

HP Server Automation Storage Visibility and Automation

for the HP-UX, IBM AIX, Red Hat Enterprise Linux, Solaris, SUSE Linux Enterprise Server, VMware, and Windows® operating systems

Software Version: 7.80

User Guide

Document Release Date: June 2009
Software Release Date: June 2009



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

© Copyright 2000-2009 Hewlett-Packard Development Company, L.P.

Trademark Notices

Intel® Itanium® is a trademarks of Intel Corporation in the U.S. and other countries.

Java™ is a US trademark of Sun Microsystems, Inc.

Microsoft®, Windows®, Windows® XP are U.S. registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

- 1 Introduction 11
 - Overview of Storage Visibility and Automation 11
 - Basic Business Needs 12
 - Feature Capabilities 14
 - Icons Used in Storage Visibility and Automation 20
- 2 Asset Discovery 23
 - Overview of Asset Discovery 23
 - Data Collection 23
 - Storage Scanners 23
 - Update from Storage Essentials 25
 - Open with Storage Essentials (Click & Launch) 26
 - Storage Device 26
 - Storage Host Agent Extension (SHA) 27
 - Server Assets 28
 - Scheduling Server Asset Discovery 29
 - Server Storage and SAN Relationships 30
 - Frequently Asked Questions 31
 - Viewing FCAs Attached to a Server 32
 - Viewing Ports and Volumes in an FCA 32
 - Viewing Port Connections and Zones 33
 - Viewing Storage Summary in a Server 33
 - Viewing File Systems Mounted on a Server 33
 - Viewing File System Properties 33
 - Viewing Volume Details 33
 - Viewing Local Disks 34
 - Viewing the Storage Supply Chain 34
 - Viewing Volume Properties 34
 - Viewing Antecedents and Dependents of Volumes for a Managed Server 34
 - Viewing Access Paths for a Volume 35
 - Viewing Volume Connectivity 35
 - Viewing Disks Utilized by a Volume 35
 - Viewing Local Disks or DAS Attached to a Server 36
 - Viewing Disk Properties 36
 - Viewing Volumes Created on a Disk 36
 - Viewing File Systems Created on a Disk 36
 - Viewing Volume Manager, MPIO, VMFS, & ASM 37
 - Viewing Volumes Available Through Manager Software 37
 - Viewing Volumes Available Through the Hypervisor 37

Viewing Unmounted Volumes Through the Hypervisor	37
Viewing Unmounted Volume Properties	38
Viewing Antecedent of Aggregate Storage Supply Chain	38
Viewing Remote SAN Based Volumes	38
Viewing Storage Targets.	38
Viewing Access Paths to Storage Targets	38
Viewing Switches Attached to a Server	39
Viewing Fabrics Attached to a Server	39
Viewing FCA Ports Attached to a Fabric.	39
Viewing Zone Membership	39
Viewing Databases Hosted by a Server	40
Viewing Hosts in the Hypervisor	40
Viewing Pools on the Hypervisor.	40
Viewing Pool Properties on the Hypervisor.	40
Viewing Pool Volumes on the Hypervisor	40
Viewing Storage Initiators & Volumes.	41
SAN Array Assets	41
SAN Array Assets and Relationships.	42
Update from Storage Essentials	42
Open with Storage Essentials.	42
Frequently Asked Questions	43
Viewing SAN Arrays	43
Viewing SAN Arrays and Storage Configurations.	44
Viewing Summary of a SAN Array	44
Viewing SAN Array Properties	44
Viewing Customers for a SAN Array.	45
Viewing the SAN Array History Log.	45
Viewing Storage Initiators Consuming Storage Resources from an Array	45
Viewing Initiator Volumes and Ports	46
Viewing Storage Volumes with Access Paths for Remote Initiators	46
Viewing Attached Fabrics in an Array	46
Viewing Fabric Ports and Connectivity in an Array	46
Viewing Zones Configured for an Array	47
Viewing SAN Switches Attached to an Array	47
Viewing Hardware Information for a SAN Array	47
Viewing Front End Fabric Ports in a SAN Array.	47
Viewing Properties of the Fabric Port in a SAN Array	47
Viewing Volumes Served by a SAN Array Port.	48
Viewing Zones Configured for an Array Port	48
Viewing Storage Summary for a SAN Array.	48
Viewing Volumes in a SAN Array	48
Viewing Volume Properties	49
Viewing the Connectivity of a Volume	49
Viewing the Composition of a Volume	49
Viewing the Access Path from the Volume	49
Viewing Array Ports from the Selected Volume	50
Viewing Storage Pools	50

Viewing Storage Pool Properties	50
Viewing Volumes for a Selected Pool	50
Fabric Assets	51
Fabric Assets and Relationships	52
NetApp Assets	52
NetApp Assets and Relationships	54
Update from Storage Essentials	54
Open with Storage Essentials	54
Frequently Asked Questions	55
Viewing NAS Filers	55
Viewing the NAS Filer Storage Configuration	55
Viewing Summary of a NAS Filer	55
Viewing NAS Filer Properties	55
Viewing Customers for a NAS Filer	56
Viewing the NAS Filer History Log	56
Viewing Attached Fabrics in a NAS Filer	56
Viewing SAN Switches Attached to a NAS Filer	56
Viewing vFilers	57
Viewing Volumes Configured for a vFiler	57
Viewing Exports for a vFiler	57
Viewing Storage Summary for a NAS Filer	57
Viewing Disks in a NAS Filer	57
Viewing NAS Filer Disk Properties	58
Viewing Filer Volumes Created by a Disk	58
Viewing Aggregates Created by a Disk	58
Viewing Plexes on a Disk	58
Viewing RAID Groups on a Disk	58
Viewing Aggregates Configured in a NAS Filer	59
Viewing Aggregate Properties	59
Viewing Filer Volumes Created by an Aggregate	59
Viewing Antecedents and Dependents of an Aggregate	59
Viewing Plexes Used by an Aggregate	60
Viewing RAID Groups Consumed by an Aggregate	60
Viewing Filer Volumes Configured in a NAS Filer	60
Viewing Filer Volume Properties	60
Viewing Antecedents and Dependents of a Volume for a NAS Filer	60
Viewing Qtrees Created in a Volume	61
Viewing Quota Created at the Volume or Qtree	61
Viewing Exports for a Volume	61
Oracle Database Assets	62
Database Assets and Relationships	62
Frequently Asked Questions	64
Viewing Databases	65
Viewing Database Storage Configurations	65
Viewing Properties of the Oracle Instance	65
Viewing Tablespaces Configured in the Oracle Instance	65
Viewing Tablespace Properties	66

Viewing Tablespace Antecedents and Dependents	66
Viewing Storage Supply Chain Connectivity of the Tablespace	66
Viewing Redo Logs	67
Viewing Redo Log Properties	67
Viewing Redologs Antecedents and Dependents	67
Viewing Redologs Connectivity	67
Viewing Datafiles	68
Viewing Datafile Properties	68
Viewing Datafile Antecedents and Dependents	68
Viewing Datafile Connectivity	68
Viewing Servers Hosting a Database	69
Viewing Utilized Fabrics	69
Viewing Utilized Switches	69
Viewing Utilized Targets	69
Viewing the Database History Log	70
Device Group Assets	70
Static Device Groups	70
Dynamic Device Groups	71
Device History Log	72
3 Storage Compliance	73
Overview of Storage Compliance	73
SA Client Integration	73
Compliance View	74
Audit Rules and Parameters	74
Audit Browser	75
Select Server	76
Audit Schedule	76
4 Storage Visualization	77
Overview of Service Automation Visualizer	77
Storage Mapping Features	77
5 Storage Reports	79
Overview of Storage Reports	79
Service Automation Reporter Client Integration	79
Database Storage Reports	80
Host and Application Storage Reports	81
Storage Array Reports	82
6 Storage Search	83
Overview of Storage Search	83
SA Client Integration	83
Storage Objects and Attributes	83
A Global File System (OGFS)	85
Overview of OGFS	85
Global Shell	85

Permissions	85
Storage Directory in the OGFS	86
/opsw/api/com/opsware/storage Directory	86
/opsw/Storage Directory	86
/opsw/Server/.../Storage Directory	86
/opsw/Server/.../Storage/Controller Directory	87
/opsw/.../Disk/Dependent Directory	87
/opsw/.../Server/Storage/Fabric Directory	87
/opsw/Server/.../Storage/FileSystem Directory	88
/opsw/Server/.../Storage/Port Directory	88
/opsw/Server/.../Storage/Service Directory	88
/opsw/Server/.../Storage/Switch Directory	88
/opsw/Server/.../Storage/System Directory	88
/opsw/Server/.../Storage/Volume Directory	89
Index	91

1 Introduction

Overview of Storage Visibility and Automation

Storage Visibility and Automation is a feature in Server Automation (SA) that offers storage management capabilities designed for server administrators by enabling end-to-end visibility and management of the entire storage supply chain. This feature is commonly known as *storage enablement in SA* and is referenced as such throughout this document.

This feature helps server administrators by providing tools that increase cost savings through application storage, dependency and visibility, storage audits, storage capacity and utilization trending, and scripting and automation.

This document describes how to use storage enablement in SA to collect, visualize, search for, report on, and troubleshoot storage objects, their attributes, and devices in your IT environment. This information is intended for a server administrator who is familiar with the storage supply chain in your environment.

A server administrator (also known as a server operator) manages and maintains servers; in particular, a server administrator performs server-related tasks, such as:

- Setting up servers to support new applications, including requesting or utilizing new storage
- Adding a new server to an existing application or application cluster and attaching to an existing storage resource or newly allocated storage resources
- Requesting or adding new storage to an existing application (In some IT environments, this might be a database administrator's responsibility.)
- Upgrading software and patching, including installation and configuration of HBA and associated drivers, and so on
- Troubleshooting server and application availability or performance problems

Ensuring compliance of existing servers with existing standards or newly defined standards



See the *HP Server Automation: Storage Visibility and Automation Release Notes* for a list of supported storage device hardware and software from various vendors.

Basic Business Needs

The following sections describe how Storage Visibility and Automation provides several key benefits that add business value.

Capacity and Utilization Reporting

Storage Visibility and Automation offers reporting of storage capacity and utilization that is intended to create sizeable cost savings by better utilization of current resources and better forecasting of future needs. This feature allows capacity and utilization information to be explored at any level in the storage stack, including per-application utilization, departmental utilization, database utilization, utilization of specific storage resource classes, and so on. Storage Visibility and Automation also provides the historical dimension by tracking and reporting historical utilization, as well as forecasting future storage needs that are based on past trends.

IT environments typically achieve less, if not far less, than 50% utilization of available storage resources. There are several reasons for this. Most IT departments still do much of their SAN configuration tracking manually through the use of LUN map spreadsheets and the like. Inevitably, mistakes are made and information becomes out of date, leading to orphaned storage. There are so many levels in the storage supply chain where storage can be underutilized, such as SAN, volume manager, file system, database, and so on. Even relatively minor underutilization at each of these layers can result in overall utilization of less than 25%.

Storage Visibility and Automation helps organizations reach 80%-90% utilization by:

- Providing the missing visibility into whether and how storage resources are being used to allow reclamation of unused volumes.
- Allowing storage administrators to identify the true utilization through all levels in the supply chain for a given server, database, or application. The storage administrator can then maintain an appropriate reserve at the right level in the infrastructure and reclaim unused storage.
- Providing the reporting and visibility necessary to better align end-user storage consumption behavior with business goals and storage costs. The storage administrator can then set up charge-backs, provide executive-level visibility into storage utilization and costs by internal customer, department or application, and so on.

Operational Insight

Storage Visibility and Automation provides storage administrators and server administrators with deep operational insight into the entire storage supply chain. This insight begins with the automatic discovery of all relevant storage resources, whether it is software, server, network-attached, or SAN resources. From a storage resource, the administrator has the ability to view and report on the storage inventory, as well as viewing and navigating dependencies throughout the storage infrastructure. Using Service Automation Visualizer (SAV), the storage supply chain can be depicted and navigated graphically as well.

The following cost savings result from this rich operational insight:

- Increases IT efficiency by streamlining troubleshooting operations, providing an accurate baseline for planning provisioning and audit activities, identifying un-utilized or underutilized resources for reclamation, highlighting configuration change in the storage infrastructure, facilitating training, and increasing situational awareness.
- Understands business applications and how these applications are built up from a variety of constituent IT resources. This understanding allows Storage Visibility and Automation to report on how storage resources are used to service specific applications and business services. This, in turn, allows you to identify storage resource conflicts and contention between applications.

- Allows you to better understand the storage configuration holistically, across IT silos. When permitted, server administrators will have limited visibility into how the servers and applications they own consume SAN resources. Storage administrators are provided limited visibility into how their storage resources are consumed by servers and applications. This holistic picture of the storage environment improves communication between teams in IT, and reduces friction and delay when provisioning storage or troubleshooting storage issues.



Storage information is protected by AAA (authentication, authorization and accounting) permissions. This ensures that sensitive storage data can only be accessed by staff members who have been granted access permissions.

Infrastructure Management

In most IT environments, no single user has a complete overview of the entire storage infrastructure. This is because much record keeping is manual, and because there are many point tools and command-line interfaces used to manage pieces of the infrastructure without being able to see or visualize the whole.

This situation has the following significant implications for an IT organization:

- It makes it difficult and costly to perform internal technology audits of the storage configuration. These audits are designed to ensure that the storage infrastructure is configured in a standardized manner, that best practices are followed, and that misconfigurations have not cropped up. These audits are necessary to ensure the integrity and reliability of the storage infrastructure. Because these audits can be extremely costly to perform, they are typically done in an ad-hoc, as-time-allows fashion, resulting in increased risk.
- The lack of visibility into the storage infrastructure and how specific storage resources are used by business applications make it extremely difficult and costly to confirm and document compliance. In the model regulatory climate, organizations are required to guarantee and document that their IT environment is secure and robust. Organizations need to be able to show that specific business applications are configured to meet defined internal standards. Without a clear understanding of what storage resources these business applications are using and visibility into the compliance status of these storage resources, the process of documenting compliance of the storage infrastructure becomes costly and confusing.

Storage Visibility and Automation fills this gap by providing the single, automated source for storage infrastructure audits and compliance documentation. This feature automatically analyzes the storage infrastructure and dependencies to ensure that a wide variety of storage best practices are properly implemented and that known misconfigurations have been prevented. This information is rolled up into the Compliance View at the device level or at the application level. The Compliance View provides a single clearinghouse that helps identify potential risks in the environment, address these risks, and document the bill of health of the storage environment and of specific critical applications.

Feature Capabilities

Table 1 describes the capabilities and benefits of operational insight, reporting, audit and compliance, scripting and automation, enterprise features, and administrative features.

Table 1 Storage Visibility and Automation Capabilities

Feature	Description	Benefit
Operational Insight		
Discovery	<ul style="list-style-type: none"> Discover servers using the Agent Deployment feature Discover SAN resources using appropriate management station Discover storage inventory details of hosts, SAN resources, and NAS resources Discover inventory changes over time Discover dependencies between storage resources throughout the storage supply chain 	<ul style="list-style-type: none"> Quick startup with minimal overhead Maintain detailed up-to-date inventory of all storage resources and configurations
Device Explorer	<ul style="list-style-type: none"> Explore server storage properties: <ul style="list-style-type: none"> Volume management software Multipath I/O software Disks File systems Mounted and available LUNs Logical and physical volumes Fibre Channel Adapter (HBA) Explore and navigate server storage dependencies: <ul style="list-style-type: none"> Attached or downstream SAN switches SAN arrays and mounted volumes NAS filers and mounted file systems Resident databases and applications 	<ul style="list-style-type: none"> Identify and troubleshoot server storage configuration Identify and navigate downstream storage dependencies Provide information to facilitate efficient provisioning and modification of storage resources Identify available and unused storage resources

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
SAN Array Browser	<ul style="list-style-type: none"> • Support for SAN arrays and NAS filers • Explore array properties: <ul style="list-style-type: none"> Volumes and RAID groups Pools Fibre Controllers and Ports • Explore and navigate array dependencies: <ul style="list-style-type: none"> SAN fabrics Dependent servers, applications and databases Attached and upstream storage switches • Support for vendor proprietary volume technologies: <ul style="list-style-type: none"> NAS: aggregates, exports, ... SAN: meta, BCVs, LUSEs, ... • Report and navigate volume mirroring and replication relationships 	<ul style="list-style-type: none"> • Identify and troubleshoot array / filer configuration • Effectively manage existing and new storage resources • Identify and navigate upstream storage dependencies • Provide information to facilitate efficient provisioning and modification of storage resources • Identify available / unused storage resources
Database Browser	<ul style="list-style-type: none"> • Explore database properties: <ul style="list-style-type: none"> Tablespaces Data files Configuration files • Explore and navigate database dependencies: <ul style="list-style-type: none"> Dependent applications Host servers Downstream storage resources and arrays 	<ul style="list-style-type: none"> • Identify and troubleshoot database configuration • Identify and navigate downstream storage dependencies • Identify available and unused storage resources • Identify how database fits into business application delivery
Visualization	<ul style="list-style-type: none"> • Dynamic visualization of the storage infrastructure devices and resources in the context of the broader IT and business-application environment: <ul style="list-style-type: none"> Expand or collapse the level of detail Dependency tree navigation • Logical, physical, and hybrid storage views 	<ul style="list-style-type: none"> • Trace and troubleshoot application dependencies • Build situational awareness of the current storage infrastructure configuration • Visually identify upstream and downstream dependencies, contention and conflicts • Document and visually compare storage infrastructure changes over time

Reporting

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
Search-Based Reporting	<ul style="list-style-type: none">• Create ad-hoc tabular reports using a simple search-based user interface:• Save reports• Export or email the results• Run reports periodically	<ul style="list-style-type: none">• Quick access to data needed for daily operational tasks• Automate both ad-hoc and recurring asset, inventory and storage configuration reports that often require manual compilation today
Inventory Reports	<ul style="list-style-type: none">• System offers built-in inventory reports on the storage infrastructure as a whole and specific elements within it• Host reports:• Array / filer reports:• Report at the host/array level or break out information by LUN• Export or email the report output	<ul style="list-style-type: none">• Simple, at-a-glance view of the storage infrastructure• Facilitates communication and troubleshooting between groups• Make data in the system available for non-users

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
Capacity & Utilization Reports	<ul style="list-style-type: none"> • System offers built-in capacity and utilization reports • Organizational reports: <ul style="list-style-type: none"> Application Department/Business Unit Storage Detail Site Storage Detail Device Group Storage Detail (applies to any arbitrary device group) • Host reports: <ul style="list-style-type: none"> Host Capacity & Utilization Host Capacity & Utilization Trend Host Capacity & Utilization Trend Detail • Array reports: <ul style="list-style-type: none"> Array Capacity & Utilization Array Capacity & Utilization Trend Array Capacity & Utilization Trend Detail • Database reports: <ul style="list-style-type: none"> Database Allocation Trend Database Capacity and Utilization Trend Database Utilization Trend Tablespace Allocation Trend Tablespace Capacity and Utilization Overview Tablespace Capacity and Utilization Trend Data Tablespace Utilization Trend • Filter or group data by Application, Department, Site or any Device Group • Timeseries reports show past information • Trend reports show future projections 	<ul style="list-style-type: none"> • Identify opportunities for reclamation of un-used or underutilized storage • Map out storage utilization for upper management • Implement charge-backs • Encourage better utilization of assigned storage resources • Analyze historical storage trends • Forecast future storage needs

Audit & Compliance

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
Compliance View	<ul style="list-style-type: none"> • Automatically audit the storage infrastructure to determine and report host compliance with key best practices, such as: <ul style="list-style-type: none"> Proper multipathing at all levels in the storage supply chain All LUNs on the SAN that are mapped to specific hosts should be properly mounted by those hosts • Configure compliance metrics to match organizational policies • Integrated with the compliance dashboard: <ul style="list-style-type: none"> Feeds into unified audit of overall application compliance 	<ul style="list-style-type: none"> • Reduce risk • Increase standardization and adoption of best practices • Reclaim resources orphaned due to misconfiguration
Scripting & Automation		
Global Shell File System (OGFS)	<ul style="list-style-type: none"> • Standard shell interface into storage inventory • Navigate all storage entities in a file system-like presentation: <ul style="list-style-type: none"> Arrays and filers Storage switches Servers Databases • Ideal for scripting • Supports all common scripting languages, such as Perl, TCL, Python, shell-script, and so on 	<ul style="list-style-type: none"> • Quickly and easily create scripts for: <ul style="list-style-type: none"> Custom reports Custom audits Ad-hoc exploration of storage information
Unified API	<ul style="list-style-type: none"> • Unified API provides full programmatic access to all data and features: <ul style="list-style-type: none"> Supports web-services programming interfaces Provides a single API for datacenter automation across storage, servers, and IP networks 	<ul style="list-style-type: none"> • Extend and customize Storage Visibility and Automation • Integrate with external systems, such as CMDB, change request, ticketing, monitoring, and so on
Enterprise Features		

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
Multimaster	<ul style="list-style-type: none"> • Connect multiple Storage Visibility and Automation instances together, spanning multiple data centers • Each Storage Visibility and Automation instance has a complete view of the entire storage infrastructure • Each Storage Visibility and Automation instance can manage local resources (with failover when an instance becomes unavailable) • Data is synchronized across instances automatically 	<ul style="list-style-type: none"> • Built-in high availability and disaster recovery • Storage Visibility and Automation scales seamlessly, both within a single datacenter and across the global IT network • Users in each datacenter have a local UI, improving UI responsiveness
Gateway	<ul style="list-style-type: none"> • Secure tunnel between system components: <ul style="list-style-type: none"> — Enable management of overlapping IP networks — Enable management across bastion hosts — Providing redundant routing between system components — Allow public networks to be used for management without security concerns 	<ul style="list-style-type: none"> • Ensure system security in distributed deployments • Manage complex network topologies seamlessly • Support fault tolerant distributed Storage Visibility and Automation deployments
AAA (Authentication, Authorization, & Accounting)	<ul style="list-style-type: none"> • User authentication: <ul style="list-style-type: none"> — Use built-in database or leverage external authentication mechanisms (such as LDAP) • Role-based user authorization: <ul style="list-style-type: none"> — Control access by server or storage device — Control access by user role to include some, all or no server/application visibility — Control access by user role to include some, all or no storage infrastructure visibility 	<ul style="list-style-type: none"> • Prevent unauthorized or improper use of the system; ensure data integrity • Ensure proper levels of separation between server and storage teams • However, facilitate an appropriate level of visibility between teams where such visibility is lacking today, which will reduce organizational friction and delay • Ensure sub-groups within server and storage teams only have access to the devices appropriate to their job function
Administrative Features		

Table 1 Storage Visibility and Automation Capabilities (cont'd)

Feature	Description	Benefit
Device Groups	<ul style="list-style-type: none"> • Create device groups containing any combination of storage arrays and servers: <ul style="list-style-type: none"> — Static groups — Dynamic groups computed automatically based on device properties — Hierarchical groups 	<ul style="list-style-type: none"> • Powerful array of grouping mechanisms allow devices to be organized to align with corporate structure, IT processes, physical device attributes, and so on • Dynamic grouping ensures groups are always up to date with no maintenance required • Groups allow operations and reports to be performed on large groups of devices easily and accurately
Server Discovery	<ul style="list-style-type: none"> • Server discovery automatically discovers and classifies all servers on the network in preparation for automated agent deployment 	<ul style="list-style-type: none"> • Streamline Storage Visibility and Automation rollout • Ensure accurate inventory information

Icons Used in Storage Visibility and Automation

Table 2 identifies icons used in the Storage Visibility and Automation user interface. See [Asset Discovery](#) on page 23 for expanded descriptions of these icons. Text in parenthesis represents the terminology used in this documentation.

Table 2 Icons Used in Storage Visibility and Automation











ICON	DESCRIPTION
	Databases (Oracle Database Assets)
	Data Files
	Device Group Assets
	Ethernet Ports
	Fibre Channel Ports (Fibre Channel Adapter Ports)

Table 2 Icons Used in Storage Visibility and Automation (cont'd)

ICON	DESCRIPTION
	File Systems
	Fibre Channel Adapters
	NAS Filer Aggregates (Aggregates)
	NAS Filers (NetApp Assets)
	SAN Arrays (SAN Array Assets)
	SAN Fabrics (Fabric Assets)
	SAN Fabric Zones (Zoning Database and Zones)
	SAN Fabric Zone Sets (Zonesets)
	SAN Switches (Fibre Switch and Fabrics)
	Server Assets
	Static Device Groups
	Storage (SAN Array Storage Systems)
	Storage Disks (Disks)

Table 2 Icons Used in Storage Visibility and Automation (cont'd)

ICON	DESCRIPTION
	Storage Initiators
	Storage Pools
	Storage Targets
	Tablespaces
	Volumes (Storage Volumes)

2 Asset Discovery

Overview of Asset Discovery

Storage Visibility and Automation collects storage information from Storage Essentials (SE) discovery and enables storage information in Server Automation (SA).

- Host server storage information is discovered by the Storage Host Agent Extension in Storage Visibility and Automation—and not by SE discovery.
- Database storage information is discovered by a Storage Scanner in Storage Visibility and Automation—and not by SE discovery.

SE discovers information about the SAN arrays, switches, fabrics, and NetApp filers in your environment. SA obtains information about these storage assets from SE. Storage assets that are enabled in SA include SAN arrays, fabrics, controllers, and zones. Storage Visibility and Automation also provides information about connectivity, such as the relationship of a switch and a managed server. Edge information is also provided; however, not all fabric switches are displayed.

Storage Visibility and Automation also provides visibility into the storage configurations in a host virtualization environment, such as a VMware ESX servers and their associated virtual machines.

Data Collection

Storage information is discovered in SE and then collected by Storage Scanners in Storage Visibility and Automation.

Information about storage arrays, switches, fabrics, and NetApp filers in your environment is collected by a component in SA called *SE Connector*. SE Connector retrieves data about the SAN infrastructure from SE. This storage data is transferred to the SA core and stored in the Model Repository. See [SE Connector](#) on page 24 for more information.

Information about Oracle storage configurations in a SAN or Network Attached Storage (NAS) is collected by a component in Storage Visibility and Automation called the *Database Scanner for Oracle*. This Storage Scanner retrieves data about Oracle instance, tablespaces, and datafiles. See [Database Scanner for Oracle](#) on page 24 for more information.

Storage Scanners

There are two types of Storage Scanners that collect storage data:

- [SE Connector](#)
- [Database Scanner for Oracle](#)

SE Connector



Storage Essentials (SE) version 6.1.1 or later and Server Automation 7.80 or later are required to configure and deploy SE Connector. SE is required to view, report, or perform any Service Automation Visualizer (SAV) and Service Automation Reporter (SAR) operation on SAN objects, such as arrays, switches, volumes, and so on. SAN objects are discovered in Storage Essentials. To enable discovered SAN objects in the SA, SAV, and SAR products, the Server Automation SE Connector component must be installed and configured. See the *HP Server Automation: Storage Visibility and Automation Installation & Administration Guide* for information about configuring and deploying SE Connector.

SE Connector collects data from Storage Essentials (SE) about SAN elements and inventory, and their connectivity. These SAN elements include storage arrays, fabrics, switches, and NAS filers.



In the user interface, SE Connector is referred to as *SE Storage Scanner*.

The frequency of collecting this storage data is configured by the user. By default, SE Connector collects SAN information from SE once every 24 hours. This action is commonly known as *a scheduled full synchronization*. You can adjust the collection frequency by modifying a setting in SE. In SE Connector, you can specify the maximum wait time (in minutes) for a scheduled full synchronization when the Get All Elements Detail (GAED) process in SE is in progress. The following property specifies the maximum wait time for a scheduled full synchronization task when GAED is in progress:

```
com.creeppath.agent.common.devices.scheduled.full.sync.max.wait.minutes=60
```

SE Connector can communicate with multiple Storage Essentials instances.



For scalability and performance reasons, it is recommended that you have *multiple* SE Connectors communicating with *one* instance of Storage Essentials. You can also have *multiple* SE Connectors communicating with *multiple* instances Storage Essentials.



When you have *multiple* SE Connectors communicating with *one* instance of Storage Essentials, each access control must be configured to collect a unique set of devices from the same Storage Essentials instance. See the *HP Server Automation: Storage Visibility and Automation Installation & Administration Guide* for information about modifying access controls.

Since, by default, storage information is collected from SE once every 24 hours, you might need updated data more frequently than once a day. You can modify this data collection frequency by modifying a setting in SE. See the *Storage Essentials SRM Software User Guide* for information about configuring the schedule for data collection. You can also request an update of SAN information from SE at the storage device level, such as arrays and NetApp filers. See [Update from Storage Essentials](#) on page 25 for more information.

When you are viewing storage information about a SAN array or NetApp filer in Storage Visibility and Automation and would also like to view storage information for the same SAN array or NetApp filer in Storage Essentials, you can launch Storage Essentials to display the data that was discovered. See [Open with Storage Essentials \(Click & Launch\)](#) on page 26 for more information.

Database Scanner for Oracle

The Database Scanner for Oracle collects data about Oracle storage configurations in a SAN or Network Attached Storage (NAS), such as the Oracle instance, tablespaces, and datafiles.

Storage elements configured for an Oracle instance can be classified as *physical database storage* and *logical database storage*:

- **Physical database storage** includes datafiles and redo logs that directly consume system storage resources (filesystems or partitions) or are built on top of ASM Files.
- **Logical database storage** includes entities such as tablespaces that are created inside the instance consuming different physical and logical storage entities.

This Storage Scanner identifies relationships between the database elements and other storage assets as described in [Table 1](#).

Table 1 Database Assets & SAN Relationships

Database Asset	External Storage Asset	Dependency	Description
SAN-based physical database storage	Server assets SAN array assets Fabric assets NetApp assets	Block storage dependency	Provides the dependency chain between the database storage elements and SAN arrays through system (server) resources and fabrics.
NAS-based physical database storage	Server assets NetApp assets	NAS storage dependency	Provides the dependency chain between the database storage elements, server resources, and NetApp.

See the *HP Server Automation: Storage Visibility and Automation Installation & Administration Guide* for information about how to configure and deploy the Database Scanner for Oracle.

Update from Storage Essentials



You must have the required SA user permissions to request an “Update from Storage Essentials”. See the *HP Server Automation: Storage Visibility and Automation Installation & Administration Guide* for more information.

“Update from Storage Essentials” is an option in the Actions menu that requests an update of storage information about an individual storage device, such as a SAN array or NetApp filer. SE Connector provides this synchronization of data for a storage system when you request an update at the device level.

SE Connector does not directly communicate with a storage device to collect data about it. SE Connector collects data about an individual storage device from Storage Essentials. This request copies data from the SE repository into the Model Repository.

You can specify the maximum wait time (in minutes) for the “Update from Storage Essentials” action in SE Connector while SE performs GAED for certain storage elements. The following property specifies the maximum wait time for the "Update From Storage Essentials" action when GAED is in progress:

```
com.creeppath.agent.common.devices.manual.full.sync.max.wait.minutes=60
```

The Jobs and Sessions user interface displays the status of the update request. A “Success” status indicates that the Model Repository has been updated. The timestamp for the job also provides the following information about the “Update from Storage Essentials” request:

- When the request was made
- When SE discovered the data
- What the status of the update job is, such as:
 - **OK**—Request was submitted successfully. SE Connector will collect the data for the requested device and send it to the Web Services Data Access Engine.

- **Error**—Request could not be submitted. Check the SE Connector logs on the managed server to troubleshoot this error. Begin with the `pam-0-0.trace` log file and then check other trace files.

See the *SA User Guide: Server Automation* for more information about Jobs and Sessions.

Open with Storage Essentials (Click & Launch)

“Open with Storage Essentials” is an option in the Actions menu that launches Storage Essentials for an individual storage device, such as a SAN array or NetApp filer. You can also launch Storage Essentials for a device by selecting it and then by right-clicking—this is commonly referred to as the *Click & Launch* user interface.

SE Connector supports Click & Launch for storage devices that are managed by Storage Essentials. This functionality requires an SE account login. When you select other storage objects and then Click & Launch from them (while in the same session), the initial SE login will persist. There is no single sign-on capability for Storage Visibility and Automation and SE.



You can launch SE only from a storage device. You cannot launch SE from a host.

Storage Device

A storage device in the Storage Visibility and Automation model is an entity that represents the volumes, extents, and other things in the real world that contain the actual bits and bytes. A storage device can be any block storage device that has a capacity measured in bytes, such as physical disks, partitions, and database tablespaces.

There are two types of storage devices:

- A **logical volume** is a block storage extent, such as partition, database tablespace, and remote initiator volume.
- A **physical disk** is a physical storage device that provides the raw storage for a storage system, such as storage array disks or spindles.

Deleting a Storage Device

This action only removes a storage device (SAN array, NAS filer, and database) from the Model Repository.

To remove a storage device from the Model Repository, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** and then select a device type, such as **SAN Arrays** or **NAS Filers**.

Or

From the Navigation pane, select, **Library** ► **Databases**.

- 2 In the content pane, select a device. You can also select more than one device.
- 3 From the Actions menu, select Delete.
- 4 In the Delete dialog box, click **Delete** to remove the device from the Model Repository.


Storage Host Agent Extension (SHA)

A Storage Host Agent Extension (SHA) runs on a managed server to discover storage related data on that managed server, such as Fibre Channel Adapter (FCA) hardware and drivers, filesystem mounts, the storage available on an ESX Server and the storage allocated to its virtual machines, and so on.

SHA is not an agent. SHA is a C++/Python module that can be executed by scheduled jobs. This module runs on managed servers and produces XML output that is sent back to the SA core. This data (XML output) persists in the Model Repository from where it is then retrieved and displayed in the user interface.

Figure 1 shows the end-to-end process of the SHA that captures storage-related data on a managed server.

Figure 1 SHA on a Managed Server



VMware ESX & ESXi Support

This feature provides visibility into the storage configurations in a host virtualization environment, such as a VMware ESX server attached to SAN. Storage Visibility and Automation provides storage dependency data for both the VMware ESX hypervisors and the guest hosts on the hypervisors. This functionality supports the following use cases:

- **Hypervisor View**—The hypervisor is treated as a server (host). The VMware ESX administrator will want to visualize where the storage is coming from and which guests are consuming which storage.
- **Guest VM View**: Each guest VM is treated as a server (host). The user needs visibility into the underlying storage for each guest host. The user will be able to see the association of the guest VM to the hypervisor.

See the *SA User Guide: Server Automation* for more information about hypervisors and VMware ESX virtual machines.

Server Assets

This feature provides visibility into the storage configuration in a managed server attached to the Storage Area Network (SAN) or Network Attached Storage (NAS). These configurations can help you understand the dependency on SAN and NAS elements.

Storage elements in a server (server assets) include logical entities, such as volumes and physical hardware. This feature allows you to view and evaluate storage-related information for a managed server, such as FCAs and FCA drivers, volume mounts, volume manager information, and so on.

The discovery of server storage elements is enabled by processes that periodically query the managed server and then push the data to the core.



Storage Visibility and Automation displays server storage assets after the hardware registration process is successfully completed. Hardware registration typically occurs within the first 24 hours of the SE Connector deployment. See the *SA User Guide: Server Automation* for information about how to initiate the hardware registration process as needed.

You can define the job schedule for capturing storage data. See [Scheduling Server Asset Discovery](#) on page 29.

With this feature, a server administrator can view:



- Storage-related information in each managed server, such as FCAs and FCA drivers, volume mounts, volume manager information, and so on
- Unmounted storage resources that the SAN is making available to the host
- Aggregate information on storage elements (such as volumes used) that are associated with an application in a server


This feature also provides quick navigation to downstream physical or logical storage elements associated with an application or a single managed server. At any level in the storage supply chain, you can drill down to the associated storage entities to view:

- **Physical entities**—fibre switches, filers, disks, and so on
- **Logical entities**—zones, LUNs, and so on

From these entities, you can drill down or drill up to see storage resources that are shared between two or more selected applications or managed servers.

Storage Visibility and Automation captures the following server assets:

- **Fibre Channel Adapter (FCA)**  —Storage information on the active Fibre Channel Adapter (FCA). In the storage industry, FCA also applies to the hardware interface that uses non-Fabric Channel protocol. Storage Visibility and Automation discovers the Manufacturer, Model, Serial Number, Software Driver Version, Software Firmware Version, and Physical Hardware Version for an FCA.
- **Fibre Channel Adapter Port**  —The port on the FCA to provide access to the outside world—storage information on the active Fibre Channel Port World Wide Name (WWN). See [Fabric Assets](#) on page 51 for information about storage in the connectivity of an FCA Port and an FC Switch Port.
 - **Base Volumes**—Basic raw storage capacity (primordial storage). This raw storage may be virtualized and consumed in the layers above (Volume Managers, File Systems). These storage elements include local storage (Root Volumes) based disks, and Direct Attached Storage (DAS).
 - **Root Volumes (Local and DAS)**—Storage exposed through disks that are local and are based on DAS as logical Root Volume. Hardware information of the spindles or disk controllers is not captured.

- **Multipath IO (MPIO)**—Information on the MPIO software and all of the Fibre Channel protocol based storage access paths to remote storage volumes. This information includes access path information for each of the LUN bounded to the server. For each path, Storage Visibility and Automation captures the following information:
 - State of the path—active or passive
 - Is Preferred Path
 - FCA Port consumed for each of the path
- **Volume Manager (Logical Volume Manager/Logical Disk Manager)**—Native and third party Volume Manager Software information. This includes information on all aspects of storage, such as capacity and storage layout, that are consumed and produced by the Volume Manager. Storage Visibility and Automation does not provide detailed information on internals of the Volume Manager. However, for each of the volumes produced by the Volume Manager, Storage Visibility and Automation provides dependencies on primordial storage, such as dependency chain that includes consumed storage. SE does not support layered volume managers and layered volumes.
- **Consumed Storage**—Storage that is directly consumed by the Volume Manager from the server primordial storage. SE discovers storage capacity if the applicable volume group (Disk Group) is the source of the consumable storage (server level primordial storage only). Storage Visibility and Automation also captures storage consumed by a volume manager, which consumes external storage. Not all storage resources are virtualized to create volume manager volumes.
- **Produced Storage**—Storage capacity when the applicable volume group (Disk Group) is the type (SIMPLE, RAID5, and so on) for the produced storage (volumes). Produced storage consists of storage units created by a hosted service, such as file systems, volume managers, and databases.
- **File Systems**  —Consumers of the produced storage. This functionality includes capturing the file systems that are based on primordial storage and storage offered by volume manager volumes. Storage Visibility and Automation discovers Block Driver, Cluster Size, Is Remote, File System Size, File System Type (nfs, ufs, vxfs, hfs, ext2, ext3, ntfs, and so on), Mount Point, and Root and Storage capacity information. Storage Visibility and Automation also discovers the dependency chain for the File System based on NetApp. See [Storage Visibility and Automation displays relationships between fabric elements and other storage asset discovery features](#). See Table 6. on page 52 for more information. For this release, Storage Visibility and Automation does not support distributed filesystem.

Scheduling Server Asset Discovery

In Storage Visibility and Automation, the discovery of server storage elements is enabled by processes (jobs) that periodically query the managed server and then push the data to the core. You can define the job schedule for capturing storage data. See the *SA User Guide: Server Automation* for more information about jobs and scheduling them.

Server Storage and SAN Relationships

Storage Visibility and Automation displays relationships between server storage elements and other storage asset discovery features. See [Table 2](#).

Table 2 Server Assets and SAN Relationships

Server Asset	External Element	Description
File System	Remote File Systems (only NFS mounted file system dependency)	Captures only NFS-mounted file system dependency information.
FCA Port	Connected FC switch port	The port on the FCA to provide access to the outside world—storage information on the active Fibre Channel Port World Wide Name (WWN).



Storage Visibility and Automation does not include relationships that are local to the discovery of a server storage asset.

Frequently Asked Questions

Table 3 shows you how to use this documentation to find answers to some frequently asked questions about finding information about server storage elements in a managed server.

Table 3 Frequently Asked Questions About Server Assets

	Question	How to Find the Answer
1.	What are the application storage configurations in a managed server?	<ul style="list-style-type: none"> • Viewing File Systems Mounted on a Server on page 33 • Viewing Databases Hosted by a Server on page 40
2.	What is causing application performance degradation? How is all of the server storage configuration impacting the application? What are the shared resources between applications?	<ul style="list-style-type: none"> • Viewing Ports and Volumes in an FCA on page 32 • Viewing Port Connections and Zones on page 33 • Viewing Access Paths to Storage Targets on page 38 • Viewing Fabrics Attached to a Server on page 39 • Viewing Storage Volumes with Access Paths for Remote Initiators on page 46 • Viewing the Access Path from the Volume on page 49
3.	What SAN elements (switches, arrays, and gateways) are applications, through server storage assets, dependent on?	<ul style="list-style-type: none"> • Viewing Antecedents and Dependents of Volumes for a Managed Server on page 34 • Viewing File Systems Mounted on a Server on page 33 • Viewing Volume Details on page 33
4.	What are the storage expansion opportunities for an application and for a managed server?	<ul style="list-style-type: none"> • Viewing Storage Summary in a Server on page 33 • Viewing File Systems Mounted on a Server on page 33 • Viewing Volumes Created on a Disk on page 36 • Viewing Unmounted Volumes Through the Hypervisor on page 37
5.	Are critical applications configured for high availability, such as multiple paths?	<ul style="list-style-type: none"> • Viewing Ports and Volumes in an FCA on page 32
6.	What are the unmounted storage resources that the SAN is making available to the managed server?	<ul style="list-style-type: none"> • Viewing Unmounted Volumes Through the Hypervisor on page 37 • Viewing the Access Path from the Volume on page 49

Table 3 Frequently Asked Questions About Server Assets (cont'd)

	Question	How to Find the Answer
7.	What is the aggregate information on storage elements (volumes used) associated with an application in a server?	<ul style="list-style-type: none"> • Viewing Antecedent of Aggregate Storage Supply Chain on page 38
8.	How do I navigate to downstream physical or logical storage elements associated with an application or a single managed server?	<ul style="list-style-type: none"> • Viewing Antecedents and Dependents of Volumes for a Managed Server on page 34
9.	How do I navigate to the associated storage entities at any level in the storage supply chain, such as: <ul style="list-style-type: none"> • Physical entities—fibre switches, filers, disks, and so on • Logical entities—zones, LUNs, and so on? 	<ul style="list-style-type: none"> • Viewing Ports and Volumes in an FCA on page 32 • Viewing Port Connections and Zones on page 33
10.	How can I view storage resources that are shared between two or more selected applications or servers?	<ul style="list-style-type: none"> • Viewing Storage Volumes with Access Paths for Remote Initiators on page 46 • Viewing the NAS Filer Storage Configuration on page 55

Viewing FCAs Attached to a Server

To troubleshoot fibre channel connectivity and identify the inventory, operational status, and storage resources consumed through a Fibre Channel Adapter, you need to know what FCAs are attached to a certain managed server.

To view FCAs that are attached to a server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server.
- 3 In the View drop-down list, select Hardware. The Hardware pane displays information about the server and the FCAs that are attached to it, such as Manufacturer, Model, Processors, Memory (RAM), and Memory (SWAP). This information is captured by Storage Host Agent Extensions (SHAs).

Viewing Ports and Volumes in an FCA

To view which ports and volumes are in an FCA, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Navigation pane, select **Inventory**.
- 4 In the Views pane, expand the **Hardware** list and then select the Fibre Channel Adapter (FCA) name. See [Server Assets](#) on page 28 for more information.

Viewing Port Connections and Zones

To view switch ports that a port is connected to and zones that a port belongs to, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Navigation pane, select **Inventory**.
- 4 In the Views pane, expand the **Hardware** list and then select the fibre channel port.

Viewing Storage Summary in a Server

To view storage summary in a managed server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage**.

Viewing File Systems Mounted on a Server

To view file systems mounted on a server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **File Systems**. The content pane displays information about all file systems that are mounted on the selected server.



For VMware ESX and ESXi servers, file systems data is not provided.

Viewing File System Properties

To view properties of a file system, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **File Systems**.
- 4 In the content pane, select a file system to display the Properties pane. This pane displays details about the file system.

Viewing Volume Details

You can view volume details, such as properties, composition, disks, access path, and connectivity by volume.

To view volume details, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.

- 3 From the Views pane, select **Inventory** ► **Storage** ► **File Systems**.
- 4 In the content pane, select a file system to display the Volumes pane.
- 5 From the View drop-down list, select Volumes.
- 6 (Optional) From the View drop-down list, select an option to display properties, volumes, disks, or connectivity by volume.

Viewing Local Disks

To view local disks used by a file system, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Navigation pane, select **Inventory** ► **Storage** ► **File Systems**.
- 4 In the content pane, select a file system.
- 5 From the View drop-down list, select Disks to display the Disks pane. This pane displays information about disks that are consumed by a file system.

Viewing the Storage Supply Chain

To view the storage supply chain as a tree, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Navigation pane, select **Inventory** ► **Storage** ► **File Systems**.
- 4 In the content pane, select a file system.
- 5 From the View drop-down list, select Connectivity to display the Connectivity pane. This pane displays Storage Supply Chain for the selected file system.

Viewing Volume Properties

To view properties of a volume, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Properties to display the Properties pane. This pane displays information about the selected volume.

Viewing Antecedents and Dependents of Volumes for a Managed Server

To view antecedent and dependents of the volume storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.

- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Composition to display the Composition pane.
- 6 (Optional) In the Composition pane, right-click on a storage resource to open its corresponding SAN Array Browser.

Viewing Access Paths for a Volume

Access paths include Target and Target Port LUN based volumes. This provides visibility into the MPIO configuration of a volume.

To view a list of access paths, perform the following steps:

- 1 From the Navigation pane, select **Servers > All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Access Path to display the Access Path pane. This pane provides visibility into the MPIO configuration of a volume by displaying Target and Target LUN based volumes in the list of access paths.
- 6 (Optional) Select a path and then select one of the following options from the Actions menu:
 - **Open Target**—Displays the SAN Array Browser or the NAS Filer Browser, depending on the AAA permission.
 - **Open Target Port**—Displays the SAN Array Ports Browser or the NAS Filer Ports Browser, depending on the AAA permission.
 - **Open Target Volume**—Displays the SAN Array Volume Browser or the NAS Filer Volume Browser, depending on the AAA permission.

Viewing Volume Connectivity

To view the connectivity of a volume, perform the following steps:

- 1 From the Navigation pane, select **Servers > All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Connectivity to display the Connectivity pane. This pane displays the connectivity information for the volume.
- 6 (Optional) Select a level and then select **Actions > Open (Enter)** to display the SAN Array Browser or the NAS Filer Browser.

Viewing Disks Utilized by a Volume

To view disks utilized by the volume, perform the following steps:

- 1 From the Navigation pane, select **Servers > All Managed Servers**.

- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Disks to display the Disks pane. This pane shows the following information about the disks that are consumed by the volume:

Viewing Local Disks or DAS Attached to a Server

To view local disks or DAS attached to a managed server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Disks**.

Viewing Disk Properties

To view the properties of a selected disk, including properties of a physical disk in the hypervisor, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select Inventory and then select **Storage** ► **Disks**.
- 4 Select a disk to display the Properties pane.

Viewing Volumes Created on a Disk

To view volumes created on a selected disk, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 Select a disk and then select Volumes from the View drop-down list.

Viewing File Systems Created on a Disk

To view file systems created on a selected disk, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select Inventory and then select **Storage** ► **Disks**.
- 4 Select a disk and then select File Systems from the View drop-down list.

Viewing Volume Manager, MPIO, VMFS, & ASM

To view Manager Software on a server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Manager Software**.
- 4 In the content pane, select an inventory element, such as an Oracle Automatic Storage Management (ASM) disk group.

Viewing Volumes Available Through Manager Software

To view volumes that are available through Manager Software, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Manager Software**.
- 4 Select a Manager Software to show detailed information in the bottom content pane.

Viewing Volumes Available Through the Hypervisor

To view volumes that are available through the hypervisor, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a hypervisor (ESX server) and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 Select a volume to show detailed information in the bottom content pane, including virtualization information.

Viewing Unmounted Volumes Through the Hypervisor

To view SAN based volumes that are accessible to the managed server but are not mounted, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** and then select **Storage** ► **Unmounted Volumes**.
- 4 (Optional) Select a unmounted volume and then select **Actions** ► **Open with** ► **Remote Terminal** to open the remote device that this (remote) unmounted volume resides on.

If a hypervisor or guest volume is removed and depending on the snapshot order between the hypervisor and the guest host, the volume will display as an unmounted volume.

Viewing Unmounted Volume Properties

To view properties of unmounted volumes, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select Inventory and then select **Storage** ► **Unmounted Volumes**.
- 4 Select a volume and then select Properties from the View drop-down list.

Viewing Antecedent of Aggregate Storage Supply Chain

To view the antecedent of the aggregate storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select Inventory and then select **Storage** ► **Unmounted Volumes**.
- 4 Select a volume and then select Composition from the View drop-down list.
- 5 (Optional) From the Composition pane, select an antecedent or dependent and then open its corresponding SAN Array Browser.

Viewing Remote SAN Based Volumes

To view remote SAN based disks, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Inventory** ► **Storage** ► **Unmounted Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Disks to display the Disks pane. This pane displays information about disks that are consumed by unmounted volumes.
- 6 (Optional) If the volume is LUN based, select then open it to display its corresponding SAN Array Browser.

Viewing Storage Targets

To view storage targets (SAN Array or NAS Filer (Block)), perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **Storage Targets**.

Viewing Access Paths to Storage Targets

To view access paths to storage targets, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.

- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **Storage Targets**.
- 4 Select a target to display the access paths.
- 5 (Optional) Select an access path and then select one of the following options from the Actions menu:
 - **Open Target**—Displays the SAN Array Browser or the NAS Filer Browser, depending on the AAA permission.
 - **Open Target Port**—Displays the SAN Array Ports Browser or the NAS Filer Ports Browser, depending on the AAA permission.
 - **Open Target Volume**—Displays the SAN Array Volume Browser or the NAS Filer Volume Browser, depending on the AAA permission.

Viewing Switches Attached to a Server

To view switches attached to a managed server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **SAN Switches**.

Viewing Fabrics Attached to a Server

To view fabrics attached to a server, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **SAN Fabrics**.

Viewing FCA Ports Attached to a Fabric

To view FCA ports attached to a fabric, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** and then select **SAN Fabrics**.
- 4 From the View drop-down list, select **Ports**.

Viewing Zone Membership

To view zones where a managed server is a member of a certain fabric, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **SAN Fabrics**.
- 4 From the View drop-down list, select **Zones**.

Viewing Databases Hosted by a Server

To view databases that a selected server hosts, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **All Managed Servers**.
- 2 In the content pane, select a server and then double-click or right-click to open it.
- 3 From the Views pane, select **Relationships** ► **Databases**.
- 4 (Optional) Select a database and then open it to display the Database Browser. You must have Read permissions on the server that is hosting the database.

Viewing Hosts in the Hypervisor

To view all hosts in the ESX hypervisor, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **Virtual Servers**.
- 2 In the content pane, select an ESX server and then double-click or right-click to open it.
- 3 Select the plus (+) symbol that precedes the ESX server name to expand the display so that you can view the hosts in the ESX server (hypervisor).

Viewing Pools on the Hypervisor

To view pools (datastores) on the hypervisor and their properties, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **Virtual Servers**.
- 2 In the content pane, select an ESX server and then double-click or right-click to open it.
- 3 Select **Inventory** ► **Storage** ► **Pools**.

Viewing Pool Properties on the Hypervisor

To view pool properties on the hypervisor, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **Virtual Servers**.
- 2 In the content pane, select an ESX server and then double-click or right-click to open it.
- 3 Select **Inventory** ► **Storage** ► **Pools**.
- 4 From the View drop-down list, select Properties.
- 5 Select a pool (datastore) to view properties in the bottom content pane.

Viewing Pool Volumes on the Hypervisor

To view pool volumes on the hypervisor, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **Virtual Servers**.
- 2 In the content pane, select an ESX server and then double-click or right-click to open it.
- 3 Select **Inventory** ► **Storage** ► **Pools**.
- 4 From the View drop-down list, select Volumes.

- 5 Select a pool (datastore) to view associated volumes in the bottom content pane.

Viewing Storage Initiators & Volumes

To view storage initiators and the volumes for each initiator, perform the following steps:

- 1 From the Navigation pane, select **Servers** ► **Virtual Servers**.
- 2 In the content pane, select an ESX server and then double-click or right-click to open it.
- 3 Select **Relationships** ► **Storage Initiators**.
- 4 Select an initiator to view the volumes associated with it.


SAN Array Assets

This feature provides visibility into various storage array configurations that are attached to a SAN. This excludes SAN array devices that are directly attached to servers. SE captures generic SAN array assets as well as vendor-specific assets.


Using this feature, a server administrator can view the following:


- SAN array storage utilization and configuration
- SAN Array security configuration
- Storage resources shared between applications or hosts

SE captures the following generic SAN array assets:

SAN Array (Storage Systems)  —SAN array configuration information. SE uses an IP-based network based mechanism to access the device or device interface, such as API, CLI, and so on. Storage Visibility and Automation discovers the Manufacturer, Model, Serial Number, Firmware Version, Serial Number, and Storage Capacity for a SAN Array.

Disk Controllers—Disk controllers are discovered for arrays. However, the Storage Visibility and Automation user interface does not display the association or relationship between a disk controller and a disk for any disk controller.


Disks  —Disk storage capacity. Storage Visibility and Automation discovers the Free Capacity, Firmware, Label, Location, Manufacturer, Model, Serial Number, and whether a disk is a spare.

Storage Volumes  —Storage volumes that are carved out of disks. Storage Visibility and Automation discovers the Storage Capacity, Block Size (if available), Storage Type, and Volume ID. Meta volumes (devices) are composite volumes that are created using two or more regular storage volumes.

Composite Volumes—Some SAN arrays support composite volumes. For these types of arrays, Storage Visibility and Automation will capture information for composite volumes. Composite volumes allow individual storage volumes to be concatenated and create a storage volume with larger storage capacity.

Port Controllers—The port controller and a list of fabric ports in the controller.

Fabric Ports—SAN array fabric ports are used to provide FC connectivity between SAN array and fibre switch (standalone or part of a fabric). Storage Visibility and Automation does not capture internal fabric ports in the system. It captures the Port Number, State (Offline, Online), and the World Wide Name.

Storage Pools —Some SAN arrays support grouping volumes together, based on common characteristics. For example, for an HP XP or Hitachi SAN array, groups of storage volumes are grouped together in an entity called Array Groups. Not all SAN arrays support storage pool concepts. For SAN arrays that do support pool concepts, Storage Visibility and Automation captures the storage type and storage capacity. For supported SAN arrays, Storage Visibility and Automation captures the free extents in the pool. If the storage pool is not supported by the device, HP creates the storage pool for the storage volumes, based on the storage type.

LUN Mapping—Lun Mapping information, such as Storage Volume, SAN Array Port, Lun ID, and Host FC Port/FCA information.

SAN Array Assets and Relationships

Storage Visibility and Automation displays relationships between SAN array elements and other storage asset discovery features. See [Table 4](#).

Table 4 SAN Arrays Assets and Relationships

SAN Array Asset	External Element	Description
Connected fabric port	Ports connectivity	Provides the point-to-point links through a fabric.
Connected device	Fabric zone/alias context	Captures devices attached to the ports that are members of a zone or alias.

Update from Storage Essentials

Since, by default, storage information is collected from SE once every 24 hours, you might need updated information about a SAN array more frequently than once a day. Check SE to verify the frequency of its data collection, and configure it if necessary. See the *Storage Essentials SRM Software User Guide* for information about configuring the schedule for data collection.

In Storage Visibility and Automation, you can request an update of SAN information from SE at the device level (such as, the SAN array). See [Update from Storage Essentials](#) on page 25 and the *HP Server Automation: Storage Visibility and Automation Installation & Administration Guide* for more information.

Open with Storage Essentials

When you are viewing storage information about a SAN array in Storage Visibility and Automation and would also like to view storage information for the same SAN array in Storage Essentials, you can launch Storage Essentials to display the data that was discovered. See [Open with Storage Essentials \(Click & Launch\)](#) on page 26 for more information.

Frequently Asked Questions

Table 5 answers some frequently asked questions on finding information about SAN arrays.

Table 5 Frequently Asked Questions About SAN Array Assets

	Question	How to Find the Answer
1.	What storage is available in a SAN array in a datacenter?	<ul style="list-style-type: none"> In the View drop-down list, select Zones to display a list of zones for the array port. on page 48
2.	What LUNs are bound to the servers?	<ul style="list-style-type: none"> Viewing Storage Initiators Consuming Storage Resources from an Array on page 45 Viewing Initiator Volumes and Ports on page 46 Viewing Storage Volumes with Access Paths for Remote Initiators on page 46 Viewing Volumes Served by a SAN Array Port on page 48
3.	Are critical applications configured for high availability, such as multiple access paths)? What are the redundant paths between a server and a SAN array for an application?	<ul style="list-style-type: none"> Viewing Storage Volumes with Access Paths for Remote Initiators on page 46
4.	What is the storage utilization of the SAN array?	<ul style="list-style-type: none"> In the View drop-down list, select Zones to display a list of zones for the array port. on page 48
5.	Are the critical applications backed up?	<ul style="list-style-type: none"> Viewing Storage Volumes with Access Paths for Remote Initiators on page 46
6.	Do the SAN arrays provide room for growth?	<ul style="list-style-type: none"> In the View drop-down list, select Zones to display a list of zones for the array port. on page 48
7.	When will the SAN arrays run out of space?	<ul style="list-style-type: none"> In the View drop-down list, select Zones to display a list of zones for the array port. on page 48
8.	What is causing application performance degradation? How is all of the SAN configuration impacting the application? What are the shared resource between applications?	<ul style="list-style-type: none"> Viewing Storage Volumes with Access Paths for Remote Initiators on page 46 Viewing the Access Path from the Volume on page 49

Viewing SAN Arrays

To view a list of SAN arrays with their hardware and status, perform the following steps:

- From the Navigation pane, select **Storage** ► **SAN Arrays**.
- Select a SAN array to view the Summary in the content pane. The content pane displays information about the SAN array hardware and status.

Viewing SAN Arrays and Storage Configurations

To view the SAN array and storage configurations, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array. The content pane displays information about the SAN array hardware and status.
- 3 (Optional) Select a target and then select one of the following options from the Actions menu:
 - **Open**—Displays the SAN array details in a separate window, depending on the AAA permission.
 - **Open with HP Service Automation Visualizer**—Opens the target with Service Automation Visualizer (SAV) and displays the storage map.
 - **Open with Storage Essentials**—Opens the target with Storage Essentials.
 - **Update from Storage Essentials**—Requests an update of storage information at the device level from Storage Essentials.
 - **Rename**—Displays an editable field to rename the SAN array directly in the table, depending on the AAA permission.
 - **Delete**—Removes the SAN array from the Model Repository.

Viewing Summary of a SAN Array

To view a SAN array summary, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 From the View drop-down list, select Summary.
- 4 In the content pane, select a SAN array and then open it to display the summary.

Viewing SAN Array Properties

To view properties for a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 From the View drop-down list, select Properties. The content pane displays properties of the SAN array.
- 4 (Optional) Select a target and then select one of the following options from the Actions menu:
 - **Open**—Displays the SAN array, depending on the AAA permission.
 - **Open with HP Service Automation Visualizer**—Opens the target with Service Automation Visualizer (SAV) and displays the storage map.
 - **Open with Storage Essentials**—Opens the target with Storage Essentials.
 - **Update from Storage Essentials**—Requests an update of storage information at the device level from Storage Essentials.
 - **Rename**—Renames the target. Changes are saved in the system.
 - **Delete**—Removes the SAN array from the Model Repository.

Viewing Customers for a SAN Array

To view customers for a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 From the View drop-down list select Properties.
- 4 From the Actions menu, select **Open**, or right-click and then select **Open**. The Properties window displays.
- 5 Click the Change link to open the Select Customer window.
- 6 Select a customer and then click **Select**.

Viewing the SAN Array History Log

To view the history log for a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select History. The content pane displays the history log for the SAN array.
- 4 In the content pane, select an event from the history log and then select the following option from the Actions menu:
 - **View Event Details**—Displays detailed information about the event.Or
 - Right-click on the event and select **View Event Details**.
- 5 (Optional) Click Export to export the history log to a .csv or an .html file.

Viewing Storage Initiators Consuming Storage Resources from an Array

To view storage initiators consuming resources from an array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Storage Initiators. The content pane displays a list of storage initiators.
- 4 (Optional) In the lower content pane, select a storage initiator and then select the following option from the Actions menu:
 - **Open Device**—Opens the Device Explorer.Or
 - Right-click on the target and select **Open Device**.

Viewing Initiator Volumes and Ports

To view initiator volumes and ports in a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Storage Initiators.
- 4 Double-click on a SAN array. The SAN Array Browser displays.
- 5 In the SAN Array Browser, select a storage initiator. The content pane displays a list of initiator volumes and ports.
- 6 (Optional) In the lower content pane, select a target and then select the following option from the Actions menu:
 - **Open Device**—Opens the Device Explorer for the selected server.
 - **Open Initiator Port**—Displays the hardware information for the selected server.
 - **Open Initiator Volume**—Displays the storage information for the selected server.

Viewing Storage Volumes with Access Paths for Remote Initiators

To view storage volumes with access paths for remote initiators not managed by HP, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Storage Initiators. The content pane displays a list of initiator volumes and ports.
- 4 In the lower content pane, select a storage initiator.
- 5 Right-click on the selected SAN array and then select **Open**. The content pane displays a list of storage volumes.

Viewing Attached Fabrics in an Array

To view attached fabrics in a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select SAN Fabrics. The content pane displays information about attached fabrics.

Viewing Fabric Ports and Connectivity in an Array

To view array ports and their connectivity to the selected fabric, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Ports. The content pane displays information about array ports.

Viewing Zones Configured for an Array

To view zones configured for an array in a selected fabric, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select SAN Fabrics.
- 4 In the content pane, open a SAN array.
- 5 In the SAN Array Browser, select a Fabric.
- 6 In the View drop-down list, select Zones. The content pane displays the information for zone sets.

Viewing SAN Switches Attached to an Array

To view SAN switches attached to an array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select SAN Switches. The content pane displays the information for SAN switches.
- 4 In the content pane, open a SAN array.

Viewing Hardware Information for a SAN Array

To view the hardware information for the SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Hardware. The content pane displays the hardware information.

Viewing Front End Fabric Ports in a SAN Array

To view front end fabric ports in a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Ports. The content pane displays the port information.

Viewing Properties of the Fabric Port in a SAN Array

To view the properties of the fibre port, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Ports.
- 4 In the lower content pane, highlight a port to select it.

- 5 In the upper content pane, open a SAN array. The content pane displays the properties of the fibre port.

Viewing Volumes Served by a SAN Array Port

To view volumes served by a SAN array port, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Ports.
- 4 In the content pane, open a SAN array.
- 5 In the SAN Array Browser, highlight a port to select it.
- 6 In the View drop-down list, select Volumes. The content pane displays a list of volumes.
- 7 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open Device**—Opens the Device Explorer for the selected server.
 - **Open Initiator Port**—Displays the hardware information for the selected server.
 - **Open Initiator Volume**—Displays the storage information for the selected server.

Viewing Zones Configured for an Array Port

To view zones configured for an array port, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array.
- 3 In the View drop-down list, select Ports.
- 4 In the content pane, open a SAN array.
- 5 In the SAN Array Browser, highlight a port to select it.
- 6 In the View drop-down list, select Zones to display a list of zones for the array port.

Viewing Storage Summary for a SAN Array

To view the summary of storage information offered by a SAN Array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the View drop-down list, select Storage and then select the SAN array. The content pane displays summary information.

Viewing Volumes in a SAN Array

To view volumes in a SAN array, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.

- 3 From the Views pane, select **Inventory > Storage > Volumes**. The content pane displays the volumes.

Viewing Volume Properties

To view properties of a volume, perform the following steps:

- 1 From the Navigation pane, select **Storage > SAN Arrays**.
- 2 In the content pane, select a SAN array and then open the SAN array.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume and then open it. The content pane displays the properties.

Viewing the Connectivity of a Volume

To view the storage supply chain connectivity of a volume, perform the following steps:

- 1 From the Navigation pane, select **Storage > SAN Arrays**.
- 2 In the content pane, select a SAN array and then open the SAN array.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Connectivity. The content pane displays the connectivity of the volume.

Viewing the Composition of a Volume

To view the antecedents and dependents of the volume storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Storage > SAN Arrays**.
- 2 In the content pane, select a SAN array and then open the SAN array.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Composition. The content pane displays the composition of the volume.
- 6 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open**—Opens the SAN array or Server Storage view.

Viewing the Access Path from the Volume

To view the access path, perform the following steps:

- 1 From the Navigation pane, select **Storage > SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the Views pane, select **Inventory > Storage > Volumes**.
- 4 In the content pane, select a volume.

- 5 From the View drop-down list, select Access Path. The content pane displays the access path.
- 6 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open Device**—Opens the Device Explorer for the selected server.
 - **Open Initiator Port**—Displays the hardware information for the selected server.
 - **Open initiator Volume**—Displays the storage information for the selected server.

Viewing Array Ports from the Selected Volume

To view the array port that is part of the LUN mapping, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the Views pane, select **Inventory ► Storage ► Volumes**.
- 4 In the content pane, select a volume.
- 5 From the View drop-down list, select Ports. The content pane displays the ports.
- 6 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open Port**—Opens the SAN arrays view.

Viewing Storage Pools

To view storage pools configured in the storage system, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the Views pane, select **Inventory ► Storage ► Pools**. The content pane displays the list of storage pools.

Viewing Storage Pool Properties

To view the properties of a storage pool, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the Views pane, select **Inventory ► Storage ► Pools**.
- 4 In the content pane, select a storage pool. The content pane displays the properties.

Viewing Volumes for a Selected Pool

To view volumes for a selected pool, perform the following steps:

- 1 From the Navigation pane, select **Storage ► SAN Arrays**.
- 2 In the content pane, select a SAN array and then open it.
- 3 From the Views pane, select **Inventory ► Storage ► Pools**.
- 4 In the content pane, select a pool.

- 5 In the View drop-down list, select Volumes. The content pane displays the volumes

Fabric Assets


This feature provides visibility into fabric configurations in a SAN. Storage elements in a fabric include physical entities, such as fabric switches, ports, and security items, such as Zones and Zonesets.

Minimal information about fabrics is displayed in the user interface, such as when there is a connection between the fabric and the server, the Device Explorer displays fabrics. To view detailed information about fabric assets, you must use the Global File System (OGFS) feature. See [Global File System \(OGFS\)](#) on page 85 for more information.

With this feature, a system administrator can view the following:

- Physical fabric configuration
- Security configuration enabled in the fibre channel network
- Servers and SAN storage devices connected to the fibre channel network
- Storage resources shared between applications or managed servers

Storage Visibility and Automation captures the following fabric information in a SAN:

Fibre Switch and Fabrics  —Identifies whether a fibre switch is a principal switch or a subordinate switch of the fabric. After the principal switch is identified, Storage Visibility and Automation gathers information on the connected fibre switches and end devices, such as servers, SAN-based storage devices, and so on. Storage Visibility and Automation lists the switch ports for all switches in the fabric. Storage Visibility and Automation collects the Manufacturer, Model, Serial Number, Firmware, Hardware Version, World Wide Name (WWN), Domain Name, IP, and the number of Ports for each of the switch. Storage Visibility and Automation identifies a physical fabric by the WWN of the principal switch. If the role of a switch changes from subordinate to principal, the identity of the fabric changes—this change is discovered by SE. Storage Visibility and Automation does not collect routing table information in each switch and does not support multiple vendor fabrics that are in an interop configuration.


Switch Enclosure—Information about the switch enclosure includes blades and slots that make up that enclosure. This feature applies to director-class switches. Storage Visibility and Automation collects the Part Number, Serial Number, Manufacturer, and Model.


Switch Blade—Blade information about the enclosure. Storage Visibility and Automation collects the Manufacturer, Model and Serial Number of the blade.


Switch Ports—The fibre channel port World Wide Name (WWN), including information about F_Ports, E_Ports, and G_Ports. For each port, Storage Visibility and Automation collects the WWN, Port Number, Port State, and Port FCID (Fibre Channel ID). If enclosure is captured, Storage Visibility and Automation also collects slot information for the ports. Storage Visibility and Automation does not collect node loop ports attached to an arbitrated loop. However, Storage Visibility and Automation collects fabric switch loop port.

E_Ports/Inter-Switch Links—Inter-Switch Link (ISL) information. SE identifies the E_Ports that participate in the ISLs for each pair of fibre switches.

Ports Connectivity and N_Ports—For each fibre switch port, Storage Visibility and Automation collects the World Wide Name (WWN) of the connected N_Port, such as FCA Port (REF-008) and Storage Array Port (REF-009). This information helps identify the dependency chain from the initiators to the targets through the fabric, and also displays information about the multi-path between the initiators and the targets. Ports that are discovered without any parent device information are marked as Unknown Connected Ports. Storage Visibility and Automation does not collect node loop ports attached to an arbitrated loop.

Zoning Database and Zones  —All zones configured in the zoning database of the principal switch. Zones contain a list of World Wide Name (WWN) and alias naming. Alias naming may not be supported by all fibre switch vendors. A zoning database may not be synchronized between switches. The administrator should verify that zoning databases across the fabric are synchronized. Storage Visibility and Automation does not capture the zoning information if the list (members) are configured as Domain:Port format.

Alias  —Storage Visibility and Automation collects all alias in a zoning database. Alias is not supported by McData-based fabric switches. Storage Visibility and Automation does not collect the alias information if the list (members) are configured as Domain:Port format.

Zonesets  —Storage Visibility and Automation collects all active and inactive zonesets in a fabric. Storage Visibility and Automation also collects all zones that make up a zoneset.

Fabric Assets and Relationships

Storage Visibility and Automation displays relationships between fabric elements and other storage asset discovery features. See [Table 6](#).

Table 6 Fabric Assets and Relationships

Fabric Asset	External Element	Description
Connected N_Port	Ports connectivity	Provides the point-to-point links through a fabric.

NetApp Assets

This feature provides visibility into NetApp storage configurations in a SAN or IP network, such as File I/O based storage entities (Filer Volumes, Qtrees) and Fabric elements (FCA/Ports).

A NetApp storage administrator can view the following:

- NetApp storage utilization and configuration
- Storage resources shared between applications or managed servers


SE captures the following NetApp information in a SAN:

NetApp Filer—Filer Name, Manufacturer, Model, Serial Number, Firmware Version, Customer, IP Address, Facility, Status, Description, Discovered on, Last Scan, Scan Status, Object ID.

Disks—NetApp disk storage capacity. Storage Visibility and Automation discovers the raw capacity of the disk and discovers whether the disk is consumed by RAID groups.

RAID Groups—NetApp creates RAID groups based on disks. RAID group is an internal entity that is not configured by the user. Storage Visibility and Automation does not display storage information for RAID groups. The storage information is not available through the API. It is derived by the size of the disk and configuration, such as RAID 4 or RAID 4DP.

Plex—NetApp creates Plex based on RAID Groups. Plex is an internal entity that is not configured by the user. Storage Visibility and Automation does not display storage information for Plex. The storage information is not available through the API—it is derived by the size of the RAID Group.

Aggregates —Aggregate is a storage entity that is introduced by ONTAP 7.0. Storage Visibility and Automation discovers aggregate caption and aggregate capacity.

File I/O—NetApp exposes storage resources for File I/O. NetApp volumes, Qtree and Directories are exposed through CIFS or NFS for File I/O. Storage Visibility and Automation collects storage information related to File I/O access.

NetApp Volumes—Filer volume information that includes storage capacity, vFiler that the volumes are associated with, aggregate that the volume consumes, volume type (Traditional or Flex), snapshot reserve, snapshot reserve size, and whether the volume is a target or a source of replication.

Traditional Volumes—Volumes based on ONTAP 6.5 and traditional volumes in ONTAP 7.0 or higher.

Flex Volumes—Flex is a storage entity that is introduced by ONTAP 7.0. In addition to filer volume information, Storage Visibility and Automation also discovers committed capacity and space guarantee (volume, none, file) for flex volumes.

Qtrees—A subdirectory in a volume that acts as a virtual sub-volume with special attributes, such as primarily quotas and permissions.

Quota—The quota against the NetApp volumes. Quota that is applied on the Qtrees is also captured at the volume level. Storage Visibility and Automation discovers the following quota information:

- Caption
- Type (tree, user, or group)
- Qtree (Qtree that is applicable)
- Unix ID (ID of the user or group, whichever is applicable)
- Windows ID (ID of the user or group, whichever is applicable)
- Space threshold
- Space threshold unit (in KB/MB, and so on)
- Quota target (Default targets are discovered as *.)

Quota will also have zero or more QuotaLimit. QuotaLimit captures the following information:


- Limit type (DiskSpace or NumberOfFiles)
- Soft Limit Value
- Soft Limit Unit
- Hard Limit Value
- Hard Limit Unit

Directory—Directory paths that are exported through CIFS share or NFS. Storage Visibility and Automation does not discover all directories that are residing in the NetApp volume.

File I/O Exports—NetApp provides exporting NetApp Volumes using File I/O through CIFS and NFS protocol.

CIFS—Common Internet File System. NetApp shares and the volume/qtree/directory path. Storage Visibility and Automation does not discover the managed servers that are consuming the NetApp resources through CIFS protocol.

NFS—Network File System. Export information for the volumes/Qtree/directory path. HP might discover stale information, based on the data available to it through the API.

NetApp vFiler —vFiler configuration. Storage Visibility and Automation discovers IP addresses and the volume path that is configured for each vFiler. NetApp requires that a MultiStore license is enabled in the filer to support virtual filers. If a MultiStore license is not enabled, Storage Visibility and Automation assumes that a default vfiler0 exists in the system.

NetApp Assets and Relationships

Storage Visibility and Automation displays relationships between NetApp elements and other storage asset discovery features. See [Table 7](#).

Table 7 Fabric Assets and Relationships

Fabric Asset	External Element	Description
Connected fabric port	Ports connectivity	Provides the point-to-point links through a fabric.
NFS exports and file I/O storage elements	Remote file system in a server	Provides supply chain from the managed server to the NetApp, based on NFS protocol.

Update from Storage Essentials

Since, by default, storage information is collected from SE once every 24 hours, you might need updated information about a NetApp filer more frequently than once a day. You can modify this data collection frequency by modifying a setting in SE. See the *Storage Essentials SRM Software User Guide* for information about configuring the schedule for data collection.

You can also request an update of SAN information from SE at the NetApp level. See [Update from Storage Essentials](#) on page 25 for more information.

Open with Storage Essentials

When you are viewing storage information about a NetApp filer in Storage Visibility and Automation and would also like to view storage information for the same NetApp filer in Storage Essentials, you can launch Storage Essentials to display the data that was discovered. See [Open with Storage Essentials \(Click & Launch\)](#) on page 26 for more information. See the *Storage Essentials SRM Software User Guide* for information about the discovery process in SE.

Frequently Asked Questions

Table 8 answers some frequently asked questions on finding information about NetApp assets.

Table 8 Frequently Asked Questions About NetApp Assets

	Question	How to Find the Answer
1.	What storage is available in the NetApp in a datacenter?	<ul style="list-style-type: none">• Viewing Summary of a NAS Filer on page 55
2.	What file systems are exported to the servers?	<ul style="list-style-type: none">• Viewing Exports for a vFiler on page 57
3.	What is the storage utilization of the NetApp?	<ul style="list-style-type: none">• Viewing Summary of a NAS Filer on page 55
4.	Does the NetApp Filer provide room for growth?	<ul style="list-style-type: none">• Viewing Summary of a NAS Filer on page 55
5.	When do the NetApp Filers run out of space?	<ul style="list-style-type: none">• Viewing Summary of a NAS Filer on page 55

Viewing NAS Filers

To view NAS Filers, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the View drop-down list, select Summary.

Viewing the NAS Filer Storage Configuration

To view information about NAS Filer storage configuration, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 Open a NAS Filer.
- 3 From the Navigation pane, select Inventory.
- 4 In Views pane, select a configuration to change the display in the content pane.

Viewing Summary of a NAS Filer

To view a NAS filer summary, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the View drop-down list, select Summary.
- 3 In the content pane, select a NAS Filer and open it to display the Summary.

Viewing NAS Filer Properties

To view the properties of a NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.

- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select a NAS Filer and then open it to display its properties.

Viewing Customers for a NAS Filer

To view customers that exist in the Model Repository, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select a NAS Filer and open it to display the Properties.
- 4 Click the “Change” link next to Customers to open the Select Customer dialog.
- 5 In the Select Customer dialog, select a customer and then click Select to display its properties.

Viewing the NAS Filer History Log

To view the history log for the NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the View drop-down list, select History.
- 3 In the content pane, select a NAS Filer and open it to display the history log.
- 4 From the View drop-down list, select a time range for this history log, such as Last Day, Last Week, Last Month, Last Quarter, or Custom Range.
- 5 (Optional) Right-click an event to open the Event Details dialog.
- 6 (Optional) Click **Export** to export the history log to a file, such as .csv or .html.

Viewing Attached Fabrics in a NAS Filer

To view attached fabrics in a NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer.
- 3 In the View drop-down list, select SAN Fabrics.

Viewing SAN Switches Attached to a NAS Filer

To view SAN switches attached to a NAS filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS filer.
- 3 In the View drop-down list, select SAN Switches. The content pane displays the information for SAN switches.

Viewing vFilers

To view the vFilers that are configured in the NetApp system, perform the following steps:

- 1 From the Navigation pane, select **Device** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer.
- 3 In the View drop-down list, select Virtualization.

Viewing Volumes Configured for a vFiler

To view the volumes that are configured for a vFiler, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Virtualization**.
- 4 In the View drop-down list, select Volumes.
- 5 In the content pane, select a vFiler.

Viewing Exports for a vFiler

To view the exports for a vFiler, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Virtualization**.
- 4 In the View drop-down list, select Exports.
- 5 In the content pane, select a vFiler.

Viewing Storage Summary for a NAS Filer

To view the summary of storage information offered by a NAS filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage**. The content pane displays summary information.

Viewing Disks in a NAS Filer

To view the disks in a NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 Select a disk and then open it.

Viewing NAS Filer Disk Properties

To view the properties of a NAS Filer disk, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 From the View drop-down list, select Properties.
- 5 Select a disk to display its properties.

Viewing Filer Volumes Created by a Disk

To view the filer volumes that are created by a disk, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 Select a disk.
- 5 From the View drop-down list, select Volumes.

Viewing Aggregates Created by a Disk

To view the NAS Filer aggregates that are created by a disk, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 Select a disk.
- 5 From the View drop-down list, select Aggregates.

Viewing Plexes on a Disk

To view the plexes that are on top of a disk in the storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Disks**.
- 4 Select a disk.
- 5 From the View drop-down list, select Plexes.

Viewing RAID Groups on a Disk

To view the RAID groups that are on top of a disk in the storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.

- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory ► Storage ► Disks**.
- 4 Select a disk.
- 5 From the View drop-down list, select RAID Groups.

Viewing Aggregates Configured in a NAS Filer

To view the aggregates that are configured in a NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Storage ► NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory ► Storage ► Aggregates**.
- 4 From the View drop-down list, select Properties.

Viewing Aggregate Properties

To view the properties of an aggregate, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Storage ► NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory ► Storage ► Aggregates**.
- 4 From the View drop-down list, select Properties.
- 5 Select an aggregate.

Viewing Filer Volumes Created by an Aggregate

To view the filer volumes that are created by an aggregate, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Storage ► NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory ► Storage ► Aggregates**.
- 4 From the View drop-down list, select Volumes.
- 5 Select an aggregate.

Viewing Antecedents and Dependents of an Aggregate

To view the antecedent and dependent of the aggregate storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Storage ► NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory ► Storage ► Aggregates**.
- 4 Select an aggregate.
- 5 From the View drop-down list, select Composition.

Viewing Plexes Used by an Aggregate

To view the plexes that are used by an aggregate, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Aggregates**.
- 4 Select an aggregate.
- 5 From the View drop-down list, select Plexes.

Viewing RAID Groups Consumed by an Aggregate

To view the RAID groups that are consumed by an aggregate, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Aggregates**.
- 4 Select an aggregate.
- 5 From the View drop-down list, select RAID Groups.

Viewing Filer Volumes Configured in a NAS Filer

To view the filer volumes that are configured in a NAS Filer, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 From the View drop-down list, select Properties.

Viewing Filer Volume Properties

To view the properties of a filer volume, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 From the View drop-down list, select Properties.
- 5 In the content pane, select a filer volume.

Viewing Antecedents and Dependents of a Volume for a NAS Filer

To view the antecedent and dependent of a volume storage supply chain, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.

- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 Select a volume.
- 5 From the View drop-down list, select Composition.

Viewing Qtrees Created in a Volume

To view the Qtrees that are created in a volume, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 Select a volume.
- 5 From the View drop-down list, select Qtree.

Viewing Quota Created at the Volume or Qtree

To view the quota that is created at the volume or the Qtree for a volume, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 Select a volume.
- 5 From the View drop-down list, select Quota.



A disk quota is a set of rules that restrict disk space and the number of files used by a user or a group. A quota can also restrict the total space and files used in a Qtree or the usage of users and groups within a Qtree. For example, you could create a user quota that specifies that a particular user can use up to 225MB of disk space in a given Qtree in a given volume before a warning is issued (soft limit) and up to 250MB maximum (hard limit).

Viewing Exports for a Volume

To view the exports for a volume, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **Storage** ► **NAS Filers**.
- 2 In the content pane, select a NAS Filer and then open it.
- 3 In the Views pane, select **Inventory** ► **Storage** ► **Volumes**.
- 4 Select a volume.
- 5 In the View drop-down list, select Exports.

Oracle Database Assets

This feature provides visibility into the Oracle storage configurations in a SAN or Network Attached Storage (NAS). These configurations include physical entities, such as fabric switches and SAN Arrays.


A database administrator can view the following:

- Oracle database storage configuration
- System resources available to the Oracle databases in one or more servers in a data center
- Servers and SAN Storage devices connected to the FC network
- Storage resources shared between Oracle Instances and/or hosts
- Network Attached Storage (NAS) resources that the Oracle Databases depends on

Storage Visibility and Automation provides the following Oracle database storage configuration in a SAN or NAS:

Oracle Instance—Storage-related information for the Oracle Instance. Storage Visibility and Automation collects the version of the Oracle database that is running the Oracle instance.


Storage Elements—Storage elements configured for an Oracle instance can be classified as physical database storage and logical database storage. *Physical database storage* includes data files and redo logs that directly consume system storage resources, such as file systems and partitions. *Logical database storage* includes entities such as tablespaces that are created inside the instance that is consuming different physical and logical storage entities.

Data Files —The file path, allocated space, and space used for each data file. Oracle data files can reside directly on the file system or they can utilize raw partitions.

- **Based on File System**—Data files can be created directly in a file system and specified to the Oracle database. Storage Visibility and Automation collect information about the file path of the data file. HP uses this information to calculate the file system that is consumed by the data file.
- **Based on Raw Partition**—Storage Visibility and Automation collects partition information. HP uses this information to calculate the storage supply chain that is consumed by the data file.

Redo Logs—Redo logs information for the Oracle instance, such as file path and allocated space for the redo log.

Logical Database Storage—Storage-related information for the tablespace and its association to the physical storage entities. Storage Visibility and Automation does not collect storage entities, such as data segments, extents, data block, schema, table, and index.

Tablespaces —Storage Visibility and Automation collects tablespace information for each Oracle instance, such as the status, allocated space, and space used for each of the tablespaces.

Automatic Storage Management (ASM)—Storage configuration and utilization of the Oracle database. ASM helps manage volumes that might be composed of block-based devices (SAN) or file-based devices (NFS), with additional data replication inside ASM.

Database Assets and Relationships

Storage Visibility and Automation displays information about relationships between the database elements and other storage asset discovery features. See [Table 9](#).

Table 9 Database Assets and Relationships

Database Asset	External Element	Description
SAN-based physical database storage	Block storage dependency	Provides the dependency chain between the database storage elements and SAN arrays through System (server) resources and fabric.
NAS-based physical database storage	NAS storage dependency	Provides the dependency chain between the database storage elements, Server resources and NetApp.

Frequently Asked Questions

Table 10 answers some frequently asked questions on finding information about Oracle database assets.

Table 10 Frequently Asked Questions About Database Assets

	Question	How to Find the Answer
1.	What are the storage resources used by an Oracle instance?	<ul style="list-style-type: none"> Viewing Tablespaces Configured in the Oracle Instance on page 65 Viewing Tablespace Antecedents and Dependents on page 66 Viewing Storage Supply Chain Connectivity of the Tablespace on page 66
2.	What is the FC SAN network configuration that an Oracle database depends on?	<ul style="list-style-type: none"> Viewing Storage Supply Chain Connectivity of the Tablespace on page 66
3.	Are critical Oracle instances configured for high availability, such as multiple access paths? What are the redundant paths between a server and a storage array for an Oracle instance?	<ul style="list-style-type: none"> Viewing Storage Supply Chain Connectivity of the Tablespace on page 66 SAN Map on page 77 Storage Map on page 77
4.	What is the storage utilization of an Oracle instance?	<ul style="list-style-type: none"> Viewing Properties of the Oracle Instance on page 65
5.	Is Oracle-consumed storage on a remote SAN array configured for replication?	<ul style="list-style-type: none"> Viewing Storage Supply Chain Connectivity of the Tablespace on page 66 Viewing Utilized Targets on page 69
6.	What is causing Oracle instance performance degradation? How is all of the SAN configuration impacting the Oracle instance? What are the shared resources between Oracle instances?	<ul style="list-style-type: none"> Viewing Storage Volumes with Access Paths for Remote Initiators on page 46 Viewing the Access Path from the Volume on page 49 Viewing Storage Supply Chain Connectivity of the Tablespace on page 66
7.	Does the SAN provide room for growth?	<ul style="list-style-type: none"> Viewing Storage Supply Chain Connectivity of the Tablespace on page 66 Viewing Utilized Fabrics on page 69 Viewing Utilized Switches on page 69 Viewing Utilized Targets on page 69
8.	What are the NAS (Network Attached Storage) resources used by an Oracle instance?	<ul style="list-style-type: none"> Viewing Storage Supply Chain Connectivity of the Tablespace on page 66 Viewing Utilized Targets on page 69

Viewing Databases

To view databases discovered by the Database Scanner for Oracle, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database. The content pane displays a list of databases.
- 4 (Optional) Select a target and then select one of the following options from the Actions menu:
 - **Open**—Opens the Database Browser with the selected view.
 - **Rename (F2)**—Provides an editable field to rename the database directly in the table. The command is enabled only if the user has permissions.
 - **Delete**—Removes the database from the Model Repository. This command is enabled only if you have permission. This command does not remove the login credentials for the Database Scanner for Oracle.

Viewing Database Storage Configurations

To view database storage configurations, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Redologs**. The content pane displays configuration information.

Viewing Properties of the Oracle Instance

To view the properties of the Oracle instance, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it. The content pane displays the properties of the selected database.

Important to Know

File ► Save—This action is enabled when the name or description is edited. This saves the changes to the system.

File ► Revert—This action is enabled when the name or description is edited. This reverts the changes to the previous settings.

Viewing Tablespaces Configured in the Oracle Instance

To view tablespaces configured in the Oracle instance, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.

- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 From the Navigation pane, select **Inventory**.
- 5 From the Views pane, select ► **Tablespaces**. The content pane displays a list of tablespaces.

Viewing Tablespace Properties

To view tablespace properties, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 From the Navigation pane, select **Inventory**.
- 5 From the Views pane, select ► **Tablespaces**.
- 6 In the content pane, select a tablespace and then review the lower part of the content pane for the properties.

Viewing Tablespace Antecedents and Dependents

To view antecedents and dependents of a tablespace, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the View command, select **Inventory ► Tablespaces**.
- 5 In the content pane, select a tablespace and then review the lower part of the content pane.
- 6 In the View drop-down list, select Composition. The content pane displays the composition.
- 7 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open**—Opens the SAN Array, Server Storage, or NAS Filer view.

Viewing Storage Supply Chain Connectivity of the Tablespace

To view the connectivity of the database storage element, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Tablespaces**.
- 5 In the content pane, select a tablespace and then review the lower part of the content pane.
- 6 In the View drop-down list, select Connectivity. The content pane displays the connectivity.

Viewing Redo Logs

To view the redo logs, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Redologs**. The content pane displays the information.

Viewing Redo Log Properties

To view properties of a redo log, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Redologs**.
- 5 In the content pane, select a redo log and then review the lower part of the content pane for the properties.

Viewing Redologs Antecedents and Dependents

To view the composition of the redo logs, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Redologs**.
- 5 In the View drop-down list, select Composition.
- 6 In the content pane, select a redo log and then review the lower part of the content pane for the composition.
- 7 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open**—Opens the SAN Array, Server Storage, or NAS Filer view.

Viewing Redologs Connectivity

To view the connectivity of the database storage element, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then open it.
- 4 In the Views pane, select **Inventory ► Redologs**.
- 5 In the content pane, select a redolog and then review the lower part of the content pane.
- 6 In the View drop-down list, select Connectivity. The content pane displays the connectivity.

Viewing Datafiles

To view the datafiles, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Datafiles**. The content pane displays a list of datafiles.

Viewing Datafile Properties

To view the properties of a datafile, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Datafiles**.
- 5 In the View drop-down list, select Properties. The content pane displays the properties.
- 6 In the content pane, select a datafile and then review the lower part of the content pane.

Viewing Datafile Antecedents and Dependents

To view the composition of the datafiles, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Datafiles**.
- 5 In the View drop-down list, select Composition.
- 6 In the content pane, select a datafile and then review the lower part of the content pane for the composition.
- 7 (Optional) Select a target and then select the following option from the Actions menu:
 - **Open**—Opens the SAN Array, Server Storage, or NAS Filer view.

Viewing Datafile Connectivity

To view the connectivity of the database storage element, perform the following steps:

- 1 From the Navigation pane, select **Library ► Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Inventory ► Datafiles**.
- 5 In the content pane, select a datafile and then review the lower part of the content pane.
- 6 In the View drop-down list, select Connectivity. The content pane displays the connectivity.

Viewing Servers Hosting a Database

To view servers hosting a database, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Relationships** ► **Severs**. The content pane displays a list of servers.
- 5 (Optional) Select a target and then select one of the following options from the Actions menu:
 - **Open**—Opens the Device Explorer.



The column in this panel is the same as the main panel for the server rendered in the All Managed Servers feature. All columns are not described above.

Viewing Utilized Fabrics

To view utilized fabrics, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Relationships** ► **SAN Fabrics**. The content pane displays a list of fabrics.

Viewing Utilized Switches

To view utilized switches, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Relationships** ► **SAN Switches**. The content pane displays a list of switches.

Viewing Utilized Targets

To view utilized targets, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **Databases**.
- 2 In the View drop-down list, select Properties.
- 3 In the content pane, select the database and then double-click or right-click to open it.
- 4 In the Views pane, select **Relationships** ► **SAN Arrays**. The content pane displays a list of targets.
- 5 (Optional) Select a target and then select one of the following options from the Actions menu:
 - **Open**—Opens the SAN Array Browser.



The column in this panel is the same as the main panel for the arrays rendered in the SAN Array feature. All columns are not described above.

Viewing the Database History Log

To view the history log, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **Databases**.
- 2 In the View drop-down list, select History.
- 3 In the content pane, select the database and then double-click or right-click to open it. The content pane displays the history log.
- 4 (Optional) Select a target and then select the following option from the Actions menu:
View Event Details—Displays the detailed information for the event.

Device Group Assets

The Device Groups feature in Storage Visibility and Automation provides a useful way for gathering storage devices into collections. Grouping storage devices enables you to view all of the devices simultaneously, instead of viewing individual devices, one at a time.

This feature allows you to include storage devices in a device group and view a history of events for a device group. You can view detailed storage device history, such as information in snapshot messages from the synchronization process. Device groups and a history log can help you troubleshoot problems.

The Device Groups feature in Storage Visibility and Automation is modeled after the Device Groups feature in SA. See the *SA User Guide: Server Automation* for more information about managing device groups.

You can create the following types of storage device groups:

- **Static Device Group**—This group contains heterogeneous storage devices, such as any combination of servers, storage arrays, and NetApp filers.
- **Dynamic Device Group**—This group contains a homogeneous set of storage devices, such as servers or storage systems (SAN Arrays and NAS filers). A dynamic device group contains only one type of storage device because it is generated based on device properties. A dynamic device group cannot contain a combination of different types of storage devices.

See the *SA User Guide: Server Automation* for information about adding a server to a static or dynamic device group.

Static Device Groups

Static storage device groups can contain heterogeneous devices—one or more SAN arrays, NAS filers, and managed servers.

Creating a Static Device Group

To create a static device group, perform the following steps:

- 1 From the Navigation pane, select **Device Groups** ► **Public**.

- 2 From the **Actions** menu, select **Device Group ► New Static Group**.
- 3 Enter a new name for the static device group that you just created. Click outside the name field to save it. The default name is New Device Group.

Adding a Device to a Static Device Group

To add a storage device to a static device group, perform the following steps:

- 1 In the Device Group Browser, select the device group.
- 2 From the Actions menu, select **Open** to display the Group Browser.
- 3 In the Views pane, select Device Membership.
- 4 In the Device Membership window, from the Actions menu, select **Add**.
- 5 In the Add Members to Static Group dialog, select one or more members, such as such as any combination of servers, storage arrays, and NetApp filers.
- 6 Click **Add to Group**.

Removing a Device from a Static Device Group

To remove a storage device from a static device group, perform the following steps:

- 1 In the Static Device Group window, select one or more storage devices.
- 2 From the Actions menu, select **Remove Members**.

Dynamic Device Groups

Dynamic storage device groups can contain only homogeneous devices—one or more storage systems (SAN arrays and NAS filers), and managed servers.

Creating a Dynamic Device Group




To create a dynamic device group, perform the following steps:

- 1 From the Navigation pane, select **Device Groups ► Public**.
- 2 From the **Actions** menu, select **Device Group ► New Dynamic Group**
- 3 Enter a new name for the dynamic device group that you just created.
- 4 Press enter to save it. The default name is New Device Group plus an incremental number, such as New Device Group 4.

Adding a Device to a Dynamic Device Group

To add a storage device to a dynamic device group, perform the following steps:

- 1 In the Device Group Browser, select the device group.
- 2 From the Actions menu, select **Open** to display the Group Browser.
- 3 In the Views pane, select Device Membership.
- 4 In the Content pane, specify the rules for the dynamic storage device group.

- 5 From the first drop-down list, select **Server** or **Storage Systems**. Depending on the attribute you select, options available for the operator and values for the rule will change. The default is **Storage Systems**.
- 6 From the second drop-down, select **Facility Name**, **Manufacturer**, **Model**, **Name**, **Object ID**, **Port**, **RAID Type**, or **Serial Number**. The default is **Manufacturer**.
- 7 Select the operator from the third drop-down list. The default is **Contains**.
- 8 Enter a value in the field, or select a value from the drop-down list, or click  to select multiple values from the **Select Values** window.
- 9 Click  to add additional rules and repeat steps 4 through 6. Click  to delete any rules.
- 10 Select the logic (**And/Or**) to be applied for every rule in the query.
- 11 Click **Preview** to view the storage devices that matched the rule—the rule that defines members of the dynamic device group.
- 12 From the **File** menu, click **Save** to save the rules.

Editing a Rule for a Dynamic Device Group

To change a rule for a dynamic device group, perform the following steps:

- 1 Open the **Group Browser** and then select a dynamic device group.
- 2 From the **Device Membership** window, change the rule and then save it.

Device History Log

A history of events for a storage device group can help you troubleshoot problems.

Viewing the History Log

To view the history log for a device group, perform the following steps:

- 1 Open a storage device group.
- 2 From the **Views** pane, select **History**.
- 3 From the **View** drop-down list, select the time period for the history log.
- 4 (Optional) Right-click an event to open the **Event Details** dialog.
- 5 (Optional) Click **Export** to export the history log to a file, such as **.csv** or **.html**.

3 Storage Compliance

Overview of Storage Compliance

Storage compliance provides storage administrators and server administrators the ability to audit storage infrastructure to identify any potential issues. A storage administrator typically needs to audit the infrastructure to identify the potential for problems or issues before they have a chance of affecting an internal customer (business owner, application, or server). A server administrator typically needs to make sure the storage provided meets certain Service Level Agreements (SLAs) for reliability and performance. From a compliance perspective, provided storage information would also help auditors, compliance officers, and managers, when applied globally to classes of servers or applications.

Because storage administrators and server administrators generally share a common goal of ensuring consistency with the unique software stack that has been approved and tested by an organization, they need to monitor firmware and other software compliance. A storage administrator needs to confirm the right firmware levels. A server administrator needs to make sure the HBAs, volume managers, MPIOs, and so on, are the correct versions. The shared end-goal is a standardized environment where interoperability between all moving parts is much more reliable.

SA Client Integration

Storage compliance in Storage Visibility and Automation is integrated with the Audit and Remediation feature in the SA Client. To implement storage compliance, you must be familiar with the basics of SA Audit and Remediation. See the *SA User Guide: Application Automation* for information about using the following Audit and Remediation capabilities:

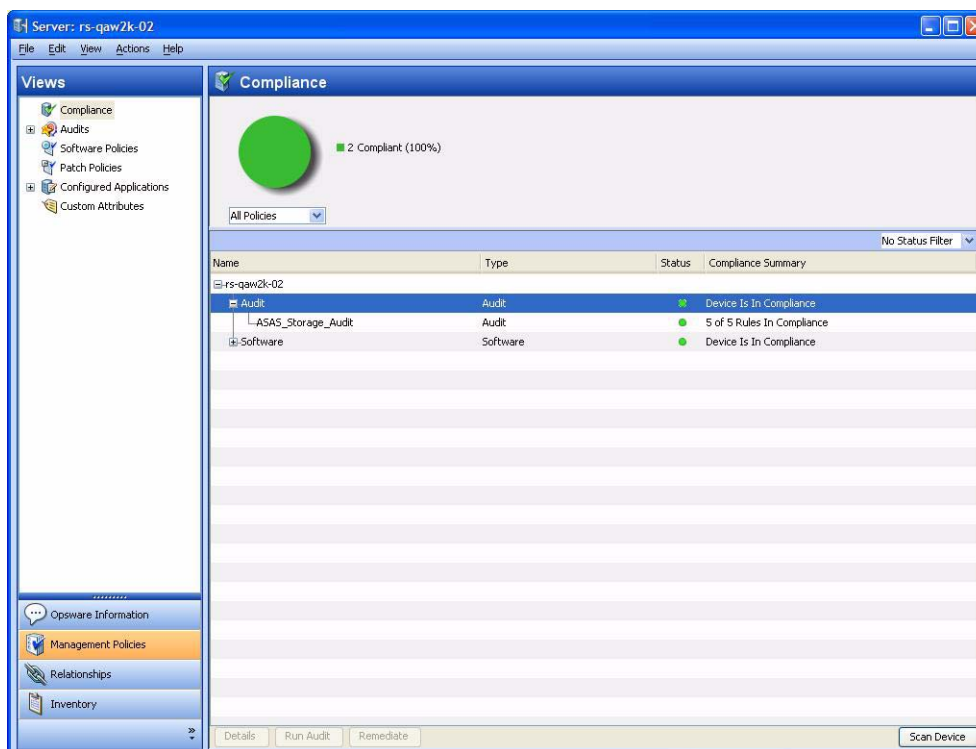
- Compliance View
- Configuring a storage audit
- Configuring specific rules
- Selecting target servers
- Configuring a snapshot specification
- Running and scheduling an audit
- Viewing audit results
- Deleting an audit

Compliance View

To help troubleshoot problems, the Compliance View allows you to view storage compliance for managed servers in a facility. In the Compliance View, you can navigate to a managed server to view the audits you created and see whether the server is compliant. Figure 2 shows storage compliance information for a managed server.

See the *SA User Guide: Application Automation* for more information about the Compliance View.

Figure 2 Compliance View for Storage Audits



Audit Rules and Parameters

The storage compliance feature provides visibility into storage configurations in a managed server or device group by allowing you to specify audit rules that check initiator ports connected to multiple ports, FCAs, LUNs available but not mounted, RAID levels, SAN storage, and so on.



Storage audit rules are run against data in the Model Repository.

Table 11 lists the audit rules and parameters for storage compliance. These rules allow you to define the check for minimum, maximum, or exact numbers.

Table 11 Storage Audit Rules & Parameters

Audit Rule	Parameter
Unmounted Volume Capacity	Acceptable total capacity of unmounted volumes in bytes

Table 11 Storage Audit Rules & Parameters (cont'd)

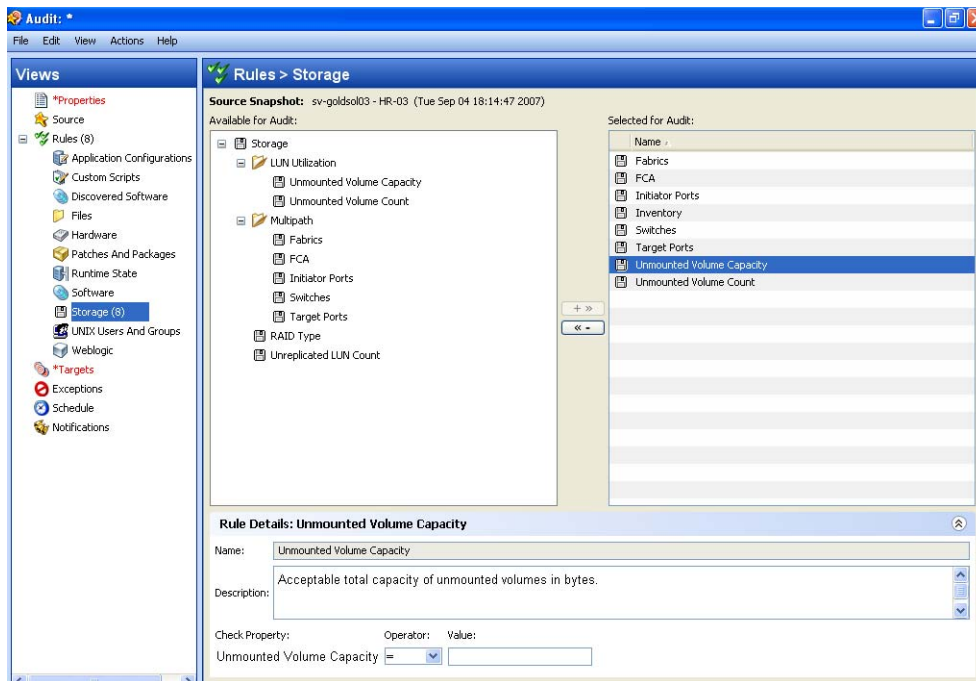
Audit Rule	Parameter
Unmounted Volume Count	Acceptable number of unmounted volumes
Fabrics	Acceptable number of fabrics
FCA	Acceptable number of Fibre Channel Adapters (FCAs)
Initiator Ports	Acceptable number of ports
Switches	Acceptable number of switches
Target Ports	Acceptable number of ports
RAID Type	Acceptable RAID types on the target storage array Note: The audit will fail if this rule is selected and no RAID type is specified.

▶ The compliance rules that involve ports, switches, or fabrics, check active ports only. These types of compliance rules do not check for physical port connectivity.

Audit Browser

Figure 3 shows storage rules selected in an audit. See the *SA User Guide: Application Automation* for information about specifying rule details.

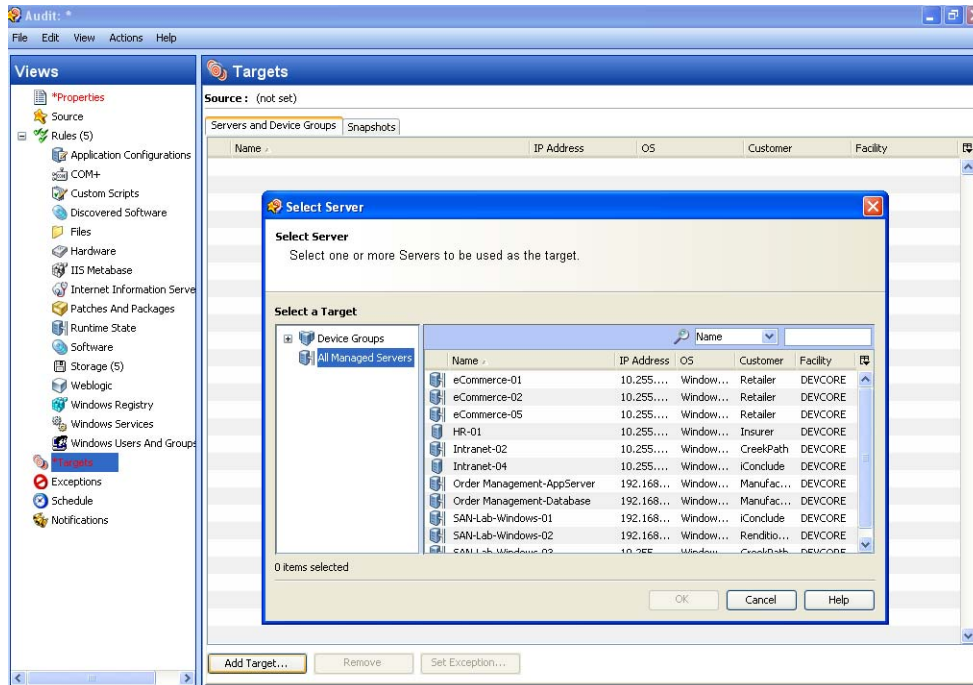
Figure 3 Audit Browser for Storage Rules



Select Server

Figure 4 shows how you specify which servers are used as targets in the audit. See the *SA User Guide: Application Automation* for information about adding audit targets.

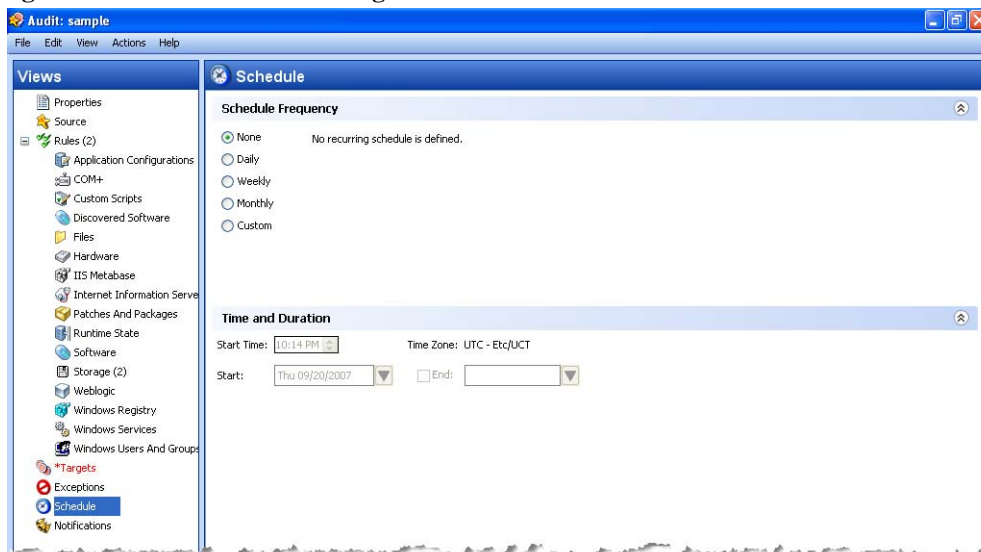
Figure 4 Target Server for a Storage Audit



Audit Schedule

Figure 5 shows how to schedule an audit. See the *SA User Guide: Application Automation* for information about scheduling audits.

Figure 5 Schedule for a Storage Audit



4 Storage Visualization

Overview of Service Automation Visualizer

The Service Automation Visualizer (SAV) feature in Server Automation (SA) allows you to visualize logical storage dependencies and physical storage connections in your data center. SAV displays how these dependencies and connections relate to your business applications in mapping diagrams. These diagrams map out the storage supply chain from servers to spindles—how servers map to the storage devices they are connected to. SAV diagrams also show file systems and their relationships to local and remote storage, FC adapters, FC switches (physical and virtual), connections and ports, disk arrays, NAS filers, LUN mappings, RAID configurations, SAN fabrics, and so on.



Storage Essentials (SE) version 6.1.1 or later is required to view, report, or perform any Service Automation Visualizer (SAV) and Service Automation Reporter (SAR) operation on SAN objects, such as arrays, switches, volumes, and so on. SAN objects are discovered in Storage Essentials. To enable discovered SAN objects in the SA, SAV, and SAR products, the Server Automation SE Connector component must be installed and configured.

See the *SA User Guide: Application Automation* for more information about SAV.

Storage Mapping Features

SAV provides visual maps that display physical and logical drawings of storage devices and connections in your environment—in a Storage Map and a SAN map.

SAN Map

The SAN map shows a superset of the Storage Map, including a graphical view of the SAN involved in a SAV snapshot, all servers and their FCAs (including fibre channel ports), and each adapter's connections to switches in the SAN. See the *SA User Guide: Application Automation* for more information.

Storage Map

The Storage Map displays storage dependencies for servers, FCAs, NAS filers, SAN arrays, fabrics, switches, ports, and so on, that you have scanned with SAV. This map provides a graphical view that displays servers, the process families running on them, including the file systems and local or remote storage devices these files are stored on and served from, either through local disks, NAS filers, or remote disk arrays. The map shows connections between process families and their open files and where the files are stored. See the *SA User Guide: Application Automation* for more information.

5 Storage Reports

Overview of Storage Reports

Storage reports provide comprehensive, real-time information about storage configurations in managed servers, and storage hardware and software in your environment. You can use these reports to help determine when an application will run out of storage, what the storage trend is for an application, and so on.

These parameterized reports are presented in graphical and tabular format, and are actionable—which means that you can perform appropriate actions on objects, such as a policy or an audit, within the report. These reports are also exportable to your local file system (as .pdf, .html, and .xls files) to facilitate use within your organization.

Storage reports are organized by the following folders:

- Reports ► Storage Reports ► Database Storage Reports
- Reports ► Storage Reports ► Host and Application Storage Reports
- Reports ► Storage Reports ► Storage Array Reports



Storage reports are available only for managed servers that have a Storage Host Agent Extension (SHA) installed and running. See the *Storage Visibility and Automation Installation & Administration Guide* for information about installing SHA.

Service Automation Reporter Client Integration

The reports feature in Storage Visibility and Automation is integrated with the reports feature in the Service Automation Reporter (SAR) Client. To use storage reports, you must be familiar with the basic capabilities of SAR reports. See the *SAR User Guide* for information about using the following report capabilities:

- Report formats
- Running a report
- Scheduling a report
- Exporting a report
- Printing a report

Database Storage Reports

Table 12 describes the database storage reports available in Storage Visibility and Automation.

Table 12 Database Storage Reports

Report Title	Description
Database Allocation Trend	This report shows how allocated and used capacity for a set of databases has changed over time and is projected to change in the future.
Database Capacity and Utilization Trend Data	This report shows how the capacity and usage information for a set of databases has changed over time and is projected to change in the future.
Database Utilization Trend	This report shows how usage information for a set of databases has changed over time and is projected to change in the future.
Tablespace Allocation Trend	This report shows how allocated and used capacity for one or more tablespaces across one or more databases has changed over time and is projected to change in the future.
Tablespace Capacity and Utilization Overview	This report lists tablespaces and their capacity and utilization information, grouped by database.
Tablespace Capacity and Utilization Trend Data	This report shows how usage information for one or more databases, or for one or more tablespaces has changed over time and is projected to change in the future.
Tablespace Utilization Trend	This report shows how usage information for one or more tablespaces across one or more databases has changed over time and is projected to change in the future.

Host and Application Storage Reports

Table 13 describes the host and application storage reports available in Storage Visibility and Automation.

Table 13 Host and Application Storage Reports

Report Title	Description
Host Capacity and Utilization Overview	This report shows detailed capacity and utilization information for multiple hosts that are potentially grouped by customer, facility or device group.
Host Capacity and Utilization Trend Data	This report shows detailed capacity and utilization trend data for a collection of hosts that are grouped by customer, facility or device group.
Host DB Storage Allocation Trend	This report shows how allocated and used database storage for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host DB Storage Utilization Trend	This report shows how the database storage utilization rate for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host File System Storage Allocation Trend	This report shows how allocated and used file system storage for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host File System Storage Utilization Trend	This report shows how the file system storage utilization rate for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host Reclaimable Storage Overview	This report shows the amount of readily reclaimed storage of various types for a set of hosts.
Host Storage Detail	This report shows a detailed inventory of all storage-related information related to a specific host.
Host Total Storage Allocation Trend	This report shows how total allocated and used storage for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.

Table 13 Host and Application Storage Reports (cont'd)

Report Title	Description
Host Total Storage Utilization Trend	This report shows how the total storage utilization rate for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host Volume Manager Storage Allocation Trend	This report shows how allocated and used volume manager storage for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.
Host Volume Manager Storage Utilization Trend	This report shows how the volume manager storage utilization rate for a set of hosts (grouped by customer, facility or device group) has changed over time and is projected to change in the future.

Storage Array Reports

Table 14 describes the storage array reports available in Storage Visibility and Automation.

Table 14 Storage Array Reports

Report Title	Description
Array Capacity and Utilization Overview	This report provides summary capacity and utilization statistics for SAN storage arrays and NAS filers, grouped by customer, facility, or device group as desired.
Array Detail	This report shows a detailed inventory of storage-related information related to a specific storage array or filer.
Array Inventory	This report provides basic inventory information and storage statistics for SAN storage arrays and NAS filers.
Storage Allocated to Hosts Unmanaged by the SA Storage	This report identifies arrays and filers providing storage to hosts that are unknown to or unmanaged by the SA Storage.

6 Storage Search

Overview of Storage Search

In Storage Visibility and Automation, you can search for storage information in your IT environment, such as objects and attributes for server, database, and storage systems.

SA Client Integration

The search feature in Storage Visibility and Automation is integrated with the SA Client Search tool in the SA Client. To use search, you must be familiar with the basics of the SA Client Search tool. See the *SA User Guide: Server Automation* for information about using the following SA Client Search capabilities:

- Performing a simple search by using keywords
- Performing an advanced search by creating search queries
- Saving a search query
- Deleting a saved search
- Running a saved search query
- Performing actions on search results
- Emailing search results
- Printing search results

Storage Objects and Attributes

By using the SA Client Search tool, you can search for the following storage objects and attributes:

- **Server**—FCA Manufacturer, FCA Model, FCA Port, FCA Version (Fibre Channel Adapter), Storage Make & Model, Storage Manager Software Manufacturer, and Storage Manager Software Type
- **Database**—Server Name and Version
- **Storage Systems**—Device Customer, Device Facility, Device Group, Discover Date, Facility Name, Manufacturer, Model, Name, Object ID, Port, RAID Type, and Serial Number

A Global File System (OGFS)

Overview of OGFS

The Global File System (OGFS) feature shows storage inventory in a file system-like presentation. The OGFS feature in Storage Visibility and Automation is integrated with the OGFS feature in the SA Client. To use the OGFS feature, you must be familiar with its basic capabilities in the SA Client. See the *SA User Guide: Server Automation* for more information.

The OGFS represents the data model as a hierarchical structure of file directories and text files. For example, in the OGFS, the `/opsw/Server` directory contains information about managed servers and includes subdirectories that represent the contents of the managed servers, such as storage, ports, and disks.

Global Shell

The Global Shell is a command-line interface to the OGFS. This command-line interface is a Unix shell, such as bash, that runs in a terminal window. The default shell in Global Shell is bash.

The OGFS unifies the data model and the contents of storage devices and managed servers, including files, into a single, virtual file system. You open a Global Shell session from within a SA Client or from a direct ssh connection in a terminal client on your desktop. With the Global Shell, you can automate repetitive system administration tasks by running scripts across multiple storage devices and servers in a secure environment. Common scripting languages, such as Perl and Python, are supported.

The OGFS represents the data model as a hierarchical structure of file directories and text files. For example, in the OGFS, the `/opsw/Server/@/e-commerce.example.com/Storage` directory contains details about storage dependencies for the server named `e-commerce.example.com`.

See the *SA User Guide: Server Automation* for more information on how to use Global Shell.

Permissions

The actions that you can perform within the Global Shell are determined by the operations specified by the `aaa` utility. The `aaa` utility grants and revokes permissions for operations that use the OGFS. For example, the `aaa` utility grants permission for the `readServerFilesystem` operation, which allows you to browse a server's file system in the SA Client. To run the `aaa` utility, you must belong to the Administrators user group. The permissions granted and revoked by the `aaa` utility are stored in the `/opsw/Permissions` directory of the OGFS. See the *SA User Guide: Server Automation* for information about how to use the `aaa` utility.



You must have permissions in the Global Shell to view the file systems for storage inventory.

Storage Directory in the OGFS

The following sections describe subdirectories that contain storage content. The directory layout is designed to reflect the natural locations of storage objects, with the intent to facilitate intuitive navigation by OGFS users.

The primary storage directory roots are `/opsw/Storage`, which contains all storage system and storage network content and `/opsw/Server/.../Storage`, which contains storage objects specific to their respective servers.

Additional storage directory trees are provided to scope the content contained within `/opsw/Storage`. Examples include the `/opsw/Customer/.../Storage` directories, which contain storage resources that are assigned to customers, and the `/opsw/Facility/.../Storage` directories, which contain storage resources that are associated with facilities.

Although storage trees exist in multiple scoped locations, each and every object has a single location at which it can be reliably located in the OGFS. These locations are collectively referred to as *canonical locations*. Within each storage directory tree, symbolic links (symlinks) are used to traverse storage topology and reference storage information located in canonical directory paths.

The canonical location for a volume of a server is in the following directory:

```
/opsw/Server/.../Storage/Volume
```

In the following example, the Volume directory contains content pertaining to a raid0 volume on the server:

```
/opsw/Server/@/myserver.opsware.com/Storage/Volume/raid0
```

If a file system is mounted on that volume, this is reflected in the OGFS as a symbolic link to the volume, which is located in its canonical location:

```
/opsw/Server/@/myserver.opsware.com/Storage/FileSystem/myfs/Volume/raid0  
-> /opsw/Server/@/myserver.opsware.com/Storage/Volume/raid0
```

`/opsw/api/com/opsware/storage` Directory

This directory is automatically provided by the existing (SA Client) UAPI infrastructure. The api directory contains executables that invoke methods on the SA API.

`/opsw/Storage` Directory

This directory contains storage system and storage network artifacts. This directory is replicated in the following paths with filtered content:

- `/opsw/Facility/<facility>/@/Storage`—This directory shows the content of `/opsw/Storage` with the top-level directories filtered by facility.
- `/opsw/Group/<Public|Private>/<group>/@/Storage`—This directory shows `opsw/Storage` content with the top level directories filtered by group.
- `/opsw/Customer/<customer>/@/Storage`—This directory shows `opsw/Storage` content filtered by customer.

`/opsw/Server/.../Storage` Directory

This directory contains server-related storage artifacts.

```
/opsw/Server/.../Storage
```

/opsw/Server/.../Storage/Controller Directory

This directory contains a subdirectory with symlinks to controller ports and a subdirectory with symlinks to volumes imported through this controller.

```
/opsw/Server/.../Storage/Controller/Port  
/opsw/Server/.../Storage/Controller/Volume
```

/opsw/.../Disk/Dependent Directory

This directory contains symlinks to dependent volumes. There are several places where the Disk directory can exist. It can occur on the server side (to show the local disks on a server), such as:

```
/opsw/Server/@/myserver.opsware.com/Disk
```

Or, it can be located on a storage system (to show the disks contained in a storage array), such as:

```
/opsw/Storage/System/emcdmx-01/Disk
```

The Dependent subdirectory can exist in many locations (one of which is a subdirectory of a disk). In each case, the Dependent subdirectory shows storage artifacts that are dependent on the object represented in the current directory, such as:

```
/opsw/Storage/System/emcdmx-01/Disk/01A\C0/Dependent
```

In this example, this path shows the objects that directly depend on disk 01A:C0.



In Storage Visibility and Automation, when disks are referred to in the context of a storage system, disk information might not be available. Disk information that is provided depends on what the settings are for discovering disks in Storage Essentials.

/opsw/.../Server/Storage/Fabric Directory

This directory contains symlinks to fabric-connected devices, zones on the fabric, symlinks to ports in the zone, zone aliases on the fabric, symlinks to zone aliases that reference the zone, zone sets, and switches on the fabric.

```
/opsw/Server/.../Storage/Fabric/ConnectedDevice  
/opsw/Server/.../Storage/Fabric/Switch  
/opsw/Server/.../Storage/Fabric/Zone  
/opsw/Server/.../Storage/Fabric/Zone/Member  
/opsw/Server/.../Storage/Fabric/Zone/ZoneAlias  
/opsw/Server/.../Storage/Fabric/Zone/ZoneSet  
/opsw/Server/.../Storage/Fabric/ZoneAlias  
/opsw/Server/.../Storage/Fabric/ZoneAlias/Member  
/opsw/Server/.../Storage/Fabric/ZoneAlias/Zone  
/opsw/Server/.../Storage/Fabric/ZoneSet  
/opsw/Server/.../Storage/Fabric/ZoneSet/Zone  
/opsw/Server/.../Storage/Fabric/Switch
```

[/opsw/Server/.../Storage/FileSystem Directory](#)

This directory contains symlinks to the volume underlying the file system.

```
/opsw/Server/.../Storage/FileSystem/Volume
```

[/opsw/Server/.../Storage/Port Directory](#)

This directory contains symlinks to a connected device, symlinks to a connected port, symlinks to a protocol controller, symlinks to the connected fabric, symlinks to volumes on this port, symlinks zones referencing this port, and symlinks to zone aliases referencing this port. This directory can reside under multiple storage domains.

```
/opsw/Port/Connected Device  
/opsw/.../Port/ConnectPort  
/opsw/.../Port/Controller  
/opsw/.../Port/Fabric  
/opsw/.../Port/Volume  
/opsw/.../Port/Zone  
/opsw/.../Port/ZoneAlias
```

[/opsw/Server/.../Storage/Service Directory](#)

This directory contains MPIO Manager and Volume Manager software, symlinks to volumes exported, and symlinks to volumes imported by each service.

```
/opsw/Server/.../Storage/Service/<service-name>/Volume  
/opsw/Server/.../Storage/Service/<service-name>/InputVolume  
/opsw/Server/.../Storage/Service/<service-name>/OutputVolume
```

[/opsw/Server/.../Storage/Switch Directory](#)

This directory contains symlinks to connected devices, symlinks to connected fabrics, and switch ports.

```
/opsw/Server/.../Storage/Switch/Fabric  
/opsw/Server/.../Storage/Switch/Port
```

[/opsw/Server/.../Storage/System Directory](#)

This directory contains storage system protocol controllers, storage system physical disks, storage system volume pools, symlinks to volumes within a storage pool, storage system ports, and storage volumes.

```
/opsw/Server/.../Storage/System/Controller  
/opsw/Server/.../Storage/System/Disk  
/opsw/Server/.../Storage/System/Pool  
/opsw/Server/.../Storage/System/Pool/Volume  
/opsw/Server/.../Storage/System/Pool/Volume/Antecedent  
/opsw/Server/.../Storage/System/Pool/Volume/Dependent  
/opsw/Server/.../Storage/System/Pool/Volume/Initiator  
/opsw/Server/.../Storage/System/Pool/Volume/Pool  
/opsw/Server/.../Storage/System/Pool/Volume/Port  
/opsw/Server/.../Storage/System/Port  
/opsw/Server/.../Storage/System/Volume
```


/opsw/Server/.../Storage/Volume Directory

This directory contains symlinks to antecedent storage volumes and physical disks, symlinks to controllers importing this volume, symlinks to dependent storage volumes, symlinks to dependent file systems, paths through which the volume is imported, symlinks to storage ports importing this volume, symlink to the service that provides this volume, symlink to the target volume in the storage system domain, symlinks to connected host initiator volumes, and symlinks to the storage pool associated with this volume.

```
/opsw/Server/.../Storage/Volume/Antecedent  
/opsw/Server/.../Storage/Volume/Controller  
/opsw/Server/.../Storage/Volume/Dependent  
/opsw/Server/.../Storage/Volume/FileSystem  
/opsw/Server/.../Storage/Volume/Path  
/opsw/Server/.../Storage/Volume/Port  
/opsw/Server/.../Storage/Volume/Service  
/opsw/Server/.../Storage/Volume/Target
```


Index

A

AAA. See Authentication, Authorization, & Accounting., 19
aaa utility, 85
aggregate capacity, 53
aggregate caption, 53
aggregates, 53
Array Capacity and Utilization Overview Report, 42
Array Detail Report, 82
Array Groups, 42
Array Inventory Report, 82
ASM. See Automatic Storage Management., 37, 62
Authentication, Authorization, & Accounting, 19
Automatic Storage Management, 37, 62

B

block-based devices, 62

C

canonical locations, 86
CIFS. See Common Internet File System., 53
committed capacity, 53
Compliance View, 13, 18, 73
composite volumes, 41
connected fabric port, 42
consumed storage, 29

D

DAS. See Direct Attached Storage., 28
Database Allocation Trend Report, 80
Database Browser, 15
Database Capacity and Utilization Trend Data Report, 80
Database Scanner for Oracle, 23, 24
Database Utilization Trend Report, 80
data files, 62
data replication, ASM, 62

datastore, 40
dependent volumes, 87
device groups, 20
 dynamic, 71
 static, 28, 29, 41, 42, 51, 52, 53, 54, 62, 70
Direct Attached Storage, 28
disk controllers, 41
disk quota, 61
dynamic device group, 70, 71

E

E_Ports, 51

F

F_Ports, 51
fabric ports, 41
fabric switches, 51
FCA. See Fibre Channel Adapter., 28
Fibre Channel Adapter, 14, 28
Fibre Channel ID, 51
fibre switch and fabrics, 51
file-based devices, 62
File I/O and Block I/O, 53
File I/O Exports, 53
flex volumes, 53
free extents, 42

G

G_Ports, 51
GAED, 25
GAED. See Get All Element Details., 24, 25
Gateway, 19
Get All Element Details, 24
Global File System, 51, 85
Global Shell, 85
Global Shell File System, 18

guest VM, 27
guest volume, 37

H

HBA. See Fibre Channel Adapter., 14
Hitachi SAN array, 42
Host Capacity & Utilization Overview Report, 81
Host Capacity and Utilization Trend Data Report, 81
Host DB Storage Allocation Trend Report, 81
Host DB Storage Utilization Trend Report, 81
Host File System Storage Allocation Trend Report, 81
Host File System Storage Utilization Trend Report, 81
Host Reclaimable Storage Overview Report, 81
Host Storage Detail Report, 81
Host Total Storage Allocation Trend Report, 81
Host Total Storage Utilization Trend Report, 82
Host Volume Manager Storage Allocation Trend Report, 82
Host Volume Manager Storage Utilization Trend Report, 82
HP XP, 42
hypervisor, 27, 36, 37

I

Inter-Switch Link, 51
ISL. See Inter-Switch Link., 51

L

layered volume managers, 29
layered volumes, 29
local storage, 28
logical database storage, 24, 62
logical entities, 28
logical volume, 26
LUN Mapping, 42

M

meta volumes, 41
Model Repository, 23, 26, 27
 Oracle setup, 85
MPIO. See Multipath IO., 29, 35
Multimaster, 19
Multipath IO, 29
MultiStore, 54

N

N_Port, 51
NAS. See Network Attached Storage., 23, 24, 28, 62
NetApp Filer, 52
NetApp VFile, 54
NetApp volumes, 53
Network Attached Storage, 23, 24, 28, 62
NFS. See Network File System., 53, 62
NFS exports, 54

O

OGFS. See Global File System., 51, 85
OGFS. See Global Shell File System., 18
ONTAP 6.5, 53
ONTAP 7.0, 53
Oracle ASM. See Automatic Storage Management., 37
Oracle instance, 62

P

physical database storage, 24, 62
physical disk, 26
physical entities, 28
Plex, 52
port controllers, 41
ports connectivity, 51
primordial storage, 28, 29
principal switch, 51, 52
produced storage, 29

Q

Qtrees, 53
quota, 53

R

RAID groups, 52
raw partitions, 62
redo logs, 62
replication, 64
root volumes, 28

S

SA. See Server Automation., 77
SAN. See Storage Area Network., 28, 62

SAN Array Browser, 15
SAR. See Service Automation Reporter., 24, 77, 79
SAV. See Service Automation Visualizer., 12, 24, 44, 77
SE. See Storage Essentials., 24, 77
SE Connector, 23, 24
server administrator, 11, 28
Server Automation, 24, 77
server discovery, 20
Service Automation Reporter, 24, 77, 79
Service Automation Visualizer, 12, 24, 44, 77
Service Level Agreement, 73
SE Storage Scanner. See SE Connector., 24
SHA. See Storage Host Agent Extension., 27, 79
SLA. See Service Level Agreement., 73
space guarantee, 53
space threshold, 53
static device group, 28, 29, 41, 42, 51, 52, 53, 54, 62, 70
Storage Allocated to Hosts Unmanaged by the SA Storage, 82
Storage Area Network, 28
Storage Essentials, 44, 87
Storage Host Agent Extension, 23, 27, 32, 79
storage pool, 42
storage target, 38
storage volumes, 41
subordinate switch, 51
switch blade, 51
switch enclosure, 51
switch port, 51
symbolic links, 86
symlinks, 87
symlinks. See symbolic links., 86, 87

T

Tablespace Allocation Trend Report, 80
Tablespace Capacity and Utilization Overview Report, 80
Tablespace Capacity and Utilization Trend Data Report, 17, 80
tablespaces, 62
Tablespace Utilization Trend Report, 17, 80
traditional volumes, 53

U

Unified API, 18

V

vFiler, 53
virtual filers, 54
VMware ESX hypervisors, 27
VMware ESXi server, 27
VMware ESX server, 23

W

World Wide Name, 28, 51, 52
WWN. See World Wide Name., 28, 51, 52

Z

zones, 51, 52
zonesets, 51, 52
zoning database, 52

