Qumulo Security Architecture and Practices

White Paper

October 2024
Version 1.3

James Walkenhorst
Sr. Technical Marketing Engineer
Qumulo

Abstract

This white paper describes the security controls and features of Qumulo's storage architecture. It also suggests effective measures that can be taken in any Qumulo deployment to minimize the risk of corporate data loss due to malware or other forms of cyberattack.

Introduction

Malware outbreaks. Phishing scams. Brute-force attacks. Not only are there more of them than ever before, but the level of sophistication and the scale of impact are increasing as well. Today's enterprises are challenged like never before to protect their systems and data

Executive summary

Every organization's CISO knows that a malware attack on their watch is a question of when, not if. A quick Google search of "recent ransomware outbreaks" brings up a sobering list of both large and small enterprises whose defenses were breached, and who suffered a varied list of consequences: lost data, stolen data, loss of revenue, ransomware payments.

Depending on the scope of the attack, the effects may impact more than just the company's own bottom line. The theft of employee or customer data – including sensitive financial or medical records – creates downstream risk and consequences beyond the company's balance sheet. Even worse, in the event of ransomware attacks on hospitals, medical treatments can be delayed or denied.

Beyond the loss of revenue and business, in some industries the failure to secure critical business systems and data may incur additional regulatory and legal penalties: fines, settlements, and even criminal liability. As if that weren't enough, the reputational impact for any storage administrator, IT director, and CSO could be long-lasting or even permanent.

Ransomware attacks have risen by 13 percent over the last five years. Additionally, for companies willing to pay a ransom in order to recover lost data, the average amount paid has increased from US$1.85 million per incident in 2023 to US$2.73 million in 2024[1].

Not all the news is bad, and the fact remains that there are effective countermeasures for most attack vectors. Even as the number of attacks has continued to rise, 97% of all impacted organizations were able to recover lost data.

The reality of that 97% figure is that most enterprises have implemented a robust security framework that both reduces the likelihood of ransomware attacks, and which ensures that critical data can be restored if the primary copy is lost. It also shows that a multi-layered security framework that recognizes where systems and data are vulnerable and applies the necessary security features, products, and practices to address those vulnerabilities can be the most cost-effective approach.

[1] Sobers, Rob. "Ransomware Statistics, Data, Trends, and Facts." Varonis. September 13, 2024. https://www.varonis.com/blog/ransomware-statistics.

Securing Qumulo systems and data

Enterprises who trust their unstructured data to Qumulo storage can leverage its built-in security controls and data-protection capabilities to safeguard against malware attacks.

The Qumulo operating environment was engineered to simplify data protection and security using a multi-layered security model that minimizes the risk of intrusion, provides tools to quickly detect and contain security breaches, and effective data-protection controls to ensure that data recovery, if necessary, be achieved quickly.

With a broad spectrum of security controls that address risks at every level, Qumulo simplifies data security, streamlines recovery, and delivers peace of mind for enterprise administrators and executives alike.

Document purpose

Qumulo provides ample storage-centric security controls that can even meet the stringent requirements of regulated enterprises. This document is intended to describe Qumulo's inherent security features as well as its ability to integrate with enterprise-class monitoring and security infrastructure to protect unstructured data against both internal and external security threats.

Audience

This document is intended for system administrators, IT management, and enterprise information and security executives. Familiarity with enterprise security concepts and practices in managing file data, storage and network platforms in an enterprise environment is assumed within the scope of this white paper.

For deeper investigation and external research, footnotes are included to external references within the body of the paper. Links to internal Qumulo resources, such as knowledge base articles and Qumulo administrative documentation, are provided in an Appendix section at the end of this document.

Table of Contents

Data and system vulnerabilities

Even in a well-protected enterprise environment, the security of critical systems and data is not guaranteed. Access badges and security guards can keep most bad actors out, but not all. Firewalls provide a high degree of protection from the outside world, and while a well-managed firewall can repel more than 99% of malicious access attempts, no security mechanism is 100% foolproof. And when even small enterprises experience hundreds to thousands of cyberattacks per day, even the <1% gap between 99%+ effective and 100% effective means a significant risk.

No matter how well implemented and managed, any one security layer can still be defeated by an opportunistic attacker or can be bypassed by a careless employee who opens the wrong email attachment. Even the tightest system security and role-based access controls won't stop a rogue administrator looking to cause harm.

When it comes to security, it is crucial to recognize that each layer is a point of vulnerability that risks exposing systems and data to attack. It is designed with layered security measures that, when combined, ensure that a breach at any one point in the environment still leaves additional barriers to overcome.

Multi-layered controls and practices

The Qumulo security model is based on this same approach. Rather than being a specific product feature, it's a matrix of security controls, features, and practices that add security at each point of vulnerability – system, network, users, and data.

Recognizing that no security systems or measures are guaranteed to be 100% effective, Qumulo's security approach has been engineered with three primary objectives:

1. Minimize the likelihood of a security intrusion
2. Detect security breaches in real time and provide tools for administrative notification and response, in ways that minimize any adverse effects
3. Provide a combination of tools, features, and mechanisms to ensure the rapid recoverability of any affected data and services

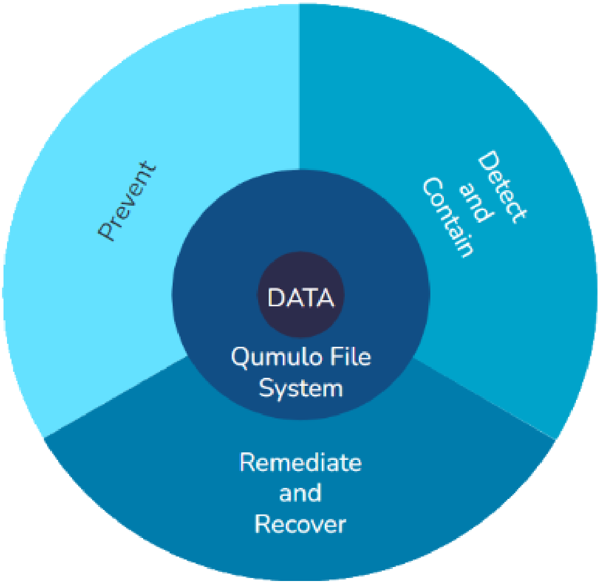

Figure 1: Qumulo's holistic security model

The overarching goal of an effective security strategy is to insert as many security layers as necessary – at every identifiable point of vulnerability – between the threat and the data. Where possible, Qumulo's simplicity-first approach has been used: innate encryption at rest; integration with standard enterprise security tools, and well-established industry processes like access control list-based data security, as well as industry-standard encryption methods for data in transit.

Part 1: Preventing Intrusion

The most critical objective – preventing attacks from reaching their intended target – is achieved through a combination of system design, inherent security controls (including security hardening settings and features in each new Qumulo software release), configurable security options, and recommended security practices. When effectively planned and implemented, these features and practices minimize the risk of a security breach against either the Qumulo storage environment itself or the data on it.

Secure system architecture

The Qumulo software stack was engineered from the very beginning for the tightest possible system security. Many of these measures are inherent in the operating environment itself – automatically made available to all Qumulo customers in all deployments.

Linux environment adaptations

While Qumulo's software runs on standard, enterprise-grade hardware paired with an Ubuntu Long-Term Support release, the underlying Linux operating system is locked down, allowing only the operations needed to perform the required supporting tasks of the Qumulo software stack. Other standard Linux services have been disabled to further reduce the risk surface for an attack.

Fully engineered software stack

Although Linux includes open-source components to provide both NFS and SMB client and server services (e.g. Samba, Ganesha, etc.), these services are not included in the hardened Ubuntu image that supports the Qumulo software stack. Qumulo develops and controls all code used for data-access protocols NFS, SMB, FTP, and S3 – in the Qumulo operating environment. Any risks that might arise from known or discovered vulnerabilities in third-party components are effectively neutralized by Qumulo's control of its proprietary code.

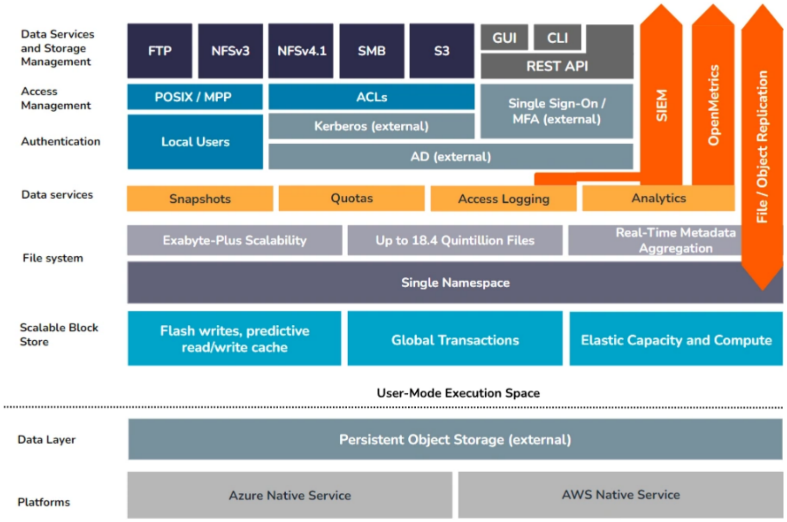

Figure 2: Qumulo software in on-premises hardware environment

User-space execution

Not only is Qumulo software developed in accordance with secure coding best practices – further reducing the potential attack surface and the risk of exploitable vulnerabilities – but it is also completely contained within the user-space of the underlying Linux operating system. If a Linux vulnerability is discovered and exploited, even if the attacker were able to seize user privileges on the underlying appliance, they would still not be able to access Qumulo's proprietary management controls or file-system data.

Level of separation from the operating system

There is no possible pathway or means for sharing users or privileges in the underlying operating system with users and privileges within the Qumulo operating environment: any user accounts in the base Linux image are not recognized by the Qumulo platform, whose user base is typically maintained either in Active Directory or in a local database within the cluster. This is a different approach from other storage vendors, where an admin- or root-level user with local access to any of the storage nodes or controllers can easily access and manipulate all data, including within the file system or on any local volume.

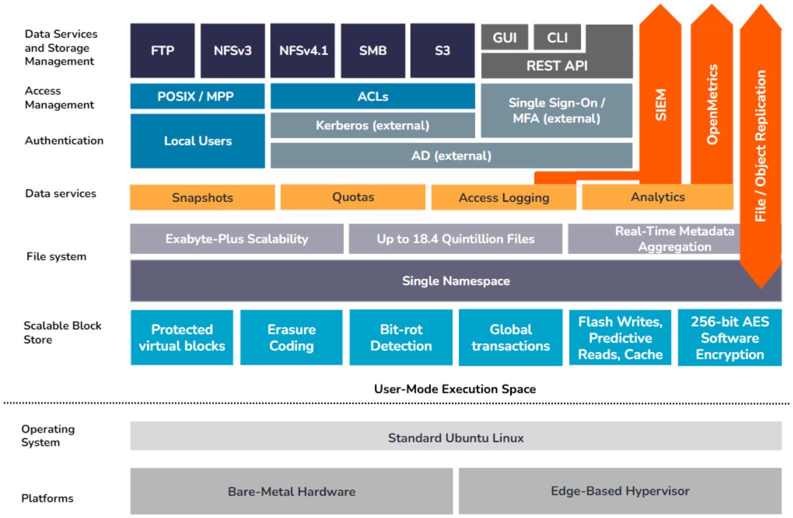

Figure 3: Qumulo software environment on cloud-based hypervisor

Instant upgrades

Qumulo's iterative development process is simple and streamlined, with new software updates released regularly. Not only does this enable rapid innovation to develop and roll out new features, but it also fosters a more secure storage platform. Once a threat has been identified its overall risk profile is assessed using the Common Vulnerabilities and Exploits Scoring System (CVSS)[2].

Each threat is given an overall assessment rating: Critical, High, or Medium, based on its potential impact. Critical and High threats are given a higher priority for remediation. Patches for these threats are developed, tested and according to policy-driven timelines, and are released outside the standard Qumulo software update cycle. Any security vulnerabilities that are assessed at a Medium level (minimal risk, no exploitability, etc.) are remediated as part of the normal release cycle.

Qumulo's container-based architecture enables a unique upgrade process that minimizes disruption to users and workflows. On a rolling, node-by-node basis, the new operating software is deployed in a parallel container to the old version. Once the new instance has initialized, the old environment is gracefully shut down and the upgrade proceeds to the next node until the entire cluster has been upgraded.

This process compares favorably to that of other storage platforms, who may release minor updates every few months and major updates annually, leading to less-frequent opportunities to close any new security holes.

As for in-the-wild exploits, implementing them on other storage systems can be a challenge when the installed storage firmware is older than the earliest compatible version for a given patch. Enterprises may find themselves having to choose whether to risk an unplanned outage from a possible failed update or risk leaving an unsecured system open to a successful attack vector.

Simple equals secure

While the motivation behind the instant-upgrade process was to minimize service impacts, the net effect is to increase the security of the system as well: customers are much more likely to deploy the latest firmware version, complete with the latest security fixes, if they can do so with minimal disruption to end users and services.

For more information, please refer to the Qumulo instant upgrades section of the Appendix at the end of this document.

[2] For more information on CVSS metrics and calculations, please visit the NIST CVSS page at https://nvd.nist.gov/vuln-metrics/cvss.

Protecting and securing Qumulo systems

In addition to its built-in security features, customers can use a number of configurable means for securing and protecting their Qumulo system from unauthorized or unwanted access.

Administrative security

Since full cluster-management privileges are granted to administrative users, the security of the cluster itself, as well as all data on the cluster, are dependent on the security of the user accounts entrusted with admin rights to the cluster.

Qumulo recommends administrators apply the "principle of least privilege[3]" for their Qumulo clusters. This section outlines available options as well as other recommended practices for maintaining a high level of security.

Administrative accounts

As with many other systems and platforms, cluster rights and privileges are granted based on membership in one or more local groups on the cluster. Administrative rights are granted to all local and domain accounts that are members of the cluster's built-in Administrators group.

Domain-level administrative users

Most enterprise security policies require that the administration and management of critical enterprise systems follow a one-user, one-account policy to ensure accurate records of system access and privilege use. The simplest method for complying with this policy is by adding the relevant Active Directory user accounts to the cluster's local Administrators group.

Local admin user

Every Qumulo cluster comes with a default account, called admin, which is automatically assigned membership in the local Administrators group, and as such has full administrative rights and privileges to the cluster.

The admin account's initial password is assigned as part of the cluster-creation process. Care should be taken to ensure that the password is well documented and is changed and secured in accordance with the existing enterprise security policy.

[3] According to NIST AC-6 standards. More information on "principle of least privilege" is available at https://csf.tools/reference/nist-sp-800-53/r5/ac/ac-6/.

Emergency use of the local admin user account

In the event of widespread facilities failure, in which Active Directory is not available to authenticate domain-level administrative accounts, the local admin account can be used to gain cluster access.

Single sign-on with multi-factor authentication

Without additional security measures, simply authenticating to the storage cluster with an administrative user's credentials grants full access to the system and the data on it. Since a single admin-level user often has admin-level access to multiple enterprise systems, a single compromised administrator-level account can pose a serious risk to systems and data across the entire enterprise.

Single sign-on (SSO) eliminates the need for an administrator to re-enter their login credentials to gain access to the system. Enterprises want SSO not just because it streamlines the login process, making it more convenient for admins to authenticate, but also because it reduces the risk of account theft via keystroke loggers or interception as the login attempt traverses the network.

Multiple-factor authentication (MFA) adds another layer of security to the login process, requiring that admin users retrieve a one-time code from either a key token or a challenge request on a separate device, neither of which would be in the possession of an intruder. Many enterprise customers have internal security policies that require MFA for admin-level accounts to log in to critical information technology systems.

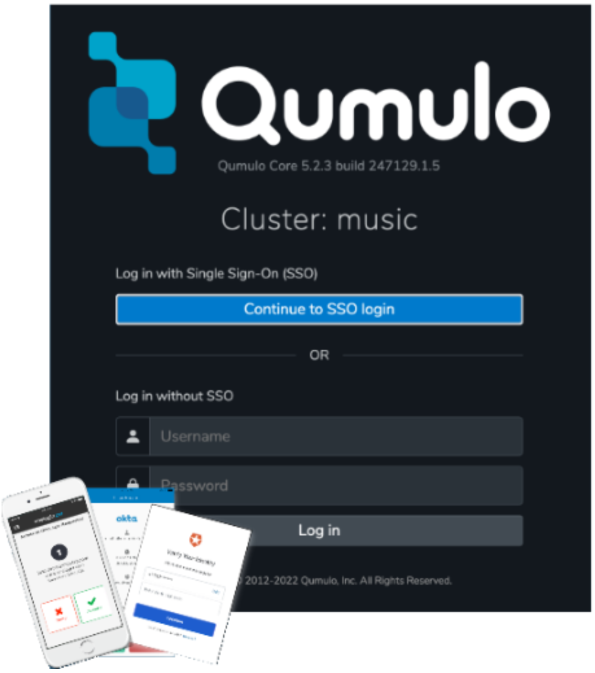
Qumulo
Qumulo Core 5.2.3 build 247129.1.5
Cluster: music
Log in with Single Sign-On (SSO)
Continue to SSO login
OR
Log in without SSO
Username
Password
Log in
© 2012-2022 Qumulo, Inc. All Rights Reserved.
okta

Figure 4: Single Sign-On with Multifactor Authentication

Qumulo's SSO solution integrates with Active Directory via Security Assertion Markup Language (SAML) 2.0. For MFA, customers can leverage any Identity Provider (IdP) that integrates with the AD domain registered on the cluster, including but not limited to OneLogin, Okta, Duo, and Azure AD.

For more information about configuring SAML SSO with MFA on a Qumulo cluster, please see the Single sign-on with multi-factor authentication section in the Appendix at the end of this document.

Access tokens

With Qumulo's API-first development and management model, all features and functionality are available through the HTTP REST API as well as via CLI and through the web user interface. While this enables enterprises to automate any or all storage and data-management functions, and opens the door to the use of third-party tools to expand its capabilities, it also raises two issues:

1. With the use of standard authentication methods, every automated workflow that leverages a Qumulo API needs to re-authenticate after 10 hours, whether active or not.
2. Having to pass user credentials again for every session renewal undermines the security of the system, since these credentials must be persisted and re-used with each login.

To resolve these issues, enterprises can generate a long-lived API token that can be used by automated workflows indefinitely, until the key is either revoked or deleted. The token is generated by an administrator via CLI, and can be attached to each API-based workflow, which can now make authenticated API calls without having to log in.

For auditing purposes, each token maps to a specific AD or cluster account. If the associated user account is deleted or deactivated, the access token will stop functioning.

More information about configuring and using access tokens to secure automated workflows is available in the Access tokens section of the Appendix at the bottom of this document.

Role-Based Access Control

Role-based access control (RBAC) allows administrators to assign fine-grained privileges to non-administrative users or groups who require elevated rights to the cluster for specific management tasks. The use of the RBAC model allows the secure delegation of privileges on an as-needed basis without needing to confer full administrative rights.

In addition to the built-in **Administrators** group, who have full access to and control of the Qumulo cluster, there are currently two more predefined roles:

- **Data Administrators**: The Data Administrators role is ideal for API/CLI user accounts. With this role, a user or group will not have access to the Web UI but will have the same file privileges as the Administrators role along with some others
- **Observers**: With the Observers role, a user or group will have the privilege to access the Web UI and read-only APIs with a few exceptions (debug APIs and authentication settings).

As with the local **Administrators**, these groups can be populated with the appropriate Active Directory user accounts to grant the necessary privileges while ensuring a verifiable audit trail of access and privilege use.

For more information about role-based access control concepts and settings, please see the Role-Based Access Control section in the Appendix at the end of this document.

Management traffic restrictions

Requiring administrative users to use their Active Directory accounts and authenticate via SSO with MFA eliminates much of the risk of access from a compromised administrator account. Some organizations, however, have additional security policies that require that admin access for enterprise systems be restricted to one or more specific networks, or VLANs.

By offering the ability to block specific TCP ports at an individual VLAN level, Qumulo allows for the segmentation of management traffic – e.g. API, SSH, and web UI, and replication – from client traffic – e.g. SMB and NFS.

Management-only VLANs help ensure that, even if an intruder were able to bypass the other available security measures from any other network location, then their attempt to gain unauthorized access to the cluster would still fail.

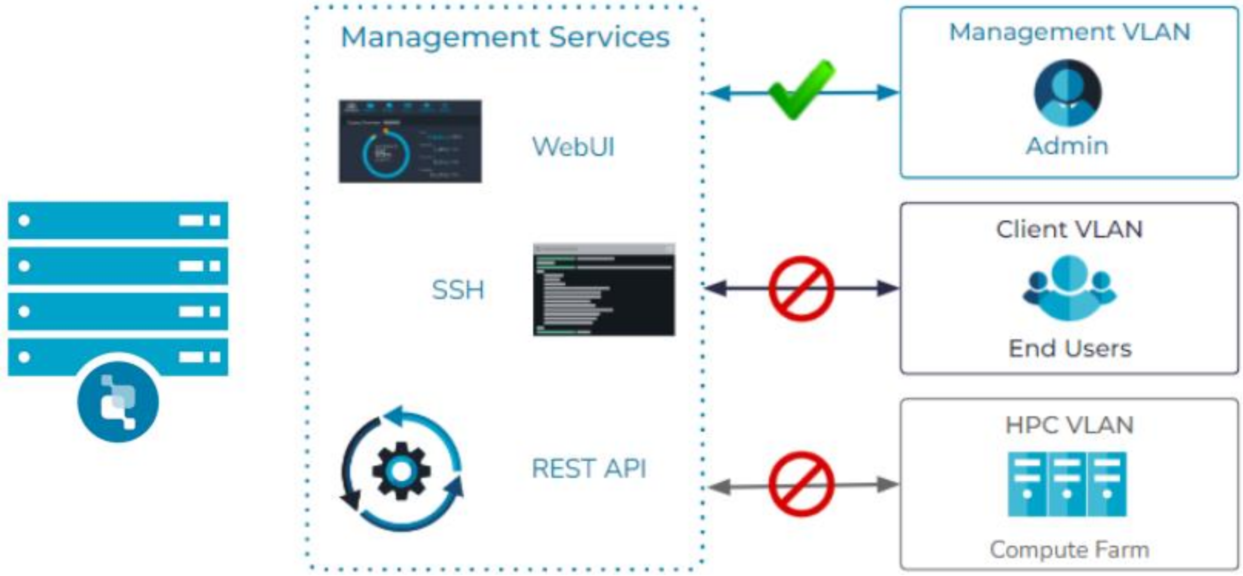

Figure 5: Management VLAN

For more information on the creation of management VLANs, please refer to the Management VLANs section of the Appendix at the end of this document.

Protecting and securing Qumulo data

Besides providing a series of security features for the purpose of safeguarding the cluster itself from unauthorized access, Qumulo's software incorporates several built-in features and configurable controls, all designed to protect the data on the cluster. This section describes these protective features.

Qumulo data-security features

Inherent in every Qumulo system, and implemented by default, are a pair of controls intended to ensure the security of all data from corruption, loss, or intrusion at the level of the media on which the data is written.

Software-based data encryption at rest

The first of these features ensures that all data on an on-premises Qumulo cluster is automatically encrypted as it's written to disk using an AES 256-bit compliant algorithm, ensuring that all data on a Qumulo system is secured against bad actors even if they were able to gain physical access to the disk itself.

For on-premises deployments, software encryption is part of the file system stack. The encryption algorithm initializes as part of the initial cluster build process and encompasses all file system data and metadata at the block level.

Keys are used to encrypt data, as well as to encrypt data keys themselves. A master key is utilized and stored on every boot drive in the cluster, in a file that only root can access, adding an extra layer of security.

Qumulo clusters in the cloud rely on block-level encryption within the cloud-storage layer, thereby ensuring that all at-rest data on any Qumulo instance is fully encrypted.

FIPS 140-2 compliant encryption

While data encryption at rest is a standard component of nearly every enterprise platform, not all encryption algorithms are engineered to the same standard. Many enterprises, including government agencies and customers in some regulated industries require compliance with Federal Information Processing Standards (FIPS) as a core component of their security policies.

Qumulo's software-based encryption module is certified as compliant with FIPS 140-2 requirements. For enterprise customers that require FIPS-compliant data services, the Qumulo security module that includes at-rest data encryption is bundled and versioned separately from the rest of the software stack. This will allow these customers to upgrade their Qumulo firmware separately from the security module and maintain their FIPS-compliant status.

For more information about Qumulo's at-rest data encryption and FIPS compliance, please see the Software-based encryption section in the Appendix at the end of this document.

If FIPS-compliant data services are required for cloud-based Qumulo deployments, please refer to the cloud vendor's specific statements regarding their FIPS status.

Active Directory integration

Beyond the need for restricting system access to only authorized accounts and securing data at the disk level via encryption, the next layer of data protection requires a secure directory of user accounts from which storage and data access rights and permissions can be managed.

Qumulo software was engineered to leverage Microsoft Active Directory (AD) for both administrative and user rights and permissions. AD accounts can be configured for both cluster management and client access. This model conveys a number of advantages over a local-only approach, including:

- A single source of record for all user accounts
- Separation of user accounts and identities from both the storage and hardware environments
- Seamless integration with both SMB and NFSv4.1 Access Control Lists (ACLs)
- Integration with Kerberos-based authentication and identity management protocols for all system and data access requests
- Integration with SSO and MFA access providers

While a Qumulo cluster starts with a single local account for initial configuration tasks, and supports the use of local accounts generally, the use of AD accounts is recommended for enabling access to both administrative and client functions.

Data security recommendations and practices

When it comes to making shared file data available to only the right users and groups, Qumulo's multi-layered model enables administrators to apply network security based on the client's network location. User and group access permissions can be set at the level of the share/export itself, as well as directly to individual directories and files.

For all Qumulo's built-in security controls, it still falls on enterprise administrators to configure their systems and data according to industry standards, Qumulo best practices, and their own internal security policies.

This section outlines the recommended configuration settings and practices for minimizing the risk of intrusion, unauthorized data access and data corruption.

Multi-tenant networking

Network multi-tenancy can be implemented on any on-premises Qumulo cluster running software version 5.3.4 or later. It leverages the same partitioning approach as Qumulo's ability to isolate management and client traffic based on VLAN.

Enterprises can use network multi-tenancy to consolidate multiple business units to a single Qumulo cluster, reducing cost and simplifying management without compromising security.

Tenant components

Once multi-tenancy has been enabled on a Qumulo cluster, individual tenants – each of which has a unique name and is assigned one or more member networks, or VLANs – can be created. While a single tenant may comprise multiple VLANs, an individual VLAN can only be mapped to a single tenant, and all clients connecting from that VLAN are treated as members of that tenant network.

All tenants share the cluster's single file system, identity providers, RBAC configuration, and any other custom global settings.

Managing tenant access

Access to a specific protocol (such as SMB and/or NFS) can be globally granted but with per-tenant exceptions, or globally denied but with tenant-level exceptions, e.g. restricting API and Web UI access to management VLAN(s) only.

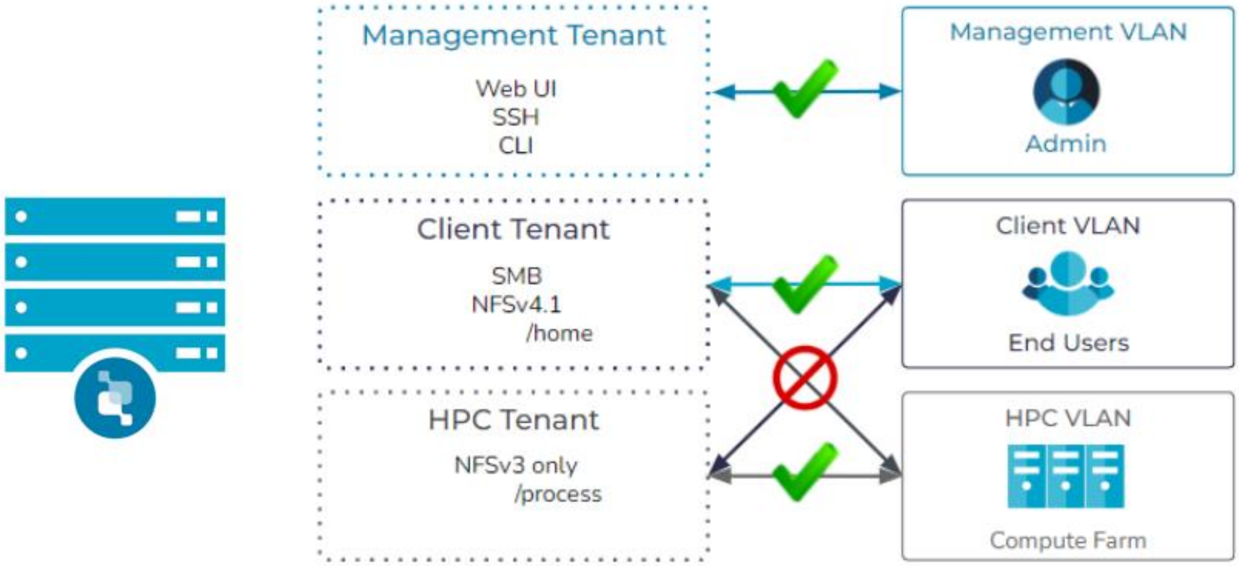

Figure 6: Network multi-tenancy

Besides granting or denying access to each tenant at a protocol level, administrators can also enable or disable access to specific shares and exports by tenant.

For more information about Qumulo's implementation of multi-tenant service management, including information on configuring and managing access to specific services, please refer to the Network multi-tenancy section of the Appendix at the end of this document.

Securing shares and exports

For scenarios in which multi-tenancy isn't feasible – such as cloud-based clusters or shares/exports that need to be accessed by some (but not all) users within a single tenant VLAN – Qumulo offers additional options to limit the visibility of shares and exports.

Share Permissions (SMB)

Access levels to SMB shares can also be managed at an individual level or based on group membership. Permissions models are simpler than the directory and file-level ACLs, offering only Read, Write, and Change permissions, as well as the option to allow or deny those permissions to individual users or groups.

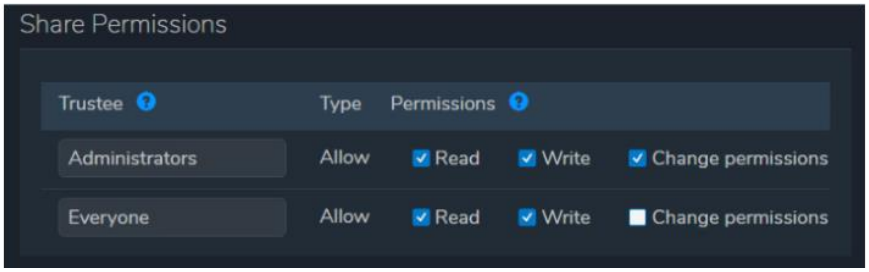

Share Permissions

Trustee	Type	Permissions		
Administrators	Allow	☑ Read	☑ Write	☑ Change permissions
Everyone	Allow	☑ Read	☑ Write	☐ Change permissions

Figure 7: Share-level permissions

Hidden shares (SMB)

Shares can also be hidden from all users, whether granted access or not. Mounting a hidden share requires explicit knowledge of the share path to block potential intruders from browsing shares.

Access-based enumeration

Qumulo allows hiding SMB Shares from unauthorized users. In addition, access-based enumeration (ABE) can be enabled for every share. By doing so, only the files and folders that a user has permission to access will be displayed to that user. If a user does not have read or equivalent permissions for a folder, the folder is hidden from the user's view.

Once implemented, only the files and folders that a user has permission to access will be displayed to that user. If a user does not have read or equivalent permissions for a folder, the folder is hidden from the user's view. Mounting the share requires explicit knowledge of the share path to block potential intruders from browsing shares.

For more information about implementing and managing hidden shares and exports, please refer to the Access-based enumeration section of the Appendix at the end of this document.

Export restrictions and host access rules

Host restrictions by client IP address range provide a good way to reduce risk surface by limiting share/export access to specific hosts, independent of the User/Group permissions of that share. Versions of this control are currently available for both SMBv3 and NFSv3.

Host restrictions (SMB)

Different address ranges may be granted full, read-only, or no access, depending on the needs of your deployment. Host permissions interact with user/group share permissions and file permissions on a “least privilege” basis, which means that in order for a privilege to be granted for a particular file, the file permissions, share user permissions, and share host permissions must all permit it.

SMB host restrictions by client IP address range are not a standard feature of the SMB protocol. By implementing this functionality Qumulo provides an additional, optional layer of SMB security.

Host access rules (NFSv3)

Host access rules can be used to govern how hosts are allowed to access an NFSv3 export. Administrators can specify one or more specific settings to comply with the export’s security requirements – including limiting access by IP address (single IP addresses, address ranges, or whole subnets); restricting clients to read-only access for the export; or mapping clients to a specific NFS username or user ID to control access levels.

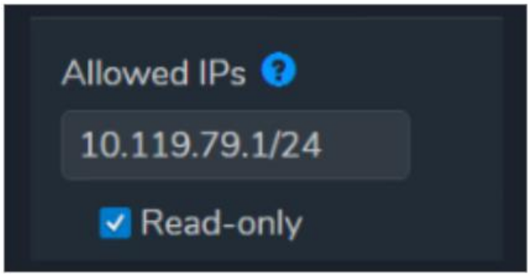

Allowed IPs ?

10.119.79.1/24

☑ Read-only

Figure 8: IP address restriction

NFSv3 access restrictions are limited to those defined by the POSIX protocol standard. NFSv4.1 offers greater control over data access settings and restrictions than NFSv3.

More information about export restrictions for the Host restrictions section of the Appendix at the end of this document.

Directory and file security

In addition to the controls described above, all of which focused on securing access to shares and exports, Qumulo also supports a number of security measures for managing access to the directories and files within each share or export.

While the Qumulo operating environment complies with the security standards defined by the different file-access protocols, and while Qumulo has also engineered a number of custom crossover features to simplify the process for maximizing data security in cross-protocol environments, there are still some protocol level differences when it comes to managing data security.

Access Control Lists

For workloads accessed via SMB and NFSv4, Qumulo supports authentication via Active Directory and Windows-style Access Control Lists (ACLs) that can be shared across both protocols.

Kerberos enhancements

All SMB and NFSv4.1 data requests, if originating from a Windows or Linux client that is joined to the same domain as the Qumulo cluster (or joined to a trusted domain), are authenticated using Kerberos-based user identity management.

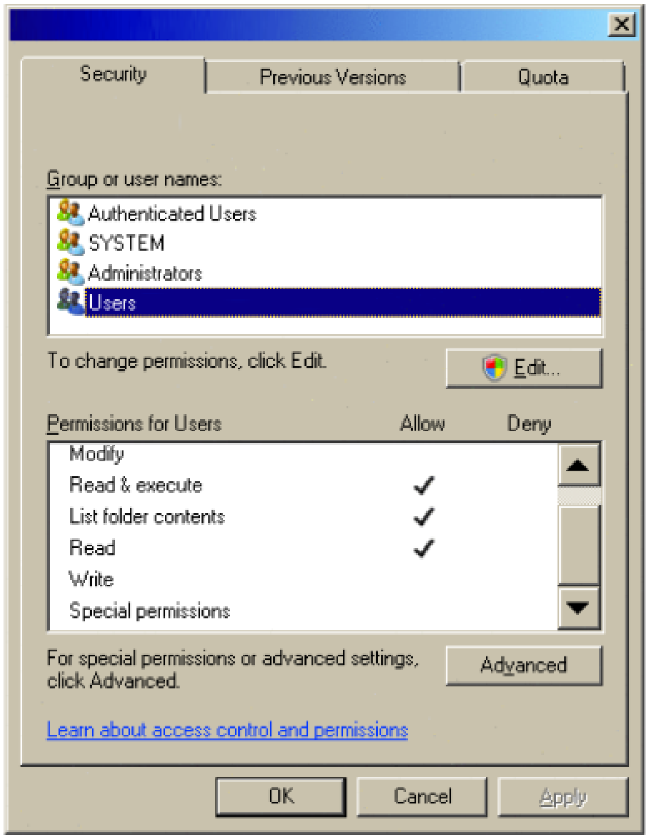

Group or user names:

- Authenticated Users
- SYSTEM
- Administrators
- Users

To change permissions, click Edit.

Permissions for Users	Allow	Deny
Modify		
Read & execute	✓	
List folder contents	✓	
Read	✓	
Write		
Special permissions		

For special permissions or advanced settings, click Advanced.

Learn about access control and permissions

Figure 9: Access Control List permissions settings

Default file permission settings, by protocol

Since every enterprise has its own particular security profile and policies, Qumulo's default security settings for new data enable data access for end users but leave it to administrators to change the default permission levels.

Default SMB permissions

For data created via SMB, the default security policy gives the cluster's local admin user "Full Control" and "Delete" permissions. Non-admin users – specifically, members of the cluster's local Users group – are given "Read" permissions to all files and subdirectories in the share created by others. They can also create their own new files and subdirectories within the share, to which they are automatically given "Modify" permissions.

Any non-administrative users will not be able to modify, change permissions, or take ownership of any files and directories unless permission is explicitly granted by the original owner of the file or directory, or by a Qumulo admin.

Default NFSv3 permissions

When a new directory is created via NFSv3, the default POSIX settings allow all users to create and delete files and subdirectories, regardless of ownership. Only the root user will be able to modify permission settings for all files and directories.

Default NFSv4.1 permissions

For workloads accessed via NFSv4, Qumulo offers full integration with Active Directory with Kerberos for user authentication, identity management, and access management. The default permissions settings match those of the SMB protocol's access control list.

Cross-protocol (SMB and NFSv3) permissions

Qumulo supports making the same data on the file system available over multiple protocols simultaneously. In many cases, an SMB share on the cluster may also be configured as an NFSv3 export, an NFSv4.1 export, and an S3 bucket. While this maximizes the cluster's flexibility, there are a number of considerations that need to be considered when managing permissions.

SMB and NFSv4.1 both use the same ACL-based permissions model, in which access is granted or denied to the user by virtue of the user's Active Directory account's membership in one or more groups whose access has been configured at the data level.

For mixed SMB and NFSv3 workloads, however, there can be a mismatch between the ACL permissions to a file or directory, and its POSIX settings. A Qumulo cluster can be configured for mixed-mode operations, in which SMB and POSIX permissions are maintained separately for all files and directories that are shared across both protocols.

For mixed-protocol workloads, Qumulo's proprietary cross-protocol permissions (XPP) model preserves SMB ACLs and inheritance even if the NFS permissions are modified.

S3 permissions

Besides sharing data via SMB and NFS protocols, Qumulo also supports making data from the file system available via S3. When a directory on the cluster is shared via S3, the directory is treated as an S3 bucket, and all subdirectories and files within that directory are treated as objects within the bucket.

Granting a user access to an S3 bucket on the Qumulo file system requires the assignment of an S3 access key, which is mapped to a known user ID – typically either an Active Directory account or a local Qumulo user. When a user attempts to access an S3 object using their S3 access key, the key's mapped user ID is checked against the object's SMB / NFSv4.1 access control list.

If anonymous access to the S3 object is required, then the object's access-control entry must also be modified to grant read access to either the cluster's Guest user ID or the Everyone group.

More detailed information on creating and managing Qumulo access keys can be found in the File and object security section of the Appendix at the end of this document.

Managing file permissions

The process for managing access to the data on a Qumulo cluster varies by protocol, since the level of access to the data also depends on the protocol used.

Qumulo provides API-based tools for modifying all access permissions, whether ACL or POSIX based, via SSH. Beyond that, access permissions can also be modified using native external tools.

- **SMB** – Access permissions for data shared via SMB can be managed using any Windows client's file permissions dialog boxes.
- **NFSv3** – POSIX permissions can be modified via SSH or any POSIX-compliant GUI
- **NFSv4.1** – Can be modified using a Windows client's dialog boxes if also shared via SMB; can also be managed using the `nfs-acl-tools` command (available in most native Linux distributions)

For more information about managing access permissions across all Qumulo-supported protocols, please refer to the File and object security section of the Appendix at the end of this document.

Data locking

Another precaution that enterprises should consider for protecting their primary data from alteration or deletion – whether via ransomware attack, rogue administrative action, or user error – is the locking of critical primary data using Qumulo's built-in file and object locking feature.

Since Qumulo supports the sharing of data via NFS, SMB and S3 protocols, locking features are enabled for both file and object data types. The specific features for file locking differ from what is supported for object data, consistent with what a file or object client will expect for a given data type.

Use cases for data locking

When it comes to file data, organizations may have financial or legal requirements for ensuring that some data types, once written, cannot be modified or deleted from primary storage. Enterprises may also be required to document that locking policies are in place to protect data from being altered or removed prematurely.

Scenarios that call for locked data may include:

- **Backup data** – Many enterprises use Qumulo as a storage target for their enterprise backup software. Since any robust ransomware recovery strategy needs to include the ability to recover lost data from a clean backup, it's essential that the backup data stored on a Qumulo instance be protected against modification or deletion from any party other than the backup software itself.
- **Archive data** – can include digital audio and visual files, software binaries, healthcare records, and other forms of intellectual property representing years, even decades, of activity. Since archived files often represent the last remaining copy of valuable old data, it can be particularly vulnerable to ransomware.
- **Audit logs** – Qumulo's virtually unlimited scalability and management simplicity make it a popular target for storing hundreds of terabytes of enterprise audit logs. As with archive data, audit logs can also contain a record of many years of corporate activity. Additionally, many industries have legal or regulatory mandates to preserve audit data over a very long period of time.
- **Medical records** – Healthcare providers and payers have regulatory and other legal requirements to maintain certain types of medical records for an indefinite timespan – e.g. as over the lifetime of the patient associated with a given record – even if the data remains on primary storage throughout its entire lifecycle.
- **Litigation and investigation** – organizations are legally obligated to ensure that any data that is part of a legal discovery effort or civil / criminal investigation cannot be modified for as long as the legal hold stands.

File locking

When applied, a file lock prevents the file from being modified or deleted. Qumulo's file-locking feature is the same over both NFS and SMB and overrides any ACL or POSIX permissions settings.

There are two types of file locking available. **Retention period** locks are enforced for a specific timeframe, which is defined at the time the lock is implemented and can be specified in days or years as appropriate. Retention periods can be extended if needed, but never shortened or removed. **Legal hold** locks are intended to last indefinitely and do not expire – they must be manually removed by an administrator.

Enabling and managing file locks

File locks can be set via CLI, REST API, or Python SDK, and require `FS_FILE_LOCK_WRITE` privileges on the Qumulo instance to enable. The ability to lock files can be combined with the Change Notify feature to lock new files once they have been created[4].

Once in place, the contents of the file, including metadata, cannot be altered. Any attempt to delete locked files from an SMB client will initially show the delete as successful; but refreshing the folder view will show the file(s) as still present. If an NFS client attempts to delete a locked file, the result will be a “Read-only file system” error. Tree-delete jobs will fail if they encounter a locked file anywhere in the tree.

File locking is an auditable event. Every attempt to access, modify, or delete locked files, whether successful or not, is written to the system’s audit logs.

Object locking and versioning

As with file locking, S3 object locking prevents the deletion of a specific version of an object for a specific period of time (**Retention period**) or indefinitely (**Legal hold**). Unlike file locking, object locking is applied and enforced per bucket rather than at the individual object level. Additionally, object locking is only available on S3 buckets that also have versioning enabled.

Object locking can be enabled using standard AWS CLI and S3 commands, either at the time of creation or afterwards. Qumulo offers the option of assigning a default locking policy to the bucket which will automatically apply to all child objects, or to enable locking at the bucket level without assigning a specific policy. Once enabled, object locking (and versioning) cannot be disabled from an S3 bucket on a Qumulo system.

Since Qumulo supports the sharing of data via both file and object protocols, object locking will prevent the impacted data from being modified or deleted via NFS or SMB access, and an error message will be returned to the client. It is possible, however, to move an object (i.e. no modification of either data or metadata in the locked object) by creating a hardlink to the new location and then deleting the old link.

Similarly, data protected by a file lock cannot be modified or deleted over the S3 access protocol.

For more information regarding the implementation and management of file and object locking, please refer to the Data locking section of the Appendix at the end of this document.

[4] Care should be taken when combining file locking with SMB Change Notify to ensure that files have been fully written to disk before they are locked. No further changes to the file are permitted once the lock is in place.

Over-the-wire data encryption

Even with the appropriate share and data-level security settings in place, some enterprises need an additional layer of data security to protect data from unauthorized access. For those environments, Qumulo also supports over-the-wire data encryption to and from supported clients.

SMB

With Qumulo support for SMB3 encryption, a cluster-wide level or a per-share level encryption can be enabled. Share encryption can be implemented cluster-wide for all SMB shares if needed, or per-share if only required for some data.

NFSv4.1

For NFSv4.1 workloads that require it, Qumulo supports two types of Kerberos-based elevated security over the wire:

- **kr5bi** – Data packets are transmitted with a hash signature that ensures data integrity. Data can be intercepted and examined, but any attempt to modify data during transit will invalidate the data.
- **kr5bp** – All data packets are encrypted in transit to eliminate the risk of data being intercepted and read in transit.

S3

Per AWS S3 standards, all data transmitted via S3 is encrypted over the wire using TLS (HTTPS).

FTP

FTP access can be enabled and encrypted using standard TLS settings if needed, for both user authentication and data transmission.

FTP services are disabled on a new Qumulo cluster by default.

More information about protocol-specific encryption options and settings is available in the Over-the-wire data encryption section of the Appendix at the end of this document.

Quotas

Like other storage platforms, Qumulo provides a quota service that can be used to manage storage utilization across the cluster. Unlike most other storage systems, Qumulo's implementation of a quota service is integrated directly into the file-system management services, eliminating the need for tree walks to track usage and enabling immediate enforcement when activated.

While the primary focus of quotas is to manage capacity utilization under normal operating conditions, Qumulo's intelligent quota services can also lower the risk of a runaway user inflicting damage to critical data. In response to a detected malware or intrusion event, a Qumulo quota can be applied at any level of the file system and used to prevent any further write operations. This can be done either manually or programmatically.

For more information on managing Qumulo quotas, please refer to the Quotas section of the Appendix at the end of this document.

Part 2: Detecting Intrusion

Despite all the available precautions, policies, and best practices designed to protect data and minimize risk, there is no fail-safe guarantee against the possibility of a malware attack, or an intruder gaining unauthorized access to critical systems and data. Qumulo's overall approach is consistent with the framework outlined by the National Institute of Standards and Technology (NIST)[5], and is intended to enable organizations to protect critical data, as well as to detect and respond to ransomware events quickly.

This approach, while designed to maximize upfront security by enabling systems and data protection across multiple layers and on multiple fronts, recognizes that intrusive events still happen. As such, Qumulo also provides the tools to enable impacted organizations to restore affected services and recover lost data

Implementing a holistic security approach that includes network, compute, device and event-monitoring techniques, together with data correlation and analysis, is preferable to siloed solutions that are embedded in the storage system alone.

Qumulo supports all major ISV security solutions through its auditing feature. In addition, Qumulo's API enables automated mitigation actions to deliver a robust and timeline response.

There are more comprehensive security-focused resources available that provide much more in-depth, actionable information for how to protect enterprise resources against viruses, malware, ransomware, and other attack vectors. The next few sections of this document will focus primarily on how to ensure that safeguards are in place to detect and contain cyberattacks and malware outbreaks, quickly and effectively.

[5] Detailed information on NIST's cybersecurity recommendations for managing ransomware risk is available in the "Basic Ransomware Tips" section on pp. 2-3 of "Ransomware Risk Management: A Cybersecurity Framework Profile," which can be found here.

Antivirus scanning

The first line of antivirus (AV) prevention should be the data center security infrastructure. This can include firewalls, network scanning, email servers and desktop clients. It is essential to understand that if malware reaches the storage system, the data can be compromised.

That said, there are methods, tools, and techniques that can be used to minimize the risk of malware reaching the storage system, including:

- **On-demand-scans**: Qumulo recommends on-demand scans using any of the major antivirus solutions on the market. These can be scheduled to run regularly – preferably during off-peak hours[6].
- **Client-side scans**: This is the optimal point for AV scanning since client systems are the likeliest point of outbreak for malicious payloads. Qumulo suggests adopting a next-generation antivirus engine, based on AI and binary fingerprinting rather than signatures that can be modified easily by advanced attackers. Qumulo also recommends a proper patch-management strategy, using a whitelisting approach that allows only legitimate, IT-controlled software to execute.

Recommended antivirus practices

To avoid the disadvantages of running antivirus scans on the storage cluster itself or using external ICAP agents[7], Qumulo recommends the following:

1. If on-access scanning is required, use client-side AV scanning agents, which are readily available for all major platforms and come with no significant downsides.
2. Use on-demand or scheduled scans of the Qumulo cluster as appropriate, using any major AV solution. These can be performed in off-peak hours during normal operations, or they can be

[6] On-demand scans on Qumulo can be completed much faster than against a scale-up NAS platform because multiple scanners can be run in parallel against multiple nodes in the cluster simultaneously.

[7] Some AV vendors offload virus scanning from both the client and the storage server using the Internet Content Adaptation Protocol (ICAP). Using this method, every time a file is opened, it is sent to an antivirus SW instance on an external host, which scans the file before the client is able to open it. The downsides to this approach include:

- Unacceptable response times: Qumulo is a scale-out NAS solution that delivers typical response times in the range of 0.5-5 ms. If an antivirus engine for scan-on-open is inserted between the storage and the client, response times would typically increase to several seconds per file, even for small files on a fast network. Large files can be expected to take even longer to open. This is usually not acceptable for a high-performance storage solution where users and applications require fast response times.
- A high number of scanning servers: On-access scans require a very large server farm for the scanners – typically 1-2 physical servers per Qumulo storage node. This is not an optimal use of enterprise resources, since viruses can be detected and contained much more efficiently on the client, firewall, or email server.

activated on-demand in response to a perceived or actual malware threat (see the Automating responses to security events section of this document for more information).

3. Use regularly-scheduled snapshots to enable rapid recovery in the event of a malware outbreak or security breach.

For more information on antivirus management tools and practices, please refer to the Antivirus section of the Appendix at the end of this document.

Ransomware

Where viruses generally set out to inflict damage by destroying data and/or systems, ransomware aims to monetize the infection. The goal of a ransomware attack is usually either to exfiltrate data and threaten to release it unless a ransom is paid, or to encrypt data on-prem and attempt to force the victim to pay for the keys to decrypt and recover the data.

Common ransomware vectors

While a ransomware attack can come from anywhere, most of them originate from one of the following common vectors, most of which are the result of user negligence rather than brute-force external attacks[8]:

- Spear-phishing emails (the most common)[9]
- Trojanized software[10]
- Web server exploits and watering-hole websites[11]
- Domain spoofing[12]
- Exploitation of unpatched operating-system vulnerabilities[13]

[8] "A Blueprint for Ransomware Defense Using the CIS Controls" CIS, 25 Sept. 2024, https://www.cisecurity.org/insights/blog/a-blueprint-for-ransomware-defense-using-the-cis-controls.

[9] "How to Spot a Phishing Email." IT Governance Blog En, 25 Oct. 2022, www.itgovernance.eu/blog/en/5-ways-to-spot-phishing-scams.

[10] "Trojan Malware." Microsoft Learn, 24 Apr. 2024, https://learn.microsoft.com/en-us/defender-endpoint/malware/trojans-malware.

[11] "Watering Hole Attacks." National Cyber Security Centre, www.ncsc.gov.uk/collection/supply-chain-security/watering-hole-attacks.

[12] "What Is Domain Spoofing? - Nilesh Parashar - Medium." Medium, 3 Dec. 2022, medium.com/@niitwork0921/what-is-domain-spoofing-5bf43fd7fd44.

[13] "Understanding Patches and Software Updates | CISA." Cybersecurity & Infrastructure Security Agency, 19 Nov. 2019, www.cisa.gov/uscert/ncas/tips/ST04-006.n

- Man-in-the-middle attacks[14]
- Cross-site scripting[15]
- SQL injection attacks[16]

While the threat of many of these vectors can be reduced through a combination of user training and regular system maintenance – including patching system security vulnerabilities before they can be exploited – these actions serve to collectively reduce risk. They do not eliminate it entirely.

Ransomware impacts

Once a ransomware attack has occurred, enterprises have a limited set of options to choose from: they can recover their data using their own internal protection measures; they can choose to lose the data; or they can pay the ransom in hopes of getting their data back.

While the potential downstream impacts of a malware or ransomware attack can be both numerous and severe, the purpose of this document is to focus rather on the options, controls, and solutions that Qumulo provides – for minimizing the risk of an attack, as well as minimizing the impact in the event that one occurs.

Detecting and containing ransomware

In order to optimize the potential responses to an intrusive event, it's important to identify the different phases of an attack, and to plan out potential responses at each phase.

When they do occur, ransomware attacks typically play out in the following order[17]:

1. **Delivery** – compromise the network by gaining access to at least one internal system
2. **Command and control** – once inside, establish a connection with the attacker's command-and-control server to receive instructions

[14] What Is a Man-in-the-Middle Attack: Detection and Prevention Tips. 24 Feb. 2022, www.varonis.com/blog/man-in-the-middle-attack.

[15] "Cross-Site Scripting (XSS) Attacks and How to Prevent Them." Splunk-Blogs, 29 Aug. 2024, www.splunk.com/en_us/blog/learn/cross-site-scripting-xss-attacks.html.

[16] "SQL Injection - SQL Server." Microsoft Learn, 3 May 2024, learn.microsoft.com/en-us/sql/relational-databases/security/sql-injection?view=sql-server-ver16.

[17] Source: "5 Stages of the Ransomware Lifecycle | JPMorgan Chase." 7 Sept. 2022, www.jpmorgan.com/commercial-banking/insights/the-anatomy-of-a-ransomware-attack.

3. **Credential access** – under stealth, obtain credentials and gain access to more accounts across the network

4. **Canvas** – search for files to encrypt, both on the local compromised system and on any networks and services to which it has gained access through lateral movement

5. **Extortion** – attacker exfiltrates and/or encrypts local and network files, then demands payment to either decrypt files or return exfiltrated data

A comprehensive prevention strategy should focus primarily on the Delivery phase of the outbreak. Considering the most common vectors of an attack, the most effective prevention steps involve securing the likeliest entry points for malware: desktop computers, email servers, and network devices. Since modern antivirus packages can also detect and eradicate malware before it reaches the end user, antivirus software is by far the most effective means of preventing outbreaks in the first place.

Security Information and Event Management

One key takeaway from this: while the Qumulo cluster may contain vast amounts of critical and sensitive data that needs to be protected against a malware outbreak, the storage layer of the enterprise is actually best protected when administrators focus on securing the perimeter and client systems as the first line of defense. Any damage that an infection may inflict is best prevented and minimized by detecting and neutralizing it on one of the systems where the malware lands initially.

Since some enterprise nodes (e.g. IoT systems, cameras, printers) within the data center environment can't be secured using local antivirus software, the next layer of security from malware/ransomware should be a robust security information system and event management (SIEM) solution.

SIEM platforms[18] capture and compile event- and security-log data from enterprise systems. When an enterprise SIEM solution is paired with an Intrusion Detection System[19] (IDS), administrators can identify traffic and activity patterns that indicate potential and actual threats: anomalies in network traffic and server behavior, as well as unusual data read and write activities on enterprise file services.

Many organizations are also using intrusion prevention systems[20] (IPS) that have advanced firewalling and exploit-detection capabilities that can help to fence off some categories of attacks.

[18] See "The Best SIEM Tools for 2024: Vendors & Solutions Ranked" for a list of recommended SIEM platforms for amalgamating enterprise system and event log data.

[19] See "The Best Network Intrusion Detection Systems Software & NIDS Tools" for recommended IDS solutions for identifying and preventing cyberattacks.

[20] See "The Best IPS Software Tools for 2024 & Guide" for the latest list of recommended IPS software tools.

Auditing Qumulo storage events

Audit logging provides a mechanism for tracking Qumulo file-system events as well as management operations. As connected clients issue requests to the cluster, event-log messages are generated describing each attempted operation. These log messages are then sent over the network to the remote syslog instance – i.e. the designated SIEM target – specified by the current audit configuration in compliance with RFC5424[21].

The advantages of this approach include:

- Qumulo's use of an industry-standard syslog format means that event log data – including audit logs – can be read, parsed, and indexed by any standard SIEM solution.
- Any and all data-access and system-management events from the Qumulo cluster can be captured, logged, and retained.
- Since the target instances of all these log events are collected and stored outside of the Qumulo cluster, they cannot be manipulated or deleted if the cluster is subjected to a security breach.
- A SIEM platform can be configured to incorporate automated or semi-automated actions in response to suspicious activity.

Among others, Qumulo-validated SIEM solutions include Splunk[22], Elasticsearch and AWS Cloudwatch. For Intrusion Detection solutions, Qumulo integrates with Varonis[23], Netwrix[24], and others. Qumulo also supports the OpenMetrics API standard for syslog exports, enabling integration with Prometheus-based monitoring solutions as well.

More information about Qumulo's integration with SIEM and other enterprise management tools can be found in the Auditing section of the Appendix at the end of this document.

Automating responses to security events

Once a suspicious event or malicious activity has been detected on the storage system, the SIEM and Qumulo systems can be configured to trigger one or more automated actions in response. As with other event types, such as platform or service outages and network issues, the SIEM platform can be configured for certain automated actions. These might include an administrator alert tree (email and/or

[21] For more information, refer to RFC5424: The Syslog Protocol.

[22] See also "Using Splunk with Qumulo Core Audit Logging" for more information.

[23] For more information on how Varonis integrates with Qumulo to identify ransomware attacks, click here, or download the Varonis solution brief

[24] More information, with downloadable solution brief, at "Netwrix with Qumulo for Intelligent File Data Security"

text messages), disabling an AD user account, or shutting off network ports when suspicious activity is detected, etc.

Leveraging the Qumulo API

There are a number of ways to leverage the Qumulo API for automated responses to security events:

1. Direct API calls
2. Use the Qumulo-provided Python libraries to simplify API script development
3. Use the Qumulo CLI

More information about managing automated responses to security events can be found in the Automating intrusion responses section of the Appendix at the end of this document.

Activating automated Qumulo responses

Qumulo's API-first development model means that literally any action on the storage cluster can be initiated and managed via API. Some examples of automated actions that can be triggered in the event of a security breach, malware detection, or other forms of cyberattack include:

- Immediately apply a 0-quota policy – applicable to a single directory, directory tree (including whole shares and/or exports, or the entire file system – that blocks all further write activity (although overwrites might still be possible)
- Set any targeted exports to read-only
- Set an IP address restriction policy to any share or export
- Remove access privileges for a user or group
- Initiate an on-demand snapshot of any suspected
- Lock one or more existing snapshots

Automating responses on other enterprise platforms

While there are many first-response actions that can be initiated on the affected Qumulo system, other platforms within the enterprise space can be leveraged to secure critical systems and data in the event of a malware attack, e.g.:

- Active Directory – disable any compromised user or service accounts as soon as hostile activity has been detected and the offending accounts have been identified

- Antivirus software – launch an on-demand antivirus scan of any and all systems, including the Qumulo cluster, which are known or suspected to be under attack

For specific information on these and other actions to third-party platforms, please refer to the appropriate vendor-provided documentation.

Part 3: Intrusion recovery

Just as intrusion protection can minimize but not eliminate risk, rapid detection and automated responses can minimize, but not completely eliminate, the impact to affected systems and data. A comprehensive plan against cyberattacks – whether malware, ransomware, virus outbreak, or other form of malicious activity – must necessarily include steps to minimize the risk of an outbreak in the first place, and a rapid-response plan to limit the scope and impact when an attack does occur. Beyond that, there must be measures in place to ensure that any impacted systems and data can be recovered.

Data never stands still in an enterprise environment, and a single Qumulo cluster can host petabytes of unstructured data. In a relatively large enterprise where daily change rates are less than 1%, a 5PB Qumulo cluster could see up to 50TB of new and altered data per day. Since data is constantly changing, a data-recovery plan needs to minimize the amount of data that can be lost to a malware outbreak.

As such, a recovery strategy can be complex. It should include all the systems and data that might be affected by a cyberattack, along with mitigation plans in place that address all possible impacts.

It is beyond the scope of this paper to provide a prescriptive recovery plan. Since every enterprise operates under unique circumstances – data volume, change rates, retention and recovery objectives, and other considerations and constraints – every recovery strategy will also be unique.

This section of the document provides an overview of the Qumulo controls and features that can be leveraged to ensure timely recoverability of any impacted data.

Snapshots

Snapshots on a Qumulo cluster can be used in several ways to protect the cluster’s data. They can be used alone for quick and efficient data protection and recovery. A snapshot of the live data on one Qumulo cluster can be replicated to a secondary cluster or to S3 storage. Alternatively, a Qumulo snapshot can be paired with a third-party backup solution to provide effective long-term protection (with more robust version control for changed files) against data loss.

Snapshot basics

By itself, a Qumulo snapshot is a very efficient method of protecting data. A snapshot can be taken at any point in time, either according to a fixed schedule, or on-demand as needed. Once taken a snapshot consumes no space initially.

A snapshot preserves everything in the file system – file data, directory entries, creation and modification times, permissions, etc. As files within the snapshot change over time, new data is written alongside the original version, and new entries are written in the file system identifying each version of the same file.

If a file or directory needs to be rolled back to a previous version, it can be recovered from a snapshot. Since each snapshot is immutable (read-only), a potential malware or ransomware will not be able to change the data within a snapshot. Snapshots can be further protected by locking, which prevents a snapshot from being deleted until a set period of time has passed.

Snapshot policy management

In order to automate a scheduled snapshot, a storage administrator must first define a named snapshot policy defining the parameters and lifecycle of the snapshot. These include:

- **Target directory** – this can exist at any hierarchical level within the file system.
- **Schedule** – can be configured using a wide range of settings, including hourly or daily; on one or more days per week, or more/less frequently as needed.
- **Naming format** – used to determine how to ensure each snapshot version has a unique yet identifiable name to facilitate recovery later if needed
- **Lifecycle** – determines whether a snapshot is allowed to expire automatically or must be manually deleted.
- **Locking** – if enabled, prevents a snapshot from being deleted before expiring.

On-demand snapshots

Besides scheduled snapshots, Qumulo also supports the use of on-demand snapshots, which can be taken at any time from anywhere in the file system. As with scheduled snapshots, an on-demand snapshot can be set either to expire automatically, or to be retained (with optional locking) until an administrator deletes it.

As mentioned earlier in the document, on-demand snapshots can be automated using the Qumulo API, and programmed into the SIEM or IDP platform in response to a detected security event.

Snapshot Locking

Snapshot immutability provides a historical reference for the state of the data at a particular point in time, rather than the actual data itself. Immutability is not a foolproof guarantee against external tampering, however. A rogue administrator, or an external actor using a compromised user account with elevated privileges and access to the web UI, CLI or API, can either change the snapshot's expiration time so it's purged prematurely, or delete the snapshot outright.

To provide added protection, administrators can apply a "lock" as part of the password's policy settings. If enabled, snapshot locking prevents the alteration or deletion of a snapshot – even by a storage administrator – prior to its expiration time.

Enabling snapshot locking

To use the snapshot locking feature, an administrator must first generate an asymmetric cryptographic key pair using a Qumulo-supported key-management service (KMS), either through their preferred KMS or preferred open-source tool, or by using the Qumulo CLI to generate one directly from the cluster. The public key is installed on the cluster – manually if the key was generated externally, or automatically if the key was created through the CLI – and the private key must be secured using the customer's own established security practices.

Managing locked snapshots

Under ordinary operating circumstances, locked snapshots will be allowed to expire on their own, and administrators can leverage them as necessary in the event of a cyberattack, secure against even administrative tampering.

If a customer wants to delete a locked snapshot prior to its expiration date – e.g. to reclaim space on a nearly-full cluster – they will need to generate a request on a separate system using the private key, then upload the signed and authorized request to the cluster before the locked snapshot can be adjusted.

For more information about managing snapshots on a Qumulo system, please refer to the Snapshot management section of the Appendix at the end of this document.

Duplicating data to secondary storage

For additional data protection, Qumulo offers a built-in replication feature to ensure that at least one copy of all data is hosted in secondary storage from which it can be recovered if data on the primary storage instance is lost.

This secondary storage instance can be another Qumulo cluster – whether on-premises, at a second data center, or hosted in one of the public cloud providers (Amazon Web Services, Google Cloud, or Microsoft Azure) – or an S3 target.

Qumulo replication

Qumulo's built-in replication service can copy data at scale between any two Qumulo storage instances. Besides protecting data against cyberattack, a secondary location with another Qumulo cluster can also serve as failover storage in the event of a site-level outage.

Since Qumulo storage can be deployed anywhere – in a second data center or in any of the public-cloud platforms (Amazon Web Services, Google Cloud, and Microsoft Azure) – and delivers the same services regardless of location, replication can be configured to run in any direction between any two Qumulo endpoints.

Continuous replication

This form of replication simply takes a snapshot of the data on the source Qumulo cluster and copies it to a directory on a target cluster. As long as the replication relationship is active, the system scans any modified files to identify and copy only the specific changes to the target.

While this does create a copy of the primary dataset on a second cluster, it is not recommended as a standalone data protection solution, since any data corruption or loss on the primary source will also propagate to the target.

Snapshot-based replication

With snapshot-based replication, snapshots are also taken of the target directory on the secondary cluster. Once a replication job has been completed, a new snapshot of the target directory is created, ensuring consistent data on both clusters, as well as maintaining a change log for each file on the target.

These snapshots can be set to use the same expiration as the source snapshot policy, or they can be configured with longer expiration times to provide a longer recovery window in the event of a malware attack.

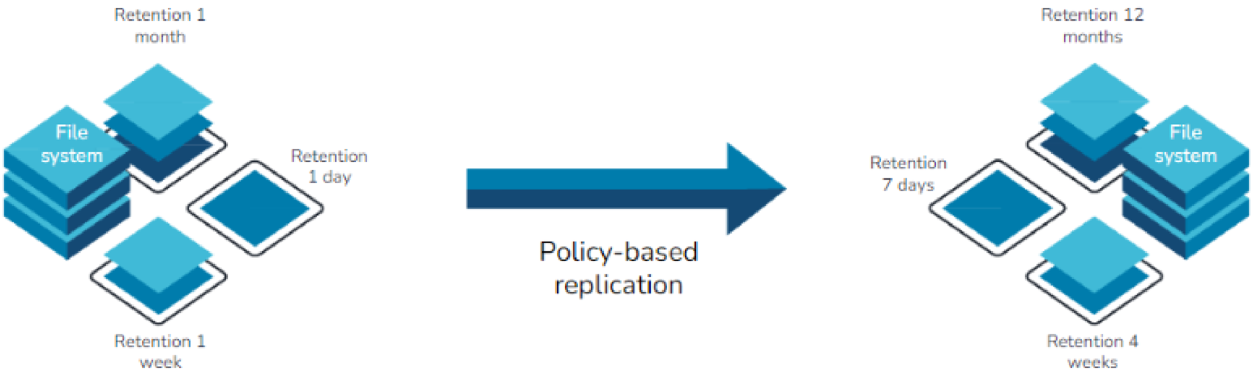

Figure 10: Snapshot-based replication

For more information on managing scheduled, on-demand, and automated snapshots, please refer to the Replication section of the Appendix at the end of this document.

Replication to / from S3 storage

Where the native replication service provides the ability to copy snapshots between two Qumulo clusters, Qumulo can also copy data in either direction between the local file system and an AWS S3 bucket.

Within the AWS S3 bucket, all files are written in native S3 format, with no gateway required. Subsequent copies are incremental so that only changed files are copied – another reliable way to protect Qumulo data offsite and away from an on-premises attack.

Depending on the workload or data change rates it is possible to copy consistent snapshots to an S3 bucket and use AWS versioning to track the multiple versions of files in the bucket. Once data has been written successfully to AWS S3 storage, intelligent tiering can be leveraged to move older files to AWS Glacier or Glacier Deep Archive Storage, providing a cost-effective data storage solution for data that is not in active use.

Like all other tasks on Qumulo storage, data movement to and from AWS S3 storage can be automated by using the API or CLI. For more information on this feature, please refer to the Copy to / from S3 links in the Appendix at the end of this document.

External backup integration

For longer-term data protection, and the ability to maintain a longer version history for critical files, Qumulo integrates with any file protocol-based[25] backup solutions.

Some vendors, such as CommVault and Atempo, use the Qumulo API to compare snapshots and identify changes, enabling them to take instantaneous incremental backups without the need for a tree walk. Another advantage to file-based backups is that the backup image is platform agnostic, meaning that, if necessary, data can be restored to any storage target.

Integrating with Qumulo at an API level like this also enables an incremental-forever backup strategy with minimal effort. More information on Qumulo's support for external backup platforms is available in the Backup software integration section at the end of this document.

Conclusion

Qumulo takes security seriously, integrating data security and protection into the storage architecture at every level:

- A compact, hardened operating environment
- Fully native software stack
- Effective isolation between the underlying Linux kernel and the storage architecture
- A regular cadence of firmware updates, deployed via an instant upgrade process
- Ability to separate administrative from user access
- Ability to hide and / or restrict access to SMB and NFS shares
- The use of file and object-based locking features
- Administrative hardening using single sign-on with multi-factor authentication
- Role-based access control
- FIPS compliant software-based data encryption at rest
- Active Directory integration with Kerberos identity management
- Robust cross-protocol permissions management in mixed-mode environments
- Optional over-the-wire encryption for SMB, NFSv4.1 and FTP traffic

[25] E.g. SMB and / or NFS capable backup software. Qumulo does not support NDMP-based backup solutions.

- Integration with Security Information and Event Management and Intrusion Detection platforms
- API-first management model for automated responses to security threats

Not only that, but Qumulo simplifies data protection by enabling instantaneous data protection through snapshots, enabling data recoverability through replication to secondary storage, and through seamless integration with major backup software platforms – all features that can be leveraged to protect not just your systems and data, but also your enterprise, from malicious attacks.

As a data-management solution that's trusted by enterprises worldwide with their most critical and sensitive unstructured data, Qumulo recognizes the importance of securing and protecting critical data in an unsecure world. To make security simpler, Qumulo's software environment has been engineered to support the holistic security model needed to protect enterprise data against today's threats.

Appendix

This section provides links to additional references and resources, both internal and external, to assist in planning and maintaining a robust security model based on Qumulo's unstructured data management solutions.

Note: access to some of the following links may require a valid Qumulo support agreement.

Additional links and references

The following links, organized largely by topic, are intended to provide more in-depth information on the concepts, practices, and recommendations discussed in this document.

In addition to the links below, the following documents and portals may also prove useful, as Qumulo's security features and documentation will continue to be updated after this document's publication.

- Qumulo Documentation: Qumulo Administrator Guide (downloadable PDF)
- Qumulo online administration reference: Qumulo Documentation
- Qumulo's Knowledge Base portal: Qumulo Care

Intrusion prevention

Qumulo instant upgrades

- Qumulo administration guide: Performing Instant Software Upgrades and Platform Upgrades

Single sign-on with multi-factor authentication

- Blog post: How Qumulo's Built-in Security Approach Makes Your Unstructured Data Simply Secure
- Qumulo administration guide: Configuring SAML Single Sign-On (SSO) for Your Qumulo Cluster

Access tokens

- Blog post: How Qumulo's Built-in Security Approach Makes Your Unstructured Data Simply Secure
- Qumulo documentation: Using Qumulo Access Tokens

Role-based access control

- Blog post: Data Protection: Preventing Malware Incidents with Built-In Security Controls
- KB article: Managing Role-Based Access Control (RBAC) with Qumulo

Management VLANs

- Blog post: How Qumulo's Built-in Security Approach Makes Your Unstructured Data Simply Secure
- KB article: Separate Cluster Management and Production Traffic

Software-based encryption

- Blog post: Software-Based Encryption Ensures Data Security and Compliance at Multi-Petabyte to Exabyte Scale
- Blog post: Qumulo Awarded the FIPS 140-2 Certificate for Qumulo Secure TLS KDF
- White paper: Qumulo Secure Software-Based Encryption
- KB article: Qumulo's Encryption at Rest
- Qumulo administration guide: Qumulo Compliance Posture
- NIST certificate of compliance for on-prem deployments: Qumulo Cryptographic Algorithm Validation

Securing and protecting shares and exports

Network multi-tenancy

- Blog: How Qumulo's Built-in Security Approach Makes Your Unstructured Data Simply Secure
- Qumulo documentation: Qumulo Administrator Guide
- Qumulo documentation: Partitioning a Qumulo Cluster into Tenants
- Qumulo documentation: Configuring Management Protocols on a Tenant
- Qumulo documentation: Configuring File System Protocols on a Tenant

Access-based enumeration

- KB article: Hide an SMB Share

Host restrictions

- KB article: SMB Host Restrictions
- KB article: Create an NFS Export

File and object security

- KB article: Qumulo File Permissions Overview
- KB article: Default File Permissions in Qumulo
- KB article: Cross-Protocol Permissions (XPP)
- Qumulo administration guide: Using Kerberos Permissions in the Qumulo File System
- Qumulo administration guide: Managing File Access Permissions by Using NFSv4.1 Access Control Lists (ACLs)
- Qumulo administration guide: Managing Access to S3 Buckets in a Qumulo Cluster
- Qumulo administration guide: Creating and Managing S3 Access Keys in Qumulo

Over-the-wire data encryption

- KB article: SMB3 Encryption with Qumulo
- Qumulo administration guide: Performing Additional Cluster Configuration after Joining Active Directory
- Qumulo administration guide: FTP: TLS Security

Data locking

- Qumulo administration guide: Managing File Locks in Qumulo Core
- Qumulo administration guide: Enabling Object Lock for S3 Buckets

Quotas

- KB article: Quotas in Qumulo

Intrusion detection

Auditing

- KB article: Qumulo Audit Logging with Splunk
- KB article: Qumulo Audit Logging with Elasticsearch
- KB article: Qumulo in AWS: Audit Logging with CloudWatch
- KB article: How Qumulo Core Integrates with Varonis
- Qumulo administration guide: Qumulo OpenMetrics API Specification
- Code / script libraries: Prometheus Github repository

Antivirus

- Blog: Detecting Ransomware Attacks in Real Time with Qumulo

Automating intrusion responses

- KB article: Getting Started with Qumulo's REST API
- Qumulo administration guide: Managing Snapshots in Qumulo
- Code / script libraries: Qumulo Github repositories

Data protection and recovery

Snapshot management

- KB article: Snapshots: Deep Dive
- Qumulo administration guide: Managing Snapshots in Qumulo
- KB article: QQ CLI: Snapshots

Replication

- KB article: Replication: Creating and Managing a Continuous Replication Relationship in Qumulo Core
- KB article: Snapshot Policy Replication

Copy to / from S3

- Qumulo administration guide: Using Qumulo Shift to Copy Objects to Amazon S3
- Qumulo administration guide: Using Qumulo Shift to Copy Objects from Amazon S3

Backup software integration

- KB article: Identify File and Byte Range Changes between Snapshots