Point of View

A Practical Guide to Zero Trust Network Access

Key considerations for adopting ZTNA and alleviating VPN struggles

Executive Summary

Enterprise IT teams are being pulled in opposite directions. On the one hand, they need to facilitate their company's digital transformation goals by providing users easy, reliable access, from anywhere, to the resources they need. On the other hand, teams are tasked with security and privacy mandates to protect and control access to those resources from anyone who doesn't actually require it. The status quo—legacy remote access virtual private network (VPN)—fails to accommodate these demands. Zero trust network access (ZTNA), a newer approach, bridges the apparent chasm without creating tradeoffs. ZTNA also fosters collaboration across the IT infrastructure, networking, and security teams—a step that further propels digital transformation with the aim to improve business agility, resilience, productivity, and ultimately, enhance a firm's competitiveness.

ZTNA products are important underpinnings of a secure access service edge (SASE) architecture, for which more than 60% of enterprises are expected to have a plan in place in the next few years. But deploying ZTNA can be a major undertaking. IT teams still grappling with the pandemic may resist such transformational change, fearing the outcome or what it might mean to their job security. For decision-makers, the right choices in three areas can help overcome potential resistance and build buy-in throughout the company:

- Give every stakeholder a seat at the table, and ensure IT, security, and business's needs are met.
- Start with a single use case or focus on one group of users.
- Choose the right ZTNA product.

Good decisions in these areas will enable the entire initiative—and therefore, the bigger transformational shifts involving SASE—to proceed much more smoothly.

Why Move to ZTNA?

ZTNA is already on most IT organizations' roadmaps. The zero trust principle is straightforward: The security infrastructure should, by default, trust no one. ZTNA is a category of products that operationalize the zero trust principle, granting access to specific resources only after authenticating the user's identity and ensuring they have explicit permission to access that resource.

This approach represents a leap forward compared with most organizations' legacy—and many organizations' current—strategy of allowing individuals and devices that gain network access to then connect freely to other resources behind the firewall. Such freedom gives attackers who breach the perimeter a great deal of latitude to move laterally, accessing applications and data throughout the network.

"When your home router is connected to the office VPN server, you may get a permanent connection. … 'Whatever's on the home network has access to the company.'"[3]

Wayne Rash, Forbes

The more a company relies on remote work, the more necessary it is to ratchet down access control. Gartner® predicts that "by the end of 2022, the share of knowledge workers working remotely will increase to 47%, up from 27% in 2019."[1] Gartner® also states that, "by 2024, at least 40% of all remote access usage will be served predominantly by zero trust network access (ZTNA), up from less than 5% at the end of 2020."[2]

ZTNA is a different type of product than a remote access VPN, deployed with different intents and often providing better outcomes, some of which we will discuss below. The two can operate in parallel, or ZTNA can replace a remote access VPN. Either way, adding ZTNA to the infrastructure provides visibility and greatly increases the granularity of control over which users and devices are accessing what resources hosted in a cloud, on-premises, or on a hybrid network. Ultimately, organizations can achieve continuous adaptive trust, an outcome that allows organizations to govern every transaction via insights gleaned from contextual information about the whos, whats, whys, wheres, and hows regarding that resource and access to it.

Where to Start With ZTNA Deployment?

For IT teams planning to deploy ZTNA, an enterprisewide rollout may be intimidating. Coordinating an entirely new approach to accessing resources across large swaths of the corporate workforce, contractors, and third-party users, all at once, creates significant risk—to IT infrastructure, networking, security, lines of businesses, worker productivity, and perhaps even the IT team's standing and trust level within the larger organization.

Instead, IT teams can and should bring ZTNA into the business incrementally. They can initiate a zero trust project with a limited budget and scope, then leverage quick wins to expand ZTNA's footprint within the organization.

The following considerations ensure that a project is successful and generates quick, convincing wins. First, identify and invite the stakeholders to the table. Second, start with a single use case. Finally, select the right vendor with which to partner.

"Zero trust does not require that you rip out all your current security controls to start fresh, and with the right approach, you can realize benefits right away."[4]

Forrester

[1] Gartner®, "Forecast Analysis: Remote and Hybrid Workers, Worldwide", Ranjit Atwal, Rishi Padhi, Namrata Banerjee, Anna Griffen, 2 June 2021.
[2] Gartner® Press Release, "Gartner Forecasts 51% of Global Knowledge Workers Will Be Remote by the End of 2021", 22 June 2021.
[3] Wayne Rash, "Your VPN May Be Your Greatest Security Risk During COVID-19," Forbes, 17 June 2020.
[4] Steve Turner, et al., "A Practical Guide to a Zero Trust Implementation," Forrester, 4 March 2021.

Consideration 1: Give every stakeholder a seat at the table.

A ZTNA deployment may eventually touch most, if not all, aspects of the organization. It can be tempting to devote too much energy to the technology and fail to appreciate its effects on teams and non-obvious processes. An IT team that makes decisions in isolation—that is, without the input of other lines of business—is setting itself up for failure, or at least for intractable resistance to the change. Instead, leaders should gather the perspectives of other groups while planning and implementing zero trust projects. Let's consider four examples.

- **Infrastructure and operations.** The company's networking team is in the business of enabling connectivity, access, and productivity, and their insights should guide much of the decision-making within the project.
- **Application developers and owners.** ZTNA operates smoothly when application owners (who are often also line-of-business leaders) cooperate with infrastructure leaders to identify private applications and maintain their access policies (the who/what/where/when/how mentioned earlier). As users change roles within the organization, their access privileges to certain applications will likely require adjustment.

  Use the application inventory process as an "opportunity to eliminate application access privileges and entitlements that are no longer relevant," as Gartner® advises.[5] This shared responsibility model can be an effective strategy, and succeeds when communication between the application owner and the infrastructure leader is routine and effective.
- **Business units.** Collaborating with line-of-business managers is absolutely necessary. As mentioned above, business units can be owners of some applications. Not only will teams need to enlist their endorsement for the transition to ZTNA, but also devise a list of which third parties, such as contractors, will need ongoing access to specific applications. It's common for very large organizations to lose track of various contractors and consultants, so expect this effort to be dismissed at first. Engage gentle, yet persistent persuasive skills.
- **Security teams.** The information security group provides overall policy guidance on security and compliance. It also advises business units and application owners about risk, but shouldn't own the risk itself: Risk is the parlance of each and every business unit. Obtain documented risk ownership signoffs from each.

SUGGESTED QUESTIONS TO ASK BUSINESS UNIT LEADERS:

What applications does your team use? Are any of those applications mission-critical to your business? Do those applications contain proprietary and sensitive information?

Who will be accessing these applications? What are their roles and responsibilities? Do your users include third parties and contractors?

What are the geographic locations of your users? Where will they be accessing information?

How often will your users be accessing these applications?

How will your users be accessing these applications? What types of devices will they be using (including operating systems)? Are these BYOD and/or corporate-issued devices?

[5] Gartner®, "Gartner® Best Practices for Implementing Zero Trust Network Access", Lawrence Orans, John Watts, and Neil MacDonald, 10 June, 2021.

Consideration 2: Start with a single use case.

Once teams have engaged with the key stakeholders and generated a preliminary application inventory, they will likely find several use cases where zero trust could add significant value to the organization's business and security transformation—maybe dozens of use cases. But the initial deployment of ZTNA should start with a single use case, until the team has gained confidence in the product and demonstrated its value. Possible use cases include:

- **Remote worker access.** More than ever before, employees need access to business-critical applications from anywhere that isn't their office. Even within the remote worker access use case, the IT team can further narrow the scope to one department, one line of business, or one user group. For example, a retail chain may start by focusing on store managers: Every store manager needs access to the inventory management system at headquarters, but other store employees likely don't. Role-based application access will ensure each individual receives only the necessary permissions.
- **Cloud migration for the DevOps team.** Organizations running applications in public Infrastructure-as-a-Service (IaaS) clouds generally establish an early habit of directly logging into the provider's console and directly into virtual machines—habits that mirror the days of traditional data centers. However, because of the sheer scale of most IaaS deployments, mistakes made during direct logins can wreak havoc. Just like ZTNA constrains users only to permitted applications, ZTNA can constrain developers only to the specific resources created for their projects. An added benefit is the performance gains this approach offers over backhauling all developer access through the corporate VPN and back out to the internet.
- **Third-party access.** Contractors, consultants, suppliers, and other constituents often require access to specific internal or private resources. Often they are using devices that are not managed by the organization—thus eliminating access methods that rely on endpoint agents. These third-party users typically need access to only one or two applications. Provisioning full remote access VPN is risky because it grants too much access to the underlying network. ZTNA constrains third-party access to specific internal applications and shields the remainder from third-party users.
- **Mergers and acquisitions (M&A).** In the event that two organizations integrate, corporate functions in one entity will need to access private information within the other entity. ZTNA can be a quick way to securely provide this access without the complexity of combining networks.

MOVING FROM A PROOF OF CONCEPT TO PRODUCTION

Resist the urge to start a ZTNA proof of concept with an application that processes highly sensitive data. ZTNA represents a new paradigm and should be approached with care. Select an application with a small but receptive user population. Educate them how ZTNA alters familiar access methods and explain the benefits. Be receptive to their feedback and fine-tune the configuration. Eventually a repeatable pattern will emerge to form the basis for expanding ZTNA to additional applications, to applications that process sensitive data, and to applications that require third-party access.

Teams should consider the following scenarios as they expand the estate of ZTNA-published applications:

- Highest risk applications, such as those supporting finance or HR
- Users for whom a reduction in latency saves the most money
- Applications that are most prone to generating performance-related help desk service requests
- Avoidance of expensive new capital infrastructure

Consideration 3: Choose the right ZTNA product.

ZTNA products usually contain some kind of trust broker, which mediates interactions between users and applications. When a user tries to access one of the organization's applications, their device connects, either via a browser or a locally installed agent, to the trust broker, which in Netskope's case resides in the cloud. The ZTNA product confirms the user's identity, verifies that the individual is authorized to access the requested application, performs posture checks to ensure the user's device is compliant, and analyzes whether contextual clues (such as user behavior, data movement, and other anomalies) surrounding the access request might indicate additional risk for which to account.

On the resource side, applications are connected to the cloud-based trust broker using a publisher (as in Netskope's case) or connector. With this type of "inside-out" connectivity mechanism, private applications are never exposed to the public internet. Resources remain hidden to anyone who hasn't been given explicit permission to access them. This design differs from remote access VPN concentrators hosted in the demilitarized zone (DMZ), which are discoverable and thus potentially vulnerable to attacks. ZTNA removes the chance that an unknown weakness in firewall or VPN security could result in an attacker moving freely around the network, thus creating a much stronger security posture.

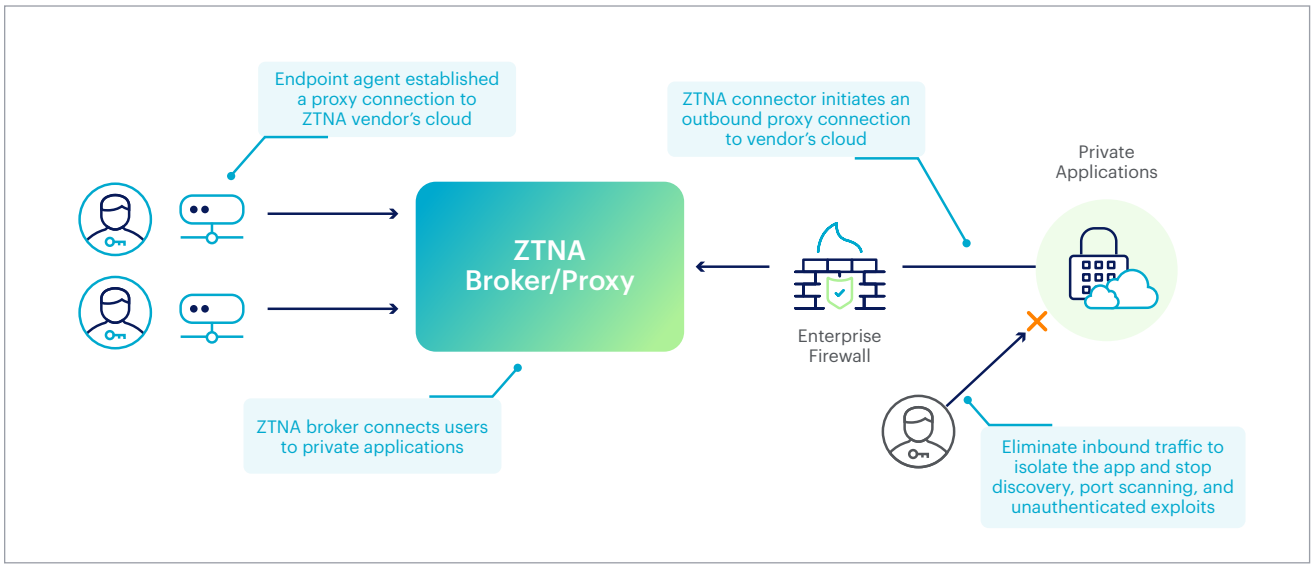

Zero Trust Network Access at a glance

ZTNA is a crowded market, with offerings that vary widely in capability and maturity. Successful ZTNA projects clearly meet the organization's security needs, pose no application performance degradation, and streamline user productivity. When evaluating alternatives, select a product that connects users to the nearest gateway and brokers the traffic from users to applications using the most efficient path. Closely examine the provider's infrastructure: Products riding a bespoke security private cloud outperform those that rely on public IaaS networks. It's also important to ensure confidentiality: Verify that all traffic is encrypted end to end by default. For scenarios where traffic inspection is required, verify that optional transport layer security (TLS) session termination doesn't downgrade to weaker protocols or ciphers.

In addition, Publishers (or connectors) should be completely managed by the ZTNA vendor and the vendor should be responsible for updating the Publisher software including the host operating system (OS) that supports the Publisher instance. IT and infrastructure teams should be weary of vendors that expect the customers to manage the host OS. Enterprises typically have a large footprint of Publishers, resulting in operational challenges for patching the host OS. Unpatched servers create an attack surface and undermine the ZTNA operation.

Likewise, teams should evaluate integrations between products on their shortlist and other components of the security infrastructure. A ZTNA product needs to work closely with an identity and access management (IAM) system that provides multi-factor authentication. Time-based and step-up authentication capabilities further strengthen access control. Strong device posture check is also important. Compare the ZTNA product's native posture check capabilities to those offered by integration with a third-party unified endpoint management (UEM) system and select the alternative that offers the degree of posture check teams will require.

Early in a ZTNA project, teams should narrow down a shortlist to include vendors whose technology platform and vision are aligned with their anticipated SASE architecture and internal roadmap.

Netskope Private Access

Tight integration with the organization's infrastructure is a critical foundation for building an effective SASE architecture. Netskope Private Access (NPA), for example, dynamically steers user traffic to the Netskope Security Cloud to enforce secure access policies and accelerate throughput. NPA supports both web applications and non-web applications utilizing Secure Shell (SSH) or Remote Desktop Protocol (RDP) connectivity, in addition to SQL Server and Active Directory.

RANSOMWARE AND OTHER THREAT PREVENTION NEEDS

ZTNA is sometimes viewed as modernizing access, but every security team—plus every organization converging its IT infrastructure, networking, and security functions into a more flexible, cloud-ready architecture—cares about enhancing security posture as a whole.

ZTNA itself does not detect or block threats, but it can hide assets from attackers and reduce the overall digital attack surface. This is especially important for threats such as ransomware. Over 4,000 ransomware attacks take place daily around the world (FBI). Applying ZTNA can deny attackers the initial attack vector, prevent the spread of ransomware across the network, while keeping operations running.

Remote desktop protocol (RDP) presents a top attack vector preferred by attackers. In a recent survey (cite: Sophos Active Adversary Playbook), remote access services such as RDP were involved at the start of almost one in three attacks, and 69 percent of attacks used RDP for internal lateral movement. Compromised RDP and VPN credentials are routinely outsourced to criminal gangs, including ransomware operators.

With the right ZTNA, however, teams can segment virtually any application located in a local data center or in a public IaaS cloud, without opening any inbound service that can be probed by attackers. There is also no need for any on-premises hardware to install, patch, and maintain, which avoids scalability issues and performance bottlenecks. Finally, a check on the security posture of the endpoint is enforced before accessing the target application.

When Trust Continuously Adapts, Networks Are More Secure

Ultimately, companies adopting ZTNA products should view user access through the lens of continuous adaptive trust. The network implicitly trusts no one. Trust is something that users may gain, on a moment by moment and access request by access request basis, through the context surrounding their interactions. User identity, data sensitivity, device trust characteristics, user and device behavior, network location, data flow analyses, and other factors can play into automated access decisions.

ZTNA products also now offer much more sophisticated decision options than the simple allow/deny dichotomy seen in earlier iterations. ZTNA can restrict users, allowing read-only access; redirect data flows to a safe location; present users with warning messages; and many more options.

IT and security teams should not delay. Attackers rarely take vacations and aren't distracted by anything like a day job. Meanwhile, businesses must seek every opportunity to build competitive advantage, accelerated by enabling the right access to the right applications and data at the right times. Even if the company's end goal is to replace VPN with ZTNA down the road, both can peacefully coexist for a time. Launching a ZTNA initiative that produces quick wins will instill confidence in IT, networking, and security teams, plus demonstrate agility and responsiveness to corporate executives. From there, ZTNA can be extended to additional use cases, and ultimately supplant the legacy remote access VPN.

Tightly controlling access to business-critical applications is the security approach of the future, but organizations needn't wait to benefit from ZTNA today.

Find out more about Netskope Private Access and Netskope's full range of SASE-ready solutions.