DELL Technologies

Whitepaper

Always Authenticated and Authorized

Zero Trust with PowerProtect Cyber Recovery

This whitepaper is intended to show how Dell PowerProtect Cyber Recovery follows the Zero Trust Architecture and is a key component for a Cyber Resilience strategy that fits into a Zero Trust Framework.

July 2025

Revisions

Date	Description
July 2025	Initial release

Acknowledgments

Author: Kevin McDonough, Advisory Systems Engineer

Support: Dell Technologies Americas Technology Office

Table of contents

Executive summary

A good cyber resilience posture starts with knowing who or what has access to your companies' vital assets. Deploying a Zero Trust Architecture (ZTA) within your infrastructure gives you peace of mind concerning the integrity of your devices, and that your data is secure. In short, a Zero Trust Framework requires all users, and devices, whether inside or outside your organization's network, to be authenticated, authorized, and continuously evaluated for secure access to applications and data.

The United States Federal Government is pushing hard for agencies and commercial organizations alike, to adopt ZTA, with guidance from the Cybersecurity, and Infrastructure Security Agency (CISA). The security model of zero trust is not new. The tenets of information assurance, the concept of protecting and defending data availability, integrity, authentication, confidentiality, and non-repudiation have not changed. What has changed is the landscape. Our modern enterprise environments now consist of many interconnected segments, corporate networks, on-premises infrastructure and applications, cloud-based infrastructure and applications, enablement of remote and mobile workforce environments. The ongoing and increasing connectivity to non-traditional Information Technology elements like IoT devices. This modern infrastructure evolution means organizations must acknowledge the reality that traditional data protection methodologies, and perimeter-based defenses, are no longer adequate.

Organizations need a comprehensive approach to cyber-risk mitigation that goes beyond threat detection. Lax security on internal networks has enabled intruders who breach the perimeter to launch a full-scale cyber-attack on strategic assets in the data center. Insider attacks- intentional or otherwise- constitute a significant percentage of enterprise breaches and must be dealt with effectively, which has been exacerbated by the extension of the workforce beyond traditional network boundaries. The introduction of Bring your own device (BYOD) policies and the extension of data centers to public clouds have further blurred enterprise boundaries, making the enterprise more vulnerable to threats, and creating the need for a more effective and resilient ZTA approach imperative for any organization.

What is Zero Trust?

Zero Trust Principles are nothing new. However, due to the increases in ransomware, security vulnerability exploits, and supply chain attacks, there is a renewed interest in Zero Trust Principles. The zero trust concept can be traced as far back as 1994 in Dr. Stephen Paul Marsh's doctoral thesis. In 2010, John Kindervag, an analyst at Forrester Research, challenged the security model of using firewalls to create a perimeter within the IT infrastructure. He coined the term "zero trust," which centered around the idea that an organization should not trust anything inside or outside its perimeter. Zero Trust has evolved further into describing the process and technologies of implementing trust on a transactional basis. The framework focuses on authentication, authorization, and ensuring there is no implicit trust as much as possible, providing granular levels of authority, enforcing least privilege policies, while maintaining the goal of IT, which is the availability of services, and minimizing delays in authentications.

As defined by **National Institute of Standards and Technology (NIST) Special Publication 800-207**, *Zero Trust*, provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least-privileged per-request access decisions in information systems and services" under the assumption that the network is compromised. A zero trust architecture (ZTA) uses zero trust principles to plan industrial and enterprise infrastructure and workflows. Zero trust assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned). Authentication and authorization (both subject and device) are discrete functions performed before a session to an enterprise resource is established.[1] A *Zero Trust Architecture* is an enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies. *Zero Trust Security*, also known as perimeter-less security, describes a comprehensive approach to the design and implementation of IT systems. The main concept behind zero trust security is "never trust, always verify."

Cybersecurity & Infrastructure Security Agency's (CISA's) Zero Trust Maturity Model is one of many roadmaps, specific to the federal government, for agencies to reference as they transition towards a zero trust architecture. The goal of the maturity model is to assist agencies in developing trust strategies and implementation plans and present ways in which various CISA (Cybersecurity and Infrastructure Security Agency) services can support zero trust solutions across agencies.

The maturity model, which includes five pillars and three cross-cutting capabilities, is based on the foundations of zero trust. Within each pillar, the maturity model provides agencies with specific examples of a traditional, advanced, and optimal zero trust architecture. CISA drafted the Zero Trust Maturity Model in June of 2021 to assist agencies in complying with Executive Order 14028 "Improving the Nation's Cybersecurity."[2]

The National Counterintelligence and Security Center (NCSC), a part of the Office of the Director for National Intelligence, provides an alternate but consistent approach in identifying key principles behind *Zero Trust Architectures.* The NCSC calls out the essential need for a single strong source of user identity, user authentication, machine authentication, additional context, such as policy compliance and device health, authorization policies to access an application and access control policies within an application.

It should be noted that employees may view a "never trust" system as cumbersome because they must constantly prove their legitimacy to an individual system. While these challenges of Zero Trust make in more difficult to maneuver, it is still, by far, the most effective and secure principle for companies to use because it adds layers of defense and is extremely successful in preventing data breaches and data destruction. In terms of attack vectors, it is the most effective tool against unauthorized lateral movement.

The Tenets of Zero Trust

NIST Special Publication 800-207 Zero Trust Architecture states that a zero trust architecture is designed and deployed with adherence to the following tenets:

- All Data Sources and computing services are considered resources
- All communication is secure regardless of network location
- Access to individual enterprise resources is granted on a per session basis
- Access to resources is determined by dynamic policy-including the observable state of client identity, application/service, and the requesting asset- and may include other behavioral and environmental attributes.
- The enterprise monitors and measures the integrity and security posture of all owned and associated assets
- All resource authentication and authorization are dynamic and strictly enforced before access is allowed
- The enterprise collects as much information as possible related to the current state of assets, network infrastructure and communications and uses it to improve its security posture

Zero Trust is not a single architecture; it is more a set of guiding principles for workflow, operations, and systems design. The importance of a strong Identity, Credentialing, and Access Management (ICAM) practice cannot be overstated. It is a key component. Without effective ICAM practice, all underlying security practices are at risk. An effective ICAM solution must include necessary tools and controls that can capture and store user login details, facilitate the assignment and revocation of user access credentials, and oversite of a central database of user roles, levels, and access privileges.

Logical Components of Zero Trust

Access is granted to a resource via a Policy Decision Point (PDP) and a corresponding Policy Enforcement Point (PEP). The basic tenets of authentication and authorization ensure that the subject is authentic and is granted access to a resource.

The implicit zone represents an area where all entities are trusted to at least the level of the last PDP/PEP gateway and applies a set of controls to all traffic beyond the PDP/PEP checkpoint. Zero trust provides a framework to move the PDP/PEP checkpoints closer to the resource with the goal of explicitly authenticating and authorizing all subjects, assets, and workflows within an enterprise. The PDP is broken down into two logical components: the policy engine and the policy administrator**. It is important to note that a critical component of a ZTA is the separation of the control plane and the data plane**. Enterprise assets can reach the PEP component, but the PEP is the only component that accesses the policy administrator.

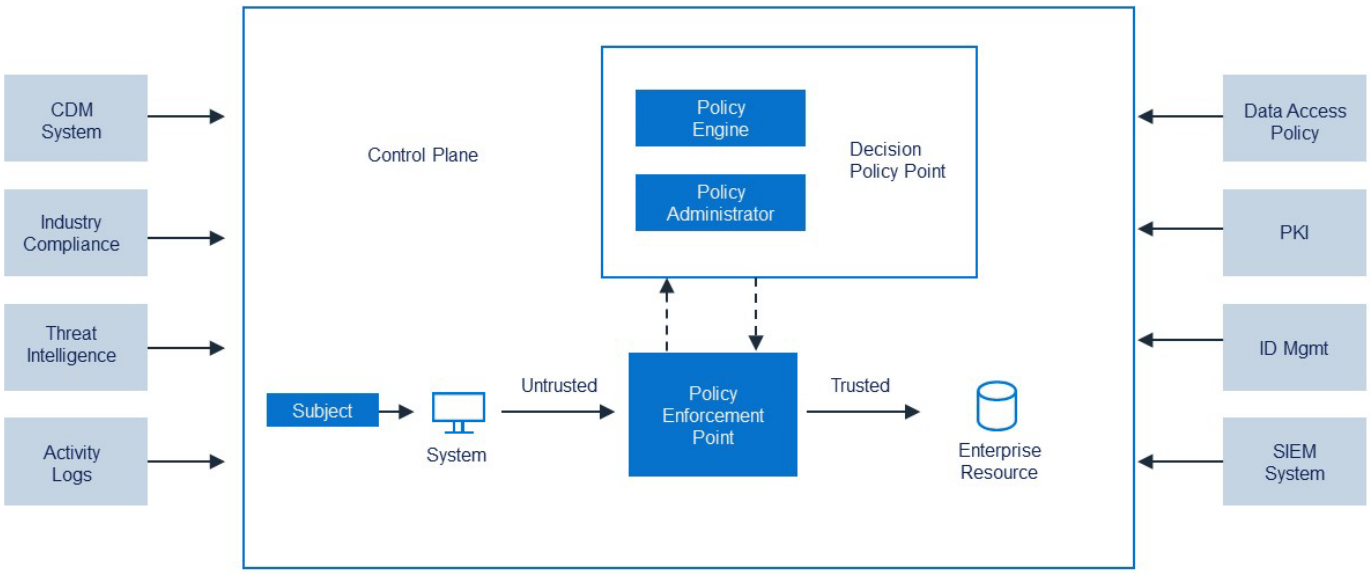

Figure 1: Zero Trust Logical Components NIST SP (Special Publication) 800-207

Dell Technologies 7 Pillars of Zero Trust

Dell Technologies follows the government architecture model and guidelines from the NIST SP 800-207 Zero Trust and information security model. The goal of this model is to make it easier for cybersecurity leaders to assess their current cybersecurity maturity level. It also provides a way to discuss Dell Technologies solutions and how they fit into the zero trust principles and architecture.

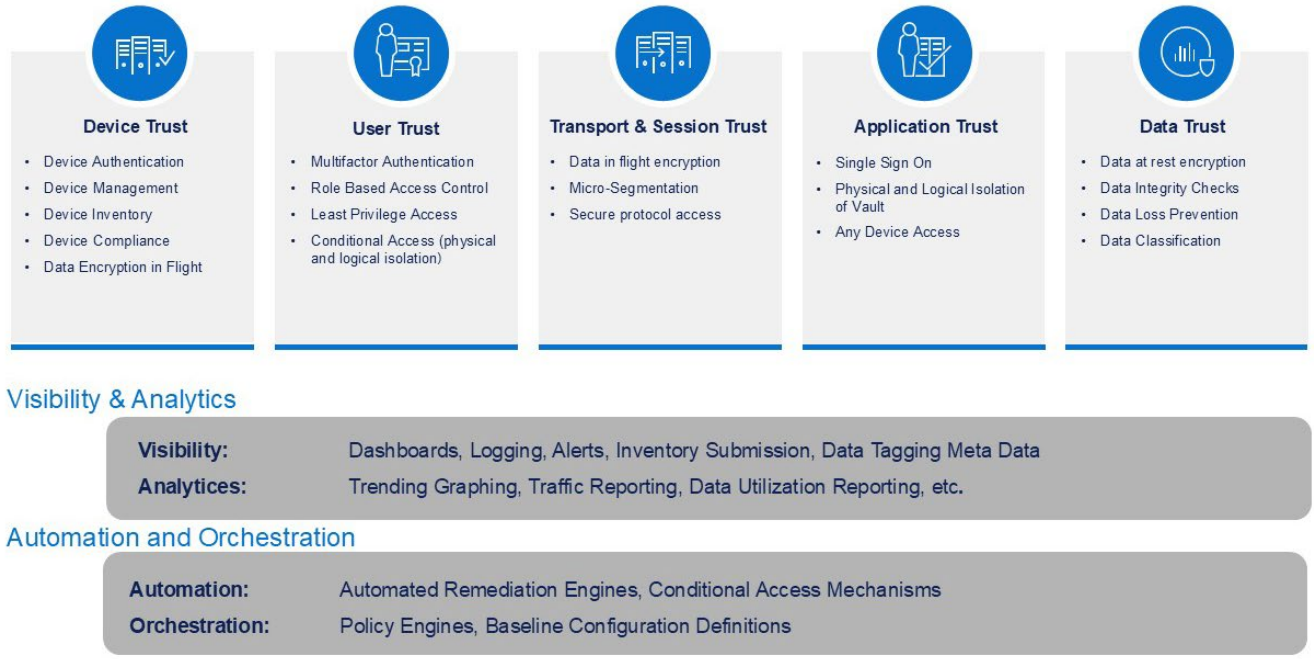

Figure 2: Seven Pillars of Zero Trust

1. Device Trust: Defined as any physical device within an enterprise.
2. User Trust: Defined as a user, administrator, and service level accounts.
3. Transport/Session Trust: Defined as the communication path utilized to move into, across, and out of an enterprise network.
4. Application Trust: Defined as both local and cloud applications that enter, work within, or leave the network for data access.
5. Data Trust: Defined by the organization as key assets used to execute the function and mission of the organization that can be held within the enterprise and extended into cloud services.
6. Visibility and Analytics: Defined by the resources from the 5-Pillars that should be enabled, to the fullest extent, to allow for analysis of the secure state and function of the pillar definition.
7. Automation and Orchestration: Defined using the visibility and analytics output to perform policy enforcement, baseline configuration definitions, automated remediation, and conditional access models.

The unfortunate reality is no enterprise can eliminate all cybersecurity risks. However, when it is combined with existing cybersecurity policies, guidance, identity and access management, monitoring, good cyber hygiene, a properly implemented and maintained zero trust architecture can reduce overall risk and protect against cyber threats.

Dell PowerProtect Cyber Recovery Vault

The PowerProtect Cyber Recovery vault provides the last line of defense against cyber-attacks. It offers multiple layers of protection to provide resilience against cyber-attacks, even from an insider threat. It moves critical data away from the attack surface, physically isolating it within a protected part of the data center and requires separate security credentials and multifactor authentication for access. Additional safeguards include an automated operational air gap to provide network isolation and eliminate management interfaces which could be compromised. PowerProtect Cyber Recovery automates the synchronization of data between production systems and the vault creating immutable copies with locked retention policies. If a cyber-attack occurs, you can quickly identify a clean copy of data, recover your critical infrastructure and data, and get your business back up and running.

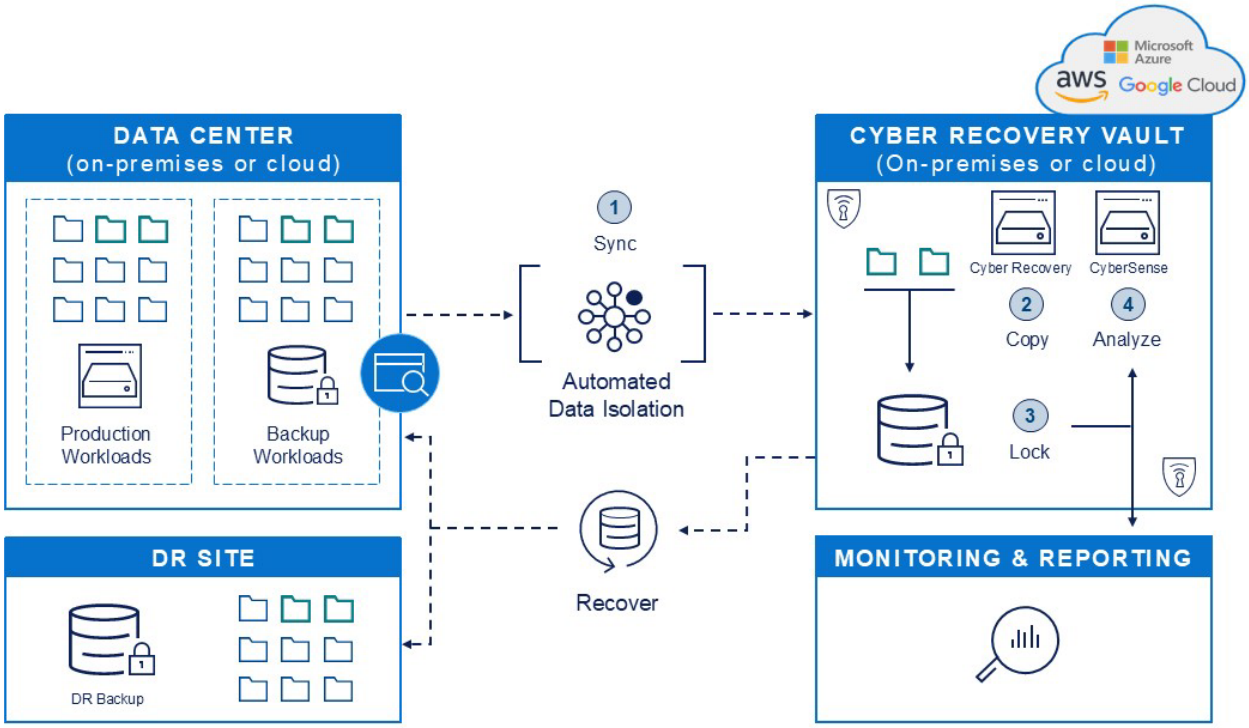

Figure 3: PowerProtect Cyber Recovery Solution Abstract

The vault operates in 4 basic steps:

1. The PowerProtect Cyber Recovery Software and PowerProtect Data Domain (DD) devices separate the untrusted zone from the implicit Trust Zone by establishing a policy decision point and a policy enforcement point. Data representing critical applications is synced through the logical air gap, which is unlocked by the management server inside the vault and replicated into the vault target storage. The air gap is then re-established. This makes access control enforcement as granular as possible, making the vault invisible to unauthorized resources (human, software, hardware, networking, etc.). The replication plane, control plane, and data plane are separated in the vault. At this point Isolation is established.

2. A copy of that data is made. Vault retention is configurable, with most customers keeping copies for 2 weeks and for as long as 1 month. This provides several effective tools in responding to a cyber-attack. Unlike a natural disaster event, you must determine the last known good copy of your data prior to recovery or restoration. Being able to roll back to the last known good copy is essential given the average adversarial dwell time.

3. The data is then retention locked to further protect it from accidental or intentional deletion. **It is important to understand that immutability does not equal invulnerability**. In simple terms, immutable storage alone is best used to prevent data alteration through “normal” means. Therefore, immutability in a production setting adds a layer of defense and should absolutely continue to be used in an overall data protection model. It cannot be viewed as or used as the last line of defense in the event of a catastrophic attack. Immutability must be tied to

isolation to ensure that your critical data is safe and available to recover from an attack. Without isolation, immutable storage is not enough; it may be enough to survive a run of the mill, unsophisticated attack, but it will certainly not stand up to a targeted and sophisticated attack. Immutability alone offers no protection against system overrides, clock-based attacks, system factory reset, altering or eliminating retention policies, kernel access, firmware corruption, boot lock, snap corruption, or physical access.

4. The data is then presented to our analytics engine, CyberSense performs a full content scan of the data (file metadata, document header and documents content) replicated into the vault. This includes scanning unstructured files, databases, and core infrastructure. CyberSense is an important component in enabling speedy recovery after an attack, by determining whether a data set is valid and usable for recovery to a 99.99% confidence in finding corruption. CyberSense uses the following steps in making this determination. CyberSense analyzes the data in its native backup software format, so there is no need for the original backup software in the vault because the analysis is done without rehydrating the backup image. This capability is critically important in defending against sleeperware or zero-day exploits. The vault is essentially a “zero oxygen environment” so any dormant malware will never be able to execute. CyberSense analytics makes over 200 observations per file. Analysis is performed by machine learning algorithms on analytics to determine if an attack on the data has occurred by looking for indicators of compromise. This process is typically performed after each new replication. This creates observation points which can be compared to previous observations to see how data has changed.

Forensic reporting and analysis tools are available to assist with the investigation process. At this point it is critically important to understand that disaster recovery is very different from cyber recovery. In a cyber recovery scenario, confidence in the integrity of the data is paramount. In other words, am I able to recover clean data? Where is my last known clean copy of data? Just as important is quickly determining the “blast” radius of the attack. What was affected? Whose credentials were used. What files were corrupted? What malware/ransomware/attack vector was used? When did the attack begin? The answers to all these questions are quickly determined using CyberSense in the vault.

After the declaration of an event, the incident response team can quickly begin mitigation and remediation measures because CyberSense has already performed the prerequisite identification and detection tasks/activities required before moving on to mitigation and remediation activities like containment and eradication, helping the incident response team quickly determine the most appropriate recovery technique. Remember that cyber restore and cyber recovery are not the same. A cyber restore does not rebuild a service; it simply restores data to an existing infrastructure. Cyber recovery includes the rebuilding of a service and all the required components and dependencies invoking the correct people, process and technology.

PowerProtect Cyber Recovery Aligned to the 7-Pillars

Organizations should use the following features to align PowerProtect Cyber Recovery optimally with the Zero Trust Pillars. In the optimal stage of zero trust maturity, the security processes are fully automated, and dynamic attribute assignments are based on automated triggers. Least privilege access is dynamically adjusted across the enterprise, ensuring just enough access for assets and dependencies. There is interoperability, continuous monitoring and centralized visibility, providing comprehensive situational awareness. The Dell zero trust strategy is to fully embrace zero trust principles and align our portfolio, services, and partner ecosystem to support our customer's zero trust journey.

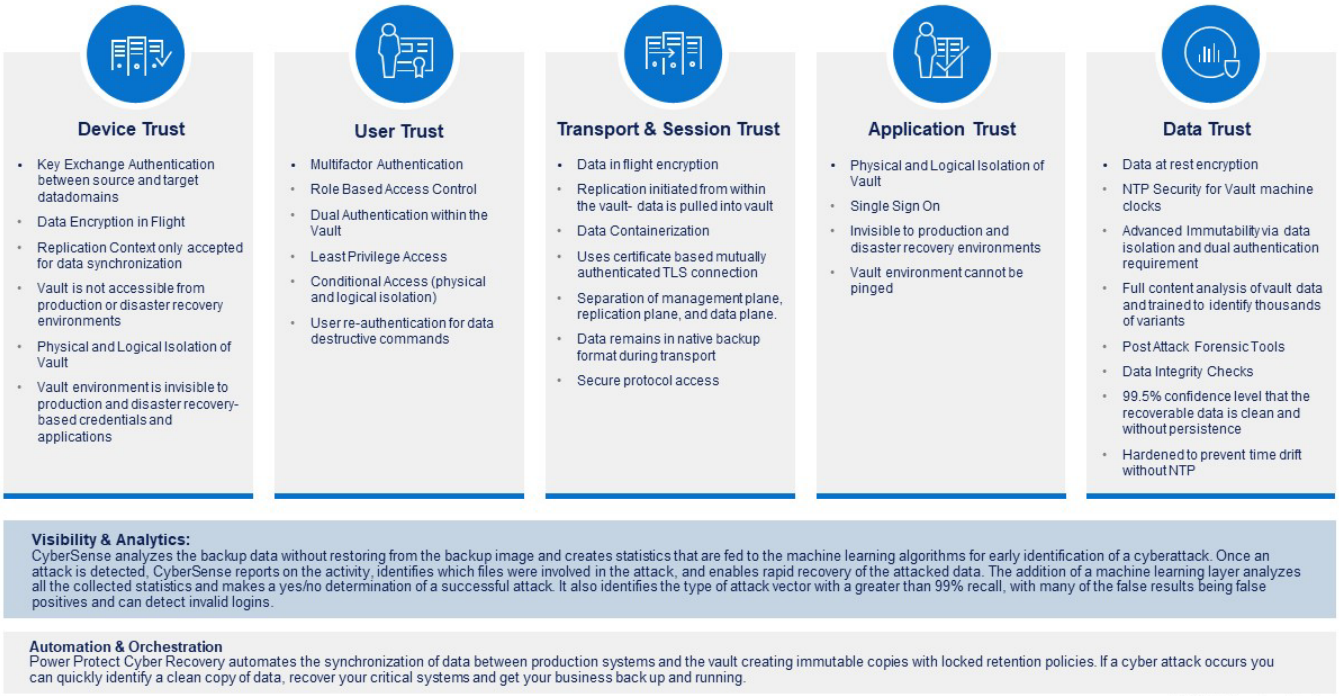

Figure 4: PowerProtect Cyber Recovery Product Feature Mapping

Conclusion

Overall, Dell distinguishes itself by providing both the strategy and the practical tools, partnerships, and blueprints organizations need to accelerate their zero trust journey, reduce complexity, and ensure that all elements of their environment are addressed in a comprehensive, integrated way.

For additional information please also reference the Dell PowerProtect Data Domain Cybersecurity Framework and Zero Trust Architecture Whitepaper

[1] https://doi.org/10.6028/NIST.SP.800-207
[2] https://www.federalregister.gov/executive-order/14028