



Technical Report

The Self-Healing Data Fabric

Solving Storage Problems in Real Time

Yaron Regev, NetApp
April 2017 | TR-4585

Abstract

No matter how hard we try, or how much innovation we throw at a problem, it is extremely difficult to build a system that can respond to every imaginable issue in a manner that will satisfy all users.

This technical report explores the concept of self-healing storage and provides practical solutions that address the NetApp® Data Fabric, whether on the premises or in the public cloud.

TABLE OF CONTENTS

1	Introduction to Self-Healing Systems.....	3
2	Self-Healing Systems Design	4
2.1	System-Monitoring Component of Self-Healing Storage.....	5
2.2	System Analysis and Alert Management Component of Self-Healing Storage	9
2.3	Remediation Orchestration Component of Self-Healing Storage	12
3	Self-Healing Systems Implementation	17
3.1	Use Case 1: Automatic Capacity Management.....	17
3.2	Use Case 2: Automatic Inode Management in High-File-Count (HFC) Environments	28
3.3	Use Case 3: Automatic Performance Management (Tiering).....	35
	Appendix A: Integrating OnCommand Manager Alerts with WFA When Monitoring Multiple Clusters.....	38
	References.....	44

1 Introduction to Self-Healing Systems

No matter how hard we try or how much innovation we throw at a problem, it is extremely difficult (perhaps impossible) to build a system that can respond to every imaginable issue in a manner that will satisfy all users. After all, we are creatures of opinions, and opinions are hard to predict and nearly impossible to automate. Of course, we wouldn't need to deal with such issues if all systems were perfect and never failed, but that is not today's reality. Perhaps it will be tomorrow.

Nowadays, most systems support some level of internal self-healing capability (some more than others). Internal self-healing requires the ability to identify aspects of the system that are performing suboptimally and to remediate those issues. Those processes are usually universal in nature, meaning that they do not depend on the type of workload for which the system is used. They are also usually fairly easy to describe in a consistent, repeatable, and predictable manner. Such internal self-healing capabilities might include RAID reconstruction, scale-out workload load balancing across compute nodes, or recovery from component failure.

The more compelling aspect of a self-healing system revolves around the customer's workload and the need for service levels. If an application is about to run out of space or reach a performance limit, should it be granted additional capacity or IOPS? What if the user paid only for a specific level of service? What if the space or IOPS is available, but allocating either will negatively affect some internal best practice of utilization threshold? If a document repository is rapidly growing to a substantial size, should we consider moving it to a more cost-effective object repository instead? If a system is delivering an acceptable level of performance but shows a consistent pattern of degradation, should empirical data be collected proactively in case a support engagement is required with the vendor? These questions are some of the many considerations that we encounter when pursuing a self-healing infrastructure.

This technical report explores the concept of self-healing storage and provides examples of practical solutions for NetApp Data Fabric systems, whether on the premises or in the public cloud.

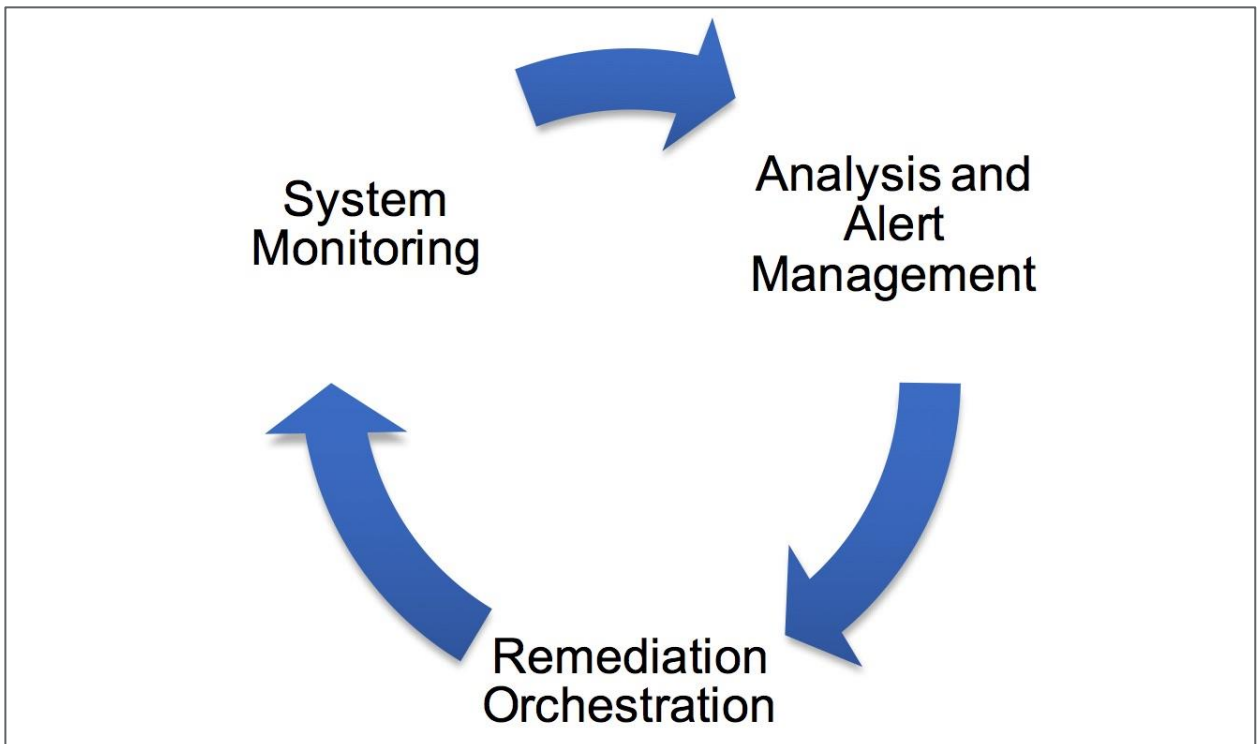
2 Self-Healing Systems Design

Every system that supports adaptive, customizable, and expandable lifecycle management of actionable events must incorporate the following:

- A component that monitors the state and usage patterns of the system
- A component that performs anomaly detection and analysis (in its simplest form—an event and alert management component)
- A component that executes a prescribed remediation activity (an orchestration component)

Figure 1 shows the relationship between the components of the self-healing system, which provides continuous, cyclical feedback.

Figure 1) Components of a self-healing system.

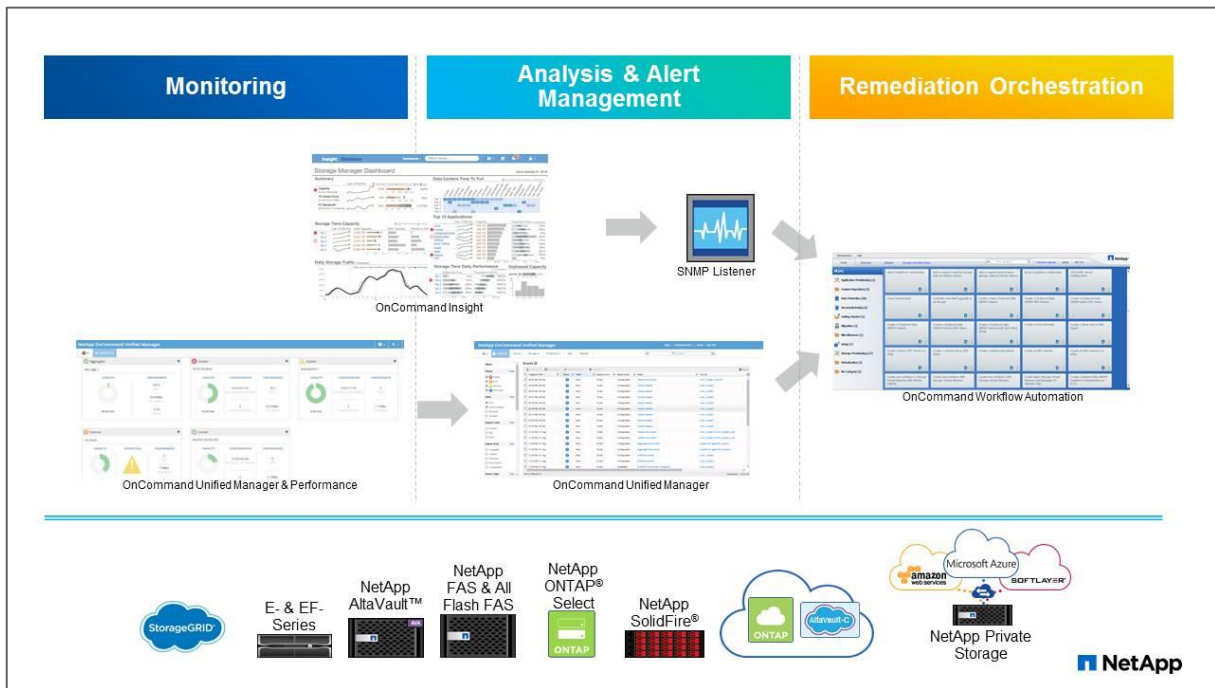


Within the NetApp OnCommand® management suite, the NetApp Data Fabric provides all the components that are needed to deliver a self-healing system:

- System monitoring can be performed by OnCommand Unified Manager, OnCommand Performance Manager, and OnCommand Insight.
- Analysis and alert management can be performed by OnCommand Unified Manager and OnCommand Insight.
- Remediation orchestration can be performed by OnCommand Workflow Automation.

Figure 2 shows the NetApp Data Fabric solution for an end-to-end self-healing storage infrastructure.

Figure 2) NetApp Data Fabric end-to-end self-healing storage infrastructure.



Note About NetApp Service Level Manager

NetApp has recently released new management software called NetApp Service Level Manager. This software optimizes storage operations for predictable service levels, defines service-level objectives (SLOs), and enables the transition to a storage-as-a-service private cloud platform.

NetApp Service Level Manager follows a methodology called MAPE (Monitor, Analyze, Plan, Execute), which is similar to the process that is described in this document.

NetApp Service Level Manager is beyond the scope of this document, however. For additional details, see the References section.

2.1 System-Monitoring Component of Self-Healing Storage

A prerequisite to self-healing automation is the ability to auto-discover the environment, continuously monitor its behavior, and understand what is considered normal (baseline) and what is acceptable deviation from normal. Many products on the market can accomplish this task—with varying levels of functionality and integration. This document explores the NetApp OnCommand Unified Manager and OnCommand Performance Manager (focused on NetApp ONTAP powered solutions) products that are available to monitor the Data Fabric.

Introduction to OnCommand Unified Manager and OnCommand Performance Manager

OnCommand Unified Manager and OnCommand Performance Manager are included with ONTAP and provide a comprehensive data management solution for NetApp FAS and All Flash FAS systems:

- NetApp OnCommand Unified Manager is specifically designed to support the innovative capabilities of ONTAP software. It facilitates monitoring the health, availability, capacity, performance, and data

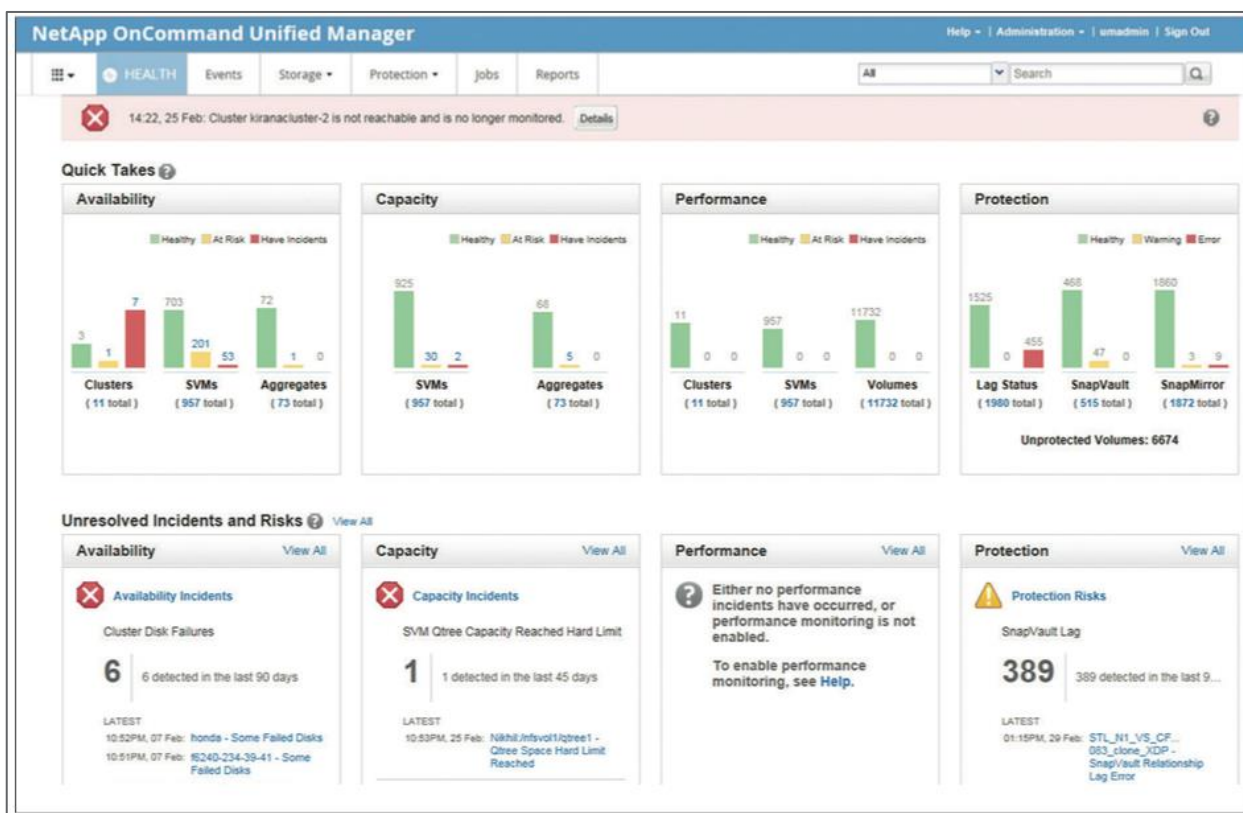
protection status of your NetApp clustered storage. In addition, it provides alerts and vital information for proactive management.

- NetApp OnCommand Performance Manager is tightly integrated with Unified Manager. It delivers comprehensive storage performance monitoring and data retention, along with notifications and alerts for proactive management. It also provides timely system analysis and assists with root-cause identification for performance issues.

Together, these products provide a complete view of your NetApp storage health. A simplified GUI enhances the overall user experience with an integrated view of how the storage infrastructure is performing. Administrators can monitor health and performance attributes from a single portal by logging in to a single URL for access to Unified Manager and Performance Manager. The Favorites dashboard (Figure 4) is a personalized feature where users can store frequently accessed views, which helps streamline monitoring of essential storage data. With this feature, you can:

- Get proactive notifications about storage events and issues.
- Quickly drill down to vital information with color-coded topology views.
- Receive suggestions for corrective actions for fast remediation.
- Use groups and annotations to filter notifications and report information.
- Access standard reports or create custom operational reports.

Figure 3) OnCommand Unified Manager dashboard.



Monitoring and alerting are fully dynamic. Trending and historical information—with graphical representations of the storage topology—help facilitate analysis. This approach optimizes use, capacity, and performance across the storage infrastructure. Full monitoring of your NetApp SnapMirror®, SnapVault®, and MetroCluster™ environments is incorporated. Simple, peered, cascaded, and fan-out

topologies for data protection are all monitored, thus keeping important data protection operations in place and completing on schedule.

Simple, Powerful Management for NetApp ONTAP

The OnCommand Unified Manager management server provides the foundation for improved availability, scalability, supportability, performance, and security of the storage infrastructure.

Maximize availability. Use Unified Manager to view overall storage health by cluster, aggregate, or storage virtual machine (SVM). From the main dashboard, you can easily see availability, capacity, performance, and protection status or drill down to see details about physical or logical objects. This information enables quick understanding of where availability is at risk, the corrective actions that can be taken to eliminate potential issues, and the steps to avoid unplanned downtime.

Maintain control over capacity. Unified Manager helps project when available capacity will be at levels that require expansion. Notifications include details with recommended actions so that you can quickly determine where to provision volumes for new workloads. It also identifies where to move volumes to minimize the risk of running out of capacity.

Automate data protection. In a busy storage environment, confidence that your data is adequately protected is essential. With OnCommand Unified Manager, you can immediately determine whether data is protected or is at risk. Health monitoring for SnapMirror, SnapVault, and MetroCluster environments identifies when data protection relationships are in jeopardy because of capacity or configuration changes. Unified Manager helps automate data protection when provisioning new storage, promoting a consistent process that follows best practices.

Get critical operational reporting. To provide vital operational information, Unified Manager comes with several standard reports. Optionally, customized reports to meet the specific needs of your organization can be generated. Reports cover topics such as configuration, capacity use, operational status, storage efficiency, and inventory. For example, you can easily determine whether configurations comply with best practices, or you can review use trends to determine when there is a risk of running out of capacity.

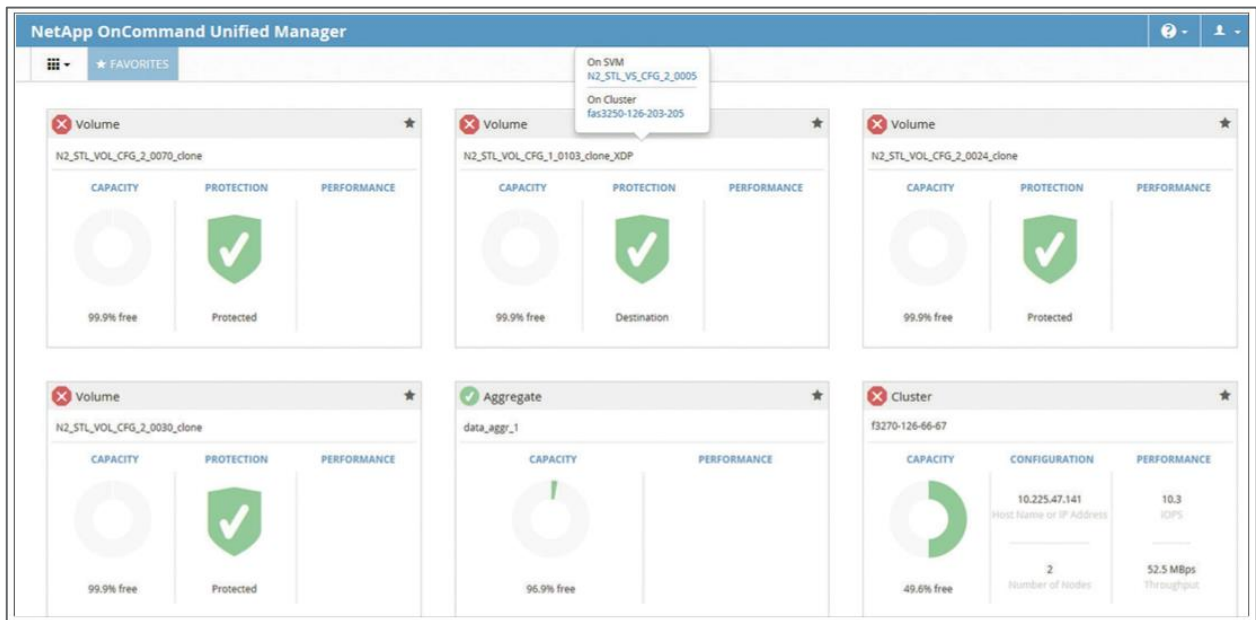
You can view reports in PDF, HTML, and CSV formats; filter, sort, and group data within the standard reports; or schedule reports to run later. Unified Manager also provides an interface that extracts the operational data and enables import to another reporting solution.

View historical data. By viewing the volume move status and history and the junction path change history, you can track cluster events better and can make better use of ONTAP features.

Integrated Performance Management

Managing performance can be a time-consuming task when it is done manually or with limited tools. OnCommand Performance Manager is an integrated component of Unified Manager that monitors systems and alerts about potential performance issues so that optimal performance can be obtained. It also helps troubleshoot, isolates potential problems, and offers directed solutions to performance issues based on system analysis.

Figure 4) OnCommand Unified Manager Favorites dashboard.



OnCommand Performance Manager is easily launched within the Unified Manager dashboard through single sign-on. By using automated analytics, it provides cluster status and cluster health performance metrics, including latency, IOPS, megabytes per second, disk use, and node use.

OnCommand Performance Manager continually monitors and analyzes performance to simplify performance management for NetApp FAS and NetApp All Flash FAS systems. The performance dashboard within Performance Manager sorts clusters by level of importance, allowing simple clicks to drill down into cluster details. The intuitive UI allows navigation, exploration of performance trends, and comparison of the performance for storage objects within clusters. You can quickly and easily view and compare multiple objects, identify areas of concern, and proactively manage and optimize storage performance.

Optimize performance. By using built-in system, dynamic, and user-defined policy thresholds, OnCommand Performance Manager detects and alerts on performance incidents. By using the performance explorer, you can easily compare performance workloads. With the suggested corrective actions, you can quickly troubleshoot performance issues and quickly resolve events. Other features include:

- **Thresholds and alerts.** Built-in system-defined and dynamic thresholds define custom policy thresholds for greater control over FAS and All Flash FAS alerts.
- **Network services.** Performance Manager automatically alerts administrators when off-storage network services—such as network virus scanning, Lightweight Directory Access Protocol (LDAP) authentication, and other network tasks—cause I/O response time to cross a threshold.
- **Policy group limit.** Performance Manager monitors and analyzes all the quality-of-service (QoS) policies that affect workloads. In addition, it indicates when an abnormal or “bully” workload is causing throttling or is affecting response time.
- **Data processing.** Performance Manager provides detailed information about data processing activity and identifies which workloads have changed and have caused a CPU bottleneck.
- **Aggregates.** Performance Manager continuously monitors storage aggregates and confirms that they provide the optimal space, IOPS, and throughput for peak performance. If an aggregate is the source of unacceptable performance, administrators are notified so that they can take corrective action quickly.

- **All Flash FAS speed.** Performance Manager optimizes NetApp All Flash FAS to deliver peak performance, combined with superior flexibility and best-in-class data management for workloads that demand exceptional service.

Streamline management with a single namespace. NetApp Unified Manager and Performance Manager support the FlexGroup capability that was introduced in NetApp ONTAP 9.1. With a single namespace, you can seamlessly scale performance or capacity. When working with multiple physical file systems, a single namespace presents one virtual container as if all the data were centrally located. This configuration allows you to significantly simplify management from one Unified Manager interface.

2.2 System Analysis and Alert Management Component of Self-Healing Storage

The analysis and alert management component is the glue that ties the other components of the self-healing system together. It is responsible for identifying an event through the monitoring system (discussed in Section 2.1) and for triggering an alarm, which in turn initiates the remediation workflow (discussed in Section 2.3).

This document explores the NetApp OnCommand Unified Manager (focused on ONTAP powered solutions) product as the event and alert management system for the NetApp Data Fabric.

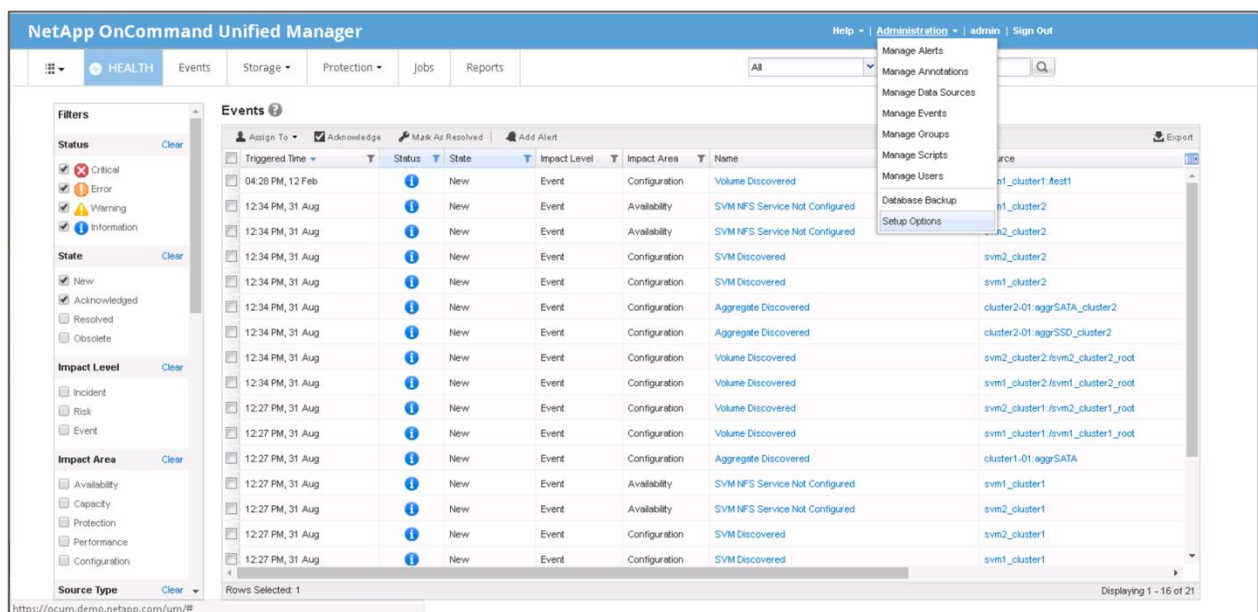
Introduction to OnCommand Unified Manager Event Management Capability

Events are notifications that are generated automatically when a predefined condition occurs or when a particular threshold is breached. These events enable you to take action to prevent issues that can lead to poor performance or service unavailability. Events include an impact area, severity, and an impact level.

Events are categorized by the type of impact area, such as availability, capacity, configuration, or protection. Events are also assigned a severity type and impact level, which help you determine whether immediate action is required.

You can access events in OnCommand Unified Manager either through the Health dashboard (see Figure 3) or through the Events tab.

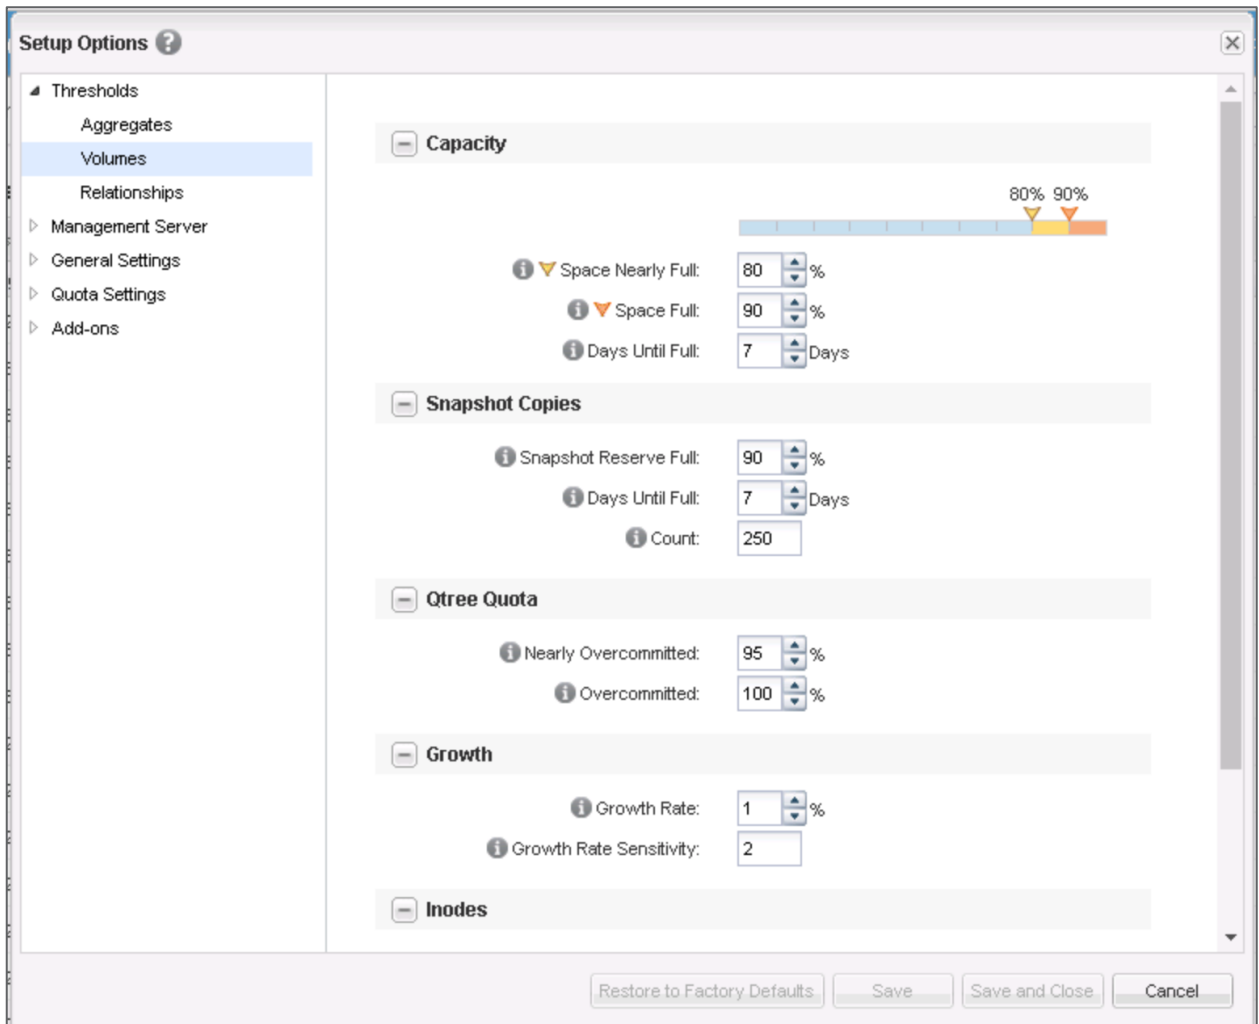
Figure 5) OnCommand Unified Manager events view.



You can configure threshold-driven events through the Setup Options menu, accessible through Administration > Setup Options (see Figure 5).

From the Setup Options menu, expand the Thresholds section and select the relevant storage object (Aggregates, Volumes, or Relationships); see Figure 6.

Figure 6) OnCommand Unified Manager events threshold configuration.



With OnCommand Unified Manager, you can create alerts to notify staff when a particular event is generated. You can create alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and you can associate a script to the alert. These scripts are executed automatically when the specific alert is generated, and they enable you to obtain information about storage objects for which the alert is generated.

OnCommand Unified Manager supports Perl, Shell, Windows PowerShell, and .bat script formats and provides the following information to the script during execution time:

- eventID
- eventSourceID
- eventSourceName

- eventSourceType
- eventState
- eventArgs

Table 1 shows sample argument values that are provided to a script at run time.

Table 1) OnCommand Unified Manager alert script sample arguments.

Arguments	"Volume Space Nearly Full" Event	"Volume Space Full" Event
args[0]	-eventID	-eventID
args[1]	50003	50003
args[2]	-eventName	-eventName
args[3]	Volume	Volume
args[4]	Space	Space
args[5]	Nearly	Full
args[6]	Full	-eventSeverity
args[7]	-eventSeverity	error
args[8]	warning	-eventSourceID
args[9]	-eventSourceID	5428
args[10]	5428	-eventSourceName
args[11]	-eventSourceName	svml_cluster2:/vol_test
args[12]	svml_cluster2:/vol_test	-eventSourceType
args[13]	-eventSourceType	VOLUME
args[14]	VOLUME	-eventState
args[15]	-eventState	NEW
args[16]	NEW	-eventArgs
args[17]	-eventArgs	volNearlyFull=80 volFull=90 dfMountedOn=vol_test dfKBytesTotal=70656 dfKBytesUsed=68500 dfKBytesPercent=96.9485960144927 5 dfInodesTotal=31122 dfInodesUsed=102 dfInodesPercent=0.32774243300559 09 dfKBytesAutoMaxSize=1258288 aggrDfBytesAvail=93327323136 isAutosizeEnabled=false autoIncrementSizeBytes=53686272 isVolSpaceGuaranteeVolume=true
args[18]	volNearlyFull=80 volFull=90 dfMountedOn=vol_test dfKBytesTotal=70656	

Arguments	"Volume Space Nearly Full" Event	"Volume Space Full" Event
	dfKBytesUsed=61344 dfKBytesPercent=86.82065217391305 dfInodesTotal=31122 dfInodesUsed=99 dfInodesPercent=0.31810294968189706 dfKBytesAutoMaxSize=1258288 aggrDfBytesAvail=93327298560 isAutosizeEnabled=false autoIncrementSizeBytes=53686272 isVolSpaceGuaranteeVolume=true	

By using scripts, integration with the Unified Manager alert mechanism, and the Workflow Automation workflow engine, you can leverage the WFA RESTful API interface.

2.3 Remediation Orchestration Component of Self-Healing Storage

In a world filled with opinions, habits, and "it depends," where operational best practices can change with every release of a new service or adoption of a new technology, organizations need simplicity and flexibility. An open orchestration layer that readily adapts to a multitude of technologies and consumption models is required. Key to automation are a decision engine and enforced boundaries that are stipulated by the infrastructure owners. These key tenets are critical whether infrastructure resides in private, public, or hybrid clouds.

Many orchestration tools are available on the market. However, one tool is not only fully integrated with all the components of the Data Fabric, but is also available to all NetApp customers with no additional software license requirements. That tool is OnCommand Workflow Automation.

Introduction to OnCommand Workflow Automation

OnCommand Workflow Automation (WFA) is a software solution that automates storage management tasks, such as provisioning, migration, decommissioning, data protection configurations, and data cloning operations.

NetApp OnCommand Workflow Automation delivers on the Data Fabric vision, providing automation and integration to meet the demands of today's evolving IT organizations. Whether managing storage on the premises or in a cloud environment, Workflow Automation makes it easy to quickly create simple or complex workflows. Storage administrators can create storage workflows for the most frequent tasks and make them available to consumers for one-click automation. Storage architects can automate time-consuming, complex processes to meet cost-saving objectives, a common goal when moving data to the cloud. OnCommand WFA helps provide consistency, predictability, and standardization across the Data Fabric.

With Workflow Automation, you can integrate storage workflows with existing IT orchestration processes for fast, seamless delivery of services to your business consumers. And by implementing Workflow Automation in storage management operations, you can also standardize and integrate processes to align with NetApp best practices for data management.

A workflow is a repetitive and procedural task that consists of sequential steps. The goal is a consistent, repeatable, and predictable result. Just a few examples of what you can do with WFA include:

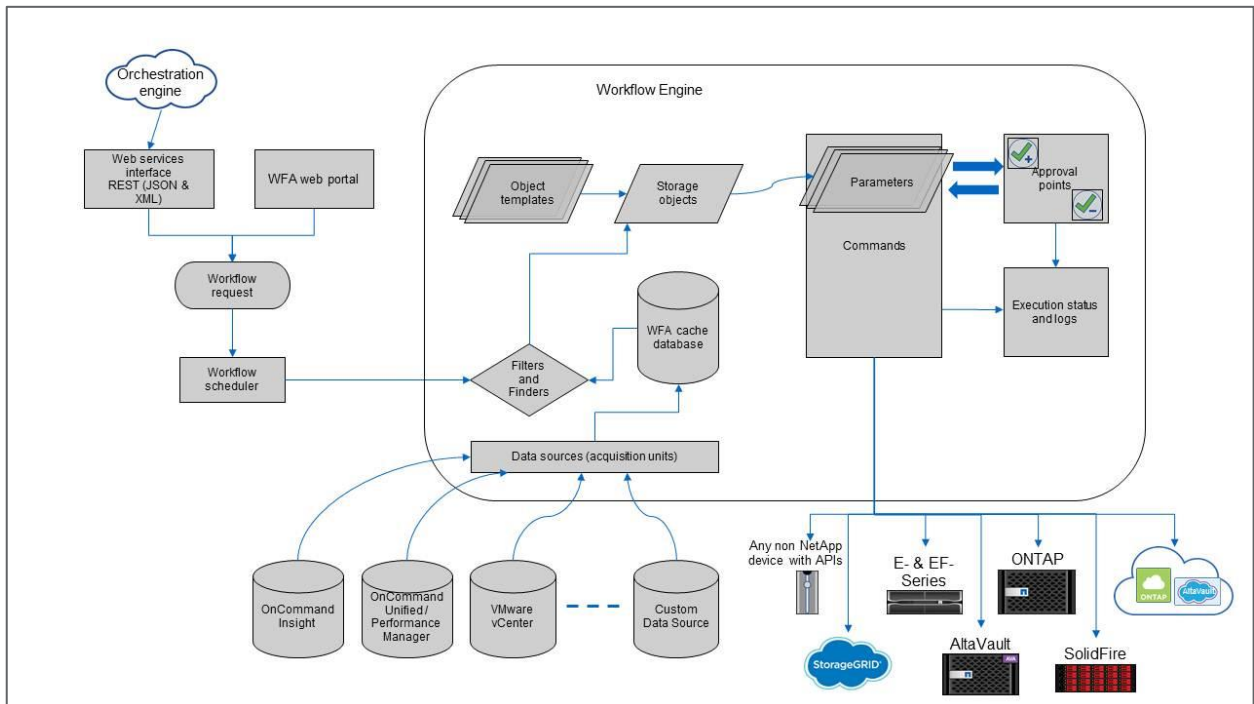
- Provision, migrate, or decommission storage for databases or file systems.
- Automate your data protection process to a disaster recovery location on the premises or in the cloud.
- Set up storage for an application as part of an end-to-end orchestration process.
- Clone and set up a new development and testing environment, then tear it down when it's no longer needed.

- Set up the FlexPod® Datacenter system or virtual desktops, or conduct a centralized NetApp SnapCenter® software activation.

Storage architects can define workflows to follow best practices and to meet organizational requirements, such as the following:

- Adhering to required naming conventions
- Setting unique options for storage objects
- Selecting appropriate resources
- Integrating internal configuration management database (CMDB) and ticketing applications

Figure 7) OnCommand Workflow Automation architecture.



Workflow Automation integrates with OnCommand Unified Manager to collect data about your storage environment (see Figure 7). With this integration, you can leverage information in numerous OnCommand Unified Manager instances to automate storage and data protection processes globally. You can also automate processes across multivendor storage environments by leveraging OnCommand Insight. Simply download the OnCommand Insight Connector from the Storage Automation Store to integrate OnCommand Insight as a data source for Workflow Automation.

In addition, Workflow Automation can connect to internal systems to collect information to use in resource selection, to open trouble tickets, and more. These capabilities allow you to accelerate storage service delivery and reduce time to market.

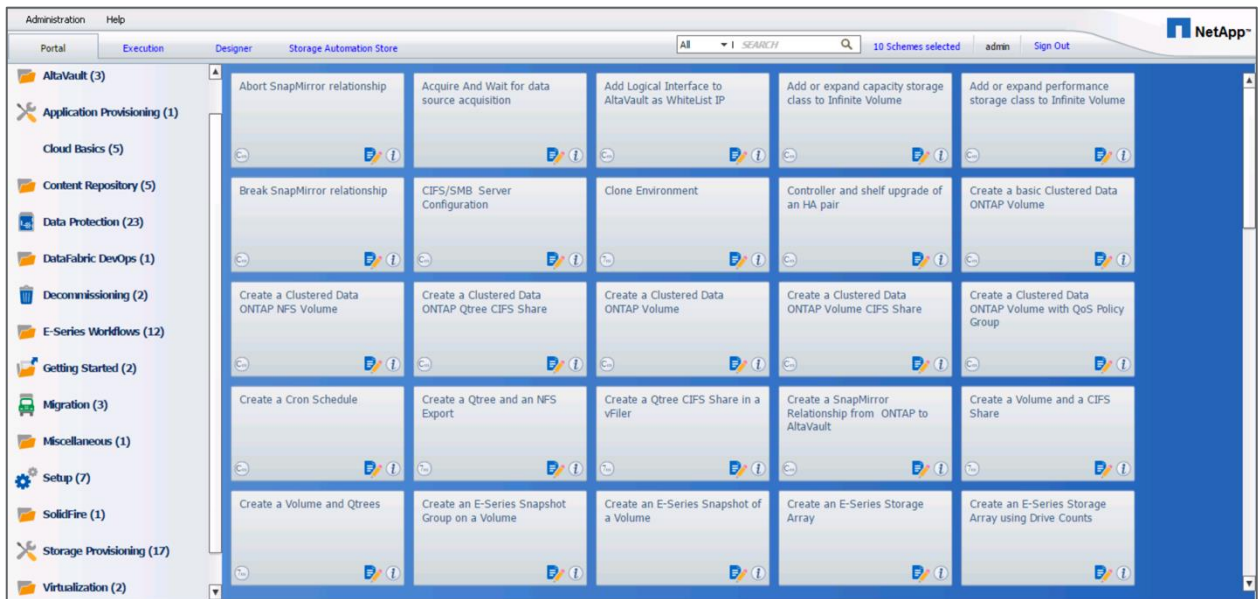
For example, Workflow Automation integrates with VMware vRealize Operations, VMware vRealize Automation, and VMware vCloud Automation Center to deliver custom IT as a service with a single click. To invoke your workflows from an external portal or from the data center orchestration software, simply use the REST APIs that Workflow Automation provides. OnCommand Workflow Automation is a flexible framework. The same approach can be taken to integrate and automate with other orchestrators such as Cisco UCS Director or Microsoft System Center Orchestrator.

OnCommand Workflow Automation features include:

-
- **An execution portal** to invoke workflows, verify the status of workflow execution, and access logs.
- **Designer portal** to build workflows. The designer portal includes several building blocks, such as commands, templates, finders, filters, and functions, that are used to create workflows. The designer enables you to include advanced capabilities with workflows such as automated resource selection, row repetition (looping), and approval points. The designer portal also includes building blocks, such as dictionary entries, cache queries, and data source types, for caching data from external systems.
- **An administration portal** for tasks such as setting up WFA, connecting to data sources, and configuring user credentials.
- **Web services interfaces** to invoke workflows from external portals and data center orchestration software, with full support for RESTful APIs that leverage JSON and XML.
- **The Storage Automation Store** to download WFA packs (workflows and commands that are provided by either NetApp, NetApp partners, or the WFA community).

When you first log in to the Workflow Automation software, you are presented with the production-ready workflows portal.

Figure 8) WFA Portal view.



As Figure 8 shows, you can divide workflows into categories that are based on context. Examples of workflow categories include AltaVault, E-Series, ONTAP, SolidFire, DevOps, setup, migration, virtualization, and cloud. The view of available workflows and categories depends on the user's account and the role-based access control (RBAC) rules that are associated with those workflows.

For information about a workflow's execution status (log), schedules, and configured data sources (acquisition units through which WFA collects information about the infrastructure), go to the Execution tab (Figure 9).

Figure 9) WFA Execution view.

Job #	Name	Start Time	End Time	Status	Completed	Submitted By	Submitted At	Execution Comment	Sched
9722	Modify_Volume_Inode_Count	02/09/17 10:13:17 PM	02/09/17 10:13:24 PM	Successful	1/1	demolayaron	02/09/17 10:13:16 PM		
9721	Modify_Volume_Inode_Count	02/09/17 10:11:20 PM	02/09/17 10:11:27 PM	Successful	1/1	demolayaron	02/09/17 10:11:20 PM		
9719	Modify_Volume_Inode_Count	02/09/17 10:04:44 PM	02/09/17 10:04:51 PM	Failed	0/1	demolayaron	02/09/17 10:04:43 PM		
9718	Modify_Volume_Inode_Count	02/09/17 9:57:44 PM	02/09/17 9:57:59 PM	Failed	0/1	demolayaron	02/09/17 9:57:42 PM		
1037	Transfer SnapMirror relationship	10/08/15 3:50:22 AM	10/08/15 3:50:27 AM	Successful	1/1	admin	10/08/15 3:50:19 AM		
1036	Transfer SnapMirror relationship	10/08/15 3:50:05 AM	10/08/15 3:50:09 AM	Successful	1/1	admin	10/08/15 3:50:02 AM		
1033	Transfer SnapMirror relationship	10/08/15 2:21:23 AM	10/08/15 2:21:28 AM	Successful	1/1	admin	10/08/15 2:21:20 AM		
1032	Transfer SnapMirror relationship	10/08/15 2:21:23 AM	10/08/15 2:21:27 AM	Successful	1/1	admin	10/08/15 2:21:20 AM		
1031	Transfer SnapMirror relationship	10/08/15 2:20:47 AM	10/08/15 2:20:51 AM	Successful	1/1	admin	10/08/15 2:20:44 AM		
1030	Transfer SnapMirror relationship	10/08/15 2:20:46 AM	10/08/15 2:20:51 AM	Successful	1/1	admin	10/08/15 2:20:44 AM		
1022	Transfer SnapMirror relationship	10/07/15 10:58:53 PM	10/07/15 10:58:58 PM	Successful	1/1	admin	10/07/15 10:58:50 PM		
1021	Transfer SnapMirror relationship	10/07/15 10:58:38 PM	10/07/15 10:58:46 PM	Successful	1/1	admin	10/07/15 10:58:35 PM		
631	Modify SnapMirror relationship	09/29/15 8:12:36 PM	09/29/15 8:12:41 PM	Successful	1/1	admin	09/29/15 8:12:35 PM		

The Execution tab is also where you can store credentials information. One of the advantages of using Workflow Automation instead of scripts is the increased security that results from not exposing user names and passwords. WFA stores credentials information in a secure and encrypted format. This information cannot be accessed directly by any user.

The Designer tab (Figure 10) is where automation architects can design and implement new workflows.

Figure 10) WFA Designer view.

Certification	Name	Scheme	Entry Version	Description
	Acquire And Wait for data source acquisition		1.0.0	This workflow enables you to Acquire and wait for data source from ONTAP instances on AWS. Acquire and wait for data source from cloud ONTAP instance involves the following steps: 1. Acquiring cloud ONTAP instances from AWS. 2. Wait to acquiring for cloud ONTAP instances from AWS. Prerequisites: 1) OCOM version: 2.1 or above 2) WFA Server: WFA 3.1 GA P1 3) Local WFA admin credentials must be added into WFA credentials (Host Name = localhost) 4) OCOM server credentials should be added into WFA credentials
	Add Logical Interface to AltaVault as Whitelist IP	altavault_cm_storage	1.0.0	This workflow Add Logical Interface to AltaVault as Whitelist IP. It adds the intercluster LIP. This workflow valid only for clusters running with Data ONTAP version 9.1.0 onwards. Add Logical Interface to AltaVault as Whitelist IP involves the following steps: 1) Create a new intercluster LIP. 2) Add intercluster LIP to AltaVault as Whitelist IP Prerequisite: 01) AltaVault up and running and it should intergate with cloud. 02) Add AltaVault credentials into WFA. 03) Add AltaVault data source script to WFA and acquire.
	Add or expand capacity storage class to Infinite Volume	cm_storage	1.0.0	Adds or expand a capacity storage class to an Infinite Volume. The capacity storage class is 1. The data constituents are provisioned on aggregates with 'SATA' type of disks. 2. Data constituents are provisioned with a 'volume' space guarantee. 3. Data constituents have deduplication enabled with efficiency priority as 'background' and scheduled to run from midnight - 6 AM.

Not every aspect of WFA can be covered in this document. Following are a few of the more important components of this tab:

- **Commands** are the building blocks of the workflow. They are the steps that the workflow executes in sequence. Workflow Automation provides many commands as part of the basic installation package.

You can add other commands by downloading them from the Storage Automation Store, by finding them through community sharing, by purchasing them from professional service organizations, or by writing them in-house. WFA commands are scripts that are written in PowerShell or Perl.

- **Filters** are used by the WFA decision-making engine to identify potential resources for a given task. For example, if a workflow needs to identify a solid-state drive (SSD) storage pool resource, a filter to find all relevant storage pools (aggregate) that meet the required criteria can be employed. These criteria can include RAID type, media type, nonroot, capacity requirements, and so on.
- **Finders** are used to identify the “best” resource among all potential resources for a task. Finders use filters to locate all relevant resources and use prioritization logic to determine the one optimal resource for the requirements. An example of an optimal resource is the aggregate with the most free space or the aggregate with the greatest amount of available IOPS.
- **Workflows** are the collection of commands with the additional logic required to define a task to be automated.

Another important feature is the Storage Automation Store (Figure 11).

Figure 11) WFA Storage Automation Store view.

The screenshot displays the NetApp Storage Automation Store interface. At the top, there's a navigation bar with 'Administration' and 'Help' links. Below it, a tabbed interface shows 'Portal', 'Execution', 'Designer', and 'Storage Automation Store'. The 'Storage Automation Store' tab is active, showing a search bar and '10 Schemes selected'. The main content area is titled 'Available Packs' and lists several packs with columns for Certification, Name, Description, Pack Version, and Author. The packs listed are: 'Create a Qtree and an NFS Export - Perl', 'Datasource Script for VMware vCenter 6.0', 'Day-0 c-mode cluster build setup', and 'E-Series Pack'. The 'Pack Details' section below shows the details for the 'Create a Qtree and an NFS Export - Perl' pack, including its version (1.0.0), description, author (NetApp), and certification (NETAPP). A 'Download Pack' button is visible at the bottom right.

Certification	Name	Description	Pack Version	Author
	Create a Qtree and an NFS Export - Perl	Create a Qtree and an NFS Export -perl workflow will create a new qtree and NFS export in an existing selected volume or in new volume based on the selection criteria. It will be exported over NFS with the appropriate access privileges to the specified hosts. Create volume, create qtree and NFS export commands has been implemented using perl.	1.0.0	NetApp
	Datasource Script for VMware vCenter 6.0	The Data source script for VMware vCenter 6.0 pack contains the data source script to acquire VMware vCenter server data into the OnCommand Workflow Automation (WFA) database for authorized users to perform heterogeneous IT automation activities. This pack enables authorized users to create custom workflows using the imported data in WFA.	1.1.0	NetApp
	Day-0 c-mode cluster build setup	The Day-0 c-mode cluster build setup workflow will enable storage administrators to configure and customize the two node cluster. You can install the configurations on a new cluster or a cluster node before deploying to the end user	1.1.0	netapp
	E-Series Pack	The E-Series workflows pack enables you to quickly configure your E-series storage system. You do not have to use the SANtricity Manager to configure your E-Series storage system. You must execute the workflows from the E-series pack in the following recommended sequence: 1. Create an E-Series Storage Array 2. Create an E-Series Storage Array using Drive Count	1.0.0	netapp

Pack Details
Create a Qtree and an NFS Export - Perl

Pack Version: 1.0.0 [View Online](#)

Description: Create a Qtree and an NFS Export -perl workflow will create a new qtree and NFS export in an existing selected volume or in new volume based on the selection criteria. It will be exported over NFS with the appropriate access privileges to the specified hosts. Create volume, create qtree and NFS export commands has been implemented using perl.

Author: NetApp

Certification: NETAPP

[Download Pack](#)

The Storage Automation Store is a portal from which you can download certified software packs to use with OnCommand Workflow Automation, OnCommand Unified Manager, and OnCommand Performance Manager to extend data management capabilities.

3 Self-Healing Systems Implementation

The use cases for self-healing storage systems are almost endless and range from resource (capacity and performance) management all the way to proactive issue avoidance activities. This technical report discusses the following sample use cases in detail:

- **Use case 1: automatic capacity management.** In this use case, the system reacts to a NetApp FlexVol® volume running out of capacity and checks the feasibility of a space allocation increase while maintaining a maximum storage pool (aggregate) utilization of X%. If the host storage pool does not have enough capacity, the system investigates similar tier storage pools; if a match is found, the volume is relocated and is then expanded.
- **Use case 2: automatic inode management in high-file-count (HFC) environments.** In this use case, the system reacts to a FlexVol volume running out of available inodes and increases the number of inodes to maintain X% inode utilization.
- **Use case 3: automatic performance management (tiering).** In this use case, the system leverages the I/O density (I/O per terabyte stored) metric to identify the proper performance tier for each FlexVol volume. If necessary, the system nondisruptively moves the volumes to a different storage pool (aggregate) or compute node type.

As stated previously, the use cases for self-healing storage systems are almost endless. Some additional use cases include:

- Proactive performance metric collection (such as the NetApp Perfstat collection tool) that is based on the performance profile of a system (for example, a performance latency trigger)
- Adaptive QoS management that is based on the I/O density of a workload
- Conformance validation of protection settings that is based on the protection tier of an asset
- Detection and remediation of configuration drift
- Detection and remediation of ransomware attacks

3.1 Use Case 1: Automatic Capacity Management

In this use case, the system reacts to a FlexVol volume running out of capacity and checks the feasibility of a space allocation increase while maintaining a maximum storage pool (aggregate) utilization of X%. If the host storage pool does not have enough capacity, the system investigates the availability of similar storage pools. If a match is found, the volume is relocated and is then expanded. You should compare and contrast this approach with the internal self-healing capability of the ONTAP native volume auto-grow feature.

Creating a Remediation Orchestration Workflow

The first step in the implementation is to design and create a remediation workflow in OnCommand Workflow Automation.

Best Practice

A best practice for all automation workflow design is to clearly specify the requirements, develop an approach that the workflow will follow, and translate it to a logic-based flowchart model before starting the actual implementation.


Table 2 lists the requirements and implementation methodology for this workflow.

Table 2) Use case 1 requirements and methodology.

Workflow Requirements	Workflow Implementation Methodology
<ul style="list-style-type: none"> • Increase the size of a FlexVol volume. • Maintain aggregate space utilization at or below X%. • If necessary, relocate the volume to a different aggregate of a similar tier. 	<ul style="list-style-type: none"> • Find all relevant information about the FlexVol volume. • Determine whether the hosting aggregate can sustain the growth of the volume without violating the utilization guideline. • If it can, increase the size of the FlexVol volume. • If it can't, find another aggregate of the same tier, move the volume, and then increase the volume size.

The flowchart in Figure 12 represents the logic of this workflow.

Figure 12) Use case 1 flowchart.




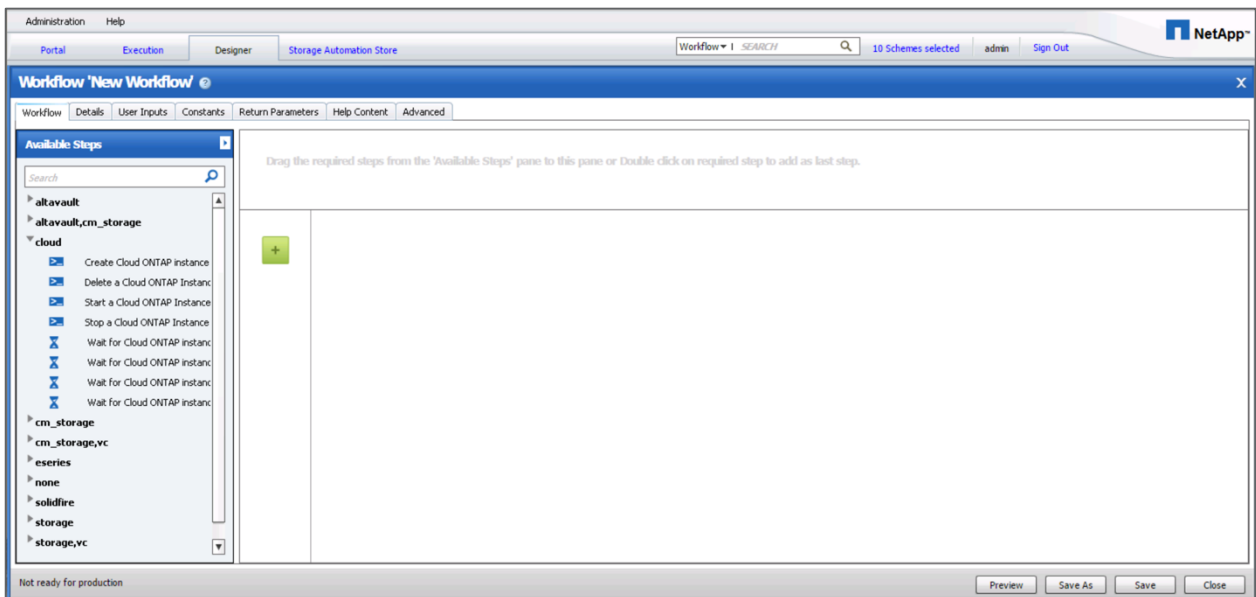

Now you are ready to create your workflow. To create a workflow, log in to the Workflow Automation portal, click the Designer tab, and click the New () button.

Figure 13) WFA new workflow design view.

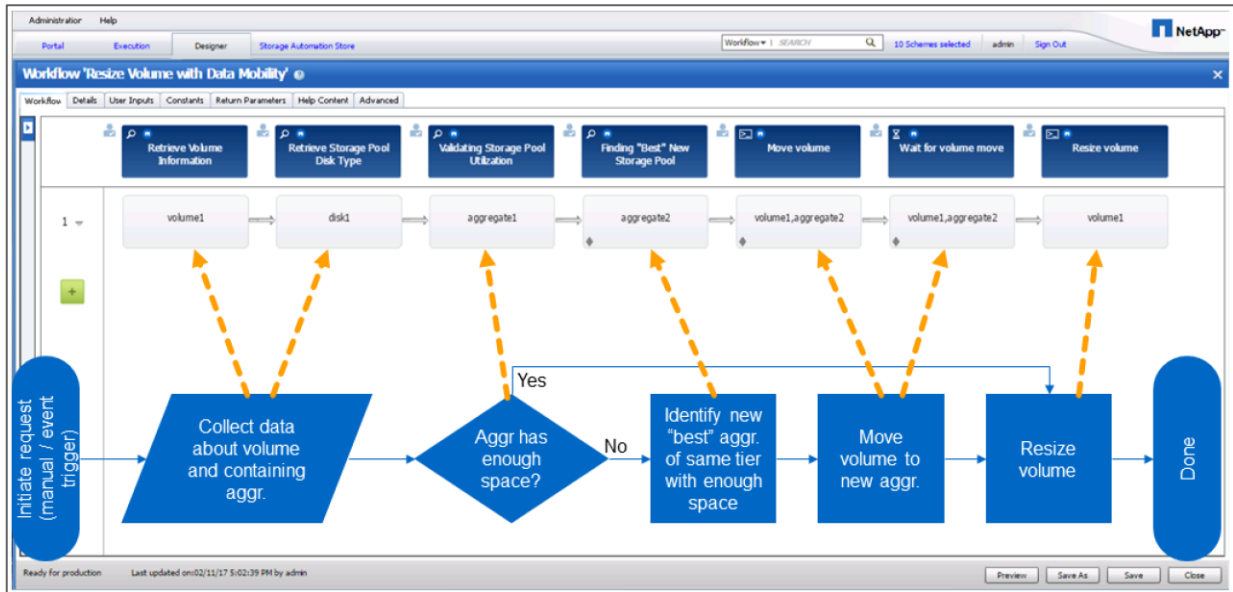


Remember that a workflow is a repetitive and procedural task that consists of sequential commands. As Figure 13 shows, all the available commands in WFA are listed on the left panel and are divided in contextual categories (AltaVault, Cloud, and so on). Building your workflow starts with the simple task of identifying the relevant commands and then dragging them in the correct order to the workflow design pane.

Figure 14 is a view of the completed workflow with the layered implementation flowchart above it. In WFA, the symbol  indicates a conditional run. In this case, the commands with that symbol run only if the

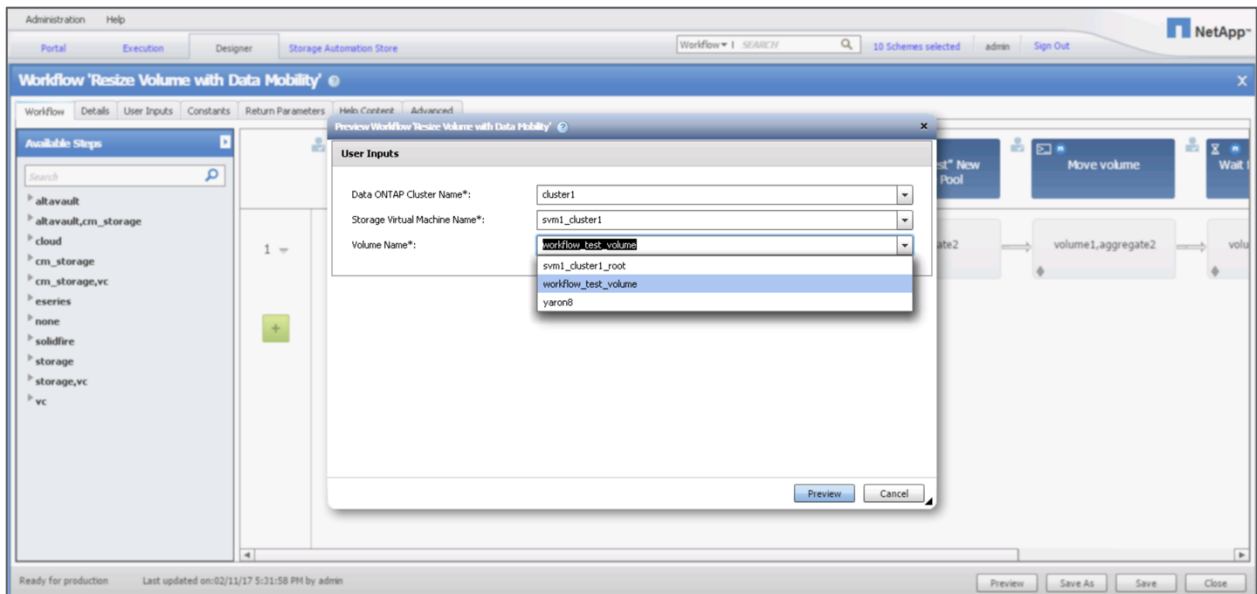
hosting aggregate space validation condition is not met. In addition, you can customize the workload to any level of utilization that you want. Further, you can use tailored heuristics in deciding how much to increase the volume's size.

Figure 14) Use case 1 WFA workflow view.



Workflow Automation makes it very easy to test a workflow before actual execution: Simply click the Preview button (bottom right of Figure 14).

Figure 15) Use case 1: previewing a workflow run.



When you click the Preview button, a pop-up menu appears (Figure 15). You are asked to provide the relevant required workflow parameters (in this example case, the cluster name, storage virtual machine name, and volume name). The options in the drop-down menus are not hard-coded (although they could be), but, rather, are data that WFA has collected about the environment from its data sources.

Once again, select the correct parameters and click the Preview button. The workflow runs, and “would be” results are displayed (Figure 16). In this case, the volume would be moved to a new aggregate before its size would be increased.

Figure 16) Use case 1 results of a dry run.



Creating a Workflow Alert Launch Script for OnCommand Unified Manager

Before creating the relevant OnCommand alert, you should create a script (both PowerShell and Perl samples are provided here) that can be attached to the alert for automated execution.

Following is the PowerShell version of the script (OCUM_WFA_Capacity_Event.ps1):

```
# 2017-02-09 yaron@netapp.com
# OCUM_WFA_Capacity_Event.ps1
#
# This script is meant to be used as part of an automated self-healing process
# that remediates high volume capacity utilization scenarios. It is meant to be used
# in conjunction with OCUM 7.x Windows-based systems and OnCommand WFA 4.1.x.
#
# IMPORTANT: This code assumes a single ONTAP cluster is present in the
#             environment for code simplification. If environment contains more than one
#             cluster see Appendix A.
#
# (c) 2017 NetApp Inc., All Rights Reserved
#
# Check the following link in case you run into SSL certificate issues
# https://d-fens.ch/2013/12/20/nobrainier-ssl-connection-error-when-using-powershell/
#=====
# CHANGE IF DESIRED PRIOR TO UPLOADING TO OCUM SERVER
$cluster_name = "cluster1"
$wfa_rest_server = "https://wfa.demo.netapp.com/rest"
$wfa_username = "admin"
$wfa_password = "Netapp1!"
#=====
# DO NOT CHANGE CODE BELOW THIS LINE
#=====
if ($args[6] -eq "Full"){
    # This is an "Inodes Nearly Full" event
    $source_name = $args[12]
    $event_state = $args[16]
} else {
    # This is an "Inodes Full" event
    $source_name = $args[11]
    $event_state = $args[15]
}
$event_id = $args[1]

# Ignore all non-new events
if (($event_state.ToString()).ToLower() -ne "new"){
    exit
}

# Extract SVM and Volume names from OCUM-provided arguments
```

```

$svm_name = $source_name.Split("{:/}")[0]
$volume_name = $source_name.Split("{:/}")[2]

# Prepare WFA credentials
$securePassword = ConvertTo-SecureString $wfa_password -AsPlainText -Force
$cred = New-Object System.Management.Automation.PSCredential -ArgumentList
$wfa_username,$securePassword

# Find RESTful execution url for WFA workflow
$restCommand = $wfa_rest_server + "/workflows?name=Resize_Volume_with_Data_Mobility"
try {
    $output = Invoke-RestMethod -Method Get -Uri $restCommand -Credential $cred
} catch {
    # An error occurred. Either workflow does not exist or credentials are wrong
    exit
}
$workflow_execution_uri = ($output.collection.workflow.link | ?{$_rel -eq
'execute'}).href

# Initiate workflow execution
$body = @"
{
    "comments":"OCUM triggered workflow. Event ID: $event_id",
    "userInputValues":[
        {
            "key":"cluster_name",
            "value":"$cluster_name"
        },
        {
            "key":"svm_name",
            "value":"$svm_name"
        },
        {
            "key":"volume_name",
            "value":"$volume_name"
        }
    ]
}
"@

$output = Invoke-RestMethod -Method Post -Uri $workflow_execution_uri -Credential $cred -Body
$body -ContentType "application/json"

# Find RESTful job monitoring link for WFA workflow
$job_status_uri = ($output.job.link | ?{$_rel -eq 'self'}).href
if ($job_status_uri.Count -gt 1){$job_status_uri = $job_status_uri[0]}

# Wait until workflow job either completes successfully or fails
do {
    Start-Sleep -Seconds 2
    $output = Invoke-RestMethod -Method Get -Uri $job_status_uri -Credential $cred
    $jobStatus = $output.job.jobStatus.jobStatus
} while (($jobStatus -ne "COMPLETED") -and ($jobStatus -ne "FAILED"))

```

Note: To use PowerShell or .bat scripts, OnCommand Unified Manager must be installed on a Windows Server. If it is installed on a Linux server, use either Perl or Shell.

The following Perl script is provided as a reference in case PowerShell cannot be used.

```

# 2017-02-09 yaron@netapp.com
# OCUM_WFA_Capacity_Event.pl
#
# This script is meant to be used as part of an automated self-healing process
# that remediates high volume capacity utilization scenarios. It is meant to be used in
# conjunction with OCUM 7.x Windows-based systems and OnCommand WFA 4.1.x.
#
# IMPORTANT: This code assumes a single ONTAP cluster is present in the
# environment for code simplification. If environment contains more than one
# cluster see Appendix A.

```

```

#
# (c) 2017 NetApp Inc., All Rights Reserved
#

use REST::Client;
use JSON;
use MIME::Base64;

# Unmark the next line if you run into SSL certificate issues
#$ENV{PERL_LWP_SSL_VERIFY_HOSTNAME}=0;

#####
# CHANGE IF DESIRED PRIOR TO UPLOADING TO OCUM SERVER
my $cluster_name = 'cluster1';
my $wfa_rest_server = 'https://wfa.demo.netapp.com';
my $wfa_username = 'admin';
my $wfa_password = 'Netapp1!';
#####

#####
# DO NOT CHANGE CODE BELOW THIS LINE
#####
if ($ARGV[6] == 'Full'){
    # This is an "Volume Space Nearly Full" event
    my $source_name = $ARGV[12];
    my $event_state = $ARGV[16];
} else {
    # This is an "Volume Space Full" event
    my $source_name = $ARGV[11];
    my $event_state = $ARGV[15];
}
my $event_id = $ARGV[1];

# Ignore all non-new events
if ($event_state == 'NEW'){
    exit;
}

# Extract SVM and Volume names from OCUM-provided arguments
my ($svm_name,$volume_name) = split /\:\/\/,$source_name;

# Find RESTful execution url for WFA workflow
my $headers = {Accept => 'application/json', Authorization => 'Basic ' .
    encode_base64($wfa_username . ':' . $wfa_password), 'Content-Type' => 'application/json'};
my $client = REST::Client->new();
$client->setHost($wfa_rest_server);
$client->GET('/rest/workflows?name=Resize_Volume_with_Data_Mobility', $headers);
my @response_json = @{ decode_json($client->responseContent()) };
my $workflow_execution_uri = '/rest/workflows/' . $response_json[0]->{'uuid'} . '/jobs';

# Initiate workflow execution
my $json_body = '{
    "comments":"OCUM triggered workflow. Event ID: ' . $event_id . '",
    "userInputValues":[
        {
            "key":"ClusterName",
            "value":"' . $cluster_name . '"
        },
        {
            "key":"SvmName",
            "value":"' . $svm_name . '"
        },
        {
            "key":"VolumeName",
            "value":"' . $volume_name . '"
        }
    ]
}';

# Find RESTful job monitoring link for WFA workflow

```



```

$client->POST($workflow_execution_uri, $json_body, $headers);
my $response = decode_json($client->responseContent());
$job_status_uri = $workflow_execution_uri . '/' . $response->{'jobId'};

# Wait until workflow job either completes successfully or fails
do {
    sleep 2;
    $client->GET($job_status_uri, $headers);
    my $response = decode_json($client->responseContent());
    $jobStatus = $response->{'jobStatus'}{'jobStatus'};
} while (($jobStatus != 'COMPLETED') && ($jobStatus != 'FAILED'));

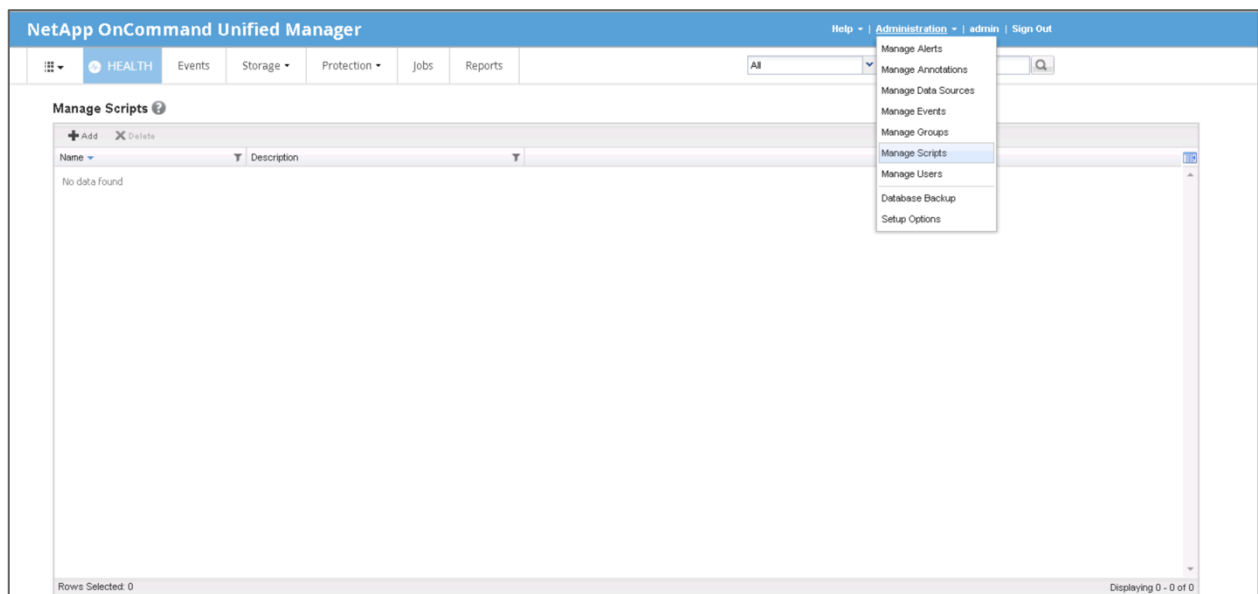
```

Creating an OnCommand Unified Manager Alert

To create an OnCommand Unified Manager alert:

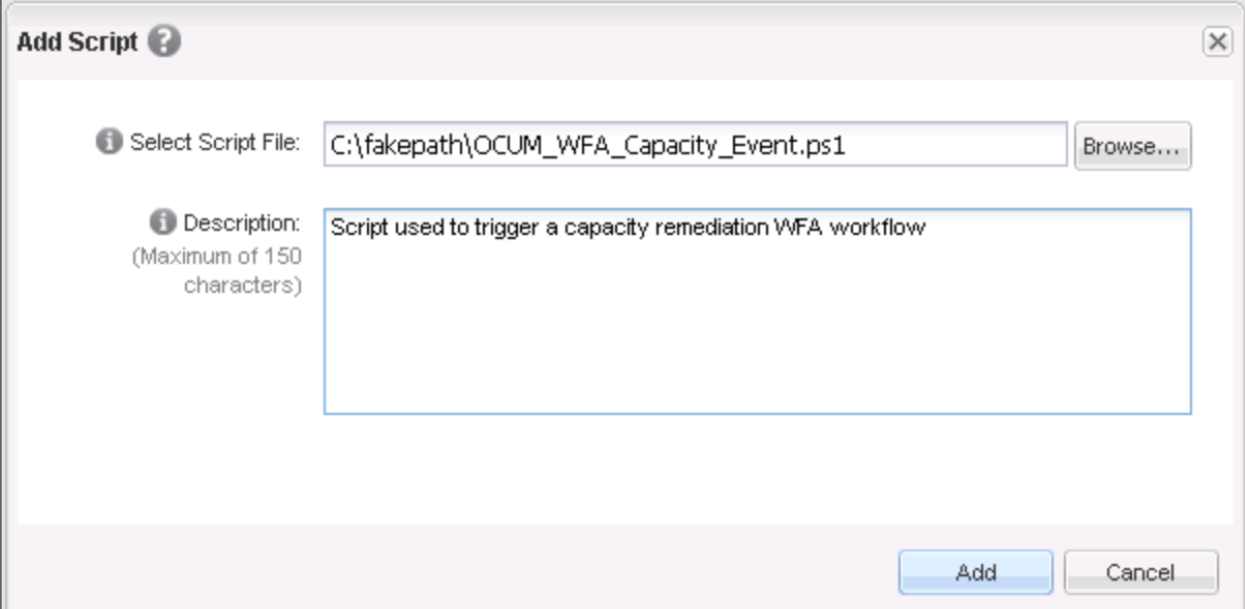
1. Start with uploading the workflow launch script to the OnCommand Unified Manager server. To upload the script to the OnCommand Unified Manager server, select Administration > Manage Scripts, and click the Add button (Figure 17).

Figure 17) OnCommand Unified Manager Manage Scripts view.



2. Browse to the file to upload and provide a brief description (optional). See Figure 18.


Figure 18) OnCommand Unified Manager Add Script form.



The 'Add Script' form is a dialog box with a title bar containing a question mark icon and a close button. It contains two main sections. The first section is labeled 'Select Script File:' and includes a text input field with the path 'C:\fakepath\OCUM_WFA_Capacity_Event.ps1' and a 'Browse...' button. The second section is labeled 'Description:' and includes a text area with the text 'Script used to trigger a capacity remediation WFA workflow'. Below the text area, it says '(Maximum of 150 characters)'. At the bottom right of the form are two buttons: 'Add' and 'Cancel'.

3. Create an alert by selecting Administration > Manage Alerts and clicking the Add button (Figure 19).

Figure 19) OnCommand Unified Manager Manage Alerts view.



- a. Provide an alert name and a description (Figure 20).

Figure 20) OnCommand Unified Manager, first step in adding an alert.

The 'Add Alert' dialog box is shown with a sidebar on the left containing 'Name', 'Resources', 'Events', and 'Actions'. The 'Name' tab is selected. The 'Alert Name' field contains 'High volume capacity utilization alert'. The 'Alert Description' field contains 'Alert used in a self healing automated workflow'. At the bottom right are 'Save' and 'Cancel' buttons.

- b. Select the resources to which the alert will be applied. In this example, all the volumes in the environment are selected (Figure 21).

Figure 21) OnCommand Unified Manager, second step in adding an alert.

The 'Add Alert' dialog box is shown with the 'Resources' tab selected. The 'Include' button is active. The 'Select the resources to be included in the alert.' section shows a dropdown menu set to 'Volume' and a search field 'Name contains' with the placeholder 'Begin typing to find by name'. Below this are two lists: 'Available Resources' and 'Selected Resources'. The 'Available Resources' list contains 9 items: cluster1-01:/vol0, cluster1-02:/vol0, cluster2-01:/vol0, cluster2-02:/vol0, svm1_cluster1:/svm1_cluster1_root, svm1_cluster1:/test1, svm1_cluster2:/svm1_cluster2_root, svm2_cluster1:/svm2_cluster1_root, and svm2_cluster2:/svm2_cluster2_root. The 'Selected Resources' list contains '<< All Volumes >>'. At the bottom, it says '9 resources found' and 'Resource included within "<< >>" represents a dynamic rule.' At the bottom right are 'Save' and 'Cancel' buttons.

- c. Select the events that will trigger this alert. In this example, the Volume Space Nearly Full (Warning) event and the Volume Space Full (Error) event are selected (Figure 22).

Figure 22) OnCommand Unified Manager, third step in adding an alert.

Add Alert ?

Name >

Resources >

Events <

Actions >

Select the events that will trigger the alert

Event Severity: Critical
Error
Warning

Event Name Contains: Begin typing to find by name

Matching Events:

- All Error Events
- Inodes Full
- Performance Error
- Snapshot Reserve Fully Consumed
- Too Many Snapshot Copies
- Volume Days Until Full
- Volume Efficiency Operation Error
- Volume Move Status - Failed
- Volume Gtree Quota Overcommitted
- Volume Snapshot Reserve Days Until Full

Selected Events:

- Volume Space Nearly Full
- Volume Space Full

Save Cancel

- d. Select the actions to be taken when the alert is triggered. This example has selected to execute the script that was uploaded earlier in the process (Figure 23).

Figure 23) OnCommand Unified Manager, fourth step in adding an alert.

Add Alert ?

Name >

Resources >

Events >

Actions <

Send Email _____

Alert these users: Type an email address or username. ▼

Name	Email
------	-------

Notification Frequency: ☒ Once

☐ Remind every 15 ▼ minutes

☐ Remind every 15 ▼ minutes between 08 and 17 hours

Issue SNMP trap: ☒

Execute Script _____

Select script to execute: OCUM_WFA_Capacity_Event.ps1 ▼

Save Cancel

That's it! You have created a fully automated process that deals with volumes running out of capacity.

Testing the Self-Healing Scenario

To test use case 1:

1. Create a 1GB volume (Figure 24).

Figure 24) Testing use case 1: Create a volume.

NetApp OnCommand System Manager

Cluster: cluster1

Storage Virtual Machines: svm1_cluster1

Storage: svm1_cluster1_root

Volumes

Name	Aggregate	Status	Thin Provisioned	% Used	Available Space	Total Space	Storage Efficiency	Is Volume Moving
automated_capacity_test	aggrSATA	Online	No	0	1023.85 MB	1 GB	Disabled	No
svm1_cluster1_root	aggrSATA	Online	No	5	18.86 MB	20 MB	Disabled	No

General

Name:	automated_capacity_test	Autogrow Mode:	Disabled
Status:	Online	Autogrow Maximum Size:	Disabled
Maximum Files:	31.12k	Autogrow Increment Size:	Disabled
Current Files:	97	Snapshot Autodelete:	Disabled
Language:	C.UTF-8 (POSIX with UT...)	Snapshot Autodelete Commitment:	NA
Unicode:	Enabled	Junction Path:	/automated_capacity_test
Cache Policy:	NA	Export Policy:	default

Nodes: +

Details Space Allocation Snapshot Copies Storage Efficiency Data Protection Volume Move Details

- Copy enough data to increase the volume's capacity utilization to above 80% and observe that an event is created (Figure 25 and Figure 26, respectively).

Figure 25) Testing use case 1: volume capacity status.

Volumes								
Name	Aggregate	Status	Thin Provisioned	% Used	Available Space	Total Space	Storage Efficiency	Is Volume Moving
automated_capacity_t...	aggr5ATA	Online	No	95	45.43 MB	1 GB	Disabled	No
svm1_cluster1_root	aggr5ATA	Online	No	5	18.86 MB	20 MB	Disabled	No

Figure 26) Testing use case 1: OnCommand Unified Manager Events view.

Events						
Triggered Time	Status	State	Impact Level	Impact Area	Name	Source
05:03 AM, 14 Feb	New	New	Risk	Capacity	Volume Space Full	svm1_cluster1/automated_capacity_test

The event triggered an alert, in turn executing a Workflow Automation workflow to remediate the issue (Figure 27).

Figure 27) Testing use case 1: WFA Execution Status view.

AdministrationHelp

Portal

Execution

Designer

Storage Automation Store

10 Schemes selectedadminSign Out

NetApp

Workflow

Execution Status

Recurring Schedules

Reservations

Schedules

Execution Status

Job

Name

Start Time

End Time

Status

Completed

Submitted By

Submitted At

Execution Comment

Sched

10064

Resize_Volume_with_Data_Mobility

02/14/17 5:03:25 AM

02/14/17 5:03:32 AM

Successful

4/4

admin

02/14/17 5:03:24 AM

OCUM triggered workflow. Event ID: 50020

The volume was increased in size. Utilization went down from 95% to 31%, and in this case, no movement to a different aggregate was necessary (Figure 28).

Figure 28) Testing use case 1: remediated volume capacity status.

Volumes								
Name	Aggregate	Status	Thin Provisioned	% Used	Available Space	Total Space	Storage Efficiency	Is Volume Moving
automated_capacity_t...	aggr5ATA	Online	No	31	2.04 GB	3 GB	Disabled	No
svm1_cluster1_root	aggr5ATA	Online	No	5	18.86 MB	20 MB	Disabled	No

3.2 Use Case 2: Automatic Inode Management in High-File-Count (HFC) Environments

In this use case, the system reacts to a FlexVol volume running out of available inodes and increases the number of inodes to maintain X% inode utilization.

Creating a Remediation Orchestration Workflow

Table 3 lists the requirements and implementation methodology for this workflow.

Table 3) Use case 2 requirements and methodology.

Workflow Requirements	Workflow Implementation Methodology
<ul style="list-style-type: none"> • Increase the number of available inodes in a FlexVol volume. • Inode utilization should be at or below X%. 	<ul style="list-style-type: none"> • Find all the relevant information about the FlexVol volume. • Calculate the new total count of inodes in the FlexVol volume by following this formula: $\text{new_total_inode_count} = \text{current_used_inodes} / X\%$ (where X is the desired maximum inode utilization and is a value between 1 and 99). • Change the total number of inodes based on the calculated value.

The flowchart in Figure 29 represents the logic of this workflow.

Figure 29) Use case 2 flowchart.


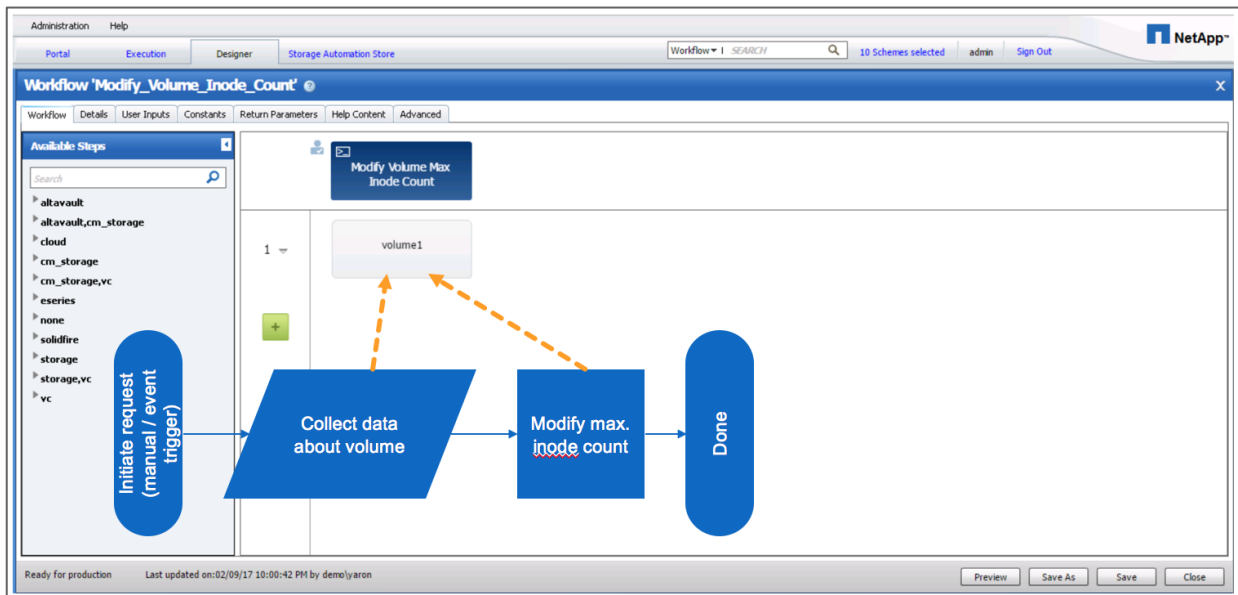


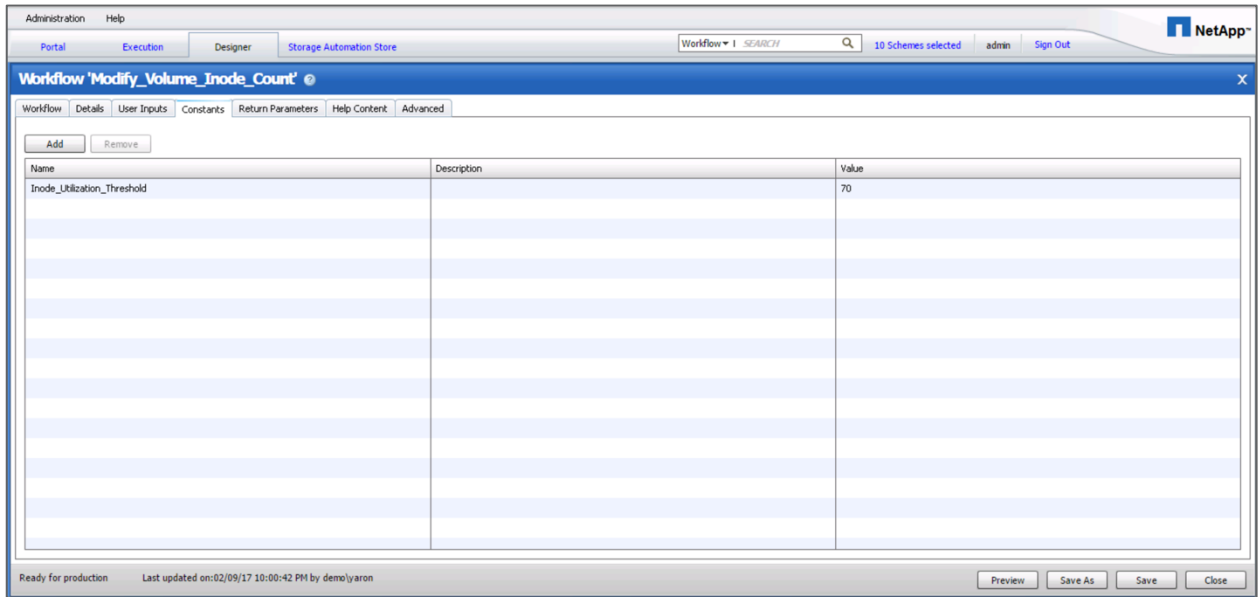
Figure 30 is a view of the completed workflow with the layered implementation flowchart above it.

Figure 30) Use case 2 WFA workflow view.



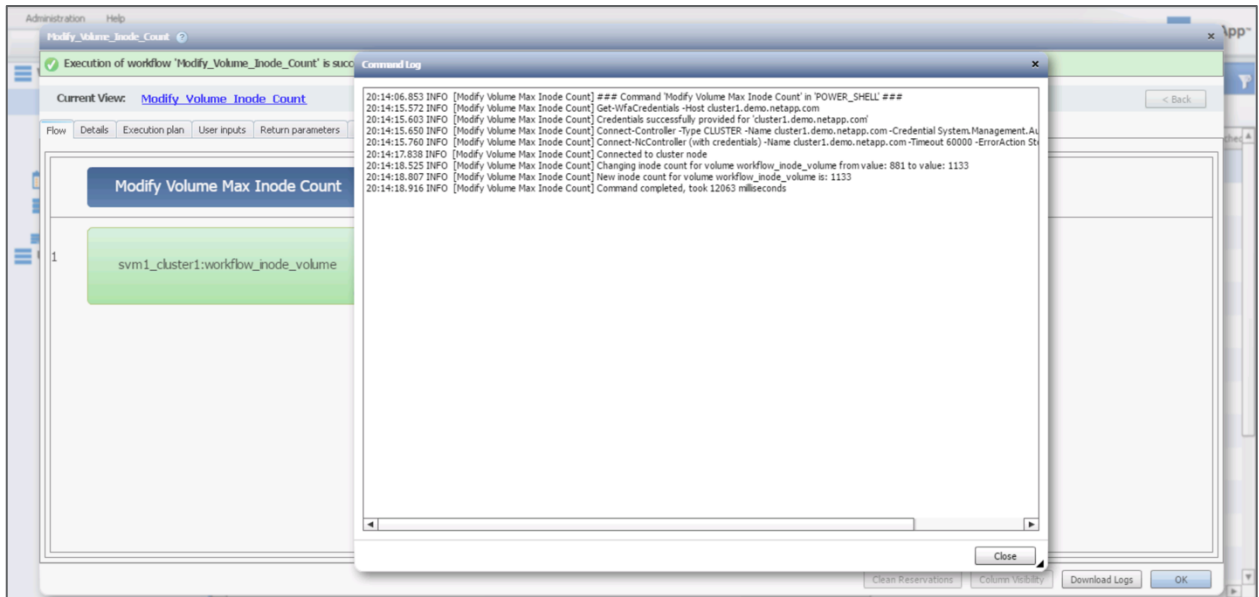
Workflow Automation enables you to define legal values for the various parameters that are leveraged within the workflow. You can also define global values that are easy to modify if requirements change. Figure 31 shows how to define a global constant for the inode maximum threshold utilization (set to 70% in the example).

Figure 31) Use case 2: customizing workflow inputs.



Execution of the workflow completes with the volume's inode count increasing from 881 to 1133 (using a default value of 70% maximum utilization). See Figure 32.

Figure 32) Use case 2 execution results.



Creating a Workflow Alert Launch Script for OnCommand Unified Manager

Before creating the relevant OnCommand alert, you should create a script that can be attached to the alert for automated execution.

Following is a PowerShell version of the script (OCUM_WFA_Inode_Event.ps1):

```

# 2017-02-09 yaron@netapp.com
# OCUM_WFA_Inode_Event.ps1
#
  
```

```

# This script is meant to be used as part of an automated self-healing process
# that remediates high inode utilization scenarios. It is meant to be used in
# conjunction with OCUM 7.x Windows-based systems and OnCommand WFA 4.1.x.
#
# IMPORTANT: This code assumes a single ONTAP cluster is present in the
#             environment for code simplification. If environment contains more than one
#             cluster see Appendix A.
#
# (c) 2017 NetApp Inc., All Rights Reserved
#
# Check the following link in case you run into SSL certificate issues
# https://d-fens.ch/2013/12/20/nobrainier-ssl-connection-error-when-using-powershell/

#=====
# CHANGE IF DESIRED PRIOR TO UPLOADING TO OCUM SERVER
$cluster_name = "cluster1"
$wfa_rest_server = "https://wfa.demo.netapp.com/rest"
$wfa_username = "admin"
$wfa_password = "Netapp1!"
#=====

#=====
# DO NOT CHANGE CODE BELOW THIS LINE
#=====
if ($args[5] -eq "Full"){
    # This is an "Inodes Nearly Full" event
    $source_name = $args[11]
    $event_state = $args[15]
} else {
    # This is an "Inodes Full" event
    $source_name = $args[10]
    $event_state = $args[14]
}
$event_id = $args[1]

# Ignore all non-new events
if (($event_state.ToString()).ToLower() -ne "new"){
    exit
}

# Extract SVM and Volume names from OCUM-provided arguments
$svm_name = $source_name.Split("{:/}") [0]
$volume_name = $source_name.Split("{:/}") [2]

# Prepare WFA credentials
$securePassword = ConvertTo-SecureString $wfa_password -AsPlainText -Force
$cred = New-Object System.Management.Automation.PSCredential -ArgumentList
$wfa_username,$securePassword

# Find RESTful execution url for WFA workflow
$restCommand = $wfa_rest_server + "/workflows?name=Modify_Volume_Inode_Count"
try {
    $output = Invoke-RestMethod -Method Get -Uri $restCommand -Credential $cred
} catch {
    # An error occurred. Either workflow does not exist or credentials are wrong
    exit
}
$workflow_execution_uri = ($output.collection.workflow.link | ?{$_rel -eq
    'execute'}).href

# Initiate workflow execution
$body = @"
{
    "comments":"OCUM triggered workflow. Event ID: $event_id",
    "userInputValues":[
        {
            "key":"cluster_name",
            "value":"$cluster_name"
        },
    ],

```

```

        {
            "key": "svm_name",
            "value": "$svm_name"
        },
        {
            "key": "volume_name",
            "value": "$volume_name"
        }
    ]
}
"@

$output = Invoke-RestMethod -Method Post -Uri $workflow_execution_uri -Credential $cred -Body
$body -ContentType "application/json"

# Find RESTful job monitoring link for WFA workflow
$job_status_uri = ($output.job.link | ?{$_rel -eq 'self'}).href
if ($job_status_uri.Count -gt 1){$job_status_uri = $job_status_uri[0]}

# Wait until workflow job either completes successfully or fails
do {
    Start-Sleep -Seconds 2
    $output = Invoke-RestMethod -Method Get -Uri $job_status_uri -Credential $cred
    $jobStatus = $output.job.jobStatus.jobStatus
} while (($jobStatus -ne "COMPLETED") -and ($jobStatus -ne "FAILED"))


```

Creating an OnCommand Unified Manager Alert

To create an OnCommand Unified Manager alert:

1. Start with uploading the workflow launch script to the OnCommand Unified Manager server. To upload the script to the OnCommand Unified Manager server, follow the same steps that are described in section 3.1 (in the subsection “Creating an OnCommand Unified Manager Alert”). See Figure 33.

Figure 33) OnCommand Unified Manager Manage Scripts view.




2. Follow the same steps that are described in section 3.1 and create an alert by using the parameters that are detailed in Table 4.

Table 4) Parameters for alert creation for use case 2.

Parameter	Value
Name	High inode utilization alert
Description	Alert used in a self-healing automated workflow
Resources	<< All Volumes >>
Events	<ul style="list-style-type: none"> • Inodes Nearly Full (Warning) • Inodes Full (Error)
Action	Execute script OCUM_WFA_Inode_Event

3. When the alert has been created, the Manage Alerts view should look like Figure 34.

Figure 34) OnCommand Unified Manager Manage Alerts view.




That's it! You have created an automated process that deals with volumes running out of available inodes.

Testing the Self-Healing Scenario

To test use case 2:

1. Create a 30MB volume. Creating a small volume yields a fairly low number of inodes by default.

Figure 35) Testing use case 2: Create a volume.



- As Figure 35 shows, there are currently 881 total inodes, with 97 of them being used.
2. Create 653 small files to increase the total number of used inodes to 750 (85% inode utilization). See Figure 36.

Figure 36) Testing use case 2: volume status after creation of files.

Volumes

Create

Edit

Delete

Clone

Status

Snapshot Copies

Resize

Storage Efficiency

Move

Storage QoS

Protect

Refresh

Name	Aggregate	Status	Thin Provisioned	% Used	Available Space	Total Space	Storage Efficiency	Is Volume Moving
automated_capacity_t...	aggrSATA	Online	No	31	2.04 GB	3 GB	Disabled	No
automated_inode_test	aggrSATA	Online	No	1	29.7 MB	30 MB	Disabled	No
svm1_cluster1_root	aggrSATA	Online	No	5	18.86 MB	20 MB	Disabled	No

General

Name:

automated_inode_test

Status:

Online

Maximum Files:

881

Current Files:

750

Language:

C LITE R/PDSIX with LIT

Autogrow Mode:

Disabled

Autogrow Maximum Size:

Disabled

Autogrow Increment Size:

Disabled

Snapshot Autodelete:

Disabled

Snapshot Autodelete Commitment:

NA

3. The high inode utilization threshold causes an event to be generated (Figure 37), in turn triggering an alert that notifies Workflow Automation to execute a remediating workflow (Figure 38).

Figure 37) Testing use case 2: OnCommand Unified Manager Events view.

Events							
Assign To	✓ Acknowledge	✗ Mark As Resolved	✗ Add Alert				
Triggered Time	Status	State	Impact Level	Impact Area	Name	Source	
05:40 AM, 14 Feb	⚠	New	Risk	Capacity	Inodes Nearly Full	svm1_cluster1:/automated_inode_test	

Figure 38) Testing use case 2: WFA Execution Status view.

Execution Status									
Job #	Name	Start Time	End Time	Status	Completed	Submitted By	Submitted At	Execution Comment	Sched
10072	Modify_Volume_Inode_Count	02/14/17 5:40:04 AM	02/14/17 5:40:11 AM	Successful	1/1	admin	02/14/17 5:40:02 AM	OCUM triggered workflow, Event ID: 50024	

As Figure 39 shows, the total number of inodes was increased to 1,070, thus lowering the inode utilization to 70%.

Figure 39) Testing use case 2: volume status after remediation.

Volumes									
Name	Aggregate	Status	Thin Provisioned	% Used	Available Space	Total Space	Storage Efficiency	Is Volume Moving	
automated_capacity_t...	aggrSATA	Online	No	31	2.04 GB	3 GB	Disabled	No	
automated_inode_test	aggrSATA	Online	No	1	29.7 MB	30 MB	Disabled	No	
svm1_cluster1_root	aggrSATA	Online	No	5	18.86 MB	20 MB	Disabled	No	

General			
Name:	automated_inode_test	Autogrow Mode:	Disabled
Status:	Online	Autogrow Maximum Size:	Disabled
Maximum Files:	1.07k	Autogrow Increment Size:	Disabled
Current Files:	750	Snapshot Autodelete:	Disabled

3.3 Use Case 3: Automatic Performance Management (Tiering)

In this use case, the system leverages the I/O density (I/O per terabyte stored) metric to identify the proper performance tier for each FlexVol volume. If necessary, the volumes are then nondisruptively moved to a different storage pool (aggregate) or compute node type.

This use case is somewhat different from the previous two in that it is not intended to be triggered dynamically through an event. Rather, it is designed to be scheduled and run once a day or once a week to revalidate and reassess the volume storage tier assignment.

An I/O density approach describes the performance of a storage service, regardless of the media or the protocol that is used to deliver that service. The key measurements for a storage service include:

- **Latency.** This is the time between when a storage controller receives a request for data and when it is able to answer that request.
- **IOPS/TB.** This term relates to the number of I/Os performed per second per terabyte (TB) of data stored.

Note: IOPS/TB is a key measure for delivering consistent storage service to the end consumer.

These service levels are typically described for the consumer in a service catalog. The service catalog defines the storage service in terms of several parameters: IOPS/TB; latency; frequency of backups; frequency of off-site replication; retention of backups; and, in most circumstances, the cost of the service for a given quantity and interval (for example, TB/month).

For more information about service-oriented delivery of storage, you can schedule a Service Design Workshop, a one-day engagement that NetApp offers to customers who are interested in optimizing the service delivery process.

Automation of performance tiers requires constant monitoring of the I/O density metric. One way of tracking this information is by leveraging either OnCommand Performance Manager or OnCommand Insight. (The two products differ in the frequency of data collection and in the granularity of historical data.) Normally we look at the 95th percentile of I/O density over a period of seven days per workload (volume). Both monitoring applications can synchronize that data with OnCommand Workflow Automation through relevant data sources.

Access to this data can help you determine the proper tiering design, as shown in Table 5.

Table 5) Services and IOPS/TB settings example.

	Silver—Tier 2	Gold—Tier 1	Platinum—Tier 0
SLO (burst QoS throttle)	512 IOPS/TB	2,048 IOPS/TB	12,288 IOPS/TB
SLA (to end customer)	128 IOPS/TB	512 IOPS/TB	6,144 IOPS/TB
Media type	SATA	SAS	SSD

The flowchart in Figure 40 represents the logic of this workflow.

Figure 40) Use case 3 flowchart.


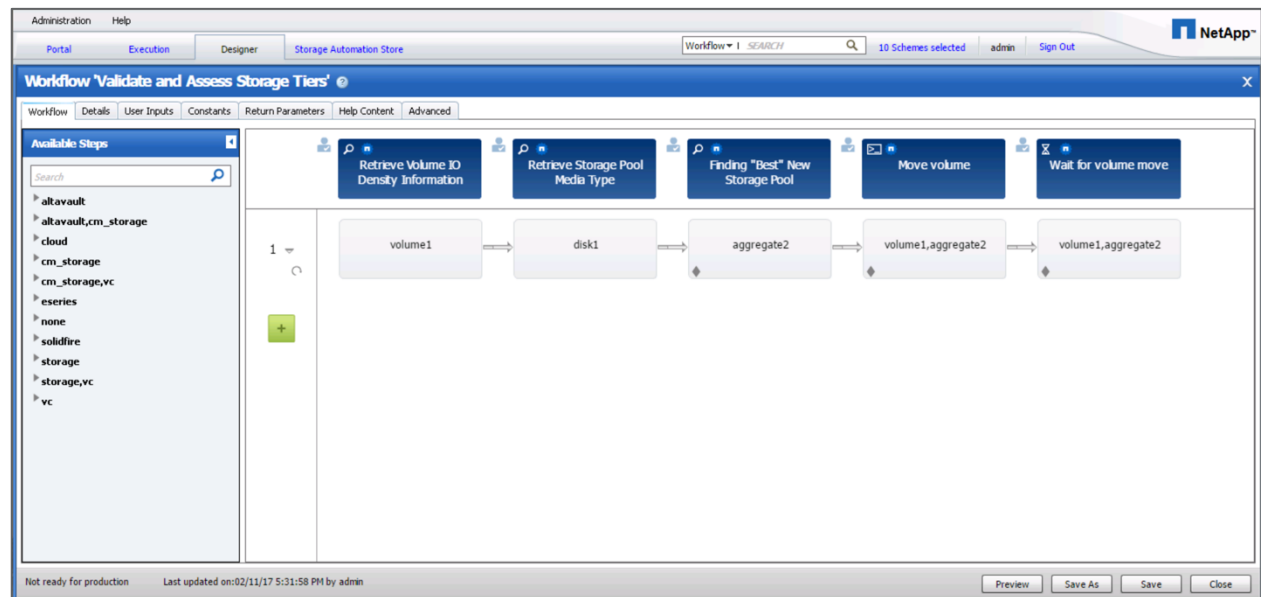


Figure 41) Use case 3 WFA workflow view.



The workflow view in Figure 41 is another indication of the power and the simplicity of OnCommand Workflow Automation. Although there is only a single row in the workflow, this row is executed against every nonroot volume in the environment through a feature called repeat rows. (The feature is represented by the curved arrow symbol next to the row indicator.) This feature uses a filter (discussed in section 2.3) to identify all the relevant resources that the workflow should run against.

The workflow can then be scheduled to run daily by using the built-in scheduler of WFA (Figure 42).

Figure 42) Use case 3 WFA scheduler.

AdministrationHelp

PortalExecutionDesignerStorage Automation Store

AllSEARCH10 Schemes selectedadminSign Out

Workflow

Execution Status

Recurring Schedules

Reservations

Schedules

Data Sources

Credentials

User Management

Users

Active Directory Groups

Schedules

Name	Description	Association Co
Tier Re-Assessment and Validation Schedule	Tier Re-Assessment and Validation Schedule. Scheduled to run nightly at 1AM.	0

Appendix A: Integrating OnCommand Manager Alerts with WFA When Monitoring Multiple Clusters

From Table 1, when a volume-based alert in NetApp OnCommand Unified Manager (OCUM) is triggered, OCUM passes the identifying details of the relevant volume to the alert script. OCUM passes this information in the form of `svm_name:/volume_name` (the value of the `-eventSourceName` argument). What is not being passed to the alert script is the containing cluster. If a single cluster is being monitored, that information is not a concern. However, if more than one cluster is actively being monitored, you need a method to identify the affected cluster.

This appendix explores one method to determine the affected cluster by using SQL queries against the OCUM database. These queries are based on the OCUM volume ID that is included in the arguments that are passed to the script (`-eventSourceID`). The goal is to find the OCUM object that represents the relevant volume, and from that object extract the volume, SVM, and cluster names that are needed to pass along to the WFA workflow.

This appendix discusses the alert script for use case 1 (automatic capacity management) and shows both PowerShell and Perl examples. If you use PowerShell, be sure to download and install Connector/Net (a fully managed ADO.NET driver for MySQL) on the WFA server from <https://dev.mysql.com/downloads/connector/net/>. If you use Perl, install the DBD package (DBI package for managing MySQL).

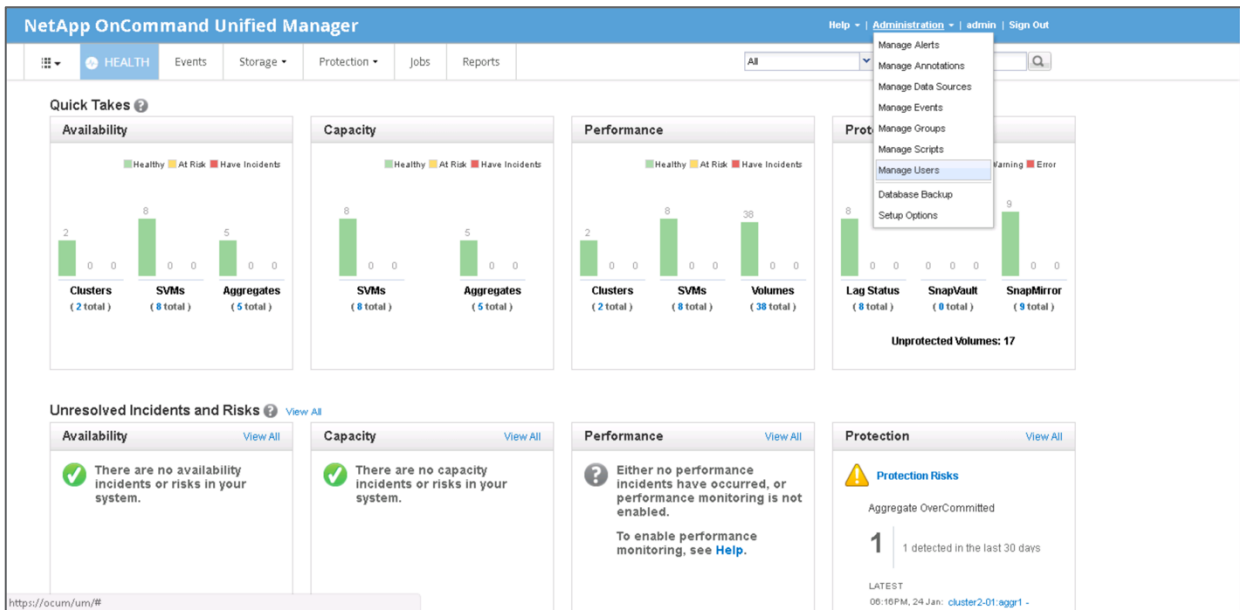
Note

The following blog post from TheTechArch is a great resource for querying the OCUM database by using PowerShell: <http://thetecharch.com/2016/03/querying-ocum-database-using-powershell-2/>.

Before you proceed with the alert script, create a database user in OCUM that can be used to query the `ocum_report` database:

1. Log in to OCUM and select Manage Users from the Administration drop-down menu (Figure 43).

Figure 43) OnCommand Unified Manager Manage Users drop-down menu.



2. In the Manage Users view, click the Add button and input the following information (Figure 44 and Table 6):

Figure 44) OnCommand Unified Manager: creating new database user form.

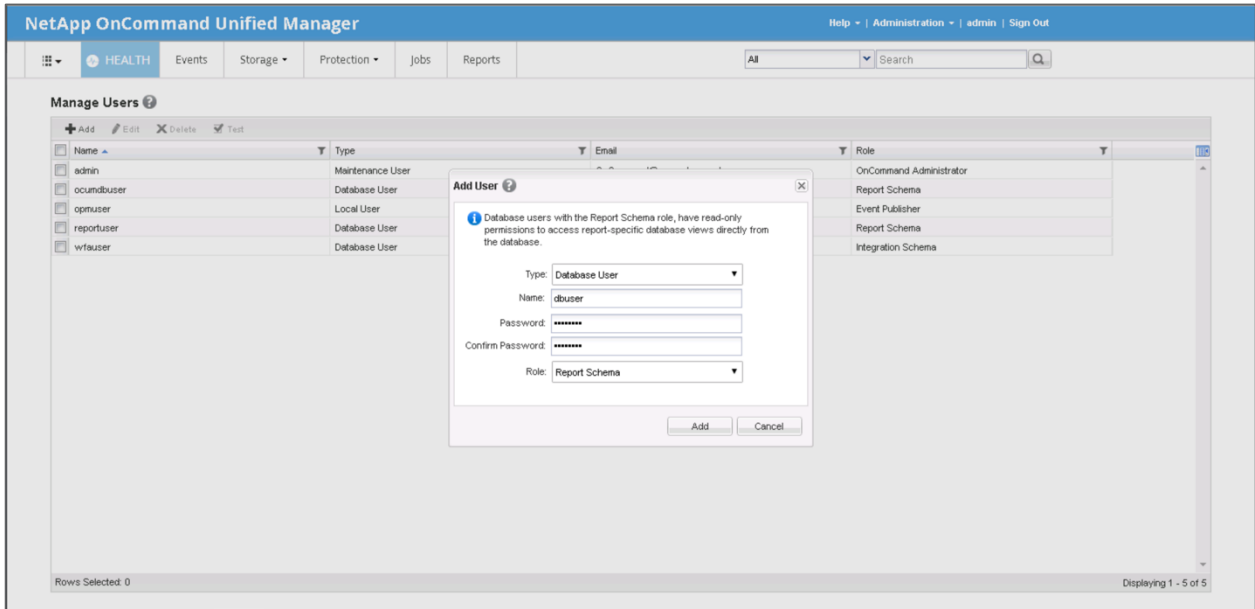


Table 6) Parameters for OCUM database user creation.

Parameter	Value
Type	Database User
Name	Name of user (dbuser in our example)
Password	User password (Netapp1! in our example)
Role	Report Schema

3. After the user has been created, proceed with the updated alert scripts.

Following is a PowerShell script example (OCUM_WFA_Capacity_Event_with_SQL.ps1):

```
# 2017-02-09 yaron@netapp.com
# OCUM_WFA_Capacity_Event_with_SQL.ps1
#
# This script is meant to be used as part of an automated self-healing process
# that remediates high volume capacity utilization scenarios. It is meant to be used
# in conjunction with OCUM 7.x Windows-based systems and OnCommand WFA 4.1.x.
#
# (c) 2017 NetApp Inc., All Rights Reserved
#
# Check the following link in case you run into SSL certificate issues
# https://d-fens.ch/2013/12/20/nobrainier-ssl-connection-error-when-using-powershell/

#=====
# CHANGE IF DESIRED PRIOR TO UPLOADING TO OCUM SERVER
$wfa_rest_server = "https://wfa.demo.netapp.com/rest"
$wfa_username = "admin"
$wfa_password = "Netapp1!"
$ocum_username = "dbuser"
$ocum_password = "Netapp1!"
$ocum_server = "192.168.0.74"
#=====
```

```

#=====
# DO NOT CHANGE CODE BELOW THIS LINE
#=====
# Function MySQL queries OCUM database
# usage: MySQL -Query <sql-query>
function MySQL {
    Param(
        [Parameter(
            Mandatory = $true,
            ParameterSetName = '',
            ValueFromPipeline = $true)]
        [string]$Query
    )

    $MySQLAdminUserName = $ocum_username
    $MySQLAdminPassword = $ocum_password
    $MySQLDatabase = 'ocum_report'
    $MySQLHost = $ocum_server
    $ConnectionString = "server=" + $MySQLHost + ";port=3306;Integrated Security=False;uid=" +
    $MySQLAdminUserName + ";pwd=" + $MySQLAdminPassword + ";database=" + $MySQLDatabase

    Try {
        [void][System.Reflection.Assembly]::LoadFrom('C:\Program Files (x86)\MySQL\MySQL
Connector Net 6.9.9\Assemblies\v4.5\MySql.Data.dll')
        $Connection = New-Object MySql.Data.MySqlClient.MySqlConnection
        $Connection.ConnectionString = $ConnectionString
        $Connection.Open()

        $Command = New-Object MySql.Data.MySqlClient.MySqlCommand($Query, $Connection)
        $DataAdapter = New-Object MySql.Data.MySqlClient.MySqlDataAdapter($Command)
        $DataSet = New-Object System.Data.DataSet
        $RecordCount = $dataAdapter.Fill($dataSet, "data")
        $DataSet.Tables[0]
    }

    Catch {
        Write-Host "ERROR : Unable to run query : $query `n$error[0]"
    }

    Finally {
        $Connection.Close()
    }
}

if ($args[6] -eq "Full"){
    # This is an "Inodes Nearly Full" event
    $source_id = $args[10]
    $event_state = $args[16]
} else {
    # This is an "Inodes Full" event
    $source_id = $args[9]
    $event_state = $args[15]
}
$event_id = $args[1]

# Ignore all non-new events
if (($event_state.ToString()).ToLower() -ne "new"){
    exit
}

# Extract Cluster, SVM and Volume names from OCUM database
$sql_query = @"
SELECT
    volume.name AS 'Volume',
    cluster.name AS 'Cluster',
    svm.name AS 'Svm'
FROM
    volume,
    cluster,
    svm

```

```

WHERE
    volume.id=$source_id
    AND cluster.id=volume.clusterId
    AND svm.id=volume.svmId
"@
$ocum_vol = MySQL -Query $sql_query
$cluster_name = $ocum_vol.Cluster
$svm_name = $ocum_vol.Svm
$volume_name = $ocum_vol.Volume

# Prepare WFA credentials
$securePassword = ConvertTo-SecureString $wfa_password -AsPlainText -Force
$cred = New-Object System.Management.Automation.PSCredential -ArgumentList
$wfa_username,$securePassword

# Find RESTful execution url for WFA workflow
$restCommand = $wfa_rest_server + "/workflows?name=Resize_Volume_with_Data_Mobility"
try {
    $output = Invoke-RestMethod -Method Get -Uri $restCommand -Credential $cred
} catch {
    # An error occurred. Either workflow does not exist or credentials are wrong
    exit
}
$workflow_execution_uri = ($output.collection.workflow.link | ?{$_rel -eq
'execute'}).href

# Initiate workflow execution
$body = @"
{
    "comments":"OCUM triggered workflow. Event ID: $event_id",
    "userInputValues":[
        {
            "key":"cluster_name",
            "value":"$cluster_name"
        },
        {
            "key":"svm_name",
            "value":"$svm_name"
        },
        {
            "key":"volume_name",
            "value":"$volume_name"
        }
    ]
}
"@

$output = Invoke-RestMethod -Method Post -Uri $workflow_execution_uri -Credential $cred -Body
$body -ContentType "application/json"

# Find RESTful job monitoring link for WFA workflow
$job_status_uri = ($output.job.link | ?{$_rel -eq 'self'}).href
if ($job_status_uri.Count -gt 1){$job_status_uri = $job_status_uri[0]}

# Wait until workflow job either completes successfully or fails
do {
    Start-Sleep -Seconds 2
    $output = Invoke-RestMethod -Method Get -Uri $job_status_uri -Credential $cred
    $jobStatus = $output.job.jobStatus.jobStatus
} while (($jobStatus -ne "COMPLETED") -and ($jobStatus -ne "FAILED"))

```

And following is the comparable Perl script example (OCUM_WFA_Capacity_Event_with_SQL.pl):

```

# 2017-02-09 yaron@netapp.com
# OCUM_WFA_Capacity_Event_with_SQL.pl
#
# This script is meant to be used as part of an automated self-healing process
# that remediates high volume capacity utilization scenarios. It is meant to be used in
# conjunction with OCUM 7.x Windows-based systems and OnCommand WFA 4.1.x.
#

```

```

# (c) 2017 NetApp Inc., All Rights Reserved
#

use REST::Client;
use JSON;
use MIME::Base64;
use DBI;

# Unmark the next line if you run into SSL certificate issues
$ENV{PERL_LWP_SSL_VERIFY_HOSTNAME}=0;

#####
# CHANGE IF DESIRED PRIOR TO UPLOADING TO OCUM SERVER
my $wfa_rest_server = 'https://wfa.demo.netapp.com';
my $wfa_username = 'admin';
my $wfa_password = 'Netapp1!';
my $ocum_username = 'dbadmin';
my $ocum_password = 'Netapp1!';
my $ocum_server = 'ocum.demo.netapp.com';
#####

#####
# DO NOT CHANGE CODE BELOW THIS LINE
#####
if ($ARGV[6] == 'Full'){
    # This is an "Volume Space Nearly Full" event
    my $source_id = $ARGV[10];
    my $event_state = $ARGV[16];
} else {
    # This is an "Volume Space Full" event
    my $source_id = $ARGV[9];
    my $event_state = $ARGV[15];
}
my $event_id = $ARGV[1];

# Ignore all non-new events
if ($event_state == 'NEW'){
    exit;
}

# Extract Cluster, SVM and Volume names from OCUM database
my $sql_query = "SELECT
                volume.name AS 'Volume',
                cluster.name AS 'Cluster',
                svm.name AS 'Svm'
FROM
    volume,
    cluster,
    svm
WHERE
    volume.id=$source_id
    AND cluster.id=volume.clusterId
    AND svm.id=volume.svmId";
my $dsn = "DBI:mysql:ocum report:" . $ocum_server;
my %attr = ( PrintError=>0, # turn off error reporting via warn()
            RaiseError=>1 # report error via die()
            );
my $dbh = DBI->connect($dsn,$ocum_username,$ocum_password,\%attr);
my $sth = $dbh->prepare($sql_query);
$sth->execute();
while(my @row = $sth->fetchrow_array()){
    my $volume_name = $row[0];
    my $cluster_name = $row[1];
    my $svm_name = $row[2];
}
$sth->finish();
$dbh->disconnect();

# Find RESTful execution url for WFA workflow

```

```

my $headers = {Accept => 'application/json', Authorization => 'Basic ' .
    encode_base64($wfa_username . ':' . $wfa_password), 'Content-Type' => 'application/json'};
my $client = REST::Client->new();
$client->setHost($wfa_rest_server);
$client->GET('/rest/workflows?name=Resize_Volume_with_Data_Mobility', $headers);
my @response_json = @{ decode_json($client->responseContent()) };
my $workflow_execution_uri = '/rest/workflows/' . $response_json[0]->{'uuid'} . '/jobs';

# Initiate workflow execution
my $json_body = '{
    "comments":"OCUM triggered workflow. Event ID: ' . $event_id . '",
    "userInputValues":[
        {
            "key":"ClusterName",
            "value":"' . $cluster_name . '"
        },
        {
            "key":"SvmName",
            "value":"' . $svm_name . '"
        },
        {
            "key":"VolumeName",
            "value":"' . $volume_name . '"
        }
    ]
}';

# Find RESTful job monitoring link for WFA workflow
$client->POST($workflow_execution_uri, $json_body, $headers);
my $response = decode_json($client->responseContent());
$job_status_uri = $workflow_execution_uri . '/' . $response->{'jobId'};

# Wait until workflow job either completes successfully or fails
do {
    sleep 2;
    $client->GET($job_status_uri, $headers);
    my $response = decode_json($client->responseContent());
    $jobStatus = $response->{'jobStatus'}{'jobStatus'};
} while (($jobStatus != 'COMPLETED') && ($jobStatus != 'FAILED'));

```

References

The following references are relevant to this TR:

- OnCommand Workflow Automation: Workflow Developer's Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2597424
- OnCommand Workflow Automation: REST API Primer
https://library.netapp.com/ecm/ecm_download_file/ECMLP2597425
- OnCommand Storage Management Community
<http://community.netapp.com/t5/FAS-ONTAP-and-OnCommand/ct-p/fas-ontap-oncommand-and-related-plug-ins>
- TR-4438: IT as a Service—Simplifying Application and Storage Provisioning Using NetApp OnCommand Workflow Automation and System Center Orchestrator 2012 R2
www.netapp.com/us/media/tr-4438.pdf
- TR-4103: IT as a Service—Storage Automation Using NetApp OnCommand Workflow Automation and HP Operations Orchestration
www.netapp.com/us/media/tr-4103.pdf
- TR-4217: Automating and Orchestrating the Software—Defined Data Center: Using NetApp and VMware to Build Your Cloud
www.netapp.com/us/media/tr-4217.pdf
- TR-4217: ONTAP Storage Service Deployment Guide—Deploying Systems Based on a Service Design Workshop
www.netapp.com/us/media/tr-4548-deploy.pdf
- TR-4572: The NetApp Solution for Ransomware
www.netapp.com/us/media/tr-4572.pdf
- OnCommand Management Software and Management Integration Tools
www.netapp.com/us/products/management-software/index.aspx
- Make Storage & Data Management Both Easier to Use & Scale with NetApp Service Level Manager
<http://community.netapp.com/t5/Technology/Make-Storage-amp-Data-Management-Both-Easier-to-Use-amp-Scale-with-NetApp/ba-p/129061>

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

Copyright Information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

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