Technical Report

NetApp Solutions for Cybersecurity

Lee Vorthman, NetApp
March 2013 | TR-4130

Abstract

NetApp delivers a comprehensive portfolio of cybersecurity technologies that enables government agency data to be protected and trusted while still remaining accessible. NetApp® storage solutions collect and store cyber-related information, analyze data to drive actionable intelligence, and allow agencies to quickly recover from cyber-related incidents similar to disaster recovery. Cybersecurity technologies help protect data, know how data is being used, and minimize the risk that data will be lost in the event of an attack.

TABLE OF CONTENTS

LIST OF FIGURES

1 The Impact of Cybersecurity on Storage Requirements

Storage represents a significant portion of IT budgets. Data retention and archiving, authentication and authorization, data loss prevention, and privacy regulations all demand appropriate security technologies within storage solutions that provide both transparency and accountability.

With the growing volume of security sensors across agency networks, NAS- or SAN-based storage will be required to correlate the growing data to identify and block threats, as well as to provide remediation and forensic capabilities. And as algorithms for data protection and information sharing (for both unclassified and classified information) become more complex, the ability to analyze larger datasets becomes critical.

2 NetApp Cybersecurity Reference Architecture Overview

2.1 NetApp's Approach to Cybersecurity

NetApp's cybersecurity technologies help government agencies protect their data, know how their data is being used, and minimize the risk that data will be lost in the event of an attack. NetApp storage solutions collect and store cyber-related information, analyze data to drive actionable intelligence, and allow agencies to quickly recover from cyber-related incidents similar to disaster recovery. The NetApp approach:

- **Is open.** Use the tools you want to access your data using standard protocols and formats.
- **Is flexible.** Adapt NetApp solutions to specific agency use cases using modular building blocks.
- **Uses best-in-class components**. Partner with leading cyber companies to develop solutions to meet agency requirements.

Figure 1) NetApp cyber-reference architecture.

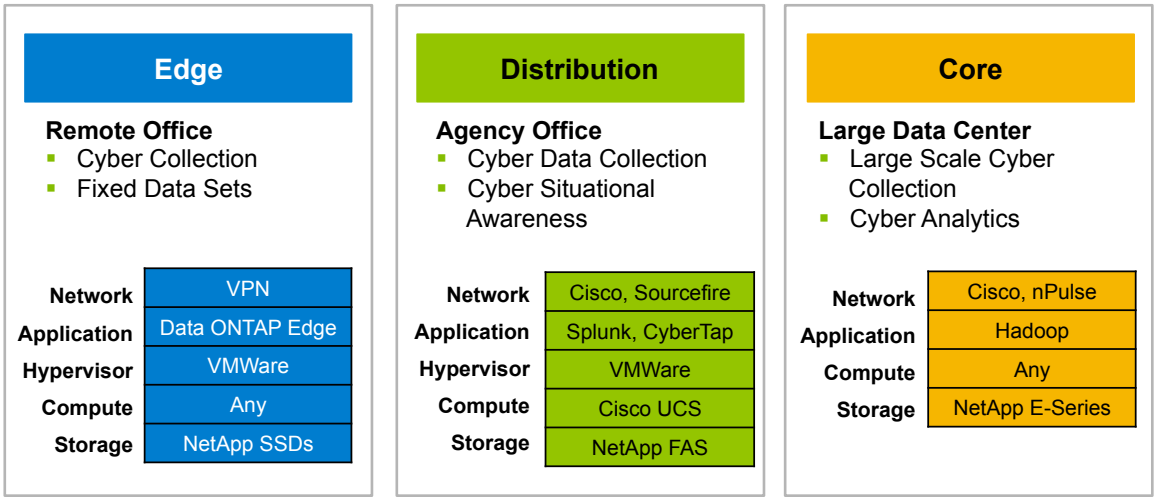

2.2 Next-Generation Cybersecurity Reference Architecture

NetApp's government customers are looking for new ways to collect, analyze, and distribute the volumes of cybersecurity data in their environments. As network speeds increase from 1 gigabit per second (Gb/sec) to 10Gb/sec, the challenge to collect all of the network traffic without dropping packets becomes more difficult. In addition, government agencies are faced with increased retention requirements for compliance and long-term threat analysis. A single 10Gb/sec link will generate 100TB of data per day.

And storing all of this data affects not only the physical footprint, but also power, weight, and cooling. And finally, the sheer volume of data and different reporting requirements across organizations makes it difficult to analyze and display the results.

NetApp's next-generation cyber-reference architecture addresses these challenges in the following ways:

- **Collection**. For customers with optical links at the gateway, an optical switch provides the flexibility to break out links for collection. Next, full packet captures in standard pcap format and NetFlow records are generated by an nPulse CPX. The nPulse CPX is connected by Fibre Channel to a NetApp E5460. The nPulse CPX accepts alerts from intrusion detection systems (IDSs) and will flag the related traffic. Analysts can then send this traffic directly to analysis software, such as Splunk or ArcSight. A single nPulse CPX and NetApp E-5460 will provide 20Gb/sec packet capture and approximately one day of storage in 5 rack units.
- **Analysis.** FlexPod® provides an ideal platform to perform cyber analytics. First, as a prevalidated reference architecture, it has been thoroughly tested for a variety of workloads. In addition, it provides a tremendous amount of flexibility. As storage or compute requirements grow, additional resources can be seamlessly added to the FlexPod configuration. For cyber analytics, analysis software such as Splunk can run in a virtual machine or on a dedicated server. Then each analysis VM or server simply accesses the stored packet captures and NetFlow records using a LUN presented by the E-5460.
- **Distribution**. The same FlexPod configuration that provides analytics can also host a virtual desktop infrastructure (VDI) to present the results of the analysis to a variety of users.

Figure 2) Next-generation cyber-reference architecture.

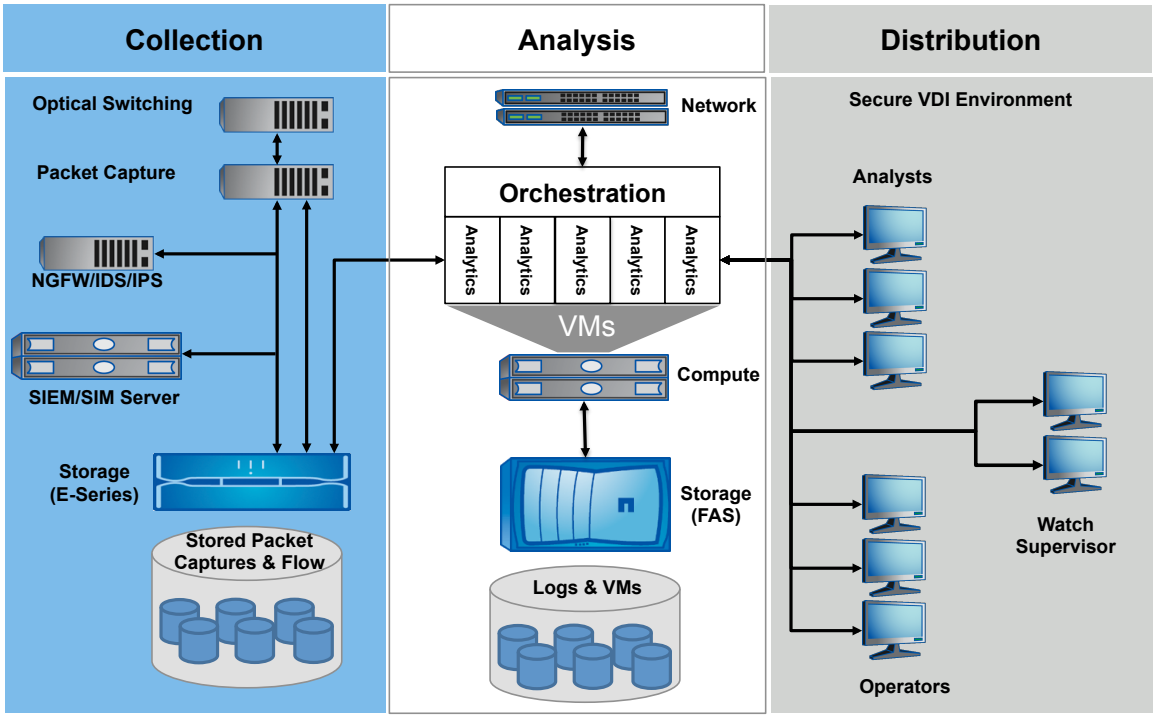

This architecture demonstrates the ability for high-speed collection, analysis, and distribution of cyber data. It is consistent with our approach to cybersecurity with the ability to collect and store data using open formats. The architecture is extremely flexible, and we have established partnership with best-in-class companies to provide unique capabilities.

3 High-Speed, High-Fidelity Packet Capture

As government agencies deploy ever-faster networks, one of the greatest challenges is enabling the network and security-monitoring infrastructure to keep up with the network itself. A key component of any monitoring infrastructure is full packet capture and storage that allows the enterprise to go back in time, examine network performance or security incidents, and answer the question "What happened?"

With the advent of 10Gb/sec IP links, the previous means of packet capture can no longer keep up without dropping packets. Further, the ingest load has grown such that the resulting packet capture and flow data must be written into a shared parallel file system for postingest processing by an open and extensible toolset.

3.1 NetApp and nPulse

NetApp and nPulse Technologies, the industry leader in high-performance flow and packet capture, have teamed to offer a packet capture solution that is setting new records for speed of capture and scalability of storage:

- Encrypting the packet and flow data as it is being stored at a rate of 24 gigabits per second
- Capturing multiple 24-gigabit data streams to a shared high-performance file system
- Making that data available for analytics in the open standard "pcap" format with the high performance required for exploiting the data

The nPulse and NetApp solution addresses both the bandwidth and content challenges, keeping networks protected while handling vast amounts of data in real time. Capturing data packets at the highest rates of speed without dropping any of the packets is essential to maintaining security across large, ultrafast networks.

Figure 3) High-speed, high-fidelity packet capture.

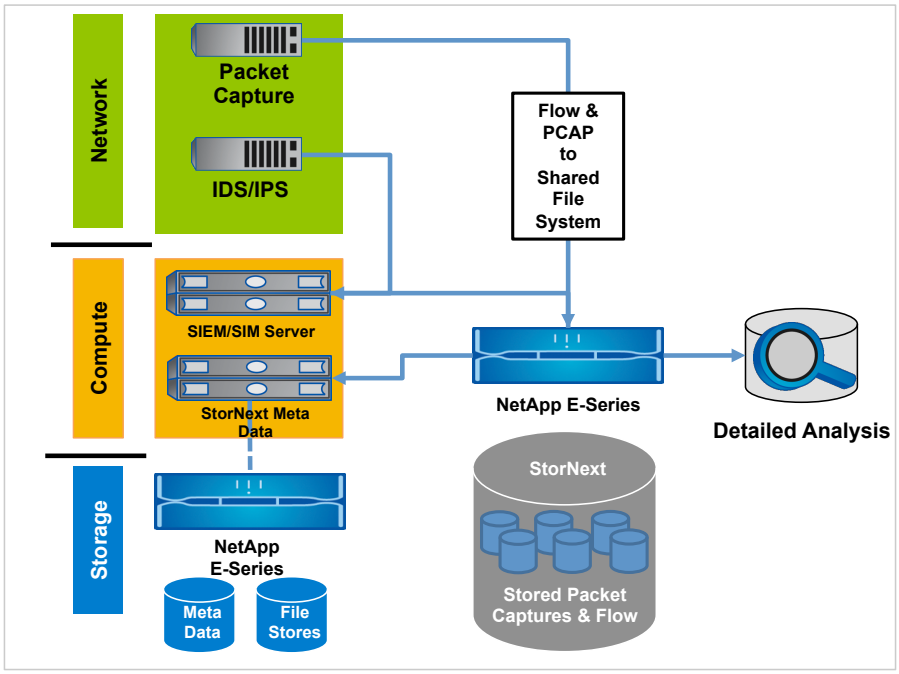

3.2 nPulse Capture Probe eXtreme (CPX)

CPX Flow and Packet Capture is a high-speed, multiterabyte traffic recording and analysis platform for network operations center (NOC) and security operations center (SOC) environments. The 20Gbps, continuous recording solution provides deep, high-fidelity indexed storage of network traffic for direct analysis or use with other security or monitoring applications. CPX delivers an easily searched, multilevel, deep-time view of network packets, trends, and events.

Full Packet Capture with Flow Indexing

Even on the busiest networks, CPX captures 100% of the traffic, time stamping every packet with nanosecond resolution, and extracting flow identification parameters. As traffic is streamed to disk, CPX generates a flow- and time-based index that allows rapid search and retrieval of targeted traffic from many terabytes of capture records. CPX's browser-based drill-down interface allows remote analysis of selected packets without the need to export entire PCAP files. Alternatively, traffic can be retrieved in industry-standard PCAP format for analysis by external tools such as Wireshark.

Key Features

- Continuous flow and packet capture with nanosecond resolution time stamping
- 100% sustained recording to disk at up to 20Gbps
- Real-time indexing of captured flows using time-stamp and flow attributes
- Real-time export of flow index in NetFlow v9 format for use with other flow analysis tools (1:1 NetFlow record generation)
- Rapid search and retrieval of target flows and packets
- Packet search on Layer 1–4 header info or user data pattern match
- Web-based, drill-down GUI for search and inspection of individual packets
- Graphical statistics and trending for link bandwidth, Layer 1–4 protocols, microbursts, and flows
- Up to 72TB of traffic recording in one appliance or expandable fiber-attached storage options
- Data export in industry-standard PCAP format
- Pivot2Pcap Web-based RESTful AI for integration and interoperability with existing monitoring or security applications

Figure 4) nPulse CPX capture statistics.

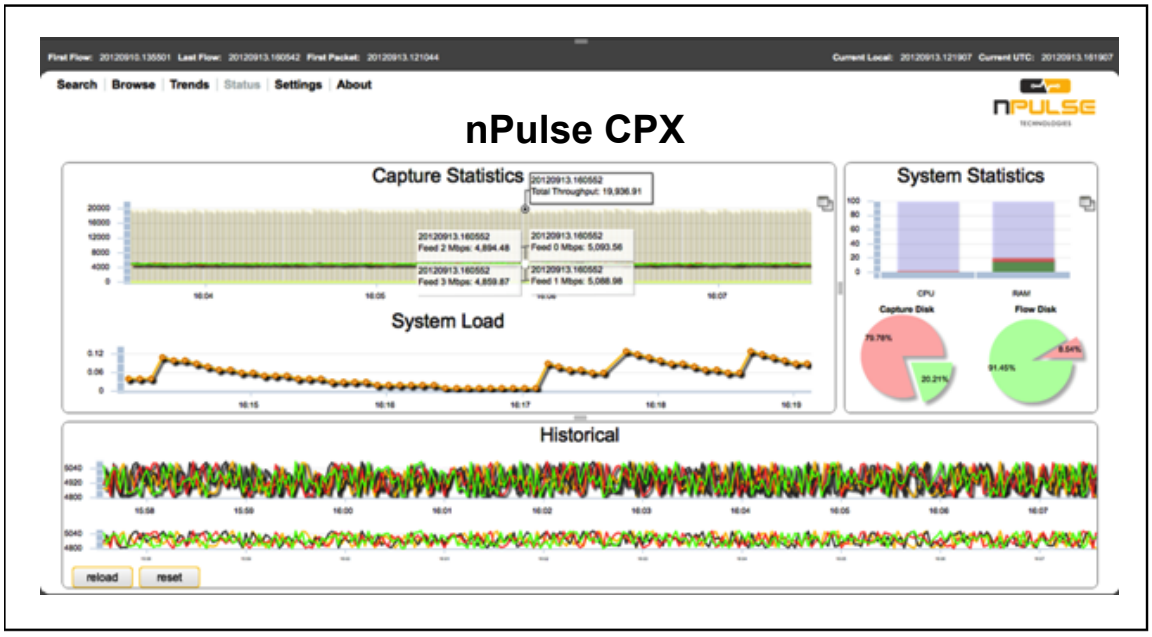

3.3 Quantum StorNext File System

The Quantum StorNext File System dramatically streamlines workflow and improves productivity by creating a shared storage pool, allowing all servers to directly access the data in disk arrays. It offers flexible, high-performance streaming, even with extremely large file sizes (100GB to many terabytes) characteristic of pre- and poststack seismic imaging data:

- Eliminates the need for local file copies.
- Enables transparent data access, high performance, and high availability.
- Provides a policy-based data path that identifies where information should reside.

The StorNext distributed file system provides high-performance, shared access to files stored on disk resources over a switched fabric. By leveraging the SAN to access multiple network destinations, multiple copies can be created and moved through the fabric at the same time from a single master file. It can also give bandwidth priority to specific jobs, which is important for time-critical tasks. This allows concurrent, high-speed access to a common content pool by multiple organizations.

3.4 NetApp E-Series Storage Platform

The NetApp E5400 meets the demanding performance and capacity requirements of HPC environments without sacrificing simplicity and efficiency. The modular design of the E5400 allows you to mix drive types in a single enclosure to address different requirements with the same system. This allows you to create a storage deployment tailored to your seismic processing requirements that will grow with your needs and keep you within your budget.

E5424: Optimized for Performance

The NetApp E5424 delivers both high bandwidth and high IOPS with leading price performance. The E5424 saves cost by consuming 50% less power using up to 24 2.5-inch SAS drives in a 2U form factor. A fully loaded rack delivers performance of more than 30GB/sec sustained write throughput in one 40U.

E5460: Delivers Industry-Leading Density

The NetApp E5460 optimizes storage density for maximum capacity with excellent performance, supporting up to 60 drives in each 4U enclosure. The E5460 supports high-capacity near-line SAS disk options that are superior to SATA drives for high capacity and lower cost per MB/sec and are an excellent choice for throughput-intensive applications. The 4U enclosure holds 60 disk drives in 5 drawers, delivering up to 3.1GB/sec of write throughput in one 40U.

E2624: Metadata Storage

Best practices dictate a dedicated storage system for client access to file metadata (such as name, size, access times, and data locations). NetApp offers the E2624, a cost-effective SAS-connected disk system that can be attached directly to the metadata server, delivering maximum performance.

4 Cyberthreat Analytics and Monitoring

The information collected and stored in security logs can provide a valuable source of intelligence, in addition to the need to maintain the data for compliance reasons. Analytic tools are an excellent way for agencies to gain a better understanding of the threats, attacks, and vulnerabilities. Armed with these valuable insights, agencies are able to strengthen security policies to better safeguard systems and data from future attacks.

- **Intrusion detection.** Incoming network traffic is monitored inline or passively using virtual IDS/IPS sensors.
- **Active threat control.** IDS sensors manage vShield™ rules to actively protect virtual machines.
- **Network forensics.** Captured packets and logs are analyzed in near real time.

Figure 5) Cyberthreat analytics and monitoring.

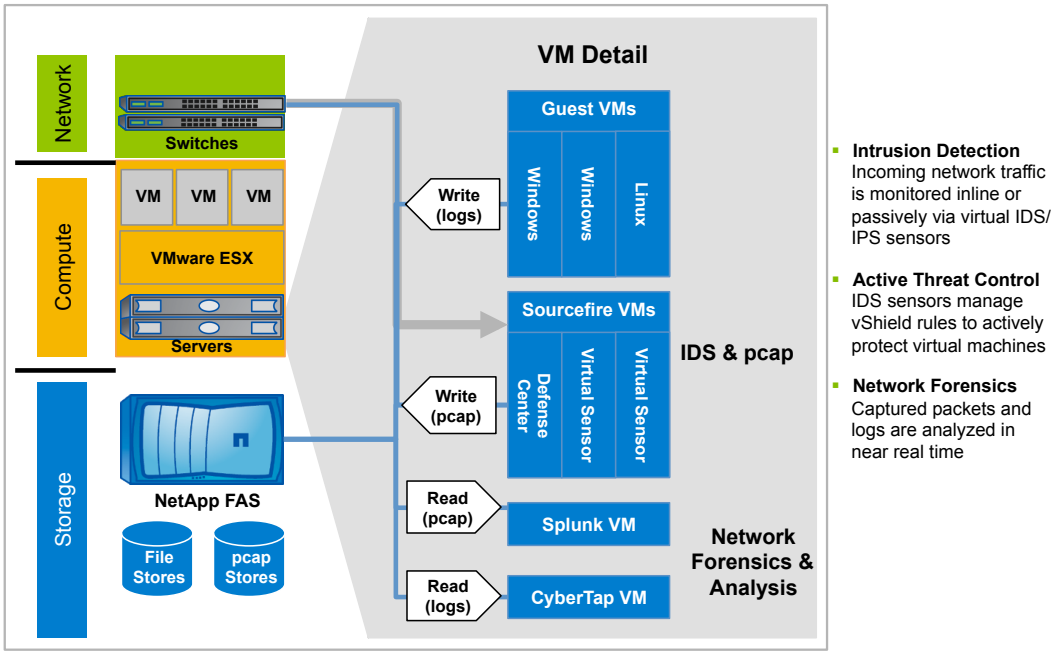

4.1 NetApp for Hadoop

To meet the continually expanding needs of big analytics, including cyber analytics, NetApp has introduced a supporting architecture that is simpler, more efficient, and more reliable. NetApp storage solutions solve the needs of big data, including several solutions designed specifically for complex data analytics:

- **NetApp for Hadoop.** NetApp for Hadoop is an optimized solution for node-storage balance, cost, reliability/serviceability, performance, and storage density. The NetApp Hadoop solution is focused on providing customers with flexible choices driven by deciding factors, such as number of clusters, servers, and support options. Organizations that are using Hadoop often use traditional server-based storage using direct access storage (DAS). The NetApp solution was designed to employ the external DAS model, which provides scalability and competitive total cost of ownership (TCO) at effective prices. The NetApp Hadoop solution provides enterprise-grade storage systems, offering data protection and addressing the single points of failure that plague current Hadoop deployments. Enterprise-grade firmware RAID provides near nonstop transparent RAID operations without the burden of server-based replication—typically, software triple mirroring.

The Apache Hadoop software library framework allows for the distributed processing of large datasets across clusters of computers using a simple programming model. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high availability, the library itself is designed to detect and handle failures at the application layer, delivering a highly available service on top of a cluster of computers.

The Hadoop project includes a number subprojects, including:

- **Hadoop Common.** The common utilities that support the other Hadoop subprojects.
- **Hadoop Distributed File System (HDFS).** A distributed file system that provides high-throughput access to application data.
- **Hadoop MapReduce.** A software framework for easily writing applications that process vast amounts of data (multiterabyte datasets) in parallel on large clusters (thousands of nodes) of commodity hardware in a reliable, fault-tolerant manner.

5 Tactical Collection Platform

5.1 Data ONTAP Edge

NetApp Data ONTAP® Edge is a low-cost remote office storage solution that runs on the VMware vSphere® platform. Designed to complement NetApp FAS and V-Series storage solutions, Data ONTAP Edge allows you to build a data center on an enterprise server using the VMware vSphere platform. Data ONTAP Edge converts the server's internal disk drives into a flexible storage platform, giving you many of the same benefits as a dedicated NetApp storage system.

Data ONTAP Edge includes the NetApp Data ONTAP-v™ operating system, whose native data deduplication and FlexClone®, SnapVault®, and SnapRestore® software leverages your central site's NetApp storage system for data backup, data recovery, and archiving of remote-site data. The result is that your remote offices can participate in your shared IT infrastructure with the same levels of efficiency and flexibility that Data ONTAP provides in your data center.

If you need to recover your branch office data, SnapRestore software uses local Snapshot™ copies to recover entire file systems or data volumes in seconds, regardless of capacity or number of files, using a permanent connection to your central site.

- Deploy consistent storage architecture from edge to core.
- Improve the flexibility of your virtual infrastructure.
- Extend the value of NetApp storage efficiency to server-attached DAS.
- Provide superior data protection and data recovery for remote offices.

Figure 6) Tactical collection platform using Data ONTAP Edge.

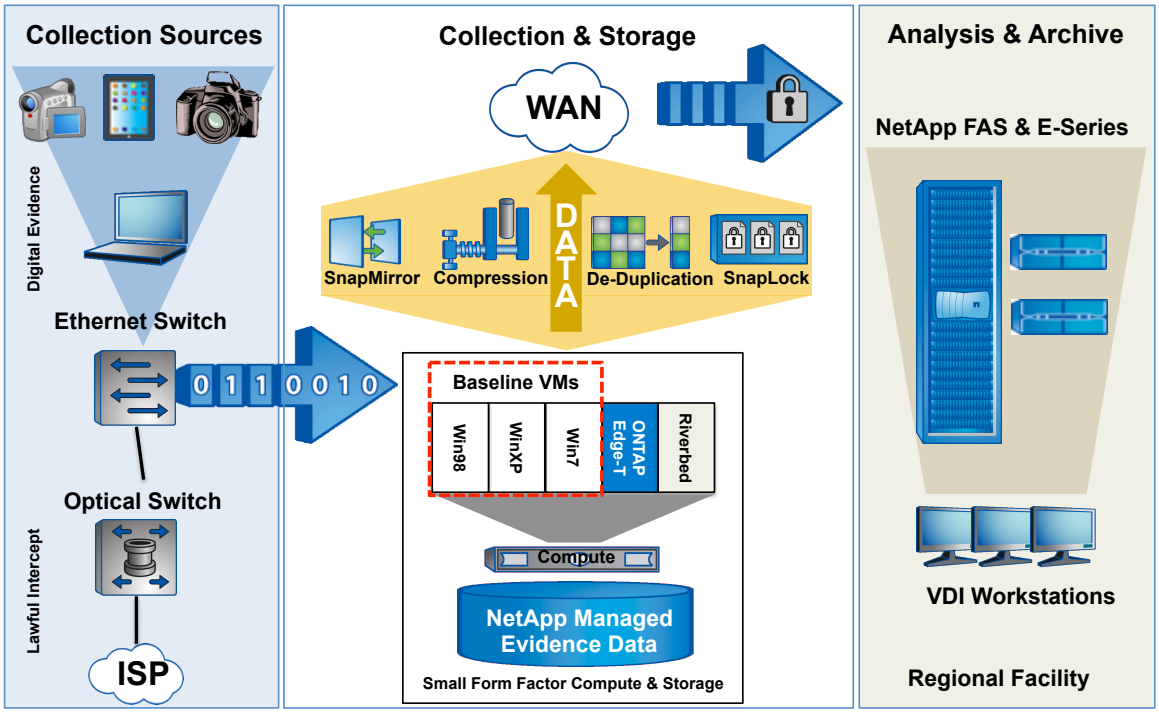

6 Data Recovery

Unfortunately, security breaches are inevitable, given the intricacy of attacks and technology available to hackers. The best way to respond is to have a plan in place before the incident even occurs. This means

establishing an incident response plan for backing up and recovering information, as well as the policies for documenting and reporting the incident required for compliance. NetApp provides the following technologies to our customers to recover from a cyber attack:

- **Snapshot**. NetApp Snapshot copies are built into Data ONTAP and enable our customers to recover their data, down to the file level, from virus, malware, and cyber attacks—instantly.
- **SnapMirror**. SnapMirror® allows our customers to replicate all or part of their environments to another storage system or even another location synchronously or asynchronously. This provides a high level of resiliency when being attacked and allows our customers to continue operating if an entire site is compromised.
- **Clustered Data ONTAP.** Clustered Data ONTAP provides the highest level of resiliency to our customers. In the event of a compromise, our customers can nondisruptively move data within the cluster, making sure they stay operational under the most difficult conditions.

7 Secure Multi-Tenancy

Security is always a challenge, regardless of whether you are deploying a cloud instance or consolidating your existing environment. NetApp provides an effective way for customers to host multiple tenants on the same architecture—securely. NetApp MultiStore® enables our customers to create multiple virtual storage systems within a single physical storage system. When deployed within a FlexPod architecture, MultiStore enables our customers to securely dedicate resources to different tenants across the stack.

8 Encryption

Among numerous security mechanisms, the most effective method for protecting any confidential information is data encryption. If the agency's first line of defense is penetrated, data encryption will make sure that confidential data can't be viewed. NetApp offers several options for encryption across all of our storage platforms:

- **NetApp Storage Encryption (NSE).** NetApp offers encrypted drives with up to AES-256 bit encryption and centralized key management.
- **NAS encryption**. NetApp partners with SafeNet to provide centralized key management and up to AES-256 bit encryption for data in flight using NAS protocols.
- **SAN encryption**. NetApp partners with Brocade to provide up to AES-256 bit encryption for data using SAN-based protocols.

9 Summary

NetApp takes a comprehensive approach to cybersecurity by providing technology and solutions that make sure data is protected and trusted, while still remaining accessible. With the creation of the NetApp cybersecurity reference architecture, government agencies can deploy solutions that include best-in-class partner technologies and allow them to collect, analyze, and distribute the volumes of cybersecurity data in their environments across high-speed networks.

NetApp storage systems and tactical collection platforms provide a cost-effective solution to collect and analyze threat data, with the encryption and secure multi-tenancy capabilities to help agencies protect their data from threats. And with integrated data protection features, NetApp business continuity solutions help maintain availability across a broad spectrum of recovery point and recovery time requirements during planned as well as unplanned downtime. Collectively, these solutions enable government agencies to achieve continuous availability and protect critical data.

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com