CyberEd.io

The Zero Trust Dictionary

A Set of Definitions to Guide Enterprise Strategies

Written by Steve King & John Kindervag

The Zero Trust Dictionary

Because there has been so much word salad thrown about these days around Zero Trust, we look to John Kindervag to provide the proper definitions behind his Zero Trust creation, so as we move toward a strategy, we have a better chance of success if we know what we are talking about and agree to a term set that tries to define the concepts into actionable behavior.

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 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 micro-perimeter that enforces controls with Kipling Method[1] policy. Layer 7 refers to the top layer in the 7-layer OSI Model of the Internet. It is also known as the "application layer." It's the top layer of the data processing that occurs just below the surface or behind the scenes of the software applications with which users interact directly - the HTTP requests and responses used to load webpages, for example, are layer 7 events. 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 IIoT (Internet of Things) devices such as point-of-sale terminals, SCADA controls, manufacturing systems, and networked medical devices.
- **Services** – These are sensitive services that are very fragile upon your business depends. 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 orders 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.

Microperimeter

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

Microsegmentation

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 Micro-Perimeter 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 preform malicious actions 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 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, 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 a 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 Environment

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 in a poem in 1902.

Since idea of WWWWHD is well known worldwide, it crosses languages and cultures and allows easily created, easily understood, and easily auditable Zero Trust policy statements for various technology. 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

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 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

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

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 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

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 documents improvements made 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.

For more information about Zero Trust, the Zero Trust Council or the CyberTheory Institute, please go to www.cybertheory.io/cybertheory-institute/ or contact Steve King at 505-795-8855 or sking@cybertheory.io.

LET'S GET SMARTER.

Talk to a CyberEd Expert today

Visit cybered.io

@cyberedio

CyberEd.io

CyberEd.io

+1-609-356-1499 • info@cybered.io • www.cybered.io

CyberEd.io