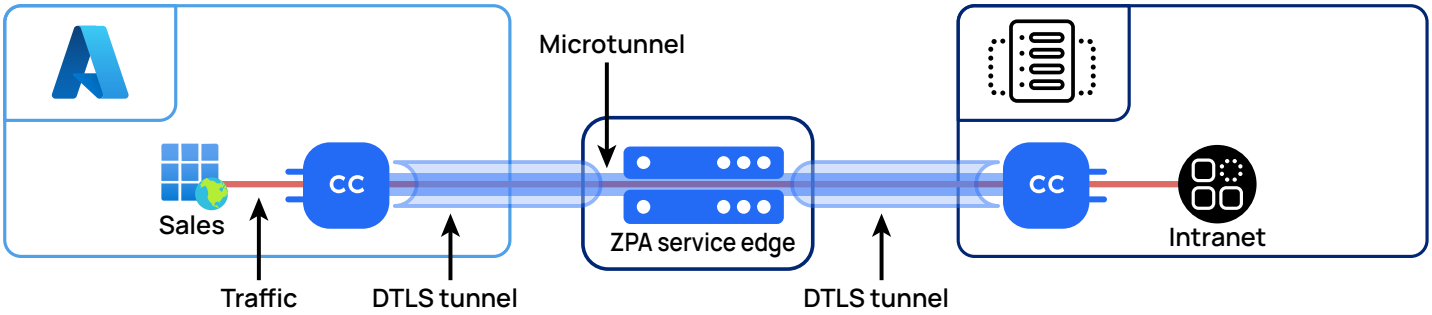

Figure 10. Microsoft Azure access to internal apps in the private data center

All communication between ZPA components travels within a mutually pinned, client and server certificate-verified TLS connection. Within this TLS-encrypted Zscaler Tunnel, a microtunneling protocol (i.e., Microtunnel) exists. Select components of ZPA run through this encrypted Microtunnel end to end. Because the client and server use pinned certificates, it is cryptographically impossible for ZPA to experience a Man-in-the-Middle (MITM) attack. The client certificates are verified against an organization's Certificate Authority (CA) and the server certificates are verified against Zscaler's CA, which cannot be spoofed by any third-party compromised CA.

ZPA only accepts connections from the Zscaler Cloud Connector and the App Connector instances that present a client certificate signed by a CA associated with each tenant. Zscaler Cloud Connector and App Connector only connect to ZPA service components that present a certificate signed by the ZPA infrastructure PKI.

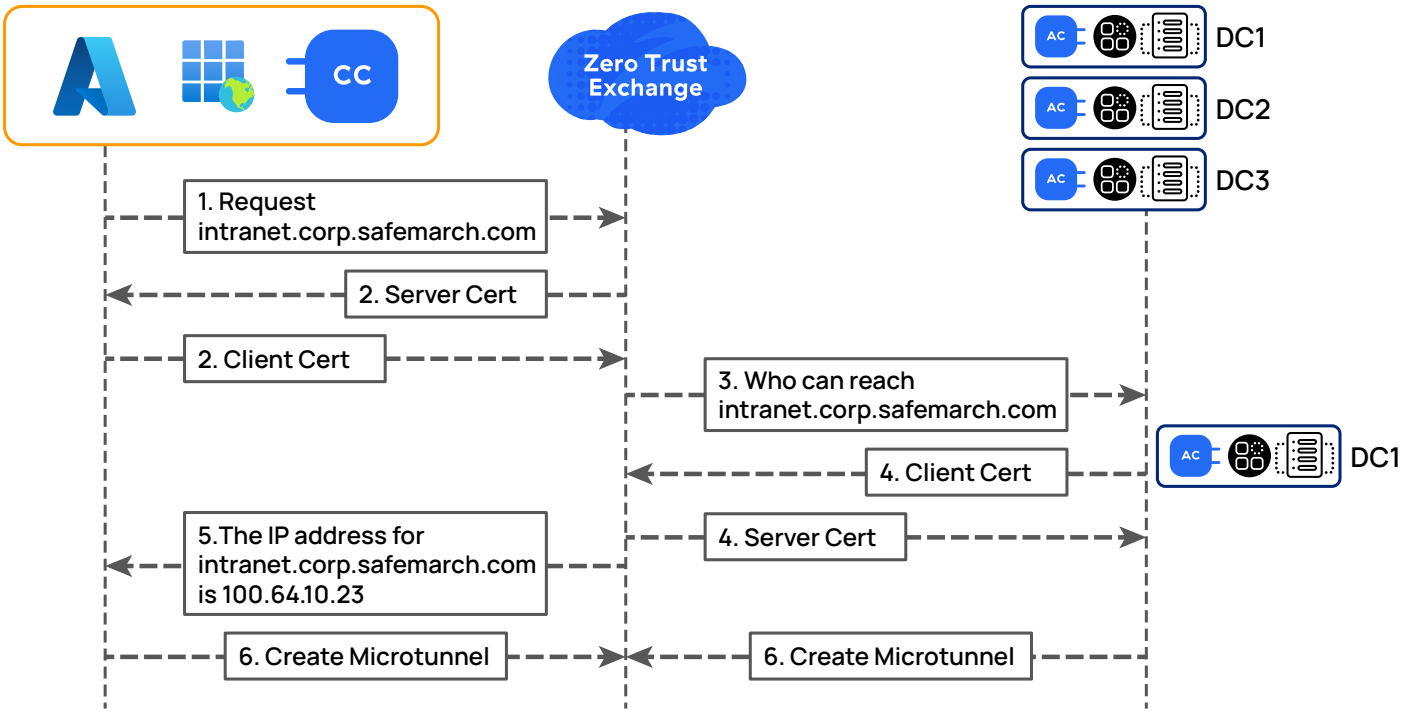

Figure 11. Authentication and tunnel setup between workloads and internal apps

1. A workload requests access to an application.
2. The ZPA Service Edge and Zscaler Cloud Connector authenticate via certificate exchange.
3. If the workload is authorized to access the requested application, the ZPA Service Edge determines which App Connector can service the request.
4. The ZPA Service Edge and Zscaler App Connector authenticate via certificate exchange.
5. The workload is presented with the synthetic IP of the application.
6. A Microtunnel is established between the Zscaler Cloud Connector and Zscaler App Connector.

Zscaler Cloud Connector recognizes the internal applications that are available via ZPA. Access to these applications is defined in ZPA based on policies. Using information received from the ZPA Public Service Edge or ZPA Private Service Edge, Cloud Connector intercepts workload requests for applications, and then forwards those requests to the ZPA cloud.

No network information is required to access available applications. To facilitate secure private connections that are abstracted from the physical network, Cloud Connector associates permitted internal applications with a set of synthetic IP addresses. When a workload sends out a DNS request, Zscaler Cloud Connector can recognize the domain as an internal application being protected by ZPA. Zscaler Cloud Connector then intercepts the DNS request and delivers a DNS response to the workload that uses the synthetic IP address associated with the internal application.

To intercept and modify DNS requests, Cloud Connector must see the initial request from the cloud workload. To facilitate this, Zscaler recommends adding a custom DNS server within the Azure cloud. Ensure internal domain requests are forwarded across the Cloud Connector.

When implementing this design option, the first step is to consider which automation technique to employ. This has been discussed in the Terraform and ARM Template sections.

- If using ARM Templates, consider deploying a second Cloud Connector appliance within a separate availability zone. This can be done by simply re-running the ARM Template workflow again. ARM Templates are not capable of instantiating Microsoft Azure Load Balancer. Azure Load Balancer needs to be set up manually or leverage Terraform scripts downloaded from the Cloud Connector portal.
- If using ARM Templates scripts, it is recommended to deploy NAT Gateway. Since NAT Gateway(s) operate within a single availability zone, Zscaler recommends creating a second NAT Gateway in a different AZ so that an infrastructure failure of one AZ does not affect both NAT Gateways. Ensure that the Cloud Connector in the first AZ is in a different subnet than the Cloud Connector in the second AZ. Then, associate each NAT Gateway to each respective subnet towards the new NAT Gateway.
- Since cloud workloads initiate requests for resources in an on-premises data center towards App Connector, you must ensure DNS requests from these cloud workloads transit the Cloud Connector.
- By default, a rule already exists for ZPA-bound traffic, but you should ensure that the Cloud Connector Forwarding Policy is correctly matching and forwarding traffic to ZPA.
- Ensure application segments have been defined within the ZPA portal and, if using a custom DNS server, ensure these same application domains are forwarded through the Cloud Connector.
- Ensure that Zscaler Private Access policy is configured to accept inbound traffic from cloud locations and allowed (or denied) access to the internet.
- When co-located with other Cloud Connectors, App Connectors must sit parallel to the Cloud Connector and not "behind" the Cloud Connector.

Use Case: Securing Traffic Between Clouds with ZPA

Multi-cloud deployments, where workloads are spread across more than one cloud provider, are becoming more common as organizations look to provide hosting across more than one vendor. You might host your cloud workloads in more than one cloud or, for redundancy or geoproximity, in multiple regions of the same cloud service provider. This use case focuses on how to solve for the challenges faced in this scenario and how we can secure this traffic using the ZPA model discussed previously.

This use case is like that of the ZPA model discussed previously but builds on the fact that remote application destinations secured by ZPA may not be in an on-premises data center. Instead, these applications exist within a different cloud region or in a different cloud service provider altogether. It is common in this scenario to see both Cloud Connector and App Connector co-located in the same workload VNet or Transit/Egress VNet. As originating cloud workloads send requests to remote applications, Cloud Connector routes them to the appropriate App Connectors in the destination cloud:

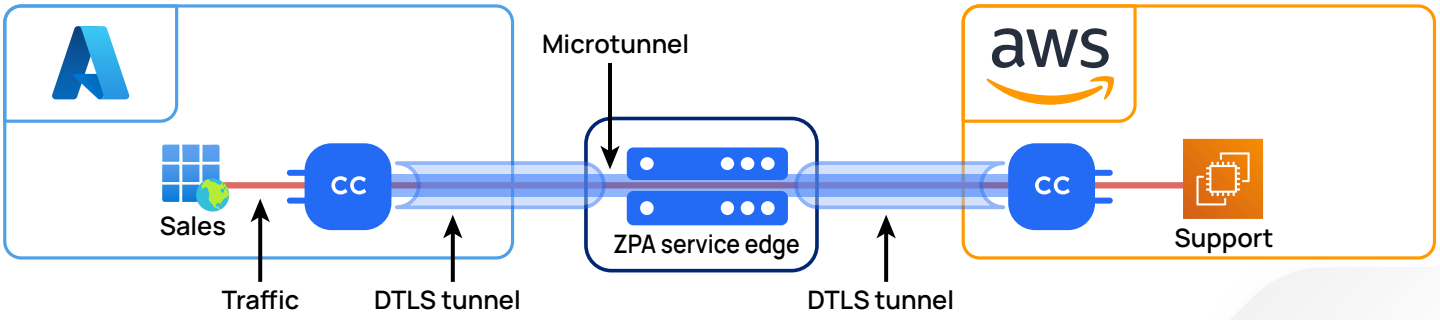

Figure 12. Workload-to-workload communication across cloud providers

When implementing this design option, the first step is to consider which automation technique to employ. This has been discussed in the **Terraform** and **ARM Template** sections.

- If using ARM Templates, consider deploying a second Cloud Connector appliance within a separate availability zone. This can be done by simply re-running the ARM Template workflow again. ARM Templates are not capable of instantiating Microsoft Azure Load Balancer. Azure Load Balancer needs to be set up manually or leverage Terraform scripts downloaded from the Cloud Connector portal.
- If using ARM Templates scripts, it is recommended to deploy NAT Gateway. Since NAT Gateway(s) operate within a single availability zone, Zscaler recommends creating a second NAT Gateway in a different AZ so that an infrastructure failure of one AZ does not affect both NAT Gateways. Ensure that the Cloud Connector in the first AZ is in a different subnet than the Cloud Connector in the second AZ. Then, associate each NAT Gateway to each respective subnet.
- Since cloud workloads initiate requests for resources in an adjacent cloud or region (towards App Connector), you must ensure DNS requests from these cloud workloads transit the Cloud Connector.
- By default, a rule already exists for ZPA-bound traffic, but you should ensure that the Cloud Connector **Forwarding Policy** is correctly matching and forwarding traffic to ZPA.
- Ensure application segments have been defined within the ZPA portal and, if using a custom DNS server, ensure these same application domains are forwarded through the Cloud Connector.
- Ensure that Zscaler Private Access policy is configured to accept inbound traffic from cloud locations and allowed (or denied) access to the internet.
- When co-located with other Cloud Connectors, it is imperative that App Connectors sit parallel to the Cloud Connector and not “behind” the Cloud Connector.
- For inbound traffic from a remote destination, App Connector must be able to resolve the FQDN of the requested host to the real IP address. So, you must ensure that App Connectors are pointed at a real DNS server that can resolve workload FQDNs. Consider this when deploying custom DNS servers.

Summary

Connecting workloads to the internet across different networks is difficult. What makes this harder is the traditional approach used by organizations to solve this challenge, such as using technologies like VPNs and firewalls. While the outcome of connecting these workloads is achieved, the cost to achieve these goals is significant:

- Risk of lateral threats and internet-based attacks by overextending the trusted network across the internet using VPN and WAN technologies.
- Complexity increases because of complicated route filtering, multiple network hops, and fragmented policy management.
- Poor visibility across application connectivity paths and increased network blind spots.
- Costs rise due to overprovisioning network services and the use of virtual appliances such as firewalls, IPs, routers, and other point products in cloud environments.
- Limited scale and performance from the increase in network and security services used in cloud environments.

As a result, there is a need for a better approach. Zscaler Cloud Connector is a cloud-native zero trust access service that provides fast and secure app-to-app, app-to-internet connectivity across multi-cloud environments. With integrated automated connectivity and security, it reduces complexity and cost, and provides a faster, smarter, and more secure alternative to legacy network solutions.

