RelianceCyber

RELIANCE CYBER WHITEPAPER

Zero Trust

Marketing hype or fundamental best practice?

reliancecyber.com

Introduction

Zero Trust Network Access (ZTNA) is part of the Zero Trust model and it refers to an IT network security solution which organisations may follow in their pursuit of securing access to Corporate resources.

It aims to provide secure access to specific applications and data, rather than granting broad network access, minimising the potential attack surface and reducing the risk of lateral movement in case of a breach.

The Zero Trust model began in the BeyondCorp initiative, which was developed by Google in 2010. Zero Trust Network Access (ZTNA) emerged on the markets in 2014 as a broad term which aimed to "solutionise" some grouped Zero Trust concepts, but it really only began to generate momentum from 2018 onwards. Its adoption was undoubtedly accelerated by the COVID pandemic and the unprecedented requirement for workers to access on-premises and cloud resources remotely.

This whitepaper provides a concise introduction into what ZTNA sets out to achieve, why it is important, and some of the solutions involved.

CONTENTS

Background

How did we get here and what are the challenges?

Since 2009, organisations have invested heavily in procuring Next Generation Firewalls (NGFWs), which are deployed within their datacentres or offices. These aimed to establish a secure perimeter around the servers and applications which host company sensitive information. In simple terms, the idea many years ago was that the Internet was the 'malicious **outside**', and we could maintain security by simply blocking **outside** network traffic from accessing the 'trusted **inside**'. Unfortunately that overly simplistic view does not account for the fact that cyber security threats can originate both internally and externally.

For example, in our experience, the majority of our new customers' user base use some form of VPN solution to access on-premises or cloud resources when working remotely, usually by authenticating them with a security device (normally a firewall) at the perimeter. The VPN therefore brings the remote user **inside** the perimeter – thereby bestowing **trust** upon them... They are **inside** and therefore – from the firewall's perspective - they are **safe**.

An endpoint could be compromised in any number of ways. Today, phishing or some form of browser based attack (such as ad poisoning or search engine optimisation poisoning) is normally the initial vectors. Thereafter the attacker would seek to gain a foothold on the endpoint in order to stage and deploy further malware – including command and control software, giving them unfettered access to the compromised host.

The downsides are obvious. If the user's endpoint is compromised, and they use a VPN to bring their host inside the perimeter, threat actors can trivially leverage these VPN tunnels to move laterally into more valuable assets, to encrypt them or to steal data. Furthermore, by automatically trusting anything that isn't the internet on the inside of our networks, we fail to properly partition them to delay or contain a threat actor in the event that an internal asset is somehow compromised.

Enter the Zero Trust model

According to recent figures from the Office for National Statistics, over 50% of UK employees perform at least some of their work remotely. Our lived experience here at Reliance Cyber is that fewer than 25% of organisations have implemented true ZTNA – instead many rely on VPN access to their environments. Where we do see ZTNA correctly implemented, the risks and attack surface are substantially reduced.

In a Zero Trust model - the solution does not trust any access until we verify and validate every user, every device, every request, no matter their location. ZTNA solutions, implemented properly, never trust the user. Any change in their behaviour or technical configuration requires revalidation of the resources to which they are permitted access.

Is Zero Trust right for your organisation?

When considering an employee's work platform as an extension of the 'trusted network', several questions start to emerge.

- How does the employee access resources in the company datacentre / office remotely?
- What devices do they use, do they have the same security posture as company-owned computers deployed in the office?
- Can the device access potentially malicious internet sites undetected?
- Is remediation possible on devices not managed by the company?
- Does the company employ a Bring Your Own Device (BYOD) policy and have adequate controls in place to maintain security?

The list above is just an example of some of the areas which IT security professionals should consider when assessing remote access, which is a key driver for many of the ZTNA solutions. It is vital to assess which devices and platforms an employee can use to connect to company resources. These devices have the potential to open unintended gaps in the organisation's security model, increasing the attack surface and introducing risk.

Additionally, organisations tend to have some element of their operations hosted in the cloud which demands a more dynamic approach to IT security than simply protecting the on-premises DC with a firewall and enabling VPN remote access. Expanding this Zero Trust to all your corporate environments is key for any organisation.

How can ZTNA help?

The Fundamentals

A ZTNA solution evaluates a multitude of factors before granting remote access to ensure that only authorised users and devices can connect to specific resources. Here are some of the key factors it considers:

Identity:

- **User authentication:** Verifying the user's identity through strong authentication methods such as multi-factor authentication (MFA).
- **Role-based access control (RBAC):** Determining if the user's role or group has the necessary permissions to access the requested resource.
- **Contextual access:** Considering additional context such as the user's location, time of day, or the sensitivity of the data being accessed to make more informed access decisions.
- **Capture:** Identity is generally captured via integration between the SSO solution and the core identity provider. In most cases, this would be Azure Entra, though other Identity and Access Management (IAM) solutions exist for non-Microsoft organisations. These can be integrated too using common protocols.

Device:

- **Device posture:** Assessing the security health of the device, including its operating system, antivirus software, and security patches.
- **Device compliance:** Checking if the device meets the organisation's security policies and standards.
- **Device trust:** Determining if the device is known and trusted by the organisation, potentially using device fingerprinting or certificates.

Network:

- **Network location:** Identifying the user's network location and assessing its risk level.
- **Network security:** Checking the security of the network the user is connecting from, such as the presence of firewalls or intrusion detection systems.

Application:

- **Application sensitivity:** Evaluating the sensitivity of the application being accessed and applying appropriate access controls.
- **Application risk:** Assessing the risk associated with the application, such as its potential for vulnerabilities or data breaches.

Behaviour:

- **User entity and behaviour analytics (UEBA):** Monitoring user behaviour to identify any anomalous activity that might indicate a compromised account or malicious intent.
- **Risk-based authentication:** Adjusting authentication requirements based on the perceived risk level of the access request.

In essence, a ZTNA solution adopts a continuous and dynamic approach to access control, evaluating multiple factors in real-time to ensure that access is granted only when necessary and under the right conditions, thus significantly reducing the risk of unauthorised access and data breaches.

How does this differ from traditional VPN usage?

The key difference between ZTNA and a traditional remote access VPN as a concept for securing access to corporate resources is that the ZTNA solution offers a more dynamic and granular approach based on factors such as where the user is connecting from and what application they are requesting. Some VPN models offer the same broad access to a user whether they connect from a company laptop or a personal laptop. Crucially, ZTNA considers every transaction as untrusted until identity has been evaluated and validated.

ZTNA provides a more scalable security model for any organisation, by validating and classifying every access request to an application; this allows the organisation to control and monitor corporate resource access to a much higher level of scrutiny. From an accessibility perspective the workforce can access these resources regardless of their location which makes this a truly scalable solution.

Integration with an Identity provider to accommodate strong access controls for corporate users and devices is imperative (MFA, SAML, certificate authentication, etc).

The following sections will highlight some of the main aspects of ZTNA solutions.

Network segmentation & discovery

The Basics of Segmentation

Network segmentation refers to logically or physically segregating an organisation's networks based on their security posture or sensitivity of the information and services hosted. Avoiding a 'flat' network is an important part of protecting key services and applications, reducing the attack surface by allowing isolation of traffic and stopping lateral movement within the overall network. Additionally, it allows the application of custom security controls relevant to each segment.

For example, we've all read about cases where entire businesses have been impacted by ransomware and have been unable to operate. One of the main reasons for which this is possible is because of 'flat' internal design, meaning that network traffic can trivially pass from one network segment (e.g. Domain Controllers) to another (e.g. Print Servers). Where organisations segment their networks, the threat actor is unable to propagate their ransomware beyond their immediate landing zone, which significantly limits the damage.

Enter Micro Segmentation

Micro segmentation extends this control to a more granular level. Each individual client, server, or application can be considered as a unique segment. ZTNA effectively provides micro segmentation capabilities because each transaction is authenticated, and the core principle is to deny everything unless it is explicitly permitted (and the identity is validated). This concept inhibits lateral movement within a network (something which is generally possible using traditional remote access VPN), further reducing the potential attack surface.

Device discovery and knowing what devices are connected to any corporate network is key for any organisation as this is the foundation for accurate and effective access controls. This in-depth knowledge is lacking in most medium to large organisations.

Many customers implement a form of network segmentation but fail to address the fact that at any point any device can be connected to the network. If network access control (NAC) or another technology is not used to discover the connecting devices, the organisation is at risk of introducing security issues within their own network.

> "ZTNA effectively provides micro segmentation capabilities because each transaction is authenticated"

Identity and device compliance

Integration with an identity provider provides another core element of ZTNA (authentication and authorisation). This entails enforcement of access policies that only permit legitimate users access to resources which they are authorised to use.

Device compliance and posture is utilised to validate that endpoint clients are allowed and up to date before accessing corporate resources. Remediation policies for non-compliant devices enforce system updates to return the device back into a compliant status before accessing corporate resources.

Enforcing strict access policies for non-corporate devices based on the organisation's BYOD policies maintains separation of security policy based on the premise of trusted vs semi-trusted.

Access policy and enforcement

This is one of the ZTNA elements that is more time-intensive to implement but once in place enhances the scalability of ZTNA. Feeding in from network segmentation and discovery, access policies can be created with more value, to create a granular policy which meets the organisation's security needs.

Creating access policies based on user roles, application types, and location, will provide the organisation with the required fine grain controls to limit unauthorised access to critical applications.

Access policies should always be based on the least privilege principle to minimise risk from internal threats as well as external. For example, a penetration test employee should have no potential to access HR systems and edit sensitive and confidential information, even when they are connected inside the network.

Software defined perimeter

Software Defined Perimeter (SDP) is a key part of a ZTNA consideration and refers to a shift away from the traditional concept of relying on hardware topology to protect access to resources. With a defined ZTNA solution, access should not be afforded to a system from an unknown identity or an authenticated identity who does not have sufficient privilege to access that resource.

Essentially, an authenticated user will be provided with a dedicated connection to that resource only with no access to the wider network as part of that transaction. This means that the perimeter should effectively be invisible and inaccessible to users and clients that are not provisioned to have access. Put simply, access to each corporate resource is evaluated on a transaction-by-transaction basis.

The SDP essentially defines how users are authenticated and the subsequent access to applications based in the cloud or on-premises.

In simple terms SDP replaces the existing VPN solution but it's the whole framework of ZTNA combined that provides such enhanced security benefits over legacy perimeter security.

Benefits of ZTNA

- Improved security to corporate resources, defining and enforcing more granular access policies based on least privilege principle.
- Improved user experience by providing secure access to applications regardless of the user location. This results in a standard approach to accessing corporate applications.
- Cost effective and scalable as most ZTNA solutions are cloud based with on-premises connectors or device agents, providing a centralised deployment that provides simplified administrative costs and is easy scalable to meet organisation's needs.
- Visibility and compliance by providing detailed access logs and applying restricted access based on user's access privilege. This visibility also provides enhanced data analytics for detection and response to possible attacks.

Summary

Some of the key elements of ZTNA covered in this document clearly provide tangible benefits to information security. A complete ZTNA solution will include NGFW capabilities (e.g. HTTPS inspection, malware detection, Application Control, etc) while integrating identity authentication and adopting the principle of least privilege. However, it is important to understand the challenges which come with migrating to a ZTNA model.

Firstly, it requires extensive insight into all of the company's assets and dataflows which can be time consuming and complex. Having this understanding is critical in a model which effectively locks down access unless explicitly permitted, to avoid breaking important access inadvertently. Operational issues may also occur due to the inherent 'gateways' introduced to workflows, which some applications may not tolerate or be compatible with. Administrators must also understand that while ZTNA can offer enhanced security and performance, the security of systems can still be undermined with badly configured policies (just as with traditional network security methods).

To summarise, ZTNA should be seen as a journey rather than a one-off project. Adopting a phased approach provides a more manageable method of working up to ZTNA and sets the foundation for what should be a lifecycle model which continues to adapt over time. Additionally, it is vital that organisations have a solid understanding of their own assets and workflows to take advantage of the benefits offered by ZTNA toolsets.

"ZTNA should be seen as a journey rather than a one-off project"

Case study

This section provides a high-level example of a ZTNA journey. The topologies below present various stages, starting from the original form through to the end state. The environment is representative of a standard setup which many organisations have adopted in recent years, with a mixture of corporate applications hosted in the on-premises DC as well as the cloud:

- Perimeter firewall cluster deployed in the on-premises datacentre for North-South inspection, NGFW capabilities and hosting the remote access IPSec VPN.
- IPSec VPN from a well-known vendor utilised for employees to gain access to services in the company's datacentre, with a flat network architecture on the LAN.
- Various on-premises applications hosted in DMZs connected to perimeter firewall.
- Microsoft services (Office 365, SharePoint) and some Azure hosted corporate applications hosted in the cloud, with Azure AD (now Microsoft Entra ID) used for single sign-on (SSO).
- Employees able to use personal devices for remote access VPN and access to cloud services with no endpoint controls or posture checking.

Original

Access to cloud services is provisioned using Azure AD (now known as Microsoft Entra ID). Remote users can access on-premises services using remote access VPN from any device by installing the vendor's VPN client on their machine. Broad firewall policies allow access to most of the LAN on a range of ports, providing a large attack surface in the event of a compromise.

While some applications are hosted within the perimeter DMZ, a lack of LAN segmentation means that sensitive applications are reachable from the main corporate network. Additionally, there is little information available concerning the types of devices connecting to resources, and the applications and services required for various roles within the company.

ORIGINAL ENVIRONMENT

- Flat LAN network without segmentation.
- Azure Active Directory for cloud services.
- IPSec VPN provides access to on-premises services filtered on IP layer.
- Personal devices in use with minimal corporate control / visibility.

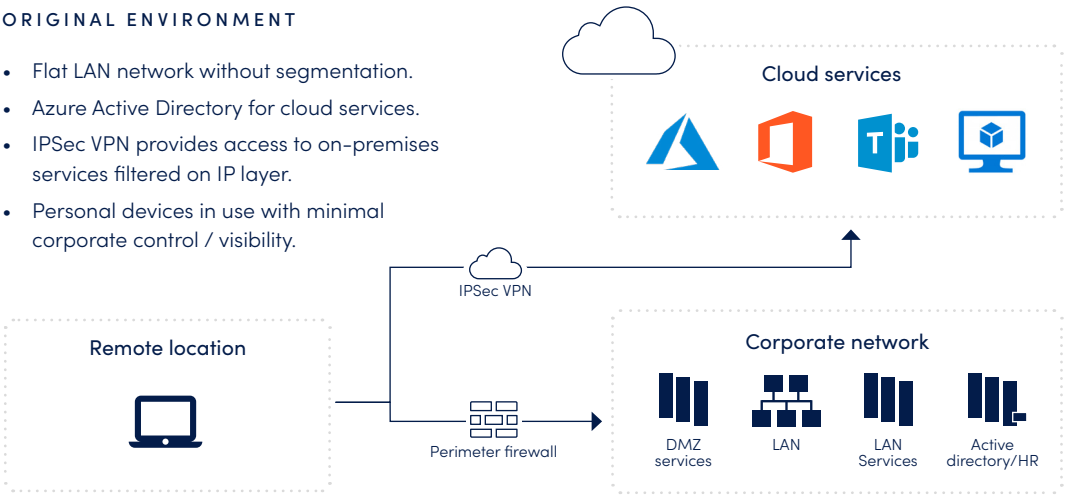

Phase 1 – Discovery

Timeframe: Generally 1 week

A key stage of achieving Zero Trust is gaining a solid understanding of which devices are connecting to corporate resources, what are your applications and their classification. This is a crucial pre-requisite to forming well-defined access controls based on which users / roles should have access to which applications, and the types of devices that should be permitted.

Some vendor firewalls provide a feature that can be enabled which provides identification of devices. This contributes to the discovery of details such as IP addresses, operating systems, hostnames and usernames connecting to various services and applications. Analysis is performed to establish what will become the foundation of the Zero Trust access policy.

Phase 2 – Segmentation

Timeframe: Generally 2-4 weeks, depending on size and complexity

The flat LAN network is redesigned using Virtual Forwarding and Routing (VRF). VRF is a networking technology that allows multiple virtual routing instances to exist within a single physical router. Each VRF has its own independent routing table and forwarding information, enabling the coexistence of overlapping IP address spaces without conflict. This facilitates network segmentation and isolation, enhancing security and enabling service providers to offer multiple VPNs to different customers using the same physical infrastructure.

There are various solutions to achieve network segmentation, in this case Shared Services and HR / Accounts VLAN's are placed in dedicated VRFs based on common purpose and security posture.

Logical interfaces are created on the perimeter firewall, acting as layer 3 gateway for each VRF. Traffic between the VRF VLAN's themselves (as well as between VRF VLAN's and global LAN) must pass the firewall for inspection.

Log analysis allows for gradual policy refinement until granular control is achieved.

SEGMENTATION

- VRF routing created to logically segment LAN networks.
- Perimeter firewall utilised to inspect and restrict inter-VRF traffic.

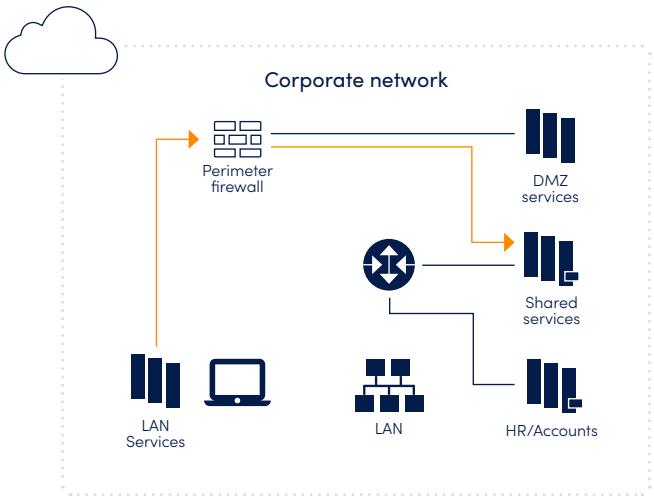

Phase 3 – Zero Trust Deployment

Timeframe: Generally 2-4 weeks depending on size and complexity

Client connectors are deployed as agents to corporate laptops, providing connectivity between users' machines and the Zero Trust provider network. On-premises / application connectors are deployed within the corporate network, providing integration between the Zero Trust provider and the on-premises environment.

There are various methods of implementation, but this example ensures consistent enforcement of the same access policies regardless of whether the employee is working remotely or within the corporate premises. The Zero Trust provider receives connection requests from users to protected resources, performs authentication via the Identity Provider (IdP), then passes the auth token required for access to the requested application. This access is brokered by the on-premises connectors which provides complete control over employee access to corporate resources. Each transaction is authenticated, and administrators have enhanced visibility over corporate traffic flows.

ZERO TRUST

- Client connectors deployed to corporate laptops.
- On-premises connectors deployed inside the network
- Location agnostic approach, the same policy is applied to internal and external users.
- Access to resources must use Zero Trust provider, ensuring identity authentication and policy enforcement for each transaction.

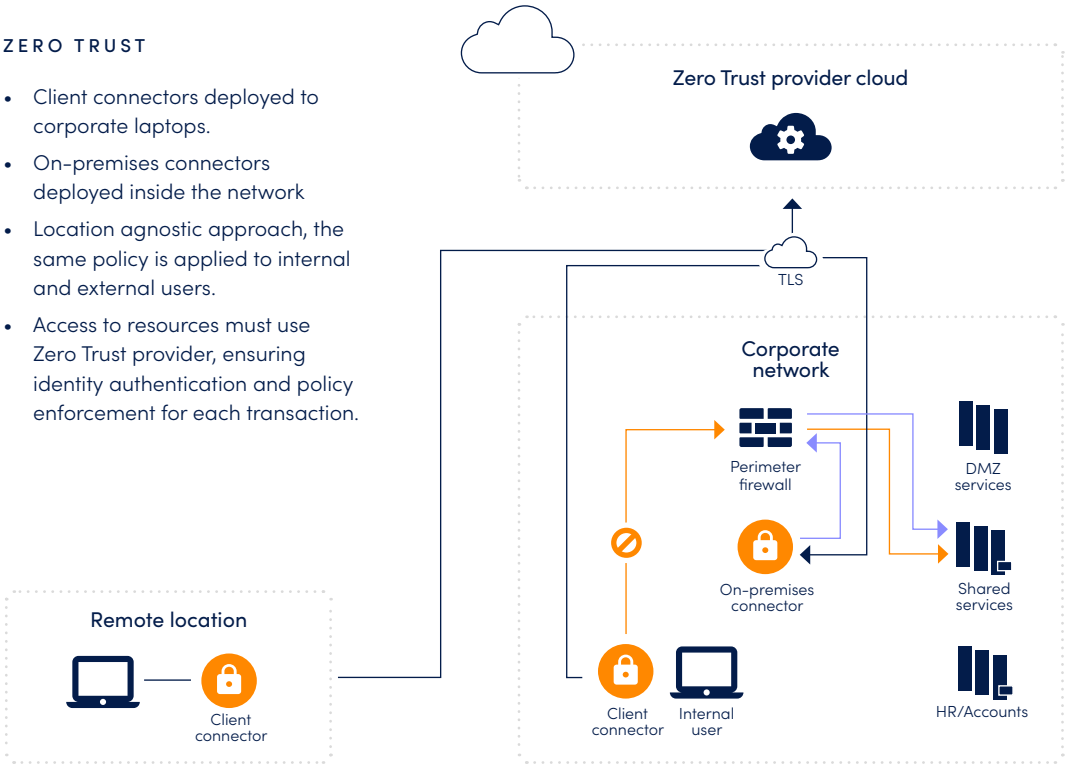

Phase 4 - Policy Refinement

Timeframe: Ongoing and continuous through life

With the foundations in place, Zero Trust access policies can be refined to achieve the level of granularity required. The Zero Trust platform is utilised to analyse logs and establish a baseline of what access is required, by whom. In most cases there is initially a lack of deep understanding concerning which services and applications are needed by which employee types, and the inter-dependency between applications. As such, semi-open rules can be placed toward the bottom of access policies in order to gradually build the permanent policy without impacting operations.

Why don't all organisations do this?

As the phases above make clear, embarking on a segmentation and ZTNA project is neither simple or fast. Vendor selection in itself is a difficult and time-consuming process, let alone the multitude of technical and architectural decisions that must be made to get there. Our assessment of the situation based on our own experience is that the friction is caused by three main factors:

- **Time:** Internal networking teams are generally busy keeping things updated, patched and adding new rules and policies to networking devices. They also are asked to respond to incidents such as power outages or device failures. "Keeping the lights on" is always the priority for stretched internal teams, placing the planning and execution effort required to deliver ZTNA out of reach.
- **Cost:** If the hardware backbone is ageing, replacing it with more modern infrastructure and licensing a product such as ZScaler is something that must be individually costed and budgeted. The projects can be expensive and require a commercial process, which again impacts time. CISOs and Heads of IT/Infrastructure sometimes struggle to justify the time and budget against other priorities. For even small businesses, the entry to ZTNA can be between £35,000-50,000, but compare that against the rising cost of security breaches and associated fines from regulators. This was estimated to be £3.4m for UK companies in 2023.[1]
- **Expertise:** Transitioning from a flat network with a traditional VPN set up, to micro segmentation and ZTNA is far from trivial. Changes must be made at every layer of the Open Systems Interconnect (OSI) Model. This deep expertise does not always exist within SMEs, making the undertaking daunting.

How can Reliance Cyber help?

Reliance Cyber have a wealth of expertise in segmentation, micro segmentation and ZTNA. We work with multiple vendors in this space and can cost a range of different solutions and options. Our Project Managers have overseen ZTNA implementation for many government and private sector customers.

We are here to help, and would be delighted to hold a free of charge roadmap consultation call with any organisation that is looking to mature in this space, but lacks the time, expertise or is concerned about the budget required to deliver this.

[1] https://uk.newsroom.ibm.com/24-07-2023-IBM-Security-Report-Cost-of-a-Data-Breach-for-UK-Businesses-Averages-3-4m#:~:text=This%20year's%20report%20shows%20a,this%20is%20still%20a%209

RelianceCyber

Get in touch

+44 (0)845 519 2946
contact@reliancecyber.com

reliancecyber.com