

Definitive GuideTM to *Zero Trust Security*

Secure Your Cloud Workloads, Applications,
and Endpoints with Micro-segmentation



Jon Friedman

FOREWORD BY:
Tony Scott

Compliments of:

 **COLORTOKENS**

About ColorTokens

ColorTokens Inc., a leader in zero trust cloud security, provides a new generation of security that empowers global enterprises to protect their most important crown jewels, decrease their attack surface, and secure their enterprise from the inside out. The proactive ColorTokens platform single-handedly secures cloud workloads, dynamic applications, endpoints, and users. Through its award-winning Xtended ZeroTrust Platform, ColorTokens delivers the only cloud-delivered solution that combines micro-segmentation, workload protection, visualization, application control, endpoint protection, and zero trust network access — all while seamlessly integrating with existing security tools.

Definitive Guide[™] to *Zero Trust Security*

Secure Your Cloud Workloads,
Applications, and Endpoints
with Micro-segmentation

Jon Friedman

Foreword by Tony Scott

With contributions from Scott Emo,
Kayvon Sadeghi, and Ajay Uggirala



CYBEREDGE
GROUP

Definitive Guide™ to Zero Trust Security

Published by:

CyberEdge Group, LLC

1997 Annapolis Exchange Parkway

Suite 300

Annapolis, MD 21401

(800) 327-8711

www.cyber-edge.com

Copyright © 2020, CyberEdge Group, LLC. All rights reserved. Definitive Guide™ and the CyberEdge Press logo are trademarks of CyberEdge Group, LLC in the United States and other countries. All other trademarks and registered trademarks are the property of their respective owners.

Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of the publisher. Requests to the publisher for permission should be addressed to Permissions Department, CyberEdge Group, 1997 Annapolis Exchange Parkway, Suite 300, Annapolis, MD, 21401 or transmitted via email to info@cyber-edge.com.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on CyberEdge Group research and marketing consulting services, or to create a custom *Definitive Guide* book for your organization, contact our sales department at 800-327-8711 or info@cyber-edge.com.

ISBN: 978-1-948939-08-9 (Paperback); ISBN: 978-1-948939-09-6 (eBook)

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgements

CyberEdge Group thanks the following individuals for their respective contributions:

Editor: Susan Shuttleworth

Graphic Design: Debbi Stocco

Production Coordinator: Valerie Lowery

Table of Contents

Foreword	v
Introduction	vii
Chapters at a Glance.....	vii
Helpful Icons	viii
Chapter 1: The Promise of Zero Trust Security	1
Once Upon a Time.....	1
The Death of Trust.....	2
Principles of Zero Trust Security	3
Putting the Principles Together	5
Benefits of Zero Trust Security.....	6
In This Guide.....	8
Chapter 2: The Path to Zero Trust Security	9
One Step at a Time	9
The Steps	10
Chapter 3: Obtaining Comprehensive Visibility	13
Discover Entities	13
Classify and Group Entities.....	14
Create and Use Visual Maps.....	15
Chapter 4: Defining Zero Trust Zones and Creating Policies	17
What Is a Zero Trust Zone?.....	17
Enforcement Requires Software-defined Micro-segmentation	18
Start with a Model	19
Define the Zero Trust Zones.....	21
Create Policies	21
Chapter 5: Enforcing Zero Trust Security	23
Observe Before You Leap	23
Turn on Enforcement.....	25
Extend Zero Trust Security to Endpoints	27
Improve Application Security	28
Chapter 6: Use Cases	29
Improving Situational Awareness.....	29
Protecting the “Crown Jewels”	30
Strengthening Compliance.....	30
Environmental Separation	31
DevOps and Continuous Delivery	31
Protecting Legacy Endpoints and Operational Technologies	32

Simplifying Remote Access and (Sometimes) Eliminating VPNs.....	33
Securing Cloud Migrations.....	33
Protecting Microservices	34
Managing Third-party Risk.....	34
Chapter 7: Selecting a Zero Trust Security Platform	35
Cloud Delivery	35
Scope of Capabilities	36
Breadth of Resources Protected	37
Ease of Implementation and Management.....	38
Total Cost of Ownership.....	39
Glossary	40

Foreword



My advice to IT security organizations: start practicing zero trust security as soon as possible.

Here's how I came to that conclusion.

I joined the U.S. government in February 2015 as the Federal Chief Information Officer, overseeing government information technology programs and more than \$85 billion annually in IT spending.

In June 2015, the U.S. Office of Personnel Management (OPM) announced it had suffered a hacking incident that ultimately resulted in the theft of more than 21 million records containing personal information about government employees and job applicants.

My team and I oversaw the investigation of the data breach and the development of the Cybersecurity Sprint and Implementation Plan (CSIP) to improve the information systems security posture of the federal government. We helped federal agencies accelerate the adoption of two-factor authentication, cut back on the number of people with privileged access to systems and networks, and improve basic cybersecurity hygiene.

When we started talking about fundamental structural changes that could make a major long-term difference to cybersecurity effectiveness, zero trust security stood out.

It is a fact that most of today's technology is designed for maximum interoperability. The default design principle is "connect to everything." Nobody starts by asking "What should we connect to?" or "How do we control access to minimize the risk of data breaches and other damage?"

Zero trust security addresses these questions. It builds security on the idea that no user or system should be allowed access to a resource until a level of trust has been established, and even then, that access should be controlled on a strict "need to know" basis.

With zero trust security, government agencies and commercial enterprises get a dramatically improved cybersecurity footprint at substantially lower cost.

Full disclosure: I was so impressed by the value of zero trust security, that after leaving the government I joined the board of directors of ColorTokens, one of the vendors developing technology in this area.

I commend ColorTokens for sponsoring this guide. It explores key concepts of zero trust security, including visibility into access between resources, micro-segmentation, and access control policies that adapt to dynamic cloud environments.

There is also a chapter on use cases, which can be very helpful. Organizations need to pick their spots, jump in, and start practicing zero trust security to build up their experience and skills. Picking compelling use cases is a great way to start.

I hope you will read this guide, share it with your colleagues, and start the journey toward safer information systems through zero trust security.

Tony Scott

Chairman, The TonyScottGroup
Formerly Federal CIO, CIO at VMware, Microsoft Corporation, and Walt Disney Company, and CTO at General Motors Information Systems and Services

Introduction

At first glance, “zero trust security” sounds like a terrible idea. After all, if we have zero trust in a person or thing, there is no point in sharing, or transacting, or really communicating at all.

But what if zero trust security doesn’t mean “never trust”? What if it means starting from a position of no trust, assessing the trustworthiness of the other, and granting just enough access to get a specific job done?

What if it means gaining visibility into how every entity on a network connects to every other entity, using that information to determine which connections are authorized and which are unauthorized, and ultimately blocking the unauthorized connections?

What if you could use these capabilities to prevent data breaches, strengthen compliance, and help your organization migrate to cloud platforms?

In that case, you might have a winning recipe for bringing cybersecurity into today’s world of credential-stealing hackers and dynamic cloud environments.

The aim of this Definitive Guide™ is to introduce you to the concepts and benefits of zero trust security, with a focus on visibility into network traffic, micro-segmentation, and protection of cloud workloads, endpoints, and applications.

If you work in cybersecurity, or rely on the people who do, please read on.

Chapters at a Glance

Chapter 1, “The Promise of Zero Trust Security,” summarizes the basic concepts and benefits of zero trust security.

Chapter 2, “The Path to Zero Trust Security,” outlines the steps for deploying zero trust security.

Chapter 3, “Obtaining Comprehensive Visibility,” explains how to create and use visual maps of connections between entities on premises and in the cloud.

Chapter 4, “Defining Zero Trust Zones and Creating Policies,” describes zero trust zones and how to use micro-segmentation to control access between them.

Chapter 5, “Enforcing Zero Trust Security,” explores best practices for testing enforcement and extending zero trust security to endpoints and applications.

Chapter 6, “Use Cases,” highlights 10 high-value scenarios where zero trust security can be leveraged.

Chapter 7, “Selecting a Zero Trust Security Platform,” reviews five criteria you should use to select a zero trust security platform that fits your organization.

The Glossary provide handy definitions of key terms (*appearing in italics*) used throughout this book.

Helpful Icons



Tips provide practical advice that you can apply in your own organization.



When you see this icon, take note as the related content contains key information that you won't want to forget.



Proceed with caution because if you don't it may prove costly to you and your organization.



Content associated with this icon is more technical in nature and is intended for IT practitioners.



Want to learn more? Follow the corresponding URL to discover additional content available on the Web.

Chapter 1

The Promise of Zero Trust Security

In this chapter

- Understand why the concepts of “trusted” and “untrusted” no longer work in cybersecurity
- Review the principles of zero trust security
- Learn about the benefits of zero trust security

“Men are able to trust one another, knowing the exact degree of dishonesty they are entitled to expect.”

– Stephen Leacock

Once Upon a Time...

Once upon a time, in a faraway kingdom, everything was either “trusted” or “untrusted,” and people were happy.

Employees were trusted. They sat in corporate offices, working on desktop computers attached to the corporate network. A simple user ID and password gave them access to all the data and applications they might desire. Everyone else was untrusted: required to prove they were friends, not foes, and restricted to isolated applications.

The gallant people in the IT department stopped intruders and malware at the perimeter of the kingdom. IT teams designed *flat network* where every resource on the network could be reached through one router or switch. These were easy to manage, and security could be handled by placing a firewall and an intrusion prevention system (IPS) where the corporate

network connected to the internet. There was no need for complex rules about which people and systems could access what parts of the network.

The Death of Trust

But that was a long time ago, when creating a flat network for trusted users (that is, everyone inside the perimeter) seemed to be a workable paradigm.

That era has passed. Among other developments:

- ✓ Employees outside the perimeter expect access to virtually all information resources.
- ✓ Customers and business partners need nearly the same level of access as employees.
- ✓ Confidential data and software applications are now scattered across cloud platforms, corporate data centers, and endpoints.
- ✓ Threat actors have proven that they can always find ways to steal (or buy) user credentials and evade perimeter security defenses.

What are the implications of these changes?

- ✓ You can't implicitly trust anyone (or anything) on the network.
- ✓ There is no single perimeter, in the sense of a defensible boundary around an organization's information resources.
- ✓ Flat networks are no longer acceptable; organizations can't afford to allow threat actors to move laterally across their entire IT infrastructure.

But all is not lost. Industry analysts and cybersecurity experts are gathering under the banner of “*zero trust security*,” a set of principles and practices about how to establish trust between entities and allow granular, secure access to information resources.

Principles of Zero Trust Security

All network traffic is untrusted

A core principle of zero trust security is that all network traffic is untrusted. The fact that some traffic may originate on the corporate network, or even in a highly secure segment of the network, does not entitle it to implicit trust. Because all network traffic is untrusted, all of it should be inspected and logged.

Micro-segmentation should be applied

All networks should be finely segmented and access control policies should be enforced between segments. This is known as *micro-segmentation*. Without this very granular segmentation, attackers who have acquired user credentials or compromised a system can roam freely across the entire infrastructure. So can malicious insiders.

As we will discuss, micro-segmentation can be applied not just to networks, but also to applications, and even individual servers, devices, endpoints, and workloads.

All entities are low trust

There should be no implicit trust between *any* entities in the IT environment. That includes users of all kinds, workloads (whether on physical servers or virtual machines or in containers, in corporate data centers, or on cloud platforms), network and security devices, and endpoints such as servers, laptops, mobile devices, kiosks, and ATMs. Connections between these entities should be allowed only after their identities have been verified and a level of trust has been established.

“Zero trust” does not mean “no trust”

The term “zero trust security” does not mean “no trust” or “never trust.” Entities start from a position of zero trust, but when they are

assessed a level of trustworthiness (or riskiness) can be established, and they can be given appropriate access to enterprise resources.

Trust should be assessed dynamically

The trustworthiness of each entity should be assessed and reassessed dynamically, based on all available information about the entity and the situation. For example:

- ✓ A user might be assessed initially based on credentials and information available at logon (such as the device used and its location), then reassessed based on behaviors, such as requests to access applications and volume of content downloaded.
- ✓ A workload might be assessed initially based on its function (say, web front end of an ecommerce application) and location (Amazon Web Services), then reassessed based on activities such as connections to workloads in the corporate data center.

Trust should always be assessed in the same way

A given entity should always be assessed using the same criteria. If using a personal mobile device represents a risk, that fact should carry the same weight whether the user is in the headquarters office or at an airport. The final result of the assessment may be different (the use of a public WiFi network at the airport increases risk), but the same tests are applied. To put it another way, if a personal device is risky, it should not be given a free pass just because it connects inside the headquarters office. This approach of providing a uniform, dynamic set of access criteria is behind the concept of “*Zero Trust Network Access*” (ZTNA).

The principle of least privilege access should be applied

The *principle of least privilege* says that entities should be given access only to those resources they need to perform their intended functions. In the context of zero trust security, that means once a level of trust has been established for an entity, it should be granted an appropriate amount of access for that level, but no more.

Putting the Principles Together

When you put the principles of zero trust security together, they add up to a world where:

1. Networks are finely segmented rather than flat.
2. All entities in the IT environment start at zero trust but gain trust through ongoing assessments.
3. Connections between any pair of entities are managed according to access control policies based on trust levels and the principle of least privilege.

Figure 1-1 illustrates the impact of moving to zero trust security. Instead of one or a few very large zones where many entities can freely access each other, the environment features many smaller zones. In addition, each zone can monitor and restrict access from other zones and entities.

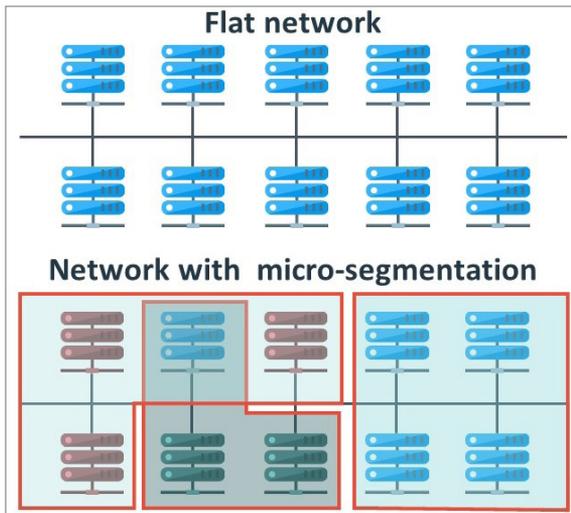


Figure 1-1: From a flat network to zero trust security.

Instead of having one ineffective perimeter at the edge of the corporate network, there are many perimeters around small zones, and even individual entities, including workloads on cloud platforms and remote endpoints.

Benefits of Zero Trust Security

The benefits of zero trust security practices extend beyond security to business issues like enabling cloud transformations and complying with regulations.

Increase visibility into lateral movement by attackers

In most advanced cyberattacks, attackers gain a foothold on one system on the target organization’s network, then move laterally to find and extract confidential data.

Figure 1-2 shows how a zero trust security platform can provide visibility into network traffic between systems and workloads. By observing suspicious connections (or using analytics tools to find them), security teams can detect and respond to advanced attacks quickly, in many cases before any damage is done.

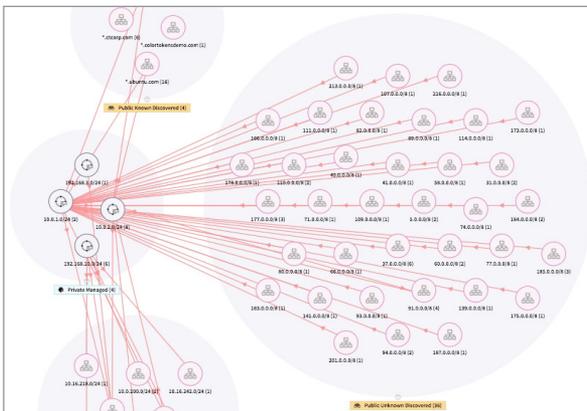


Figure 1-2: A zero trust security platform provides visibility into suspicious connections between entities in data centers and on cloud platforms. (Source: ColorTokens)

Prevent data breaches by isolating high-value assets

All organizations have applications, databases, and repositories that store or process confidential, high-value data such as credit card and financial account information, customer and

employee data, and intellectual property (IP) such as software, engineering designs, and business plans.

By implementing zero trust security, IT teams can restrict access to these “crown jewels” to users with a real need to know and to software services, computing platforms, and devices that have a genuine “need to access.”

For example, suppose an organization’s human resources (HR) system contains confidential information about employees. The HR system should not be accessed by users or systems unless they are explicitly authorized to view this data. Communications from unexpected sources, say the customer relationship management (CRM) application (as illustrated in Figure 1-3), should be investigated.

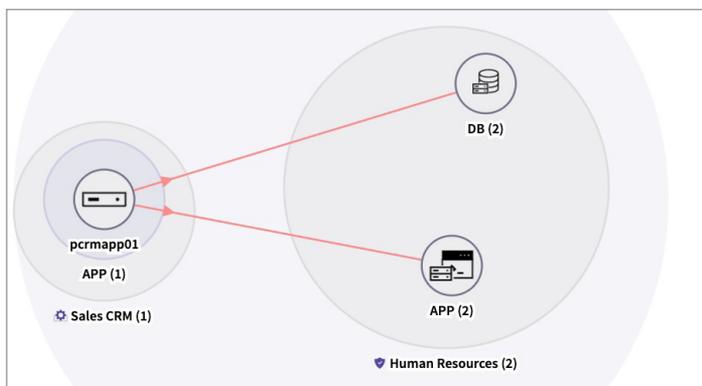


Figure 1-3: Unexpected communication with systems holding confidential data should be investigated. (Source: ColorTokens)

Strengthen compliance

Zero trust security solutions can strengthen compliance by:

- ✓ Isolating cardholder data environments (CDEs), healthcare applications, and customer databases
- ✓ Preventing unauthorized access by insiders and business partners as well as outside attackers
- ✓ Generating reports that document enforcement of privacy regulations and need to know standards

- ✓ Ensuring that software development and test environments can't access production databases
- ✓ Narrowing the scope and cost of audits by showing exactly what systems process protected data

Enable migration to cloud platforms

Zero trust security platforms make cloud migrations more secure by giving IT teams the ability to visualize, monitor, and control network traffic to workloads on cloud platforms, including those running in virtual machines and containers.

Integration with *orchestration* and cloud management tools can ensure that security policies are applied to workload instances as soon as they are created, and that when workloads are moved policies move with them.

Zero trust security also facilitates identifying and blocking command and control (C&C) traffic between attackers on the internet and workloads hosted on cloud platforms.

In This Guide...

Zero trust security is a broad (and rapidly evolving) subject. In this guide we focus on major topics related to:

- ✓ Visibility into suspicious network traffic
- ✓ Micro-segmentation to isolate and protect critical applications and high-value assets
- ✓ Cloud workload protection
- ✓ Endpoint and application security
- ✓ Practices to implement zero trust security without disrupting the business

We will not have enough space here to cover a few related topics such as dynamic authentication, but you can find articles and reference works about those topics on the web.

Chapter 2

The Path to Zero Trust Security

In this chapter

- Review challenges involved in deploying zero trust security
- Learn about creating maps and defining zero trust zones
- Understand the basics of creating and enforcing policies

“To achieve great things, two things are needed; a plan, and not quite enough time.”

– Leonard Bernstein

One Step at a Time

We have discussed the principles and benefits of zero trust security, but how do you go about deploying it?

You should know at the outset that some significant challenges are involved. For example, you will need to:

- ✓ Discover and classify hundreds of resources on the corporate network and on cloud platforms
- ✓ Avoid disrupting business processes when you begin to enforce security policies and block access to resources
- ✓ Show benefits early, to generate momentum for the project

Some people have the impression that the journey to zero

trust security is daunting. Fortunately, it is very manageable with a step-by-step approach and the right tools.

The Steps

Figure 2-1 shows the steps you can use to implement zero trust security.

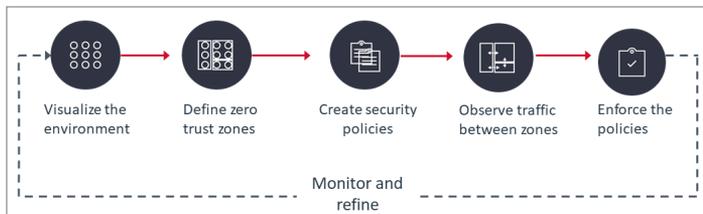


Figure 2-1: Steps on the path to zero trust security

Visualize the environment

Start by visualizing the entities in the IT environment and their relationships. This involves:

- ✓ Discovering and classifying the entities, including users, applications, resources, and network segments
- ✓ Observing network traffic and mapping connections between the entities

Discovery and mapping can be performed largely with automated tools. Figures 1-2 and 1-3 in the previous chapter are examples of the output of such tools.

Visualizing the environment provides a lot of value right away. Observing network traffic can reveal indicators of attacks. A map of connected entities can be used forensically to trace lateral movement by attackers. Also, with the right tools you can detect vulnerabilities caused by open ports and misconfigured systems. We will go into more depth on visualization in Chapter 3.



Don't try to map the entire IT environment at once. Start with one standalone application or high-value asset so you can prove the concept and learn by doing.

Define zero trust zones

The second step in the process is to define *zero trust zones*: collections of information resources that should be segmented, monitored and protected as a unit, using one set of access policies. Zones can range from the IT infrastructure for an entire country, to complete data centers and cloud platforms, to all the components of one application, down to individual workloads running on one server, endpoint, or container.

Activities for defining zero trust zones can be automated. For example, analytics tools can recognize workloads that are in the same network segment or part of the same application. However, human administrators should be involved to ensure that zones are aligned with business processes. We will describe the definition process in Chapter 4.

Create policies

The next step after defining zero trust zones is to create security policies to control access between them.

Zero trust security solutions usually employ techniques to create and manage policies that are very different from firewall and VLAN rules. Some of these are:

- ✓ A whitelist approach that minimizes the number of rules that need to be created and maintained
- ✓ Visual policy editors that leverage the maps created in the early phases of the process
- ✓ Templates that allow policy sets to be reused

We will explore these techniques in Chapter 4.

Observe traffic between zero trust zones

If you give Murphy's Law free rein, the first time you enforce security policies one of them will prevent a top executive from performing an urgent task.

To avoid this, you need an observation period during which you observe production network traffic and generate alerts based on your policies, but do not actually block any connection. You investigate the alerts to find out which are critical and which are false positives. We will review some of the details in Chapter 5.



If you use simulated traffic flows for this exercise, Murphy will get you! No matter how hard you try to replicate real production traffic, you will miss edge cases. When enforcement is turned on, you will inevitably interfere with some important business process.

Enforce the security policies

Finally, the biggest payoff from zero trust security! When you enforce the security policies, you protect the zero trust zones from attackers moving laterally and from insiders trying to reach resources they are not authorized to use.

We will review the benefits of policy enforcement in Chapter 5. We will also look at how very similar principles can be applied to protecting endpoints and applications.

Monitor and refine the zero trust zones and policies

Organizations are constantly adding and changing applications, modifying their infrastructure, and giving users new roles. It is important to monitor those changes and continuously update zero trust zones with dynamic policies that protect them.



Look for zero trust security platforms that automate as much of the monitoring process as possible and dynamically update policies.

Chapter 3

Obtaining Comprehensive Visibility

In this chapter

- Learn about the process for discovering and classifying entities
- Discover how visual maps can be used to find indicators of attack, support incident response, and assess risk

“I wisely started with a map.”

– J. R. R. Tolkien

Comprehensive visibility is a prerequisite for zero trust security. To segment meaningful zero trust zones and create policies to control access between them, you must be able to visualize all the resources and entities in the IT environment and their connections.

But you don't have to wait for security policy enforcement to begin generating wins for your zero trust security initiative. Visibility by itself can deliver major benefits for security and compliance.

Discover Entities

To make zero trust security work, you need to discover many types of entities in your IT environment, including users, network segments, systems and devices, applications, workloads, and endpoints. Fortunately, most of the process can be automated.

Data about entities can be obtained from:

- ✓ Directories, asset inventory systems, configuration management database (CMDB) products, and cloud platform services such as AWS Systems Manager and Azure Resource Manager
- ✓ Lightweight agents distributed to endpoints, containers, on-premises servers, and cloud platforms
- ✓ Scanners that monitor network traffic and discover new entities

Classify and Group Entities

As you discover entities, you want to classify (tag) and group them. Tag entities based on factors such as membership in a user, application, or device group, workload type (web front end, database, etc.), and location in a container, data center, or cloud platform.

With the right tools, much of this classification can be done automatically, based on known attributes of the entities. However, administrators or analysts should verify and elaborate on the classifications to make sure they are meaningful.

Visualization tools

Your maps can involve hundreds of entities and thousands of connection paths. To help answer questions and solve problems, zero trust security platforms usually include a visualization tool that makes it easy to:

- See what entities (including entities on cloud platforms) belong to the logical groups created

during the classification process

- Zoom in to focus on small corners of the map, and zoom out to obtain a “big picture” view of the groups and traffic between them
- Filter and organize views based on criteria such as zero trust zone, type of resource, location, and relevance to a compliance standard



TIP Look for a zero trust security solution that can use tags to classify entities and to update zones and policies automatically when conditions change.

Create and Use Visual Maps

The next step in the process is to monitor network traffic and create maps showing connections made between entities. The maps give you comprehensive visibility into the topology of your infrastructure and the relationships and interactions among users, applications, workloads, devices, and endpoints.

Find indicators of attack

Maps will show you connections that indicate malicious activities. These include connections to key resources by users and systems that do not have a need to access them, and possible C&C traffic between internal workloads and external systems on the web (Figure 3-1). If your zero trust security platform has built-in analytics, you can use statistics like flow data and failed connection attempts to flag attacks.

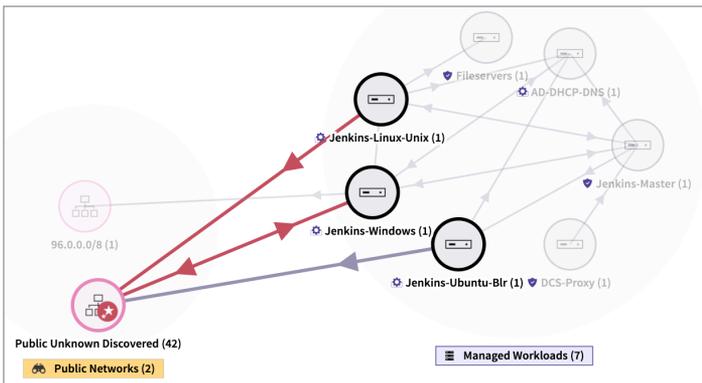


Figure 3-1: A map showing internal workloads connecting to an external server. (Source: ColorTokens)

Support incident response and forensics

A visual map can provide valuable insight into the activities and movement of an attacker within the organization's IT environment. Incident response and forensics teams can "pivot" from an indicator to work back to the attack's point of origin and to understand how it has traversed the network.

For example, if a user account is flagged for suspicious behavior (indicating that it has been compromised), maps can show external servers that communicated with that user's devices and the resources on the network that the user has accessed. This information helps security teams contain attacks faster.

Find compliance violations

Industry standards and government regulations often mandate that access to information be controlled on a need to know basis. Your map will show all entities that have accessed applications and databases that handle protected information.

In the event of a data breach, visual maps can show not only what entities have accessed protected data, but what entities have *not* had access. This information can reduce the scope and cost of breach notification efforts.

Assess exposure and risk

Maps can help assess the “reachability” of high-value assets by documenting how many entities have access to them and the exploitable paths that untrusted sources could use to connect. Some zero trust security solutions go even farther by generating residual risk scores for systems based on reachability plus the number and severity of vulnerabilities identified by vulnerability scanners. This information can be used to prioritize remediation and add security controls to protect the organization's most vulnerable assets.

Chapter 4

Defining Zero Trust Zones and Creating Policies

In this chapter

- Understand zero trust zones
- Learn about software-defined micro-segmentation
- See how to define zero trust zones and create policies for them

“Prepare and prevent, don’t repair and repent.”

– Anonymous

What Is a Zero Trust Zone?

A zero trust zone is a segmented collection of IT resources protected according to a set of access policies based on zero trust security principles.

As we mentioned earlier, zero trust zones can be all the systems and workloads in a data center or on a cloud platform, or all the web servers in one location, or one physical server, or all the workloads and services in one application, or a single workload.

Zero trust zones can be nested. As illustrated in Figure 4-1, a zone containing all production systems might include a zone for a three-tier application, which could include zones for web, app, and database tiers.

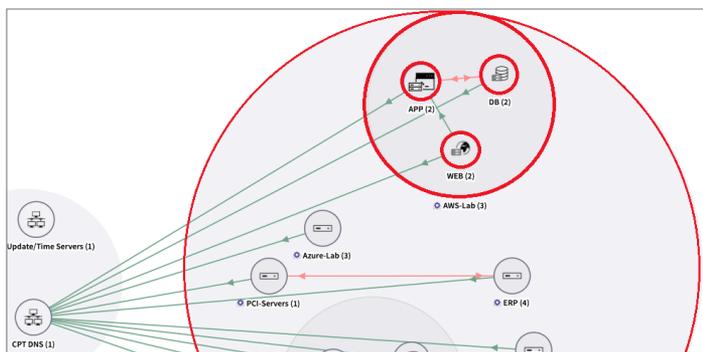


Figure 4-1: Zero trust zones can include many types of IT resources and can be nested. (Source: ColorTokens)

In this case, one set of access rules applies to the production systems zone, another, more-restrictive set applies to the application, and even more restrictive rules apply to the individual application tiers.

Zero trust zones can be defined by physical location, but are more useful when defined logically. Zero trust zones can cross physical boundaries, which is critical, because today one application can be spread across multiple cloud platforms, company data centers, and remote offices.

Enforcement Requires Software-defined Micro-segmentation

You might think that organizations could implement zero trust security using conventional firewalls or virtual LANs (VLANs) and access control lists (ACLs) to segment their environments. However, using that approach has proven to be very difficult in practice, because:

- ✓ Change management processes for firewalls and ACLs are typically manual, slow, and error prone.
- ✓ Conventional firewalls are too expensive to make fine-grained segmentation affordable.
- ✓ Rules are based on IP addresses and network-level constructs like VLAN membership, which makes changes very difficult to manage.

The last point is the most important. Applications running in virtual environments and on cloud platforms are highly dynamic. Workload instances are constantly being moved, and new ones spun up in new locations. Static firewall rules and ACLs cannot possibly be updated quickly enough to keep pace with these changes.

In dynamic environments, zero trust security policies can only be enforced with software-defined micro-segmentation.

Say what?

Software-defined micro-segmentation? Let's break down that term.

The original definition of *micro-segmentation* in IT was the division of a network into small segments or zones with barriers or access controls between them. With the widespread adoption of virtualization and cloud platforms, the term has evolved to mean the division of a network *and the resources on it* into zones (e.g., a single server or a workload can be a zone).

Software-defined micro-segmentation means that entities and zones are described in terms of logical attributes abstracted from underlying hardware and networks. Entities described logically are much easier to manage in a dynamic environment.



Suppose you define a zone as the web tier of a CRM application in a production environment, then apply a set of access control policies to that zone. If the application is moved to a new virtual machine, or a different data center, or a cloud platform, the policies can follow it automatically (assuming you have a zero trust security platform with orchestration capabilities). Nobody has to manually update security tools with the new IP address or subnet of the application.

Start with a Model

You can jump in and start defining zero trust zones from the ground up with the entities that have been discovered and classified. However, it is a better practice to start by creating a high-level model.

For example, for a CRM system you might start by scoping out the elements shown in Figure 4-2 to identify:

- ✓ Application modules such as the web front end, the business logic tier, and the database
- ✓ Users of the application, for example, the sales team and authorized resellers
- ✓ Services that support the application, such as DNS and back-up and recovery services

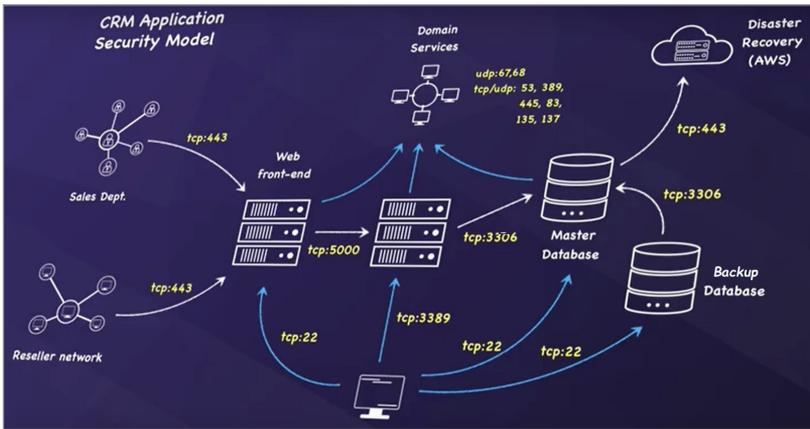


Figure 4-2: A high-level model shows the main elements of the zone, users, supporting resources, and related information.



As you develop the model, document related information such as the ports and protocols used to connect the elements. You can also ask users about expected usage, exceptions, and “abuse cases” (who might misuse or compromise the application, and how).

Define the Zero Trust Zones

Exactly how you define zero trust zones will depend on the zero trust security platform you are using. The platform may suggest possible groupings based on the attributes and tags assigned to entities during the classification process. After that you should be able to choose entities either by picking them from a list or selecting them from a visual representation like a map.

Zones are defined by attributes, which can be assigned by you, or the zero trust security platform, or both. Tags assigned during discovery and classification can become attributes. Attributes might include:

- ✓ Membership in a resource group (applications, shared services, users, network, endpoints, etc.)
- ✓ Application, container, or workload name
- ✓ Application role (web front end, application logic, database, etc.)
- ✓ Location (country, city, subnet, data center, cloud platform, etc.)
- ✓ Environment type (dev, test, staging, production, etc.)

Create Policies

Everything not permitted is forbidden (whitelists)

Two principles of zero trust security are that all entities start as low trust and least privilege access should be applied at all times. In practice this means applying a whitelist approach to policy creation. Specifically:

- ✓ Each zone starts with a “deny all” policy.
- ✓ An appropriate security policy must be created for each entity that needs to access the zone.
- ✓ All access that is not explicitly allowed should be blocked (although not necessarily right away).

Whitelists versus blacklists

Firewalls typically are configured with hundreds of rules to allow and deny access to entities. These rules quickly become outdated and are extremely difficult to maintain. Zero trust security platforms start with a deny all posture and require administrators to explicitly

allow connections. This approach is more secure and much easier to keep current. The trick is to avoid blocking unusual but legitimate connections. In the next chapter we will describe how to succeed at this.

Defining access

Creating policies involves defining access parameters and allowed connections between two entities. Parameters usually include the user, role, protocol, and ports allowed for connections between the entities.

Depending on the zero trust security platform you are using, you may be able to define access parameters using a table, matrix, or visual policy editor.

Templates

Policy templates can greatly simplify the policy creation process. For example, you might want to ensure that external entities can access the front ends of three-tier web applications, but never access workloads in the business logic or database tiers. If you capture the necessary rules in a template you can apply the template to each new three-tier application with only a few clicks, as shown in Figure 4-3.



Figure 4-3: Templates make it easy to quickly deploy policies for protecting a three-tier web application. (Source: ColorTokens)

Chapter 5

Enforcing Zero Trust Security

In this chapter

- Understand why observation comes before policy enforcement
- Explore best practices for turning on enforcement
- Learn about zero trust security for workloads, applications, and endpoints

“Good fences make good neighbors.”

– Robert Frost

Observe Before You Leap

You have finished defining zero trust zones and creating security policies based on all known legitimate and malicious traffic patterns. But it’s not quite time yet to start blocking unauthorized entities from accessing resources.

No matter how much time you put into developing the policies, there is a risk that enforcing them will interrupt some business process or prevent people from reaching data and applications they need to do their jobs. You are likely to encounter problems because of:

- ✓ Unusual and infrequent business processes
- ✓ Unforeseen user needs for data
- ✓ Subtle changes to applications

To avoid unintended consequences of policy enforcement, establish a period for observing but not blocking traffic. During this time a zero trust security platform can provide visibility into connections that comply with policies and those that violate them (shown by different colored lines in Figure 5-1), and can generate alerts for the violations.

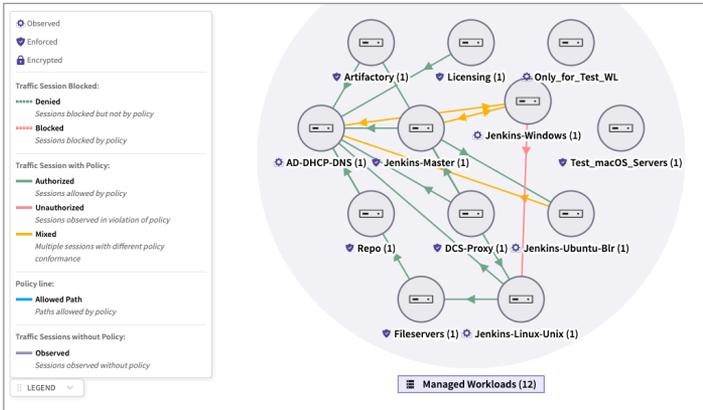


Figure 5-1: During the observation period, a zero trust security platform can generate alerts for connections that violate policies. (Source: ColorTokens)

Investigate the alerts and categorize the connections as:

- Rare but authorized
- Suspicious
- Unauthorized

Tweak the policies so that authorized connections will continue to be allowed, suspicious ones will be monitored to determine if they are malicious, but not blocked right away, and unauthorized connections will be blocked.

Use observation, not simulation

There is a critical difference between observing results with production network traffic and running a simulation with traffic that approximates or samples actual flows. Simulations are great

for analyzing typical behaviors, but almost inevitably fail to include all the edge cases and exceptional events – which are precisely what you need to test before turning on enforcement.

Turn on Enforcement

Go zone by zone

When the observation period is complete, you can gradually turn on enforcement one zone or use case at a time. For example, you might want to start by protecting a high-value application or by enforcing separation between development, test, and production environments. A phased approach allows you to catch mistakes, work the kinks out of the processes, and show early successes. You can pick up the pace as you become more proficient.



TIP

Be prepared. Don't forget to develop a problem reporting plan and to alert your end users, technical support group, and IT staff when you plan to start enforcing policies. Everyone needs to know what to do if a misconfiguration causes a service interruption. Fixing issues quickly to avoid disrupting business must be the highest priority.

Use orchestration for DevOps

Automation and orchestration are essential to protect applications and resources hosted in dynamic cloud and virtual environments. You can ensure that security policies are applied to workloads as they are provisioned and follow the workloads when they are moved by using a zero trust security platform that is integrated with *DevOps* and cloud infrastructure tools such as Chef, Ansible, Puppet, Terraform, AWS CloudFormation, and Azure Resource Manager. (Figure 5-2)

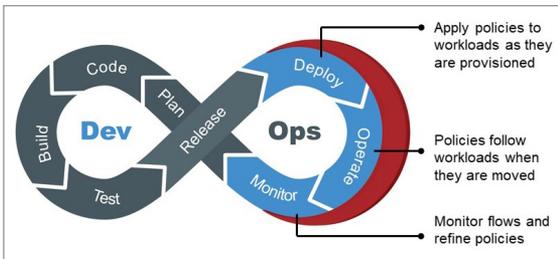


Figure 5-2: Integration with orchestration tools ensures that policies are applied to cloud workloads and move with them.

Monitor and refine policies

Use your zero trust security platform to monitor connection activity. As new threats emerge and applications and infrastructure change, the security team needs to look for both policy violations and new connections that require policies. Policies should be updated and created on a regular basis.



Make someone explicitly responsible for monitoring and refining policies. Unless these tasks have a clear owner, they will fall through the cracks when the team is busy.

Generate security and compliance reports

A zero trust security platform that monitors and logs connections can generate reports that provide valuable information in areas like:

- ✓ Regulatory compliance (demonstrating that confidential data is being protected and access is being logged)
- ✓ Data breach notification (documenting what entities did, and did not, access confidential data)
- ✓ Incident response and forensics (tracing lateral movement by attackers back to the source of compromise and forward to targeted resources)
- ✓ Attack surface reduction (pinpointing users involved in risky or suspicious activities and systems unnecessarily exposed to outside threats)

Extend Zero Trust Security to Endpoints

So far, we have been focusing on zero trust security as it applies to the micro-segmentation and control over network access. However, the same principles can be applied to controlling the behavior of processes on endpoints (including servers, laptops, PCs, and mobile devices). This is called “*zero trust execution*.”

Whitelisting processes

We mentioned in the previous chapter that zero trust security principles align with a whitelisting approach to network access. The same idea can be applied to processes running on individual endpoints. If you only allow authorized processes to run, you reduce the risk of harm from malicious software.

Allowing only authorized processes to run:

- ✓ Protects against file-based malware such as viruses, Trojans, keyloggers, and ransomware
- ✓ Thwarts fileless attacks that utilize tools like Microsoft PowerShell to run in memory without installing a file on the endpoint
- ✓ Blocks malicious processes from spawning and infecting other applications

Advantages over AV products

Because zero trust execution does not require malware signatures, it offers several advantages over conventional antivirus (AV) products:

- ✓ Protection against zero-day and other unknown attacks
- ✓ Ability to run on systems with limited CPU power and memory
- ✓ Effectiveness at remote sites with limited network connectivity (because large AV signature databases don't need to be downloaded)

- ✓ Ability to safeguard legacy and hard-to-patch endpoints such as ATMs, point of sale (POS) systems, kiosks, and industrial control systems

Improve Application Security

Many organizations today are redesigning their security programs to focus on protecting applications and data rather than infrastructure. Zero trust security platforms support application-focused security in several ways.

Discovery

By analyzing application transaction information, zero trust security platforms can discover and classify application elements like databases, web front ends, files, and external infrastructure services. In some cases, they can recognize entities within databases, such as individual tables.

Visualization

Zero trust security platforms map application elements and their data flows. This mapping enables software developers and security teams to understand the complete structure and functioning of the applications, including APIs and code paths. Analysts can use this information to identify vulnerabilities, exposed elements, and areas where security controls are weak.

Identification and prevention of web attacks

Some zero trust security platforms can inspect HTTP traffic, then analyze URLs, headers, parameters, and other details in relation to application code, and use the analysis to block attacks. An example would be recognizing that input to a web application contains an improperly formatted SQL query, and based on this recognition blocking a SQL injection attack. Like a web application firewall (WAF), a zero trust security platform can protect against many of the OWASP Top 10 web attacks and similar threats.

Chapter 6

Use Cases

In this chapter

- Explore the value of zero trust security platforms for 10 common use cases

“Distrusting me was the wisest thing you’ve done since you climbed off your horse.”

– Petyr Baelish (“Littlefinger”), Game of Thrones

A use case is a typical scenario for using a product or technology to perform an activity or achieve an objective. Use cases help teams think about processes and concentrate their efforts on producing specific, high-value results. In this chapter we look at 10 use cases that can give focus to zero trust security programs.

Improving Situational Awareness

As we discussed in Chapter 3, organizations can obtain substantial benefits from zero trust security simply by mapping the resources in their environment and monitoring information flows. These benefits include:

- ✓ Finding indicators of compromise
- ✓ Supporting incident response and forensics
- ✓ Finding and remediating compliance violations
- ✓ Assessing exposure and risk

Protecting the “Crown Jewels”

Most organizations have “crown jewel” applications that are vital to the success of the enterprise. Zero trust security can protect these applications by:

- ✓ Helping security teams understand the elements of the applications and monitoring their data flows
- ✓ Observing which people and resources access them
- ✓ Uncovering vulnerabilities and weaknesses in networks, applications, and endpoints
- ✓ Using micro-segmentation to “ring-fence” them and limit access to authorized users and systems

Strengthening Compliance

Use cases relating to regulatory compliance often involve multiple applications and data stores. Organizations may need to protect complex cardholder data environments (CDEs), electronic protected health information (ePHI), and personally identifiable information (PII) stored in multiple locations. Also, the costs of documenting compliance and passing audits can be considerable.

Zero trust security platforms can safeguard sensitive data and reduce the costs of compliance by:

- ✓ Providing visibility into both authorized and unauthorized paths to protected data
- ✓ Determining which users and systems have a legitimate need to access systems with protected data
- ✓ Using micro-segmentation to isolate protected environments and stop attackers from moving laterally
- ✓ Accelerating incident response by logging all connections into and out of protected environments
- ✓ Simplifying compliance by generating reports that document how the organization is controlling access
- ✓ Reducing the scope of audits by showing exactly which systems process or store protected information and which ones fall out of scope

Environmental Separation

Zero trust security can ensure that development environments are more secure and more compliant with regulations. A zero trust security solution can isolate code repositories and prevent access to proprietary software by:

- ✓ Outside attackers, including competitors and hostile state-sponsored hackers
- ✓ Unauthorized insiders, including systems administrators and other privileged users who don't have a real need to know

Another important use case is enforcing separation between software development, test, staging, and production environments. For example, Figure 6-1 is a visual map showing developers on a test system using real customer credit card data from a production database. A zero trust security platform can alert security and compliance teams to this violation of PCI DSS regulations, or simply block the network path.

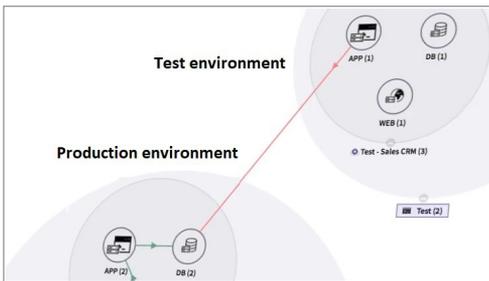


Figure 6-1: Detecting a test system accessing real customer credit card data in a production database. (Source: ColorTokens)

DevOps and Continuous Delivery

DevOps techniques and software orchestration tools enable organizations to deliver software far faster than in the past. New functionality can be coded, tested, integrated into applications, containerized, and pushed to multiple production servers and cloud platforms in hours.

But changes to software functionality often require modification of security policies. Firewalls, VLANs, and ACL rules can never be changed fast enough to keep up. Zero trust security platforms, particularly when they are integrated with orchestration tools as described in Chapter 5, allow access policies to be associated with workloads during development. The policies then follow the workloads through test, staging, deployment, and migration across environments.

Protecting Legacy Endpoints and Operational Technologies

Many organizations have legacy and specialized endpoints that are no longer supported or are difficult or impossible to patch. These include kiosks, ATMs, POS devices, and systems running old operating systems that cannot be updated. These can be a security nightmare.

Similar challenges face manufacturers, utilities, energy companies, and others that rely on operational technologies (OT) to run factories, electrical grids, and pipelines. Their infrastructure contains industrial control systems (ICS), SCADA systems, and process control networks (PCNs) that are connected to networks but are hard to patch or simply were not designed for cybersecurity.

Zero trust security platforms can dramatically improve security and compliance in these situations by:

- ✓ Isolating ATM, POS, and OT networks and tightly controlling access to and from them
- ✓ Locking down unpatched systems at the process level (zero trust execution, as described in Chapter 5)



Standards such as PCI DSS, NIST SP 800-82, NERC CIP, and ISA99/IEC62443 mandate strict control over access to environments with sensitive information and critical industrial systems. Zero trust security platforms can help define the scope and boundaries of these environments, demonstrate control of network traffic in real time, and log access to key systems by users and applications.

Simplifying Remote Access and (Sometimes) Eliminating VPNs

As we discussed in Chapter 1, the concept of Zero Trust Network Access (ZTNA) dictates that trust should be assessed the same way whether a user is in the headquarters building, in a branch office, or traveling. This means that, regardless of location, users can expect the same basic process of authentication and that appropriate access controls will always be enforced. By strictly controlling access, some organizations have found they are able to dispense with virtual private network (VPN) technology for remote and mobile users.

Securing Cloud Migrations

Organizations migrating applications to cloud platforms are often concerned about their ability to monitor and control activities there. A zero trust security platform can discover workloads running on cloud platforms, map connections to other workloads in the cloud and in on-premises data centers, and use micro-segmentation to enforce access policies.

For organizations with a multicloud infrastructure, a zero trust security platform can provide a single tool for visibility and access control across multiple cloud platforms and on-premises data centers. (Figure 6-2).

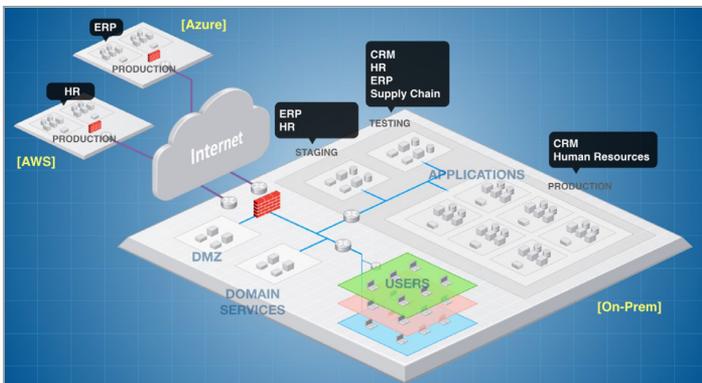


Figure 6-2: A zero trust security platform can provide visibility and access control across a multicloud infrastructure.

Protecting Microservices

Microservice architectures and containers allow development teams to create software that is extremely modular, portable, and easy to maintain. However, they can make it impossible to enforce access control with conventional firewalls and ACLs. Zero trust security platforms can track vast numbers of workloads in virtual environments and containers and ensure that security policies move with them and are applied consistently.

Managing Third-party Risk

Today's organizations want to make sure business partners and other third parties (and the cybercriminals who compromise them) can't move from one or two authorized applications into the rest of the network. Zero trust zones are an ideal solution. They can isolate third-party-accessible applications from the rest of the infrastructure, and when that is not possible, monitor unusual connections between workloads and generate alerts to SOC teams so they can investigate.

Chapter 7

Selecting a Zero Trust Security Platform

In this chapter

- Review five criteria that can help you select a zero trust security platform that fits your organization

“Out of this nettle, danger, we pluck this flower, safety.”

– William Shakespeare (*Henry IV, Part 1*)

We have discussed the principles and benefits of zero trust security, the path to a successful implementation, and important use cases. But how do you select the zero trust security platform that best fits the needs of your organization? In this final chapter we suggest capabilities you should look for and criteria you can use to compare alternatives.

Cloud Delivery

If your organization has migrated applications to the cloud, you should look for a zero trust security platform that operates on a cloud platform.

When application activity spikes, your cloud service provider will quickly spin up more instances of the software to handle the increased demand (that’s one of the reasons you moved the application to the cloud). You need your security tools to ramp up performance just as fast to protect the new instances. Cloud delivery ensures that your zero trust security platform can scale with your applications. Cloud hosting also eliminates

the costs and complexities of buying and managing hardware and installing and upgrading software for the zero trust solution.

Scope of Capabilities

Zero trust security platforms should provide two core capabilities:

1. Discovery and visibility (the subject of Chapter 3)
2. Micro-segmentation and policy enforcement (the focus of Chapter 4)

However, as illustrated in Figure 7-1, some products go beyond these core capabilities and add two more:

3. Endpoint protection and zero trust execution (discussed in Chapter 5)
4. Application security and the prevention of web attacks (also in Chapter 5)



Figure 7-1: Four types of coverage provided by zero trust security platforms.

Offering all four capabilities in one solution means:

- ✓ Fewer agents to distribute to endpoints and workloads
- ✓ Fewer dashboards and tools to learn
- ✓ More information, so analysts can better detect patterns and trace attacks

Breadth of Resources Protected

Zero trust security platforms should be able to create zero trust zones and control access for a wide range of resources:

- ✓ Individual workloads on bare metal servers, virtual machines, and containers in corporate data centers
- ✓ Workloads on cloud platforms
- ✓ Workloads and processes on endpoints
- ✓ Servers, endpoints, and specialized devices
- ✓ Complete applications
- ✓ Complete environments, such as development, test, and production environments, or POS and card data environments
- ✓ Complete locations, such as branch offices and data centers
- ✓ Customer-defined entities with combinations of the above

Ideally a zero trust security platform should be able to “zoom in” to very granular elements like individual workloads and executing processes, and also “zoom out” to very large constructs such as environments and business organizations.



TIP

Look for the ability to define users and user groups as entities and create policies that control what resources they can access.

Ease of Implementation and Management

Discovery, mapping, and policy creation tools

An organization may need to observe and manage hundreds or thousands of resources on the corporate network and on cloud platforms. That means it needs good tools and well-designed user interfaces to simplify tasks such as discovering and classifying entities on the network, grouping entities based on common characteristics, creating connection maps, and creating policies.



Look for good visualization tools, dashboards, and mechanisms for creating and using templates.

Ultra-lightweight agents

To simplify deployment and ongoing administration, a zero trust security platform should offer ultra-lightweight agents that are easy to install and don't affect the performance of endpoints and servers.

It is also advantageous if the agents support network micro-segmentation, workload protection, endpoint protection, and web application control, so that those functions can be monitored through a single management console.

Integration with other security tools

A zero trust security platform should be able to share information with other security tools, including:

- ✓ Cloud service provider security, management, and logging tools, so the platform can identify and protect cloud-based workloads
- ✓ SIEM systems, so alerts generated by the platform can be delivered quickly to security analysts
- ✓ Orchestration and automation tools, so the platform can support DevOps techniques and work effectively with cloud platforms

Total Cost of Ownership

To evaluate the total cost of ownership (TCO) of a zero trust security platform, take into account:

- ✓ Licensing and maintenance costs for the platform
- ✓ The effort required for the initial implementation, including the discovery and classification of resources, policy creation, and observation prior to enforcement
- ✓ Ongoing monitoring of connections and refinement of policies

Of course, you will want to set these costs against the potential savings from preventing expensive data breaches and reducing the effort required to manage firewalls, ACLs, and other legacy security tools. The money you save from preventing just one or two data breaches in a key application may be enough to justify the entire zero trust security program.

Glossary

DevOps: practices that combine software development processes (Dev) with IT operations (Ops), allowing new software functionality to be coded, tested, and put into production in hours.

flat network: a non-segmented network, where all resources can be reached through one router or switch.

micro-segmentation: the practice of dividing an IT environment into many segments or zones, with access to each segment or zone controlled by policies. Segmentation can be enforced by a variety of means, including network devices, firewalls, and zero trust security platforms.

orchestration: the automation of workflow processes, for example, automating the process of assigning access policies to workflows when they are spun up.

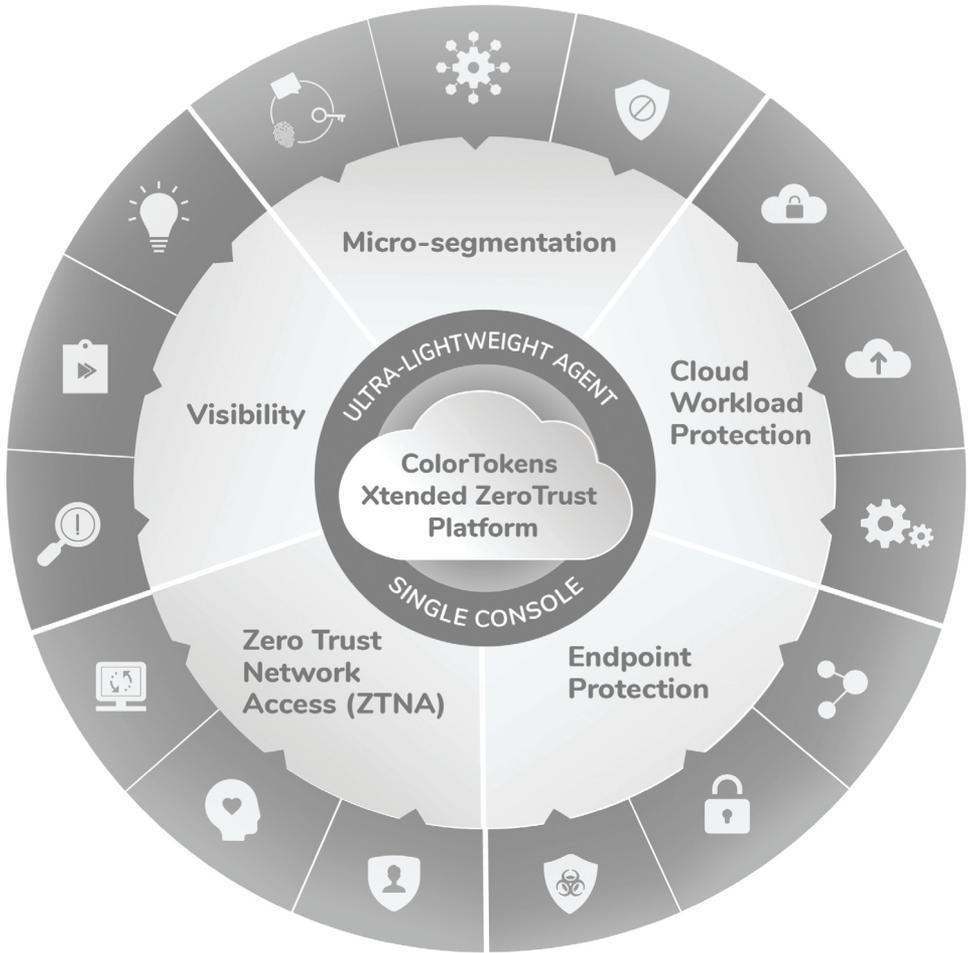
principle of least privilege: the principle that entities should be given access only to the resources and data they need to perform their intended functions.

zero trust execution: permitting only authorized (whitelisted) software to execute on a host or endpoint.

zero trust network access (ZTNA): controlling access to networks using zero trust security principles, especially in reference to access by remote and mobile users.

zero trust security: access control based on the principles that no network or entity is inherently trustworthy, that networks should be segmented, that connections between entities should be allowed only after their trustworthiness has been established, and that access should be the minimum appropriate to that level of trust.

zero trust zone: a collection of information resources that are segmented, monitored, and protected as a unit, using one set of access policies.



**Today, no person or device can be trusted implicitly.
But you can control access to systems and data so
people see exactly what they need — and no more.**

Zero trust security gives people and processes access to just the information resources they are authorized to use and prevents cybercriminals from running loose on your networks. It protects your critical workloads, applications, and data in data centers, on cloud platforms, and on thousands of endpoints. Learn how to apply the principles of zero trust security in your organization.

- **Zero trust security** — understand the concepts and benefits
- **Visibility** — see how you can monitor connections between entities on your network and in the cloud to find indicators of compromise
- **Zero trust zones** — learn how to create micro-segments for workloads and data on servers, cloud platforms, and endpoints
- **Micro-segmentation** — examine how to stop lateral movement by bad guys without disrupting users
- **Use cases** — explore 10 ways zero trust security can help your organization
- **Selection criteria** — review criteria for selecting the zero trust security platform that meets your organization's needs

About the Author

Jon Friedman has spent over 20 years working in industry analysis and marketing for software and IT services companies. He has described cutting-edge technologies and their business benefits for more than 40 high-tech companies. Jon has a BA from Yale and an MBA from Harvard.



CYBEREDGE
P R E S S

Not for resale

ISBN 978-1-948939-09-6



9 781948 939096 >