Akamai

WHITE PAPER

A Blueprint for Zero Trust Architecture

Actionable Implementation Guide

by Akamai Chief Technology Officer Charlie Gero

Table of Contents

Who should read this guide?

Network architects, security engineers, CTOs, CISOs, and other IT and security decision-makers can all benefit.

For those responsible for scoping, configuring, deploying, implementing, and managing such a Zero Trust framework, the blueprint will provide a tiered methodology for adoption and rollout.

Leaders will find guidance and discussion points to help them better position Zero Trust internally and more efficiently realize a Zero Trust architecture.

A Brief History of Network Architecture

For all of the change that has happened over the past few years, there is one thing that has remained stubbornly constant: the basic hub-and-spoke network architecture that most companies utilize.

This architecture used to make sense. Long ago, before the internet was a bustling place of business and core infrastructure, companies placed their workloads in data centers. These data centers housed the critical infrastructure and applications. As branch offices, retail storefronts, and satellite locations came online, they too needed access to the centralized applications. Companies built out their networks to mirror that need, with all networking backhauling to their core data centers. After all, the data center was the central location where all the action occurred.

As time progressed, the internet began to emerge as a commercially viable disrupter. Naturally, businesses and carriers that had been in the practice of building complex private global networks serviced these private requests by doing what they knew best. They deployed these corporate and consumer services in the same data centers their internal applications were hosted in, and purchased internet links to provide a route to them. This fortuitously served a double purpose: Outside consumers could get in, but internal employees spread out across myriad branch offices could now get out. For the time being, hub-and-spoke was still the reigning champion of network architectures.

Over time, threat actors began to capitalize on this architecture, causing a whole new industry to be born: the data center security stack. Since the hub-and-spoke architecture funnels internet traffic at data centers, large powerful boxes began to be developed to protect those high-capacity lines. Firewalls, intrusion detection, and prevention ruled inbound traffic while secure web gateways (SWGs) enforced acceptable use in the outbound direction.

The proliferation of these security systems being deployed at centralized choke points further cemented hub-and-spoke as the dominant network architecture. For a time, the castle-and-moat approach to security seemed viable, and the notion of a network perimeter where everyone outside is bad and everyone inside is good remained dominant.

But the cloud changed all of this. Security measures must now meet users and applications where they are, outside of the perimeter.

Traditional Perimeter Security

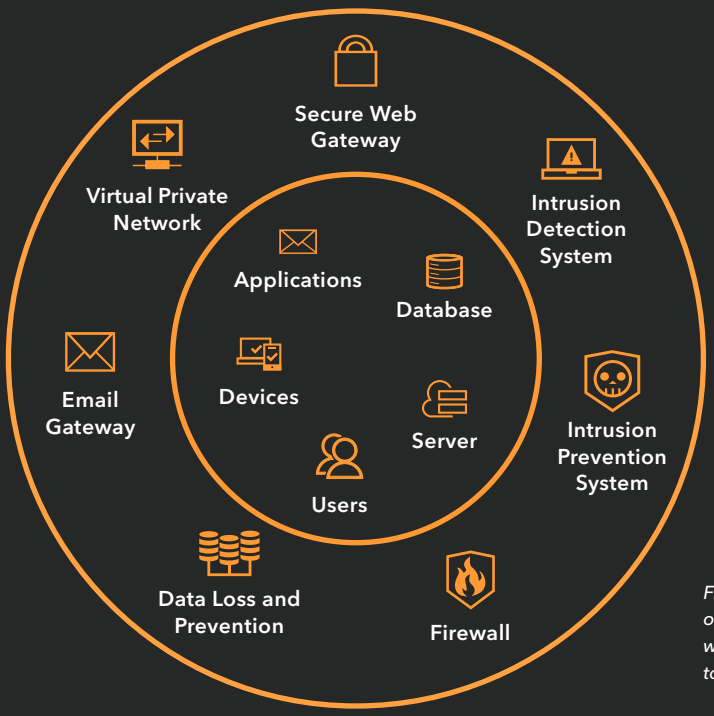

Fig. 1: With traditional perimeter security, one bad actor can exploit one perimeter weakness and, once inside, move laterally to access resources

The Rise of Cloud

Simply put, your applications are on the move. But they are not alone. Today's workforce and workplace have changed – when, how, and where people do their work have moved beyond the four walls of an office.

As a result, the network perimeter no longer exists. At least not in any recognizable form. Your users and applications are just as likely to be outside of the moat as they are inside. And with advanced persistent threats, you are highly likely to inadvertently let the malicious actors inside the perimeter with full access to your most valuable assets.

In today's modern world, utilizing a security and access approach that made sense 20 years ago is at best misaligned and at worst perilous. As Forrester Research says:

"The data economy renders today's network, perimeter-based security useless. As businesses monetize information and insights across a complex business ecosystem, the idea of a corporate perimeter becomes quaint – even dangerous."

– Forrester, Future-Proof Your Digital Business With Zero Trust Security

And this isn't just theory. This is evident in the massive amount of data breaches we've seen in the past five years, the vast majority of which happened as a result of trust being abused inside the network perimeter.

Further exacerbating this problem: Applications that were designed to live inside a network perimeter often have the worst security profiles. After all, if you were a developer 10 years ago and assumed that only authorized employees with good intentions could reach your system, would you have been as defensive as the coder today who knows vast armies of hackers will try to exploit his or her internet-based application?

So what are you to do?

A Zero Trust Security Architecture

John Kindervag, a thought leader in this space and a Forrester analyst at the time, proposed a solution that he termed "Zero Trust." The principle behind it is quite simple, but very powerful: Trust is not an attribute of location. You shouldn't trust something simply because it is behind your firewall. Instead, you should take a very pessimistic view on security where every machine, user, and server should be untrusted until proven otherwise.

The method of proof for this is strong authentication and authorization, and no data transfer should occur until trust has been established. In addition, analytics, filtering, and logging should be employed to verify correctness of behavior and to continually watch for signals of compromise.

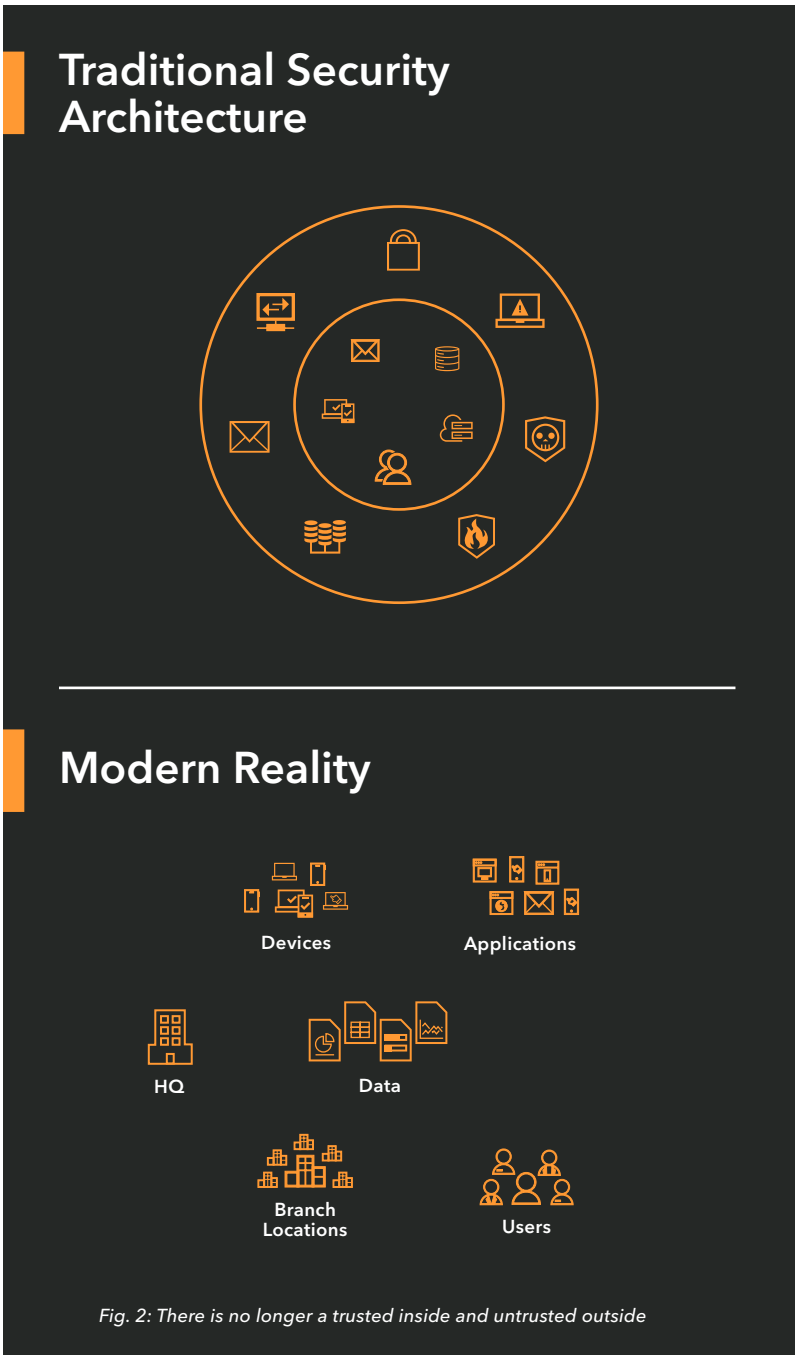

Fig. 2: There is no longer a trusted inside and untrusted outside

This fundamental shift defeats a vast amount of the compromises we have seen in the past decade. No longer can attackers spend time exploiting weaknesses in your perimeter, and then exploit your sensitive data and applications because they made it inside the moat. Now there is no moat. There are just applications and users, each of which must mutually authenticate and verify authorization before access can occur.

The Principles of Zero Trust

The network is always assumed to be hostile

External and internal threats exist on the network at all times

Network locality is not sufficient for deciding trust in a network

Every device, user, and network flow is authenticated and authorized

Policies must be dynamic and calculated from as many sources of data as possible

How does one accomplish this?

Akamai Edge Security Services

There are varying methodologies for delivering a Zero Trust architecture. One popular strategy is an access proxy architecture run entirely within your own DMZ. But this limits the ability of the cloud to absorb attacks, provide infinite bandwidth for caching, and autoscale resources as needed.

Akamai is a cloud-native company that has operated at the edge since our inception, more than 20 years ago. We have designed a Zero Trust Network Access (ZTNA) technology – an identity-aware proxy. But it lives in the cloud, scales on demand, executes CPU-heavy resources on our platform instead of your equipment, absorbs attacks, and delivers cached content closest to users with client and clientless methodologies – depending on the application type. We call this Enterprise Application Access, and it looks as follows:

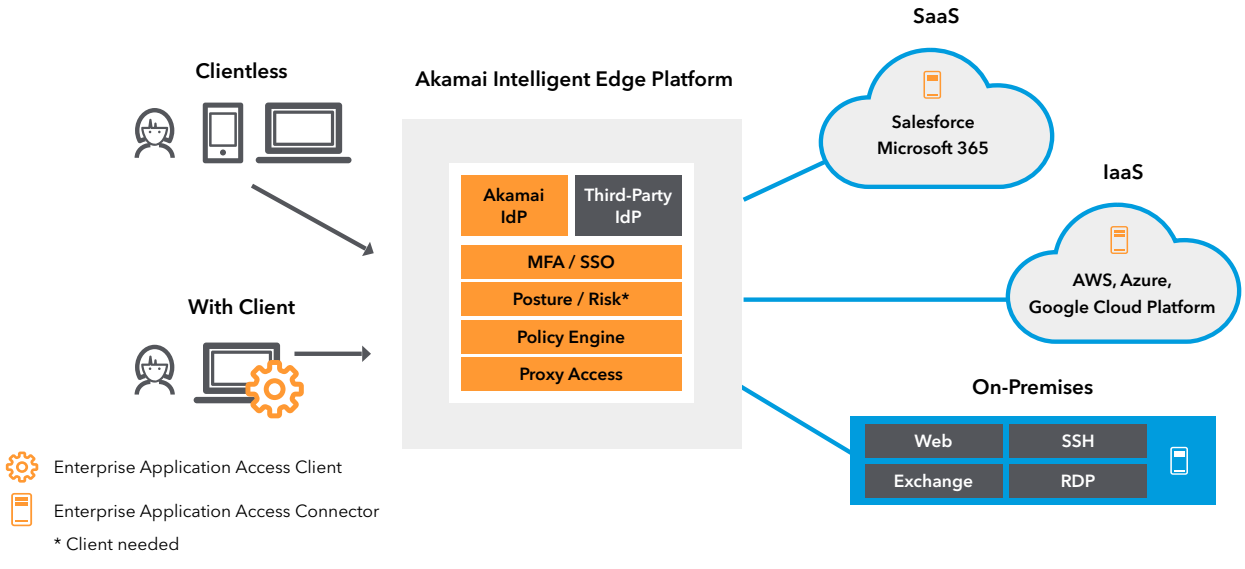

Fig. 3: Akamai's cloud-based ZTNA solution continuously and dynamically enforces access decisions

In this architecture, you provide access to your applications and not the entire corporate network. However, instead of placing your access proxy in the DMZ, you run a small virtual machine called an Akamai Enterprise Application Access Connector behind the firewall. It does not need to be, nor should it be, inside the DMZ. Its address should be on private IP space and not directly reachable from the internet. In fact, it should look exactly like any other application you would place behind the firewall.

When the Enterprise Application Access Connector starts up, it immediately establishes an encrypted connection to the Akamai platform. Once connected to Akamai, it downloads its configuration from Akamai servers and is ready to service connections.

When a user of internal applications attempts to access a service, they are directed to Akamai via a DNS CNAME and are connected to the Akamai Intelligent Edge Platform. Assuming your end user and their device pass all checks, they are then routed for authentication, multi-factor authentication (MFA), and single sign-on (SSO), and then device identity functions are performed. It is also at this Akamai point of presence (POP), as close to the user as possible, where Akamai can layer in additional services such as a web application firewall (WAF), bot detection, behavioral analytics, and caching. This is designed to provide best-in-class performance as well as the ability to keep potential threat actors as far away from your physical locations, applications, and data as possible.

After the user and machine are authorized, the connection from the client is stitched together with the outbound connection from the Enterprise Application Access Connector. Traffic from the user session flows through this stitched identity-aware proxy connection to the Enterprise Application Access Connector, which then connects to the requested application or service. At that point, a complete data path is established, and all access decisions are then continuously and dynamically enforced based on identity, device, and user context.

Capabilities: Zero Trust Network Access

Control network flows between all assets

Application access vs. network access

VPN elimination

Improved application performance

Grant identity and access to the cloud

Least-privilege user access to all applications (IaaS, SaaS, and on-premises)

Service insertion

Improved security posture against advanced threats

Authentication and authorization

Security at the edge

There are distinct and significant advantages to this method of access. The activities that are most performance- and security-sensitive take place at the edge, closest to the end user, where Akamai has more than 4,100 POPs spread around the globe. Additionally, the sensitive ingress path into the application happens over a reverse application tunnel, effectively removing the IP visibility of the perimeter and reducing the risk of volumetric attacks.

Application Access vs. Network Access

Readers might be inclined to think about this method of access as a VPN, but they would be doing themselves a great disservice. VPNs provide network-level access. Akamai Enterprise Application Access eliminates the need for a user to access the network to obtain access to an application.

If you opt for a simple VPN setup, you probably do what many companies do – you allow logged-in users to have IP-level access to your entire network. We know how dangerous this is. Why should call center employees have IP access to source code repositories? Or why should a contractor using your billing system have access to the credit card processing terminals? Access should be provided to only those things needed to perform a role.

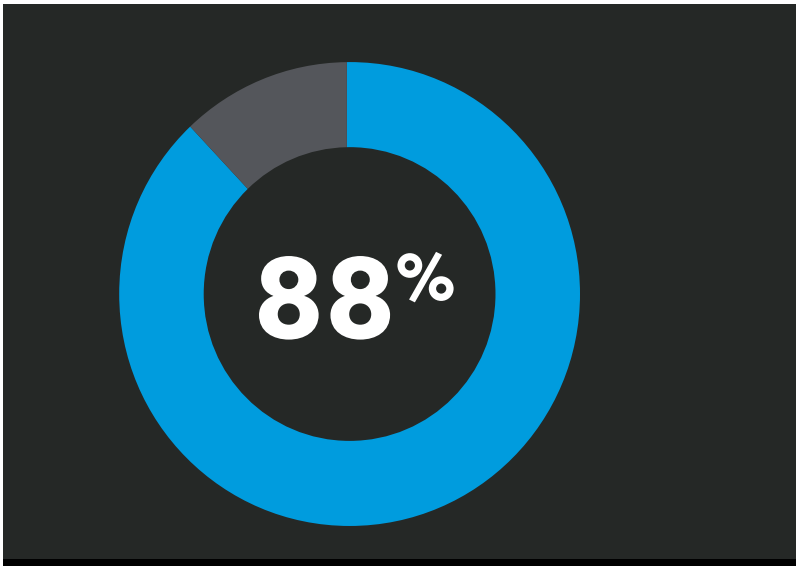

The percentage of CISOs who are not confident or are unsure of their employees' ability to apply cyber judgment

Gartner, Securing the Fully Remote Workforce

To fix this, you could begin to partition applications via VLANs onto separate segments behind a firewall and enforce archaic IP range-based rules for individual users or groups at the VPN aggregator. However, this is brittle and very prone to errors. Maybe someone is doing maintenance and moves machines to a new rack or needs to re-IP them to a new range. All of a sudden, users are locked out and support calls come rolling in. Or perhaps an application's architecture changes during a software upgrade and users are redirected to another machine as part of the workflow, but that machine is inaccessible to certain users or groups because the firewall rules were not updated. This architecture is exceedingly complex and requires all changes to have a very high degree of communication between application owners, network administrators, and security groups to ensure zero downtime.

Historically, we have significant evidence of what often happens when the above coordination fails. Administrators want to follow best practices, but in times of desperation, the dreaded IP ANY/ANY ALLOW rule gets added as a quick fix to allow affected users to access everything until the problem underneath can be diagnosed and repaired. But there often isn't time to go back and fix past holes. Again, to overcome the security downsides of unfettered horizontal access, significant complexity and operational overhead needs to be introduced when using a VPN, and that complexity often results in holes and quick fixes that worsen security posture over time.

The same trade-off happens regarding performance. In a VPN's simplest form, all traffic is directed back toward data center infrastructure. This can result in extremely slow access to internet properties and SaaS due to the hairpin effect, as well as increased congestion on internet uplinks within the data center for nonbusiness traffic such as Facebook and YouTube. Why backhaul normal internet access for a user who is already off-premises?

The Consequences of a Breach

$3.92 million
Average cost of a data breach to businesses

25,575
Average number of records in a breach

$150 globally / $242 U.S.
Average cost of each record

Source: IBM Security/Ponemon Institute, 2019 Cost of a Data Breach Study

"Network security architectures that place the enterprise data center at the center of connectivity requirements are an inhibitor to the dynamic access requirements of digital business."

– Gartner, The Future of Network Security Is in the Cloud

To overcome this performance burden, administrators often deploy split tunnels, again marking which IP ranges should travel down the VPN and which should egress directly to the internet. This is very effective and simple when you only have one internal perimeter. However, it begins to get much more complex as you add multiple data centers and virtual private clouds (VPCs) in IaaS providers. Administrators must then determine if they are going to install VPN aggregators in every data center and how to manage multipoint split tunnels effectively.

Again, all of these solutions are theoretically possible. You may already be using some combination of them. The problem is that the tasks required to do them well, maintain them, and provide proper security and performance over the lifetime of their implementations are often far too operationally complex to continually get correct. In many cases, companies convince themselves that because employees can access their applications, everything must be working optimally. They then find themselves caught off guard when one of the quick fixes above results in a catastrophic breach or a performance degradation that causes an outage or vastly decreased employee productivity.

This is not to say VPNs don't provide value. Far from it, in fact. Site-to-site access for multiple data center infrastructures is one case where they shine. It is simply to say that network-level access is not the correct paradigm to use when discussing users accessing applications. This is because network-level access enforces an unnatural compromise in simplicity vs. security and performance.

What makes proxy-based approaches like Akamai Enterprise Application Access so appealing is that they provide application-level access. With application-level access, performance and security are decoupled from complexity. This is inherently obvious in how easy it is to explain.

You simply take all applications that have locality with one another (all hosted in the same data center or same VPC, for instance), place them into a private network IP space or a restricted VLAN, place an access proxy in that micro-perimeter, and you're finished.

Application owners set their own security policies on the access proxy — who can access and why — and even more compelling, users can be anywhere. There is no distinction of on-premises vs. off, because there is no network perimeter that includes the end users. An employee working remotely or in a coffee shop is equal to an employee in an office. All that matters is whether the user is authorized and whether the machine is safe.

With application-level access, performance is best in class, despite the ease of deployment and use. Users simply utilize the internet to access applications directly, no matter where they are hosted or where the user is, allowing the internet to route packets to their destination without having to go through aggregators or intermediaries that aren't in path.

With application-level access, performance and security are decoupled from complexity. And an employee working remotely or in a coffee shop is equal to an employee in an office. All that matters is whether the user is authorized and whether the machine is safe.

In fact, with application-level access, internal networks often dissolve into simple guest Wi-Fi. Remember, for Zero Trust to truly be effective, you cannot treat internal users any differently from external. They are all untrusted by default.

Planning for Zero Trust

Revisit Key Goals

Know who your users are.

Don't trust any of them.

Protect your assets.

Authenticate/authorize all transactions.

Protect your users.

Prevent malware from infecting users.

Consider:

What challenges do you face?

What are key success factors?

What do you hope to solve?

Desired End State

Application Access

In a modern Zero Trust deployment, all applications should be segmented into their own micro-perimeters based on location and purpose. These micro-perimeters should be on private VLANs with private IP space, and directly inaccessible from anything other than the access proxy that fronts them.

We leave it to the discretion of each company to determine if it wishes to provide even further micro-segmentation within the micro-perimeters to stop horizontal escalation and movement if an application is compromised. We see theoretical value in such exercises, but in practice find that the level of complexity in getting such segments right is often far too high to justify their added security. Anyone who has tried to unwind the touchpoints that a complex application has behind the perimeter can testify as to how brittle it can be to segment applications from one another. However, if your environment allows for this in an easy and maintainable way, we encourage its adoption.

All users, whether they are on- or off-premises, should be required to access all applications through identity-aware access proxies, such as Akamai Enterprise Application Access – which perform not only standard authentication, but MFA as well. Additionally, there should be robust device posture capabilities that obtain device criteria to procure access to certain applications.

In addition, we strongly believe that Zero Trust does not end with authentication and authorization. A modern stack will also require some level of inspection and analytics of the traffic passing through the access proxies. This will help shield against advanced persistent threats and willfully malicious end users.

> *Endpoints aren't 'yours' anymore. BYOD and the increasingly mobile workforce have eliminated the control IT used to have over endpoints that attach to enterprise networks and access data.*
>
> *– Forrester, A Practical Guide to a Zero Trust Implementation*

One crucial security system that should be layered on top of your access proxies is a WAF, to ensure that application-level attacks are not being launched (intentionally or inadvertently) from your end users toward your internal applications. Other advanced systems such as human/bot detection can be leveraged for non-API sites to help ensure malware is not masquerading behind valid endpoints.

As you bring your applications online and make them accessible through access proxies, DDoS prevention becomes even more important. You should align yourself with providers who can absorb attacks against your micro-perimeters and access proxies, allowing continued operation under intense loads.

And finally, to ensure performance is best in class for your applications and that users not only accept this shift in access, but champion it, your access proxies should be fronted by networks that can provide performance benefits. Specifically, CDNs and internet routing overlays should be part of your arsenal to not only make access available, but make it more performant than prior methodologies have ever allowed.

Threat Protection

Solutions like Akamai Enterprise Application Access can protect your applications from malicious actors. But what about protecting users from inadvertently becoming those very actors through compromise – a device infected by malware or credentials stolen via a phishing link and landing page? This is where prevention and detection become crucial for web traffic.

One approach is to deploy a cloud-based secure web gateway solution such as Akamai Enterprise Threat Protector. This inspects every web request that users make and applies real-time threat intelligence and advanced malware analysis techniques to ensure that only safe content is delivered. Malicious requests and content are proactively blocked. And to provide consistent protection wherever your users are connecting from, you can install a lightweight agent on laptops.

Architectural Visualization

Putting it all together, we expect your users and applications in a fully realized Zero Trust world to look something like the following:

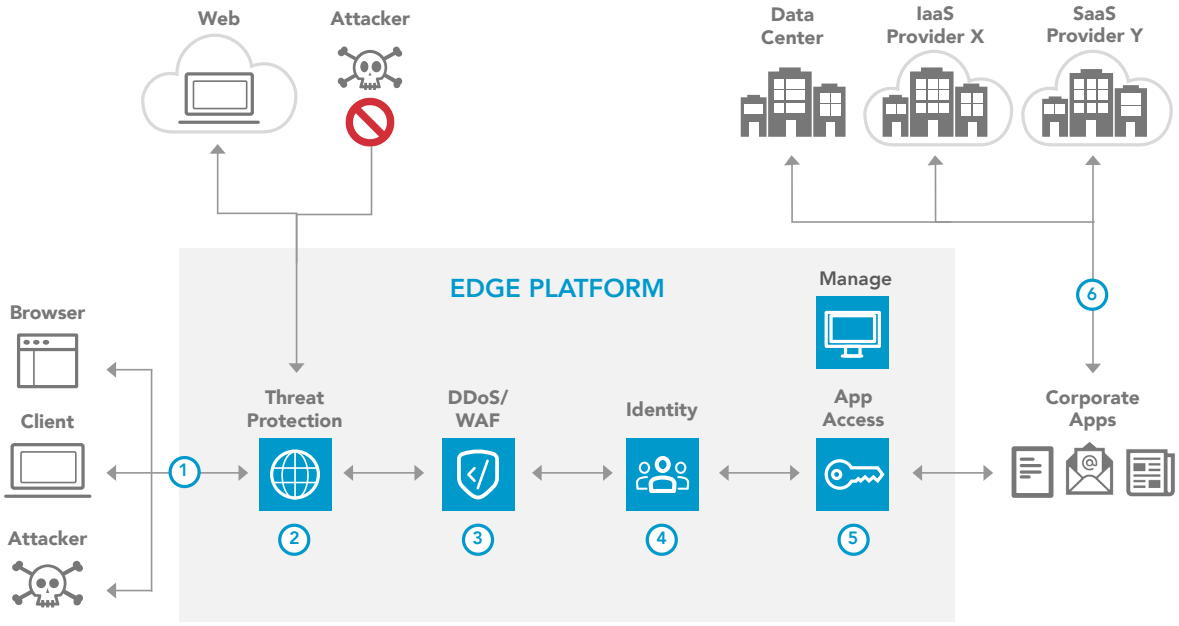

Fig. 4: A Zero Trust architecture dictates that only authenticated and authorized users and devices can access applications and data. At the same time, it protects those applications and users from advanced threats on the internet.

Getting From A to B

We have taken a pretty comprehensive look at Zero Trust thus far, including how a complete Zero Trust-enabled enterprise should look. However, much like security is best realized when it is easy, deployment of new architectures must also be simple and phase-able. No company of any significant size would be able to convert its entire infrastructure overnight to this vision. As Zero Trust itself is a strategy, so too is its deployment.

We believe the best approach to reaching a complete Zero Trust architecture is to begin transitioning applications to a Zero Trust Access model in small, manageable batches. The batch size will vary by the amount of resources you can dedicate, with some companies only able to do one at a time. Regardless, each application being transitioned will go through several stages along its path to complete Zero Trust. At each stage, you will verify that the application is functioning correctly and that authorization is being enforced as expected.

Enterprise Application Access supports HTTP(S), VNC, RDP, and SSH applications without the need for a client to be installed. Device posture assessment and additional protocols through the use of a client are also available. Access decisions can be made based on third-party threat signal intelligence, such as whether a device has been compromised or is compliant with endpoint security policies.

Bring Your Zero Trust Strategy And Roadmap Right To The Board

To do this successfully, you must:

Be clear that Zero Trust is what will ultimately get you customer trust. Zero Trust is an architectural concept designed to ultimately protect your most valuable asset, hence protecting your employees, customers, and society.

Translate technology needs to business benefits. Demonstrate how your Zero Trust initiative enables business initiatives like digital transformation, cloud migration, and enabling a remote workforce.

– Forrester, A Practical Guide to a Zero Trust Implementation

Pre-Staging Assumptions

The following assumptions are made with respect to the pre-staging state of your network:

- Applications and users within your current perimeter may reach one another directly over IP
- You have installed an Akamai Enterprise Application Access Connector into your current perimeter
- You have created a walled-off VLAN for eventual placement of corporate applications
- You have installed an Akamai Enterprise Application Access Connector into the walled-off VLAN
- Your off-premises enterprise users will continue to have a VPN installed during the phasing in order to reach those applications not yet transitioned

With these prerequisites met, we now can begin discussing the stages that each application will journey through on its way toward a complete transition. These stages are meant to reflect the highest level of granularity, and as such, you may wish to combine some of them as you become more accustomed to the process.

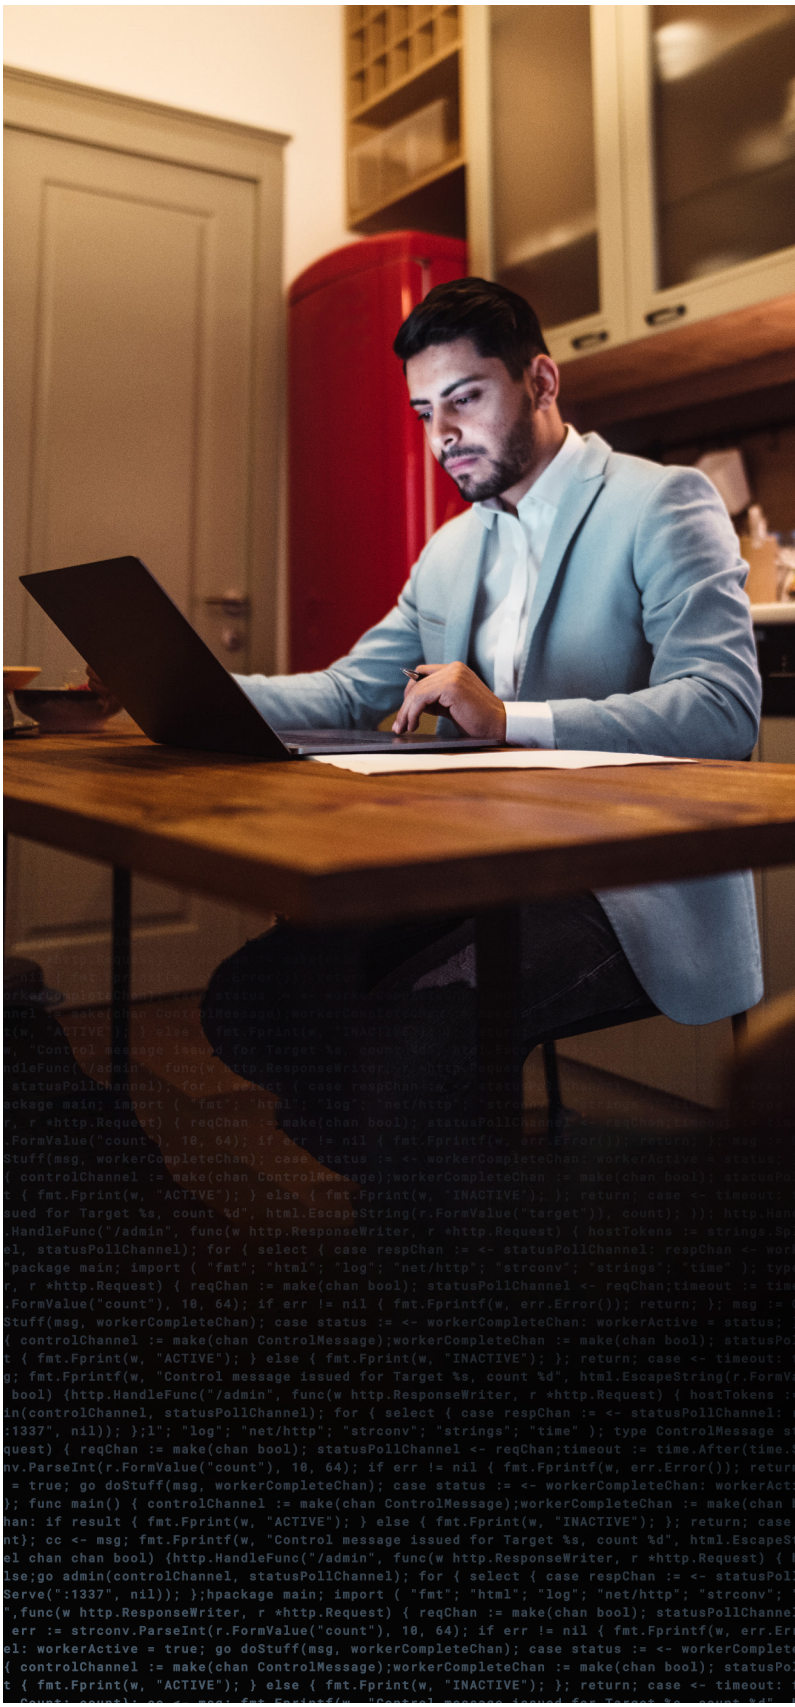

User Grouping

Part of our application staging process will involve slowly phasing this feature out to select groups of users.

We propose having three general sets of users defined for each application:

Test Lab

These users will be responsible for verifying the functional integrity of the application when it is first made accessible through Enterprise Application Access. They should be familiar with the application's operational semantics enough to test both standard behavior as well as edge cases.

As certain additional security features and performance features are layered on, these users should also have the background to test that security is behaving as expected and performance gains are truly being realized.

We encourage organizations that may have a large number of applications to consider investing in tools that will allow standard non-application–specific features, such as generic performance testing and certain security checks to be performed regularly and automatically without the need for heavy human oversight. Specific security checks that should be performed with such tools are SQL injection, cross-site scripting, bot detection, and command injection.

External Users

This will be the first set of production users that Enterprise Application Access is rolled out to and consists of all valid users who will be accessing the application outside of your network perimeter. This group will operate in a dynamic fashion. In other words, as users travel off-premises, they will dynamically enter this group, and as they travel back within the network perimeter, they will dynamically leave.

Internal Users

This will be the final set of production users that Enterprise Application Access is rolled out to, completing the migration to all users of the application. This set consists of all valid users within the network perimeter who will be accessing this application. Like the External Users group above, membership in this group is dynamic. Users will join this group automatically upon entering your internal network perimeter and will join the External Users group upon leaving.

User Grouping Methodology

In our phasing example, we are going to be utilizing Akamai Enterprise Application Access product as our access proxy. Akamai onboards applications into its distributed network through the use of DNS: Application hostnames are CNAME'd to Akamai to route users to our platform.

Therefore, users who do hostname lookups that result in the CNAME chain to Akamai will be on-ramped to Akamai's global network. Users whose lookups result in the original internal A record being returned will continue to access the application using the outdated perimeter method. We will take advantage of this architecture as a way to phase the different user groups onto Enterprise Application Access for each application by controlling which groups follow the CNAME chain and which users receive the internal A record.

We will achieve this through the use of DNS views, also known as split-horizon DNS, to define the above three sets of users. A DNS view is a method by which two or more users querying the exact same hostname can receive completely different answers depending on what their source IPs are. For example, a user in China who queries www.foo.bar could receive one distinct IP, while a user in the United Kingdom could receive a completely different IP, even though they queried the exact same hostname.

We will utilize this method by organizing our different user groups by their source CIDR blocks. All Test Lab users will reside in one CIDR block, Internal Users will be within the CIDR block that defines your entire internal network, and External Users will be sourced from all IPs that don't match the first two sets.

All common DNS servers that enterprises deploy in modern environments today support this feature. Now that we understand how to partition the various user groups, we can begin discussing the actual application rollout stages.

This user grouping methodology is supported by all common DNS servers deployed in modern environments by today's enterprises.

Application Rollout Stages

The following pages describe the eight stages of a successful application rollout, along with the factors you need to consider at each stage.

If at any point you have questions about your particular application requirements or migration methodology, please reach out to your Akamai Account Executive or contact our security experts here.

1. Application Precheck Stage

In this stage, you check to make sure that the application meets the requirements of the access proxy you have deployed. In the case of Akamai Enterprise Application Access, you will make sure that the application is a supported protocol.

Clientless protocols include HTTP, HTTP(S), RDP, or SSH. TCP and UDP applications are supported with a client.

Consider:

Which applications present the biggest pain point for remote access?

What group(s) of users use this application?

Where are the applications hosted?

What virtual environment is used?

What would be a successful outcome during a proof of concept (POC)?

2. Access Proxy Preparation Stage

Next, you will configure your access proxy to be aware of the application as well as its specific security and access rights. In the case of Akamai Enterprise Application Access, you will configure this in the cloud.

The configuration will be pushed to all of your Akamai Enterprise Application Access Connectors instantly, as well as to the entire Akamai platform, such that all Akamai POPs are capable of servicing your end users.

Consider:

Are there any outbound proxies?

What directory sources do you use for user authentication?

Do you have MFA associated with your applications?

Are you looking to extend SSO for your on-premises, cloud, and public applications?

3. Test Lab Enrollment Stage

Now that the access proxy is aware of the application in question, you can begin onboarding users. In this stage, assuming you are utilizing the Akamai Enterprise Application Access proxy, you will modify your DNS views as noted above such that members of the Test Lab group get CNAME'd to Akamai for the specific application hostname upon lookup.

Immediately following this action, members of the Test Lab will be directed into Akamai's global network whenever they attempt to access the given application.

During this time, Test Lab members should verify authentication is working correctly, MFA is appropriately configured, and SSO works across all other previously onboarded applications. More important, in-depth tests around the application's functional correctness should be run at this time.

Consider:

How are users accessing these applications today?

Which group(s) of users would be testing?

How will success be determined?

How long will users be given to test?

4. Security Upgrade Stage

Now that members of the Test Lab are safely able to access the application using Enterprise Application Access, you should consider enabling advanced security features that otherwise were impossible in the traditional perimeter model. In this section, we are going to reference specific Akamai security products that work well with, and are enabled by, the Akamai Enterprise Application Access proxy. However, whether you use these specific products or not, the general recommendations and deployment stage remain valid.

We recommend a minimum of enabling the Akamai Kona Site Defender product in order to get WAF functionality. Once this is engaged, members of the Test Lab should verify that SQL injection, cross-site scripting, and command injection attacks against the application are rejected by the WAF platform.

It is in this stage that application owners should also consider bot vs. human detection as a simple yet incredibly powerful addition to security. If the application in question is not an API server and should never be accessed programmatically, you can enable this feature through the Akamai Bot Manager product. You can instruct the platform to reject connections that cannot be determined – with a high degree of certainty through advanced analytics – to have originated from a human. This goes a long way toward shutting down malware and other advanced persistent threats that masquerade behind valid user sessions.

This stage is also where you will determine if the given application should be restricted to managed devices only. If so, and you deploy certificates to your end devices, you can upload your public Certificate Authority certificate to Akamai Enterprise Application Access so that it can reject all connections originating from machines that are unmanaged.

Finally, geography and IP-based restrictions can be enabled as well. If you know of specific CIDR-based allowlists or denylists that should be used, or of geographic regions that should be denied access, this can be enabled and tested at this time.

Now that you've secured the applications you own, we suggest considering applications you don't own, specifically those applications that are out on the internet. How are you securing the traffic that is accessing such web-based corporate applications as G Suite or Dropbox? We recommend enabling Akamai Enterprise Threat Protector, a cloud-based secure web gateway that proxies all internet traffic for analysis and inspection.

Regardless of the features you enable, it is during this stage that your Test Lab will be expected to ensure that the security options are not only working, but also are not inhibiting the functional correctness of the application.

Consider:

Are any geo-restrictions required?

How is DNS provided to your users?

What type of mobile device management (MDM) solution is in place?

5. Performance Upgrade Stage

With the application onboarded and security now implemented, the next thing to consider is if performance is deficient. In traditional perimeter-based access and security, enterprises are often limited in performance by data center locations and the enterprise's associated links between branch locations.

Sometimes features such as on-premises caching can be deployed to mitigate these issues, but they fall far short when employees go off-premises, as caching elicits the greatest performance upgrades when it is closest to the end user.

In this stage, you will assess if your applications need a performance upgrade, and if so, utilize caching and other techniques available to you. If you are using Akamai Enterprise Application Access, one obvious upgrade would be to enable Akamai's CDN through our Ion product. This will enable caching worldwide across Akamai's quarter million servers, delivering best-in-class performance to your end users, regardless of where they are located. Additionally, by enabling this, you will gain access to Akamai's network, where forward error correction (FEC), route optimization, and packet replication provides near-zero packet loss and performance-based routing.

It should be noted that this stage may optionally be delayed until the below External User Enrollment Stage occurs, as External Users may be best positioned to assess the relative performance gains while outside of the network perimeter.

Either way, we recommend performance instrumentation at this phase, as well as prior to application enrollment, in order to accurately assess the performance gains. This may be through the use of browser plug-ins, waterfall graphs, or third-party performance monitoring services.

Consider:

Where are your users located?

Where are the applications located?

6. External User Enrollment Stage

By the time you have reached this stage, the application has been onboarded and made accessible via Akamai Enterprise Application Access, had significant security upgrades applied, optionally been made much more performant, and been tested by an internal Test Lab to gain confidence that it is functioning as expected.

It's now time to roll it out to a wider audience. At this point, we recommend phasing this to External Users. They are the ones for whom this style of access will eventually lead to VPN removal. Additionally, they are the users most often affected by performance issues and are in the most hostile environments, where their very location puts your applications and data at risk. Utilizing better performance and security with this set of users is where the greatest benefits of Akamai Enterprise Application Access will be immediately visible.

We recommend that, prior to entering this stage via DNS view modification, notification is issued to these users so that they are not caught off guard by the transition. If all has been configured appropriately up to this point, the transition should effectively be nearly invisible to them, with the exception of increased performance. However, advanced notice will alert users to keep a particularly close eye on functional correctness such that if anything is amiss, they will know why and be prepared to contact the appropriate administrator.

Consider:

Are there third parties that currently access your network?

How is remote access provided?

What is challenging about this existing model?

Which remote user group has the least need but the most access?

7. Internal User Enrollment Stage

In this stage, External Users have been enrolled for some time and all bugs in operation have been addressed. At this point, you can remove all references to the application from your DNS view-specific configurations and add it as a CNAME entry to the common view. All users should then immediately begin accessing this application using your Akamai Enterprise Application Access proxy.

All procedures that were used to notify External Users should be applied to Internal Users in this stage as well. It is hoped that through the prior six stages, any errors or misconfigurations will have been discovered and remedied. Assuming this is correct, it is at this point that the application will have successfully been operationally transitioned to Akamai Enterprise Application Access, and all users should be enjoying the benefits of easier, faster, and safer access.

Consider:

What other groups need access to this application?

How is remote access provided today?

What is challenging about this existing model?

Do groups outside of these particular users have acccess to this application?

8. VLAN Migration Stage

After an appropriate amount of time has passed in the Internal User Enrollment Stage, you can finally shift the application into the walled-off VLAN. Before this occurs, even though all valid users are being directed through Akamai Enterprise Application Access using DNS, the application server itself is still reachable directly through its IP address and thus vulnerable to malware within your network perimeter.

This final stage removes all direct IP access, effectively walling the application off from anything but the access proxy itself.

Once this is complete, the application is officially and completely transitioned to Akamai Enterprise Application Access.

Consider:

How are network controls put into place?

What type of network-level access exists for these applications today?

Can these applications be segmented?

Who are the network stakeholders?

Post-Staging Operations

Once all applications have been transitioned to Enterprise Application Access, you can begin investigating the complete removal of VPN clients from end-user systems.

In addition, you may now consider taking the action of converting your internal user network into guest Wi-Fi, as all application access is now transitioned to a Zero Trust Access model.

Finally, it is in this phase that you should consider protecting the connectivity back to your data center from advanced DDoS attacks. The Akamai Prolexic product is one such solution that can assist here.

Summary and Next Steps

Traditional hub-and-spoke networking architectures, along with the castle-and-moat security perimeter they utilize, simply cannot effectively provide performance or security in today's cloud-and-mobile world. This is a problem all companies must begin facing, lest they be left behind in a vulnerable state. Failure to transition to safer enterprise security architectures is the number one cause of corporate breaches today, and it's only going to get worse. Simply put, you are not safe behind the perimeter, because the perimeter itself no longer exists.

How do you start transitioning to a Zero Trust architecture?

Akamai's cloud security services can be combined to build a comprehensive Zero Trust architecture, not only enabling safe application access in a cloud-native world, but leveraging the cloud to remove the need for internal corporate networks almost completely.

By utilizing our advanced distributed ZTNA solution, along with the power of the global Akamai Intelligent Edge Platform, you can finally move to a perimeterless world in an incredibly easy way, phasing in applications, reducing your migration risk profile to near zero, and leveraging Akamai's extensive 22-year history of proven performance and security solutions.

As you continue on your Zero Trust architectural journey, you can rest assured that Akamai will be there with you at each step, helping you to transform your network to an architecture that not only provides access to your applications and data, but does so in an easy-to-manage way while maintaining the highest levels of security and performance.

Akamai's security experts can work with you to develop a phased, customized path to Zero Trust. Schedule a 30-minute workshop to discuss your current state, desired end state, business priorities, vulnerabilities, and top concerns. Or visit akamai.com/zerotrust for more information.

About the Author

Charlie Gero, Chief Technology Officer, Enterprise and Advanced Projects Group, Akamai

Charlie is co-founder of Akamai Labs and heads research and development efforts in the enterprise and cloud networking spaces, with a special focus on performance, access, and security.

Akamai secures and delivers digital experiences for the world's largest companies. Akamai's intelligent edge platform surrounds everything, from the enterprise to the cloud, so customers and their businesses can be fast, smart, and secure. Top brands globally rely on Akamai to help them realize competitive advantage through agile solutions that extend the power of their multi-cloud architectures. Akamai keeps decisions, apps, and experiences closer to users than anyone – and attacks and threats far away. Akamai's portfolio of edge security, web and mobile performance, enterprise access, and video delivery solutions is supported by unmatched customer service, analytics, and 24/7/365 monitoring. To learn why the world's top brands trust Akamai, visit www.akamai.com, blogs.akamai.com, or @Akamai on Twitter. You can find our global contact information at www.akamai.com/locations. Published 02/21.