appgate

ZERO TRUST NETWORK ACCESS

EVERYTHING YOU NEED TO KNOW

ABSTRACT

Zero Trust Network Access (ZTNA) has rapidly become the de facto enterprise standard for secure access control, especially as cloud, remote work, and edge computing trends turn the traditional security perimeter inside out. As the threat environment heats up at the same time, traditional network access control, based on a “connect first, authenticate second” approach, exhibits too many weaknesses to remain in use. The Zero Trust approach starts from a default deny posture and then extends limited, earned trust, which is continuously reevaluated. From this basis, ZTNA enables operational efficiencies with fewer tradeoffs between security, convenience, and agility. This eBook provides a high-level overview of ZTNA. It is intended to give the reader an understanding of the architectural choices available for ZTNA, along with all of the elements one must consider when selecting and implementing a ZTNA solution.

TABLE OF CONTENTS

Introduction

The traditional security perimeter has eroded as cloud, remote work, and edge computing trends accelerate. As the threat environment intensifies in parallel, traditional network security controls are continually the source of major security breaches and are therefore being rendered deficient. The "connect first, authenticate second" approach can no longer defend digital assets from malicious actors. Zero Trust Network Access (ZTNA) offers a solution and has become the standard for secure enterprise access control.

A Zero Trust architecture hinges on securing network access.

ZTNA starts with a "default deny" approach to access using identity and risk factors to establish trust, which limits the privileges of users and devices even after they're authenticated. It involves implementing a network architecture that uses dynamic, context-sensitive policies for stronger access controls without impeding agility and experience.

This eBook provides a high-level overview of ZTNA including:

- The available architectural models
- Best practices for implementation
- Top considerations as organizations begin their Zero Trust journey

Implementing Zero Trust is not a "one and done" project. It is an ongoing journey toward strong, adaptive, and risk-based access controls embedded into the fabric of distributed, agile, and hybrid IT.

A Brief Overview of ZTNA

Zero Trust access is based on the fundamental principle that no user, human, or machine should be automatically granted access to *anything*. It is the ultimate extension of the "Principle of Least Privilege" approach to security. With ZTNA, a user is denied access to networks and digital assets by default. Then, they are only permitted access after their identity (user + device + context) is extensively authenticated. Dynamic policies and entitlements are then granted to the identity, provisioning limited access to authorized resources. These surgical entitlements are conditional and based on context and risk tolerance defined by the enterprise.

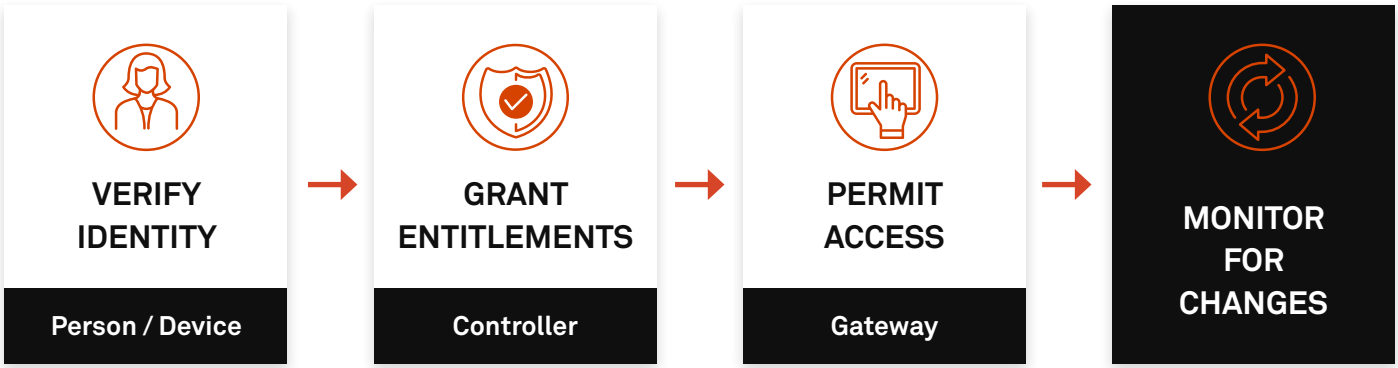

Authenticating the user's identity and access authorization is a multidimensional process. As depicted in the figure above, ZTNA starts by verifying the identity of the **user/device** to determine the right entitlements. Access is only granted to approved resources based on the context the user presents when they are connecting. In this way, the **Controller** is acting as a Zero Trust Policy Decision Point (PDP) and the **Gateway** as a Policy Enforcement Point (PEP). It's infinitely more secure than using an IP address and username/password combo because the theft of basic credentials, IP spoofing, and brute force attacks have made these traditional authentication methods vulnerable. ZTNA is a more dynamic solution that takes into account contextual factors.

After the user has been given access, ZTNA continues to monitor to determine if access privileges should be adjusted or entirely revoked. It continuously evaluates the user and device in context, including the user's role, device security posture, location, time and date, and a range of other conditional requirements. This makes it possible to immediately interrupt suspicious behavior before it causes harm.

In addition to improved secure access, another critical aspect of ZTNA is its ability to cloak the entire infrastructure and minimize the attack surface. This means that all resources secured with ZTNA are 100% invisible to malicious actors and only visible to authenticated and authorized trusted users, providing yet another layer of security.

The Zero Trust approach starts from a default deny posture and then extends limited, earned trust, which is continuously reevaluated.

ZTNA permits access only after considering the complete context. The platform continues to monitor the environment for changes after access is granted and reevaluates a user's conditions and privileges to ensure real-time protection.

The ZTNA model was originally known as the Software-Defined Perimeter (SDP) and the names are used interchangeably. By using these security architectures, enterprise organizations can modernize network security and:

- Strengthen and simplify access controls
- Reduce the attack surface
- Remove policy management complexity for admins
- Improve the end-user experience
- Unleash operations with integrations and automation

DRIVERS OF ZTNA ADOPTION

The evolving nature of the enterprise perimeter is the main factor driving ZTNA adoption. It's where the tailwinds of digital transformation efforts meet the headwinds of failing legacy network security solutions.

The most common initiatives driving ZTNA adoption include:

- Secure remote access and workforce enablement
- Privileged user and third-party risk reduction
- Secure multicloud/hybrid IT access and DevOps
- Café-style networking and overall network transformation

The days of siloed, traditional network security solutions are over. These outdated solutions no longer support the security and agility requirements of digital businesses. The enterprise perimeter has been turned inside out, driving demand for an access solution capable of protecting all resources in a lightweight and flexible manner. That's what ZTNA provides.

ZTNA vs. Other Access Control Paradigms

ZTNA is different from the legacy access control paradigms, such as Virtual Private Networks (VPNs), Network Access Control (NAC) solutions, trusted Local Area Networks (LANs), firewalls, and so forth. However, adopting ZTNA does not mean completely ripping and replacing these technologies. Remember, Zero Trust is a journey and in some cases, ZTNA works in conjunction with these other solutions. Using a layered strategy can make it easier for an organization to transition from a familiar approach, like VPN, to a complete end-to-end ZTNA environment.

The most basic difference between ZTNA and VPNs, NACs, and trusted LAN has to do with "default allow" versus "default deny." At the risk of overgeneralizing, most VPNs, NACs, and trusted LAN solutions allow access by default to any user or device which meets the requirements of a statically defined policy. If the user conforms to the policy, such as having a specific IP address or a username/password pair, they are let onto the network and given broad access.

From there, the user's access to digital assets is limited only by other controls, like application logins. This is a less than optimal approach. Malicious actors can impersonate real users and abuse network access to move laterally within networks. The ZTNA solution's default-deny posture mitigates this risk.

In terms of other security solutions, ZTNA is not a complete replacement for firewalls, including Next-Generation Firewalls (NGFWs) and Web Application Firewalls (WAFs). Instead, ZTNA works alongside firewalls, enabling them to become more agile while simplifying rules and policies. Though firewalls remain a critical network security solution, they still only rely on a simple and insecure IP-centric method for assuming trust. Another challenge with using firewalls as the single barrier to control network access is that it is relatively difficult to make changes to access rules. It can take a long time to process firewall changes, and the revocation process may be neglected as resources migrate or users change roles, and so forth. ZTNA eliminates the need to make constant access changes to firewalls by adding another layer of robust, identity-centric policies that are more easily managed.

Implementing ZTNA won't disrupt your environment. It can work alongside your existing security architecture to improve your overall security posture.

ZTNA and SASE

ZTNA is a working component of the SASE (Secure Access Service Edge) framework, which can make things a little confusing.

- **SASE** comprises inbound (ingress) and outbound (egress) access.
- **ZTNA** is how organizations approach ingress SASE. ZTNA does not apply to egress SASE requests for access to external internet facing resources.

In the SASE model, a ZTNA solution handles the private access requests (ingress) and works alongside a Cloud Access Service Broker (CASB) or a Secure Web Gateway (SWG) to handle internet bound traffic (egress).

ZTNA Implementation

ARCHITECTURAL APPROACHES

There are two architectural approaches to ZTNA, client-based and browser-based. Both approaches offer benefits over traditional access solutions, but the differences in their architectural makeup are worth considering when determining the right ZTNA solution for your organization's current and future requirements.

CLIENT-BASED:

This architecture is similar to the Cloud Security Alliance (CSA) Software-Defined Perimeter architecture. In this setup, a client is installed on a server or users' device and the endpoint initiates the connection process.

- **Pro:** The advantage of this approach is the ability to apply Zero Trust for all private resources including custom or legacy applications, as well as enhance device posture checking as criteria for trusted access.
- **Con:** There can be situations where access from an unmanaged device is required to connect and a client can't be installed. For example, this can occur when looking to set up third-party or contractor access.

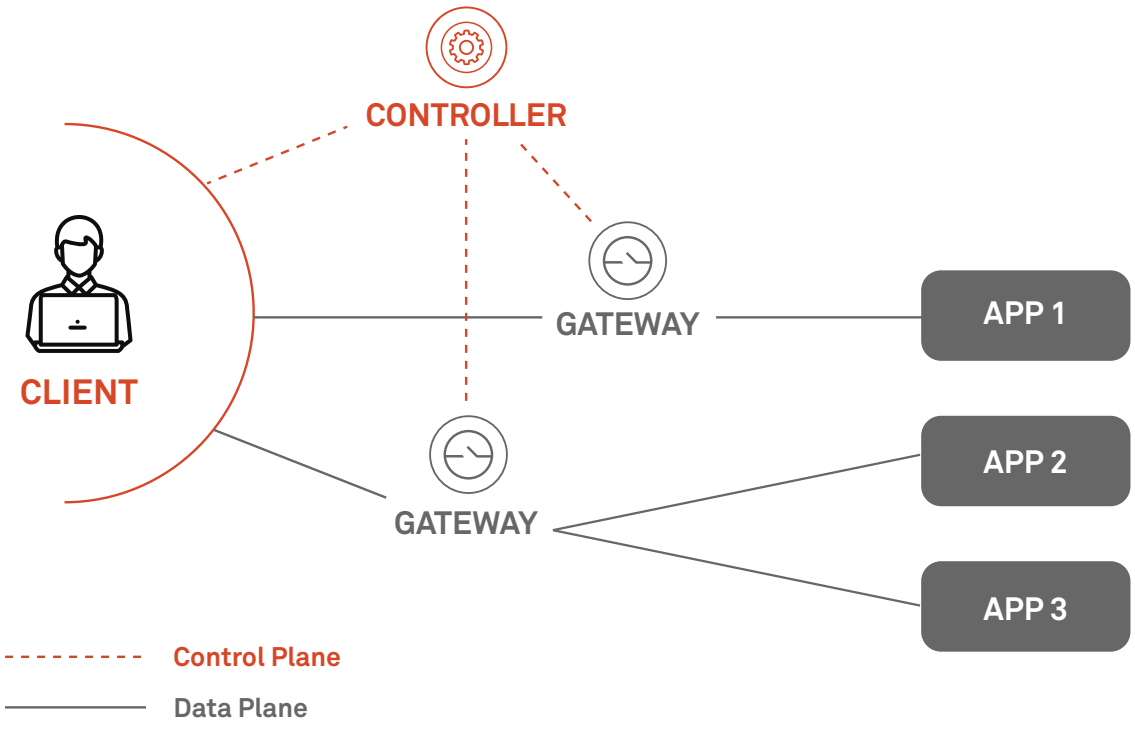

BROWSER-BASED:

This architecture is similar to the Google BeyondCorp vision, which is essentially a "clientless" or browser-based approach. In this instance, the user connects to the web application via a common internet browser.

- **Pro:** This provides frictionless web application access, which is a viable option for unmanaged devices where the installation of a client isn't allowed or possible. This is best used for smaller, less complex, web-only applications or a one-off use case.
- **Con:** This architecture only supports web applications and web protocols, which obviously has significant limitations. Therefore, browser-based is not a viable solution for a full Zero Trust deployment, which requires integration with non-web-based applications.

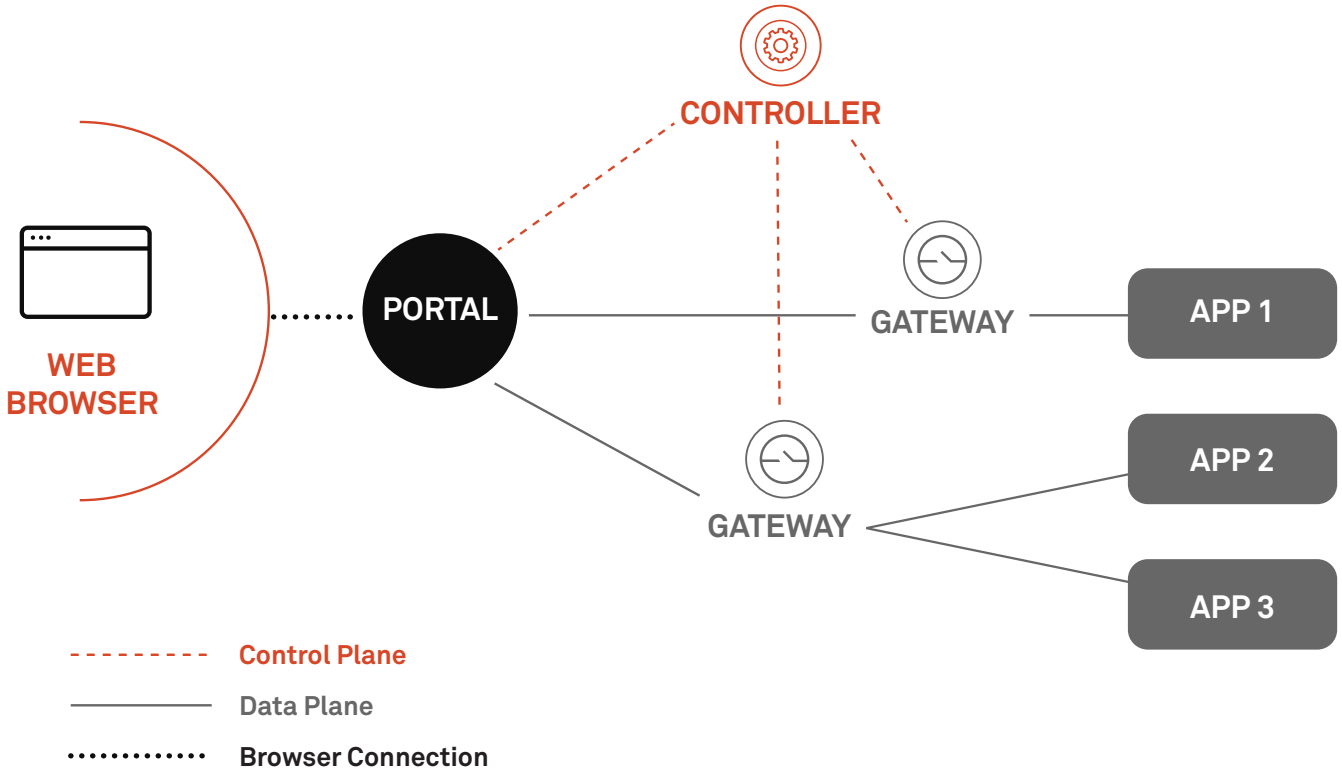

HYBRID APPROACH:

There are a few/select ZTNA providers that offer a hybrid model, which allows an enterprise to blend the advantages of both approaches depending on the use case. For example, this might include selecting clientless for third-party access to a web-based application and then using the client for critical and legacy applications.

KEYS TO ZTNA SUCCESS

Before embarking on your Zero Trust journey and selecting a ZTNA solution, it's important to consider your current and future-state requirements. No two organizations are alike, so look for the solution provider that offers robust feature sets and the flexibility required to meet all your long-term enterprise needs.

QUESTIONS TO SELECT THE RIGHT ZTNA SOLUTION

How will multiple, disparate identity providers be managed?	Consolidating identity providers is the goal of any large enterprise, but it is a very complex and challenging project. The reality is that most enterprises today are dealing with multiple identity stores. These identity providers might contain different users, hosted in different locations, that support different technologies. The ZTNA solution has to be able to work with all the disparate heterogeneous identity providers and their directories to reduce complexity and provide users a seamless experience.
What is your Zero Trust roadmap?	Another foundational question to ask when embarking on a ZTNA journey is "What do we want to protect?" Most organizations take an incremental approach to ZTNA, so the answer to this question will likely evolve as implementation progresses. Early in the journey to Zero Trust, a ZTNA solution might protect a limited set of digital assets for a defined user group or role or an area such as finance, and then expand protection from there.
Where do your resources reside?	While common practice is to start small and then expand, getting a complete sense of where everything is located at the start of the process makes the transition smoother. For instance, are there digital assets on-premises, in data centers, in one or more clouds (multicloud), or a combination of all three (e.g., a hybrid architecture)? Knowing the answers to these questions may impact which ZTNA solutions can be implemented and how they are deployed. You may need a unified private access solution that applies ZTNA policies across a complex hybrid IT environment.

How do I want to deploy?

Before adopting ZTNA, it's important to make a decision about deployment options. Some solutions are hosted by the vendor as-a-service, while others are self-deployed, and others offer a choice or hybrid approach. This will have an impact on the variety of ZTNA solutions you can choose from. It's important to consider things like whether it makes sense to have full control over ZTNA deployment or if it would be better to let the vendor manage the ZTNA infrastructure due to resource constraints or in-house skill-set limitations.

Which network traffic flows need to be protected?

To which flows of traffic do you intend to apply Zero Trust secure access methodologies? North/south or east/west? Which is more critical at your current stage of the Zero Trust journey? Most long-term Zero Trust journeys incorporate all traffic flows across the network by building a true Zero Trust café-style network and then applying the principles of least-privilege access between client-to-server, server-to-server, and service-to-service traffic. Service-to-service refers to machines that communicate with one another but are unguided by a user (e.g., microservices that use APIs). This type of traffic flow should still abide by the principles of Zero Trust. That way, if a machine gets compromised, there is a much smaller chance for lateral movement within the network.

Ultimately, the flow of network traffic will affect ZTNA architectural choices. Some ZTNA solutions only protect north/south traffic, others only east/west. If the protection of digital assets will encompass east/west and north/south network traffic, you need to select a ZTNA architecture that supports both traffic flows with a robust and unified policy engine.

What types of applications need to be protected?

In today's enterprises, many critical systems are dated or custom built, and refactoring these systems would be an expensive and resource-intensive project. This is particularly true in the finance, government, and other sectors where organizations have defined operations using legacy applications. This is a problem for some ZTNA solutions because these older systems may not support Security Assertion Markup Language (SAML) or other modern modes of authentication and enabling Single Sign On (SSO). Some ZTNA solutions are built just for HTTP/HTTPS, and sometimes Secure Shell (SSH). However, these standards only apply to web-based applications, not legacy and custom apps. To avoid this, it is necessary to have a ZTNA solution with broad protocol support.

Security as a Differentiator or a Hindrance

Making the case for Zero Trust Network Access goes beyond the inherent security benefits, which should be considered table stakes when evaluating ZTNA solutions.

A dynamic ZTNA solution has the ability to unlock business agility and accelerate digital transformation efforts. Using APIs, scripting, running security as code, leveraging metadata, and integrating ZTNA into existing business processes, systems, and workflows removes traditional security barriers and unleashes operational efficiencies.

Examples of these efficiencies include, but are not limited to:

1. DevOps teams no longer needing to switch between VPNs when working across multicloud and hybrid environments
2. Automated access permissions triggered by opening and closing tickets in your ITSM
3. Simplified onboarding and offboarding of users via your identity management solution
4. Avoiding costly refactoring projects by securing legacy applications
5. Unified policy management across hybrid IT environments
6. Auto-scaling policies with new cloud instances using metadata

IMPORTANT CONSIDERATIONS

Answering the basic questions of ZTNA implementation narrows down the list of available ZTNA solutions that meet your core requirements. From that list a variety of additional considerations will emerge that impact the implementation process and final results.

Exposing Ports or Hidden Infrastructure

Attack surface reduction is a critical component of Zero Trust Network Access and is the first line of defense against adversaries running port scans as part of their reconnaissance campaigns. Your ZTNA solution should actively cloak your ports, making them and the resources they allow entry to invisible. This is achieved using Single Packet Authorization (SPA), which uses proven cryptographic techniques to make internet-facing servers invisible to unauthorized users. Only devices that have been seeded with the cryptographic secret will be able to generate a valid SPA packet, and subsequently be able to establish a network connection.

Dealing with Device and User Risks

The granularity of access control is a key feature of ZTNA. One can surgically remove access to only critical or privileged data based on the level of risk, while traditional tools adopt an all-or-nothing approach to access. If a user or device presents itself as a risk, ZTNA can be configured to completely block or deny all access. However, using an all-or-nothing approach could prevent a user from getting any access, hold them back from completing work, and create a big headache for the help desk. Instead, ZTNA takes a precision-based approach that weighs the factors that should be taken into consideration to make precise decisions around access control. Ultimately, it comes down to control and using risk management to determine what course of action should be taken.

These decisions typically revolve around risky users and risky devices. A person might behave in a suspicious way, such as logging in at 2 a.m. when the workday is 9 to 5. Does this mean the user's account has been compromised, or perhaps they're just working late? A fine-grained ZTNA policy can calibrate and grant appropriate entitlements based on the limited risk. For example, if a user is not behaving as expected, ZTNA can automatically restrict access to all but the most basic digital assets, such as email—but prevent the user from accessing sensitive data or resources until the security team can further determine if the user presents a greater risk. Devices can be handled in a similar way. An infected or stolen device might exhibit suspicious behaviors, and the ZTNA solution can dynamically restrict access to minimize risk exposure from a potentially compromised device.

"Up rules" vs. "Down rules"

Access requirements can get complicated as enterprises are complex entities that each have a unique set of circumstances, requirements, regulations, etc. Many ZTNA solutions work well in use cases that require user/device policies for resource interactions, also known as "Up rules." For example, if a user's mobile device needs to access a database.

However, most sophisticated security teams will have to support the opposite: "Down rules" that deal with interactions between a server, service, or resource "down" to the user device. Remote desktop support and centralized endpoint protection platforms (EPP) are good examples as they must securely update clients. This is also the case for Voice over IP (VoIP), where access control has to flow in both directions. All ZTNA solutions support "Up rules," but not all support "Down rule" policies. If you require "Down rules" today as a part of your future Zero Trust strategy, look for a ZTNA platform that supports both.

Broader Security Ecosystem Integration

A security solution like ZTNA will invariably be part of a much broader, cohesive collection of tools that comprises an organization's complete security and IT operation. ZTNA does not operate in a silo … at least, it shouldn't. ZTNA needs to integrate with threat intelligence tools, Security Incident and Event Management (SIEM) solutions, Endpoint Detection and Response (EDR) platforms, IT Service Management (ITSM) solutions, and more.

Integrating ZTNA requires a solution that is programmable and extensible. In particular, the ZTNA platform needs an API that can connect it with other solutions or scripting capabilities to further enable bidirectional interoperation with the broader security ecosystem. It's even possible to achieve "security-as-code" leveraging some ZTNA solutions. In this scenario, reference code is stored in a secure repository such as Github and is used to configure, manage, and run all ZTNA operations. Running ZTNA-as-code can be used to scale or deploy new infrastructure as well as configure all access policies and entitlements. This approach provides more agility and can support rich integrations.

Scale

The ability to scale should be a major consideration with ZTNA implementation. While it may be a best practice to start with a single-use case and deploy ZTNA incrementally, at some point, the solution needs to be able to handle the full access control load for the entire organization. It must also be ready to handle increased load levels within an expanding footprint, whether planned or that arise out of unforeseen circumstances. The ZTNA solution must be able to keep up and not cause a bottleneck on network access or drag on performance. A "complete" ZTNA solution must be able to scale up linearly and handle the entire enterprise employee base for all applications across the entire network and cloud ecosystem.

The pandemic has made it clear that ZTNA scalability is a non-negotiable requirement.

It must be able to support remote users, new devices, and applications even when these needs arise unexpectedly.

The Appgate Solution

Appgate offers a ZTNA solution called Appgate SDP. It's a comprehensive, highly scalable, enterprise-grade solution for trusted access. Covering every ZTNA use case, it effectively eliminates the attack surface by rendering resources invisible until the user is authenticated and authorized. Plus, with extensive customization capabilities, security teams build a solution perfectly fit for their needs.

With Appgate SDP, access is conditional and contextual, based on the user's multidimensional identity profile and device posture. These live entitlements dynamically adjust in real time based on risk and changing conditions. The solution offers flexibility and can be deployed on-premises, in any cloud environment, or as-a-service, as well as via a client-based model or service-based model for complete user population coverage.

Appgate SDP is known for its ability to simplify operations, configuration, management, and compliance. It unifies access, offering a consistent experience and configuration across hybrid IT environments. The advanced solution simplifies policy definition and enforcement by applying a single framework to all users, devices, networks, and resources using a micro-segmentation approach, combined with detailed activity logs to simplify compliance and audits.

Benefits of Appgate SDP include:

- Control access with identity-centric micro perimeters
- Secure access for all users, devices, and hybrid workloads on a single platform
- Provide a seamless experience with concurrent access
- Make exposed ports invisible to reduce your attack surface
- Restrict access for risky devices with posture checking
- Keep policies in sync with dynamic infrastructure
- Integrate secure access into the fabric of your organization with APIs

Appgate SDP is a scalable, customizable, industry-recognized ZTNA solution built for the enterprise.

Conclusion

ZTNA offers a proven way of managing access control in the distributed modern world. It augments, and at times supplants, the older "connect first, authenticate second" approaches to network and application access. Implementing a ZTNA solution involves considering your current and anticipated security and IT ecosystem and the resulting platform requirements. Robust solutions, like Appgate SDP, prove that it is possible to address all of the top considerations and allow enterprises to implement ZTNA for a robust, scalable, and sustainable high-security posture.

ZTNA as a whole should be viewed as more than just a component of cybersecurity. It is also a driver of digital transformation. At the same time, it's important to understand the different flavors of ZTNA available in the marketplace, their benefits, and how they work into the long-term security and IT vision of your organization.

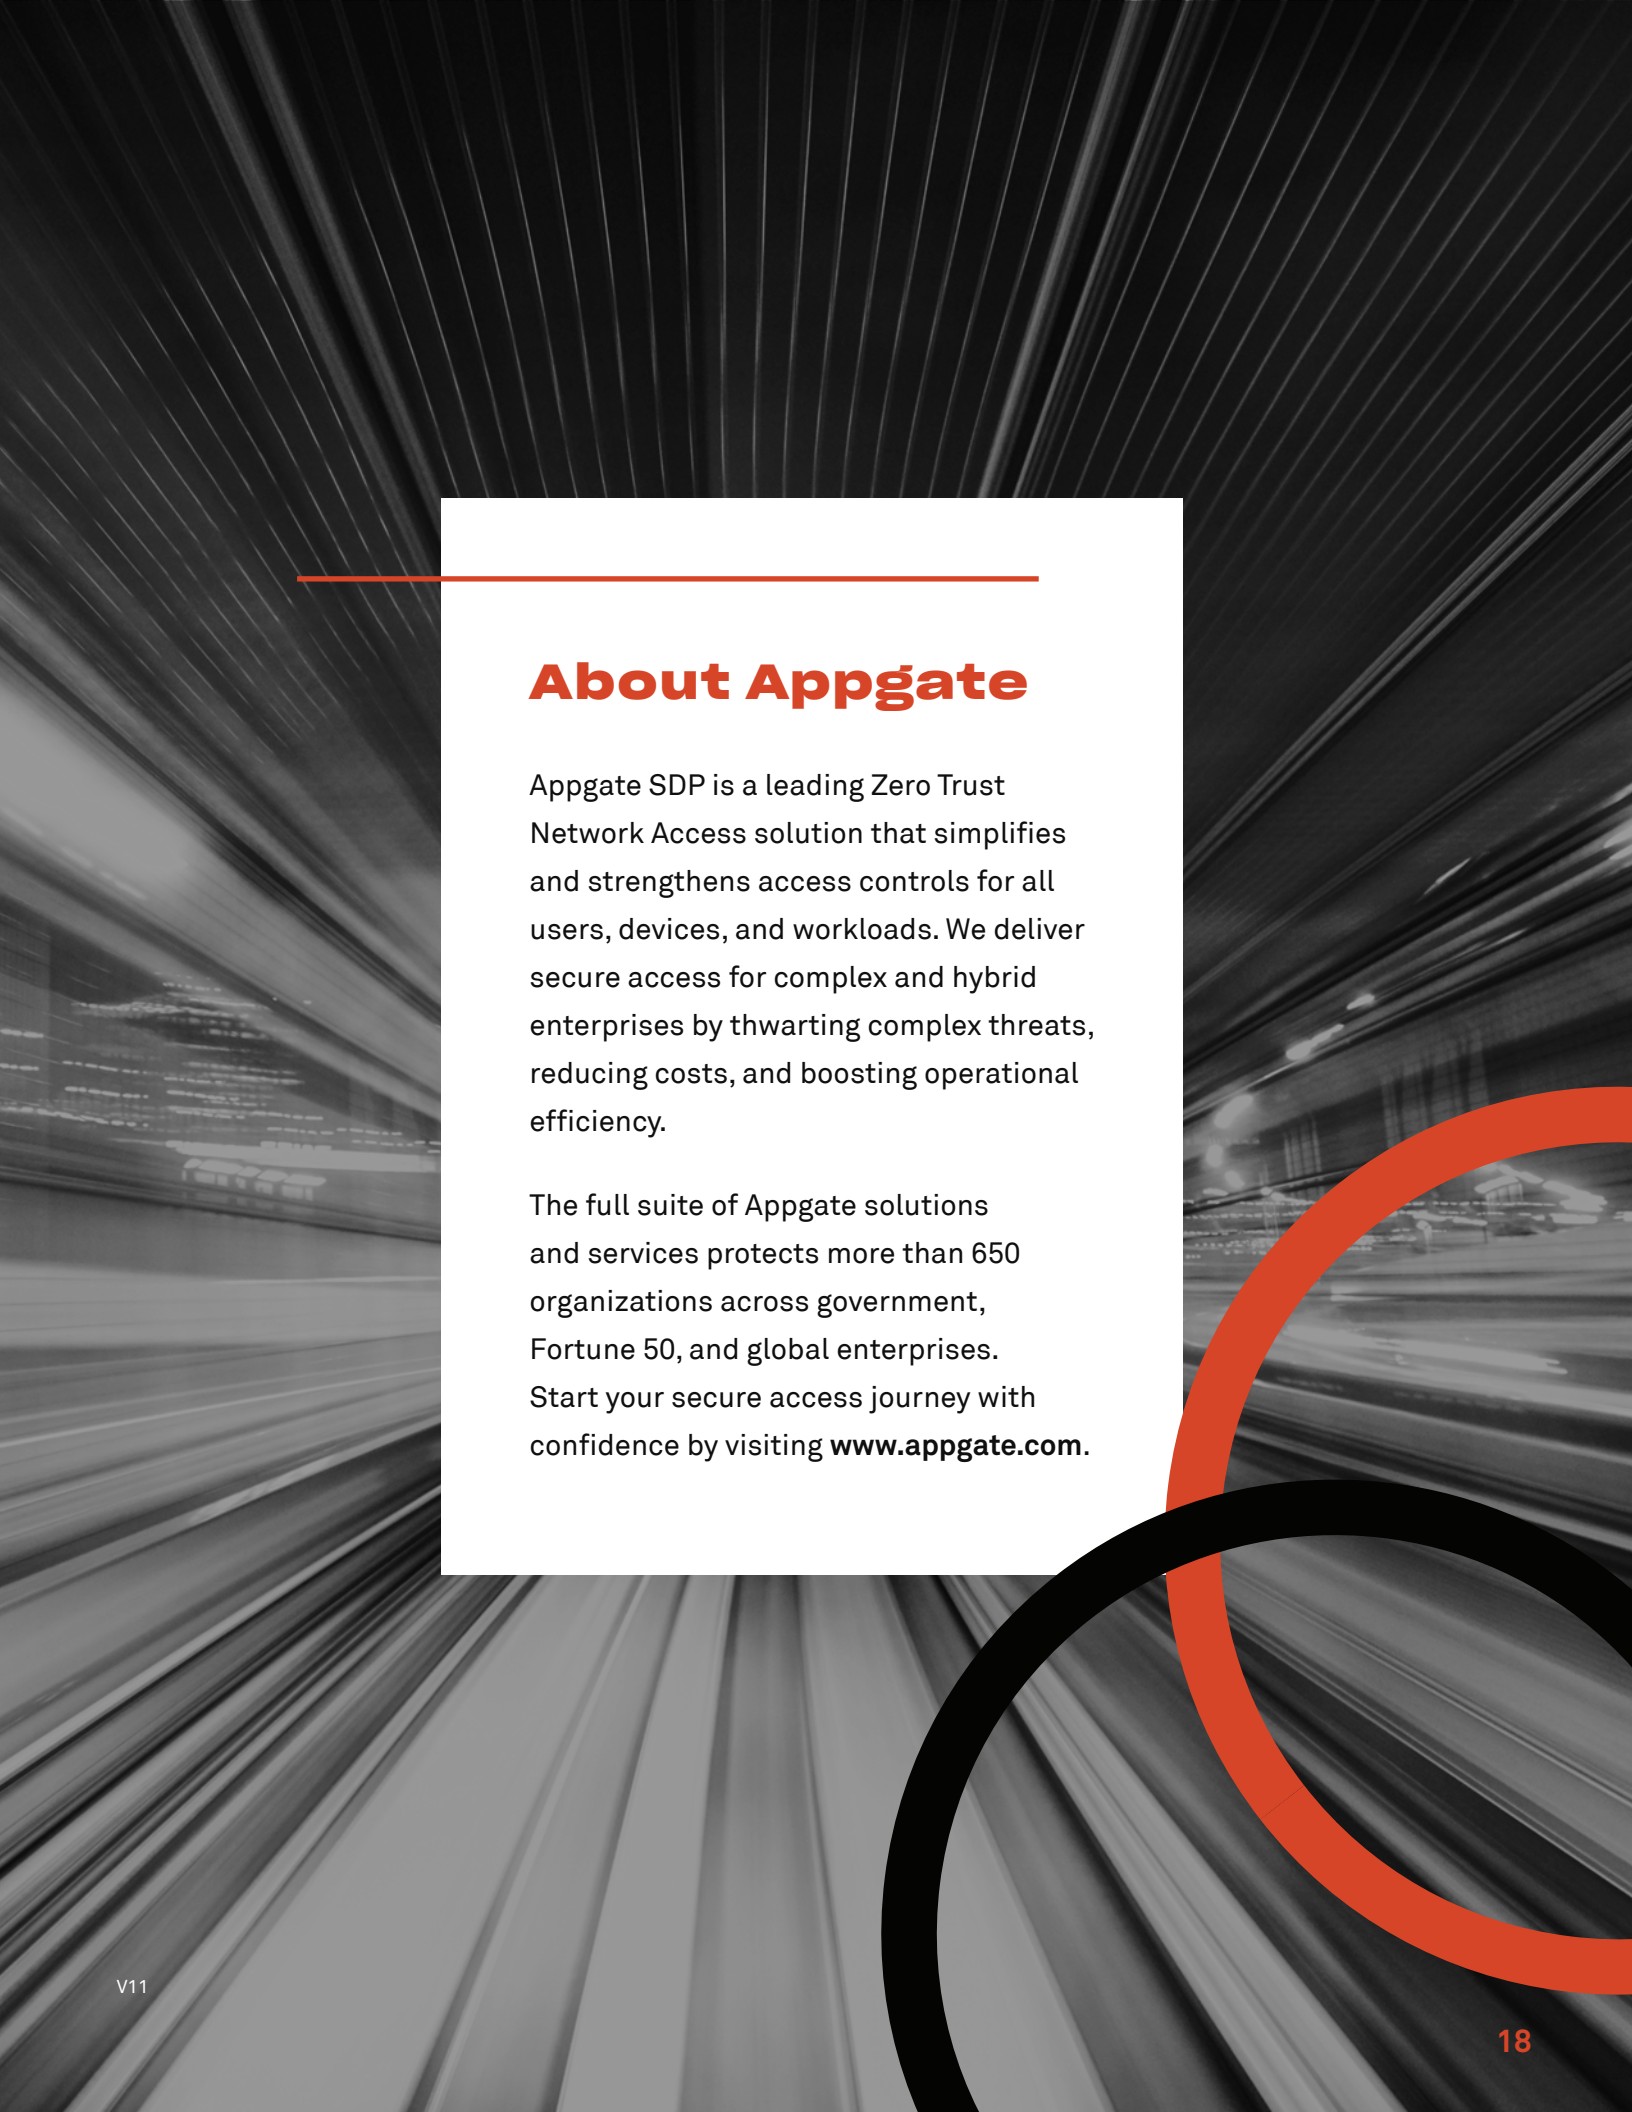

About Appgate

Appgate SDP is a leading Zero Trust Network Access solution that simplifies and strengthens access controls for all users, devices, and workloads. We deliver secure access for complex and hybrid enterprises by thwarting complex threats, reducing costs, and boosting operational efficiency.

The full suite of Appgate solutions and services protects more than 650 organizations across government, Fortune 50, and global enterprises. Start your secure access journey with confidence by visiting **www.appgate.com**.