CyberCX

Down with trust: a practical, actionable approach to achieving Zero Trust

CyberCX

Contents

INTRODUCING ZERO TRUST

The modern approach to cyber security

Trust no one and no thing. That is the central thrust of the aptly named Zero Trust paradigm. In stark contrast to traditional castle-and-moat approaches, Zero Trust is built on the idea that everything and everyone must be considered suspect – inside and outside the castle. With Zero Trust we must interrogate, investigate and cross-check until we are 100% positive the access is safe to be allowed.

Born out of the need to mitigate the risk of prolific and evolving attacks in the complex, dispersed modern environment, Zero Trust continues to gain momentum for its hard-line security stance and rigorous approach to countering today's escalating cyber threats.

Still, after almost 10 years since it was first coined[1], Zero Trust both excites and confuses. Excites because by restricting access to sensitive data, applications, and devices to the bare minimum, it provides a high level of security assurance. Confuses because there is no single specific technology or vendor associated with Zero Trust. Rather, it's a holistic approach that can require mixes of several different principles and technologies.

In this paper, we explore the fast-emerging world of Zero Trust beyond the initial definitions of what it is and isn't. We'll look at the challenges and how it should work in the real world. We also put the 4 most popular vendor approaches under the microscope, illuminating each approach and evaluating their upsides and downsides, giving you a solid overview of your Zero Trust options.

A definition:

Zero Trust Architecture

Zero Trust provides a collection of concepts and ideas designed to reduce the uncertainty in enforcing accurate, per-request access decisions in information systems and services in the face of a network viewed as compromised. It is an enterprise cyber security plan that utilises zero trust concepts and encompasses component relationships, workflow planning, and access policies.[2]

[1] https://www.darkreading.com/attacks-breaches/forrester-pushes-zero-trust-model-for-security/d/d-id/1134373

[2] NIST (2020), Draft (2nd) NIST Special Publication 800-207 Zero Trust Architecture. Available at: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207-draft2.pdf

IMPLEMENTATION

Putting the Zero Trust model into action

With the rise in mobility and remote working, people are creating, storing and accessing data from anywhere, meaning critical information is distributed across cloud, datacentres, campus and users. As we've briefly discussed though, legacy controls on the perimeter of the network don't protect against attacks within the enterprise, so all communication needs to be untrusted by default.

With Zero Trust, the end goal is clear - prevent unauthorised access to network resources which by the Zero Trust definition includes all data sources and computing services. Trust no one. This is turn leads to the question of "how".

One fundamental method of achieving Zero Trust is by making access control enforcement as granular as possible - a concept known as micro-segmentation. In its simplest form, this means breaking up a network into lots of smaller logical pieces, each with their own access control. In this way, a breach in one area doesn't automatically mean your entire network is compromised, however there is a clear cost/benefit trade-off with micro-segmentation.

While segmentation and micro-segmentation are what most people think of when they first consider Zero Trust, it's not the only approach and in certain situations it may not be the best approach either.

For example, how many segments are enough to achieve the stated goal of Zero Trust? Take it too far and the additional complexity, cost and overhead required to implement and maintain micro-segmentation can outweigh the benefits.

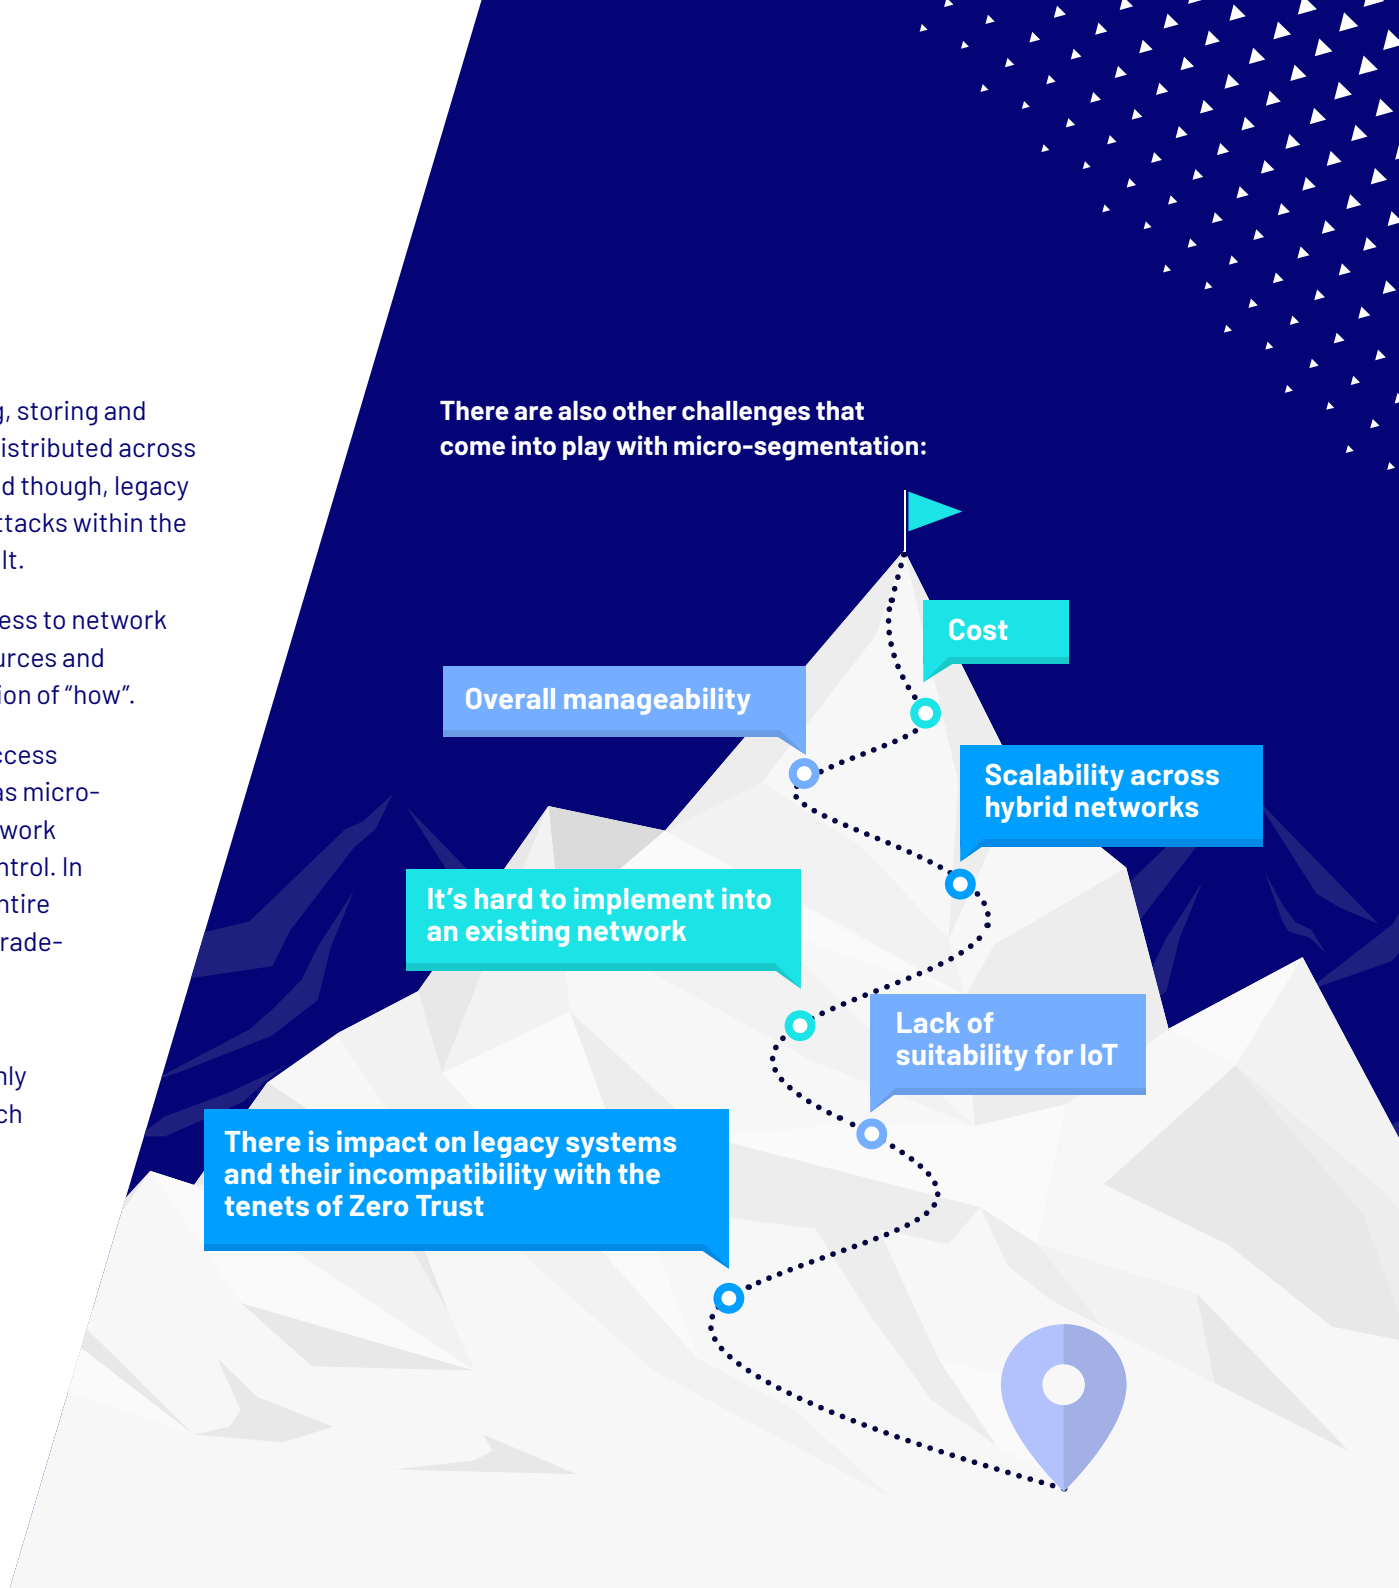

If you only implement segmentation and micro-segmentation, the increase in security effectiveness starts to quickly level off as you increase the degree of segmentation, yet the costs keep rising to an unsustainable level.

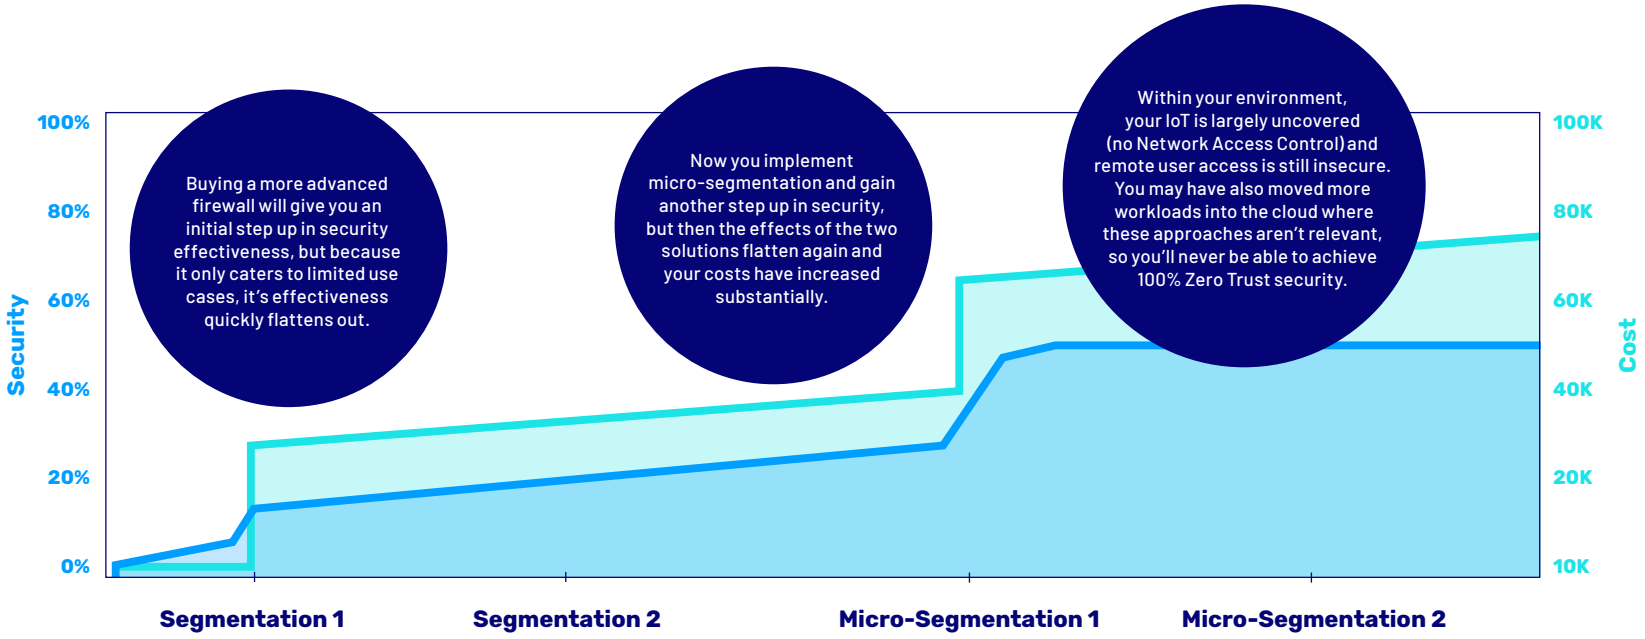

A better approach relies on performing some quantitative risk analysis to help you choose the right Zero Trust method for the right environment. It may still include some level of micro-segmentation, but now you have a way of continuing to increase your security effectiveness without an unsustainable increase in cost.

How Zero Trust *should* work

Despite these challenges, there is a logical, step-by-step approach to implementing Zero Trust into an existing perimeter-based network.

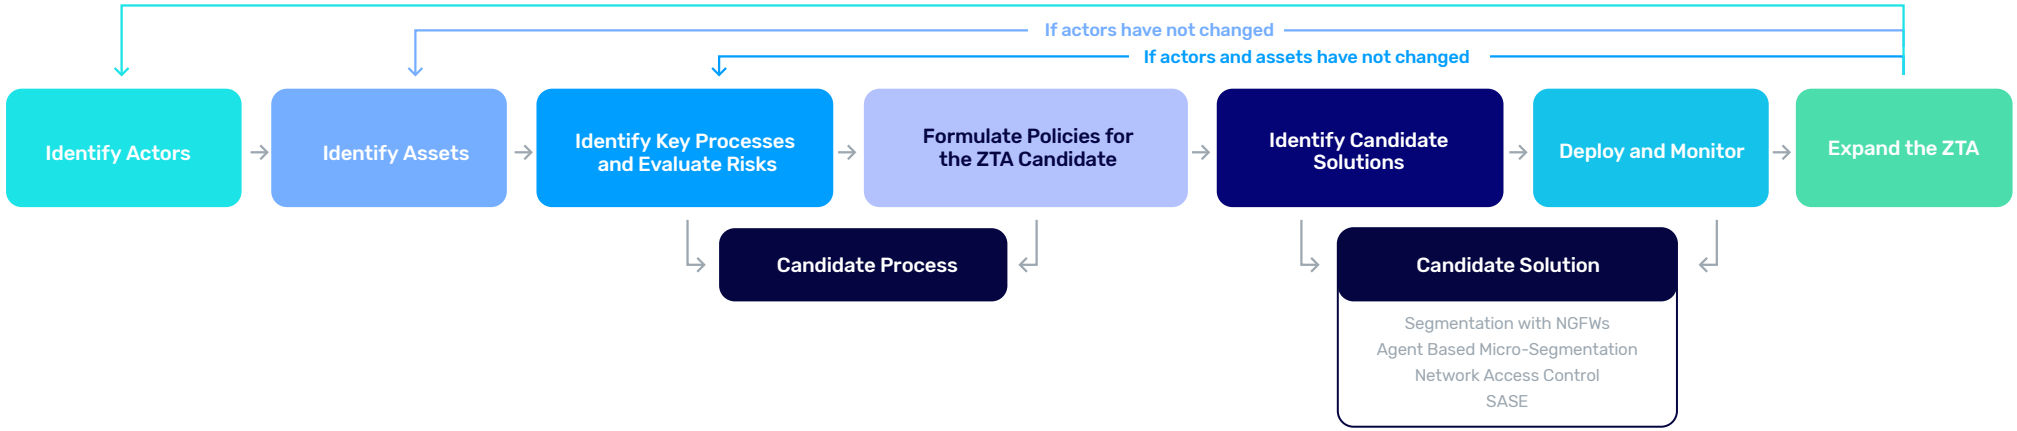

Identify Actors and Assets
The first step is to identify who is active in the network. This includes standard users, privileged users, service accounts, etc. and the assets they use such as hardware, applications, etc.This is a complex task and requires the interrogation of multiple systems and interviewing relevant subject matter experts.

Identify Processes, Evaluate Risk and Create Policies
This is a challenging step that requires mapping flows and interactions in processes and assigning risk to each process. Any prior business impact analysis can help in this step while application dependency mapping tools are essential to ensuring processes and applications are understood. Once the risks associated with a process are defined, suitable controls can be applied resulting in a set of high-level policies.

Identify, Deploy and Monitor a Solution
Once the policies have been defined, and an appropriate process has been selected as a starting point based on your assessment and appetite for risk, a candidate Zero Trust Architecture solution can be selected.

Some vendor solutions and deployment models are designed for different situations and use cases so the inputs to this step are critically important to ensure the right product/solution is chosen. Subsequent monitoring is then essential in ensuring the assumptions underlying the chosen model are playing out as predicted.

Expand the Solution
If the previously deployed solution meets specified success criteria, it can be expanded to other applications or processes to which it is suitable. If not, the process must begin again at identifying and deploying a more appropriate alternative.

In the second half of this paper, we will break out some of these vendor solutions and approaches to provide a high-level view of these differences - and the pros and cons of applying them.

Application dependency mapping

To build effective security policy in a Zero Trust enterprise, you need a clear understanding of your applications and their dependencies. Thoroughly identifying and mapping all the dependencies of your applications can be a time-consuming process – especially if you try and do it manually.

A manual approach also means you risk missing key applications or introducing errors into your process. There are various tools that can help with this, including Illumio ASP. It generates an interactive, real-time map that automatically identifies your high-value system, connections and dependencies, giving you deep visibility into users, devices, applications and more.

Evaluating Risks

Selecting a target process or workflow for Zero Trust relies on evaluating the risk of that process and assessing the resulting value to the organisation in mitigating that risk. Organisations already manage risk day to day, such as aligning cyber security policies with a framework like NIST, and performing annual risk assessments.

However, the measurement of risk is left up to the user, and many rely on qualitative assessment of risk, which makes determination of a risk mitigating solution difficult. There are no metrics to compare.

Instead, we would suggest that quantitative risk evaluations be performed as they are more rigorous, and the data generated is more defensible. Approaches such as the FAIR quantitative risk analysis model could be helpful.

SOLUTION A

Segmentation with next-generation firewalls

In a nutshell:
Protect the security of sensitive data and critical applications by dividing the network up into smaller, more secure enclaves.

Suits:

- Networks in data centres or the cloud that require a moderate degree of segmentation.
- Traffic or communications that must be secured with next-generation firewall capabilities.
- North-south security

Notable vendors:
Palo Alto Networks, Cisco Systems

Slicing up a network into different zones and segments based on usage, function, or location is nothing new. Organisations have been isolating and controlling the network like this for many years. For a Zero Trust solution, the next-generation firewall becomes the segmentation gateway, creating a micro-perimeter around the protect surface. This means policies are enforced in the next-generation firewall at the centre of the network, as close as possible to the resources they protect.

Layer 2 VLAN with next-generation firewall default router

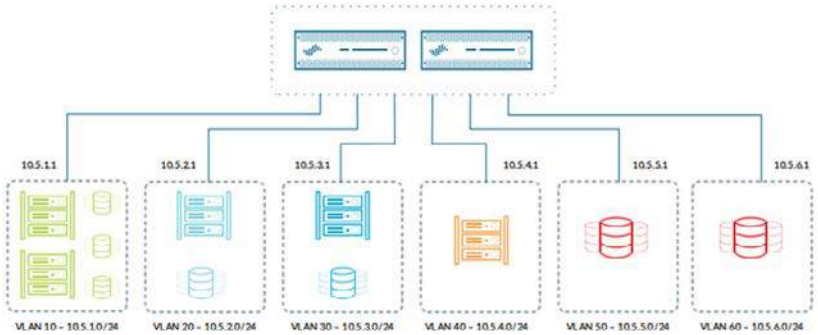

Layer 3 VRF with next-generation frewall integrated through OSPF

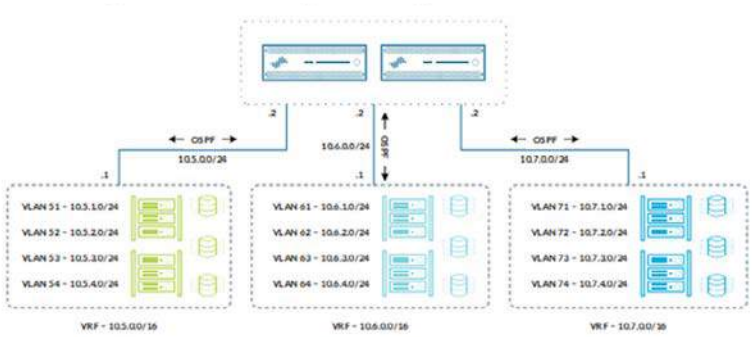

https://www.paloaltonetworks.com/resources/guides/securing-data-center-public-cloud

UPSIDES

- Network traffic and behavioural analytics capabilities to quickly find and stop the stealthiest network threats.
- An intrusion prevention system examines network traffic flows to detect and prevent vulnerability exploits.
- Network security optimisation with deep packet inspection to identify and filter out malware and other unwanted traffic.
- Future-proofed with ability to update as the security landscape evolves.
- Some protection against lateral movement attacks.

DOWNSIDES

- Can be expensive
- May have capacity limits due to the solution being hardware bound - requires a forklift upgrade.
- Less suitable for east-west communications.
- Requires significant knowledge of the network architecture upfront (an Application Dependency Map must be done prior to this).
- Requires careful planning to implement and scale for large hybrid environments.

SOLUTION B

Agent based micro-segmentation

In a nutshell:
Puts a stop to lateral movement inside hybrid environments with agent based micro-segmentation.

Suits:

- Networks that require a high degree of east-west micro-segmentation
- Cloud or data centre
- Works better with managed workloads

Notable vendors:
Illumio

Segmentation has traditionally been used to boost network performance. Today, with the perimeter-based security approach no longer able to protect the enterprise, micro-segmentation - the more dynamic and granular cousin of segmentation - has become the foundation of data centre and cloud security. Here's where agent based micro-segmentation comes in. It decouples security segmentation from network infrastructure and pairs security segmentation with real-time application dependency mapping, to effectively prevent lateral movement inside data centres and cloud environments.

Illumo ASP Architechture

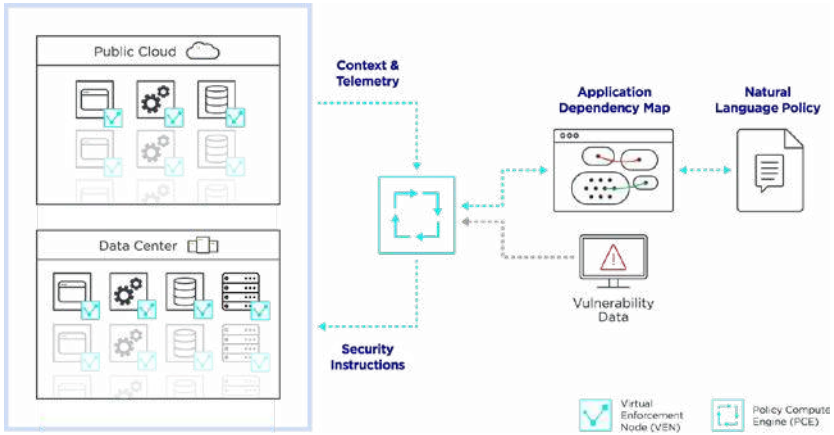

https://www.illumio.com/product/architecture/adaptive-security-platform

UPSIDES

- Micro-segmentation to enhance the efficiency and security of both on-premises or cloud-based network.
- Simple Application Dependency Mapping as a consequence of having agents on the workloads sending telemetry to a controller.
- Hugely scalable.
- Excellent protection against lateral movement attacks.
- Provides more control over workloads.

DOWNSIDES

- Because this approach relies on the underlying workload operating system, firewall capabilities are limited.
- Limited to workloads able to receive an agent such as managed workloads and not for IoT/OT.
- Requires an agent on workloads.

SOLUTION C

Network Access Control for IoT environments

In a nutshell:
Rethink the concept of identity and expand your Zero Trust initiative to protect unmanaged and unmanageable IoT infrastructure.

Suits:

- IoT/OT Networks
- Unmanaged workloads and endpoints with older operating systems
- Situation where you have a limited ability to deploy Policy Enforcement Points at some locations.

Notable vendors:
Aruba, Forescout, Cisco

Securing any network begins with knowing. You must know every connected user and device, and the data they are trying to access. For Zero Trust, or any security measure, you can't create appropriate enforcement policies and controls without this knowledge. Which makes the sheer variety and volume of unsecured IoT devices a massive challenge. Securing IoT with a Zero Trust framework involves understanding the identity of every device that touches your network - including context, traffic flows and resource dependencies - before making access control decisions based on these parameters.

Building automation network withIoT devices

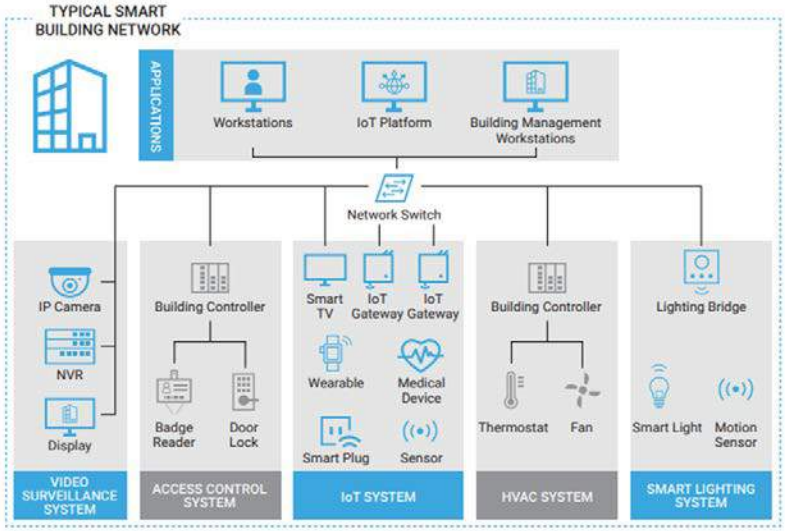

https://www.forescout.com/company/resources/iot-research-report-transforming-cyber security-strategy-for-the-age-of-iot(page 8)

UPSIDES

- Provides the ability to manage access to resources for a vast array of workload types.
- Delivers visibility of endpoints and to a lesser extent, communications.
- Offers endpoint discovery and profiling to lower risk profiles of endpoints prior to network access.
- Excellent control of endpoints.

DOWNSIDES

- The solution has a lack of next-generation capabilities. E.g. deep packet and application inspection, intrusion prevention.
- Can be difficult and time consuming to design and deploy an effective policy.
- Potentially limited usefulness against lateral movement attacks.
- Works better with an agent on managed endpoint workloads.

SOLUTION D

Secure Access Service Edge (SASE)

In a nutshell:
Protect yourself from the security pitfalls of the cloud with a cloud-aligned networking infrastructure.

Suits:

- BYOD
- Global organisations with dispersed locations (if using a SaaS option)
- Environments needing to regulate access in hybrid cloud scenarios.

Notable vendors:
Palo Alto Networks, Netskope, Zscaler

Organisations are operating in a time where accessing an ever-increasing number of applications outside the traditional network perimeters is the norm. Software, platforms, infrastructure and application workloads are often in the public cloud. The challenge for organisations is to support the user experience while protecting their users, applications, and data from security threats. A portal-based model is a purpose-built cloud-delivered infrastructure that can identify users, devices, and applications, no matter where they're connecting from. This approach ensures networking for all applications and consistent security that ensures policies are enforced at all times.

SASE delivers advanced network and security capabilities in a convereged cloud-delivered solution.

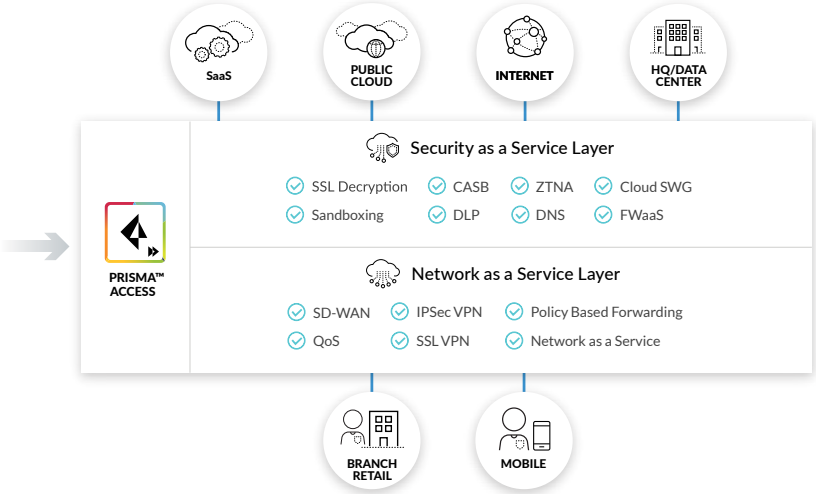

https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/ebooks/the-10-tenets-of-an-effective-sase-solution.pdf (page 12)

UPSIDES

- No need for agents on endpoints.
- Highly scalable and vendor managed if SaaS option is selected.
- Predictable cost when delivered as-a-Service.

DOWNSIDES

- Almost no control over endpoints.
- May not have full visibility or arbitrary control as can only scan and analyse assets and devices once they connect to the portal.
- Possible suboptimal communications by forcing traffic through the portal.
- Potentially susceptible to DDoS attacks
- May not be able to continuously monitor devices for malware and appropriate configuration.
- Limited information can be inferred from devices requesting access.

THE CYBERCX SOLUTION

The making of a future-state Zero Trust organisation

Zero Trust is not a singular product, rather it is an architected solution. Just as organisations are different, so is their ideal Zero Trust model. This is because the ultimate solution must be adapted to the complexity of the environment, while effectively protecting people, devices, applications and data.

At CyberCX we design hybrid Zero Trust frameworks that consolidate integrated, real-world capabilities - with an eye to delivering the greatest value with respect to cost and granularity to organisations on their Zero Trust journey. We simplify the challenging dependency mapping and policy creation steps of migrating to Zero Trust, ultimately developing a solution that is scalable across any environment.

Our solutions cater to a 'typical' enterprise, drawing the right components from each of the outlined solutions with the following characteristics:

- Hybrid multi-cloud (utilising segmentation with NGFW and agent based micro-segmentation)
- Remote sites & users (utilising SASE)
- IoT (utilising Network Access Control)
- SaaS (utilising SASE)

The solution would work as follows:

1. Deploy agent based micro-segmentation to secure ALL communications (east-west, and north-south) with whatever capabilities are available on the host firewalls. Telemetry gathered from this deployment will also give you the dependency mapping needed for the introduction of NGFWs in the environment. This solution can be deployed on both physical and cloud infrastructure, including containers.
2. Deploy NGFWs for broad segmentation of the network and regulation of north-south traffic, or traffic control that requires next-gen capabilities. Leverage the dependency maps from the prior micro-segmentation to develop policies on the NGFWs. Let them take over north-south traffic regulation (and potentially some east-west).
3. Deploy SASE to secure access to internal networks for both individual users and facilities.
4. Deploy NAC to regulate access to the edge of the network for users and devices, including IoT.

The above solution can be vastly enriched by the addition of other functions like SIEM, IAM, etc.

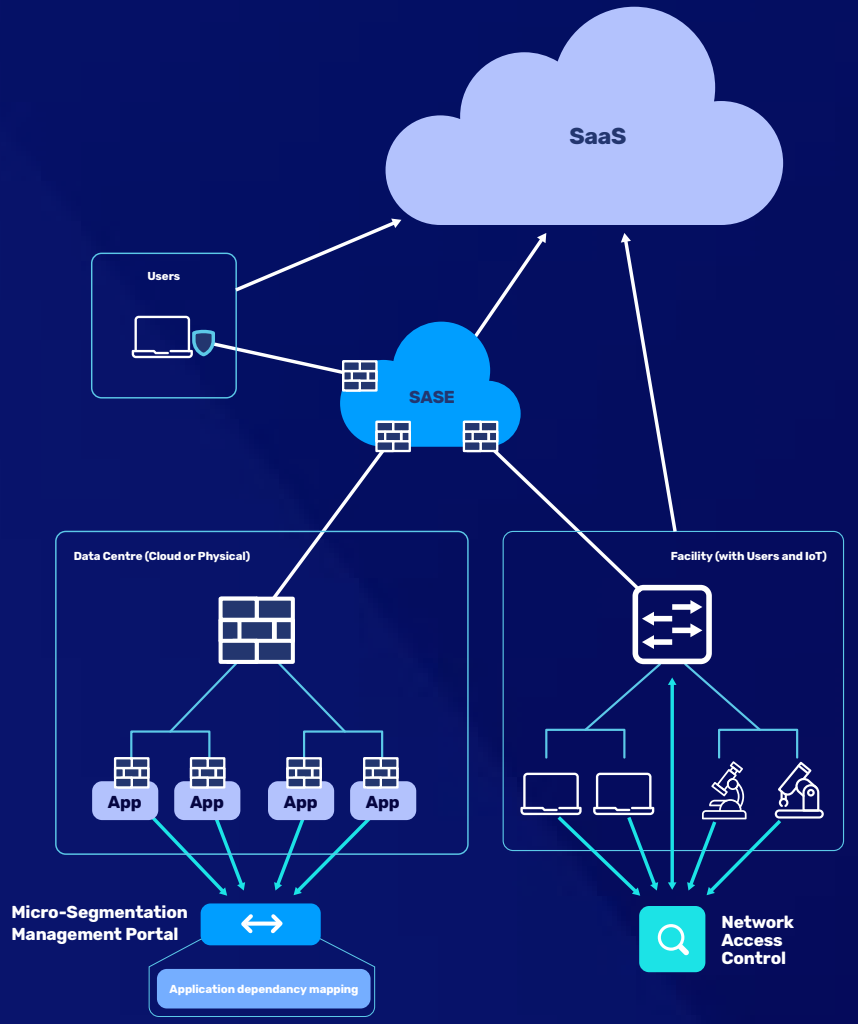

As your organisation evolves, so does Zero Trust

With the foundations of Zero Trust in place and confidence gained, continuous monitoring and assessment of the network to enhance visibility and gather insight facilitates actions to further elevate trust with extensions. Recent changes at your organisation - such as new device rollouts or software upgrades - may result in workflows or policy updates too.

As you continue to make decisions around Zero Trust, Forrester's Zero Trust eXtended[3] framework allows you to map technology purchases to the continued implementation of your Zero Trust system.

As the following excerpts from the NIST Zero Trust Architecture[4] paper demonstrate, there are also numerous logical components in the Zero Trust model that further enhance your ecosystem such as:

ID Management System (IAM/IDM)
This is responsible for creating, storing and managing enterprise user accounts and identity records (e.g. lightweight directory access protocol server). This system contains the necessary user information (e.g. name, email address, certificates) and other enterprise characteristics such as role, access attributes and assigned assets. This system often utilises other systems (such as PKI) for artefacts associated with user accounts. This system may be part of a larger federated community and may include non-enterprise employees or links to non-enterprise assets for collaboration.

Threat Intelligence Feeds
This provides information from internal or external sources that help the policy engine make access decisions. These could be multiple services that take data from internal and/or multiple external sources and provide information about newly discovered attacks or vulnerabilities. This also includes blacklists, newly identified malware and reported attacks to other assets that your organisation will want to deny access to from enterprise assets.

Continuous Diagnostics & Mitigation (CDM) System/ Industry Compliance System
This gathers information about the enterprise asset's current state and applies updates to configuration and software components. An enterprise CDM system provides the policy engine with the information about the asset making an access request, such as whether it is running the appropriate patched operating system and applications or whether the asset has any known vulnerabilities. The industry compliance system ensures that the enterprise remains compliant with any regulatory regime that it may fall under (e.g. FISMA, healthcare or financial industry information security requirements). This includes all the policy rules that an enterprise develops to ensure compliance.

Network and System Activity Logs / SIEM System
This is the enterprise system that aggregates asset logs, network traffic, resource access actions and other events that provide real-time (or near-real-time) feedback on the security posture of enterprise information systems. The Security Information and Event Management (SIEM) system collects security centric information for later analysis. This data is then used to refine policies and warn of possible attacks against enterprise assets.

[3] The Zero Trust eXtended (ZTX) Ecosystem
[4] Draft (2nd) NIST Special Publication 800-207 Zero Trust Architecture

Identity Access Management (IAM) is a foundational aspect of Zero Trust and should be at the core of any security strategy.

Policies

Further customise the authentication experience with solutions such as Okta and Duo. This allows you to enrich authentication policies and processes and better accommodate every use case, no matter the complexity.

Privileges

Eradicate weaknesses and extend privilege management with advanced, fine-grained access that better conforms to your unique access control standards.

TRUSTED PARTNERS

Take the Zero Trust journey with CyberCX

The CyberCX group unites our country's most trusted cyber security companies to deliver a comprehensive cyber security capability for Australian enterprises and governments.

With a workforce of 600+ cyber security professionals, a footprint of more than 20 offices across Australia and New Zealand, and a global presence in Europe and the US, CyberCX offers the ultimate end-to-end cyber security service.

CyberCX

National Headquarters
Level 4, 330 Collins Street,
Melbourne, 3000

1300 031 274
info@cybercx.com.au

cybercx.com.au