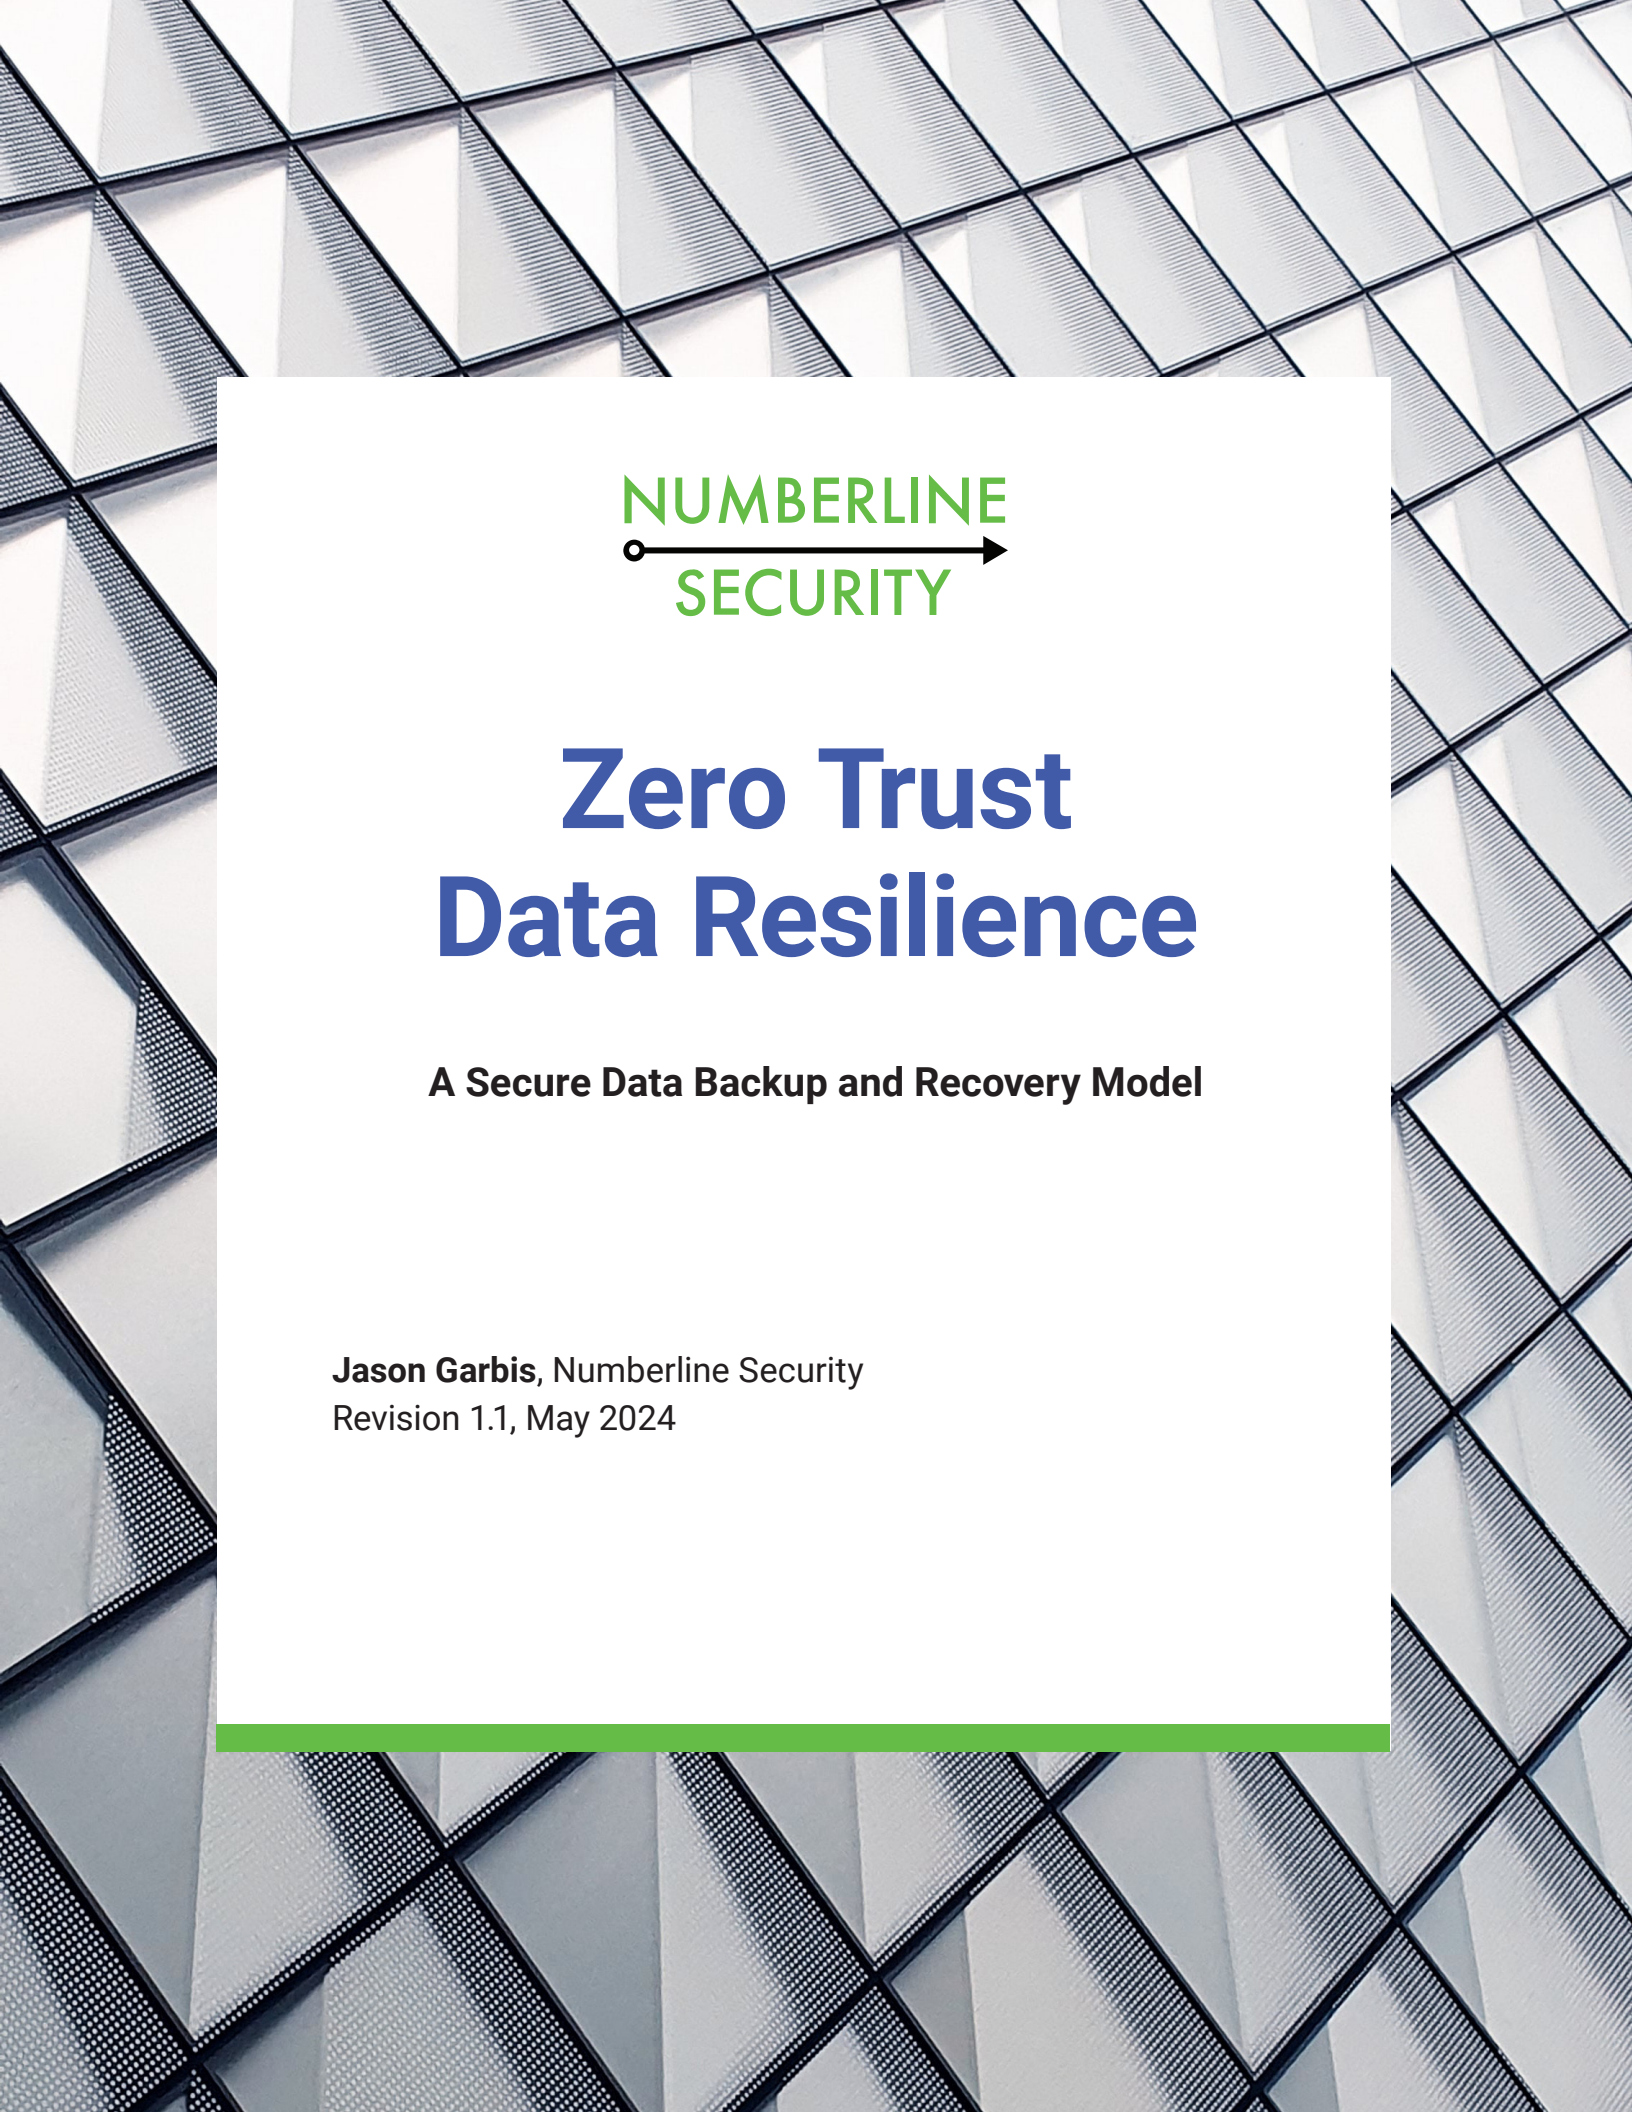

NUMBERLINE SECURITY

Zero Trust Data Resilience

A Secure Data Backup and Recovery Model

Jason Garbis, Numberline Security
Revision 1.1, May 2024

Executive Summary

Enterprises today face ongoing and significant challenges around safeguarding their data and networks from malicious actors, in particular against ransomware and data exfiltration attacks. To address these concerns, a strategy known as Zero Trust has gained significant traction in the information security industry and is being widely adopted by enterprises worldwide.

However, even the most broadly used Zero Trust models lack comprehensive guidelines in certain important areas, especially around data backup and recovery. Recognizing the importance of filling this gap and applying Zero Trust principles to this area, we introduce the concept of Zero Trust Data Resilience. This comprises of a set of requirements, an architecture, and an extension to existing Zero Trust Maturity Models.

Specifically, enterprises must use a data backup and recovery system that provides immutable data storage and configuration while enforcing contextual and strongly authenticated access to source data in production and backed up data. This system must also seamlessly support the hybrid architectures that are common in today's enterprises, and flexibly handle recovery to dissimilar environments.

By implementing a Zero Trust architecture that meets these requirements, enterprises will better protect their data, networks, and applications against malicious actors. Zero Trust provides demonstrably better security compared to traditional approaches, and organizations have an obligation to adopt it. The new data resilience requirements proposed in this whitepaper enhance and extend Zero Trust, and should be considered mandatory as part of any enterprise's security strategy.

Introduction

Zero Trust is a security strategy, and by necessity is broad in scope. However, the Zero Trust models and frameworks that are in widespread use do not include everything.[1] This can lead to corresponding gaps or omissions in enterprise security architectures. Specifically, data backup and recovery systems are not included in commonly used Zero Trust frameworks. This is an unfortunate gap, since enterprise data is very frequently the primary target of malicious actors in both ransomware and data exfiltration attacks.

Data backup and recovery systems are critical elements of enterprise IT, and must be treated as such. They have read access to everything of importance in order to back it up. They also need the ability to write data into production environments in order to perform their data restoration function. They also contain a full copy of the enterprise's most important data (granted, typically in an encrypted format). Taken together, all these attributes underscore the importance of data backup and recovery systems, and highlight their value as a target for malicious actors.

Of course, data backup and recovery systems have been part of IT's responsibility for decades, but have often not been included in the scope or responsibility of security teams.[2] However, given the level and sophistication of security threats that enterprises currently face, taking only a network and IT infrastructure perspective on data backup and recovery is no longer sufficient. In practice, we have encountered enterprises where these systems were poorly configured and unmonitored, and therefore caused significant risk.

Modern, effective security is based on Zero Trust principles, so it's time to take a fresh look at data backup and recovery systems through that lens. This whitepaper accomplishes this by proposing a new concept of Zero Trust Data Resilience. By embracing this approach, enterprises will have a clear and concrete pathway to having stronger defenses, more efficient operations, and faster recovery.

1 - The CISA ZTMM document states "While the ZTMM covers many aspects of cybersecurity critical to federal enterprises, it does not address other aspects of cybersecurity such as…recovery."

2 - Other than involvement in backup data encryption key management

Approach

The classic foundational elements of information security – the CIA triad of Confidentiality, Integrity and Availability — are all applicable to data backup and recovery. Enterprises need to avoid data exfiltration (Confidentiality), block ransomware from encrypting data (Integrity), and ensure that systems are both protected against attack and can be rapidly restored after an attack (Availability).

Core Zero Trust principles are certainly relevant to this domain, and should be applied to user and enterprise IT system access, as well as to data backup and recovery systems. These principles include the elimination of implicit trust and unsegmented neworks, controlling all access by dynamic and contextual policies via Policy Enforcement Points (PEPs), requiring appropriately strong authentication of all subjects, assuming breach, and ensuring and validating system and data integrity. Throughout this whitepaper, we'll see how these principles flow through to the proposed new set of requirements for a Zero Trust Data Resilience architecture.

The de facto standard framework for looking at Zero Trust maturity is the CISA Zero Trust Maturity Model [3] depicted in Figure 1, which defines five core pillars: Identity, Devices, Networks, Applications & Workloads, and Data. It also defines three cross-cutting capabilities: Visibility and Analytics, Automation and Orchestration, and Governance.

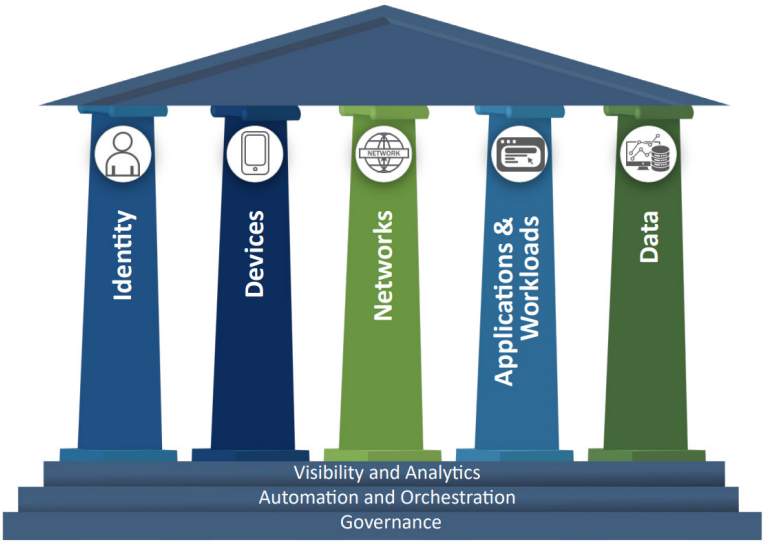

Figure 1: CISA Zero Trust Maturity Model

3 - https://www.cisa.gov/resources-tools/resources/zero-trust-maturity-model

Within the Data pillar, the CISA model identifies five detailed functions, with expected capabilities and attributes for each maturity level.

The functions are:

- Data Inventory Management
- Data Categorization
- Data Availability
- Data Access
- Data Encryption

However, within these functions, the topic of data backup integrity and recovery is minimal, and CISA points readers to a 2020 NIST document that's not connected to Zero Trust. In summary, the CISA Zero Trust model is silent on requirements and maturity levels for data backup and recovery systems. Because this area is so important for enterprise confidentiality, integrity, and availability, we believe that this gap needs to be addressed.

In order to do so, we're introducing the concept of Zero Trust Data Resilience, which includes principles, a reference architecture, and a new set of capabilities for the Zero Trust Maturity Model. Taken together, these represent an extension and enhancement of Zero Trust, and will result in a stronger enterprise security stance.

Zero Trust Data Resilience: Principles

The core principles of Zero Trust Data Resilience (ZTDR) are:

- Segmentation and Least Privilege Access
- Immutability
- System Resilience

Let's discuss each of these in turn.

Segmentation and Least Privilege Access

This principle is central to Zero Trust, and is a required part of any Zero Trust architecture. However, it's worth examining its applicability to the specifics of ZTDR, since it applies at multiple levels. From a networking perspective, the backup management system itself must be isolated on the network so that no unauthenticated or unauthorized users or devices can access it. Likewise, the backup storage system must be isolated. This prevents malicious actors from discovering either system through network reconnaissance or exploiting a vulnerability.

Legitimate and authorized access to the backup system must only occur through a Zero Trust Policy Enforcement Point (PEP) with appropriately strong authentication and device posture checks. The Zero Trust PEP must also control access to the source data (i.e., the data being backed up), with appropriate authentication and some level of device or system validation to ensure that the backup management system is what's reading production data, rather than a malicious system or process.

Access from the backup management system to the backup storage must also be controlled by a PEP and segmented from the rest of the network with appropriately strong authentication. Note that we'll be revisiting this requirement in the architecture diagram below, since it's an important one — the backup software must be segmented (separated) from the backup storage.

Immutability

The concept and requirement for immutable backup data has become widely adopted in recent years, in conjunction with the growth in the prevalence and sophistication of ransomware. An immutable backup is defined as backed-up data utilizing a storage mechanism that, once written, cannot be altered. The premise is that, even if a malicious actor were present on the network and able to take control of the backup system and have access to backup storage, they would be unable to delete or modify (encrypt) the backed-up data. Some immutability comes from physical properties of storage media, such as Write-Once-Read-Many optical disks, while newer technologies use media with immutability enforced at hardware, firmware, or software layers. Most recently, major cloud service providers have added immutable storage capabilities to meet enterprise compliance and archiving requirements.

> Note that requirements for immutability extend beyond the stored data, and must include data retention periods too. Some immutable data may be configured for indefinite storage, while others may have a defined retention period, such as one or five years. Data that ages beyond its retention period can be deleted, so the data storage system must also make the data's retention period immutable. This eliminates malicious shortening of retention periods.

System Resilience

We take a fairly broad view of system resilience, and believe it must be applied not just to the backup infrastructure itself, but to the entire ecosystem of tools, technologies, and processes related to data backup and recovery. Specifically, the backup infrastructure must be resilient to failure and attack, such as component or network unavailability, or network time server (NTP) manipulation in order to maliciously expire backed up data. It must also be easy to configure the use of distributed and heterogeneous backup data storage, such as across geographies or infrastructure types. Resilience is also improved by separating the backup data from the backup management system so that compromising the backup system will not also compromise the data storage. In fact, look for a backup management system that, in the event of compromise or failure, can be reconstituted without impacting your ability to access and restore backed up data.

The system also needs to be resilient to both expected and unexpected changes in the enterprise environment. Expected changes include planned addition or removal of infrastructure components, including the adoption of hybrid or cloud-based applications and data. That is, the backup system must be able to efficiently capture and store enterprise data, regardless of its source location or technology. Unexpected changes typically occur during incident response or disaster recovery (DR), and are most often categorized as support for recovery into dissimilar environments. When an organization is recovering data, it's entirely possible that the recovery environment will be running in a different location or infrastructure type. For example, a flooded on-premises data center may necessitate recovery into a cloud-based environment, with ongoing operations being there for an extended period of time. Therefore, the backup system must support both recovery into this different environment and new backups from this production environment moving forward.

The backup data storage system itself, in addition to providing immutable data storage, should be easily hardened. This may take the form of a pre-hardened appliance or an admin-configurable system with clear hardening recommendations, which will be more suitable for sophisticated enterprises.

Platform Requirements

In addition to these principles, Zero Trust Data Resilience has two additional Platform Requirements, which should be part of your criteria for selecting vendors and systems. Specifically, the vendors and tools you're evaluating should be viewed from the perspective of how they can ensure that your organization and your teams can meet these requirements.

- Proactive Validation
- Operational Simplicity

Proactive Validation

Ensuring proper system operation requires that the system be monitored, and all functional aspects and processes be validated. This has two aspects to it. First, the backup system should be monitored for network, performance, and security. That is, this system should be treated like any other high value production system.

Second, and most importantly, the validity of the backed-up data — and the reliability and efficacy of the recovery processes — must be regularly validated. By definition, recovery of backed-up data is going to occur at unexpected times, and likely in a high-stress environment. It's important that the organization have a well-understood, well-documented, and well-rehearsed process. There also needs to be multiple people capable of performing this to account for staff vacations, unavailability, and turnover.

Keep in mind that, although this requires an investment of time and energy, this demonstrates operational maturity, and is an "insurance policy" in case of disaster. Also note that "disaster" doesn't have to mean a literal disaster, or a major event such as a data center flooding. For example, one enterprise we worked with experienced a runaway automated workflow due to a programming error, which resulted in the deletion of significant amounts of production data in their financial management system. This wasn't a literal disaster, and was prevented from becoming a figurative disaster by using their (validated) data recovery processes.

Also, the backup management system should have the direct or indirect capability to organize backups across a malware infection timeline. That is, it should be able to detect (or be informed of) malware infections and categorize backups as clean, questionable, or compromised, depending on when they were captured.

> Note that data validation and recovery processes must also respect any data privacy and data residency requirements. This can add complexity and risk, so it must be done thoughtfully, with knowledge of both the data contents and the organization's legal and compliance obligations.

Operational Simplicity

Our final principle is Operational Simplicity, which we define as a system that is easy enough for your organization to confidently operate while still providing enough capability, scalability, and sophistication to fully meet your enterprise's needs. That is, a system that's appropriate for your organization.

This is important – we've seen enterprises struggle to make use of and operationalize systems that are too complex for their organization's size, team, skills, and needs. This results in limited benefits, frustration, and an inability to deliver either security maturity or business value. One set of attributes to look for in a backup vendor is their relative strength at orchestration and automation. Vendors with strong capabilities in their platforms will be faster and easier to operationalize.

Wrapping up this section, each of these principles are woven into the new Maturity Model extensions discussed later in this document, and will also be apparent in the reference architecture we discuss next.

Zero Trust Data Resilience: Reference Architecture

Data backup architectures will, by necessity, vary across different enterprises, given the enormous variability of network, application, and data infrastructures, among other factors. Even so, there are common architectural elements due to common Zero Trust principles, which must be present in any Zero Trust Data Resilience architecture.

Our reference architecture is shown in Figure 2, and illustrates the key requirements in this kind of system. Note that this depicts the environment from the perspective of the backup management system. Regular, everyday access by users and systems into the production systems would also be controlled by Zero Trust PEPs, but this is omitted in the diagram for clarity.

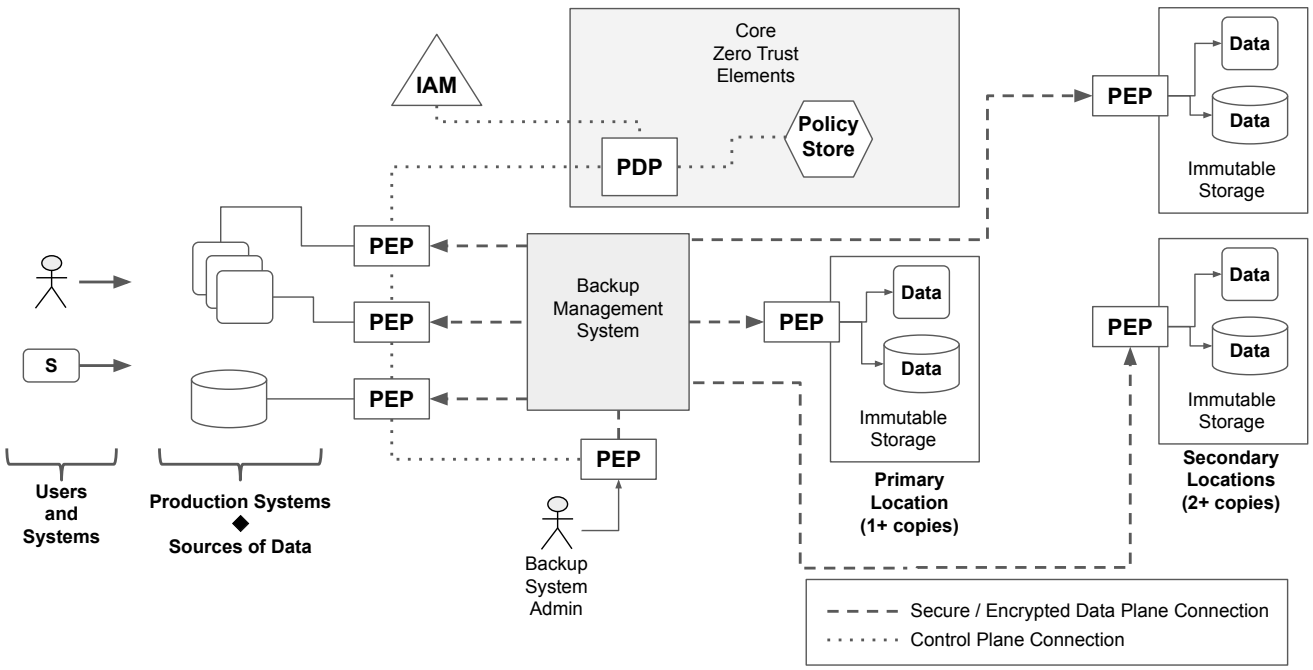

Figure 2 : Zero Trust Data Resilience: Reference Architecture

First, note the core parts of any Zero Trust architecture – the centralized Policy Decision Point (PDP), which delegates identity authentication to the enterprise Identity and Access Management (IAM) system. The PDP relies on its policy store to make access decisions for authenticated identities, including both human and non-person (system) identities. In this architecture, the PDP is making access decisions for the backup management system. These decisions are communicated via the control plane (shown as dotted lines) with the Policy Enforcement Points (PEPs), which logically sit inline between the backup management system and the sources of data to be backed up and the target backup locations.

The architecture also includes a recommended structure for backed up data. In addition to the data immutability requirement, enterprises should aim to keep at least one copy in a primary location, which has a low-latency network connection to the intended restoration site. This allows for rapid backup snapshots, that encourage more frequent recovery points and faster recovery times. Of course, the primary location is often co-located with production systems, so our reference architecture also illustrates the goal of having at least 2 copies of the data in secondary locations.[4] These must be geographically isolated from the primary location to achieve resilience against a regional disaster. The likely tradeoff is a slower network connection, which may result in lower frequency recovery points and longer recovery times.

Note that the backup management system is deliberately separate from its storage tiers. This allows the backup system to seamlessly distribute the backed up data across multiple immutable, and geographically distributed repositories. It also enables enterprises to select backup storage repositories that provide the best combination of performance, price, and operational simplicity for their unique requirements. It also provides an additional layer of security by controlling the communication via a PEP.

4 - There are various schools of thought around the number of backups in various locations, often referred via mnemonics such as 3-2-1 or 3-2-1-1-0.

Zero Trust Data Resilience: Extended Maturity Model

While the principles and reference architecture we've proposed for Zero Trust Data Resilience are universally applicable, they cannot be fully and immediately applied into most enterprises. LIke most aspects of Zero Trust, they must be planned for and adopted incrementally. The standard way to model and communicate this is through a maturity model. As we mentioned in the introduction, we're following the de facto standard CISA Zero Trust Maturity Model framework and extending it with four new functions that comprise our principles and requirements.

These new functions are:

- Access to Enterprise Data and Systems
- Access to Backup Storage and Data
- System Resilience
- System Monitoring and Validation

These ZTDR extensions to the maturity model are depicted in Figures 3 through 6, which show how each of the four new functions should be advanced across the standard maturity tiers: Traditional, Initial, Advanced, and Optimal.

For each of the functions, we've identified expected attributes for each maturity level. The model thereby depicts the improvements and changes an organization needs to make in order to advance in maturity for each function. Next, we examine each of the functions in turn as it progresses through the maturity tiers..

Access to Enterprise Data and Systems

This function is defined as the means and mechanisms by which the backup management system (BMS) has access to the source data that it's responsible for backing up.

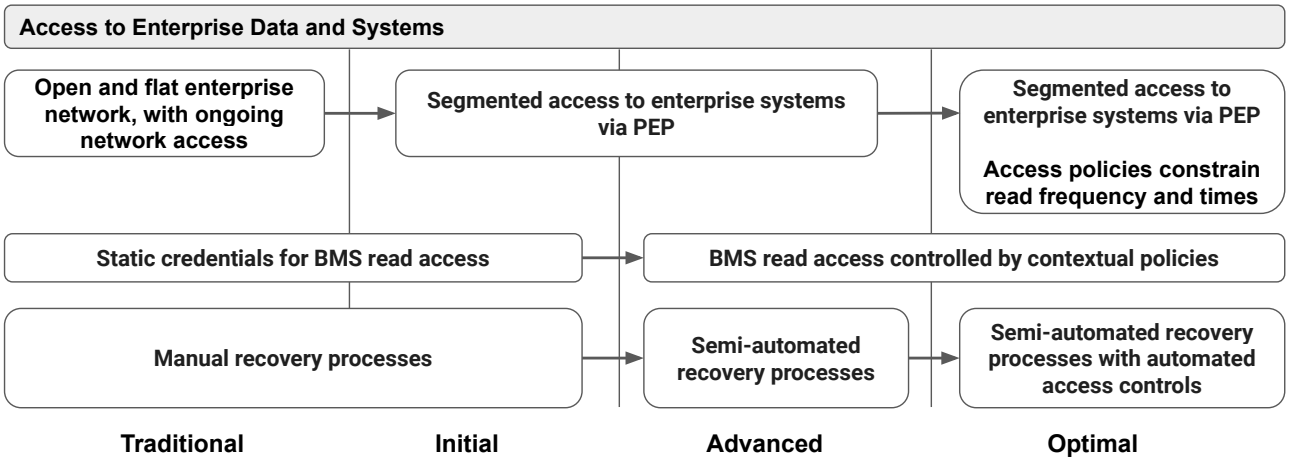

Figure 3 - Access to Enterprise Data and Systems: Maturity Model

At the **Traditional** level of maturity, the enterprise has a flat, open network, and the backup management system has ongoing and unimpeded network access to the source systems. The BMS uses static credentials, like an API key, stored username/password, or a certificate, in order to authenticate and read the source data. When the enterprise uses the BMS to recover a system, they rely on manual processes.

In order to advance to the **Initial level**, the enterprise must begin enforcing better network segmentation, and restrict BMS access to enterprise systems via a Zero Trust Policy Enforcement Point, introducing the principle of least privilege.

When the enterprise is at the **Advanced** level, they will have introduced contextual access policies for BMS access to enterprise data and systems, thus better utilizing dynamic Zero Trust policy enforcement capabilities. They will also have started to use automated recovery processes with some manual steps for initiation and process validation.

At the **Optimal** level, the organization will have enhanced their usage of access policies, to constrain BMS access to only permitted time periods or active recovery events. This further enforces the principle of least privilege.

Access to Backup Storage and Data

This function is defined to be the means and mechanisms by which the backup management system has write and read access to the backup storage, and the data stored there.

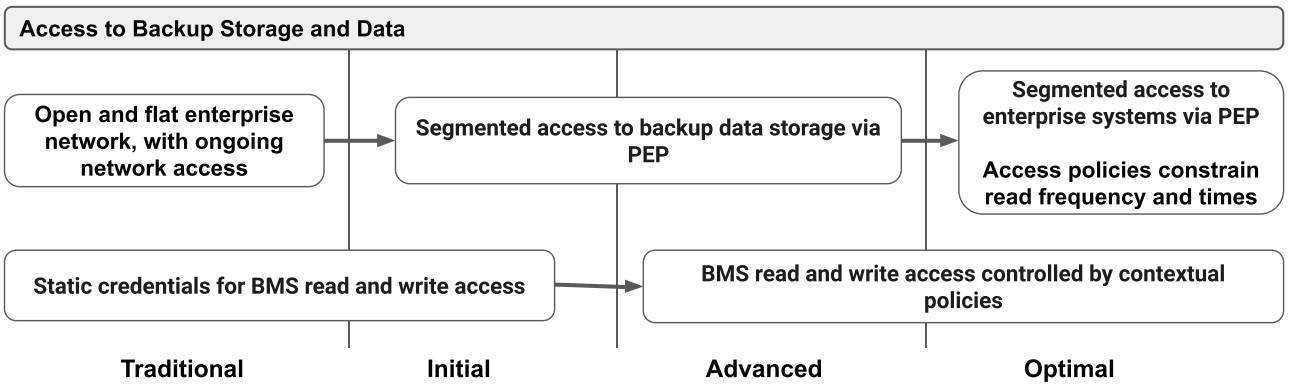

Figure 4 - Access to Backup Storage and Data: Maturity Model

At the **Traditional** level of maturity, the enterprise has a flat, open network, and the backup management system has ongoing and unimpeded network access to the backup storage system, and to the backed-up data stored there. The BMS uses static credentials, such as an API key, stored username/password, or a certificate, in order to authenticate and write to storage, and read the stored data.

In order to advance to the **Initial** level, the enterprise must begin enforcing better network segmentation, and restricting BMS access to the backup storage and stored data via a Zero Trust Policy Enforcement Point, enforcing the principle of least privilege.

When the enterprise is at the **Advanced** level, they will have introduced contextual access policies for BMS access to the backup storage system, and the stored data. This better utilizes dynamic policy enforcement capabilities within the enterprise.

At the **Optimal** level, the organization will have enhanced their usage of access policies, to constrain BMS access to storage, to only permitted time periods, or during active recovery events. This further enforces the principle of least privilege.

System Resilience

This function is defined to be the characteristics of the backup system with respect to its resistance to system failure, component failure, or malicious activity.

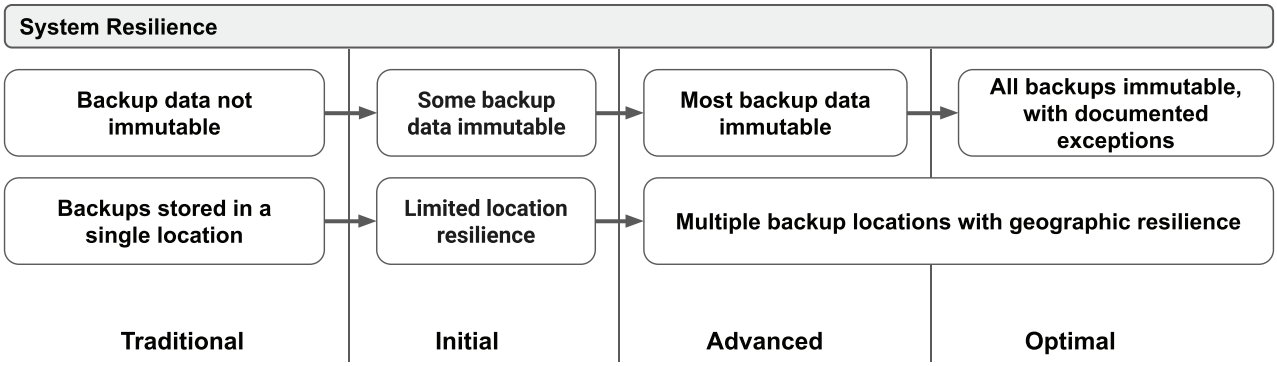

Figure 5 - System Resilience: Maturity Model

At the **Traditional** level of maturity, the organization uses mutable storage for backup data, putting its integrity and availability at risk. They are also typically storing backups in just a single location, thus subjecting the organization to complete loss in the event of a regional disaster.

As the organization moves to the **Initial** level, they must begin utilizing immutable storage for some of their data backups, and introduce some limited location resilience for those backups.

At the **Advanced** level, the organization will mostly use immutable backup storage, ideally prioritized by data sensitivity and criticality. They will also have introduced and operationalized the use of multiple backup storage locations, across distributed geographies.

When the enterprise is at the **Optimal** level they will have shifted over to fully utilizing immutable backup storage with any exceptions documented and approved. New data sources and applications will by default use immutable backup. This level provides the organization with maximum resiliency against regional disasters and malicious actors.

System Monitoring and Validation

This function is the tools and processes by which the enterprise ensures that their backup management system and backup storage is operating correctly, and that the enterprise is capable of executing a recovery process when needed.

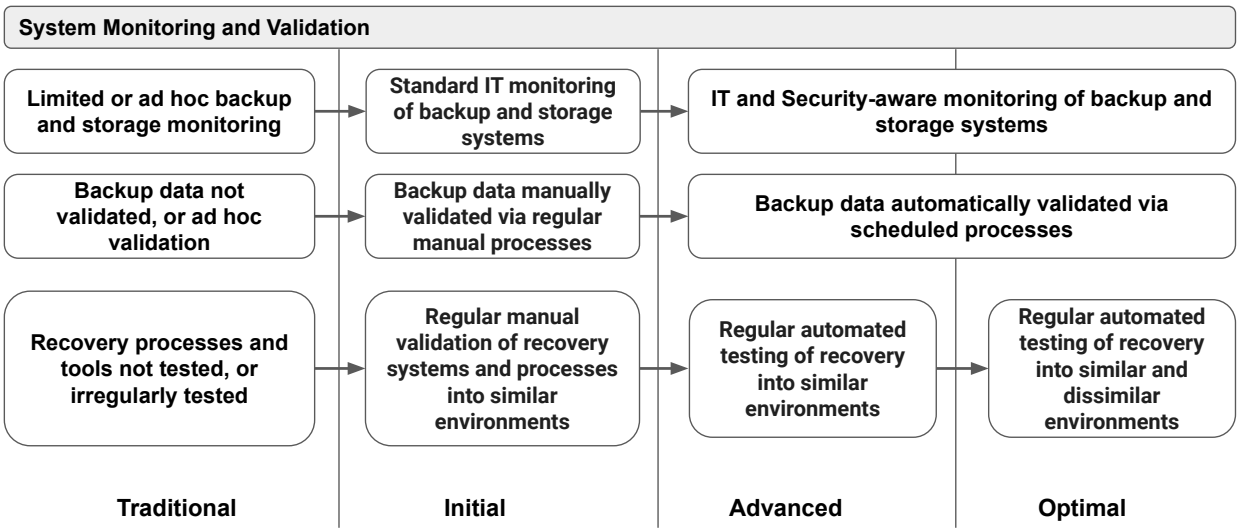

Figure 6 - System Monitoring and Validation: Maturity Model

At the **Traditional** level of maturity, the enterprise will only perform basic monitoring of the backup and storage infrastructure, often reflecting lower overall IT and operations maturity. The organization may not validate the backed up data or only perform periodic (i.e, manual and infrequent) checks. Plus, the enterprise will not be regularly testing recovery tools and processes to make them well understood, documented, and repeatable.

At the **Initial** level, they will have adopted a standardized level of IT and operational monitoring of the backup and storage system. They will also institute regular validation of backed up data via manual processes. They'll also have implemented regular (manual) validation of recovery processes to ensure institutional knowledge and comfort with them.

At the **Advanced** level, organizations will have deployed both IT and security monitoring tools and processes for backup and storage systems. And, they will automatically validate backed up data with scheduled checks that report and escalate any anomalous results. This will include automated testing of recovery tools and processes into environments similar to production.

At the **Optimal** level, the organization will have enhanced the sophistication of their recovery testing to test it for recovery into dissimilar environments.

Maturity Model Summary

Taken as a whole, these new functions define a set of capabilities and an expected set of competencies mapped across the four Zero Trust maturity levels. They provide a practical roadmap and guide for enterprises who are looking to bring their data backup and recovery systems into their Zero Trust initiative.

Conclusion

Zero Trust is a demonstrably better way to approach information security, and as security leaders, we have an obligation to bring this strategy into our enterprises. Current Zero Trust architectures and maturity models are solid starting points, but are incomplete. In particular, data backup and recovery requirements and approaches are absent from them.

Traditionally, enterprises have treated backup and recovery as being within the domain of IT, but the prevalence of ransomware and the nearly complete digitization of business necessitates that security leaders broaden their scope to include this.

In this whitepaper, we introduced the concept of Zero Trust Data Resilience, with a set of core principles, a reference architecture, and extensions to the Zero Trust Maturity Model. We believe that by embracing this Zero Trust Data Resilience approach, enterprises will have a clear and concrete pathway to stronger defenses, more efficient operations, and faster recovery. Enterprise data is too important for us to not apply security best practices, and Zero Trust is the most effective way to do so.

Note to Revision 1.1

Based on reader feedback, this revision includes a reorganization of the ZTDR principles into three core principles, and two platform requirements. The content of these sections is essentially unchanged. This revision also includes minor clarifications and rewording in a few areas.