Zero Trust Data Security

For dummies®
A Wiley Brand

- How to get started with Zero Trust Data Security
- Why your organization needs Zero Trust Data Security
- What comprises a Zero Trust architecture

Rubrik Special Edition

Lawrence Miller

About Rubrik

Rubrik delivers relentless data security, giving you the power to recover from any threat, ransomware attack, or business interruption. No matter where your data lives, Rubrik's Zero Trust Data Security Platform ensures your data is never lost and always ready. Rubrik makes your business unstoppable.

Zero Trust Data Security

Rubrik Special Edition

by Lawrence Miller

for
dummies
A Wiley Brand

Zero Trust Data Security For Dummies®, Rubrik Special Edition

Publisher's Acknowledgments

Table of Contents

Introduction

Cyber threats are increasing rapidly and targeting organizations in every industry across the entire supply chain. Ransomware attacks alone surged by 105 percent globally in 2021 according to SonicWall. While recent high profile ransomware attacks have targeted critical infrastructure and the supply chain, there have been many others — jeopardizing critical business operations for all types of organizations.

Despite massive investments in perimeter, endpoint, and application-layer security defenses, attackers are still gaining access to data.

Zero Trust is commonly implemented as a network security model, but Zero Trust principles also apply to data security and security architecture in general. In this book, you'll discover how Zero Trust Data Security can bolster the effectiveness of your organization's defenses against modern cyber threats.

About This Book

Zero Trust Data Security For Dummies, Rubrik Special Edition, consists of four chapters that explore the following:

>> Why you need to implement Zero Trust Data Security for your organization (Chapter 1)

>> What comprises a Zero Trust architecture (Chapter 2)

>> How to get started with Zero Trust Data Security (Chapter 3)

>> Ten keys to effective incident response (Chapter 4)

Each chapter is written to stand on its own, so if you see a topic that piques your interest feel free to jump ahead to that chapter. You can read this book in any order that suits you (though we don't recommend upside down or backwards).

Foolish Assumptions

It's been said that most assumptions have outlived their uselessness, but we assume a few things nonetheless!

Mainly, we assume that you are a chief information officer (CIO), chief information security officer (CISO), vice president, architect, engineer, administrator, or analyst responsible for backing up and recovering your organization's critical data. As such, this book is written primarily for technical readers with at least a basic understanding of security and backup technologies and challenges.

If any of these assumptions describe you, then this is the book for you! If none of these assumptions describe you, keep reading anyway! It's a great book and you'll learn quite a lot about Zero Trust Data Security.

Icons Used in This Book

Throughout this book, we occasionally use special icons to call attention to important information. Here's what to expect:

REMEMBER This icon points out important information you should commit to your nonvolatile memory, your gray matter, or your noggin.

TECHNICAL STUFF This icon explains the jargon beneath the jargon and is the stuff legends — well, legendary nerds — are made of.

TIP Tips are appreciated, but never expected — and we sure hope you'll appreciate these useful nuggets of information.

WARNING These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice.

Beyond the Book

There's only so much we can cover in this short book, so if you find yourself at the end of this book wondering, "Where can I learn more?" go to `rubrik.com`.

Chapter **1**

Recognizing the Need for Zero Trust Data Security

I n this chapter, you learn about Zero Trust Data Security, why legacy backup architectures are vulnerable to modern threats including ransomware, data security mandates, and the bene-fits of a Zero Trust architecture for data protection.

Understanding Zero Trust Data Security

Traditional perimeter-based IT security defenses that protect an organization's "trusted" internal network from "untrusted" external networks (such as the Internet) are failing. The net-work perimeter has all but disappeared with the proliferation of endpoints — including desktop and laptop PCs, smartphones and tablets, and Internet of Things (IoT) devices — and the rapid adoption of remote work from home (WFH) and work from any-where (WFA) models in the wake of the global pandemic. As a result, threat actors are breaching network security controls and bypassing endpoint protections with ease.

This modern threat landscape is driving many organizations to adopt a Zero Trust approach to cybersecurity. The Zero Trust

Security model is based on the concept of "never trust, always verify." That is, no user, device, or resource (including users, applications, services, databases, and so on) is inherently "trusted" simply because it is "on the network." Instead, the identity of every user, device, and resource must be positively verified every time it connects to the network and granted only the minimum level of permissions necessary to perform an authorized function for a limited period of time.

Although Zero Trust is not a new concept, it has garnered a lot of attention in recent years as more organizations recognize its effectiveness against modern threats and technology vendors seek to capitalize on this trend. Unfortunately, this situation can often create confusion as vendors sometimes attempt to redefine a trend to fit their product offerings. To avoid this bias, many vendors — including Rubrik — follow the Zero Trust model as defined by the U.S. National Institute of Standards and Technologies (NIST) in Special Publication (SP) 800-207, *Zero Trust Architecture.*

REMEMBER

As defined by NIST, Zero Trust comprises "an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources." In the NIST view, a Zero Trust architecture uses Zero Trust principles in the creation of enterprise infrastructure and workflows and the focus is on protecting users, devices, and resources rather than an arbitrary network perimeter. According to NIST, a Zero Trust architecture adheres to the following seven basic tenets of Zero Trust:

» All devices and services that connect to the network and send, receive, or process data should be treated as resources to be verified and protected.

» Regardless of the location and ownership (that is, "on" or "off" the network and enterprise, personally, or third-party owned) of a resource, all communication is protected using the most secure manner available.

» Least-privilege access to individual enterprise resources is granted on a per-session basis after trust is verified and is not transferable to other enterprise resources.

» Dynamic policies are used to determine whether access to a resource is granted, based on behavioral and environment attributes such as software version, network location, date and time of request, and others.

>> All resources that connect to the enterprise network are continuously monitored and evaluated to ensure the enterprise network security posture is not compromised.

>> Resource authentication and authorization (including re-authentication and re-authorization) is dynamic and strictly enforced — using technologies such as multi-factor authentication (MFA) and continuous monitoring — before access is granted.

>> As much data as possible is collected about the current state of resources, network infrastructure, and communications to improve the enterprise network security posture.

The logical components that make up an enterprise Zero Trust architecture include the following (see Figure 1-1):
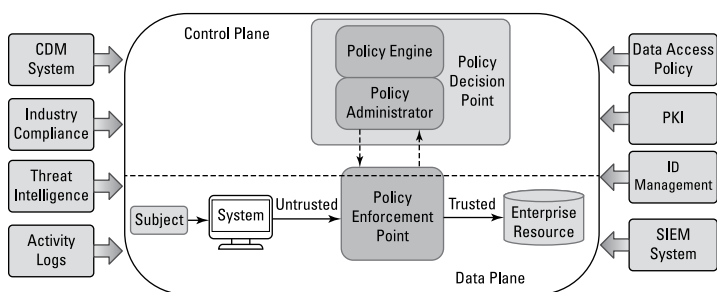
FIGURE 1-1: Logical components of a zero trust architecture. (Source: NIST SP 800-207, *Zero Trust Architecture*)

>> **Control plane**

Policy decision point consisting of a *policy engine* responsible for granting access to a resource and a *policy administrator* responsible for generating any session-specific authentication and authorization tokens (or credentials) used by the policy enforcement point to allow communication between users/devices and enterprise resources.

>> **Data plane**

Policy enforcement point responsible for enabling, monitoring, and terminating connections between users/devices and enterprise resources.

» Supporting components

Continuous diagnostics and mitigation (CDM) system responsible for collecting information about the current state of an enterprise resource and updating configuration and software components.

Industry compliance engine responsible for custom enterprise policy rules that help ensure any applicable regulatory compliance.

Threat intelligence consisting of live feeds from internal and/or external sources.

Activity logs including network and system logs to provide real-time or near-real-time feedback on security posture.

Data access policy including attributes, rules, and policies used by the policy engine to manage access.

Public key infrastructure (PKI) responsible for managing resource, subject, service, and application certificates issued by the enterprise.

Identity management system that provides identity and access management services for the enterprise.

Security information and event management (SIEM) system that collects and aggregates security events from numerous enterprise data sources and generates alerts.

Identifying Challenges with Legacy Backup

Traditional backup architectures are vulnerable to ransomware and other cyberattacks. Typical components of a legacy backup architecture include a Windows or Linux server running the backup application. The server is deployed on the same network as the backup clients (that is, other servers in the data center) and has access to the Internet to enable software updates and remote management. The credentials for the backup server are stored in Active Directory, along with other privileged accounts, and assigned a role such as backup operator. Backups are often stored on a network file system (NFS) or server message block (SMB) volume used as the backup repository.

DO REGULAR BACKUPS PROTECT AGAINST RANSOMWARE?

A common misconception about ransomware is that "we are protected because we back up all of our data." Regular backups are essential, but they can lull organizations into a false sense of security. Traditional backup strategies may fail to protect against ransomware attacks for two important reasons:

- **Cybercriminals have developed ways of encrypting or corrupting online backups.** More and more ransomware attacks are targeting backups. Often, when attackers gain a foothold on a network, one of their first actions is to start encrypting or corrupting backup data. Typically, this occurs two or more weeks before they begin encrypting production data. Thus, when organizations realize they are under attack they often find their backups are unusable.

- **Recovering backup data can take so long that organizations are forced to pay ransoms anyway.** Backing up data to tapes and storing them off site used to provide a gold standard of data protection, but this is no longer the case. It can take several days to find the latest tapes, bring them back onsite, mount, and run them. Also, restoring selectively from tape is difficult; you have to restore all files, not just the ones that were encrypted. Many organizations cannot afford to wait for these operations to complete before restarting critical business systems and therefore end up paying the ransom demand.

Because the backup architecture is not isolated from the production network it is similarly vulnerable to ransomware. Therefore, ransomware has multiple opportunities to infect or corrupt backup jobs either on the backup server or on the backup repository, or both.

Looking at Data Security Issues

Data is an organization's most valuable asset and is therefore targeted by cybercriminals. In addition to protecting the confidentiality, integrity, and availability of sensitive data, organizations must also ensure the privacy of their data.

Data privacy regulations such as the General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), and California Privacy Rights Act impose strict requirements on organizations to ensure "privacy by design" and "privacy by default" in their data architectures. This requirement extends not only to data in production environments, but anywhere data is stored and used in an environment — including backups.

Realizing the Benefits of Zero Trust Data Security

REMEMBER

The benefits of a Zero Trust Data Security architecture include the following:

» **IT teams** can protect critical data from ransomware attacks, giving their organization the ability to recover data and applications quickly — without paying a ransom.

» **Security teams** can confidently leverage secured backup data to perform attack forensics and initiate recovery operations directly from their security operations center (SOC).

» **Application owners** can rest easy knowing that business data is protected, and that if a ransomware attack were to occur, applications can be restored quickly to maintain business continuity.

» **Chief information officers (CIOs) and chief financial officers (CFOs)** can be assured that ransomware recovery plans are supported by a Zero Trust architecture that enables the organization to minimize cyber insurance costs and avoid reputation damage resulting from ransomware attacks.

IN THIS CHAPTER

» **Exploring Zero Trust Data Protection**

» **Investigating ransomware**

» **Discovering sensitive data**

» **Containing ransomware to accelerate incident response**

» **Orchestrating application recovery**

Chapter **2**

Understanding Zero Trust Data Security Architecture

I n this chapter, you learn how Zero Trust Data Protection solutions support a Zero Trust architecture to help customers safeguard their backups from ransomware attacks and other threats.

Zero Trust Data Protection

Zero Trust Data Protection delivers intelligent data management for hybrid and multi-cloud IT environments in a single software platform that securely delivers backup, recovery, analytics, and compliance across data centers and clouds.

Key Zero Trust Data Protection capabilities include:

» **Service-level agreement (SLA) automation:** Replace hundreds or thousands of backup jobs with just a few policies that can be applied across all your workloads.

- **»** **Rapid recovery:** Search across your global environment and find exactly what you're looking for. Click recover and mount directly for near zero recovery time objectives (RTOs).

- **»** **Application programming interface (API)-first:** Any behavior that can be triggered from the user interface (UI) can be scripted, automated, and integrated into tools you may already own.

- **»** **Secure by design:** Once backup data is written, it can never be changed meaning your backups are safeguarded against ransomware and other threats.

Ransomware Investigation

Speed is everything when ransomware attacks happen, but identifying the scope of impact from an attack can be challenging. Quick recoveries are often a pipe dream and it takes, on average, more than seven days to recover from a ransomware attack.

Ransomware investigation tools analyze changes across backup data to detect malicious changes and assess the blast radius. These insights enable incident responders to accelerate recovery. Key ransomware investigation capabilities include:

- **»** **Identify anomalies:** Analyze backup data for unusual behavior and changes and get alerts on suspicious activity.

- **»** **Assess impact:** Understand the scope of files and applications that were affected by ransomware and where they are located.

- **»** **Inform rapid response:** Leverage impact analysis to granularly restore only the files and applications most likely to have been affected for faster incident response.

ASL AIRLINES FRANCE BUILDS RANSOMWARE DEFENSE STRATEGY

ASL Airlines France (ASL) is a cargo and passenger airline based in Tremblay-en-France at Bâtiment Le Séquoia. Their main base is Charles de Gaulle airport, Europe's second busiest air traffic hub. A majority of ASL's fleet operates on behalf of delivery services throughout the night, including Amazon, FedEx, DHL, UPS, and La Poste. In 2017 alone, ASL carried 712,000 passengers and 38,600 tons of cargo.

Fabrice De Biasio, chief information officer at ASL Airlines, oversees the operational infrastructure of 3,000 employees and is responsible for ensuring always-on data availability and meeting strict security standards. In 2018, with the threat of cyberattacks on the rise, ASL partnered with Rubrik to proactively address the threat of ransomware with Ransomware Investigation.

Challenge

Ransomware attacks are intensifying in scale and sophistication. A recent NTT Security survey revealed that ransomware attacks rose 350 percent in 2017 over the previous year. Nearly 75 percent of companies infected with ransomware suffer two days or more without access to their files while 33 percent go five days or longer. According to Cyber Security Ventures, ransomware is likely to cost victims more than $250 billion annually by 2031.

ASL is required to maintain 99.9 percent availability — a maximum of 60 minutes of allowed downtime per year. If ASL's IT system is down for more than 15 minutes, airplanes cannot take off, customers cannot receive their cargo, and the airline is at risk of being hit with massive fines. "In our business, you cannot have downtime," said De Biasio. "Ransomware can quickly cripple an airline and prevent its ability to fly, period."

Solution

ASL's previous solution was not built for a strong defense against the rapidly growing threat of ransomware. "The cargo airline industry is a common target for ransomware, and we experience a minimum of one attack per month," said De Biasio. "In the past, we managed to recover by using a multitude of scripts to identify and erase infected

(continued)

files manually. This was an incredibly painful, time-consuming experience that killed our team's productivity for days."

By enabling fast recoveries and providing detailed impact assessments, Ransomware Investigation enables enterprises to significantly minimize downtime, cost of recovery, and reputational damage following an attack.

Results

Prior to Rubrik, the threat of ransomware was keeping De Biasio up at night. Now, with Ransomware Investigation's multi-level defense, De Biasio has peace-of-mind and realized the following benefits:

Operational Savings

- **15 to 100+ hours of IT admin time saved in case of an attack:** "We experience a minimum of one ransomware attack per month. Before Ransomware Investigation, the team spent 15 hours to recover from a minor ransomware attack. If we had been hit with a major attack, I fear recovery could've taken weeks."

- **25 percent IT admin time savings (40+ hours saved per month):** "Our team used to spend up to two hours per day monitoring our applications for ransomware. Now, we only need to spend a few minutes per day checking Ransomware Investigation, so our team can spend more time on initiatives that deliver value back to the business."

- **Automatic recovery and no downtime:** "Before Ransomware Investigation, we managed to recover from attacks with several scripts and by identifying and erasing bad files manually. That was an incredibly painful experience. Our IT Admin loves Ransomware Investigation because it does all that work automatically. Ransomware Investigation discovered a bad file, alerted him, and he just ticked a few boxes to restore to a clean state."

Business Impact

- **Global visibility and instant threat response:** "With Ransomware Investigation we can follow server activity in real time and react fast. If something is not normal, we know about it."

- **Ability to protect our business against catastrophic risk with cyber insurance:** "Because the cargo airline industry is a common target for ransomware attacks, it's incredibly difficult for airlines to

get cyber insurance. If we did not have Ransomware Investigation, we would not have been approved for a cyber insurance contract."

- **Millions of euros in potential savings in case of an attack:** "Ransomware Investigation will help us protect our bottom line and potentially save us millions of euros in case of an attack."

With Ransomware Investigation's machine learning-powered anomaly detection and accelerated recovery, ASL's team is now confident in their ability to quickly restore to the pre-infected state in the event of a threat. "Rubrik's native immutability coupled with the AI-driven alerting and detection of Ransomware Investigation are the most critical data protection and business continuity tools in my arsenal against today's intensifying cyber threats," said De Biasio.

Sensitive Data Discovery

Strong preparedness and response strategies are critical to mitigating ransomware and other disasters. But a lack of visibility into sensitive data can lead to vulnerabilities and unnecessary incident response costs.

Locating sensitive data in files and applications is vital to helping you stay compliant and in control. Here are some key sensitive data discovery capabilities your platform should have:

- ›› **Automate policy enforcement:** Your data security platform should allow you to select the types of personally identifiable information (PII) and other sensitive data you want to monitor for automated policy enforcement. It should also have pre-defined templates or create custom policies to quickly identify and classify sensitive data without impacting production performance.

- ›› **Assess data exposure:** Identify sensitive data that may be exposed during data exfiltration or other unauthorized access. Search the index using keywords or custom fields such as name, social security number, or credit card number to identify potential exposure and risk.

- ›› **Simplify compliance:** Document where sensitive data is and who has access to it, to maintain regulatory requirements and get alerts when data might violate policies. You can also

schedule periodic reports to ensure ongoing compliance with the General Data Protection Regulation (GDPR), Payment Card Industry Data Security Standards (PCI-DSS), Health Insurance Portability and Accountability Act (HIPAA), and Gramm-Leach-Bliley Act (GLBA).

CITY OF SIOUX FALLS MINIMIZES DATA BREACH RISK WITH SENSITIVE DATA DISCOVERY

Sioux Falls is the largest city in South Dakota with over 182,000 citizens. With a 4 percent annual growth rate, the city is quickly evolving into a major metropolitan area of the Midwest. With a growing population of citizens, the city is faced with an exponentially increasing data footprint and massive data sprawl, making it harder to discover and protect sensitive employee and citizen data.

Brandon Morris, Systems Administrator, is part of a 28-person IT team supporting the city's central services for its 1,200 employees and citizens, including managing the control center for fire, police, and rescue. "We have access to all kinds of citizen data, such as health clinic records, permits, and PII. We must ensure its privacy and confidentiality. As a result, it is critical that we implement strong and continuous data governance processes," said Morris. "Rubrik's [Sensitive Data Discovery] furthers our data governance strategy by providing visibility into where sensitive content resides and within previously dark repositories, allowing us to minimize risk of data exposure and breaches and drive continuous compliance."

Challenge

In the past, sensitive data discovery and classification tasks were extremely manual and labor-intensive, requiring dedicated teams of several full-time engineers. "Prior to Rubrik, we didn't have any tools to search for and categorize sensitive data. When we had to find documents that contained certain types of content, several employees would have to manually script queries for each keyword one at a time for each possible location. I could never run it across our entire environment as this only worked on a small data set and required us to know where we suspected sensitive content to reside. On top of that, the interface was very difficult to use," said Morris.

"The last time we had to run through this process, it was very painful and time-consuming, requiring 40 to 80 man-hours. Each search query would take about one to two weeks with several days to perform the scan of our applications for each keyword. We then had to manually categorize all that data into spreadsheets. The amount of manhours is too expensive," said Morris. "With [Sensitive Data Discovery], we can completely automate the process and search on-demand for hundreds of queries across all our file servers. Now, it takes just 1 hour to perform a search query. That means we can run data classification seamlessly in the background without dedicated teams and multiple full-time engineers, giving us massive productivity gains and freeing up employees for higher-value work."

Solution

One of the key objectives City of Sioux Falls looked for in Rubrik's Sensitive Data Discovery was to gain more visibility and control into where sensitive content resides. With Sensitive Data Discovery, they were able to identify where data was overexposed, such as employees and citizens' credit card numbers and social security numbers, and access levels were higher than expected. "[Sensitive Data Discovery] gave us visibility into data repositories we could never look at before. As a result, we were able to avoid data exposure and proactively create remediation plans to address high-risk incidents. The benefit to the business is huge. Data exposure can mean negative PR and impact to our citizens' well-being if there is a data breach. [Sensitive Data Discovery] helps us minimize the risk of data exposure while providing substantial management time savings."

Results

City of Sioux Falls has realized numerous benefits from the Sensitive Data Discovery solution including:

- **Up and running in less than 30 minutes without additional infrastructure:** "We saw immediate value with [Sensitive Data Discovery]. Since it leverages our existing Rubrik Cloud Data Management. . . deployments, we didn't need to purchase additional servers or storage and simply toggled on the app from the. . . UI to get started."

- **Continuous data governance without impact to production:** "[Sensitive Data Discovery] is continuously monitoring all our

> existing backup data in the background to alert us to any viola-
> tions or sensitive data stored in wrong locations without using
> agents or touching our production data."
>
> - **On-demand access to data for audits:** "We have to comply with
>   audits and regulations. With [Sensitive Data Discovery], we can
>   provide an auditor or our internal legal team access to the data
>   they need."

Incident Containment

Before an organization can recover from a ransomware attack,
they first have to contain the threat. Otherwise, they risk reinfec-
tion during recovery. But identifying which systems have been
affected by a ransomware attack can be challenging.

Your data security platform should analyze backup snapshots
and provide insights that help avoid malware reinfection during
recovery. Key incident containment capabilities include:

» **Scan for threats:** Scan backups using file patterns, file
hashes, and YARA (Yet Another Recursive Acronym) rules to
look for any indicators of compromise (IoCs) across all
objects in the backup.

» **Identify recovery points:** Analyze a time-series history of
backup snapshots to pinpoint a clean uninfected snapshot
for recovery.

» **Avoid malware reinfection:** Leverage insights to quickly
recover with less risk of reintroducing malware and provide
forensic evidence during internal and external investigations.

Orchestrated Application Recovery

Ensuring security and resiliency for data and business services
in the face of cyberattacks and other disaster events is a criti-
cal responsibility of modern digital enterprises. However, install-
ing and maintaining new infrastructure and software for disaster

recovery (DR) can be expensive and time consuming. Executing manual plans for applications with multiple tiers and interdependencies slows down the recovery process and introduces opportunities for error.

Enterprises can avoid these burdens by using orchestrated application recovery for a tightly integrated and automated DR service.

Orchestrated recovery is delivered as a software-as-a-service (SaaS)-based application. It provides orchestration of DR failover/failback, testing, and together with application-focused Ransomware investigation will radically simplify recovery for business services running in VMware vSphere environments. As a result, IT organizations can eliminate multiple point solutions, management complexity, and avoid unnecessary costs.

Key orchestrated application recovery capabilities include:

» **Simplify orchestration:** Recover to your on-premises DR site, VMware Cloud on AWS, or Azure VMware Solution.

» **Prove DR readiness:** Confirm application availability, demonstrate compliance, and prove DR preparedness.

» **Recover from attacks:** Identify encrypted data and recover the most recent clean state from your uncompromised backups.

IN THIS CHAPTER

» **Understanding the core components of a Zero Trust architecture for data protection**

» **Leveraging machine learning**

» **Classifying sensitive data**

» **Proactively hunting for threats**

» **Automating and orchestrating recovery workflows**

Chapter **3**

Getting Started with Zero Trust Data Security

In this chapter you learn about the foundational components of a Zero Trust architecture for Zero Trust Data Security. You also discover how machine learning, data classification, threat hunting, and automation and orchestration of recovery workflows all contribute to a comprehensive modern data protection strategy.

Safeguarding Data with Immutability and Data Availability

One of the reasons enterprises are unable to recover from a ransomware attack is that backups become compromised, forcing them to either pay the ransom or restore from offsite backups. Be cautious of data protection vendors that advise offsite backups as the primary recovery option as this can take weeks to months to restore and is often subject to data integrity challenges, leading to longer recovery time objectives (RTOs). Additionally, some backup

vendors advise implementing an isolated recovery to address ransomware. While this is a viable option, it comes with a large cost burden and management complexity to implement. Think of it as equivalent to the operational and financial overhead of a disaster recovery infrastructure.

To effectively protect your data from ransomware attacks, your backup and recovery system must be capable of creating immutable copies of your data — that is, backups that cannot be encrypted by ransomware. Zero Trust Data Security provides the foundation for a modern backup and recovery system with the following core capabilities and elements:

» **Reduce the risk of intrusion.** All system interfaces are secure, role-based, least privileged, and protected by multifactor authentication (MFA).

» **Secure the data.** Data is always encrypted in-flight and at rest, and backup data is stored in a purpose-built append-only file system. Backed up data is always logically air-gapped so it's offline and not accessible through standard network protocols.

» **Detect and alert anomalous behavior.** Attacks are detected and the SecOps team is alerted so that a clean recovery point can be quickly and confidently identified.

» **Enforce compliance.** New workloads are automatically protected with lock retention and the ability to find certain exposed sensitive data that may have been exfiltrated.

TECHNICAL STUFF

Data in an immutable format cannot be read, modified, or deleted by an external client once it has been ingested.

TIP

Underlying the Zero Trust architecture is a core set of technologies in backup and recovery solutions, which supports a purpose-built file system that never exposes backup data via open network protocols. The Zero Trust architecture creates a logical airgap that blocks data from being discoverable or accessible over the network and is comprised of the following components (see Figure 3-1):

» **Immutable data platform:** Once ingested, no external or internal operation should be able to modify the data. Data managed by the platform should never be available in a

read/write state to the client. Because data cannot be overwritten, even infected data later ingested by the platform cannot infect other existing files or folders.

» **Declarative policy engine:** This tool allows administrators to abstract away much of the low-end fuss required to build and maintain data protection, so they can focus on adding value at a more strategic level across the organization.

» **Threat engine:** As each backup snapshot's metadata is collected, machine learning builds out a full perspective of what is going on with the workload. The network is trained to identify trends that exist across all samples and classify new data by their similarities without requiring human input. You will be able to detect anomalies, analyze the threat, and help accelerate recovery with a few clicks.

» **Secure application programming interface (API) first architecture:** Having an API-driven architecture means that every action in your platform user interface (UI) has a corresponding API that is documented and available for use. Or in other words, if you can do it through the UI, you can programmatically do the same through the API that's secured by role-based access and API tokens.
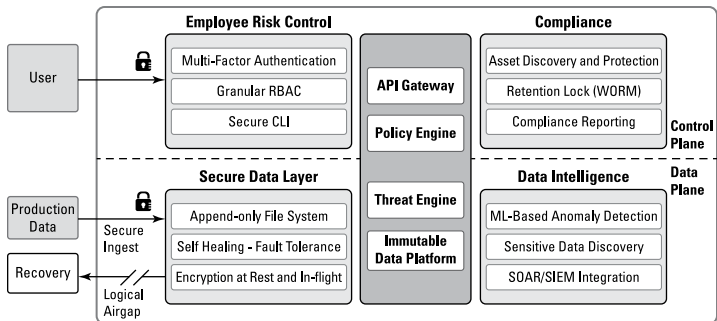

FIGURE 3-1: Core Zero Trust architecture components.

With an immutable platform, once data is written to the system, it cannot be modified, deleted, or encrypted by an attack, ensuring that a clean copy of data is always available for recovery.

CHAPTER 3 **Getting Started with Zero Trust Data Security** **21**

Discovering Data Anomalies with Machine Learning

Another important element of a Zero Trust Data Security protection strategy is machine learning tools built into the backup and recovery system. These tools monitor application metadata to detect and alert you to signs of anomalous activity that might be indicative of a ransomware attack.

Ideally, these tools would provide insight at a very granular level so you could quickly identify specific files that had been compromised — and then quickly restore just those files rather than entire virtual machines (VMs), with a single click.

REMEMBER

This combination of machine learning tools operating in real time, coupled with the ability to recover infected data quickly from immutable backups, should be part of your Zero Trust Data Security protection strategy.

Classifying Data and Assessing Exfiltration Risk

Ransomware attacks have evolved, and many cybercriminals are now targeting victims with "double extortion ransomware." In a double extortion ransomware attack, a copy of the victim's data is exfiltrated before the original files are encrypted. The victims are then threatened with two negative outcomes: losing their production data, and having sensitive information leaked on the web such as personally identifiable information (PII) of customers and employees, financial accounts and social security numbers, product designs and other intellectual property, proprietary software, and potentially embarrassing internal emails and documents.

WARNING

Paying a ransom does not guarantee that a cybercriminal will provide you the keys necessary to decrypt your data, although they typically do. (The ransomware business model would fail if there were not a reasonable expectation that paying the ransom

would enable you to unlock your data.) However, paying the ransom does not necessarily mean that a cybercriminal will not keep a copy of your data and use it to extort more ransom from you at a future time of their choosing. Alternatively, a cybercriminal might simply choose to sell their stolen copy of your data (such as customer social security numbers and credit card numbers) on the dark web.

To limit the potential impact of double extortion ransomware, organizations should identify their data to enable the following:

» **Rapid determination of whether sensitive data has been compromised** in the event of a ransomware attack (or any cyberattack, for that matter). This determination will also help the organization address any disclosure or notification requirements. For example, the General Data Protection Regulation (GDPR) requires prompt notification of the Data Protection Authority in the country or region that is impacted by a data breach, as well as notification of affected data subjects.

» **Appropriate implementation of security controls** for different classification levels of data commensurate with its sensitivity level and value. These protections may include measures such as limiting access, encrypting files, increasing logging/auditing, and backing up data sets more frequently.

If an organization can quickly determine that sensitive data was not, in fact, compromised, then it may be able to limit regulatory penalties and costs related to breach notification. And of course, knowing what data has been encrypted by ransomware helps recovery teams appropriately prioritize their efforts based on the needs of the business and reduces the time, effort, and expense associated with recovery operations.

Organizations should also consider measures that protect the entire environment, such as expanding the use of multi-factor authentication (MFA). Many ransomware campaigns take advantage of weak or stolen credentials or use brute force password cracking techniques to access target networks and systems. MFA can prevent or limit the spread of ransomware by mitigating inherent vulnerabilities associated with static passwords.

Hunting for Threats to Prevent Reinfection

To prevent reinfection, organizations should have the ability to search their backups without having to restore them to identify potentially compromised backups. By looking back in time at your VMs and file sets, you can pinpoint when an infection started so that you avoid malware reinfection during recovery.

Recovering Apps and Data with Guided Workflows

As organizations increasingly rely on IT generalists to support their technology needs, traditional backup management strategies that relied on scripting skills to schedule and manage backup jobs have become too complicated to sustain.

Modern backup and recovery systems that support a declarative management model support the need in organizations today for a system that is simple and intuitive enough for practically anyone to manage. In a declarative management model, an administrator enters the desired state for a workload into a policy engine. After a policy is set, the system automatically and intelligently executes the jobs that need to be performed to achieve that state (see Figure 3-2).
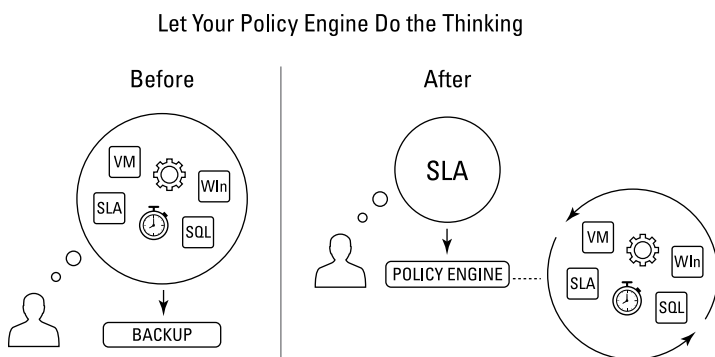

FIGURE 3-2: Collapse multiple manually implemented settings into a single, easy-to-configure and zero-maintenance declarative policy.

A strong policy engine can facilitate other aspects of service auto-mation as well, reducing the number of manual steps that an IT administrator might have to complete to accomplish a task. If the backup and recovery solution has an API-first architecture, the organization gains even greater benefits. An administrator could use these capabilities to:

» Integrate backup and recovery into an IT service catalog (for example, ServiceNow and VMware vRealize Automation or vCloud Director)

» Simplify management of large, distributed environments via configuration management or infrastructure as code (IaC) tools (for example, Puppet, Chef, SaltStack, or Ansible)

» Automate lifecycle data-management workflows, and centralize monitoring and reporting (for example, Splunk or a custom monitoring dashboard)

Additionally, automation enables regular backup validation — a requirement to mitigate the "testing gap" risk, as shown in Figure 3-3. If backups aren't tested regularly, IT cannot guaran-tee their validity to the business.
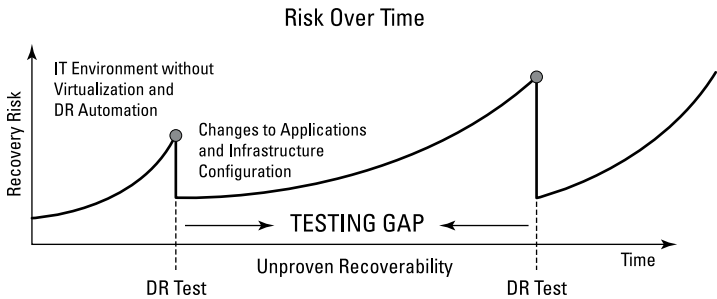

FIGURE 3-3: Without regular testing, guaranteeing reliable restores is impossible.

Chapter **4**

Ten Keys to Response Preparedness

Implementing Zero Trust Data Security helps your organization mitigate threats to your critical data and ensure you can recover from a ransomware attack (as well as other cyberthreats and disasters). Here are ten additional best–practice tips for response preparedness:

» **Keep your incident response, business continuity, and disaster recovery plans current and test often.** Creating these critical plans "on-the-fly" during an actual ransomware attack is a recipe for, well — disaster. Each of these plans needs to be a living document that is reviewed, updated, and tested regularly, as well as anytime a significant event happens such as a major hardware or software upgrade, a merger or acquisition, business growth into new markets, or a drastic workforce change — for example, a shift to a

remote work-from-home (WFH) or work-from-anywhere (WFA) model. These plans should address all areas of business operations, not just IT and testing should include tabletop exercises, structured walkthroughs, and full simulations.

» **Deploy multi-factor authentication (MFA).** Simply put, passwords are not effective in securing user accounts. MFA requires users to sign into their accounts with a combination of username, password/passphrase, and a third "factor" such as a one-time passcode sent to a registered smartphone or a security token.

Although MFA provides far more robust account security than usernames and passwords, it's not impenetrable. Attackers use sophisticated techniques to defeat MFA, particularly one-time passcodes that are sent via text messages to smartphones.

» **Implement least privilege and create separate admin accounts.** Conduct regular access reviews for all your users to ensure they have only the privileges they need to perform their job functions. Remove any unnecessary rights and carefully audit group memberships. Create separate admin accounts (or, better yet, implement privileged access management, or PAM) for users that require privileged access so that they can use a standard user account for day-to-day activities that don't require admin privileges.

» **Segment and lock down your network.** Once a targeted organization has been breached, attackers take advantage of more or less unrestricted lateral movement in the relatively flat networks that are common in traditional perimeter-based security models. Segment your network with next-generation firewalls and other security tools to inspect and control east-west traffic within your network and north-south traffic across different networks. Also remove any unnecessary network services, such as external access to remote desktop protocol (RDP).

» **Isolate your backup and recovery system from the rest of your network.** To help ensure the viability of your data backups in the event of a cyberattack, isolate your backup and recovery system from the rest of your network.

» **Set up a 3-2-1 backup strategy and test your recovery capabilities.** At a minimum, maintain three copies of your

critical data (one production and two backups) on two different storage devices (or media) and keep one copy offsite (for example, in the cloud or a separate data center). Also, remember that all the backups in the world are of no use if you can't recover your data. Be sure to regularly test your recovery capabilities to ensure your team knows what to do and your backups are reliable and sufficient to achieve your organization recovery time objectives (RTOs) and recovery point objectives (RPOs).

» **Audit your logs for unusual activity.** Far too often, organizations only discover that their system and network logs are incomplete, inadequate, or inaccessible during or after an attack. These logs provide critical data that enables security and information event management (SIEM) platforms to generate accurate events and alerts and helps train machine learning models to better detect anomalous behavior.

TIP

Your logs also provide important information to help you answer the "who, what, and when" questions in a forensic investigation, such as whether sensitive (and regulated) data — including protected health information (PHI), personally identifiable information (PII), or cardholder information — has been compromised or exfiltrated.

» **Secure your endpoints.** Endpoints — including desktop and laptop PCs, smartphones and tablets, and Internet of Things (IoT) devices — represent an organization's largest and typically most vulnerable attack surface. Today's increasingly remote workforce must make important endpoint security decisions — such as whether to grant a downloaded app access to data on a mobile device — without fully understanding the risk that these decisions may pose to the organization as a whole. Deploy (and regularly update) endpoint protection solutions — including anti-malware, email protection, data loss prevention (DLP), and endpoint detection and response (EDR) — to help secure your endpoints.

» **Train your end users.** People are traditionally the weakest link in any security strategy but, when regularly and effectively trained, they can become a force multiplier for your security team. Phishing simulations are a great example of interactive and engaging end-user security awareness training that has helped organizations reduce risks

associated with email phishing campaigns. Extending this training method to include other modern security threats, including ransomware, helps build a more secure culture in your organization.

» **Ensure you're adequately insured.** Even with an effective incident response plan that includes a robust backup and recovery solution, a cyberattack is costly. These costs often include business interruptions (and lost revenue), rebuilding or replacing systems, third-party forensic services, legal fees, civil and regulatory penalties, customer notifications, brand reputation damage, and more. The cost of downtime alone is estimated to average $300,000 for a single hour of downtime according to the ITIC *2021 Hourly Cost of Downtime Survey* (www.itic-corp.com/tag/hourly-cost-of-downtime/). Cyber insurance (also known as cyber risk insurance) can help organizations recover many of these costs to reduce the financial impact of an attack.

TIP

For many enterprises, their cyber insurance provider plays a major role in setting policies about how to respond to ransomware attacks and under what circumstances to pay ransoms. It's critical to get input from the insurance company either directly or through someone in the legal or governance, risk, and compliance (GRC) group who has a detailed knowledge of its practices and requirements.

Take a zero trust approach to your data security

A thriving ransomware economy has emerged. Hackers are now going straight for your backup data. And while legacy backup solutions are good at recovering from natural disasters and IT failures, ransomware recovery requires you to rethink your security strategy. If data is the target, then security defenses must begin at the point of data. With this book, you gain a clear understanding of how a zero trust architecture keeps data secure so attackers can't hold data hostage. You also learn how to evaluate a vendor for zero trust security capabilities.

Inside…

- Explore data security and privacy issues
- Understand the benefits of zero trust
- Secure your critical apps and data
- Locate, classify, and report on sensitive data
- Assess impact from cyber attacks
- Find malware and avoid reinfection
- Accelerate recovery from ransomware

rubrik

Lawrence Miller served as a Chief Petty Officer in the U.S. Navy and has worked in information technology in various industries for more than 25 years. He is the co-author of *CISSP For Dummies* and has written more than 200 *For Dummies* books on numerous technology and security topics.

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

9 781119 882398

for dummies®
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

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