



White Paper

Data-centric Zero Trust

An evolution in cyberdefense

Justin Spears and Matt Trudewind, NetApp

May 2024 | WP-7366

Abstract

This white paper defines data-centric Zero Trust and describes its core concepts. It also explains how NetApp® ONTAP® built-in features and capabilities, along with adjacent portfolio technologies, can help organizations secure their data by using a data-centric Zero Trust approach that aligns with the seven DoD Zero Trust pillars: User, Devices, Applications and Workloads, Data, Network and Enablement, Automation and Orchestration, and Visibility and Analytics.

TABLE OF CONTENTS

| | |
|---|-----------|
| What is data-centric security? | 3 |
| An analogy..... | 3 |
| What is Zero Trust? | 4 |
| What is data-centric Zero Trust? | 4 |
| Core concepts | 4 |
| Process for deploying a data-centric Zero Trust architecture | 5 |
| How NetApp secures your data using data-centric Zero Trust | 6 |
| Data microsegmentation..... | 6 |
| Data-centric access policies | 7 |
| Fine-grained encryption..... | 8 |
| Dynamic Access Control | 8 |
| Data monitoring and analytics | 10 |
| Automation and orchestration..... | 11 |
| Data lifecycle management | 12 |
| Summary | 13 |
| About the authors | 13 |

LIST OF FIGURES

| | |
|--|---|
| Figure 1) The seven DoD Zero Trust pillars. | 6 |
|--|---|

“This is the way the world ends. Not with a bang but a whimper.” That famously melancholy line from T.S. Eliot accurately represents the state of the cyberworld today. We live in a digitally connected world that is under attack from every angle. It’s no longer a question of *if* you will be attacked but *when*—and will you be ready for it? However, these attacks are coming not from bullets whizzing past our ears but instead manifest in the form of seemingly benign hiccups in our daily computing that eventually turn out to be cataclysmic losses that we never saw coming. How can we fight what we cannot see?

Although there are a variety of approaches to cyberdefense, NetApp has chosen to create an entirely new paradigm of cybersecurity focused on protecting data where it resides through a layered approach to security coupled with a mindset of least privileged access. We call this paradigm *data-centric Zero Trust*.

To understand the nuances of this approach, let’s first dive into the concepts of data-centric security and Zero Trust.

What is data-centric security?

Data-centric security is an approach to information security that emphasizes protecting the data itself rather than focusing solely on securing the underlying systems or network infrastructure. Traditional security approaches typically focus on securing the perimeter of the network and the endpoints, relying on technologies like firewalls, intrusion detection systems, and antivirus software. Although these measures are crucial, they often fall short of adequately protecting sensitive data from various threats.

The fundamental idea behind data-centric security is that data is an organization’s most valuable asset, and protecting it should be the primary objective of any security strategy. This approach recognizes that data exists and flows across multiple systems, networks, and devices and that data security must be applied consistently and persistently, regardless of its location.

The key difference between data-centric security and traditional security is the shift in focus. Traditional security concentrates primarily on securing the network, systems, and applications from external threats. In contrast, data-centric security emphasizes protecting the actual data throughout its lifecycle, regardless of how it is accessed, where it resides, or how it moves throughout the hybrid cloud. It recognizes that data is a valuable target for attackers and implements security measures that persistently follow and protect the data, regardless of its location or state. This approach provides an added layer of security and can better protect against data breaches, unauthorized access, and data leakage.

An analogy

Data-centric security can be likened to the protection of a highly valuable and sensitive object, such as a precious gem, in a museum. Imagine the museum as an organization, the precious gem as the data, and the visitors as users and systems interacting with the data. Traditional security would focus on securing the museum’s entrance, the doors, and the hallways, making sure that unauthorized individuals don’t enter. This approach is akin to securing the network perimeter and endpoints.

In contrast, data-centric security is analogous to placing an impenetrable, transparent, and high-security glass case around the precious gem. This glass case represents the encryption, access controls, and other protective measures applied directly to the data. It ensures that even if an intruder manages to breach the museum’s perimeter security or gain access to the area, the precious gem remains securely locked in the glass case.

Furthermore, the glass case would have additional security features, such as biometric scanners that allow only authorized personnel with the correct fingerprint or iris scan to access the gem. This approach is analogous to the granular access controls and strong authentication mechanisms used in data-centric security.

This analogy highlights the fundamental difference between traditional security and data-centric security. Traditional security focuses primarily on fortifying the museum’s physical structure, while data-centric

security prioritizes safeguarding the valuable asset itself by applying protective measures directly to the data, regardless of the environment in which it resides.

What is Zero Trust?

In 2010, John Kindervag, then a Forrester Research analyst, promoted the idea that an organization should not extend trust to anything inside or outside its perimeters. In that process, he created the concept of Zero Trust. At its core, Zero Trust is a way to think about a security strategy based on the idea of “trust no one, verify everything.” Often called the Zero Trust security model or the Zero Trust framework, it is an approach to designing and implementing a security program based on the idea that no user, device, or agent should have implicit trust. Instead, anyone or anything — a device or system — that seeks access to corporate assets must prove that it should be trusted. Zero Trust assumes that there is no traditional network edge; networks can be local, in the cloud, or a combination or hybrid with resources anywhere, as well as workers in any location.

The rise of 21st-century IT architecture, with cloud computing and an explosion of endpoint devices that require access to enterprise systems from outside the corporate IT environment, killed the concept of traditional perimeter security. With the value of Zero Trust at the forefront of modern security thinking, this approach has become an evolution of defense in depth, which is necessary to secure a global hybrid multicloud infrastructure.

What is data-centric Zero Trust?

Now, with a clear understanding of data-centric security and the fundamentals of Zero Trust, we can begin to conceptualize how these two concepts can merge and culminate in the new paradigm of data-centric Zero Trust. This information security approach combines the principles of data-centric security and Zero Trust architecture. It shifts the focus from placing trust solely on network boundaries to placing trust on individual data elements themselves. The key idea behind data-centric Zero Trust is to ensure that data is protected regardless of access method, its location, or the networks it traverses.

In a traditional Zero Trust architecture, the assumption is that no entity, whether internal or external, should be inherently trusted by default. Every access request is verified and authenticated, and access is granted based on the principle of least privilege and assuming the worst. However, in data-centric Zero Trust, the emphasis is on applying Zero Trust principles specifically to the data, rather than just to the access requests.

Core concepts

The core concepts of data-centric Zero Trust include:

- **Data microsegmentation.** Data-centric Zero Trust employs data microsegmentation, which involves breaking down data assets into smaller units or elements and placing strict access controls and encryption mechanisms around each element. Each data element is treated as an independent entity with its unique access permissions, regardless of its location.
- **Data-centric access policies.** Access to data is determined by specific policies that are applied directly to the data elements themselves. These policies can be based on factors such as user identity, device characteristics, location, and the sensitivity of the data. Access is granted in a granular manner, ensuring that only authorized users can access specific data elements.
- **Fine-grained encryption.** Encryption plays a vital role in data-centric Zero Trust. Each data element is encrypted individually, and access to the encrypted data is controlled by encryption keys. This approach ensures that even if unauthorized individuals gain access to the data, they cannot view or manipulate it without the appropriate decryption keys. It is essential to ensure encryption both at rest and in transit.

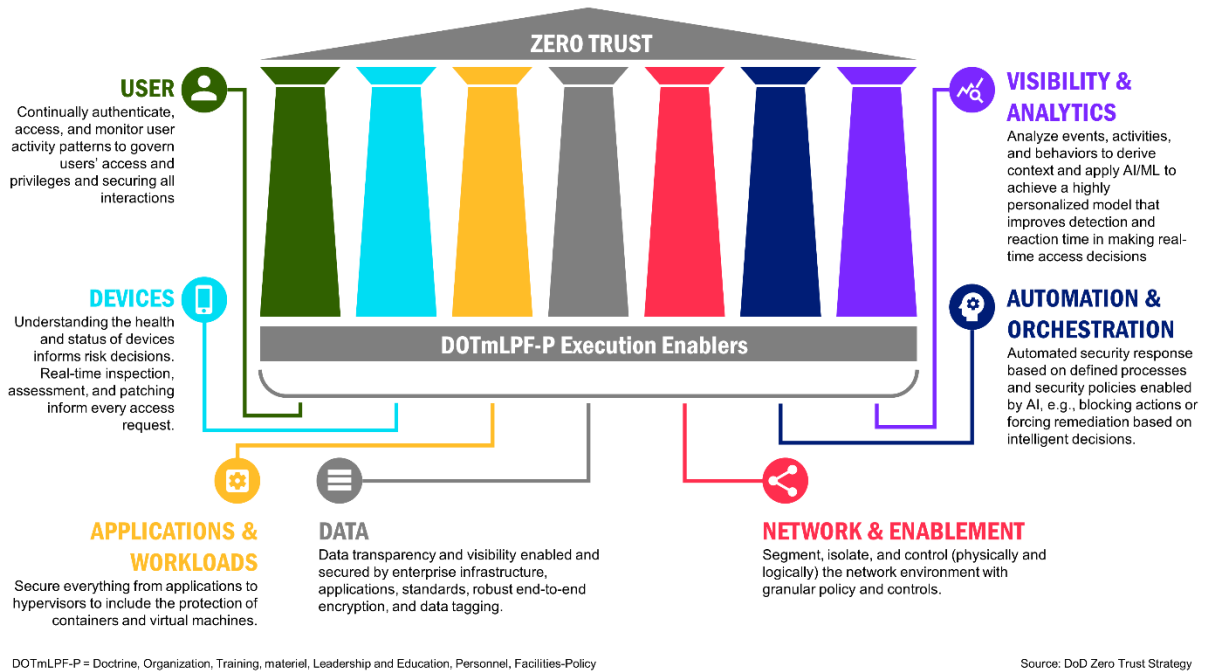
1. **Dynamic Access Control.** Data-centric Zero Trust employs Dynamic Access Control, which can adapt based on the context and state of the data and the accessing entity. These controls constantly evaluate trust levels, user behavior, and other contextual factors to determine whether access should be granted, denied, or adjusted in real time.
- **Data monitoring and analytics.** Real-time monitoring and analytics are crucial components of data-centric Zero Trust. This approach includes continuous monitoring of data access, usage patterns, and anomalies. Machine learning, artificial intelligence, and behavioral analytics can be used to detect and respond promptly to unauthorized activity or potential data breaches.
- **Automation and orchestration.** Automated security response based on defined processes and security policies enabled by AI or in response to triggered events, such as blocking actions or forcing remediation based on intelligent decisions.
- **Data lifecycle management.** Data-centric Zero Trust also considers the entire lifecycle of data, from creation to deletion. It includes identifying sensitive data, secure data storage, secure data transmission, secure backup and recovery processes, and secure data disposal.

The key difference between data-centric Zero Trust and traditional Zero Trust is the focus on protecting the data itself rather than only verifying and controlling access requests. By applying Zero Trust principles directly to the data, organizations can achieve a higher level of security and ensure the confidentiality, integrity, and availability of their data regardless of its location or the networks it traverses. To here

Process for deploying a data-centric Zero Trust architecture

At this point we've demonstrated the importance of applying data-centric Zero Trust security, but how does an organization deploy such an architecture? Recently, the Department of Defense (DoD) published its Zero Trust strategy, which is an evolution of the CISA Zero Trust strategy. The strategy projects that in 5 years the DoD will have adopted a risk-based Zero Trust framework across the defense ecosystem. This includes integrating Zero Trust principles across the five cyberfunctions that many of us know from the [NIST Cybersecurity Framework \(CSF\)](#): Identify, Protect, Detect, Respond, and Recover. This strategy is oriented around seven DoD Zero Trust pillars, as shown in Figure 1.

Figure 1) The seven DoD Zero Trust pillars.



As shown in the graphic, data is the central pillar of the DoD Zero Trust strategy. However, our data-centric Zero Trust approach to security goes beyond this pillar and helps provide security controls in almost all of the other pillars. Using the DoD model as a North Star and our innovative data-centric approach to Zero Trust, an organization can begin the process of deploying a data-centric Zero Trust security architecture; but that is certainly easier said than done.

How NetApp secures your data using data-centric Zero Trust

Often, the most difficult part of completing a job is having the right toolset. Without it, you might spend an excessive amount of time trying to complete the job, and in some cases, without success. The same is true for deploying a data-centric Zero Trust security model to secure your data. You need the right tools.

This is where NetApp can help. As the creator of the world's most secure storage, NetApp provides a robust set of features and capabilities that significantly ease an organization's deployment of a data-centric Zero Trust security model. Many of our security tools are built in and are included in our data-centric secure storage operating system, NetApp® ONTAP®, at no additional cost.

This section describes how the core tenets of NetApp's data-centric Zero Trust architecture can be implemented and how they map to the DoD Zero Trust strategy.

Data microsegmentation

Typically thought of as a network concept, microsegmentation concepts can also be applied to data. Much as networks are segmented either physically (switches, routers, access points, firewalls, etc.) or logically (VLANs), data can be segmented in the same manner. NetApp ONTAP separates data logically by using storage virtual machines (SVMs). You can think of the SVM as a policy enforcement point (PEP).

An ONTAP system can contain many different SVMs, each with its own storage pool allocation. SVMs limit data access to specific storage clients, like a single host mounting a LUN or a domain user group

accessing a specific file share from the SVM. Additionally, an administrator managing one SVM can be prevented from managing a different SVM on the same storage platform.

Taken a step further, it's also possible to limit which networks a particular SVM can access by using ONTAP IP Spaces. With IP Spaces, data access for a specific SVM is limited to a particular set of networks and interfaces that the IP Space has access to. If an attacker gains admin access to an SVM admin account, their access would be limited to that SVM and what that account's credentials could do. Other SVMs on the same ONTAP storage system remain protected, because they are part of different IP Spaces (when configured).

Basic support for networking segmentation concepts such as VLAN tagging and firewall service policies to limit data port access are also built into NetApp ONTAP storage systems.

ONTAP SVMs, IP Spaces, VLAN tagging, and built-in firewall service policies make excellent PEPs, preventing access to resources (data) in the trust zone through microsegmentation.

DoD Zero Trust alignment

Microsegmentation aligns well with the Data pillar of the DoD Zero Trust framework, and it also aligns with the Network and Environment pillar by creating a foundation of segmentation for high-level activities.

Data-centric access policies

Data access policies can restrict access to resources (data) by using rules, policies, and attributes. All of these items together can be considered a rule set (containing rules, policies, and attributes for access to resources/data). In a data-centric Zero Trust architecture, all rule sets should be defined with least privilege access. NetApp ONTAP enables this rule set for both file- and block-based data workloads.

How do these rule sets function in regard to file-based workloads? ONTAP supports both SMB/CIFS and NFS file-sharing protocols, along with their basic file-based permission access. NTFS access control lists (ACLs), NFSv4 ACLs, SMB share permissions, and NFS export policies are all supported. However, applying these file permissions as part of a data-centric access policy provides only the most basic protection with some rules and policies in the rule set. This is because access is based primarily on user account permissions, group membership, and perhaps client IP address. To complete the data-centric access policy rule set enabling data-centric Zero Trust, attributes must also be used as decision points in rules and policy settings.

Applying rules and policies for accessing files based on attributes is sometimes referred to as attribute-based access control (ABAC). ABAC is applied differently for SMB/CIFS access versus NFS, but in both cases, elements such as device identity, location, and data sensitivity are used to define the data access policy in the rule set. For more information about how ABAC is used specifically in NetApp ONTAP to define access to data resources, see "Dynamic access control, later in this document.

Rule sets made up of other policies and rules also apply to block-based workloads. These workloads include data accessed via Fibre Channel Protocol (FCP), iSCSI, and NVMe variants such as NVMe /TCP (over TCP/IP) and NVMe-oF (over fabrics). LUNs are the data resources that are accessed in these workloads by initiators (clients). NetApp ONTAP includes a LUN permissions concept known as *initiator groups* (igroups). An igroup allows access to block-based LUN data only if the clients/hosts initiator name matches the WWPNs, IQNs, or EUIs specified in the igroup. However, using a single rule set or control for access to LUNs is not enough to be considered as part of a data-centric Zero Trust architecture. Therefore ONTAP provides authentication mechanisms for LUN access as well. For iSCSI CHAP authentication using SHA-256, hashing is supported. For NVMe, DH-HMAC-CHAP is also used with SHA-256 hashing. For FC workloads, NetApp partner switches enable additional security by using the concept of *zoning*, where only allowed initiators in the target zone can access the data.

With the built-in capabilities of ONTAP for file and share access permissions, ABAC, igroups, and advanced block-based authentication and partner ecosystem support, both file- and block-based workloads can be protected by access policies that enable data-centric Zero Trust architecture.

DoD Zero Trust alignment

The use of access policies crosses multiple pillars of the DoD strategy. In addition to Data, these policies provide security controls that can be applied to the Applications and Workloads pillar as well as provide fundamental elements needed for Automation and Orchestration.

Fine-grained encryption

Encryption of data is crucial to preventing malicious actors (both internal and external) from accessing data they have obtained either by intercepting the data over the wire as it crosses networks or by gaining physical access to drives that data resides on. NetApp ONTAP enables you to encrypt both data at rest and data in flight in a variety of ways.

When accessing file data, ONTAP supports multiple methods of encrypting data in flight. For SMB/CIFS, ONTAP supports the SMB 3.0 encryption protocol, protecting the data from man-in-the middle (MITM) attacks. For NFSv3, IPsec can be used to encrypt all traffic between the client and the ONTAP SVM. NFSv4 supports encryption over the wire with both Kerberos (krb5p) and IPsec.

For block-based workloads (SAN), LUNs can be protected in transit with IPsec as well as when using the iSCSI protocol or the NVMe/TCP protocol. For FC LUN access, ONTAP supports FC-SP2 when using appropriate partner-supported FCP switches.

Replication of data between storage systems and tiering of cold data are other avenues for potential MITM attacks. ONTAP supports TLS 1.3 encryption by using Perfect Forward Secrecy (PFS) cipher suites to protect replicated and tiered data while in transit to their destination. All encryption over the network uses FIPS 140-2 validated encryption algorithms.

Encryption in flight is important for data access and replicated data, and also for management of data storage systems. NetApp ONTAP also supports TLS 1.3 encryption for all management traffic, whether you are managing an ONTAP system via the System Manager GUI or sending a command through the REST API. All CLI connections are also encrypted via SSHv2. Connections made to off-box destinations like key management servers using the Key Management Interoperability Protocol (KMIP), or connections to NetApp BlueXP™, also use TLS 1.3 for encrypted communication.

Regardless of the data access protocol used (SMB, NFS, iSCSI, FCP, etc.) to access data, all data can be encrypted at rest with multiple layers of FIPS 140-2 validated encryption algorithms. NetApp offers a physical layer with self-encrypting drives (SEDs) using the NetApp Storage Encryption (NSE) solution, which includes all storage efficiencies: deduplication, compaction, compression, etc. NetApp Volume Encryption (NVE) offers an additional layer of FIPS 140-2 validated software encryption, providing granular volume-based encryption at rest that includes all data storage efficiencies. NetApp NVE and NSE encryption-at-rest solutions allow organizations to take full advantage of AES-256-bit encryption technologies while also enjoying the benefits of space savings with storage efficiencies.

DoD Zero Trust alignment

The use of encryption crosses all seven pillars of the DoD strategy. No matter how the data is accessed, the encryption should be end to end, including User, Devices, Applications and Workloads, Data, Network and Environment, Automation and Orchestration, and Visibility and Analytics.

Dynamic Access Control

Some data is more sensitive than others. Data might be a general memo meant for the entire organization to view, or it could contain proprietary “need to know” information that only a select few people in the organization should be able to access. This is where Dynamic Access Control (DAC) plays a key role in enabling a data-centric Zero Trust architecture.

DAC adapts access to data based on the context and state of the data and the accessing entity. Information about data context, state, and accessing entity resides in user and device attributes. As

mentioned earlier in “Data-centric access policies,” attributes are a key component in applying rules and policies for accessing files. This is known as attribute-based access control, and a solid DAC mechanism relies on ABAC. NetApp ONTAP can apply ABAC to access files by using DAC features.

ABAC is applied differently for SMB/CIFS access versus NFS. We’ll start with the SMB/CIFS example that uses Microsoft Windows Active Directory (AD). These AD attributes include restricting access to the resource that is using the device identity, device claim, and user claim. For example, in addition to requiring an ACL to access a file, you can also require that the user account has a certain job title (user claim or attribute); that the device has a certain health status (device claim or attribute); and that the device is running a particular software version (device identity). Microsoft refers to this as Dynamic Access Control (DAC); NetApp ONTAP integrates and supports DAC for creating data-centric access policies.

ABAC for NFS client access policies can be applied by using Security-Enhanced Linux (SELinux) clients and NFSv4.2 MAC (mandatory access control) labels. MAC is part of the kernel in SELinux. It prevents access to file system objects based on local security policies, such as what processes are running on the local system and what actions the user account is attempting to take on the object. ONTAP fully supports assigning appropriate NFSv4.2 MAC labels to files in the export.

DAC is not limited only to data access; it must also be applied to data management. Data administrators are company insiders who may also represent a substantial threat to organizations. With virtual “keys to the kingdom,” data-centric Zero Trust principles must be applied to storage and data administrators as well. DAC should be applied to management tasks such as moving data from one location to another, providing access to data, or even deleting data.

NetApp ONTAP DAC enables features for management activities, including dynamic authorization (dynamic auth) and multi-admin verification (MAV).

Dynamic auth uses a trust score system to determine whether an administrator’s intended actions on data management should be allowed (for a high trust score), interrupted for additional verification and authentication (for a medium trust score), or denied (for a low trust score). Much like DAC applied to data access, dynamic auth goes into effect only when the data being taken action on is sensitive in nature. For example, if an admin user tries to create a new replication relationship to move data to another location, dynamic auth checks the admin’s trust score. Factors such as trusted device, location, time of day, and authorization history all factor into the admin’s trust score. A decision is then made based on the score. If the score is high (trusted), then the action is allowed. If the score is medium (somewhat trusted), the user must reauthenticate. And if the score is low (not trusted), the request is denied. These dynamic actions happen in real time, protecting the underlying data from suspicious data management actions.

Multi-admin verification also applies dynamically to data management actions that could be destructive. It requires that one or more additional admins must approve the intended action before it is allowed to proceed. MAV takes place at a layer below dynamic auth, so that the user must prove they are trustworthy before they can start an action that would invoke MAV. For example, if an administrator tries to delete a critical data volume, then after they have passed dynamic authorization checks, a request to delete the data volume must be approved by another admin of the same privilege level. This procedure ensures that disruption to the organization cannot occur even if the action is initiated by insiders with access.

Data management is not limited only to actual users, administrators, and their accounts. Machine accounts also manage data, often via an API like REST. This is quite beneficial for automated tasks that would be time consuming for a real-world administrator to oversee. However, much like administrator user accounts, machine accounts must also have built-in DACs to ensure that data is being appropriately handled to minimize deletion, exfiltration, and other malicious activities. Machine accounts can also be subject to dynamic auth and MAV, but there is an additional data-centric Zero Trust mechanism that can be applied to machine accounts in the form of token-based authentication.

Built on the OpenID Connect and Open Authorization 2.0 protocols, token-based authentication negates the need for a username and password and instead allows machine accounts to authenticate with a

token. The token can have very granular permissions for which actions can be performed, and the validity time period can be shortened, unlike a basic user account.

NetApp ONTAP provides token-based authentication for all REST API actions taken to manage data on the system. It accomplishes this authentication with support for OAuth 2.0 by using mutual TLS to securely issue commands to the storage system. Tokens can be issued only as needed for specific times and with specific permissions, which limits the attack surface.

By integrating with Active Directory DAC and NFSv4.2 MAC labels to apply ABAC, NetApp ONTAP provides Dynamic Access Control for data access. MAV and dynamic auth also provide additional DAC for data management, protecting organizations from insider threats. Token-based authentication offers further security enhancements for automated workloads, providing dynamic access only when needed. Together, all of these NetApp ONTAP capabilities enable DAC as a key component for a data-centric Zero Trust architecture.

DoD Zero Trust alignment

Dynamic Access Control is a key component of the Data pillar of the DoD Zero Trust framework. It also aligns with the User, Devices, and Application and Workloads pillars by providing signals for when harmful actions are taken against applications and workloads in the organization by using a user account or corporate device.

Data monitoring and analytics

Data anomalies are a key indicator of malicious activity. These anomalies indicate when data is abnormally changed, moved, or destroyed. Detecting anomalies as early as possible is crucial to prevent interruption in service. Extended downtime because of malicious activity can be devastating and costly to the organization. Automated anomaly detection is therefore a crucial feature in data monitoring to provide a data-centric Zero Trust architecture.

NetApp ONTAP includes Fpolicy, a built-in API that provides a mechanism for anomaly detection. FPolicy provides user and entity behavior on all data access that is to be consumed by off-box analytics servers. These off-box servers provide user behavior and entity analytics (UEBA) that can be used for forensic analysis or to make automated decisions about data access in real time. Read more about the automated aspects in the next section, "Automation and Orchestration."

NetApp provides an off-box server that uses the FPolicy API, an SaaS application called Cloud Insights Storage Workload Security (CI SWS). When an anomaly is detected through the FPolicy API, the application can display what data was accessed, by what user account, and by what IP address, and indicate whether the data has changed in any way.

User behavior is not the only way to notify about anomalies. An attacker might employ several stolen account credentials to do a similar data exfiltration, but might do it very slowly to avoid user behavior detection. This possibility makes analyzing the data itself for anomalies a logical step.

Autonomous Ransomware Protection (ARP) is a feature of ONTAP that detects anomalies in the data itself by analyzing the data directly at the file system layer. ARP examines the data for data entropy (Minimum and Average), file extension (Safe and Known ransomware extension list), and file header parser to detect whether the file header is encrypted.

In addition to User Behavior alerts from the FPolicy API, CI SWS also displays alerts from ARP, allowing simple and efficient management of data monitoring for anomalies from a single location.

In addition to monitoring for and reporting on data anomalies, traditional file and system log auditing still has its merits as well. Particularly for larger organizations that maintain a security operations center (SOC), sending file and system logs to a security and information event management (SIEM) system, where they can be consumed and acted on appropriately by the SOC, is a benefit. NetApp ONTAP

supports native file auditing or auditing through the FPolicy API. A SIEM system can ingest these audit logs. System audit logs can be sent to a remote syslog server by using TLS for encrypted transport.

Data monitoring and analytics is a core tenet of a data-centric Zero Trust architecture. NetApp ONTAP ARP, FPolicy, and built-in auditing, along with CI SWS, enable a simplified and easily consumable way to deploy such a data monitoring and analytics system. However, it's not enough to monitor and analyze data looking for abnormalities. Automated actions to protect the data when anomalies are detected are also necessary. For more information, see the next section, "Automation and Orchestration."

DoD Zero Trust alignment

Data monitoring and analytics is a significant feature of the Data pillar of the DoD Zero Trust framework. It also aligns with the User, Devices, Applications and Workloads, Network and Environment, and Visibility and Analytics pillars by providing insights into malicious actions that will affect other non-data critical infrastructure.

Automation and orchestration

When threats to data and anomalies are detected, it's not enough to simply generate alerts and set an alarm for admins and users to remediate the issue. Automated actions must kick off in real time to stop the active threat, mitigate the blast radius, and eliminate costly downtime. For the maximum positive effect, multiple orchestrated automated actions should also occur to truly enable a robust data-centric Zero Trust architecture.

As mentioned earlier, in "Data Monitoring and Analytics," NetApp ONTAP includes Fpolicy, a built-in API that provides a mechanism for anomaly detection. Cloud Insights Storage Workload Security uses this FPolicy API to detect and alert on data anomalies, and also to take automated actions to actively stop these threats.

When an anomaly is detected through the FPolicy API, CI SWS can block that user account from continuing to take harmful action against the data. It can also take other automated actions, such as sending an alert about the activity and making an automatic backup, to ensure a close recovery point. All of these automated actions are orchestrated by the CI SWS service. For example, if a user account starts to move large amounts of data in a way that is abnormal for the account, it is probably being used to exfiltrate data to an external network. CI SWS could block the user account, effectively stopping the exfiltration, while alerting the security team to the data anomaly and making a NetApp snapshot™ copy to provide a very close recovery point.

The Autonomous Ransomware Protection (ARP) feature of ONTAP detects anomalies in a different way than CI SWS. When ARP detects an abnormality, it automatically makes a backup copy and sends an alert to the security team. Orchestration is further enhanced by CI SWS, which monitors automated actions and alerts for ARP as well as ONTAP Fpolicy.

NetApp ONTAP ARP and CI SWS prevent both internal and external threats from continuing to damage organizational data by using automated and orchestrated actions that quickly remediate the situation. This Automation and Orchestration plays an important role as a complimentary component to data monitoring and analytics in a data-centric Zero Trust architecture.

DoD Zero Trust alignment

Automation and Orchestration plays a key role in the Data pillar of the DoD Zero Trust framework by proactively detecting and mitigating cyberthreats. It also aligns with the User, Devices, Applications and Workloads, Network and Environment, and Visibility and Analytics pillars by taking automatic actions to prevent malicious actors from affecting the non-data critical infrastructure.

Data lifecycle management

So far we have covered six of the key aspects to enabling a data-centric Zero Trust architecture, but perhaps the most encompassing aspect is data lifecycle management (DLM). DLM takes into account the entire lifecycle of the data, starting from when the data is created all the way until it is destroyed. Sensitive data must be identified and managed appropriately. Data must be protected at all times, using the six components of data-centric Zero Trust security covered in the previous sections.

To properly protect data throughout its lifecycle, it must first be accurately identified to determine its asset class and whether it is sensitive in nature. Treating all data the same can lead to improper controls on sensitive and critical data or too many controls on nonsensitive data, increasing the organization's costs.

The NetApp BlueXP™ unified control plane includes a service, BlueXP Classification, that makes identifying sensitive data relatively easy. BlueXP Classification can scan an organization's entire estate and search for data elements that contain personally identifiable information (PII), such as social security numbers or employee records that may need to be protected under regulations like GDPR and HIPPA. BlueXP Classification also supports custom scanning for data elements that may contain proprietary information about company trade secrets. BlueXP Classification can also be used to identify noncritical data like MP3 music files and other files that are taking up valuable corporate storage space. Data can be further divided into asset classes and types ranging from noncritical and low importance to highly critical and highly sensitive.

Once it has identified data by type, class, and sensitivity, an organization can determine the best way to secure the data throughout its lifecycle. To protect its asset and class type, data must be placed in the appropriate container by using data microsegmentation. Sensitive data must have the appropriate data-centric access policies in place, using rule sets that employ ABAC and DAC to prevent unauthorized access. Furthermore, sensitive data must use fine-grained encryption techniques, whether that data is at rest waiting to be accessed, in transit to another location, or actively in use by applications.

Data of all types and classes, including noncritical and sensitive data, must be actively monitored by analytics that can take automated and orchestrated actions when threats are detected. Many cyberattacks and threats start with low-level access to noncritical data and quickly spread. That's why it's important to equally apply data monitoring, analytics, automation, and orchestration to all data types and asset classes.

An important aspect of data lifecycle management is data recovery. If data is prematurely deleted or modified in a way that makes it unusable and inaccessible, rapid data recovery from backup is necessary. Data backups must also be protected from being deleted prematurely or modified in a way that makes them unusable and inaccessible. Backup data is a tempting attack vector for cybercriminals, who know that when the backups are compromised, they can inflict the most damage and potentially make the most money because their victims are not able to recover on their own.

NetApp ONTAP provides rapid recovery of data (petabytes in seconds) with immutable Snapshot backup copies, which can be further protected from deletion by using Tamperproof Snapshot copies. These copies cannot be deleted by any administrator account until the retention period on the Snapshot backup copies expires, making the copies indelible.

Data spills are another aspect of data lifecycle management that must be considered. Sensitive data could be accidentally moved or recovered to a less protected and less secure environment, making it an incorrect asset class for its current location. This is considered a data spill. Simply deleting the data does not remove all possibility of recovering and accessing it from that insecure location. In these cases, the data must be shredded securely. Usually this means a wipe of both the sensitive data and the other noncritical data residing in the same location, which can be time-consuming and potentially disruptive.

The Secure Purge feature of NetApp ONTAP shreds data by using AES-256-bit encryption techniques that allow the noncritical data to remain, while the sensitive data is removed and made permanently inaccessible. This feature eases the burden on administration teams when cleaning up data spills.

At some point in the data lifecycle, it's necessary to securely destroy and shred the data, to make sure that organizational data of any kind cannot be stolen and recovered later by a malicious actor. NetApp ONTAP securely sanitizes data on disk to NIST SP 800-88 standards, preventing recovery of the deleted data.

NetApp BlueXP Classification is the first step in getting data started correctly on its lifecycle journey. When combined with NetApp ONTAP capabilities in data microsegmentation, data-centric access policies, fine-grained encryption, Dynamic Access Control, data monitoring and analytics, and automation and orchestration, organizations can fully protect data throughout its lifecycle. Lifecycle management is critically important because it incorporates and relies on all of the other components that make up and fully enable a data-centric Zero Trust architecture.

DoD Zero Trust alignment

Data lifecycle management has perhaps the most crucial role in the Data pillar of the DoD Zero Trust framework by bringing together all of the other aspects to securely protect the data throughout its lifecycle. It also aligns with three other pillars in the DoD Zero Trust framework, because throughout its lifecycle data will be accessed by Users, Devices, and Applications and Workloads, which are also subject to lifecycle management.

Summary

The NetApp portfolio, through ONTAP and beyond, contains dozens of features and capabilities centered on data-centric v Zero Trust. From our robust access controls to industry-leading encryption to more dynamic features that assess user behavior and file operations in real time, NetApp offers the ability to secure your data wherever and however you use it, through a data-centric Zero Trust approach.

About the authors

Justin Spears is a consummate technical executive with more than 20 years of experience in cybersecurity, product development, data science, and engineering. Currently residing at the intersection of security, AI, and data, Justin's primary focus is driving the cybersecurity roadmap for all core NetApp products. In previous roles, he was responsible for product and engineering work at Dell and EMC and spent several years in academic research at his alma mater, the University of Pittsburgh. When not trying to circumvent virtual evildoers, you can find Justin powerlifting and renovating houses with his wife and four daughters.

Matt Trudewind, now on his second tour across NetApp in 12 years, is a senior product manager for security with a primary focus on cybersecurity, cyber resilience, and data-centric portfolio security. This focus includes but is not limited to Zero Trust, data governance and privacy frameworks, security tools, ransomware protection, and security best practices. Matt has 25 years of IT experience, with 19 of those years in the storage industry.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2024 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.