AN AUTHENTIC ZERO TRUST GUIDE

- AN INTRODUCTION TO AUTHENTIC ZERO TRUST BY JOHN KINDERVAG
- AUTHENTIC ZERO TRUST KEY CONCEPTS
- ZERO TRUST TERMS EXPLAINED
- ZERO TRUST MATURITY MODEL

INDEX

AN INTRODUCTION TO ZERO TRUST

by ON2IT Senior Vice President John Kindervag

JOHN KINDERVAG
@Kindervag

John Kindervag is founder of Zero Trust SVP Cybersecurity Strategy and Group Fellow ON2IT.

When I worked as a security analyst, I became fascinated by how people and businesses anthropomorphized their digital environments by applying the concept of trust to computing—that somehow a device could be trusted and that it cared that it was trusted. Back then, many CISOs and CIOs adhered to the idea that what's inside the corporate firewall can be trusted.

This concept of inside versus outside became a variable that was used to determine security policy, with many organizations operating under the adage "trust, but verify." In the trust-but-verify model, trust is the default. When identity is verified, trust is assumed and access is granted.

But trust applies only to people—not digital environments. Identity credentials can be stolen, networks can be hacked, and insiders with bad intent are often in positions of trust. This means it's impossible to know with certainty that the originator of network traffic can truly be trusted: An asserted identity is only an assertion, not an actual person.

In response to what CISOs and CIOs told me about their cybersecurity strategies, I created the concept of Zero Trust, which is framed around the principle that no network user, packet, interface, or device—whether internal or external to the network—should be trusted. Some people mistakenly think Zero Trust is about making a system trusted, but it really involves eliminating the concept of trust from cybersecurity strategy. By doing this, every user, packet, network interface, and device is granted the same default trust level: zero.

'TRUST APPLIES ONLY TO PEOPLE — NOT DIGITAL ENVIRONMENTS.'

Zero Trust should be thought of as a strategy or framework. It requires companies to rethink their philosophy and approach to trusted network users and devices. Zero Trust is not a product, although Zero Trust-based security infrastructures can be implemented by using many different products. Nor does Zero Trust require organizations to rip and replace existing security infrastructure—rather, it leverages existing technology to support the Zero Trust mindset, with new tools added as needed.

The hallmark of Zero Trust is simplicity. When every user, packet, network interface, and device is untrusted, protecting assets becomes simple. To reduce the complexity of cybersecurity environments, organizations can prioritize security technologies and tools that support simplicity by automating repetitive and manual tasks, integrating and managing multiple security tools and systems, and autoremediating known vulnerabilities.

Zero Trust is a journey best taken one step at a time. I recommend that organizations begin by prioritizing the smallest possible protect surfaces—a single data set, asset, application, or service—depending on the level of sensitivity or business criticality. Then, they can create a microperimeter around each protect surface and granularly control the traffic allowed into the perimeter.

I encourage security teams to learn and practice on less sensitive protect surfaces, moving to protect increasingly more sensitive or valuable ones as they fine-tune their approaches and their confidence increases. Over time and with lots of practice, they'll be ready to migrate their most critical assets to the Zero Trust environment. Finally, when high-value assets are protected, teams can focus on less important assets. And by continuing to maintain a Zero Trust mindset, organizations can protect themselves even as security technologies and tools evolve.

'ZERO TRUST IS A JOURNEY BEST TAKEN ONE STEP AT A TIME..'

AUTHENTIC ZERO TRUST KEY CONCEPTS

Zero Trust is a powerful concept, but the recent hype surrounding it has led to numerous interpretations. Some vendors have redefined the meaning of Zero Trust to fit the limitations of their products. Though there are products that work well in Zero Trust environments, if a vendor tries to sell you their ‘Zero Trust’ product, that’s a clear indicator that they don’t understand the Zero Trust concept

ZERO TRUST

Zero Trust is a strategic initiative that prevents successful data breaches by eliminating the need for digital trust in your organization. Rooted in the principle of 'never trust, always verify', Zero Trust is designed as a strategy that resonates with the highest levels of any organization yet is tactically deployed using off-the-shelf technology. Zero Trust strategy is decoupled from technology. While technologies 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 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 combines the tools and technologies used to deploy and build your Zero Trust environment. This technology depends entirely on the Protect Surface you are protecting. This is because Zero Trust is designed from the inside out, starting at the Protect Surface and moving outwards. Typically, the protect surface will be protected by a Layer 7 segmentation gateway that creates a microperimeter to enforce Layer 7 controls with the Kipling Method policy. Every Zero Trust architecture is tailor-made for an individual protect surface.

ZERO TRUST DESIGN PRINCIPLES

DEFINE BUSINESS OUTCOMES

Ask, “What is the business trying to achieve?” The answer aligns Zero Trust to the Grand Strategic outcomes of the organization. It makes cybersecurity a business enabler instead of the business inhibitor it’s often regarded as.

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 to get their job done. Known as Least Privilege, it is 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.

THE 5 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. To secure the system, understanding how the network works is imperative to a successful Zero Trust deployment. 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. 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 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, create Zero Trust as a Layer 7 Policy Statement. This 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 Zero Trust design principles is to inspect and log all traffic, all the way through Layer 7. The telemetry provided by this process prevents data breaches and other significant cybersecurity events and provides 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 in the long term.

DATA, APPLICATIONS, ASSETS AND SERVICES (DAAS)

DAAS stands for Data, Applications, Assets, and Services. These define the sensitive resources that should go into individual Protect Surfaces.
DAAS elements include:

DATA

This is sensitive data that causes problems for an organization 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

These are typically 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 the highly fragile services on which your business depends. The most common services to protect in a Zero Trust manner include DNS, DHCP, Active Directory®, and NTP.

ZERO TRUST TERMS EXPLAINED

So, what is Authentic Zero Trust? The definitions in the following pages are based on the original research by John Kindervag and the ongoing evolution of Zero Trust he and ON2IT have been and still are working on.

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 reduces the overall attack surfaces in orders of magnitude to something tiny and easily known. Each Protect Surface contains a single DAAS element. Each Zero Trust environment has multiple Protect Surfaces.

SEGMENTATION GATEWAY

A segmentation gateway (SG) is a Layer 7 gateway designed to segment networks based on users, applications, and data. Segmentation Gateways are the primary technology for enforcing Layer 7 policy in Zero Trust Environments. Segmentation Gateways can be Physical (PSG) when used in traditional on-premises networks or Virtual (VSG) when used in public or private clouds. Next-Generation Firewalls traditionally function as Segmentation Gateways when deployed in Zero Trust Environments.

MICROPERIMETER

When a Segmentation Gateway connects to a Protect Surface and a Layer 7 Kipling Method Policy is deployed, a Microperimeter is placed around the protect surface. The Microperimeter ensures that only known, approved and validated traffic has access to the protect surface, based on policy. One Zero Trust architectural principle 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 creates 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 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 generates 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, “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. Mandating a least-privilege, or need-to-know policy, severely limits a user’s ability to perform malicious actions. This mitigates against both stolen credentials and insider attacks.

GRANULAR ACCESS

Granular access control is the outcome of an explicitly defined Zero Trust Kipling Method policy statement. Multiple access control criteria provide a fine-grained policy for access to a Protect Surface, making it substantially more challenging 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 outside the corporate networks are considered 'untrusted'. Those inside the corporate networks are 'trusted'. This flaw underpins Zero Trust. Trust is a human emotion introduced into digital systems for no technical reason. It is not measurable. Trust is binary. All successful cyberattacks exploit trust in some manner. This makes trust a dangerous vulnerability that must be mitigated. With Zero Trust, all packets are untrusted and treated identically with 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 sensitive data becomes 'toxic' to your organization if it has been stolen or exfiltrated from your networks or systems into the control of malicious actors. This exfiltration harms the business. The data becomes toxic because its theft leads to lawsuits or regulatory action against 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 this simple framework.

1/2

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 in a poem in 1902. Since the idea of WWWWH is well known worldwide, it transcends languages and cultures and allows easily created, easily understood, and easily auditable Zero Trust policy statements for various technologies. A KMP determines the traffic that 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.

The Kipling Method enables you to create Zero Trust policy effortlessly by answering the following questions:

WHO

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

What application is the asserted identity allowed to use for access to 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 in traditional firewall rules.

WHEN

When defines a timeframe. When is the asserted identity allowed to access the resource? It is common to create rules 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. This makes the attacks more difficult to discover.

2/2

KIPLING METHOD POLICY (KMP)

WHERE

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 rule.

WHY

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

How 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 available on an individual control.

ZERO TRUST MATURITY MODEL

Because Zero Trust is a strategic initiative, it’s important to benchmark your Zero Trust journey and measure your maturity over time. The maturity model records improvements to your individual Zero Trust environments. Designed using a standard Capability Maturity Model, the Zero Trust Maturity Model leverages the 5-step methodology for implementing Zero Trust and should be used to measure the maturity of an individual protect surface containing a single DAAS element.

ZERO TRUST MATURITY MODEL

STEPS	INITIAL (1)	REPEATABLE (2)	DEFINED (3)	MANAGED (4)	OPTIMIZED (5)
	The initiative is undocumented and performed on an ad hoc basis, with processes undefined. Success depends on individual efforts.	The process is documented and is predictably repeatable, using lessons learned in the initial phase.	Processes for success have been defined and documented.	Processes are monitored and controlled. Efficacy is measurable.	Focus is on continuous optimization.
1. DEFINE YOUR PROTECT SURFACE Determine which single 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 FLOWS Mapping 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.	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 known 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.

ZERO TRUST INNOVATORS

This is a publication of ON2IT B.V.

ON2IT B.V.
Hogeweg 35
5301 LJ Zaltbommel
The Netherlands

on2it.net