PANASAS®

SOLUTION BRIEF: **SECURITY**

Safeguarding HPC Storage and Customer Data with PanFS

PanFS High Performance Parallel File System

Now in its 9th generation, PanFS® — the Panasas parallel file system — delivers the highest performance among competitive HPC storage systems at any capacity, and takes the complexity and unreliability of typical high-performance computing (HPC) storage systems off your hands, and it does so using commodity hardware at competitive price points.

PanFS orchestrates multiple computers into a single entity that serves your data to your compute cluster. Through sophisticated software, multiple computers that each have HDDs and/or SSDs attached to them will work together to provide hundreds of Gigabytes per second (GB/s) of data being read and written by your HPC applications.

PanFS manages this orchestration without manual intervention, automatically recovering from any failures, continuously balancing the load across those computers, scrubbing the stored data for the highest levels of data protection, and encrypting the stored data to protect it from unwanted exposure.

To safeguard the HPC storage and customer data, PanFS supports features that prevent unauthorized data access while the parallel file system is online, access control lists (ACLs) and SELinux, and one while the system is offline, hardware-based encryption at rest.

Data is the Most Valuable Asset

High-performance computing (HPC) systems aggregate resources and use parallel processing and parallel file systems like PanFS to run complex and intensive applications quickly and reliably. These HPC systems are critical in commercial, research, and government organizations, but the data they hold could be the most valuable asset an organization has, e.g., in 2020, United Airlines put a $22B valuation of their customer data.

From industry to university, research and government labs, HPC systems are critical to the productivity of nations. Along with enterprise data centers, cybersecurity is essential in HPC environments but HPC systems are under attack not by run-of-the-mill hackers, rather sometimes by nation-state actors with destructive intent. In fact, in Dec. of 2020, the US Department of Energy uncovered evidence that hackers had breached the National Nuclear Security Administration with activities at both the Sandia and Los Alamos National Laboratories.

Today's HPC environments can have tens of thousand of nodes with combinations of CPUs and GPUs; scientific computing, CAD/CAM and engineering simultations, data analytics, and AI/ML applications; and high-performance, scalable HPC storage parallel file systems like PanFS for quenching their enormous and exponentially growing data appetite.

Security in High-Performance Computing Environments

Dr. Sean Peisert who leads computer security R&D at Lawrence Berkeley National Lab published in the Sept 2017 Communication of the ACM an article titled "Security in High-Performance Computing Environments." In it, he wrote about the "thousands of probes, scans, brute-force login attempts, and buffer overflow vulnerabilities that continue to plague high-performance computing facilities today."

He also noted that while HPC systems have similarities to traditional IT systems, two different challenges are the "paramount priority of high-performance in HPC" and "the need to make some HPC environments as open as possible to enable broad scientific collaboration."

As data is more and more becoming the most valuable treasure and competitive advantage a company has, and to meet the challenges of performance and access that a secure HPC system requires, Panasas is working with security ecosystem partners in a multi-layer security framework.

Multi-Layer Security (Defense in Depth)

Likely coined in the National Security Agency (NSA) document of the same name, Defense in Depth is a best practices approach for protecting information and highly networked information systems from adversarial attacks as well as the effects from non-malicious events. The original concept was achieving Defense in Depth through a balanced focus on people, technology and operations.

In evaluating technology, the NSA detailed a layered defense approach, e.g., a first line followed by a second line of defense, with each layer or technology providing a unique obstacle to the adversary. This layered defense approach is thought to have given rise to a multi-layer security model where each layer provides a different defense mechanism, but all layers work together building on each other and providing effective and efficient security coverage for a networked environment.

One model is the seven-layer approach of policies, physical, perimeter, network, host, application, and data layers.

Policies Layer At the policies layer, security measures are certifications, compliances, audits, rules and procedures, and data handling policies including those based on file attributes.

Physical Layer Physical layer security is about the location and property, environmental and entry controls, surveillance, guards, and enforcement of authorization and authentication.

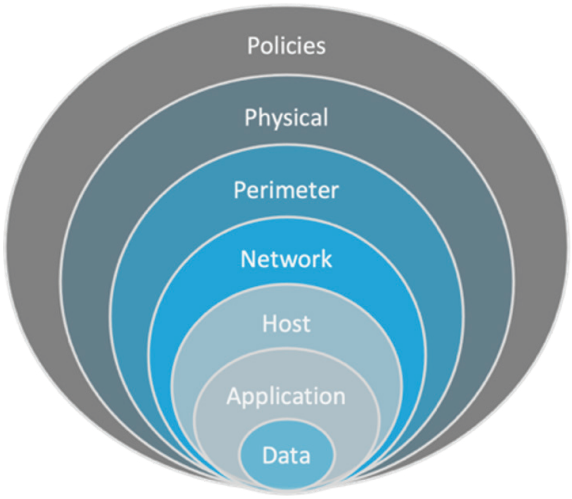

Figure 1: Multi-Layer Security Model.

Perimeter Layer At the perimeter of the network, security focuses on protecting the network boundary including with trusted and untrusted zones, firewalls, and network and filesystem ACLs.

Network Layer Network layer security includes intrusion detection/prevention systems (IDS/IPS), deep packet inspection, and security information event management (SIEM) solutions.

Host Layer Security Security measures at this layer focus on the servers, virtual machines, containers, operating systems, anti-virus/anti-malware software, and logins.

Application Layer Here focus is on securing the application during development, test, and security patch stages and securing single sign-on access, authentication, and monitoring.

Data Layer At the data layer, security measures include encryption with hardware-based disk encryption, ACLs, and content-based security such as SELinux security labels.

PanFS security features that secure the HPC storage and safeguard the vital data held there are delivered across three keys layers:

1. The policy layer with SELinux security labels to enable SELinux and Multi-Level Security (MLS) policies,
2. The perimeter layer through granular filesystem ACLs, and

3. The data layer with hardware-based encryption-at-rest.

SELinux Security Labels

Originally introduced to the Linux community as a development project from the National Security Agency (NSA) and others in 2000, SELinux was integrated upstream into the 2.6x kernel as a Linux Security Module (LSM). Much of the work was jointly driven by the NSA, Red Hat, and the SELinux development community.

What SELinux provides is a kernel-implemented set of mandatary access controls that confine user programs and system services to the minimum level of file and data access that they require. To enable that, filesystems must store and provide access to a security label stored in an extended attribute called "security.selinux" that is associated with all files and directories in the namespace.

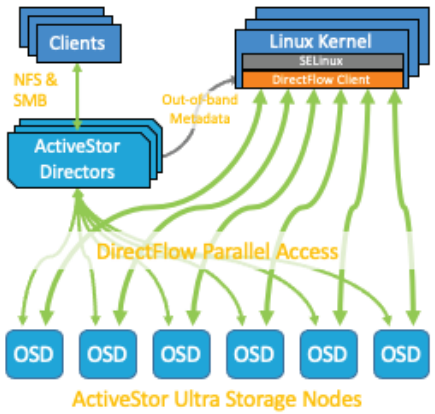

Figure 2: PanFS DirectFlow Client and SELinux Policy Engine Integration.

PanFS efficiently treats security.selinux as an additional inode field tightly integrated into the filesystem, eliminating the round trips and potential sync issues of treating the label as a traditional Linux extended attribute. By integrating the Panasas DirectFlow client with the SELinux policy engine inside the Linux kernel and providing that policy engine low-latency access to the security.selinux extended attribute, Panasas can now support high-performance workflows that rely on this capability and enable improved security at the Policies Layer.

By design the storage server does not enforce the security, rather it is enforced through security policies implemented on the clients. There are multiple SELinux policies or sets of rules for the SELinux security engine including:

i. Targeted, the most common type used,
ii. Minimum, a stripped down version of Targeted,
iii. Multi-Category Security (MCS), an enhanced version allowing users to label files as categories, and
iv. Multi-Level Security (MLS), the strongest version of Targeted, sometimes used by governments.

When compared to having to deploy completely separate "siloed" compute and storage systems for each project, users and data in different security levels and compartments can share the same compute and storage resources that help to

- Control hardware and staffing costs,
- Reduce security maintenance costs and licensing fees, and
- Ensure efficient system access for all the users.

Filesystem ACLs

There are two primary types of ACLs:

- Filesystem ACLs that manage permissions within a file system, and
- Networking ACLs that manage access to the network.

Filesystem ACLs are tables that tell the operating system the privileges users and groups should be give, granting or denying them read, write and/or execution permissions on files or directories beyond the traditional Linux permissions set of mode bits such as -rwxr-xr-x. ACLs are actually a list of Access Control Entries (ACEs) that each describe a finer grained control over which user accounts can perform which operations on each file or directory than the mode bits do.

ACLs are good at separating groups of people and setting appropriate permissions for each of those groups. For example, if an organization has an AI/ML research group and a life sciences research group, it is straightforward to

- Grant full read/write permissions to anyone in the AI/ML group for any data in the AI/ML group's directory tree,
- But only grant them read-only access to the data in the life sciences group's directory tree, and
- Then to deny both groups any access to the finance group's directory tree.

It is quite easy to set up those types of user groups in both Windows ActiveDirectory and Open-Source LDAP solutions, and then to create the ACLs in storage to implement the above types of policies.

Hardware-based Encryption-at-Rest

Panasas uses industry-standard self-encrypting drives (SEDs) in the ActiveStor Ultra storage appliance. The hardware-based encryption algorithms built into the drives are designed so that encryption does not reduce performance.

With PanFS, NIST-approved AES-256 encryption-at-rest of the data being held can be easily enabled and controlled. The Panasas encryption-at-rest solution also supports complete cryptographic erasure of a drive upon command, enabling secure repurposing of the storage system without risking exposure of the data last on the system.

Figure 3: Self-Encrypting Drive (SED) with 256-bit AES Encryption Controller.

Additionally, PanFS supports the Key Management Interoperability Protocol (KMIP), allowing customers to use well-established and proven cyber security key management solutions such as CipherTrust Manager from Thales Trusted Cyber Technologies (TCT). PanFS integrates with CipherTrust Manager to mitigate the threat of authorized access to encrypted data while providing centralized and simplified key management, e.g., key generation, escrow, and recovery.

To show what might look somewhat complex is actually not, following is a walk through of an example configuration. To enable encryption-at-rest on a PanFS realm,

1. A license for a Thales TCT CipherTrust Manager instance that's going to securely store and manage the keys and act as a KMIP server for PanFS to talk to will first need to be obtained - it's available as a VMware guest instance or as an extra-secure physical platform.
2. A high-availability configuration will be needed so that in the event of a failure, the keys will still be available while the failed unit is fixed.
3. Next, the PanFS realm will need to be told about the KMIP server; primarily its IP address and a shared SSH key so the PanFS realm can connect via KMIP to the CipherTrust Manager server.
4. The final step is to tell PanFS to enable encryption-at-rest.

After that, PanFS will connect to the KMIP server, run a small test to verify that everything has been configured correctly, ask the KMIP server to generate high-entropy keys for each of the HDDs and SSDs, and then apply those keys to each of the drives. During normal runtime, PanFS will periodically verify that the KMIP server is alive and well and contains all the right keys, and will raise an alert if there is any type of problem.

As shown, PanFS makes hardware-based encryption-at-rest an easy and secure solution to worries about data breaches resulting from drives escaping an organization's control with readable data on them.

Safeguarding HPC Storage and Customer Data

Safeguarding HPC storage and customer data with zero performance impact is a pillar of the Panasas PanFS architecture. Through an extensive and efficient multi-layer security approach and with our security ecosystem partners, PanFS release 9 delivers on data security with SELinux per-file security label support, fine-grained filesystem ACLs, and fully managed hardware-based encryption-at-rest. And with data security becoming an ever more important characteristic of storage systems, Panasas will continue developing solutions that protect your data from unauthorized access.