Transform source IP–address–based application access with Zscaler

WHITEPAPER

Many enterprises employ source IP address identification to control access to applications. But when those organizations adopt SaaS applications, migrate internal applications out of data centers, and support remote work, source IP address identification becomes less effective as a means to secure access to corporate resources.

The way of work has evolved to be cloud-first, device-agnostic, and remote. The means of securing it must evolve too. Enterprise cloud transformation demands new identity-based authorization mechanisms (such as multi-factor authentication, or MFA).

For some organizations, the path from IP address-only to MFA with inline proxy security carries high switching costs. IP address controls may be hard-coded in legacy applications, embedded in internal websites as geo-restrictions, mandated by regulatory requirements, or simply deeply ingrained in the corporate IT security culture.

The majority of user work is now done in the cloud or on the internet, often remotely from "hotspots" and/or on personal devices. Source IP address identification by itself is no longer a reliable nor enforceable security control for governing access to enterprise resources. Zscaler protects the new way of work. This white paper outlines use cases, deployment considerations, and best practices for layering existing source IP address identification controls with Zscaler's state-of-the-art, multi-tenant security model.

The history (and limitations) of IP-address controls

Restricting access to applications or resources based on IP address is a control conceit from an era when users and applications both sat within a perimeter defense. The IP address number identifies the host device seeking access to the application or resource over a corporate network. The security "challenge" is basic: Is this device's IP address within an acceptable set range of numeric values? If yes, lower the draw-bridge. If no, don't answer the doorbell.

In such an enterprise environment, IP addresses are classified into so-called "security zones," in which each zone has an assigned level of security sensitivity. An enterprise device can then access resources based on the privileges afforded to its particular zone. A zone range may be discontinuous, and some host devices might be assigned to a default security zone — assuming the interface is not already explicitly associated with an existing security zone.

IP address controls rely on permissions. To allow access, an application or service compares the source IP address number of the inquiring device to an approved list of numbers (e.g., within an authorized security zone), also known as an "allowlist," and based on the result of the comparison, allows, denies, or challenges the access request. If challenged, the host device may have to provide additional authorization details. (In legacy data center environments, such challenge capability is atypical: Access is usually determined solely by IP address.) If the host device is rejected, its number is added to the deny list. (Note that the deny list also works the other way: IT security may restrict access to a specific URL or IP address range as a destination due to security risk, real or perceived.)

Source IP address-based access controls are fairly easy to implement. If only they were effective on their own. When it comes to securing the new enterprise way of work, source IP address-based access controls have limitations:

- **Poor authentication:** As an identity mechanism, IP address controls recognize a device, not the device's user. (This prevents the application of least-privileged permissions, a key component of zero trust policies.) If any device within an authorized security zone is compromised, everything accessible to that device is vulnerable to attack.

- **Complexity:** IP address management is exceedingly complicated. Improperly configured IP ranges can inadvertently lock out access to admin sites.

- **Ineffective for remote work:** When used for geo-restriction (e.g., specific ranges assigned based on geography), source IP address controls fail when users access resources from new, "out-of-geo" locations.

- **Poor performance:** Source IP restrictions force users to VPN in from remote work locations just so they can egress to the internet via a known IP. That backhauling adds latency.

- **Vulnerable to compromise:** IP addresses can be easily spoofed. One common attack-vector scenario: An open (or weak WEP encryption-based) Wi-Fi network in an allowed address space can easily be exploited to hijack connections and gain access.

Enterprises that “anchor” source-IP addresses to control access to applications and resources must reinvent their approach to protect the new (cloud-first, device-agnostic, remote) way of work. (Their employees are working that way already.) But moving beyond source-IP address controls as an exclusive means of securing access isn’t trivial, and such efforts can incur switching costs.

Zscaler’s cloud-based security services can pair security with dedicated IP addresses, acting as a layer in a cloud security service stack to solidify an enterprise’s threat protection posture; as well as provide a migration path to a stronger security architecture.

Zscaler’s cloud-based security services can pair security with dedicated IP addresses, acting as a layer in a cloud security service stack to solidify an enterprise’s threat protection posture; as well as provide a migration path to a stronger security architecture.

Zscaler + source-IP controls: a practical approach to layered security

Zscaler was founded on the notion that cloud and mobility would disrupt traditional network and security architectures. That disruption is evident in the need for enterprises to move forward from source-IP address-based security controls to identity-based authentication.

To secure the new way of work, enterprise IT leaders must start with an assessment. To what extent does the organization depend on source IP address as an access control mechanism? What’s the gap between existing and ideal security state? And what evaluation criteria (cost, complexity, improved security posture metrics) will help sell such an initiative internally?

That evaluation is the initial stage of enterprise security strategic planning:

1. **Audit use of source IP addresses to allow/restrict access to internal and external resources.**

- Can source IP address control use be modified? If so, what’s the scope of those modifications (on a case-by-case basis)?
- Are approved-security-zone IP addresses mandated by an outside third party (like a government regulator using IP addresses to determine in-geo access)?
- Can internal sites with embedded legacy IP-address coding be updated to more modern (and dynamic) authentication mechanisms like MFA?

2. **Based on assessment findings, prioritize a security migration.**

- Zscaler + source IP address controls: Which enterprise operations should be layered with Zscaler inline-proxy cloud-based security?

- Zscaler-only: Which enterprise operations can be modified to ease dependence on IP addresses as a control mechanism?

Zscaler Internet Access (ZIA) secures user internet egress, protecting an enterprise from both external threats (phishing, ransomware, or other malware attacks) and data exfiltration. Zscaler guarantees a unique egress IP while accessing specific SaaS or other hosted applications.

ZIA and source-IP address-based controls

Enterprises use ZIA (among many other functions) to progress from legacy internet egress methods to local internet breakouts. In the legacy model (hub-and-spoke corporate network with castle-and-moat perimeter security), users connect — often via VPN — to a central web gateway, and then move from there on to the internet. Traffic is backhauled, gateways become bottlenecked, and connectivity performance lags.

Contrast that with the ZIA model, where users go online at the nearest internet onramp, and enjoy direct, secure, fast, and optimized access to internet resources, including SaaS applications like Office 365. In this new model, the concept of connecting to a corporate network (and then to the internet) goes away. ZIA acts as an inline proxy: Zscaler terminates the original connection from the customer's device or network and initiates a new, direct connection to the destination content server on behalf of the user. The source IP address seen by the content server is a public-egress IP address from the Zscaler data center, and not the original IP address of the enterprise user's device.

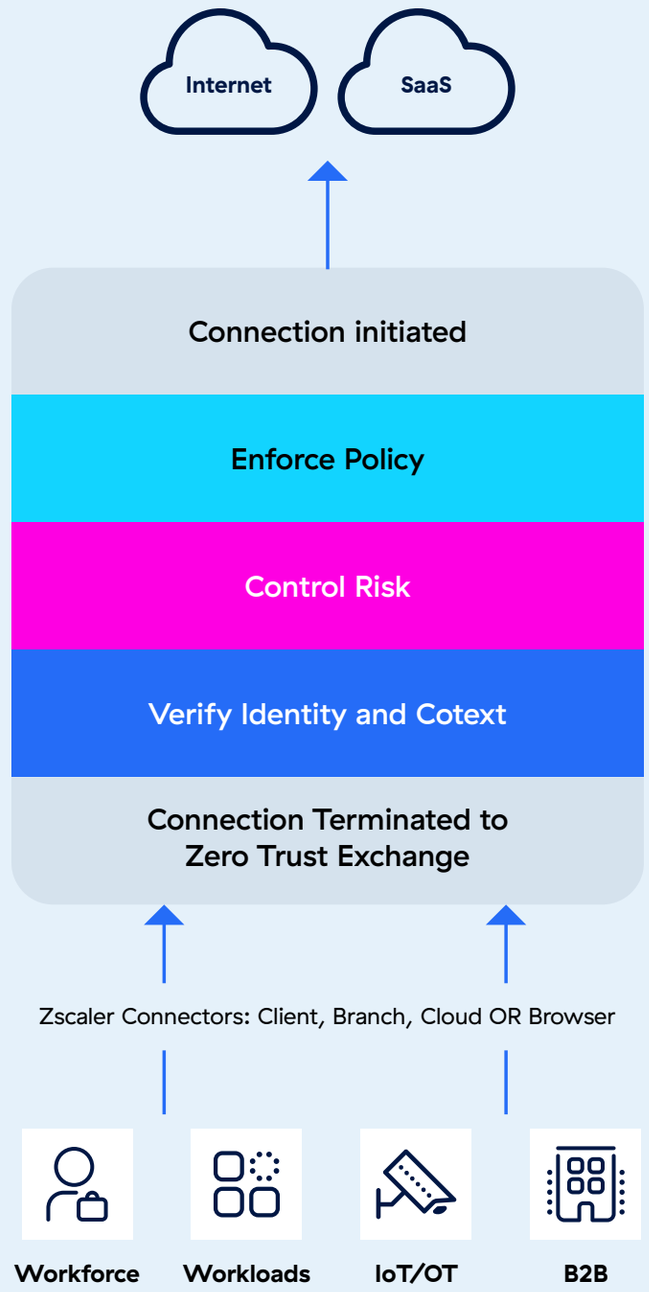

Figure 1. ZIA offers direct, secure, fast, and optimized access to internet resources, including SaaS applications like Office 365.

The ZIA proxy function allows Zscaler to inspect all content traversing from client to server and back, and protect the user if the user visits a potentially-malicious (or compromised) destination. The use of Zscaler IPs on the egress acts as a form of network address translation (NAT) protection, shielding device IP address from the destination content server. (Note that device IP address is inserted into the XFF header.)

For enterprises that still rely on source-IP address allowlisting, NAT address-masking can interfere with application access, since a destination application won't recognize a Zscaler IP address as being within an acceptable "security zone" range.

FEATURE	DETAILS
CAPABILITIES	
URL filtering	Allow, block, caution, or isolate user access to specified web categories or destinations to stop web-based threats and ensure compliance with organizational policies.
SSL inspection	Get unlimited TLS/SSL traffic inspection to identify threats and data loss hiding in encrypted traffic. Specify which web categories or apps to inspect based on privacy or regulatory requirements. Integrate inspection with developer tool to secure developer workflows.
DNS security	Identify and route suspicious command-and-control connections to Zscaler threat detection engines for full content inspection.
Dedicated IP	Provide access to applications that allow-list IP addresses with IP addresses dedicated to your organization. Also enables an easy transition from legacy architecture to zero trust.
File control	Block or allow file download/upload to applications based on app, user, or user group.
Bandwidth control	Enforce bandwidth policies and prioritize business-critical applications over recreational traffic.
Country-based logging	Store and manage logs within a specific country's borders to meet compliance with data sovereignty requirements that mandate data related to citizens be processed according to local laws.
Advanced threat protection	Stop advanced cyberattacks like malware, ransomware, supply chain attacks, phishing, and more with proprietary advanced threat protection. Set granular policies based on your organization's risk tolerance.
Inline data protection (data in motion)	Use forward proxy and SSL inspection capabilities to control the flow of sensitive information to risky web destinations and cloud apps in real time, stopping internal and external threats to data. Advanced inline protection is provided whether an app is sanctioned or unmanaged without requiring network device logs.
Out-of-band data protection (data at rest)	Use API integrations to scan SaaS apps, cloud platforms, and their contents to identify sensitive data at rest and remediate automatically by revoking risky or external shares, for example.
Intrusion prevention	Get complete threat protection from botnets, advanced threats, and zero-days, along with contextual information about the user, app, and threat. Cloud and web IPS works seamlessly across Firewall, Sandbox, DLP, and CASB. Deploy tailored threat signatures using Cloud Custom IPS to detect and stop targeted attacks.
Dynamic, risk-based access and security policy	Automatically adapt security and access policy to user, device, application, and content risk.
Traffic capture	Seamless Packet Capture: easily capture decrypted traffic via specific criteria within Zscaler policy engines, supporting efficient security forensics without requiring additional appliances.
Malware analysis	Detect, prevent, and quarantine unknown threats hiding in malicious payloads inline with advanced AI/ML to stop patient-zero attacks.
DNS filtering	Control and block DNS requests against known and malicious destinations.

Zero Trust Browser (Web isolation)	Make web-based threats obsolete by delivering active content as a benign stream of pixels to the end user's browser.
Correlated threat insights	Speed investigation and response times with contextualized and correlated alerts with insights into threat score, affected asset, severity, and more.
Application isolation	Allow safe, agentless unmanaged device access to SaaS, cloud, and private apps with granular control over user actions like copy/paste, upload/download, and print to stop sensitive data loss.
Workload-to-internet communication protection	Prevent compromise and stop lateral movement for workload-to-internet communications. Includes SSL inspection, IPS, URL filtering, and data protection for all communication.
IoT Device Visibility	Gain a complete view of all IoT devices, servers, and unmanaged user devices across your business, with automated discovery, continuous monitoring, and AI/ML classification with industry-leading auto-labeling capabilities
Role Based Access Control (RBAC)	Right-sized permissions to control what administrators can edit and view policy and analytics reporting within the Zscaler platform to prevent conflicts and improve governance.

Source IP address-based access controls: legacy use cases

Source IP address as an application access control mechanism can be categorized into four primary enterprise use cases:

1. **Controlling access to an external SaaS application**
2. **Using source IP address as a step-up authentication policy attribute**
3. **Allowing/restricting incoming connections at a perimeter firewall**
4. **Geo-locating based on source IP address**

1. Controlling access to an external SaaS application

Many applications — including SaaS applications, remarkably — continue to use IP address as an authorization criterion for access to an application server. When an inbound connection request is detected, the app compares source IP address to the allowlist (e.g., approved "security-zone" range of numbers) and allows or rejects access.

In many cases — including most SaaS application access methods — this form of access typically does not supplant application-level authorization, yet it is still commonly used. It can be employed in enterprise environments as a supplemental access mechanism for applications (often ones that carry legacy IP address-based access control code) that have been migrated from a protected data center to the cloud or internet.

Ensuring the application does not grant access to an unauthorized user requires a more modern security approach like MFA. And most SaaS applications now support single sign-on (SSO) and Security Assertion Markup Language (SAML).[1] In some SaaS applications, the source IP address will be used for tenant and authentication scheme identification, and will be mandatory as it allows the service to select which tenant and Identity Provider (IdP) to use for an incoming connection.

In a legacy network design, the egress IP address will be the public IP or IP ranges of all the customer locations, after traversing the NAT boundary in the location's firewall. A legacy network will have a relatively small number of egress locations and IPs, making the management of allowlists on cloud and partner applications manageable. When roaming users want to access these applications, they must VPN into a location, so they can egress through that location's IP to be authorized for access.

Zscaler adoption transforms an enterprise corporate network. With ZIA, the egress IP seen by the application will be in Zscaler egress IP ranges (which includes tens of thousands of IPs in hundreds of subnets), which makes the simple IP-allowlisting approach ineffective, since multiple Zscaler customers are assigned IP addresses within discontinuous ranges.

While some applications offer IP address allowlisting as an optional layer of security, some services (common in B2B situations) require source IP-allowlisting as mandatory and will not onboard a partner without an explicit list of source IP addresses. For example, Company A may have contractors working on systems in Company B, and are required to initiate connections from a predesignated IP range in order to gain access to Company B's systems. Other common examples are VAT declaration applications and others hosted by government organizations, research terminals hosted by Bloomberg or Thomson Reuters, and banking applications.

2. Using source IP address as a step up authentication policy attribute

Source IP address can be used as a decision criterion to escalate authentication challenges. For example, in an enterprise environment that uses source IP address for allowlisting (see use case #1 above), an incoming device connection would be allowed based on a single factor of authentication (the IP address number itself) if the source IP address is within an acceptable range. But if there are situations in which an outside-the-range IP-addressed device might need access, then the IP address check becomes a challenge. If the source device's IP address is not within allowable range, then a second factor (or more) of authentication (either a one-time password or RSA key entry) is required.

This use case is allowlisting (use case #1) plus an additional authentication challenge to allow for unrecognized IP address access. In that way, it enables slightly better support for remote work, since a user can theoretically log on (with second-factor authentication validation, of course) from a new device. But it still is a rather selective form of authentication (tied to device, not user), and the workflow conflicts with ZIA, since ZIA connections would all appear to come from unknown source devices.[2]

[1] SaaS vendors have widely adopted Security Assertion Markup Language (SAML) allowing Identity providers (IdP) to pass user credentials to service providers (SaaS). Historically, MFA was viewed as a clunky, enterprise-class solution: hard to implement, difficult to roll out. Administration was difficult, and end users disliked having to carry a token around for each service. But with the advent of mobile smartphones, applications have emerged to make it easier to generate MFA pins or tokens. Further, many web applications now enable MFA as an add-on capability, reducing some administrative overhead. This ease of use and setup is a big driver in the increasing popularity of this security feature. Zscaler recommends deploying Identity Federation using SAML for provisioning and authenticating users.

[2] This approach can impact access to SaaS apps like Office 365. Microsoft offers guidance on bypassing MFA for direct connection to O365 servers here: https://blogs.technet.microsoft.com/latam/2018/08/18/skip-mfa-o365/

3. Allowing/restricting incoming connections at a perimeter firewall

Some enterprises — when migrating applications or data from internal networks or data centers to public IaaS clouds — seek to restrict access to the virtual networks that host the relocated applications. In this model, IT essentially extends a perimeter firewall around a virtual network, virtualizing a castle-and-moat secured network (with all its known security limitations) in the cloud. Inbound access is allowed through a VPN (with access to the VPN granted based on source IP address, of course) or based on allowlisted IP addresses within a select range (configured in the virtualized firewall rule set).

Rarely are cloud-migrated, custom-developed internal applications "hardened" or tested to a point that they can be opened to the internet. Refactoring is expensive (and not always practical), and in this use case example, source IP address-based access control — in spite of its limitations — is entrenched (literally and figuratively) as a critical security feature.

4. Geo-locating based on source IP address

Some websites present dynamic content based on IP address geo-location. Others — including many media services and government sites — use source IP address identification to restrict access to content. ZIA, which employs NAT to reassign IP address to egress traffic, can alter perceived geo-location. For instance, a user in Canada may log on to Zscaler via a nearby server over the border. The destination site recognizes the reassigned IP address, and presents content to what it thinks is a device located in the U.S.

Unfortunately for the sites that rely on it, IP address-based geo-location isn't particularly accurate anymore:

- **"Anycasting" can obfuscate device pinpointing:** When an IP address prefix is simultaneously announced from multiple locations, it is said to be "anycast," a connectivity optimization technique commonly used by CDNs, DDoS mitigation services, and DNS providers to route traffic to its destination in the fewest network hops. But that prefix can appear to be in multiple locations at once (depending on vantage point), making source device nearly impossible to accurately geo-locate.

- **Mobile devices are, well, mobile:** Users working remotely may move and stay connected. When device locations move over a relatively short period of time (e.g., taking the train from one city to another), it's hard for destination sites/content servers to definitively associate an IP address to a specific place.

- **Subscription data services route traffic through their own gateways:** Many service carriers (think mobile device providers) use centralized gateways as onramps to the public internet. That misdirection can easily confuse destination sites into thinking source device access is coming from the location in which the gateway resides.

More modern techniques like GPS-pinpointing, cell tower triangulation, and even physical billing address correlation can help improve user geolocation accuracy. But those options are not always available to enterprise IT managers.

Employing Zscaler with source IP–address access controls: solution and deployment considerations

For organizations that need to retain source–IP address application access controls, Zscaler's massively–distributed, inline–proxy, cloud–based, edge service provides a necessary added layer to the security stack.

Enterprises seeking to use their own IP addresses — or an allowlisted subset — with Zscaler should consider six potential solution approaches:

1. **Selectively allowlist Zscaler IP addresses**
2. **Zscaler Managed Dedicated IP**
3. **Customer Managed Dedicated IP**
4. **Leverage XFF headers**
5. **Use Zscaler private–cloud infrastructure (ZIA Service Edges, formerly known as "VZEN/PZENs")**
6. **Forward traffic (after ZIA) to a northbound web proxy**

1. Selectively allowlist Zscaler IP addresses

Allowlisting select IPs (or even a few data centers) for access to cloud applications can preserve existing processes (or business logic, or site code). But incorporate additional authentication mechanisms like MFA to further reduce attack surface (and the possibility of other Zscaler customers inadvertently gaining access to your resources).

2. Zscaler Managed Dedicated IP

For businesses relying on applications with strict IP allow–listing requirements, maintaining a consistent and trusted network identity is essential. This is enabled at scale by Zscaler Managed Dedicated IPs.

Zscaler–managed Dedicated IP uses Forwarding policies to steer traffic processed by ZIA to the internal or external destination servers of your choice via Dedicated IP addresses. This approach ensures that

- Zscaler secures the traffic
- the source IP address seen by the destination is dedicated to the customer,
- the IP addresses are either owned by Zscaler or owned by the customer, and
- the dedicated source IP addresses are hosted and provisioned by Zscaler for your organization

Users can configure granular policies in the ZIA Admin Portal to forward the selected traffic with Dedicated Source IPs through ZIA's threat and data protection engines.

Organizations can view the IP addresses provisioned for their organization in their ZIA tenants after the Dedicated IPs are provisioned at a Zscaler data center.

Zscaler Managed Dedicated IP Solution

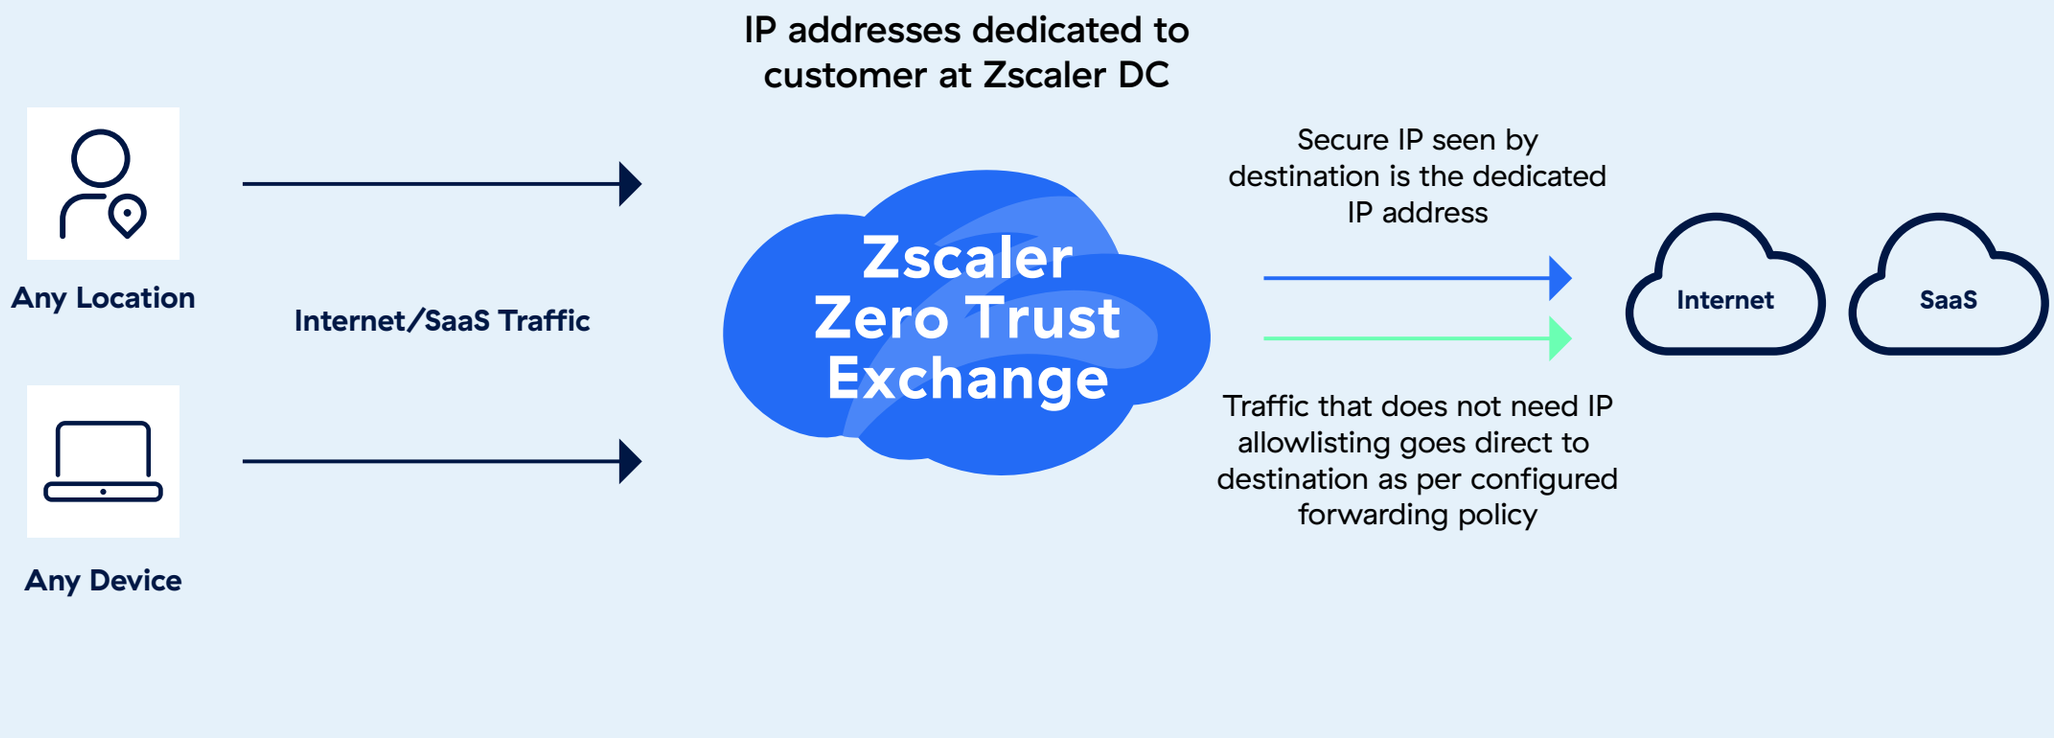

3. Customer Managed dedicated IP

Customer Managed Dedicated IP aka Zscaler Source IP Anchoring allows the configuration of forwarding policies in ZIA to forward traffic destined to specific applications to Zscaler Application Connectors. Zscaler Application Connectors are a widely-deployed and essential component of Zscaler's ZPA solution. Zscaler Application Connectors route traffic to a specific destination based on business policy, in effect enabling the preservation of source IP address controls in parallel with Zscaler security controls.

As shown in Figure 3 below, in a Zscaler Source IP Anchoring workflow, devices and/or users forward the traffic to their nearest Zscaler data center. Applications that need a known source IP address are forwarded to the Zscaler App Connector through the ZPA broker after any security policies configured for the traffic have been applied. In this scenario, the enterprise has the flexibility to select traffic to be sent to the connector (applying policy by user, location, group, destination IP, and/or application type) and steer traffic to the Zscaler Application Connector hosted in the customer's VPC or on-premises data center. When the traffic reaches the destination application or URL hosted in the internet, the source IP seen by these applications is the IP address of the App Connector. Traffic to the internet or SaaS applications that do not need a known source IP address is sent directly from the Zscaler data center after the configured and relevant security policies have been applied.

The Zscaler Application Connector IP address will be the user's designated IP address for application access. In this way, the enterprise benefits from ZIA security scanning without conflicting with existing source IP address-based access controls.

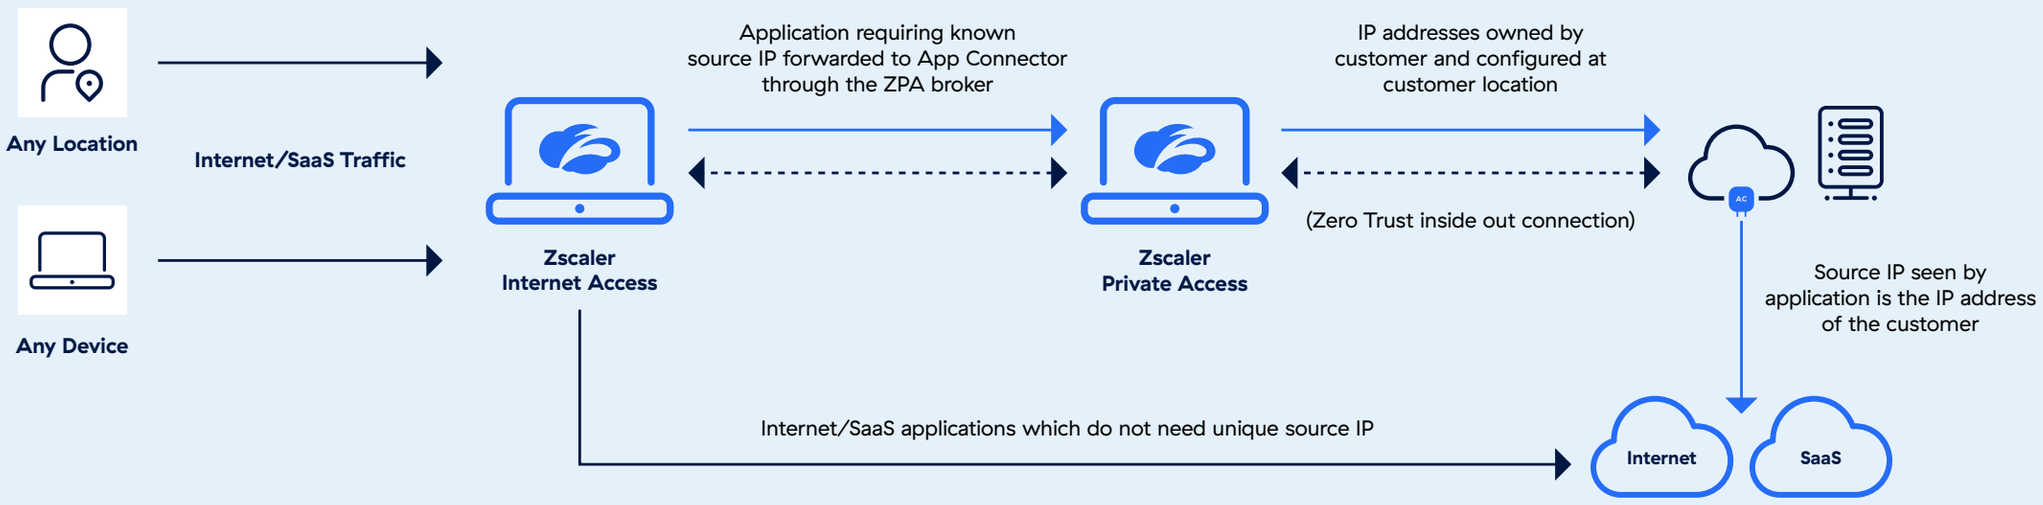

Zscaler Application Connectors can be deployed in several ways. Zscaler distributes a standard virtual machine (VM) image for deployment in enterprise data centers, local private cloud environments such as VMware, or public cloud environments such as Amazon Web Services (AWS) EC2. Zscaler Application Connectors can also be installed on supported Linux distributions.

Zscaler Application Connectors can be co-located with enterprise applications, or deployed in any location with connectivity to the applications. Typically, they are deployed on network segments that can access secured applications and the Zscaler cloud simultaneously, such as in a DMZ. Notably, Zscaler Application Connectors connect only outbound. They do not require reserved inbound open ports to function. Zscaler Application Connectors are always active, and are deployed in a redundant configuration. They never communicate with each other.

4. Leverage XFF headers

"XFF" is the "x-forwarded-for" web proxy function. Internal enterprise content servers employ it to indicate the original source IP address of a device (before the data is forwarded or routed via a proxy like Zscaler). XFF is inserted by default for all HTTP traffic going through Zscaler. If the destination application or content server can read and interpret the incoming XFF header, it can apply its source IP address-based application access rules. This satisfies use cases 1, 2, and 3 above, and the enterprise enjoys the added security Zscaler provides.

Unfortunately, many applications are not able to read or act on XFF information. In addition, as enterprises establish local internet breakouts for users, and allow remote access to applications, IP addresses grow, and the number of IPs seen in XFF headers can quickly become unmanageable.

5. Use Zscaler private-cloud infrastructure (ZIA Service Edges, formerly known as “VZEN/PZENs”)

Zscaler offers ZIA Service Edges for enterprises seeking to establish a private-cloud infrastructure. ZIA Service Edges are physical or virtual enforcement nodes that can be deployed on premises or in a private cloud. ZIA Service Edges can be configured to act as security proxies for some or all of an enterprise’s data traffic. Traffic egressing a corporate network is forwarded to the organization’s private ZIA Service Edge, where it picks up the external IP address—as assigned by the organization.

The ZIA Service Edge approach allows an enterprise to preserve source IP address–based application access controls, as the enterprise hosts the ZIA Service Edges and provides public IP addresses from its own external address ranges (public IP addresses for physical ZIA Service Edges, internal IP addresses for virtual). This model extends the ZIA dataplane on premises but within the oversight of the enterprise’s IT team, and requires additional deployment considerations (like racking, power, public IP assignments, and infrastructure budgeting). Remote user traffic must be backhauled to the private service edges before reaching restricted destinations.

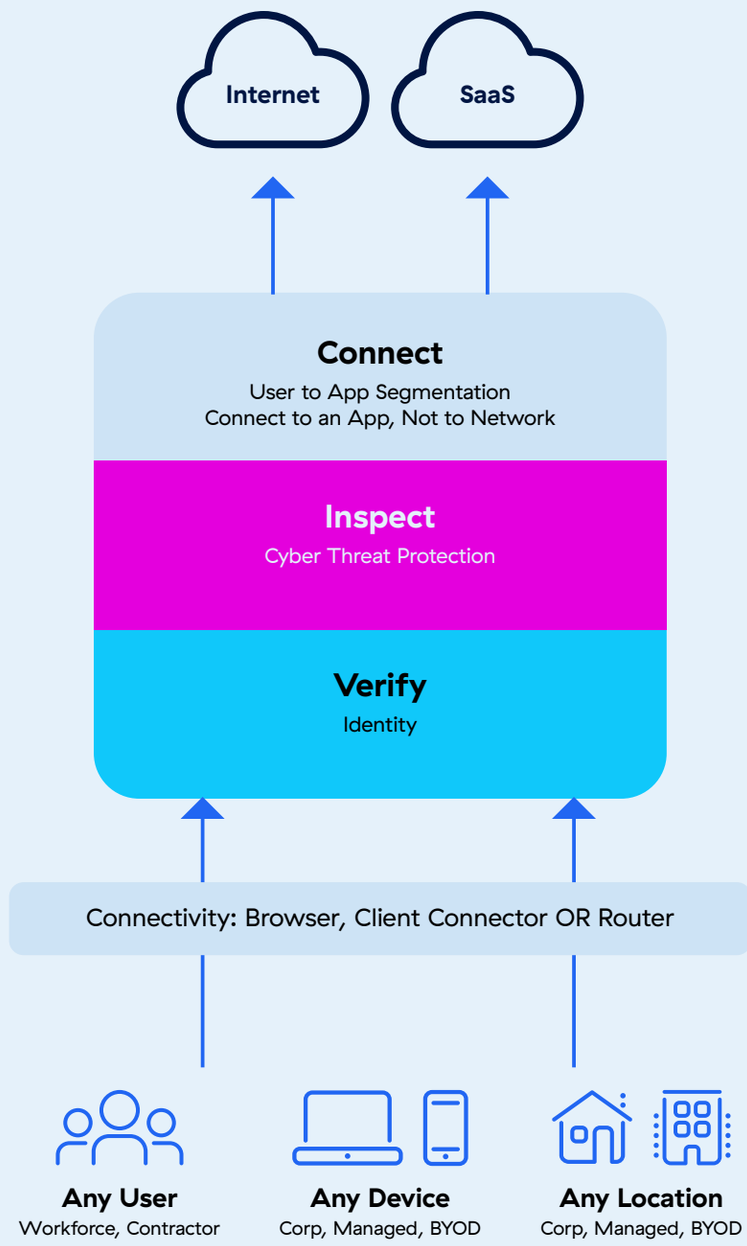

Figure 4. The Zscaler Service Edge model enables an enterprise to preserve source IP–address–based application–access controls

Virtual ZIA Service Edges are commonly used in cases where regulatory requirements such as those in the public sector that require private address space usage and the traffic cannot be routed through physical ZIA Service Edges.

6. Forward traffic (after ZIA) to a northbound web proxy

Zscaler enables proxy-chaining, the capability for an enterprise to selectively forward specific traffic that requires a known source IP — while still maintaining the ability to scan the traffic — to another authenticated web proxy (such as a Squid forwarding/caching proxy). The additional “northbound” proxy may be on premises or in the public cloud. It receives the incoming HTTP traffic, and can assign an acceptable IP address to the data before sending it on to a destination site or application.

In addition to the potential performance impact of lengthening data-travel distance, this solution introduces additional risk: the added proxy extends potential attack surface exposure beyond internal and Zscaler controls. To mitigate that, Zscaler supports the additional proxy only in northbound egress from Zscaler.

Solution Reference Summary Table

help.zscaler.com/zpa/connector-deployment-prerequisites

SOLUTION	DEPLOYMENT CONSIDERATIONS	RECOMMENDATIONS	IMPLEMENTATION SIMPLICITY VS. LEVEL OF CONTROL
Selectively allowlist Zscaler IP addresses	• Could inadvertently provide access to other customers using the same IPs/ranges. • Large IP address ranges (such as /23) to allowlist may not be acceptable to some. • Applications need to be capable of supporting this mechanism. • May not scale well across numerous internet breakout locations.	• Simplest to implement. • As best practice, recommend a secondary authentication mechanism such as MFA. • Use selectively based on application and deployment.	High / Low
Zscaler Managed Dedicated IP	• Zscaler hosted IP addresses dedicated to the customer • Either Zscaler owned or Customer-owned IP addresses	• Simple to implement • Provides redundancy within a Datacenter and across Datacenters	High / High
Customer Managed Dedicated IP	• Leverages ZPA App Connectors to provide unique source IPs • Cutomer owned and assigned IP addresses • Destination applications/ FQDNs need to be identified as part of configuration	• Provides dedicated IPs in regions where Zscaler does not have a presence • Redundant IP addresses • Existing AppConnector deployment can be leveraged	Medium/ High

Leverage XFF headers	Same as above	Same as above	Medium/ High
Zscaler Private Infrastructure	• ZIA Service Edges require several public IP addresses. • Virtual infrastructure requires VMware ESXi or Azure cloud for hosting.	• Provides flexible deployment options with redundancy and addresses all use cases if deployment considerations can be met. • Can be used in cases where private IP space limitations do not allow forwarding traffic to Zscaler public cloud ZENs.	Medium/ High
Forward traffic (after ZIA) to a northbound web proxy	• Commonly used proxy-chaining feature. • Recommended for use selectively for applications such as socialware for http/s only traffic. • Can be implemented specifically for select applications and users using new forwarding policy options (by user, department, location, destination IP and application type).	Recommended for use only for subsection of traffic that is destined to applications that require source IP authorized access and can be routed via the Zscaler public cloud.	Medium / Medium

Next steps: SAML, MFA, and (eventually) refactoring

Enterprises that must employ source IP address as a mechanism for governing access to applications should add Zscaler as a layer in their security stack. Zscaler recommends four best practices for deploying ZIA/ZPA with source IP–address–based application–access controls:

- **Deploy enterprise SAML capabilities for SaaS access.** This is a valuable step toward migrating from device authentication to user validation.

- **Add MFA for all application access.** Coupled with SAML, MFA provides an essential security layer for the new way of work, enabling SaaS application access to users anywhere, anytime.

- **Bypass Zscaler if an application is not required to go through security assessment.** Example: Bloomberg terminals.

- **Establish a migration plan for applications that still require source IP address validation.** In the short term, deploy Zscaler and apply the solution considerations above. In the long term, consider refactoring or recoding to reduce dependence on source IP address–based access controls.

Due to legacy or compliance requirements, companies will continue to use source IP based access for application access in the foreseeable future. To help customers manage these deployments, Zscaler allows for multiple mechanisms, each with its own considerations and security implications.

Zscaler: the path from source IP to the cloud

Source IP address identification is no longer a reliable means to secure enterprise data traffic. The approach cannot scale, can be easily compromised, raises risk, extends vulnerable threat landscape, and does not protect the new way of enterprise work (cloud–first, remote, and device–agnostic).

But source IP address controls remain firmly entrenched in many organizations. For enterprises that must preserve some level of source IP address controls, Zscaler offers a compelling path to a better security model…one that's cloud–based, with an inline proxy, highly–distributed, near to every user, at the edge. More importantly, it's a security model that layers nicely — given solution considerations, of course — on an existing source IP–address–based application–access control environment. And that enables enterprises to accelerate migration to more secure means of authenticating users.

About Zscaler

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

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

zscaler™

Zero Trust Everywhere