THE ZERO TRUST DICTIONARY

- Key concepts
- Design principles
- Steps to implementation
- Zero Trust terminology
- Maturity model

on2IT
ZERO TRUST INNOVATORS

MANAGEMENT SUMMARY

Zero Trust is a powerful concept, but the recent hype surrounding it has led to numerous interpretations. Agreeing to a term set that defines the concept will greatly improve the ease with which we can then implement this Zero Trust strategy. This is why we'd like to introduce the Zero Trust dictionary, an authorative lexicon with definitions and terminology defined by John Kindervag, the Creator of Zero Trust.

TABLE OF CONTENTS

ZERO TRUST

Zero Trust is a strategic initiative that helps prevent successful data breaches by eliminating digital trust from your organization. Rooted in the principle of "never trust, always verify," Zero Trust is designed as a strategy that will resonate with the highest levels of any organization, yet can be tactically deployed using off-the-shelf technology. Zero Trust strategy is decoupled from technology, so while technologies will improve and change over time, the strategy remains the same.

ZERO TRUST ENVIRONMENT

A Zero Trust environment designates the location of your Zero Trust architecture, consisting of a single protect surface containing a single DAAS element. Zero Trust environments are places where Zero Trust controls and policies are deployed. These environments include traditional on-premise networks such as data centers, public clouds, private clouds, on endpoints, or across an SD-WAN.

ZERO TRUST ARCHITECTURE

Your Zero Trust architecture is the compilation of the tools and technologies used to deploy and build your Zero Trust environment. This technology is fully dependent upon the protect surface you are protecting, as Zero Trust is designed from the inside out, starting at the protect surface and moving outwards from there.

Typically, the protect surface will be protected by a Layer 7 segmentation gateway that creates a microperimeter that enforces Layer 7 controls with Kipling Method Policy. Every Zero Trust architecture is tailor made for an individual protect surface.

ZERO TRUST DESIGN PRINCIPLES

There are four design principles of Zero Trust:

DEFINE BUSINESS OUTCOMES

Ask the question “What is the business trying to achieve?” This aligns Zero Trust to the grand strategic outcomes of the organization and makes cybersecurity a business enabler instead of the business inhibitor that it is often seen as today.

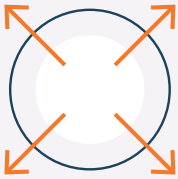

DESIGN FROM THE INSIDE OUT

Start with the DAAS elements and the protect surfaces that need protection and design outward from there.

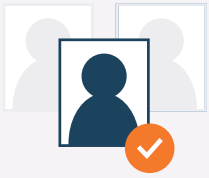

DETERMINE WHO OR WHAT NEEDS ACCESS

Determine who needs to have access to a resource in order to get their job done. Known as least privilege, it is very common to give too many users too much access to sensitive data for no business reason.

INSPECT AND LOG ALL TRAFFIC

All traffic going to and from a protect surface must be inspected and logged for malicious content and unauthorized activity, up through Layer 7.

DATA, APPLICATIONS, ASSETS AND SERVICES (DAAS)

DAAS is an acronym that stands for Data, Applications, Assets, and Services, which define the sensitive resources that should go into individual protect surfaces. DAAS elements include:

DATA

This is sensitive data that can get an organization in trouble if it is exfiltrated or misused. Examples of sensitive data include payment card information (PCI), protected health information (PHI), personally identifiable information (PII), and intellectual property (IP)

APPLICATIONS

Typically these are applications that use sensitive data or control critical assets.

ASSETS

Assets could include IT (information technology), OT (operational technology), or IoT (Internet of Things) devices such as point-of-sale terminals, SCADA controls, manufacturing systems, and networked medical devices.

SERVICES

These are fragile, sensitive business critical services. The most common services that should be protected in a Zero Trust manner include DNS, DHCP, Active Directory®, and NTP.

PROTECT SURFACE

The protect surface is the inversion of the attack surface which is massive and includes the entire internet. Using a Zero Trust strategy, the overall attack surfaces can be reduced by order of magnitude to something very small and easily known.

Each protect surface contains a single DAAS element. Each Zero Trust environment will have multiple protect surfaces.

SEGMENTATION GATEWAY

A segmentation gateway (SG) is a Layer 7 gateway designed to segment networks based upon users, applications, and data. Segmentation gateways are the primary technology used to enforce Layer 7 policy in Zero Trust environments.

Segmentation gateways can be physical (PSG) when used in traditional on-premise networks, or virtual (VSG) when used in public or private clouds. Next-Generation firewalls traditionally function as segmentation gateways when they are deployed in Zero Trust environments.

MICRO-PERIMETER

When a segmentation gateway connects to a protect surface and a Layer 7 Kipling Method Policy is deployed, then a microperimeter is placed around the protect surface.

The microperimeter ensures only known approved and validated traffic have access to the protect surface, based upon policy. One architectural principle of Zero Trust is to move your SG as close as possible to the protect surface for the most effective preventative controls enforced by the microperimeter.

MICRO-SEGMENTATION

Microsegmentation is the act of creating a small segment in a network so that attackers have difficulty moving around and accessing internal resources.

Many networks are “flat,” meaning that there are no internal segments, so if an attacker gets a foothold in the network, they can move around unnoticed to attack resources and steal data.

A microperimeter is a type of microsegment. The microperimeter defines a Layer 7 boundary for protections of a DAAS element. Some organizations may choose to use Layer 3 microsegmentation technology inside of a microperimeter.

ASSERTED IDENTITY

Identity is always an assertion of the abstraction of a user on a network. The identity system "asserts" that a device is generating packets under the control of the asserted identity. The asserted identity is the validated and authenticated 'Who' statement that is part of the Kipling Method Policy assertion: 'Who' should have access to a resource?

LEAST-PRIVILEGED ACCESS

Least-privileged access asks the question "Does a user need to have access to a specific resource to get their job done?" We give too much access to most users based upon the broken trust model.

By mandating a least-privilege, or need-to-know, policy, the ability of a user to perform malicious actions on a resource is severely limited. This mitigates against both stolen credential and insider attacks.

GRANULAR ACCESS CONTROL

Granular access control is the outcome of an explicitly defined Zero Trust Kipling Method policy statement. Multiple access control criteria provide fine-grained policy for access to a protect surface, making it substantially more difficult to perform a successful attack against that protect surface.

TRUST LEVELS

The existing cybersecurity paradigm is based upon a broken trust model where all systems external to the corporate networks are considered "untrusted" and those inside the corporate networks are known as "trusted."

It is this flaw that undergirds Zero Trust.

Trust is a human emotion injected into digital systems for no technical reason. It is not measurable. Trust is binary. All successful cyberattacks exploit trust in some manner, making trust a dangerous vulnerability that must be mitigated.

In the Zero Trust arena, all packets are untrusted, and are treated exactly the same as every other packet flowing across the system. The trust level is defined as zero, hence the term Zero Trust.

DATA TOXICITY

Data toxicity is the doctrine that defines sensitive data as "toxic" to your organization if it has been stolen or exfiltrated from your networks or systems and is in control of malicious actors.

This exfiltration leads to a negative impact on the business. The data has become toxic as its theft leads to lawsuits or regulatory action on the organization.

Every organization has both non-toxic and toxic data.

An easy way to recognize toxic data types is to remember the 4Ps of toxic data: PCI (credit card data), PII (personally identifiable information), PHI (patient health information), and IP (intellectual property). Most toxic data falls into these simple categories.

THE FIVE STEPS TO IMPLEMENTING ZERO TRUST:

1. DEFINE THE PROTECT SURFACE

Identify the DAAS elements: Data, Applications, Assets, and Services, that you want to protect.

2. MAP THE TRANSACTION FLOWS

Zero Trust is a system, and in order to secure the system, understanding how the network works is imperative to a successful Zero Trust deployment. The mapping of the transactions flows to and from the protect surface shows how various DAAS components interact with other resources on your network and, therefore, where to place the proper controls. The way traffic moves across the network, specific to the data in the protect surface, determines the design.

3. BUILD A ZERO TRUST ARCHITECTURE

Part of the magic of the five-step model is that the first two steps will illuminate the best way to design the Zero Trust architecture. The architectural elements cannot be predetermined. Each Zero Trust environment is tailor-made for each protect surface. A good rule-of-thumb in design is to place the controls as close as possible to the protect surface.

4. CREATE ZERO TRUST POLICY

Ultimately, we need to instantiate Zero Trust as a Layer 7 policy statement. Therefore, it requires Layer 7 controls. Use the Kipling Method of Zero Trust policy writing to determine who or what can access your protect surface.

5. MONITOR AND MAINTAIN THE NETWORK

One of the design principles of Zero Trust is to inspect and log all traffic, all the way through Layer 7. The telemetry provided by this process will not just help prevent data breaches and other significant cybersecurity events, but will provide valuable security improvement insights. This means that each protect surface can become more robust and better protected over time. Telemetry from cloud, network, and endpoint controls can be analyzed using advances in behavioral analytics, machine learning, and artificial intelligence to stop attacks in real-time and improve security posture over the long term.

KIPLING METHOD POLICY (KMP)

Zero Trust policy is known as The Kipling Method, named after the writer Rudyard Kipling who gave the world the idea of Who, What, When, Where, Why and How (WWWWWH) in a poem in 1902.

Since WWWWWH is well known worldwide, it crosses languages and cultures and allows easily created, easily understood, and easily auditable Zero Trust policy statements for various technologies. A KMP determines what traffic can transit the Microperimeter at any point in time, preventing unauthorized access to your protect surface, while preventing the exfiltration of sensitive data into the hands of malicious actors.

True Zero Trust requires Layer 7 technology to be fully effective. The Kipling Method describes a Layer 7 Zero Trust granular policy.

Using the Kipling Method, you can create Zero Trust policy effortlessly by answering the following questions:

WHO should be allowed to access a resource? The validated "asserted identity" will be defined in the 'Who' statement. This replaces the source IP Address in a traditional firewall rule.

WHAT application is the asserted identity allowed to use to access the resource? In almost all cases, protect surfaces are accessed via an application. The application traffic should be validated at Layer 7 to keep attackers from impersonating the application at the port and protocol level and using the rule maliciously.

The 'What' statement replaces port and protocol designations found in traditional firewall rules.

WHEN defines a timeframe. 'When' is the asserted identity allowed to access the resource? It is common for rules to be instantiated 24/7, but many rules should be time limited and turned off when authorized users are not typically using the rule.

Attackers take advantage of these always on rules and attack when approved users are away from the system, making the attacks more difficult to discover.

KIPLING METHOD POLICY (KMP)

WHERE is the resource located? The location of the protect surface could be anywhere data is stored or assets are deployed. The 'Where' statement replaces the destination IP Address in a traditional firewall

WHY is the user ('Who' statement) allowed to access the resource? In most instances, the reason for putting data or an asset into a protect surface is because of its sensitivity. The sensitivity may be defined by a compliance mandate or by a business driver.

There are often ways of tagging a packet to identify those sensitive data or systems. This tagging creates metadata that various controls can use to inform or automate policy statements. This defines the 'Why' statement in the policy.

HOW is the tuple that defines the criteria used to allow the asserted 'Who' statement to access a resource. It answers the question "How should the traffic be processed as it accesses a resource?"

These criteria often apply additional controls or inspection on the packet as it accesses the resource. Controls that once were separate products deployed individually are now delivered as a service. These advanced services can be applied to individual rules as needed.

These advanced controls include IPS, DLP, sandboxing, decryption, and other features that are available on an individual control.

THE ZERO TRUST MATURITY MODEL

STEPS	INITIAL (1) The initiative is undocumented and performed on an ad hoc basis, with processes undefined. Success depends on individual efforts	REPEATABLE (2) The process is documented and is predictably repeatable, using lessons learned in the initial phase.	DEFINED (3) Processes for success have been defined and documented.	MANAGED (4) Processes are monitored and controlled. Efficacy is measurable.	OPTIMIZED (5) Focus is on continuous optimization.
1. DEFINE YOUR PROTECT SURFACE Determine which DAAS element will be protected inside the defined protect surface.	The DAAS element is unknown or discovered manually. Data classification is not done or is incomplete.	The use of automated tools to discover and classify DAAS elements has begun but is not standardized.	Data classification training and processes have been introduced and are maturing. Protect surface discovery is becoming automated.	New or updated DAAS elements are immediately discovered, and classified as assigned to the correct protect surface in an automated manner.	Discovery and classification processes are fully automated.
2. MAP THE TRANSACTION FLOW Mapping the transaction flows to and from the protect surface shows how various DAAS components interact with other resources on your network and, therefore, where to place the proper controls.	Flows are conceptualized based on interviews and workshops.	Traditional scanning tools and event logs are used to construct approximate flow maps.	A flow mapping process is in place. Automated tools are beginning to be deployed.	Automated tools create precise flow maps. All flow maps are validated with system owners.	Transaction flows are automatically mapped across all locations in real time
3. ARCHITECT A ZERO TRUST ENVIRONMENT A Zero Trust architecture is designed, based upon the protect surface and the interaction of resources based on the flow maps.	With little visibility and an undefined protect surface, the architecture cannot be properly designed.	Protect surface is established based on current resources and priorities.	The basics of the protect surface enforcement are complete, including placing segmentation gateways in the appropriate places.	Additional controls are added to evaluate multiple variables (e.g., endpoint controls, SAAS and API controls).	Controls are enforced using a combination of hardware and software capabilities.
4. CREATE ZERO TRUST POLICY Create Zero Trust policy following the Kipling Method of Who, What, When, Where, Why and How.	Policy is written at Layer 3	Additional 'Who' statements are starting to be identified to address business needs; User IDs of applications and resources are known, but access rights are unknown.	The team works with the business to determine who or what should have access to the protect surface.	Custom user-specific elements are created and defined by policy, reducing policy space and number of users with access.	Layer 7 policy is written for granular enforcement. Only know traffic and legitimate application communication is allowed.
5. MONITOR AND MAINTAIN Telemetry from all controls in the protection chain are captured, analyzed and used to stop attacks in real-time and enhance defenses to create more robust protections over time.	Visibility into what is happening on the network is low.	Traditional SIEM or log repositories are available, but the process is still mostly manual.	Telemetry is gathered from all controls and is sent to a central data lake.	Machine learning tools are applied to the data lake for context about how traffic is used in the environment.	Data is incorporated from multiple sources and used to refine steps 1-4. Alerts and analysis are automated.

ON2IT: EMBRACE ZERO TRUST WITH EASE

Easily embrace the protective power of Zero Trust security through ON2IT's Zero Trust-as-a-Service managed approach. ON2IT applies true Zero Trust strategies in 100% of implementations through proprietary technology. Combining the innovative AUXO ™ platform with Zero Trust security deployed across policy, architecture, and operations, ON2IT effectively safeguards your networks with proven Zero Trust strategies in a convenient and holistic managed services model.

Certified 24x7, Zero Trust-Based Managed SOC

ON2IT's 24x7 mSOC™ —powered by AUXO™ —provides a dedicated team of seasoned cybersecurity analysts to continuously monitor your network, take proactive measures, and calculate and respond in real-time to threats. Leverage the experience, efficiency, and latest tooling ON2IT's Zero Trust experts offer 24x7.

Original Zero Trust & MSSP Innovators

ON2IT has been modernizing managed services by developing true Zero Trust- and SOC-as-a-Service since 2005. Home to the original inventor of Zero Trust, ON2IT provides every partner and client exclusive 1:1 consulting with John Kindervag, SVP of Cybersecurity Strategy at ON2IT, alongside a team of innovative and specialized cybersecurity Zero Trust experts.

AUXO™: Proprietary Zero Trust Platform

Created by ON2IT, the AUXO™ Zero Trust platform is the first to offer the best of SIEM, SOAR, and MDR all in 1 automated managed service platform. ON2IT uses advanced automation to monitor networks, manage alerts, and guarantee proactive measures are taken from implementation to deployment and beyond.

The Only MSSP with Zero Trust as a Service

Zero Trust as a Service includes:

- A transparent view of all protect surfaces and their DAAS elements, traffic patterns, preventive measures and threats per protect surface
- Active management of all controls per protect surface
- Zero Trust Fitness: a real-time view of the state of your cybersecurity

This is a publication of ON2IT

ON2IT

ZERO TRUST INNOVATORS

5717 Legacy Drive, Suite 250, Plano, TX 75024, USA
Email: sales@on2it.net - Phone: (+1) (214) 206 8446

Hogeweg 35, 5301 LJ Zaltbommel, The Netherlands
Email: sales@on2it.net - Phone: (+31) 088 22 66 200

FOR MORE INFORMATION

www.on2it.net