

LEARNING MADE EASY

Illumio Special Edition

Zero Trust Segmentation

for
dummies[®]
A Wiley Brand



Build a resilient
response to cyberattacks

Achieve safe, powerful
cyber transformation

Stop ransomware
before it spreads

Brought to you
by

 **illumio**

Steve Kaelble

About Illumio

Illumio, the Zero Trust Segmentation company, stops breaches and ransomware from spreading across the hybrid attack surface. The Illumio ZTS Platform visualizes all traffic flows between workloads, devices, and the Internet; automatically sets granular segmentation policies to control communications; and isolates high-value assets and compromised systems proactively or in response to active attacks. Illumio protects organizations of all sizes — from Fortune 100 companies to small businesses — by stopping breaches and ransomware in minutes, saving millions of dollars in application downtime, and accelerating cloud and digital transformation projects.



Zero Trust Segmentation

Illumio Special Edition

by Steve Kaelble

for
dummies[®]
A Wiley Brand

Zero Trust Segmentation For Dummies®, Illumio Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2023 by John Wiley & Sons, Inc., Hoboken, New Jersey

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, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY 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 ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

ISBN 978-1-394-18168-1 (pbk); ISBN 978-1-394-18169-8 (ebk)

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Elizabeth Kuball

Acquisitions Editor: Traci Martin

Editorial Manager: Rev Mengle

Client Account Manager:

Cynthia Tweed

Production Editor:

Mohammed Zafar Ali

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	2
Where to Go from Here	2
CHAPTER 1: Containing the Breach	3
Entering the Breach Containment Era	3
Understanding Breaches	5
Assuming There's a Breach	6
Trusting Nothing	7
Making Containment Happen	8
CHAPTER 2: Introducing Zero Trust Segmentation	9
Appreciating Segmentation	9
Understanding Zero Trust Segmentation	11
Using Zero Trust Segmentation	12
Transforming the Business	14
Establishing Security Objectives	16
CHAPTER 3: Seeing the Attack Surface	19
Realizing the Vulnerabilities	19
Seeing What You're Protecting	20
Putting It All in Context	22
Identifying Assets	23
Prioritizing Risks	24
CHAPTER 4: Responding to and Isolating Ransomware	25
Moving Laterally	25
Identifying the Unexpected	27
Stopping the Spread	28
Eliminating unnecessary connections	28
Using visibility	29
Improving responses to intrusions	30
Building Resilience	30
Containing Operational Technology Breaches	31

CHAPTER 5:	Protecting Applications	33
	Shrinking the Blast Radius.....	33
	Figuring Out Where to Start.....	34
	Making Quick but Smart Decisions	35
	Maintaining App Availability.....	36
CHAPTER 6:	Keeping Ahead in the Game	37
	Building a Loop of Optimization.....	37
	Validating Security Outcomes.....	38
	Deciding What's Next	39
	Shifting Security to the Left.....	39
CHAPTER 7:	Ten (or So) Use Cases for Zero Trust Segmentation	41
	Isolating Ransomware	41
	Securing the Cloud.....	42
	Migrating to and from the Cloud.....	42
	Integrating Information Technology and Operational Technology.....	42
	Securing the Supply Chain	43
	Protecting High-Value Assets.....	44
	Helping Out with Compliance.....	44
	Responding Effectively.....	44

Introduction

Here's a reality check that you may not want to hear: You can't hide from cybercrime. Sooner or later, your organization will be hit with a ransomware attack or some other act of cyber malfeasance (and odds are, you've already had that disruptive and frightening experience).

But that's not the bad news that it sounds like. In fact, when you've accepted the fact that breaches are inevitable, you can focus more attention on surviving those attacks and keeping the business moving forward. That doesn't mean you stop trying to keep the bad actors out of your network — it means you also prepare to deal with them when they do get inside.

That's the aim of Zero Trust Segmentation (ZTS) — building defenses at the workload level, cutting off unnecessary communication, and verifying everything that you purposely let through. You don't know what's benign and what's not, so you throw implicit trust out the window.

ZTS is a remarkably effective way to build resilience in the face of increasingly scary threats. The objective is to prevent any attacks that you can, while limiting the damage from the ones that inevitably will get through even the best defenses. You can't afford to let malicious actors shut down your operations, and ZTS ensures that business can go on as usual, even in the face of a successful attack.

About This Book

Zero Trust Segmentation For Dummies, Illumio Special Edition, is your guide to this powerful concept in cyber resilience. It explains why the concepts of least privilege and Zero Trust should be part of any security architecture and how segmenting your information and operational technologies can stop malware from spreading. This book helps you understand your attack surface, reduce the impact of an attack, keep your applications running, and stay ahead of the threats for good. You'll be pleased to learn that ZTS is a well-defined and achievable concept, and it's remarkably straightforward to adopt and easy to use.

Foolish Assumptions

In preparing this book, I made a few assumptions about you, the reader:

- » You may be a C-suite leader or someone who works day-to-day in information technology (IT) security or network infrastructure.
- » You want to learn how Zero Trust principles and segmentation can solidify your security posture.
- » You fully grasp the need to build resilience against cyber threats so you can keep operations moving forward.

Icons Used in This Book

In the margins of this book, you see lots of little pictures called *icons*. These icons offer clues about the paragraphs next to them.



REMEMBER

The Remember icon flags material so important that you should commit it to memory.



TIP

The Tip icon highlights anything that will make your life a little easier — at least when it comes to ZTS.



TECHNICAL
STUFF

Sometimes I get into the weeds on a subject, and when I do, I use the Technical Stuff icon. If you're in a hurry, you can skip these paragraphs without losing the point of the subject at hand.



WARNING

You already know the dangers of cyberattacks. The Warning icon underscores a threat you should definitely not ignore.

Where to Go from Here

After reading this book, you're bound to have an appetite for more information about ZTS. Head to <https://illumio.com/zts> for a ZTS solution that's both powerful and easy to use. You can also find a vast library of information on cybersecurity and ZTS at <https://resources.illumio.com>.

IN THIS CHAPTER

- » Facing the reality about breaches
- » Finding out more about breaches
- » Lowering your trust level to zero
- » Preparing to contain the inevitable breach

Chapter 1

Containing the Breach

A dose of truth is not always pleasant, but it can certainly be powerful medicine. The medicine this chapter asks you to swallow is the reality that your best cyber defenses are not good enough.

This chapter offers scary details about breaches and asserts that breaches *will* happen, no matter how careful you are. Then it explains why it's actually empowering to assume breaches will occur, lower your trust threshold accordingly, and learn how to contain the breaches that happen.

Entering the Breach Containment Era

Cybersecurity incidents tend to keep a lot of people awake at night, and not just people who work in information technology (IT). The headlines are filled with horror stories of costly attacks and incredibly disruptive ransomware incidents, which can be unspeakably painful for business operations and organizational reputations. Executives right up to the CEO are losing sleep, too, and it's safe to say that most organizations would give just about anything to prevent such nightmares.



REMEMBER

But here's the sobering news — relying on prevention alone is no longer a winning strategy. Attacks are practically inevitable these days. In fact, roughly three-quarters of organizations have been attacked by ransomware in the past two years. On average, an attack happens every 11 seconds.

When you think about it, that's hardly surprising. These days, pretty much everything is connected to the internet, and the IT environments that used to be mostly on-premises are quickly migrating to hyperconnected, cloud-first, hybrid architectures. Digital transformation is a convenience and a huge opportunity; it also greatly expands the possibility of attack.

With today's hybrid IT, you've got multiple clouds, lots of end points, data centers, containers, virtual machines, and mainframes. But with rapid digital transformation underway in every industry, it's not just these typical IT systems that are increasingly interconnected. Operational technology (OT) — such as patient health-care devices in a hospital, pumps on an oil pipeline, or machine controllers on a factory production line — are ever more densely connected to traditional IT. That further expands both the potential attack surface and what could be impacted by a breach. Interconnectivity means attackers have lots of places to go after they get in the door.

Back in the 2000s, security efforts focused on prevention, locking that door and “keeping them out.” Your systems lived relatively safely inside a moat. A decade later, attackers were becoming increasingly adept at getting across the moat, so the focus shifted to detection and “finding them quickly.”

Now, another decade has passed, and the picture seems bleaker than ever. If you're into betting, you can pretty much bet on breaches happening. That means it's time for another shift in IT security focus, to the goal of “minimizing the impact.”



REMEMBER

That sounds terribly pessimistic, but in a way, it can be almost a liberating thought. For sure, you still want to prevent breaches to the extent that you can, but once you start assuming they're going to happen, you can turn the bulk of your attention to limiting and containing the impact. And the good news is, you've got a lot more options than you may realize today, in the “breach containment era.”

Understanding Breaches



REMEMBER

A *security breach* is when an attacker makes it past an organization's cybersecurity defenses at such places as end points, the network perimeter, within data centers, or in the cloud. That, essentially, is getting the attacker in the door, with access to the corporate network.

Sometimes a breach happens unintentionally when employees accidentally leak info to third-party sources, perhaps by downloading something incorrectly or failing to secure a device. Usually, though, security breaches happen through the purposeful actions of an attacker.



REMEMBER

A *data breach* is what can happen next, when the attacker moves laterally inside the environment, reaches the prize (sensitive data), and then steals or exfiltrates the information. That data can then be sold for profit on the dark web, or it may be locked down in the hopes of extracting ransom.

Understanding how security breaches happen and trying to prevent them is important. It's harder than ever to prevent every breach, but it's still important to try. Here are some of the reasons that security breaches happen:

- » **Employee error:** This is one of the biggies, responsible nearly half the time. Maybe an employee was too generous with permissions, didn't configure security tools properly, accidentally left a sensitive document open, or didn't properly secure files or folders.
- » **Malware:** Cybercriminals may use malware to open the door, often by stealing credentials. Malware can be downloaded onto one computer and then moved laterally to infect others on the network.
- » **Phishing:** This is a common way that malware gets downloaded in the first place. Hackers create legitimate-looking emails from sources that appear trustworthy, and when the employee opens the email or attachment, or clicks a link in the email, it triggers an infection.



TIP

Notice a common thread? **Employee errors and inadvertent actions.** Many security breaches can be prevented by training employees to be on the lookout for suspicious emails, and never click anything unless they're certain it's legit. Also be sure they know how to create safe and unique passwords and use multifactor authentication when available.



TIP

Here are some other ideas for good security hygiene to reduce the likelihood of a breach:

- » **Be aware of everything that must be protected.** Ensure you have visibility into all vulnerable end points and the software running there — security teams don't like to be surprised.
- » **Keep an eye on those vulnerabilities.** When you have a list of end points to secure, monitor them all for changes that cause risk, as well as patches that beef up security.
- » **Be stingy with privileges.** The concept of least privilege is vital here. Put tight limits on administrative privileges and monitor them carefully.
- » **Benchmark your configurations.** With your assets all inventoried and a plan in place to protect and monitor them, it's wise to check those configurations and be sure they live up to what industry standards recommend.
- » **Exercise your situational awareness.** Things change all the time, with updated software, new threats, and revised priorities. Your security team should keep up to date on the changing environment and expectations.

Assuming There's a Breach



REMEMBER

It's no longer realistic to think you can prevent all breaches, or to hope that you can find breaches fast enough to fix them. These days, though your security goals still should include prevention and detection, they must go much further — you must *assume breach* and focus on limiting and containing.

Your first thought may be, “That sounds like setting myself up for crippling paranoia” by starting out assuming breach. You may imagine some sort of IT security equivalent of walking down a crowded sidewalk fearing that every passerby is plotting to do you harm. That would be enough to send a lot of people trembling back into their homes, never to emerge in public again. But that's really not the right metaphor to conjure in your mind in the era of breach containment.

Think instead about some favorite animated story of a superhero, going about a daily routine. There are bad actors all around, popping into the plotline almost incessantly. The hero notices each

one very quickly, sometimes out of the corner of their eye, and pulls off some swift and often dazzling move that neutralizes the threat. The superhero barely breaks stride as they move forward, dealing with everything that comes their way.

Okay, that's animated fiction. But that's closer to reality than you may think. Like that fictional character, your organization is capable of moving ahead through the plotline of daily operation with the situational awareness needed to spot the bad actors and the swift moves that can stop the threat.



REMEMBER

That fearless way forward involves a concept known as *Zero Trust Segmentation (ZTS)*. It's not science fiction or animated superhero fantasy, and the tools for making it happen are surprisingly within reach and easy to build into your systems.

Trusting Nothing



TIP

What you must build is a *Zero Trust architecture*, which is pretty much what its name suggests — it doesn't implicitly trust anything.

A Zero Trust architecture is harder to breach, and it also makes it harder for a successful breach to spread. With this mindset, every interaction between people and workloads and networks and data and devices must be verified before it can proceed. And the principle of *least privilege* grants devices and users only the minimum access needed to get the job done.

Zero Trust is an upgrade to your access control and much more. Before heading down this path, you'll want to apply *multifactor authentication* at all access points and be sure all connected devices are regularly updated and well maintained. You'll need regular, thorough monitoring to be sure your access control is tight. And you must improve management by limiting access to individual components in the network.



REMEMBER

You also must really embrace the principle of least privilege as a cybersecurity best practice. The basic idea: Workloads are granted only the permissions they need to perform an authorized task, and nothing more. That limits the potential attack surface, because any particular workload doesn't have all the keys to the castle. By the way, least trust applies not just to applications, devices, and integrated systems, but also to the permissions established for people.

There are five basic steps for a Zero Trust strategy:

1. Determine the attack surface.

The attack surface is all the vulnerabilities a threat actor may go after, and on a hybrid network there may be lots of vulnerabilities. You need a map of the attack surface that considers people, devices, the network, workloads, and data.

2. Create the policies.

Before launching Zero Trust, you have to fully define it, asking and answering all the pertinent questions about how the network will be used, by whom, where, and more.

3. Define access controls.

This step involves establishing access and permission levels for each user or user type. This should also take into account the context, such as user identity, device, location, type of content, and what app is being requested.

4. Choose Zero Trust solutions.

There are tools to make this easier, but every network has its own unique characteristics. *Microsegmentation* is considered a primary Zero Trust security control, because it helps you separate your hybrid infrastructure into different areas and figure out the protocols for each.

5. Monitor.

Implementing Zero Trust is important, but it's not the end of the story. You must constantly monitor network activity for weaknesses and unusual behavior.

Making Containment Happen



REMEMBER

By combining the “assume breach” mentality with Zero Trust architecture, you have the best chance to effectively contain breaches when they happen, in minutes. You’re more likely to see all risks as you visualize communication and traffic across the hybrid attack surface. And because of your Zero Trust focus, communication is controlled and restricted carefully.

You’re able to contain, or stop the spread, much more effectively. You can proactively isolate high-value assets and, during an attack, reactively isolate any compromised systems.

IN THIS CHAPTER

- » Understanding why segmentation is important
- » Applying Zero Trust to segmentation
- » Learning how to transform the business
- » Establishing your objectives

Chapter 2

Introducing Zero Trust Segmentation

Connectivity has been one of the keys to business transformation in recent years, with powerful applications linking to customers and with each other across hybrid information technology (IT) environments. That's both empowering and risky.

This chapter introduces the concept of *Zero Trust Segmentation* (ZTS), which allows the necessary connections while interrupting the malicious ones. Read on to find out how it works, why it's important to trust nothing, how the concept clears the way for business transformation, and how to set the right objectives.

Appreciating Segmentation

Much about the way our lives work today would not be possible without the many interconnections that have made the world a smaller and much more convenient place. A lot has been written about the importance of breaking down organizational walls and siloes.



REMEMBER

That's good advice for many parts of the organization, but when it comes to IT security, tearing down all the walls isn't the best way to go. On the contrary, there are really good reasons to break larger networks into smaller pieces, sometimes as small as the host and workload itself.




REMEMBER

I'm referring to the practice of segmentation. It's essential for helping to prevent attacks and threats from moving laterally, or as some describe it, moving east–west across data networks, clouds, or campus networks. If your network is segmented (sometimes called *microsegmented* depending on how it's achieved), a small security incident is contained to the place where it happened, so it won't turn into a bigger security incident.

To grasp the value of segmentation, it's helpful to consider the metaphor of the submarine. These seafaring vessels are typically divided into compartments to ensure that they'll remain seaworthy even if there's a breach in the hull. Because the submarine is compartmentalized, any problem stays isolated in just a small part of the vessel, as demonstrated in Figure 2-1.

Without Microsegmentation



With Microsegmentation




FIGURE 2-1: How segmentation can save a submarine.



TECHNICAL
STUFF

There are a number of ways to accomplish segmentation. You can rely on the network itself or deploy hardware firewall appliances. You can also enforce segmentation on the host workload itself, which gets the job done without touching the network. Here's a rundown of some of the options:

- » **Network segmentation:** This involves creating segments in networks, using subnets or virtual local area networks (VLANs). VLANs are a layer 2 approach to network segmentation, while IP subnetting is a layer 3 approach to solving the same problem. The drawbacks are that the management of this approach can be quite complicated, and networks often need to be rearchitected.
- » **Firewall segmentation:** Firewalls deployed inside a network can segment functional areas from one another, which limits attack surfaces. You can, for example, keep your engineering applications separate from finance. Firewalls are common around the perimeter, but using them to segment internal networks can be expensive, require thousands of firewall rules, and run the risk of breaking applications due to misconfiguration.
- » **Software-defined networking (SDN):** It's possible to achieve segmentation through SDN network overlay implementation, creating policies that funnel packets through a distributed set of firewalls. This approach can introduce a lot of complexity, especially if applications don't fit into network boundaries.
- » **Host-based segmentation:** This approach enforces a segmented network using the host workload rather than subnets or firewalls. It's considered the best path to implementing microsegmentation, also called ZTS.

Understanding Zero Trust Segmentation

There's segmentation, and then there's ZTS. All the approaches to segmentation outlined in the preceding section have the same general goal, and all can work, but they're definitely not all the same.



REMEMBER

ZTS is rooted in microsegmentation, which is sometimes referred to as *host-based segmentation*.

It's an approach that relies on *allowlist* models to block all traffic other than what's specifically permitted. That's the Zero Trust security model at work, in which nothing is trusted and everything is verified. Microsegmentation partitions the network down as far as individual workloads, allowing you to employ Zero Trust

concepts to their maximum potential to ensure those segments are protected separately. *Voilà!* It's ZTS!

This kind of segmentation establishes a map of the cloud and on-premises compute environments and applications. The map shows what needs to be protected and is used to put automated segmentation policies in place.



TECHNICAL
STUFF

These policies make use of human-readable labels, as opposed to IP addresses or hostnames, and you're able to enforce segmentation down to the process level.

So, why is now the time for this kind of segmentation? Lots of reasons, but here are three biggies:

- » **Cyberattacks are unavoidable.** The stats show this to be true, yet for many organizations there's a surprising lack of preparedness.
- » **Cybersecurity mindsets are often outdated.** Despite continued investment in perimeter controls, organizations still get breached. When you concede that breaches are inevitable and start to assume breach, you can then focus on isolating breaches and stopping their spread. ZTS is the fastest and easiest way to do that.
- » **ZTS works today and tomorrow.** It's a flexible and forward-thinking approach that withholds trust by default, enforces the least-privilege concept, provides for comprehensive security monitoring, and achieves the best outcomes going forward.

Using Zero Trust Segmentation

There are plenty of use cases for which ZTS is an ideal solution, and you can find more details on them in Chapter 7. Uses include very specific situations, as well as a more general achievement of continual security and the ability to respond quickly to threats.



REMEMBER

Ransomware containment is of ever-growing importance and is a use case where ZTS really shines. It's a key to successful incident response, allowing you to quarantine parts of the infrastructure that may be compromised and create clean bubbles within the environment.

This capability is essential for preventing the lateral spread of malicious activity and ensuring that a small breach remains small. The more effectively you can contain the damage, the shorter the road to recovery.



REMEMBER

With ZTS, it's possible to act incredibly quickly to deal with a breach. Whereas organizations in the past spent valuable time trying to understand the intention of malware, the faster approach is just to slam the barn door shut, and then deal with the intruder. ZTS means it's only a very small barn whose door has been shut.

Indeed, it's not just a matter of quickly slamming the door shut when a breach is discovered. ZTS ensures that the barn door is open only when it needs to be, and only for the right, predetermined reasons.

All this helps tremendously in keeping the organization as a whole up and running while the malicious activity is purged from the compromised area. The scary headlines you read usually involve attacks with a massive impact, those that shut down a whole supply chain or force a health-care system back to the days of paper charting for days or weeks.

Perhaps the best way to describe this use case is a desire for greater resilience. That's often what executives on the business side are hoping to achieve with this kind of initiative, usually after they hear about something that happened to someone else in the industry.

The desire is to make your organization's architecture not just more secure but more resilient, able to recover more quickly from the inevitable attack. For a health-care organization, that means ensuring that an incident sidelines just one medical device, not all of them. For a beer manufacturer, it's making certain the hops keep hopping at most of the production plants, even if there's an infiltration into one area.



REMEMBER

Beyond enhancing resilience in the face of ransomware incidents and other attacks, ZTS is handy for a number of important use cases. For example:

- » **Heading to the cloud:** ZTS helps your organization safely migrate workloads from an on-premises data center to the cloud, or from one cloud location to another. You can

make the move in a secure manner, with full visibility and the assurance that you'll maintain consistent security policies from one location to the next.

- » **Integrating IT and operational technology (OT):** The convergence of IT and OT is inevitable. Manufacturing systems, industrial equipment, energy systems, health-care technology, and other operational technologies are constantly becoming more integrated into IT, which is essential for improving operations but potentially risky from a cybersecurity perspective. ZTS helps ensure that neither IT nor OT is overexposed as they become more integrated.
- » **Narrowing the trust radius:** Back in the day, trust stopped where the internet began, enforced by a firewall, router, or switch. Today, the perimeter has withered and IT environments are much more complex — and it's clear that attacks can find their way inside through often unwitting but honest insiders or unintentional gaps created by rapid expansion and digital transformation. For this reason, building a least-privilege access model is essential, narrowing the blast radius to the workload level.
- » **Aligning security hygiene:** You want to ensure your organization follows best practices for vulnerability management, risk mitigation, path management, and other activities. ZTS plays an important role in this.
- » **Keeping watch over software interactions:** There's a lot of open-source code out there, amid custom builds and as part of commercial software. There's also lots of connectivity between components and infrastructure, and the pandemic accelerated the evolution. You need visibility into and control over how software interacts with the infrastructure where it lives, and ZTS offers that.

Transforming the Business



REMEMBER

The most successful businesses are the ones that are able to transform themselves, their operations, and their missions to meet the changing needs and expectations of their customers. IT has been a powerful enabler of transformation, and that has never been clearer than during the recent pandemic years, when

the transformation of many common daily activities accelerated dramatically.

Suddenly, large percentages of employees were working from home. New models allowed delivery of practically anything to nearly everywhere. Cashless and contactless interactions became highly valued. There was increased automation in banking, and more people were downloading digital hotel room entry cards and digital rental car keys.



WARNING

But IT transformations also can open the door to new and sometimes not well understood risks. Businesses sometimes transform more quickly than is safe, automating and connecting their way into new vulnerabilities.

The role of those working in cybersecurity is all about protecting the business from cyber incidents, but these professionals also must balance security priorities with business objectives. The competition isn't waiting, so the security team can't always say "no" — in short, they must ensure high levels of security and increased resilience, without standing in the way of transformation.



REMEMBER

There are four basic ways the business transformation picture can pan out:

- » **Digital conservatives:** These organizations put a low priority on digital transformation, and because of that feel like they're pretty safe from a cybersecurity perspective. Spoiler alert: Not only are they light on transformation, but if they haven't focused on cyber resilience, they're not as safe as they think.
- » **Security blockers:** These are the organizations that take cybersecurity resilience seriously, and as such they're safer than the digital conservatives. But their views on what it takes to be safe hinder their ability to transform.
- » **Cyber sprinters:** These are the organizations so gung ho on digital transformation that their revolution gets out ahead of their security. They're rocking it with their transformative ideas, while rolling into significantly risky territory. The challenges and resulting opportunities of the pandemic certainly encouraged a lot of cyber sprinting.

» **Cyber leaders:** These organizations have found the right balance to transform business successfully and safely. They're pursuing powerful digital transformation but are also serious about cybersecurity resilience. ZTS is an ideal way to achieve the balance needed to be a productive but safe business — a true cyber leader.

How does ZTS strike this balance and enable transformation? By helping the security deployment mature from a reactive approach to being proactive.

With a least-privilege model, the security team gains an understanding of what access is needed for a transformative capability to function, and then transforms access to meet that requirement. The security posture is moving hand-in-hand with the technology capability posture as it evolves.

Compared with the security blocker organizations, the cyber leaders realize they're likely to be hit with some sort of attack, so they plan for it. The blockers aim to stop all attacks and end up stopping transformation, too. The cyber leaders figure out how to be resilient so they can carry on when an attack happens.

Establishing Security Objectives

You generally don't start a journey without knowing where you're going. For businesses that have determined to go down the ZTS path, one of the first tasks is to establish security objectives.



REMEMBER

Your organization must determine your security posture in order to strengthen it. And you'll never be sure if you're winning if you don't first figure out what success looks like. This is one of those places where you shouldn't let the perfect be the enemy of the good, because implementing ZTS is really a matter of evolution rather than revolution.

Yes, there may be a very impressive, grand end state that you'd like to reach. But many projects fall flat because they focus on that end state as the only measure of success. It's better to look for some smaller wins before aiming for the championship. There's no shame in going for some low-hanging fruit.



TIP

Visibility is a great initial aim. It's an essential part of improving your security posture and moving down the road from reactivity to a more proactive posture. Visibility helps you assess your risks, and for a lot of organizations, gaining visibility everywhere feels like a big win in and of itself.

Another good step is achieving an improvement in ransomware protection. Still another is establishing environmental segmentation. For some organizations, the ultimate success means ensuring that every single application is properly ring-fenced, but that kind of success is likely to be further down the road.



TIP

What level of protection is desired is a key question to determine when setting objectives. A ZTS solution will establish rules for what's allowed and what's not, and program that security policy at the appropriate control point to protect individual workloads.

Full protection means only explicitly allowed connections are able to be established, while all other unneeded communication is denied. Needless to say, full protection is much more comprehensive and involved, and it takes longer to get there.

The point is, every organization has its own concept of what success looks like. Success in ZTS doesn't necessarily mean every access point is immediately operating on a least-privilege model, but it does generally mean the infrastructure is moving in the direction of increasingly less implicit trust.

It's also possible to reach a point where the cost exceeds the benefit. At that point, the organization decides it has done enough, and success has been achieved.

Establishing security objectives involves a conversation across multiple parts of the organization. All the various stakeholders need to align their goals and objectives to avoid frustration or dissatisfaction with the process.

Becoming more resilient is often a general aim, and it's a good one, but it's a journey. The key is not to get too hung up on the perfect state. With just a series of relatively simple steps, you can greatly reduce risk, proactively limit communication laterally to reduce the impact of a cyber event, and boost your ability to respond to threats.

ZTS AND THE NIST FRAMEWORK

The National Institute of Standards and Technology (NIST) is a federal agency within the U.S. Department of Commerce. Its work is all about innovation and competitiveness, and its experts have a keen interest in helping transformation happen safely.

No surprise, then, that NIST has something to say about Zero Trust architecture. And given that NIST is a highly technical government agency, it's also no surprise that its advice carries titles you would expect on an important government document: Special Publications 800-207 and 1800-35.

NIST has identified core logical components of a Zero Trust architecture:

- **The policy engine:** This is what decides what kind of access to grant a specific resource for a given subject. The policy engine is guided by enterprise policy, as well as external input from such sources as diagnostic tools and threat intelligence services.
- **The policy administrator:** This component is guided by the policy engine to either establish or shut down communication paths between subjects and resources. The administrator and engine together make up what's sometimes known as the *policy decision point*.
- **The policy enforcement point:** This is like the actual guard at the door, and it takes orders from the policy administrator. When the administrator decides to allow or deny a session, the enforcement point enables or terminates the connection.

Your ZTS solution should have an architecture that maps into these NIST-outlined components. It needs a "brain" that will serve as the decision point, doing the work of the policy engine and policy administrator. And that brain should be able to control the enforcement points in the environment.

IN THIS CHAPTER

- » Getting eyes on the attack surface
- » Paying attention to context
- » Knowing what assets you're protecting
- » Deciding how to prioritize risks

Chapter 3

Seeing the Attack Surface

It's war out there, war that pits you against nefarious cyber actors. As with any war, you need to be able to defend yourself, and to do that, you need to know where you're vulnerable. This chapter discusses the attack surface that you're trying to protect, what the assets are that the bad actors are seeking, and how the context of the activity fits in. It also helps you pinpoint which risks are at the top of the list.

Realizing the Vulnerabilities



REMEMBER

The whole pursuit of cybersecurity is preventing and dealing with attacks, and the *attack surface* is pretty much the sum total of the things that you're trying to protect from an attacker. The attack surface refers to all of the organization's information technology (IT) assets that are potentially exposed to the bad actors.

What do I mean by “exposed”? I'm talking about assets that may have physical or digital vulnerabilities that an unauthorized user can leverage to gain access to the network and extract or harm the data there.

By the way, people can be considered part of the attack surface, too. That's because they're frequently targeted using phishing emails and other social-engineering efforts. Just like your IT assets, your people need to be ready to defend themselves and your organization.



REMEMBER

Here are the types of attack surfaces you need to be aware of:

- » **Digital attack surface:** All computers and servers and other devices that are connected to the internet are exposed to attack, which means the things running on them are part of the digital attack surface. Examples include websites, databases, operating systems, applications, cloud resources and workloads, and the services of your third-party providers.
- » **Device attack surface:** This refers to all of your organization's hardware and physical devices, and any employee devices that may connect to the corporate network. This includes workstations and laptops, mobile devices, routers and switches, TVs and security cameras, even printers. Accessing a device can potentially allow an attacker to move laterally across the network to access other targets.
- » **Social-engineering attack surface:** Your people, even your most honest employees, can be security risks when they're targeted by *social-engineering attacks*. That's a fancy term for tricking your people into doing something they wouldn't otherwise do. Email phishing tries to get them to open a malware-infected attachment or click a malicious link. An employee may inadvertently let an attacker in the door literally, too, disguised as a repairperson or custodian. Or someone may plant a USB drive that an employee plugs in, wanting to see what's on it and not realizing that the answer is malware.

Seeing What You're Protecting



REMEMBER

You may find it a bit scary to consider just how broad the attack surface is. But as with so many other parts of life, you can't really address a problem until you recognize that you have a problem and understand the extent of it. That's what *visibility* is all

about — fully understanding what you have to protect and gaining the ability to see what’s happening.

The journey to greater security begins with a careful mapping of the attack surface, painting a comprehensive picture of the potentially vulnerable places and documenting just what makes them vulnerable. That means taking into account all potential *attack vectors*, which are the paths that attackers use to breach the network. Here are some examples:

- » **Compromised credentials:** When usernames and passwords fall into the wrong hands, bad things happen. The most common problem is when employees become phishing victims and enter their logins on fake websites.
- » **Weak passwords:** It could be the old standby passwords of “123456” or “password” or other weak efforts. Or when people use the same password in too many places.
- » **Insiders with ill intent:** This issue is less common than unintentional actions by honest employees, but sometimes unhappy employees expose confidential information on purpose.
- » **Ransomware:** This kind of attack usually locks down your data so only the attacker can unlock it, after you pay the ransom, of course.
- » **Firewall-related errors:** Misconfigure a firewall, and bad actors can get inside. The same problem can happen if a publicly facing workload is misconfigured. It’s important to regularly audit firewalls and publicly facing workloads.



TIP

A good understanding of the attack surface puts the spotlight on better ways to protect the organization by reducing the attack surface. Here are some of the potential approaches:

- » Putting segmentation in place to curb lateral movement by attackers who find their way inside
- » Getting a handle on overly permissive access rules, so employees are able to access only the assets they need to get the job done
- » Training employees to recognize and resist social-engineering efforts targeting them, filtering their internet activity to keep them away from dangerous sites, and requiring safer, more complex passwords along with MFA

- » Scanning and updating regularly to proactively address vulnerabilities
- » Adding or increasing encryption
- » Finding and fixing misconfigured cloud security

Putting It All in Context

Context is key when you're aiming to establish the right balance between accessibility and security. Your Zero Trust Segmentation (ZTS) solution needs to consider much more than just who's talking to whom.



WARNING

If all you know is that this particular resource is communicating with that other resource over there, you can't really draw any solid conclusions about whether that interaction is allowable and safe. You need to know what the resources are, what function they're serving, what the application or user is, where it is, and more.

Consider, for example, the Remote Desktop Protocol (RDP) port. A lot of vital interactions make use of the RDP port, but it also happens to be a popular pathway for ransomware. If you can remove RDP or effectively restrict its access, you can stop a whole lot of ransomware in its tracks.

That's where context comes into play. It's how your system determines whether there's a valid reason for an RDP connection. Here are some of the most common context considerations:

- » **Role:** What function does this server have in the application stack?
- » **Application:** What's the business application or function?
- » **Environment:** Is it in production now, in development, or staging?
- » **Location:** Where is the physical location of the resource or user?

With effective ZTS, all your resources are labeled by these kinds of context details. The point is that if any of these things change, the context changes. And if the context changes, whether an interaction is allowable can change.

That makes context much more useful than, say, an IP address or a hostname of a workstation. A firewall rule may apply to a particular workstation but doesn't take into account whether the user changes workstations, job roles, or workloads.

Security policy that pays attention to context can do a much better job of allowing the right activity and stopping the harmful activity. It allows you to write rules tightly, so potentially risky connections can be restricted only to authorized sources for authorized functions.

Identifying Assets

Ultimately, if you're going to protect your applications and data from intruders, you have to fully understand how they interact and interdepend. And you need to understand it better than your would-be foes do.



TIP

That requires a comprehensive effort to identify assets and map application dependencies. It's a key part of the visibility I talk about in this chapter.

In today's complex, hybrid architectures, your assets are all over the place. Your watchful eye needs to gaze over data centers, bare-metal servers, endpoints, the cloud, all your operational technologies, and everything in between. It needs to know everything there is to know about highly dynamic environments, including virtualized, containerized, and cloud platforms across which applications can migrate.

The process of *application dependency mapping* provides full visibility into the application topology. It lets you know where your highest-value assets are sitting within your environments, and what kinds of traffic flows can reach them. It also shines the light on security policies that are monitoring and controlling the traffic flows, which helps you see where the holes are.

Taking an application-sensitive approach to security is far more powerful than simply scanning the infrastructure. Understanding how applications work and interact can provide greater insights into how they may be targeted, which allows better planning for isolating and protecting them.

Prioritizing Risks

If you've made it this far in the book, you get the point that risks are all over the place. Across a big and complex network, many highly important interactions are taking place — interactions that are not only crucial for innovation but also invitations to trouble.

Some threats are more pressing than others, and you'll get the most out of your security efforts if you assess vulnerabilities to prioritize the ransomware risks. Your IT team has only so much bandwidth, so it needs to establish risk priorities.



TIP

Augmenting your visibility with rich data about communications helps in this effort. Stopping malware in its tracks means blocking the ports it wants to use to jump from one application or machine to another. But you also need to know which ports must stay open between systems to keep essential services operating.

Real-time visibility into communication between assets will help identify needlessly open ports. It can also help identify which ports carry the greatest risk. For example, highly connected ports used by Microsoft Active Directory and other core services carry a high level of risk. Systems polling and reporting on IT infrastructure can also bring risk that needs to be prioritized. Peer-to-peer ports also can be significantly risky — that includes the ones used for such things as remote desktop management and file-sharing applications.

Another high-priority issue to address is older communications protocols such as File Transfer Protocol (FTP) and Telnet. These protocols may not be used a lot anymore, and security policy may even limit their use. Yet they tend to stay on by default, which makes them attractive targets.

Vulnerability mapping should yield *vulnerability exposure scores*. That approach makes it easy to determine where the greatest risks are — pointing out where first to focus segmentation efforts.

IN THIS CHAPTER

- » Seeing how cybercriminals move across the network
- » Achieving visibility
- » Stopping the spread of ransomware
- » Becoming more resilient
- » Containing the damage

Chapter 4

Responding to and Isolating Ransomware

If you're going to thwart attackers, you need to know how they move. And you must have a good handle on what's happening across your network.

This chapter discusses lateral movement, the preferred path of cybercriminals, and how you can stop their ransomware from spreading. It shows the power of visibility into your information technology (IT) environment, explains how the overarching aim is resilience, and explores how to contain damage in operational technology (OT), too.

Moving Laterally

Having a cybercriminal break into your network is bad enough. But what's really bad is what can happen next: *lateral movement*. *Lateral movement* is just a fancy way to say cybercriminals and malware want to move or spread across your organization to find valuable data and assets.



WARNING

After they step through an endpoint and get inside the network under the guise of an authorized user, hackers aim to move deeper into the system. They're looking for sensitive data, intellectual property, or other high-value assets. They're moving from one asset to the next, trekking through the compromised system, often stealing advanced user access privileges along the way.

That's lateral movement, and it's a core tactic of cyberattackers. It's one of the ways that today's cyberattacks can be more complicated and damaging than the breaches of yesteryear. One of the big aims of a breach containment strategy is detecting lateral movement quickly and putting a stop to it.



TIP

To understand the problem of lateral movement, it helps to consider what attackers are looking for as they make their lateral moves. Here are some of the common aims and things they're seeking:

- » They may be stealing intellectual data such as a project's source code from a developer's work device.
- » They may be trying to tap into an executive's email to steal banking details or learn sensitive company information that could be used to manipulate or take advantage of stock prices.
- » They may be looking for other credentials or avenues for upping their ill-gotten privileges.
- » They may be looking for customer data, including payment card information.
- » There may be some other kind of company-specific asset or valuable payload they're after.



WARNING

Whatever it is they're up to, they're making moves through the network toward their ultimate goal. In making those moves, they're typically following three common stages of lateral movement:

- » **Reconnaissance:** This is an early stage in which the attacker is basically looking around. That means exploring and mapping the network, devices, and users. It's all about becoming familiar with network hierarchies, host naming conventions, operating systems, the location of payloads, and insights for making additional lateral moves.

- » **Gathering privileges:** The best way for an attacker to get around a network without being noticed is to look like they belong there, which means obtaining valid login credentials. Illegally obtaining network credentials is often called *credential dumping*, and cybercriminals often achieve this through phishing attacks or *typosquatting*, which involves setting a trap with a URL that looks legit but may be a slight misspelling or typo of the real site.
- » **Hopping around:** Once inside, the attacker gains access to other communication and computing points in the network, bypassing security controls and compromising even more devices. That is to say, moving laterally again and again, until the attacker either gets the job done or is detected and gets stopped.

That middle step of gathering additional privileges can help the threat actor hang around in the system for an extended period of time, even if the IT team has discovered the initial infection of the system. It allows free movement with ongoing access and no detection.



WARNING

Hanging around and moving freely is what's known as *dwelt time*, and you'd be shocked how long that dwell time can be. Lateral movement through the system can sometimes go on for weeks or even months after the initial breach, with the system open to data theft. One study found that it can take 200 days or longer to detect a phishing attack and many more days to contain it.

Identifying the Unexpected

That possibility for a leisurely dwell time is one of the big reasons attackers use lateral movement as a tactic. There can be lots of time to look around, tap into treasures, and gather new credentials. And with those ill-gotten but valid credentials, it can appear that it's a legitimate user making the rounds of the network, rather than an intruder.

Gaining full visibility into the attack surface is one of the best ways to identify these unexpected and unwelcome threats. Understanding all the systems, devices, and applications across your network will help you gauge what needs to be protected and properly

segmented. You need to be fully aware of potential attack paths, as well as exposed credentials and any misconfigurations.



REMEMBER

This is what *risk-based visibility* is all about. It's identifying which systems and applications are made vulnerable by too much unnecessary communication, as well as noncompliant data flows.

Application dependency maps paint a picture of the application topology, making it easy to see relationships between applications and understand how protocols are working across the production environment. Vulnerability data can provide further insights by generating quantitative risk scores.

Stopping the Spread

Ransomware is scary and dangerous, but the good news is that it's also predictable in the way it moves. That predictability makes it possible to take proactive steps to stop its spread.

As outlined earlier, malware gains entry through some vulnerable pathway. Then, if it's unnoticed and unchecked, it can spiderweb its way laterally across networks, devices, and servers. When it's in place, it gets activated and starts to wreak its havoc.



TIP

Here are three primary steps for stopping the spread of ransomware.

Eliminating unnecessary connections

This is how you isolate and protect your critical assets. The aim is to eliminate unnecessary paths between devices and networks, allowing only necessary communications, and tightening workflows.

For example, consider a videoconference. The host's laptop has no need to talk directly with another device logged into the meeting via Remote Desktop Protocol (RDP) or Server Message Block (SMB) ports. Closing those ports will strengthen security without having any impact on the videoconference.

You can achieve this aim by blocking all communications across individual ports, as suggested in the previous example. You can do that for a single application or in a certain geographic location, or you can do so across the whole network.



REMEMBER

When it comes to closing inbound and outbound ports, you can do so both proactively and reactively. Working proactively is kind of like locking the door at night — you're taking action just in case, preventing something that might happen.

After you've identified your highest-value assets and applications, isolate them in protective rings. That's a proactive move that keeps them safe from malware that intrudes elsewhere.

Working reactively means having policies in place in case of a breach or suspicious activity and then responding when incidents occur. Note that while the action happens reactively, this still requires proactive planning.



TECHNICAL
STUFF

A few quick thoughts:

- » **There's rarely any need for nonmanagement servers to use vulnerable ports that enable peer-to-peer communication, such as RDP and SMB.** Those ports are the most common vectors for malware, so shutting them down will do a world of preventive good.
- » **Ports used by databases and core services, such as those bundled into Linux applications, are often old and vulnerable.** Risk-based controls and policies on the inbound side of machines in the data center can help fix those vulnerabilities. And most data center servers have little or no need to talk to the internet.
- » **When you tighten communications to prevent unauthorized data from leaving the organization, that can stop ransomware's command-and-control function.** That keeps the ransomware from being triggered from the outside. You can do the same thing with cloud systems and user access permissions that are broader than necessary.

Using visibility

It's important to gather connection data and flow information from your infrastructure, including on-premises routers and switches, as well as clouds and other end-user systems. That helps you create application dependency maps.



REMEMBER

With this visibility, administrators can make solid decisions about which assets should be talking to each other and which should not. That leads to proactive policies that serve to contain critical assets and systems. These policies can work across machines, network switches, cloud-native firewalls, and other areas.

The end goal is ensuring that if there's a breach between two users, it won't affect any other users or assets in the cloud or the data center.

Improving responses to intrusions

A lot of what I outline earlier is preventive in nature. But you also need to be able to react quickly and effectively when intrusions happen — after all, you're working on the presumption that they *will* happen.



TIP

If you spot suspicious activity, you may want to immediately put up barriers around core databases, payment systems, medical records, or other sensitive things. That's a potentially more restrictive containment capability than you'd want to run day-to-day, and it requires coming up with secondary policies to activate as part of incident response, to stop malware in its tracks.

This kind of capability is a powerful blend of proactive and reactive. Making it happen requires planning in advance for the day it occurs. It also requires the ability to spot traffic flowing in unexpected ways. That means visibility into your workloads wherever they're running, whether in the data center, the cloud, or wherever. A visual map shows you where you need to slam the door.

Building Resilience

Not to beat a dead horse, but it's worth reiterating that stuff's gonna happen. Yes, it'll always be important to guard the gate, but the real name of the game going forward is resilience — being able to withstand an attack and bounce back quickly. There are four core principles for building resilience:

- » **Watch your internal communication flows.** Visibility into communication flows is the key to spotting ransomware attacks early, before they spread too far. It's also a great way

to spot pathways you didn't realize were open, before the bad actors notice.

- » **Block ransomware pathways.** Most ransomware attacks exploit a small set of high-risk pathways, including RDP and SMB. If these services are left open unnecessarily, they offer attackers an easy path into and around the network. So, block them when you can. Only open the ones that really need to stay open, keep a close watch on them, and be ready to lock down the environment when you spot suspicious activity.
- » **Protect your assets.** The assets you value the most are also the ones attackers value highly. They'll start with a lower-value asset, work their way across the network, and often hide out until they find the gold. So, always protect those high-value assets by segmenting their environment, surrounding them with rings and fences, and putting roadblocks in the paths to them.
- » **Use the right tools.** Manual firewalls and similar network segmentation tools were created years and years ago, for architectures and threats very different from today's. Those three recommendations listed directly before this one? Older tools usually don't do those things. You need modern tools designed to protect today's environments from today's (and tomorrow's) threats.

Containing Operational Technology Breaches



REMEMBER

With all the talk about IT security, sometimes OT is overshadowed. But depending on the nature of your business, you may have lots of special-purpose devices on IT networks doing mission-critical things. They may be sensors, motors, programmable logical controllers, remote terminal units, and a host of other devices that drive operations in such places as factories, medical facilities, power plants and utility grids, and other vital places.



WARNING

If they do vital work, you can bet they're valuable targets for cybercriminals. Sometimes OT devices are *air-gapped*, meaning they're not physically connected to IT networks. But these days there are a lot of powerful uses for OT that require them to be connected. In other words, they require IT/OT convergence.

The same kinds of principles that protect IT environments can protect the OT world, too. That includes visibility into activities, communication, and configuration details. That's a key for spotting and preventing lateral movement in the OT environment.

You also need a Zero Trust Segmentation policy engine that can separate everything that needs to be separated, IT or OT. And then you need to be able to enforce the policies without shutting down business-critical IT and OT networks. The goal is to keep the business running, even in the face of trouble, without having to put into place controls that get in the way of performance.

- » Limiting the damage
- » Getting started with Zero Trust Segmentation
- » Acting with urgency
- » Keeping apps available

Chapter 5

Protecting Applications

Your applications and the data they process are highly valuable to you and to the criminals preparing to attack you. Your business must keep moving forward even as attacks happen, so you've got to protect your applications.

This chapter shows how to shrink the impact of an attack and how to know where to start your journey into Zero Trust Segmentation (ZTS). It outlines how to act quickly and smartly and points out the importance of maintaining app availability even as you boost security.

Shrinking the Blast Radius



WARNING

As I mention in Chapter 4, the things you value the most will also be of the most value to potential attackers. This includes applications delivering critical services, the places where you keep your sensitive data or personally identifiable information (PII), and other resources that are regulated by various compliance mandates such as the Sarbanes-Oxley Act or the Health Insurance Portability and Accountability Act (HIPAA).

This is why application segmentation is critical. It's how you control lateral communications between applications or application

tiers. Ring-fencing protects high-value resources running on bare-metal hypervisors or containerized workloads in your data centers, public clouds, or hybrid environments.



If it sounds complicated, it can be — but it doesn't have to be. With the right tools, it's surprisingly simple to create policy and deliver a single control plane for creating and operationalizing security across network perimeters. You can decide how granular your segmentation needs to be — coarse-grained or micro-segmented. And you can meet compliance requirements without re-architecting the network.

Figuring Out Where to Start

Let's just say up front that perfection takes time. But ZTS isn't an all-or-nothing thing. As mentioned in Chapter 3, it's not only possible but often advisable to begin with high-value applications, as well as the easiest wins.



The most obvious high-value assets are things such as customer account details and other PII, as well as payment systems and other financial assets. Also high on the list should be resources that fall under any kind of regulatory compliance mandate. These are all the kinds of things that, if exposed or exploited, could cause significant financial loss, operational disruption, or reputational harm.

But don't forget about assets that may be high-value but temporary. Businesses have, after the fact, realized that something like a holiday discount coupon or Cyber Monday promotion can be leveraged for nefarious purposes. These things may not normally be included on the list of the "crown jewels," but they may still have significant value to an attacker.

Not only that, but because they're a bit more fleeting than a permanent application or infrastructure element, they may not get as much testing for security before they head into production. These assets may not need the same level of protection as other items high on the list, but don't let them fall through the cracks.

Making Quick but Smart Decisions

In the world of oncology, it's important to make the right diagnosis in order to figure out the right treatment. That may require imaging, biopsies, and other diagnostic approaches, all of which can take some time.

IT security has in the past sometimes had a similar mindset. Spot something suspicious, and then try to figure out what exactly is going on in order to plan the most appropriate response. That approach may seem sensible, but these days, it's not fast enough.



TIP

Shouldn't the first priority be stopping an attack from spreading further? The ZTS approach does just that. High-value assets across the organization are protected individually, in a way that is as manageable as it is comprehensive.

Not only that, but with a containment switch, the right tool can very quickly slam the doors shut at the first sign of potential attack and keep an attack from moving anywhere else in the network. When suspicious activity is detected, that switch can be flipped — either manually by the security team, or as part of a script, such as a Security Orchestration, Automation, and Response (SOAR) playbook.

Flipping that switch can, in seconds, isolate the attack right at the point of entry. If it's ransomware, it'll get no further, giving analysts time to figure out what it is and how to remediate it.

And in the meantime, an application dependency map shows what traffic is needed to keep vital business operations up and running while the attack is mitigated. That traffic can continue to flow, and life can go on.

It's a simple-sounding concept, but of course, an idea this powerful is bound to be complicated. There are lots of policies to create and enforce, to ensure both proactive protection and reactive capabilities.

That's where the right tool can make all the difference. Protection must be comprehensive, but if it's not manageable, rules could be overlooked or misunderstood. When choosing a solution, pay attention to its design to learn how straightforward and manageable it will be to implement.

Maintaining App Availability

Toward the beginning of this book, I talk about the old days of defending the perimeter, building a moat around the premises that would effectively keep bad actors out. That was a perfectly fine approach at the time.

Today, of course, continual business transformation requires applications to cross that moat and live all over, not just on-premises but in cloud or hybrid environments. They must be easily accessible to people as well as other apps, instantly at-the-ready and absolutely always available.



WARNING

These are the apps that are critical to the business. You're sunk if they're compromised by bad actors. But if your security compromises the accessibility of these apps, you're just as sunk. Your aim is to improve security around these critical apps that are already running across your organization, in ways that won't compromise availability.

ZTS meets that need from a number of perspectives. As outlined elsewhere, it allows high levels of protection along with the ability to know what communications are essential and let them through. It has compartment walls to limit the spread of malware, and carefully controlled doors that maintain availability and accessibility.

Beyond that, a solution with good visibility allows organizations to proceed with their security initiatives in a way that lets them see the impact of steps they're about to take. Before they press the "go" button, they can determine whether a policy will succeed in its security aims without breaking the application.

IN THIS CHAPTER

- » Improving your protection, then improving some more
- » Knowing that your work is successful
- » Gauging where to go next
- » Starting early with security

Chapter 6

Keeping Ahead in the Game

Zero Trust Segmentation (ZTS) can put you in a much better and more resilient place, but it takes constant effort to keep from falling behind. This chapter spotlights the need for continuous security improvements, discusses how you measure success, offers advice for next steps, and shares insight on why security needs to be an early consideration in the development cycle.

Building a Loop of Optimization



REMEMBER

Like most positive change — getting more exercise, improving your diet, making better financial decisions — implementing ZTS is not a one-and-done thing. It's an ongoing pursuit. From this point forward, you'll identify new opportunities, design new remedies, test the efficacy of your microsegmentation, and implement and validate your solutions. And then you'll repeat the whole cycle over and over again.

A number of factors are driving this need for a loop of optimization. The environment is constantly changing. Servers are coming

online and offline. Developers are continually improving your apps and occasionally adding new ones. There may be mergers and acquisitions that alter the landscape.

All these kinds of changes alter your security picture and require a rethink of policies. Just because your security controls are valid today, doesn't mean they'll be valid tomorrow. As the infrastructure evolves, so will your security. The good news is, as you bring new elements into the picture, you can protect them from the start.

Beyond that, increased visibility drives better awareness of the environment. The more you look, the more you see that may need attention. That means more policy changes and more trips around this loop of optimization.

Then there is the reality that your Zero Trust implementation isn't just a matter of flipping one switch and it's in place everywhere. Your work will mature over time; you'll get greater resilience in place in some key applications and then drive it to others.

Validating Security Outcomes

How do you gain confidence that the security measures you've put in place will achieve their desired outcomes? Good question. Certainly, every day that passes without your organization being debilitated by an attack is a hopeful sign. But living life on the edge awaiting potential catastrophe is not a fun way to live.



REMEMBER

Fortunately, there are many ways to validate your security outcomes and prove that they're delivering the results you're expecting. This can range from basic testing of security controls all the way through penetration testing to see how far individuals may be able to get through your network. For example, through pen testing you can check to see how your limits on Remote Desktop Protocol (RDP) will prevent an intruder from moving from low-value assets to a higher-value database or active directory.

Validating security outcomes requires a full grasp of adversarial tactics, techniques, and procedures that bad actors typically follow when they mount an attack. Visibility into the attack surface and lots of relevant event data can help you determine how successful your platform is in detecting and defeating attacks.



TIP

Continuous testing of this nature is an essential part of your Zero Trust DNA. It's worthwhile to occasionally bring in an independent "red team" to play the role of attacker and see how your "blue team" is able to respond. Indeed, it's now possible to automate this kind of red team versus blue team exercise, so you're always monitoring the effectiveness of security controls and seeing how they would fare against sophisticated adversaries.

Deciding What's Next



TIP

The farther down this path you go, the easier it will be to find your way and figure out what needs to be improved next. For starters, though, set up security milestones for your project and work toward those outcomes.

What you'll inevitably find is that as you gain visibility, you'll start identifying other things in your environment that you weren't aware of before. That's a sign you're on the right path.

You may, for example, gain a perception that a particular service is no longer being used. You may gain visibility into technical debt or perhaps a risk you didn't know about. From that will emerge additional policy enforcements and then more awareness.

In Chapter 3, I discuss risk prioritization, and that's another road map for deciding what comes next. You're prioritizing your work based on risk and exposure and the probability that a particular risk will be exploited.

And then, as I mention earlier, the constant change will rear its head and drive your decision on what's next. Stay observant and flexible, and the way forward will become apparent.

Shifting Security to the Left

Security has long been something of an orphan child in the information technology (IT) environment. It has been kept in the background, rarely invited to be in the first phase of a deployment. Developers have had a wary view of security, feeling like it's always slowing them down. One day a breach happens, and then people wonder why better protection wasn't in place.



TIP

That's why security needs to keep shifting toward the left, becoming more involved in the earlier stages of the IT development life cycle. Developers may or may not realize it at first, but ZTS has the opportunity to be more of a blessing than a bane.

Consider the reality that developers may tap into open-source code as part of their work and may do so without really knowing for sure if malware has been deployed within it, asleep and waiting to be awakened. Clearly, you can't live indefinitely with malware lurking in the shadows.

The good news is that with Zero Trust considered from the beginning of the development cycle, you're moving through the process with a resource already protected. As you stand up environments, you're doing so with security already in place.

That can take the worry level down a notch, knowing that the spread of any malware will be halted right away. And it can free developers to develop more quickly, without being bogged down by security solutions that get in the way of app availability or performance.

In addition, developers can benefit from the visibility that's enabled by a strong ZTS solution. Visibility brings a greater understanding of risk from the get-go, which enables the safest possible decisions on the left side of the life cycle.

In short, this is all part of the move from being reactive to becoming proactive. You're modernizing the way you develop applications by including and implementing security controls as part of the application development life cycle.

Shifting left means moving security from a thing that gets bolted on later to an integral element that's incorporated throughout the life cycle. Virtual machines, for example, are brought up with least privilege in place from the start, and applications are launched from inside a safe ring-fence. Security is no longer an afterthought or a latecomer — everything is born secure.

IN THIS CHAPTER

- » Keeping ransomware from getting ahead
- » Working in a secure cloud environment
- » Protecting operational technologies and supply-chain integrity
- » Securing your most important assets
- » Proving your compliance
- » Responding effectively to an attack

Chapter 7

Ten (or So) Use Cases for Zero Trust Segmentation

What's Zero Trust Segmentation (ZTS) good for? In a general sense, creating a safer, more secure information technology (IT) environment and increasing your organization's resilience. To get a bit more specific, there are several key use cases for which ZTS is a great approach. Here's a rundown of commonly appreciated use cases.

Isolating Ransomware

It's going to happen. Ransomware will get past even the best prevention tool. Ransomware may even enter from inside the organization. When it does, it could be months before it's detected, and during that time it'll head east or west, moving laterally in search of ever-higher-value targets. That is, unless you can contain it.



REMEMBER

Isolating ransomware is the first and most important use case for ZTS. By segmenting off workloads from one another and from the outside world, it's possible to contain ransomware near its point of entry. It's quite possible that the initial breached workload will

be a malware victim — but through segmentation, blocking the most vulnerable ports, and strictly managing communications, the blast radius will be small and the rest of the business will go on uninterrupted.

Securing the Cloud

In today's diverse and often hybrid environment, you're likely to encounter different segmentation tools from one data center and cloud to another. That causes siloes in your security operations, making the task of security too complex and more prone to potential errors.

The right ZTS solution will focus directly on the workload and be agnostic to the fabric. Any data center SDN security tools won't work in the cloud, and cloud enforcement endpoints are likely to be breached at some point. So, your solution should be able to live-migrate workloads from one place to another, carrying ZTS enforcement with them.

Migrating to and from the Cloud

You'll find that there are cloud workloads for which you can't deploy agents, or for various reasons you don't want to. That doesn't mean you don't need workload visibility and enforcement.



TIP

The right ZTS solution will be able to extend label-based visibility and policy from the data center to the cloud, without agents. And it'll centrally manage them in the same way as it would workloads with an agent.

Integrating Information Technology and Operational Technology

Your IT and operational technology (OT) are increasingly living in connected worlds. Telemedicine and other connected concepts are driving this convergence in health care. The supply chain is increasingly automated, deliveries are optimized, production is

nonstop, and your enterprise resource planning (ERP) is integrated all over the place.

It's a powerful trend, but problematic. For example, a breach in hospital OT can shut down critical services, and attacking a utility's IT or OT can shut down the electric grid or a pipeline. Protecting this integrated world is a powerful use case for ZTS.



TIP

The idea is to protect assets, not just the network. This is done by collecting connectivity data, mapping interdependencies, applying risk and function labels, and applying least-privilege, asset-based segmentation.

Securing the Supply Chain

A malware attack can bring vast supply-chain operations to a halt, causing everything from extra hassles to dangerous shortages of vital products.

One way this can happen is if malware gets deployed early in software development or is already deployed in open-source code. The solution is shifting left, to detect the problem early in the continuous integration/continuous delivery (CI/CD) cycle.



REMEMBER

ZTS can block embedded malware that gets activated during the build phase. And by shifting left, you'll also be able to microsegment every workload as it moves all the way into deployment and operation, which means you're preventing any possible lateral propagation right out of the gate.

Protecting High-Value Assets

The security details required by high-value assets will vary from one situation to another. ZTS should allow you to ring-fence these assets at any scale, enforcing least-privilege access between them.

You must be able to define macrosegments or microsegments, around groups of workloads or any kind of subset, right down to specific processes on individual workloads. And you shouldn't have to depend on the underlying network or cloud fabric to do so. After all, if you're segmenting down to the workload, you could

pretty quickly arrive at a larger scale than a lot of network or cloud segment tools can handle.

Your solution should create enforcement boundaries that, by default, will deny all connections. From there, you create exceptions to allow the connections your use requires.

Helping Out with Compliance

Compliance can be a major headache, and the penalties for being out of step with regulations and standards are no joke. Staying in compliance is one part of the headache; another part is ensuring that you can demonstrate your compliance.

A helpful solution is a ZTS tool that uses labels to enable visibility along business functions. For example, do you need to ensure that there's no traffic between production and development? This solution will let you know if there are violations and generate reports to provide evidence of compliance.



TIP

Those tasked with compliance also are typically concerned with risk management and vulnerability mitigation. A ZTS tool can analyze vulnerabilities and assign risk scores, helping you develop policies that will reduce those scores.

Responding Effectively

The faster you spot a breach, the quicker your response can be. Another key ZTS use case is the ability to automate response by sharing workload information with external security information and event management tools, which can then trigger automated alerts to close risky ports.

The response system is informed by cloud object metadata and flow telemetry, along with context and telemetry from managed workloads. Logs are analyzed and policy applied as appropriate. With this setup, response can be fully automated and nearly instantaneous, without requiring any manual steps.




Assume breach. Minimize impact. Increase resilience.

Stop breaches and ransomware from spreading across the hybrid attack surface with the industry's first platform for breach containment.



illumio
The Zero Trust
Segmentation Company



Learn more at illumio.com/zts

Thrive despite the inevitable security breach

Even with the best cyber defenses, you will experience a breach. An “assume breach” mindset isn’t surrender but rather the first step toward building a more resilient organization through Zero Trust Segmentation (ZTS). ZTS contains the spread of breaches and ransomware across the hybrid attack surface, keeping your business moving forward when an inevitable breach happens. This book details this powerful concept and offers a road map to a more secure future.

Inside...

- Making the shift from prevention to containment
- Understanding why it makes sense to “assume breach”
- Seeing that “least privilege” means trusting nothing
- Stopping ransomware with segmentation
- Safely achieving cyber transformation
- Seeing why resilience is the ultimate aim

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!



Steve Kaelble is the author of many books in the *For Dummies* series, and his writing has also been published in magazines, newspapers, and corporate annual reports. When not immersed in the *For Dummies* world or writing articles, he engages in health-care communications.

ISBN: 978-1-394-18168-1

Not For Resale



for
dummies[®]
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.