



Technical Report

# Software-Defined Storage with NetApp and VMware

Andrew Sullivan and Shrivatsa Upadhye, NetApp  
Bhumik Patel and Jerry Mayfield, VMware

January 2016 | TR-4308

## TABLE OF CONTENTS

<b>1</b>	<b>Joint Solution Overview .....</b>	<b>5</b>
1.1	Summary.....	5
1.2	SDDC Overview .....	5
1.3	SDS Overview.....	6
1.4	NetApp SDS in SDDC.....	7
1.5	Automation and Orchestration in SDDC .....	9
<b>2</b>	<b>NetApp and VMware Solution Components .....</b>	<b>10</b>
2.1	SDS with OnCommand WFA 3.1 .....	11
2.2	OnCommand Unified Manager for NetApp Storage Management .....	12
2.3	VMware vRealize Orchestrator .....	13
2.4	VMware vRealize Automation 6.2 in SDDC with SDS .....	15
<b>3</b>	<b>Optional Component: Performance Monitoring and Problem Remediation Management .....</b>	<b>22</b>
3.1	OnCommand Insight .....	22
3.2	vRealize Operations.....	26
3.3	Blue Medora vRealize Operations Management Pack for NetApp.....	28
3.4	vRealize Log Insight.....	28
3.5	NetApp vRealize Log Insight Content Pack .....	30
<b>4</b>	<b>Optional Component: Disaster Recovery in a NetApp and VMware SDDC Environment.....</b>	<b>30</b>
4.1	NetApp SnapMirror .....	31
4.2	VMware vCenter SRM .....	32
4.3	NetApp SRA and SVMs .....	32
4.4	Disaster Recovery in a NetApp and VMware SDDC Environment.....	32
<b>5</b>	<b>Terminologies and Functionality .....</b>	<b>33</b>
5.1	NetApp and VMware SDS Overview.....	35
5.2	Roles and Functionality.....	36
<b>6</b>	<b>NetApp SDS.....</b>	<b>36</b>
6.1	NetApp SDS Overview.....	36
6.2	SVMs .....	37
6.3	NetApp Storage Technologies for Cloud Storage .....	37
6.4	Accelerated Performance.....	39
<b>7</b>	<b>SDS Orchestration and Automation Workflows.....</b>	<b>40</b>
7.1	Where to Begin with Your First Workflow.....	40
7.2	Benefits .....	40

7.3 Architecture Capabilities .....	41
7.4 Understanding Workflow Concepts .....	41
7.5 Security and Workflows .....	43
<b>8 Conclusion .....</b>	<b>44</b>
<b>References.....</b>	<b>44</b>
Supporting Documents .....	44
<b>Version History .....</b>	<b>45</b>

## LIST OF TABLES

Table 1) Terminology and technology used in this document.....	33
Table 2) Workflow components. ....	43

## LIST OF FIGURES

Figure 1) SDS overview.....	6
Figure 2) SDDC layers. ....	8
Figure 3) NetApp technologies for SDS.....	9
Figure 4) SDS environment. ....	11
Figure 5) OnCommand WFA 3.1 overview.....	12
Figure 6) NetApp WFA and vRealize Orchestrator in SDS.....	15
Figure 7) Example service catalog (graphic provided by VMware).....	17
Figure 8) vRealize Automation uses update profiles to maintain consistency (graphic provided by VMware). ....	18
Figure 9) Model new applications with existing services using external service (graphic provided by VMware). ....	19
Figure 10) Services whose delivery can be automated through vRealize Automation (graphic provided by VMware). ....	20
Figure 11) Advanced service designer workflow (graphic provided by VMware).....	21
Figure 12) vRealize Automation example configuration. ....	21
Figure 13) NetApp and VMware SDDC solution with performance monitoring and management.....	22
Figure 14) NetApp OnCommand Insight visualization of complex storage environment. ....	23
Figure 15) Bully workloads. ....	24
Figure 16) Victim workloads. ....	24
Figure 17) Capacity and performance data for NetApp devices. ....	25
Figure 18) Dashboard showing capacity trending, forecasting, and performance. ....	26
Figure 19) vRealize Operations in the cloud (graphic provided by VMware). ....	27
Figure 20) Blue Medora vRealize Operations plug-in for NetApp storage.....	28
Figure 21) vRealize Log Insight for real-time log management (graphic provided by VMware).....	29
Figure 22) NetApp and VMware SDDC with DR. ....	31
Figure 23) NetApp and VMware SDS overview.....	35
Figure 24) NetApp QoS with different SLAs for different tenants. ....	39
Figure 25) Series of workflows. ....	42

Figure 26) Storage service catalog.....	43
---	----

# 1 Joint Solution Overview

## 1.1 Summary

Companies need to deliver applications and IT services to their customers more quickly and efficiently. Competitive pressures and the accelerating pace of business no longer accommodate time-consuming, siloed, manual processes. IT customers have come to expect the same levels of service in their business lives that they get in their personal lives. They expect to go online, request a new application or compute resource, and receive it in a matter of minutes, not days or weeks. If IT cannot deliver this level of service, businesses look for alternatives that can.

The majority of IT organizations still struggle to shift away from a reactive mode of operations. Organizations operating reactively typically have inflexible IT systems and outdated processes that prevent IT from satisfying users' needs for rapid delivery of new services and applications: services and applications that are necessary to grow the business and be more competitive. Reactive IT organizations face a fundamental challenge: Because approximately 70% of their resources are devoted to operating IT, managing infrastructure, and supporting applications, few resources remain for building innovative or enhanced applications that are capable of generating or seizing new business opportunities.

For IT to advance from a reactive operating model, organizations need to redefine the way that they operate and rethink the technology they leverage for business advantage. Business investments at this stage bring greater order to preexisting cloud computing investments and make sure that the need for rapid innovation does not create silos in the future. To evolve beyond this stage, businesses should do the following:

- Continue to invest in private cloud resources to increase the impact cloud computing can have in driving a more efficient, reliable, and agile infrastructure.
- Establish a self-service environment to deliver IT services to stakeholders in a more efficient on-demand model.
- Begin to quantify the true costs of IT services.
- Create aggressive standardizations throughout the organization to enable business policies that can then be turned into automation workflows.

Identify core processes that can be automated by leveraging the policy-based nature of their virtualized and cloud computing environments. NetApp and VMware—leaders in storage and data management and virtualization and cloud solutions, respectively—share a common vision to accelerate the customer's path to the cloud through advanced virtualization, automation, and self-service. VMware cloud infrastructure, management, and orchestration solutions on the NetApp® Unified Storage Architecture provide customers with a clear path to transform their data centers into a private cloud environment delivering IT as a service.

This technical report describes the joint NetApp and VMware solution to implement a software-defined storage (SDS) solution, which is a critical component in achieving an entire software-defined data center (SDDC). For more information about automating and orchestrating SDS, see [TR-4306: Building Automation and Orchestration for Software-Defined Storage with NetApp and VMware](#).

## 1.2 SDDC Overview

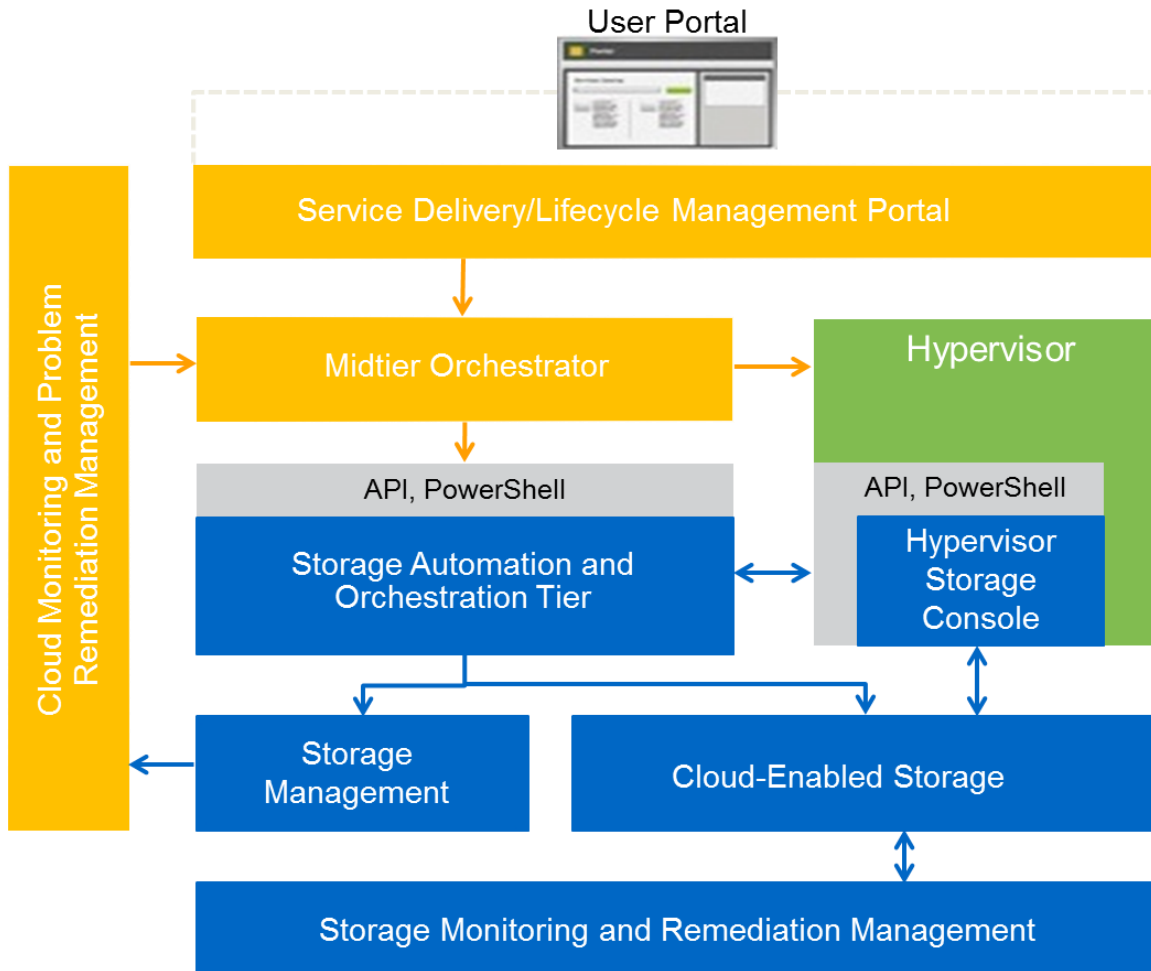
The SDDC is a unified data center platform that provides unprecedented automation, flexibility, and efficiency to transform the way administrators deliver IT. Compute, storage, networking, security, and availability services are pooled, aggregated, and delivered as software and managed by intelligent, policy-driven software. Self-service, policy-based provisioning, automated infrastructure, and application and business management complete the picture. The result is a data center that is optimized for the cloud era, providing unmatched business agility, the highest service-level agreements (SLAs) for all applications, dramatically simpler operations, and lower costs.

In the SDDC, all resources are virtualized so they can be automatically deployed with little or no human involvement. Applications can be operational in minutes, shortening time to value and dramatically reducing IT staff time spent on application provisioning and deployment.

### 1.3 SDS Overview

In order to understand the architecture behind a NetApp and VMware joint SDS environment, it is important to understand the parts of a generic solution. Figure 1 shows an overview of an SDS solution.

Figure 1) SDS overview.



The generic SDS solution consists of the following parts:

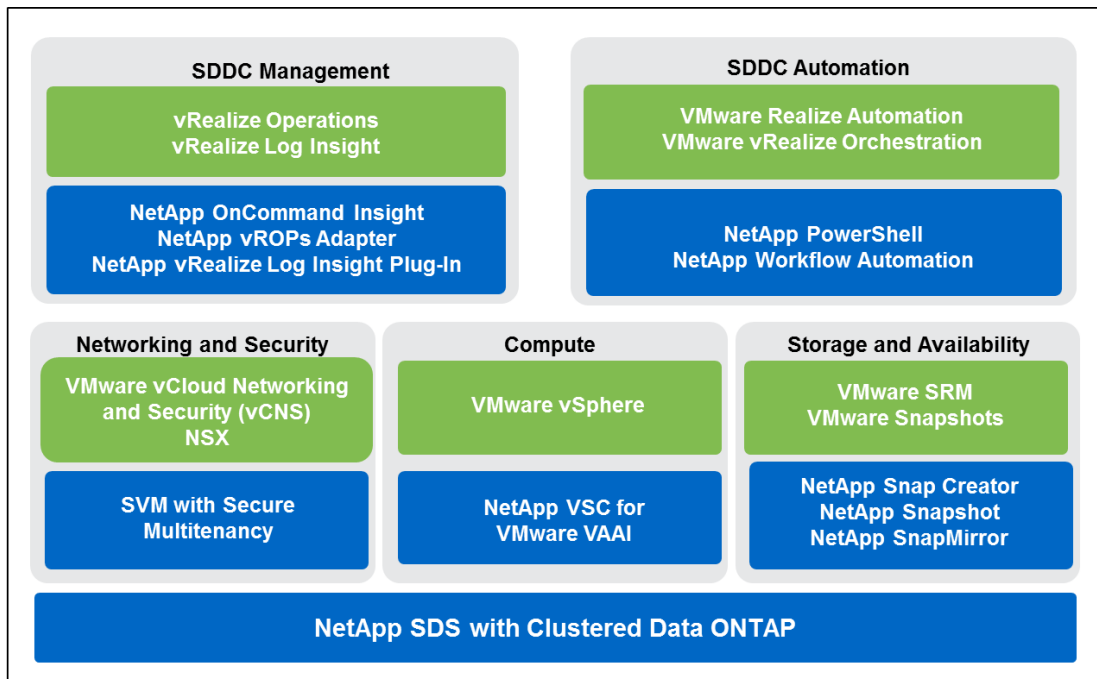
- **Cloud-enabled storage.** An enterprise storage environment that can support orchestration of its storage and has the following characteristics:
  - **Storage abstraction layer.** A true storage abstraction layer acts very similarly to a VMware virtual machine (VM), acting as its own complete storage entity. This storage VM should have the capability to move from one storage controller to another, seamlessly and without downtime, supporting continuous uptime of a cloud environment.
  - **Hypervisor integration.** Integration with the hypervisor layer supports storage and VM creation and performance and capability enhancements included with many hypervisors.

- **VM cloning support.** Cloud-enabled storage must have the capability to clone VMs at a rapid pace and then deploy those VMs to the hypervisor. At the very least, it should support the built-in VM cloning technique of the hypervisor.
- **Orchestration and automation of storage components.** Each of the features in the following list should have the capability to be orchestrated and be placed into an automation framework. Additionally, orchestration and automation of the following storage capabilities must be enabled:
  - Storage creation
  - Storage capabilities enablement
  - Storage removal
  - Backup and recovery
  - Disaster recovery (DR)
  - Storage network configuration
- **Storage management.** This layer is responsible for key management and monitoring capabilities that assist in managing the availability, capacity, protection, and performance risks of the storage systems. The information from this layer can also be used by the storage automation and orchestration tier to discover and identify the resource availability before making changes to the infrastructure.
- **Storage monitoring and remediation management.** This layer is responsible for providing storage resource management with end-to-end monitoring and brokering across heterogeneous environments. Administrators can assess the performance of applications, VMs, and storage infrastructure as well as make decisions about resource investments for the future.
- **Storage automation and orchestration tier.** The automation or orchestration framework that sits close to the storage and allows the creation, updating, deletion, and reading of the storage as needed. This framework allows storage administrators to orchestrate or automate specific functions within their storage environment and also allows higher-level orchestration and automation engines to access it in order to create storage for other processes associated with virtualization.
- **Hypervisor.** This important layer creates and runs VMs in SDDC. The following integrations help make an SDS environment easier to implement:
  - **Hypervisor storage console.** This layer assists in end-to-end VM lifecycle management along with optimization and migration of VMs.
- **API.** The capability is provided to use APIs as a methodology to automate at this layer along with the storage. This capability can be leveraged with a development language such as Microsoft PowerShell.
- **Midtier orchestration.** An orchestration layer that hypervisors use in order to create and update hypervisor objects such as VMs, networking, and VM cloning. This layer calls the SDS orchestration layer to create or alter storage in conjunction with other creation and update procedures.
  - **Cloud monitoring and problem remediation management.** A problem remediation tool that assists in detecting problems within the environment and then either resolves those problems or informs someone about the problem so that that person can resolve it.
- **Service delivery and lifecycle management portal.** The automation layer that calls other orchestrators. At this level, all processes are managed, and a user portal is made available for users to request the resources they need. Also at this level, the lifecycle of the objects being automated is managed. This layer acts as the master control level, and all other orchestration is driven from this level.

## 1.4 NetApp SDS in SDDC

NetApp is uniquely positioned to deliver on the SDS component through innovations provided by the NetApp clustered Data ONTAP<sup>®</sup> operating system. Furthermore, NetApp offers integrated or complementary technologies at every facet of the SDDC, as shown in Figure 2.

Figure 2) SDDC layers.



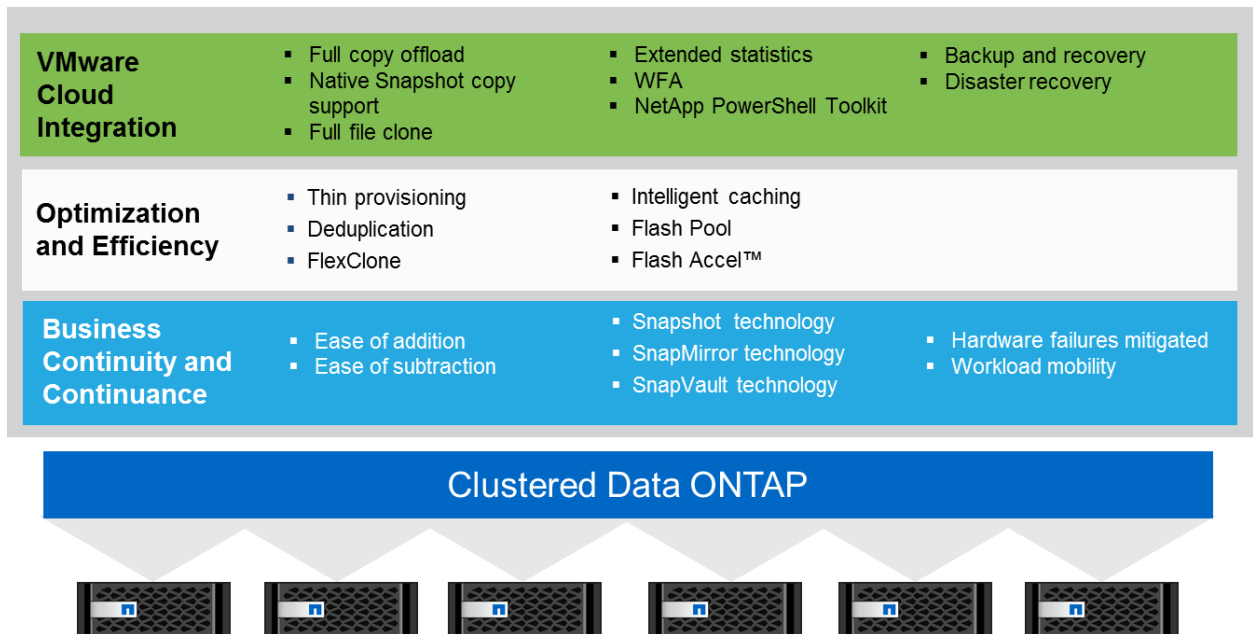
The NetApp and VMware SDDC is composed of five layers. At each layer NetApp and VMware work to bring the following integrations and complementary technologies together:

- **SDDC management.** The management layer consists of VMware vRealize Operations and VMware vRealize Log Insight for managing and troubleshooting the SDDC environment. NetApp integrates with these two technologies with an adapter by Blue Medora for vRealize Operations adapter that gives deep-dive information for the SDS. It also integrates with a plug-in for vRealize Log Insight that allows data collection for NetApp storage.
- **SDDC automation.** VMware vRealize Automation and vRealize Orchestration work in conjunction with NetApp OnCommand® Workflow Automation (WFA) and the NetApp PowerShell toolkit to create an integrated and automated SDS environment.
- **Networking and security.** At the networking and security layer, VMware and NetApp technologies work together to form a full-stack, multitenancy environment. This is done by using either VMware vCloud Network and Security (vCNS) or NSX for software-defined networking in conjunction with NetApp storage VMs to segment storage, networking, and compute into distinct and separated network stacks.
- **Compute.** NetApp has integrated with vSphere using Virtual Storage Console (VSC) and VMware API for Array Integration (VAAI).
- **Storage and availability.** At this level, NetApp provides SDS needed for the SDDC. Backup and recovery are provided at the vSphere layer with NetApp Snapshot® copies and NetApp Snap Creator®, which manages the full gamut of cloud-based backup and recovery tasks. NetApp also provides full-site failover integration with VMware vCenter Site Recovery Manager (SRM) and NetApp Storage Replication Adapter (SRA) for SRM, which works in conjunction with NetApp SnapMirror® technology.

NetApp has done for storage what VMware has done for servers by providing SDS for the enterprise with the clustered Data ONTAP operating system. The following Data ONTAP features help deliver a unified storage platform that enables unrestricted, secure data movement across multiple cloud environments and paves the way for the SDDC. Figure 3 lists the technologies and advantages that NetApp clustered Data ONTAP leverages for the SDS environment.



Figure 3) NetApp technologies for SDS.



SDS by NetApp gives customers the following advantages over other storage providers:

- **Autonomy for application owners.** Administrators can instantaneously deploy new applications and services and can respond dynamically to shifts in demand.
- **Dynamic storage provisioning based on priority and service level.** NetApp supports this capability across both SAN and NAS architectures by using virtualized storage services. The NetApp storage virtual machine (SVM, formerly known as Vserver) technology enables data access and services to be separated from the underlying hardware. This abstraction allows storage resources to be assigned and reassigned based on the needs of the user or the application over the course of its lifecycle.
- **Policy-based automation.** This allows storage to be deployed using policy-based security and delegation. Security models and delegation of storage rights are automatically configured on storage objects during storage object deployment.
- **Deployment on platform of choice.** Clustered Data ONTAP can operate on third-party storage arrays through the V-Series storage controllers. Data ONTAP Edge provides SDS functionality for commodity disks on servers that are deployed at branch locations or in cloud applications.
- **Delivery of services without compromise.** Clustered Data ONTAP excels at driving on-demand storage services and workflow automation by providing programmable APIs and OnCommand WFA, which supports integration with VMware vSphere and the VMware vCloud Suite.

## 1.5 Automation and Orchestration in SDDC

Policy-based automation and management are critical for realizing the benefits of SDDC built on top of a virtualized infrastructure. The VMware Cloud Automation solution provides the following capabilities needed to provide enterprises with the ability to automate delivery of IT services:

- **Provisioning and management of application services.** Accelerates application deployment by streamlining the deployment process and eliminating duplication of work by using reusable components and blueprints.
- **Infrastructure delivery and lifecycle management.** Automates the end-to-end deployment of multivendor infrastructures, breaking down internal organizational silos that slow down IT service delivery.

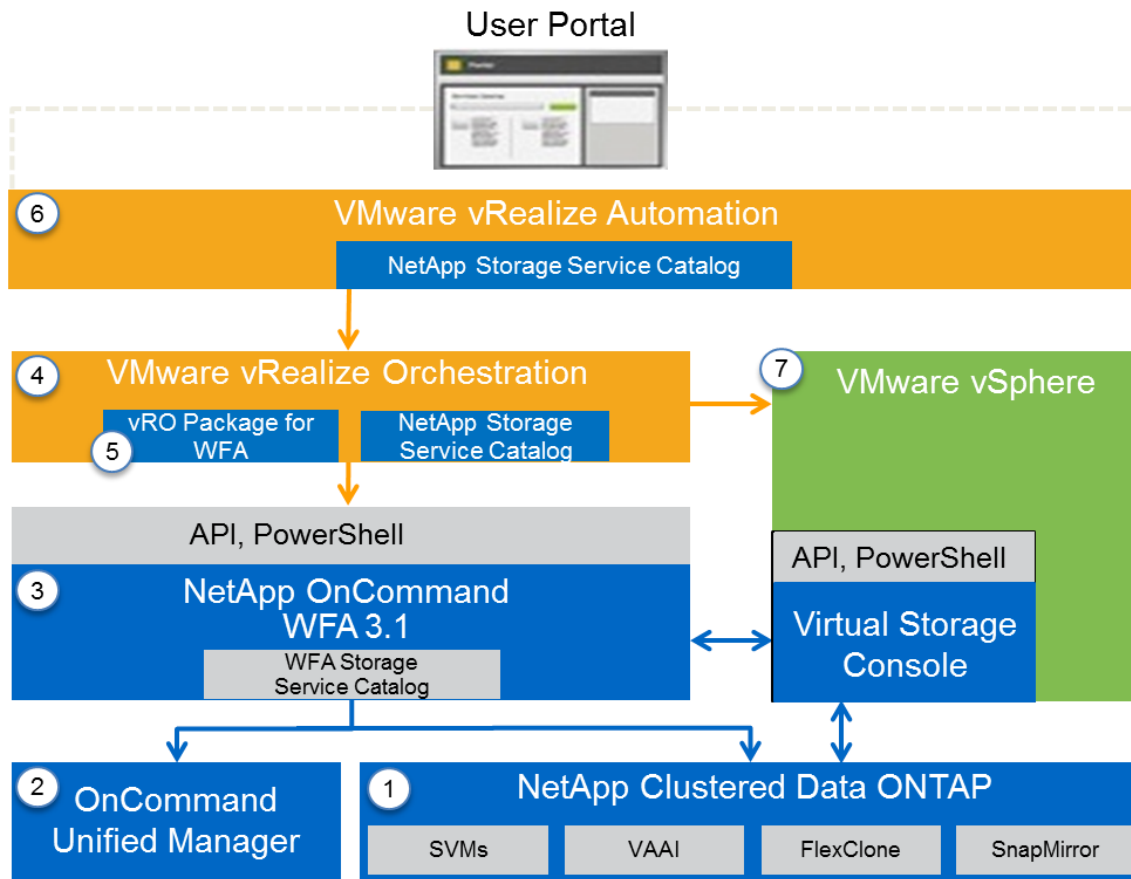
- **Personalized, business-aware governance.** Enables IT administrators to apply their own way of doing business to the cloud without changing organizational processes or policies. Enterprises gain the flexibility needed for business units to have different service levels, policies, and automation processes as appropriate for their needs.
- **Extensibility by design.** Provides a full spectrum of extensibility options that empower IT personnel to enable, adapt, and extend their cloud to work within their existing IT infrastructure and processes. This eliminates expensive service engagements while reducing risk.

## 2 NetApp and VMware Solution Components

The following required components, which correspond to the numbers in Figure 4, are used to create the NetApp and VMware SDS solution:

1. NetApp clustered Data ONTAP 8.3.1.
2. NetApp OnCommand Unified Manager 6.3.
3. NetApp OnCommand WFA 3.1 is a powerful framework for automating storage management tasks. It enables agility, flexibility, efficiency, and repeatability of best practices in NetApp clustered Data ONTAP storage environments. WFA has been designed to provide deep integration into the architecture of the clustered Data ONTAP system and assist in deploying best practices for cloud storage. Leveraging a standards-based web services interface, WFA allows higher-level orchestration products to leverage inherent knowledge with minimum added complexity.
4. VMware vRealize Orchestration 6.0 is a development and process automation platform that provides a library of extensible workflows that allow the creation and deployment of automated, configurable processes to manage the VMware vSphere infrastructure as well as other VMware and third-party technologies. Orchestrator exposes every operation in the vCenter Server API, allowing administrators to integrate all of these operations into their automated processes. Orchestrator also allows administrators to integrate with other management and administration solutions through its open plug-in architecture.
5. NetApp OnCommand WFA Package for vRealize Orchestration 3.0.1.
6. VMware vRealize Automation 6.2 provides the agility that a business needs by automating the delivery of personalized IT services. It provides a single solution for application release automation and support for various dev/ops automation tools abstracted from diverse infrastructure services. Through a self-service catalog, users request and manage a wide range of multivendor and multicloud applications, infrastructure, and custom services. Policy-based governance assures that users receive correctly sized resources across the service lifecycle. A flexible automation approach enables agility in the deployment of new IT services, while leveraging existing investments by mapping into the current infrastructure, processes, and environments.
7. VMware vSphere 6.0 is the industry-leading virtualization platform that provides high availability (HA) and resiliency for on-demand infrastructures that are suitable for building cloud environments. vSphere enables IT to meet SLAs for the most demanding business-critical applications at the lowest total cost of ownership (TCO).

Figure 4) SDS environment.



## 2.1 SDS with OnCommand WFA 3.1

A workflow is a repetitive process that consists of procedural tasks, such as:

- Provisioning, migrating, cloning, or decommissioning storage for databases or file systems
- Setting up a new virtualization environment, including storage switches and datastores
- Setting up storage for an application as part of an end-to-end orchestration process

To complete these tasks, storage architects can define workflows to follow best practices and meet organizational requirements.

NetApp OnCommand WFA is a software solution that offers one-click storage orchestration for NetApp SDS. It automates storage management tasks such as provisioning, migration, decommissioning, cloning, and building of specific workflows that processes require.

WFA provides a robust workflow design interface and a rich set of predefined content to simplify the workflow development process. Additionally, the [NetApp OnCommand WFA Community](#) site has downloadable content that can be customized to meet an organization's requirements.

WFA comes with supported, predefined building blocks that allow the system to be customized and extended to suit specific needs. It also features a web service API to allow WFA workflows to be triggered from virtually any enterprise orchestration application. Workflows are made available through a web services interface for use by orchestrators such as vRealize Orchestrator and vRealize Automation by using representational state transfer (REST) protocol requests. Additionally, the NetApp WFA workflow packages for WFA provide integration or a bridge between WFA, vRealize Orchestrator, and vRealize

Automation. This allows WFA workflows to be called by vRealize Orchestrator or by vRealize Automation through vRealize Orchestrator.

WFA relies on OnCommand Unified Manager to maintain status information about the storage infrastructure—available storage controllers, the presence and state of aggregates, volumes, and so on—and the allocation of storage into resource pools. WFA collects this data from OnCommand Unified Manager and stores it in a local cache on a regular basis. WFA can also collect and make use of data from vCenter and vCloud Director.

At its core, OnCommand WFA serves as a NetApp SDS engine. It is a software solution that helps automation and orchestration of storage management tasks such as provisioning, migration, decommissioning, and cloning. Administrators can use WFA to build workflows to complete tasks for unique processes.

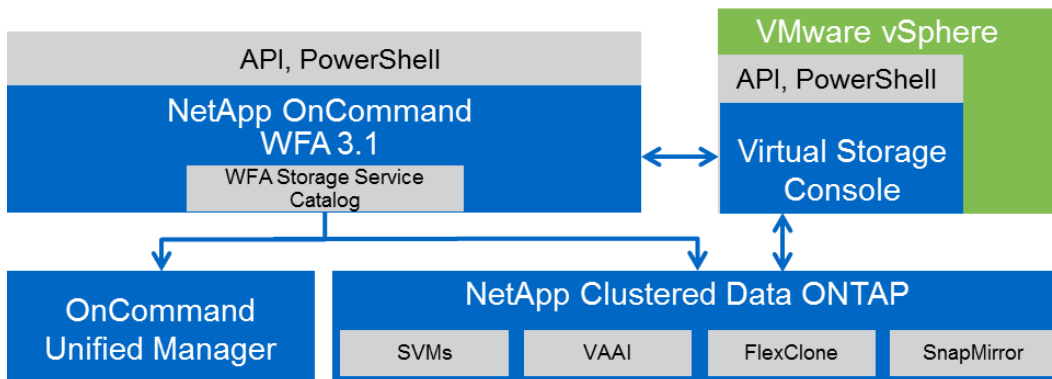
## OnCommand WFA Features

OnCommand WFA includes the following features:

- **Web services interface.** Invokes workflows from external portals and data center orchestration software.
- **Designer portal.** Includes several building blocks, such as commands, templates, finders, filters, and functions that are used to create workflows. The designer portal enables workflows to include advanced capabilities such as automated resource selection, row repetition (looping), and approval points.
- **Execution portal.** Executes workflows, verifies the status of workflow execution, and accesses logs.
- **Administration portal.** Sets up WFA, connects to data sources, and configures user credentials.
- **Storage automation store.** Downloads additional WFA packs.

Figure 5 provides an overview of OnCommand WFA 3.1.

Figure 5) OnCommand WFA 3.1 overview.



## 2.2 OnCommand Unified Manager for NetApp Storage Management

NetApp OnCommand Unified Manager is the primary monitoring and reporting tool for NetApp storage systems. It provides both a graphical user interface (GUI) and an API for integration with other management software. The API uses well-defined XML objects for communication, which provides transparency and interoperability.

Data collected by OnCommand Unified Manager is queried by WFA to provide a subset of relevant information for use with automation workflows. Multiple OnCommand Unified Manager systems can feed into a single WFA installation. The storage administrator can directly access OnCommand Unified Manager to create storage management reports and perform monitoring activities.

## 2.3 VMware vRealize Orchestrator

VMware vRealize Orchestrator 6.0 is drag-and-drop workflow software that helps automate IT service delivery and integrates VMware vRealize Suite with the rest of a business's IT environment. The software, included at no extra charge with VMware vCenter, enables administrators to develop complex automation tasks using the workflow designer. The workflows can be quickly accessed and launched directly from the VMware vSphere client, from VMware vRealize Automation, or through various triggering mechanisms.

### Key Benefits

vRealize Orchestration provides the following key benefits:

- Integrate VMware solutions into the IT environment and processes to rapidly accelerate IT operations.
- Reduce IT operating expenses and TCO of VMware solutions.
- Automate the cloud and accelerate transition to an IT-as-a-service (ITaaS) model.

### How VMware vRealize Orchestration Works

vRealize Orchestration is a powerful automation tool designed for system administrators and IT operations staff who must streamline tasks and remediation actions and integrate these functions with third-party IT operations software.

VMware vRealize Orchestration leverages the following features and components:

- **Workflow development and management.**
- **Workflow designer.** Produce simple or complex workflows by using easy-to-use, drag-and-drop workflow elements. Using Windows, Linux, and Mac OSs, build complex automation routines with just a few clicks.
- **Workflow engine.** Create workflows by using custom building blocks initiated by the user or by using plug-ins from vRealize Orchestration or from the VMware Solution Exchange.
- **Built-in workflow library.** Use tasks in the vRealize Orchestration library of workflows, including the following capabilities:
  - Take a Snapshot copy across a large number of VMs using a particular resource pool.
  - View resource utilization.
  - Notify the administrator by e-mail when a VM is powered on.
  - Create new or advanced workflows with building blocks.
  - Link actions to produce specific workflows.
- **Scripting engine.** Create new building blocks for the vRealize Orchestration platform. The scripting engine is enhanced with basic version control, variable type checking, namespace management, and exception handling. Create building blocks for actions, workflows, and policies.
- **Versioning.** All vRealize Orchestration workflows have version history, packaging, and rollback capabilities that allow basic change management during distribution of processes to different project stages and locations.
- **Checkpointing content database.** Each step in the workflow is saved in a content database, enabling server restarts without loss of state and context. This feature is especially useful for long-running processes.
- **Central management.** vRealize Orchestration provides a central way to manage processes and make sure that administrators use version-controlled scripts. Because it contains the scripts and process-related primitives in one place, administrators avoid releasing scripts without versioning and proper change control across the environment.

- **Diagnostics.** The vRealize Orchestration client's new debugging and failure-diagnostic capabilities provide infrastructure and operations workflow developers with a simpler, more efficient development experience.
- **Plug-in SDK.** The vRealize Orchestration plug-in Software Development Kit (SDK) jumpstarts those who are new to vRealize Orchestration plug-in development. It also enables advanced developers to integrate key SDK features into their existing development environments. The SDK consists of the following components:
  - Eclipse add-on for vRealize Orchestration plug-in development
  - Command-line utilities
  - vRealize Orchestration sample plug-ins
  - Developer guide
  - Best practice guide
- **Performance and scalability.** The software can execute hundreds or thousands of concurrent workflows on single or distributed architectures. By using the out-of-the-box clustering capabilities, architects can plan vRealize Orchestration deployments with cloud scalability in mind. The vRealize Orchestration cluster mode provides greater availability and enables dynamic scale-up and scale-down of orchestration capacity when used with external load balancers. If vRealize Orchestration servers become unavailable midway through a workflow execution, another vRealize Orchestration node completes the workflow with no service interruption.

## Security

vRealize Orchestration provides the following advanced security functions:

- Public key infrastructure to sign and encrypt content imported and exported between servers
- Digital rights management to control how exported content can be viewed, edited, and redistributed
- Secure Sockets Layer–encrypted communications between the desktop client and the server and HTTPS access to the web front end
- Advanced access rights management to provide control over access to processes and objects

## Integration

vRealize Orchestration provides the following integration features:

- **Cloud readiness.** vRealize Orchestration provides 100% coverage of the vSphere and VMware vCloud Director APIs and is integrated with other VMware solutions such as vCloud Director, VMware vCenter Update Manager, and vRealize Automation. With this type of coverage, administrators or developers can automate virtually any task in their cloud and virtual infrastructure.
- **Ecosystem.** Use plug-ins and workflows published on VMware Solution Exchange, a marketplace of extensible solution plug-ins, to connect to multiple VMware and third-party applications.
- **Virtual appliance.** The vRealize Orchestration virtual appliance significantly reduces the time and skill required to deploy vRealize Orchestration and provides a low-cost alternative to the traditional installation based on Windows. Distributed in open virtualization format, the appliance is prebuilt and preconfigured with all of the components included in the traditional installation. Additionally, the product can use the prebuilt directory services, a database, or external ones such as Active Directory or Oracle Database.
- **REST API.** vRealize Orchestration provides a REST API to enable other applications to execute workflows. The REST API has enhanced JavaScript Object Notation support and simplifies integration with vCenter single sign-on. These enhancements are configured to rapidly deploy vRealize Orchestration instances for test and development purposes and to scale up automation capacity as demand increases. Additionally, these enhancements leverage vRealize Orchestration workflows in a localized environment if the dedicated property files are used for the specific language.

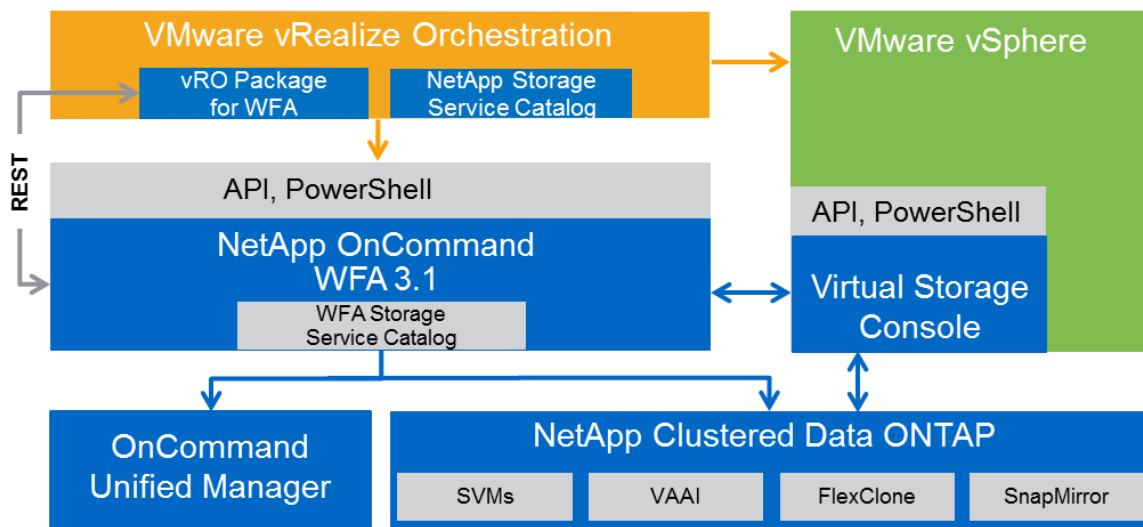


## VMware vRealize Orchestrator Support for SDS Using NetApp WFA

In this solution, vRealize Orchestrator acts as the central engine for orchestration and automation. Workflows for vCenter and vRealize Automation live here and are executed at this level. vRealize Orchestrator calls the workflows in WFA to deploy NetApp storage on an as-needed basis. Depending on the use case, vRealize Orchestrator can also act as the user and administration portal.

Figure 6 shows the vRealize Orchestration and NetApp WFA-based SDS architecture.

Figure 6) NetApp WFA and vRealize Orchestrator in SDS.



### 2.4 VMware vRealize Automation 6.2 in SDDC with SDS

VMware vRealize Automation provides the agility business needs by automating the delivery of personalized IT services. It provides a single solution for application release automation and support for various dev/ops automation tools, abstracted from diverse infrastructure services. Through a self-service catalog, users request and manage a wide range of multivendor and multicloud applications, infrastructure, and custom services. Policy-based governance assures that users receive correctly sized resources across the service lifecycle.

A flexible automation approach provides agility in deploying new IT services, while leveraging existing investments by mapping into the current infrastructure, processes, and environments.

#### Key Benefits

vRealize Automation 6.2 provides the following key benefits:

- Automate IT service delivery (applications, infrastructure, desktops, and any custom service) to rapidly respond to business needs.
- Enforce application deployment standards, resource quotas, and service levels with personalized, business-relevant policies.
- Protect investments in current and future technologies through broad multivendor, multicloud support and extensible design.
- Improve IT service delivery while lowering costs.

## What vRealize Automation Delivers

vRealize Automation accelerates the deployment and management of applications and compute services, thereby improving business agility and operational efficiency. The following capabilities empower IT to quickly demonstrate the value of deploying an automated, on-demand cloud infrastructure:

- **Comprehensive, purpose-built functionality.** vRealize Automation is a purpose-built, enterprise-proven solution for the delivery and ongoing management of private and hybrid cloud services, based on a broad range of deployment use cases from the world's most demanding environments.
- **Personalized, business-aware governance.** Enable IT administrators to apply their own way of doing business to the cloud without changing organizational processes or policies. Enterprises gain the flexibility needed for business units to have different service levels, policies, and automation processes, as appropriate for their needs.
- **Provisioning and management of application services.** Accelerate application deployment by streamlining the deployment process and by eliminating duplication of work using reusable components and blueprints.
- **Infrastructure delivery and lifecycle management.** Automate the end-to-end deployment of multivendor infrastructure, breaking down internal organizational silos that slow down IT service delivery.
- **Extensibility by design.** vRealize Automation provides a full spectrum of extensibility options that empower IT personnel to enable, adapt, and extend their cloud to work within their existing IT infrastructure and processes, thereby eliminating expensive service engagements while reducing risk.

## Key Features

vRealize Automation 6.2 provides the following key features:

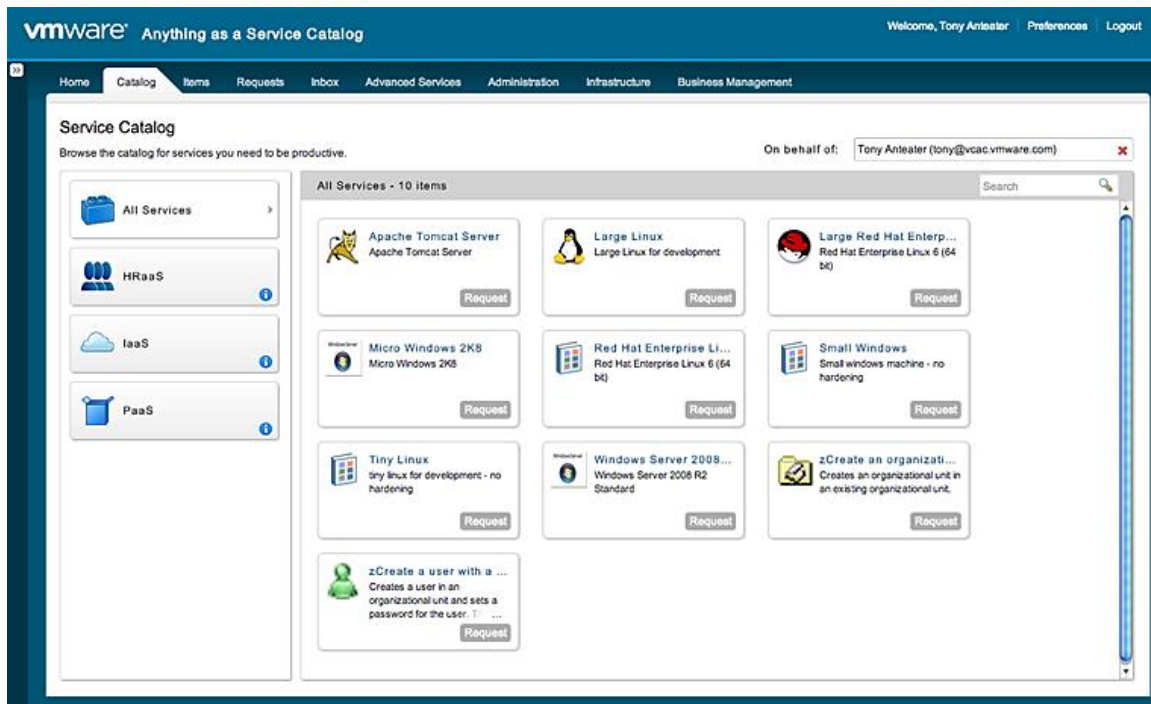
- **Unified IT service catalog.** Empower users to request and manage a variety of IT services that seamlessly span multivendor, multicloud platforms.
- **Automated service delivery.** Automate the end-to-end deployment and lifecycle management of infrastructure and application services leveraging existing infrastructure, tools, and processes.
- **Accelerated application deployment.** Leverage a single solution for application release automation and support for various dev/ops automation tools, abstracted from diverse infrastructure services.
- **Policy-based governance.** Enable lines of business and users to request and receive approved IT applications and resources based on business and IT policies.
- **Intelligent resource management.** Provide higher IT resource utilization through intelligent placement and reclamation of resources across virtual, physical, and cloud resources.
- **Flexible automation platform.** Rapidly configure, adapt, and extend VMware's comprehensive, purpose-built functionality to meet unique IT and business needs.
- **Ability to deliver anything as a service.** Provide a flexible platform that allows IT to design and automate the delivery of application and IT services.

## Unified Service Catalog

vRealize Automation 6.2 supports a unified service catalog and app store ordering experience in which users select from a personalized collection of applications, infrastructure services, and custom IT services. Service entitlements and optional approvals allow IT administrators to deliver a personalized IT service catalog that can be optimized to the specific needs of individual users or groups of users. Each tenant can have its own specific branding and user authentication directory services (for example, Active Directory-compatible or another Lightweight Directory Access Protocol-compatible directory service). Figure 7 shows an example service catalog.



Figure 7) Example service catalog (graphic provided by VMware).



## Simplified Application Release Automation


Ideally, development, test, staging, and production environments are in sync so that an application behaves the same way in production as it does in the testing environment. However, keeping different environments in sync is a very difficult task because production environments are modified independently because of emergency fixes or security issues.

IT needs a way to provide the exact same infrastructure and middleware configurations across all environments. vRealize Automation assists IT in achieving this by encapsulating frequently used updates, scripts, and properties in its update profile and promoting changes across environments.

When the update profiles feature is used, changes in one environment can be accurately replicated in another to provide consistency. These updated profile promotions can be applied to deployments of the same application in any environment. Notifications about promoted updates are available to the application deployers throughout the development lifecycle, letting them decide whether to apply the promoted update or not.

Figure 8 shows the use of the update profiles feature to save frequently used update scripts and properties and promote the profiles to different environments.

Figure 8) vRealize Automation uses update profiles to maintain consistency (graphic provided by VMware).



## Updates and Rollbacks

vRealize Automation also supports other day 2 operations on an application deployment through updates and rollbacks. Supported updates include configuration changes, code modifications, and scale-in/scale-out capabilities. Updates can be rolled back to a previous successful state with custom rollback scripts that cover the dependent components in an application model. Custom tasks can also be added in the update and rollback execution flows. A clustered deployment can scale out to allocate more resources or scale in to save unused compute resources if the application is not load constrained.

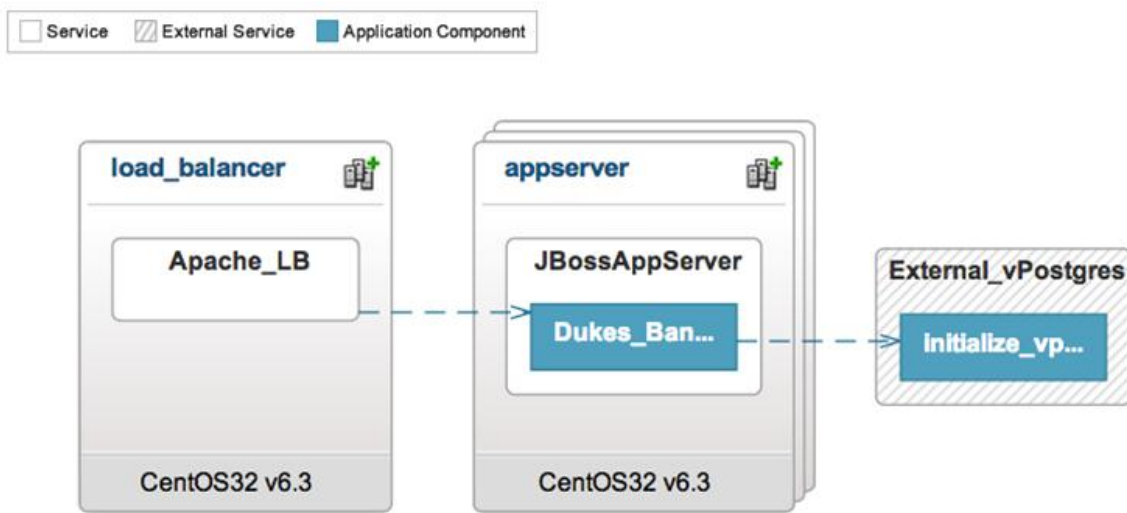
## Deploying New Applications and Connecting Them to Existing Resources

Not all applications are provisioned into new components and resources. IT might have some preconfigured, hardened services such as databases or firewalls that all applications must use and typically do not permit important services to be newly provisioned in production. vRealize Automation 6.2 provides enhanced modeling capabilities that allow applications to be modeled and provisioned into existing services.

When modeling a new application, existing services must be modeled in the blueprint. vRealize Automation enables this through a feature called external service, which is represented in a blueprint. In addition, there can be different instances of an external service in each environment. The external service profiles map the abstract external service in an application blueprint to the actual service instance in an environment. This process provides a way to model an application that uses existing services along with enough flexibility to separate out the differences in each environment. Reusing existing resources enables shorter setup time for an application, cost savings, and efficient use of resources.

Figure 9 shows the use of the external service feature to model a new application with existing services.

Figure 9) Model new applications with existing services using external service (graphic provided by VMware).



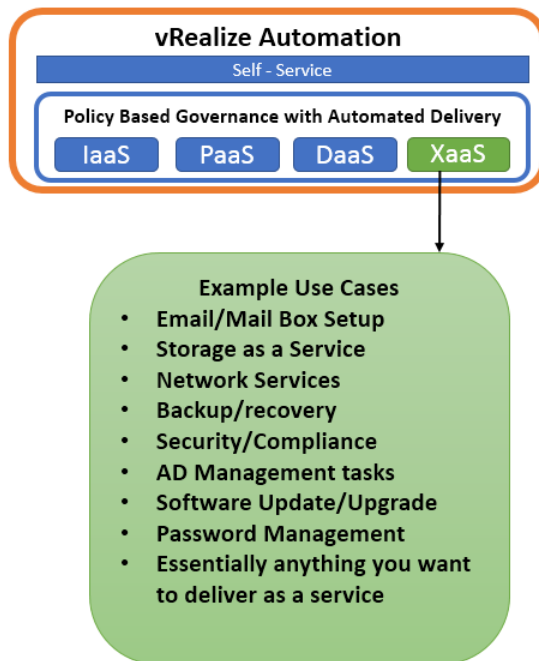
## Delivering Anything as a Service

The ideal cloud automation platform is one that can deliver any IT service on demand. vRealize Automation 6.2 is one step closer to realizing that goal, providing the capability to rapidly deploy new IT services with a wizard-driven service designer.

vRealize Automation provides standard capabilities for automating the delivery and management of applications, desktops, and infrastructure services. All of these capabilities are available by configuring purpose-built functionality to deliver personalized, business-relevant services to customers.

However, what if administrators want to deliver additional IT services not covered by out-of-the-box functionality? What if, for example, administrators want to automate the delivery of all tasks associated with onboarding a new employee? These can often include things such as updating Microsoft Active Directory, configuring a mail account, providing access to internal systems, and a number of other tasks. vRealize Automation can help automate the delivery of these and other IT operations through the advanced service designer feature. Figure 10 shows examples of services whose delivery can be automated through vRealize Automation.

Figure 10) Services whose delivery can be automated through vRealize Automation (graphic provided by VMware).



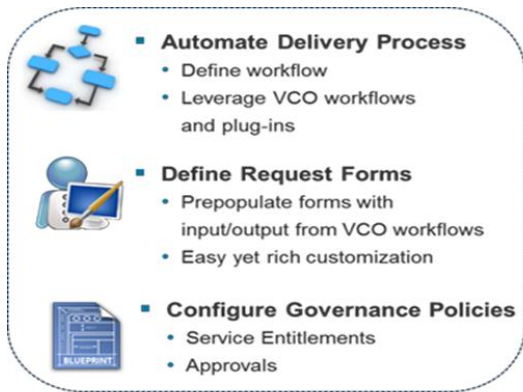
## Advanced Service Designer

The advanced service designer feature is a wizard used to design the end-to-end process for delivering a service from request to automated delivery. These custom services can be published in the vRealize Automation catalog along with other application and infrastructure services. The wizard helps administrators define service delivery capabilities, user interaction, and entitlements using the following method:

1. The administrator defines the automated workflow that is used to deliver the service. This could include existing vRealize Orchestrator workflows and plug-ins, as well as custom scripts.
2. The wizard prepopulates a request form with information that the workflow needs in order to configure and automate the delivery of the service being requested. The request form is a simple-to-use WYSIWYG form designer with capabilities to allow sophisticated interaction with the requestor.
3. The administrator defines service entitlement policies, including optional approval workflows if necessary.

The advanced service designer allows administrators to rapidly deploy new IT services, leveraging VMware and partner-supplied vRealize Orchestrator workflows and plug-ins. The services that can be automated are limitless. Figure 11 shows the method used to design service automation workflows.

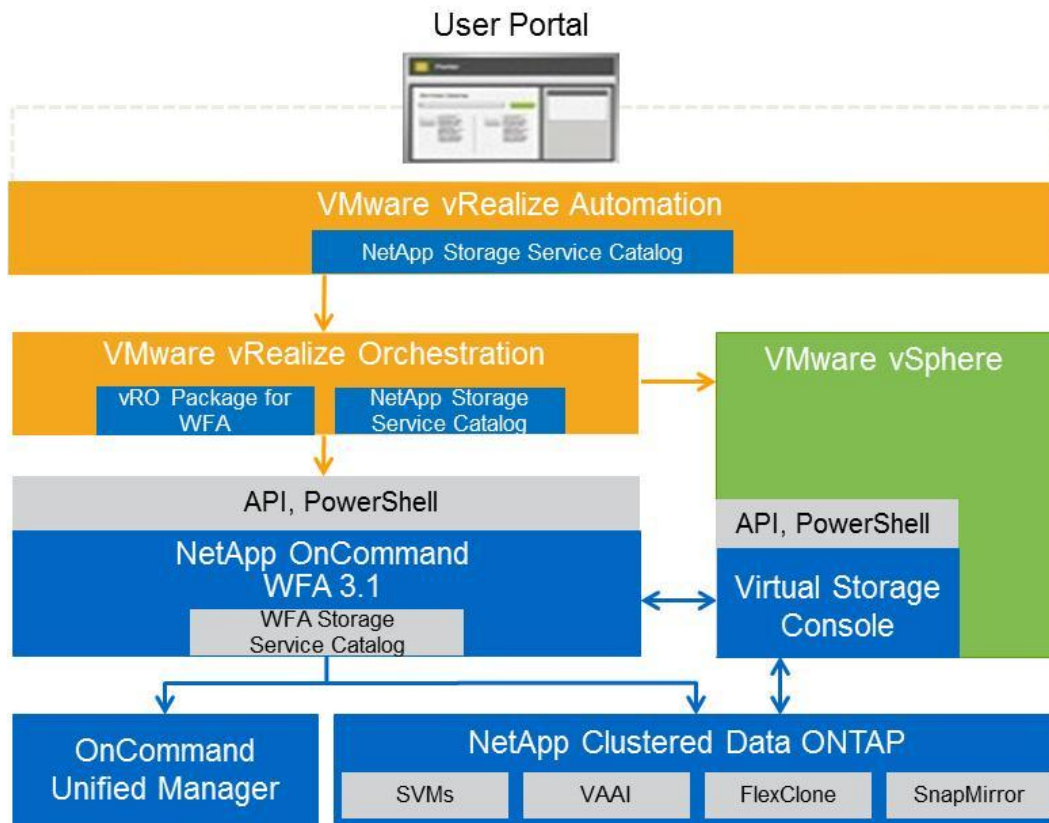
Figure 11) Advanced service designer workflow (graphic provided by VMware).



## Requesting SDS with vRealize Automation

Figure 12 shows an example configuration in which vRealize Automation acts as the central portal for both users and administrators to request services in the cloud. vRealize Automation calls vRealize Orchestrator and uses workflows contained there to perform orchestration and automation tasks inside the cloud environment. vRealize Automation also serves as the lifecycle management tool for the components.

Figure 12) vRealize Automation example configuration.



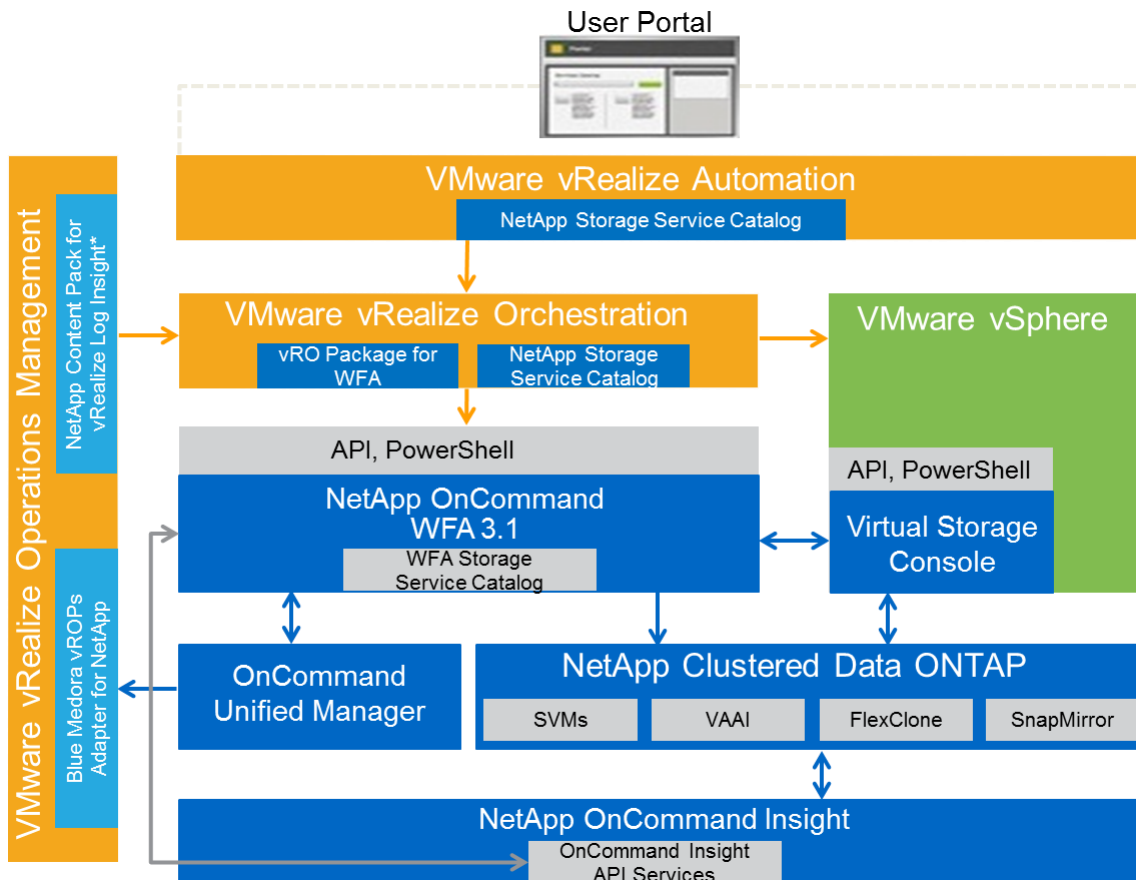
### 3 Optional Component: Performance Monitoring and Problem Remediation Management

Although performance monitoring and problem remediation management are optional components of the NetApp and VMware SDS solution, they are also an important part of the solution. The following components provide insight into the overall solution and work together to provide information for analysis, reporting, trending, and capacity planning:

- NetApp OnCommand Insight
- VMware vRealize Operations
- Blue Medora vRealize Operations Management Pack for NetApp
- vRealize Log Insight
- NetApp vRealize Log Insight Content Pack

Figure 13 shows an example of the NetApp and VMware SDDC solution with performance monitoring and management.

Figure 13) NetApp and VMware SDDC solution with performance monitoring and management.



#### 3.1 OnCommand Insight

OnCommand Insight is part of the NetApp OnCommand product portfolio. It helps control, automate, and analyze your storage infrastructure. OnCommand Insight optimizes the existing infrastructure for maximum benefit, while simplifying the process of determining what and when to buy. It also takes the



risk out of complex technology migrations. OnCommand Insight lets you manage storage as an end-to-end service and integrate storage into your entire IT service delivery chain.

## Key Business Drivers

OnCommand Insight enables the following key business drivers:


- Quick time to market and deep visibility into complex traditional and cloud deployments for heterogeneous storage and virtual environments
- Complex multitenant cost accounting, showback, and billing metrics provided in cloud service provider and enterprise environments
- Delivering big data analytics to mobile platforms of choice: Analyze billions of data points for quick, fact-based decision making

## Key Features

OnCommand Insight leverages the following key features:

- **Visibility.** OnCommand Insight generates easy-to-understand visualizations of complex storage environments to help understand how physical and virtual entities are interdependent, as shown in Figure 14.

Figure 14) NetApp OnCommand Insight visualization of complex storage environment.



- **Automated workload analysis (bullies versus victims).** The end-to-end configuration discovered by OnCommand Insight is used to track changes in VMs down to storage volumes to assist with troubleshooting exercises when multiple workloads compete with each other in a shared multitenant infrastructure. There are out-of-the-box capabilities that identify bully and victim workloads to eliminate manual data mining from millions of capacity performance records collected across the entire presentation path. Figure 15 and Figure 16, respectively, show bully and victim identification capabilities.

Figure 15) Bully workloads.

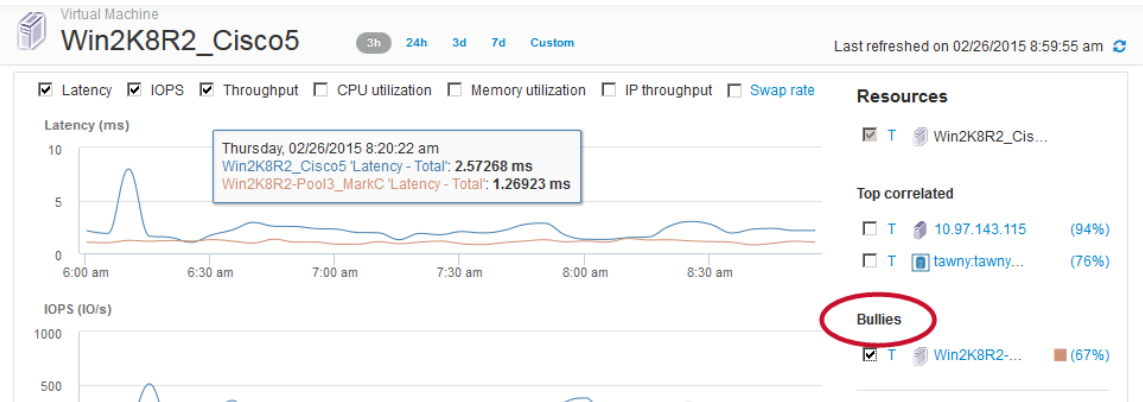
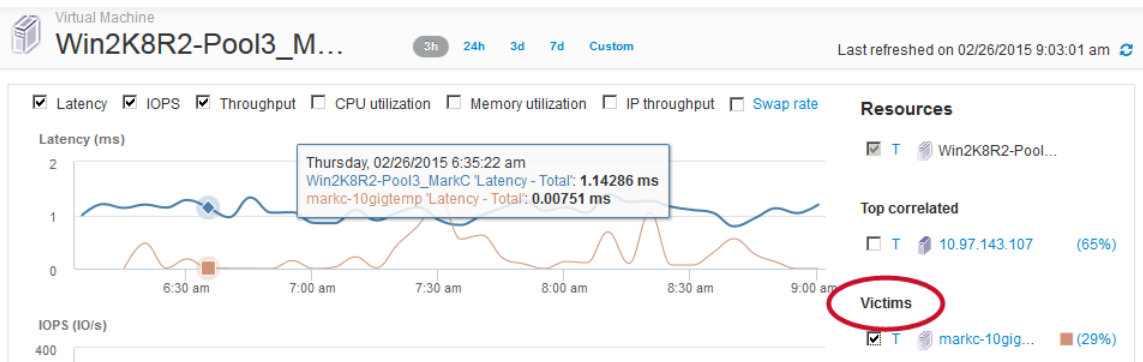


Figure 16) Victim workloads.

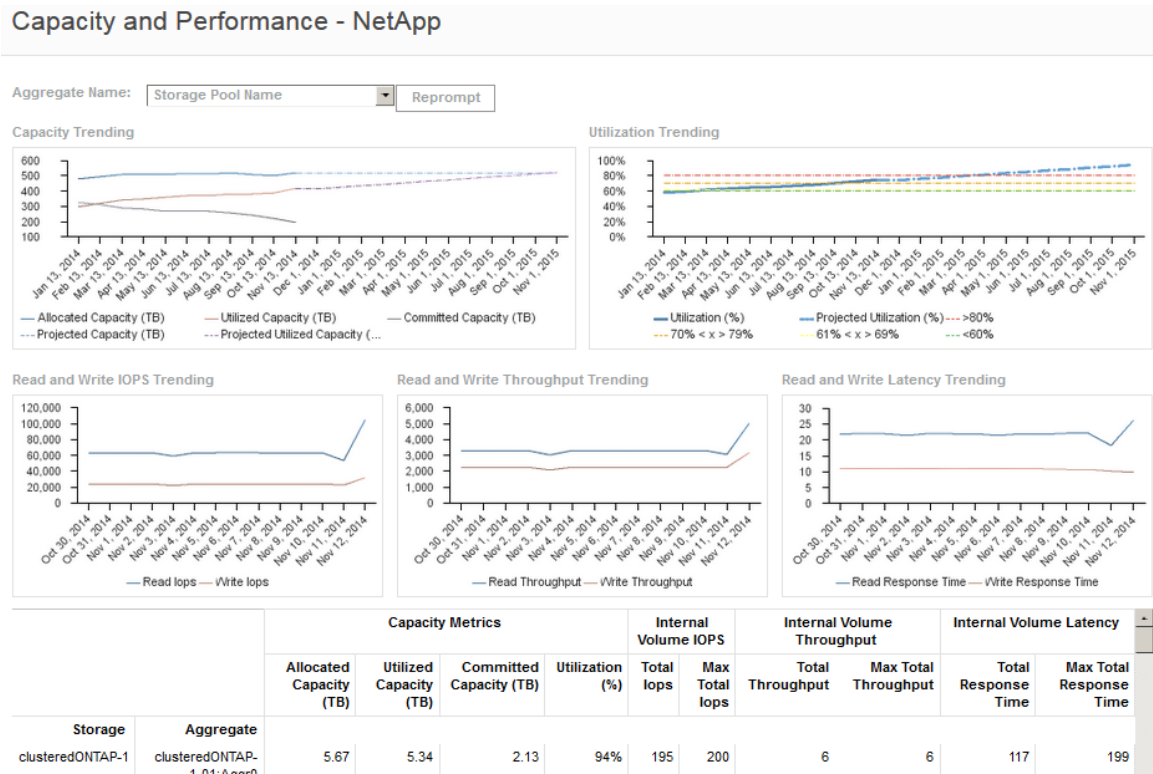


- Risk mitigation, capacity, configuration, and performance monitoring capabilities for NetApp storage.** OnCommand Insight helps customers proactively identify issues by monitoring VMs, hypervisor hosts, multiprotocol SANs (FC, iSCSI, and FCoE) and switches, storage virtualizers, storage arrays, controllers, storage pools, storage volumes, and physical disks. Across the entire storage presentation path, OnCommand Insight helps administrators mitigate risk by:
  - Automatically monitoring configuration changes: New datastores are provisioned, new zones are created, new LUNs are provisioned, new VMs are added to hypervisors, or any other decommissioning operations for the monitored environment are performed.
  - Proactively highlighting vulnerabilities in the environment based on best practice, rules-based validation: OnCommand Insight shows volumes and replicas that share the same physical disks so you can identify volumes that are not shared by all hosts assigned to an application.
  - Alerting the user about policy breaches based on out-of-the-box thresholds: Single points of failure, missing redundancies, and missing virtual cluster paths are noted.

Figure 17 shows the deep capacity and performance trending capabilities for NetApp devices.

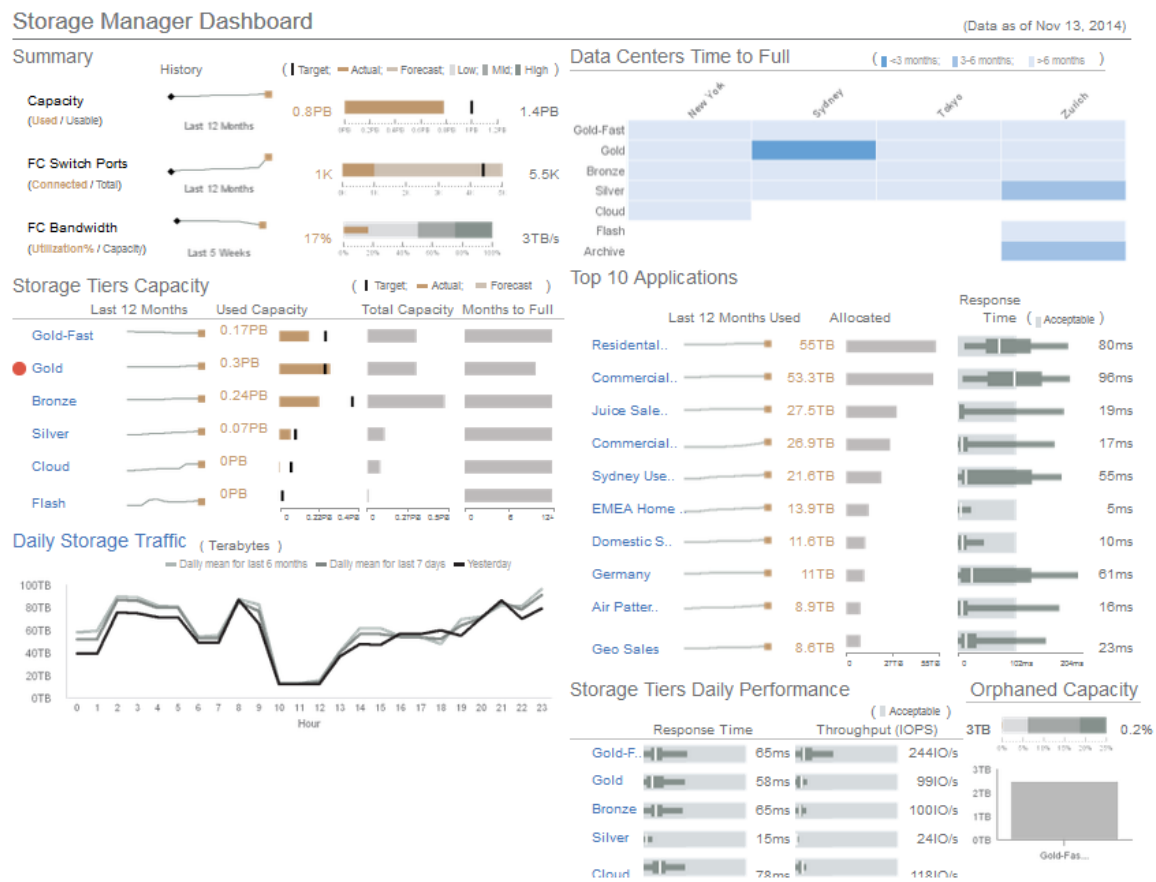


Figure 17) Capacity and performance data for NetApp devices.



- **Capacity consumption, trending, forecasting, and efficiency.** OnCommand Insight is used in large enterprise deployments for historical trending, forecasting, billing, and showback use cases. In a shared infrastructure, it is imperative to have the ability to show capacity allocation across different physical and logical entities that consume storage resources, as well as their accurate fractional billable resource usage. OnCommand Insight addresses these complex requirements with a flexible cost allocation model. Figure 18 shows an out-of-the-box dashboard that displays capacity trending, forecasting, and performance. Customers can easily find out when they will run out of capacity across data centers and tiers, as well as utilization trending across applications.

Figure 18) Dashboard showing capacity trending, forecasting, and performance.



- **Business value of OnCommand Insight for cloud deployments and SDS.** OnCommand Insight helps to monitor the infrastructure and platforms for the enterprise hybrid cloud infrastructure in a single, simple-to-deploy, agentless solution and is integral to successful SDS deployments.

Providing the ability to serve your business without worrying about vendor-specific technologies and products is one of the key building blocks for a true SDS strategy. OnCommand Insight delivers this capability out of the box. Regardless of a customers' storage, switch, server, or platform vendor, that customer has a single management platform. Another building block of a true SDDC is independence from any cloud deployment model. As customers go through transitions from a traditional to a private cloud and from a private to a hybrid cloud, OnCommand Insight gives them the ability to monitor on- and off-premises resources while continuing to assist with traditional storage resource management use cases.

### 3.2 vRealize Operations

vRealize Operations delivers intelligent operations management across physical, virtual, and cloud infrastructures from VMware vSphere and Microsoft Hyper-V to Amazon Web Services. It correlates data from applications to storage in a unified, easy-to-use management tool that provides control over performance, capacity, and configuration with predictive analytics driving proactive action and policy-based automation.

vRealize Operations Manager also enables the following capabilities:

- Proactive identification and remediation of emerging performance, capacity, and configuration issues
- Comprehensive visibility across applications and infrastructure in a single console

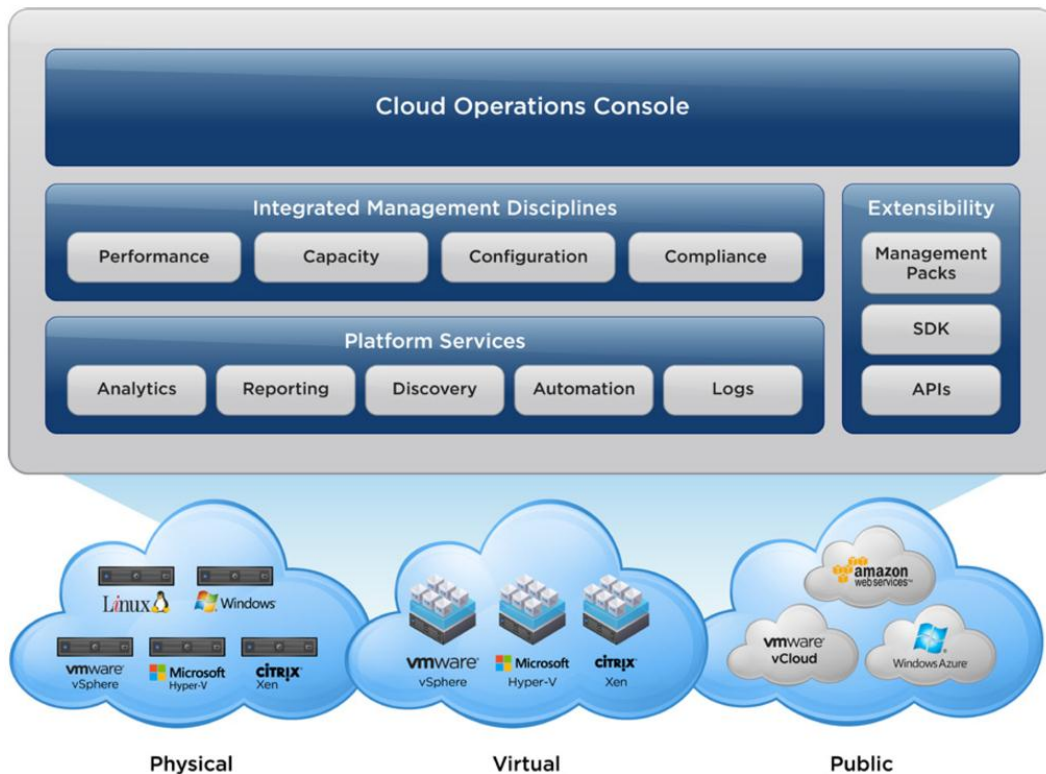
- Automated capacity optimization and planning
- Enforcement of standards for continuous compliance
- Open and extensible platform that leverages existing investments with third-party management packs for Microsoft, SAP, and more

Traditional operations management systems do not meet the requirements of today's virtual and cloud infrastructures. They make IT too reactive because they lack the intelligence to aggregate, correlate, and analyze metrics across applications and infrastructure stacks.

vRealize Operations is built on a scale-out, resilient platform designed to deliver intelligent operational insights to simplify and automate management of applications and infrastructure across virtual, physical, and cloud environments from vSphere to Hyper-V to Amazon Web Services and more.

With vRealize Operations, IT organizations of all sizes can improve performance, avoid business disruption, and become more efficient with comprehensive visibility across applications and infrastructure in one place. Figure 19 shows the capability of the Cloud Operations Console.

Figure 19) vRealize Operations in the cloud (graphic provided by VMware).



vRealize Operations delivers:

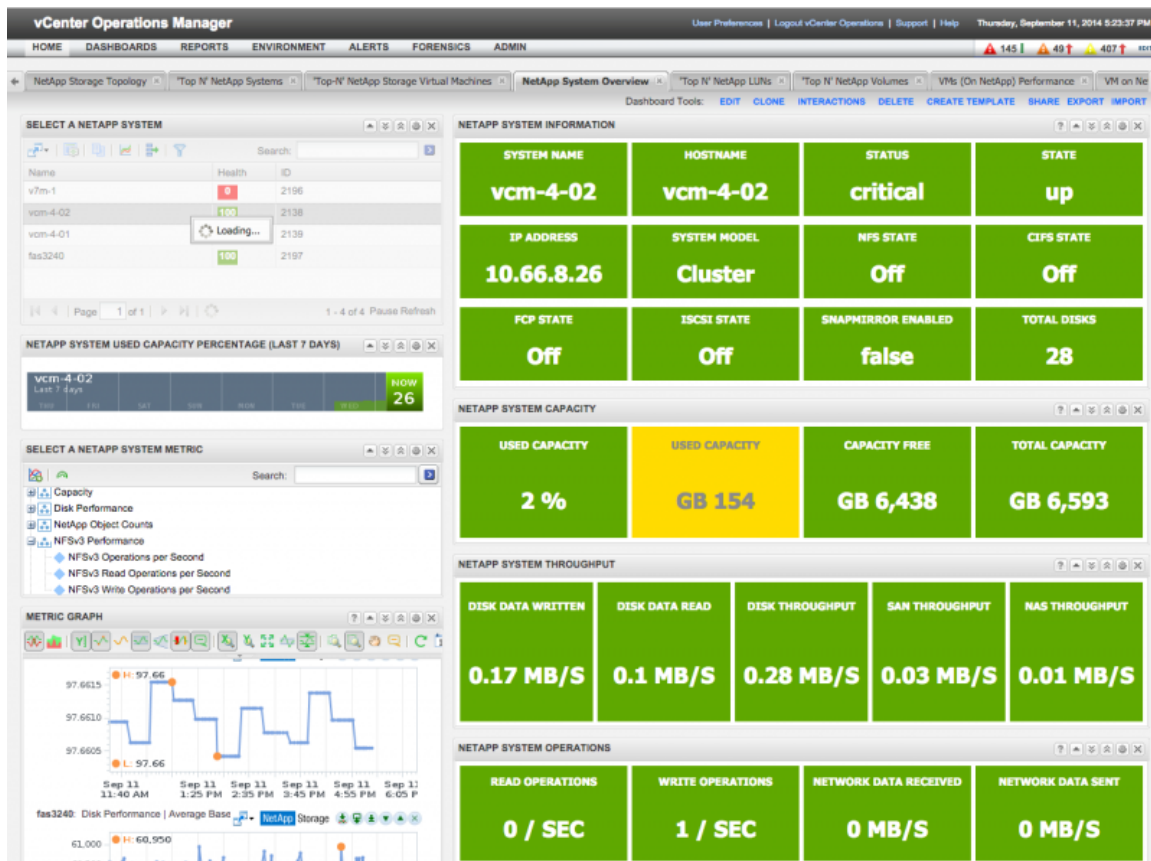
- **Intelligent operations.** Self-learning tools, predictive analytics, and smart alerts about application and infrastructure health enable proactive identification and remediation of emerging performance, capacity, and configuration issues.
- **Policy-based automation.** Out-of-the-box and customizable policies for critical IT operations are associated with smart alerts, guided remediation, and compliance standards to deliver recommendations or trigger actions that optimize performance and capacity and enforce configuration standards.

- **Unified management.** An open and extensible platform, supported by third-party management packs for Microsoft, SAP, and others, provides complete visibility in a single console across applications, storage, and network devices.

### 3.3 Blue Medora vRealize Operations Management Pack for NetApp

Blue Medora vRealize Operations Management Pack for NetApp storage provides end-to-end visibility, analytics, and capacity planning for workloads running on NetApp storage. The Management Pack for NetApp storage uses the vRealize Operations analytics engine to unveil deep insights into NetApp storage capacity and performance issues. Automated analysis exposes early indications of impending performance and capacity issues before they can negatively affect the business. Figure 20 shows the information provided by Blue Medora vRealize Operations Management Pack for NetApp storage.

Figure 20) Blue Medora vRealize Operations plug-in for NetApp storage.



vRealize Operations is the tool used by VMware administrators. The Blue Medora plug-in links the VMware asset information to the NetApp asset information.

IT shops in which the storage team is the same as the VMware team use vCenter and OnCommand System Manager. The ability to share data in context between larger teams is what really reduces mean time to identification and repair.

The Management Pack for NetApp is developed and distributed directly by Blue Medora.

### 3.4 vRealize Log Insight

VMware vRealize Log Insight delivers real-time log management for VMware environments, with machine learning-based intelligent grouping, high-performance search, and better troubleshooting across physical,

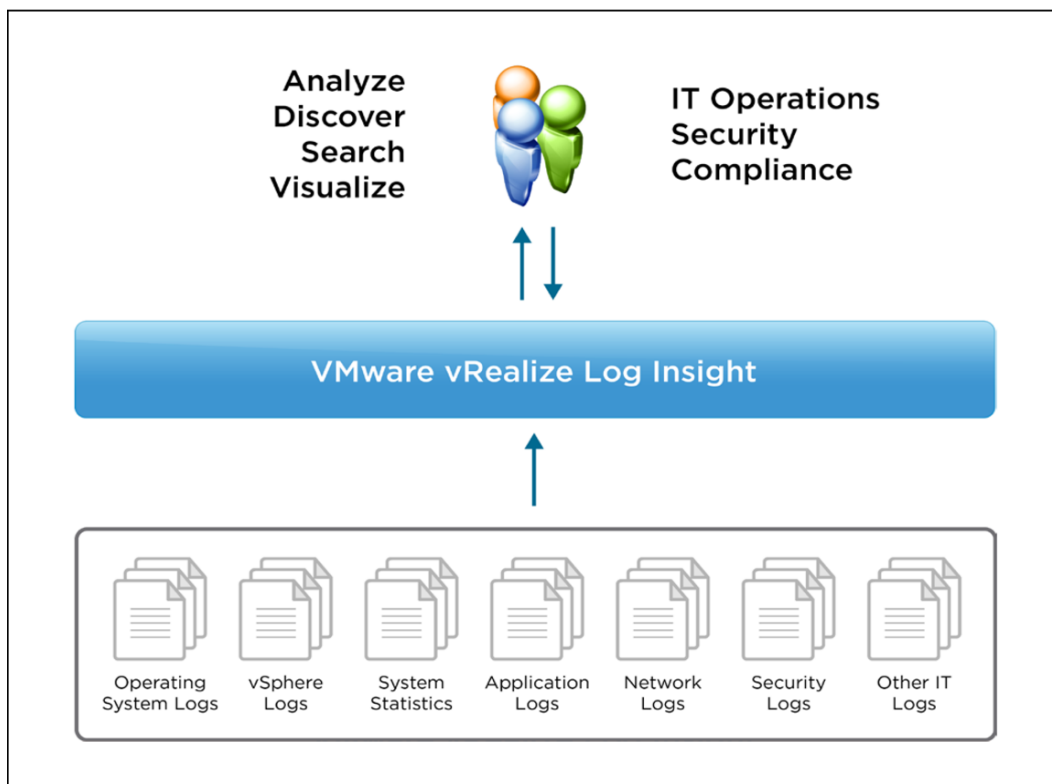
virtual, and cloud environments. With an integrated cloud operations management approach, vRealize Log Insight provides the operational intelligence and enterprise-wide visibility needed to proactively enable service levels and operational efficiency in dynamic hybrid cloud environments.

## Key Benefits

vRealize Log Insight provides the following key benefits:

- More than six times faster speed than the leading solution. In recent internal testing, vRealize Log Insight was found to be six times faster than the leading solution in query tests across 100 million log messages.
- Intuitive, easy-to-use, graphical interface for simple interactive searches as well as deep analytical queries. This benefit provides quick insights, enables fast time to value, and provides improved operational efficiency.
- Built-in knowledge and native support for VMware vSphere. This benefit makes vRealize Log Insight the solution best suited for your VMware environment. Figure 21 shows the flow of information to and from vRealize Log Insight.

Figure 21) vRealize Log Insight for real-time log management (graphic provided by VMware).



## Key Capabilities

vRealize Log Insight includes the following key capabilities:

- Collects and analyzes all types of machine-generated log data: application logs, network traces, configuration files, messages, performance data, system state dumps, and more.
- Enables administrators to connect to everything in their environment—operating systems, applications, storage arrays, firewalls, and network devices—providing a single location to collect, store, and analyze logs at scale. vRealize Log Insight features an intuitive GUI that makes it easy to run interactive searches as well as deep analytical queries for quick, actionable insights.

- Includes new cluster technology to allow your log analytics efforts to scale to new heights.
- Automatically chooses the best visualization for your data, saving you valuable time.
- Adds structure to all types of unstructured log data, enabling administrators to troubleshoot quickly, without needing to know the data beforehand.
- Delivers real-time monitoring, search, and log analytics, coupled with a dashboard for stored queries, reports, and alerts, enabling correlation of events across multiple tiers of a hybrid cloud environment.
- Helps users find virtual infrastructure problems with the addition of intelligent grouping, a machine-learning technology that can group related data together to aid in detecting problems rapidly.
- Comes with built-in knowledge and native support for vSphere and other VMware products, such as Horizon View, vRealize Operations Suite, and vRealize Automation.
- VMware and third-party extensions are available on the VMware solution exchange.
- Integrates with the VMware vRealize Operations platform to bring unstructured data (such as log files) together with structured data (such as metrics and KPIs), offering a significantly enhanced solution for end-to-end operations management.

vRealize Log Insight provides the following business benefits for IT organizations:

- Faster time to value, enabled by a powerful, intuitive, easy-to-use interface
- Lower operating expenses and higher operational efficiency, gained from a significant reduction in troubleshooting times, improved mean time to resolution, and reduced manual effort
- Predictable pricing model that includes an unlimited amount of data and does not require buying licenses based on peak usage and worst-case scenarios
- ROI and value maximization, achieved by using Log Insight with vRealize Operations, enabling you to leverage best-in-class tools for a significantly enhanced operations management experience

### 3.5 NetApp vRealize Log Insight Content Pack

The NetApp Data ONTAP Content Pack for VMware vRealize Log Insight enables users to take advantage of VMware vRealize Log Insight. It delivers automated log management that helps provide operational efficiency in dynamic hybrid cloud environments. With the NetApp Data ONTAP Content Pack, you can perform root cause analysis and troubleshoot your entire VMware environment on NetApp storage from a single interface. The NetApp Data ONTAP Content Pack for VMware vRealize Log Insight provides the following features:

- A set of read-only dashboards
- Predefined fields containing data extracted from the clustered Data ONTAP logs
- Predefined queries
- Alerts using the NetApp Event Management System (provided with Data ONTAP)

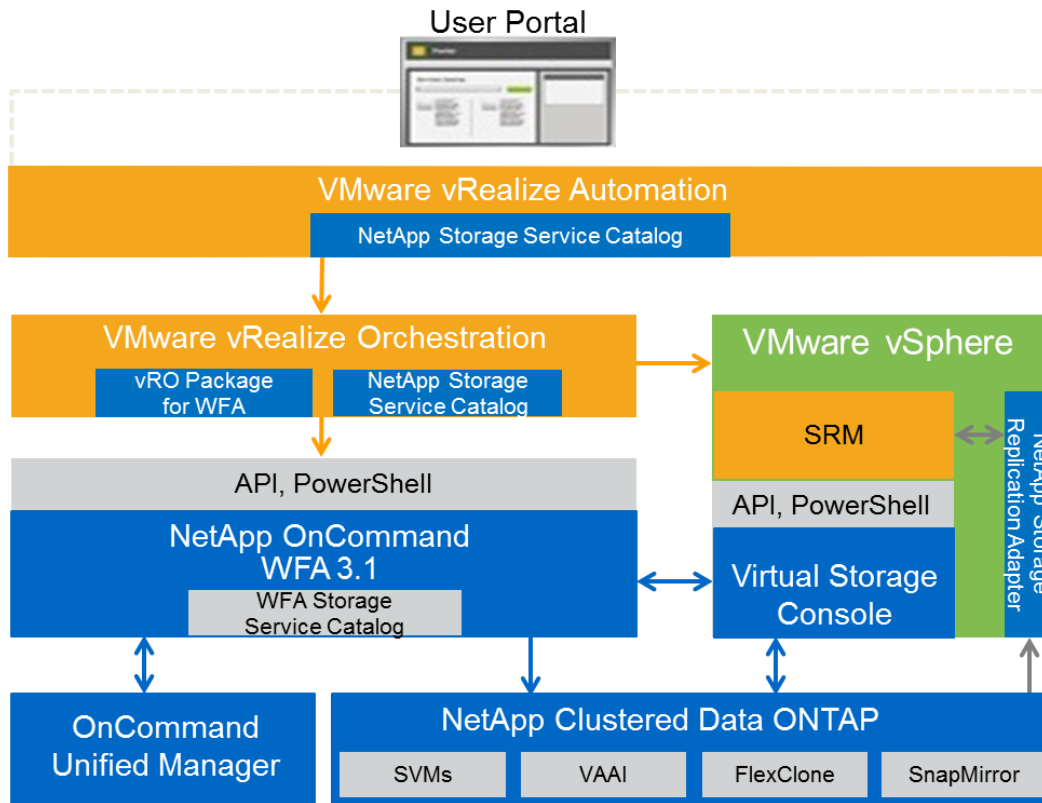
Although VMware delivers a content pack for VMware-related information, the NetApp Content Pack is custom designed by NetApp to provide information specific to NetApp. When used with Log Insight, the NetApp Content Pack provides monitoring and analysis of NetApp clustered Data ONTAP logs.

## 4 Optional Component: Disaster Recovery in a NetApp and VMware SDDC Environment

A virtualized environment using VMware vCenter SRM with NetApp storage provides an infrastructure with unparalleled opportunities to implement real, working DR processes that are quick and easy to test, consume little additional storage, and significantly reduce the recovery time objective and recovery point objective. Figure 22 shows the NetApp and VMware SDDC with DR.



Figure 22) NetApp and VMware SDDC with DR.



## 4.1 NetApp SnapMirror

SnapMirror provides data replication in a VMware vCenter SRM and NetApp environment. Built on NetApp Snapshot technology, SnapMirror replication is extremely efficient because it replicates only the 4KB blocks that were changed or added since the previous update. SnapMirror is easily configured either by using NetApp OnCommand System Manager or the Data ONTAP command line interface.

For cases in which the primary storage is not completely lost, SnapMirror provides an efficient means of resynchronizing the primary and DR sites. SnapMirror can resynchronize the two sites, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationships.

SnapMirror in clustered Data ONTAP provides asynchronous volume-level replication based on a configured replication update interval. SnapMirror uses NetApp Snapshot technology as part of the replication process.

Clustered Data ONTAP 8.1 or later provides the following replication capabilities:

- **Data protection mirrors.** Replication to create a backup copy within the same cluster (intracluster) or to create a DR copy in a different cluster (intercluster).
- **Load-sharing mirrors.** Replication from one volume to multiple volumes in the same cluster to distribute a read-only workload across a cluster.

**Note:** The load-sharing mirrors' capability is supported in SRA 2.1.

## Basics of SnapMirror Replication

When a scheduler triggers a replication update, the following operation occurs:

1. A new Snapshot copy is created on the source volume.
2. The block-level difference between the new Snapshot copy and the last replication Snapshot copy is determined and then transferred to the destination volume. This transfer includes other Snapshot copies that were created between the last replication Snapshot copy and the new one.
3. When the transfer is complete, the new Snapshot copy exists on the destination volume.

A SnapMirror destination volume is available for read-only access if it is shared by using the CIFS protocol or exported by using the NFS protocol. A LUN in the replicated volume can be made available to a client that supports a connection to read-only LUNs by using NetApp FlexClone® technology.

**Note:** Replication occurs at the volume level. Qtrees can be created in clustered Data ONTAP and replicated along with the replicated volume; however, individual qtrees cannot be separately replicated.

VMware vCenter SRM recovery plans can be resynchronized in either direction after a failover without recopying the entire volume. If a relationship is resynchronized in the reverse direction, only new data that was written since the last successful synchronization of the Snapshot copy is sent back to the destination.

SnapMirror relationships in clustered Data ONTAP 8.1 must be managed by a cluster administrator; administration cannot be delegated to an SVM administrator. For clustered Data ONTAP 8.2 and later, a cluster administrator can delegate the management of SnapMirror relationships to an SVM administrator.

VMware vCenter SRM Array Manager still must be managed by a cluster administrator; administration cannot be delegated to an SVM administrator.

**Note:** For up-to-date technical specifications, refer to the NetApp [SnapMirror Data Replication](#) website.

## 4.2 VMware vCenter SRM

One of the most time-consuming parts of a DR failover in a VMware environment is the execution of the steps necessary to inventory, register, reconfigure, and power up VMs at the DR site. VMware has solved these problems with vCenter SRM. Implementing vCenter SRM in an environment using NetApp clustered Data ONTAP provides business continuity and a DR solution that helps administrators plan, test, and run the recovery of VMs between a protected vCenter Server site and a recovery vCenter Server site. vCenter SRM works in conjunction with NetApp SnapMirror replication technology through NetApp SRA. NetApp SRA is a small software package that is installed on each vCenter SRM server and is available to any customer using vCenter SRM.

## 4.3 NetApp SRA and SVMs

NetApp SRA is a storage-specific plug-in for vCenter SRM that enables interaction between SRM and NetApp SVMs. The adapter interacts with the SVM on behalf of vCenter SRM to discover the SVMs and their associated datastores and raw device mapping (RDM) devices that are connected to vSphere. The adapter manages failover and test failover of the VMs associated with these storage objects. For more information about the new and enhanced features of vCenter SRM, refer to the [vCenter Site Recovery Manager documentation](#).

vCenter SRM manages DR operations by pairing two NetApp SVMs. It offloads all of the storage tasks to the NetApp SRA, uses Array Manager to pair the protected site and the recovery site SVMs, and identifies all SnapMirror relationships between the paired SVMs.

**Note:** Clustered Data ONTAP pairs SVMs, but does not pair controllers. Therefore, communication is with the SVM, and vCenter SRM is not aware of the individual controllers in a cluster.

## 4.4 Disaster Recovery in a NetApp and VMware SDDC Environment

In this solution, the following components must have DR enabled:



- An orchestration and automation environment that consists of vRealize Automation, vRealize Orchestrator, and WFA, specifically
- A monitoring and management environment that is made up of vRealize Operations, OnCommand Insight, and OnCommand Unified Manager

Although these components are part of the overall solution, with regard to DR, they can be considered subenvironments of the solution. However, both components can have DR enabled in a very similar manner.

In order to protect the orchestration and automation subenvironment, you must protect vRealize Automation, vRealize Orchestrator, and WFA and the vCenter to which they refer, specifically. Verify that these individual components are connected to vCenter SRM and fail them over together as a group as follows:

1. Add an overall management vCenter and vCenter SRM at both your primary and secondary sites.
2. Make sure another vCenter is running inside the management vCenters (which means that your vCenter is either an appliance or a VM running vCenter).
3. Make sure that you have vRealize Automation, vRealize Orchestrator, and WFA pointed at the virtualized vCenter (or vCenter appliance) inside the management vCenters for the VMs (in this case vRealize Automation, vRealize Orchestrator, and WFA) that you want to fail over using vCenter SRM.

By using this methodology you can safely fail over your orchestration and automation environment.

The methodology for enabling DR on vRealize Operations, vRealize Log Insight, and OnCommand Insight is very similar:

1. Make sure that you have two management vCenters and vCenter SRM at each site.
2. Make sure that you have a virtualized vCenter that serves as the vCenter environment that vRealize Operations and vRealize Log Insight use.
3. Fail over vRealize Operations, Log Insight, and NetApp OnCommand Insight at the same time using vCenter SRM, NetApp SnapMirror, and NetApp SRA for vCenter SRM.

## 5 Terminologies and Functionality

Table 1 defines the terminologies and technologies referred to in this document. Figure 4 illustrates where each technology is placed within a generic solution.

Table 1) Terminology and technology used in this document.

Terminology	Primary Technologies	Secondary Technologies	Definition
Cloud-enabled storage	NetApp clustered Data ONTAP	None	<p>An enterprise storage environment that supports orchestration of its storage. The orchestration must include the following components, at the very least:</p> <ul style="list-style-type: none"> <li>• Storage creation</li> <li>• Storage capabilities enablement</li> <li>• Storage removal</li> <li>• Backup and recovery</li> <li>• DR</li> <li>• Storage network configuration</li> </ul> <p>This is made much easier by a storage abstraction (or storage virtualization) layer that is part of each of the components being orchestrated.</p> <p>A true storage abstraction layer acts very similarly to</p>

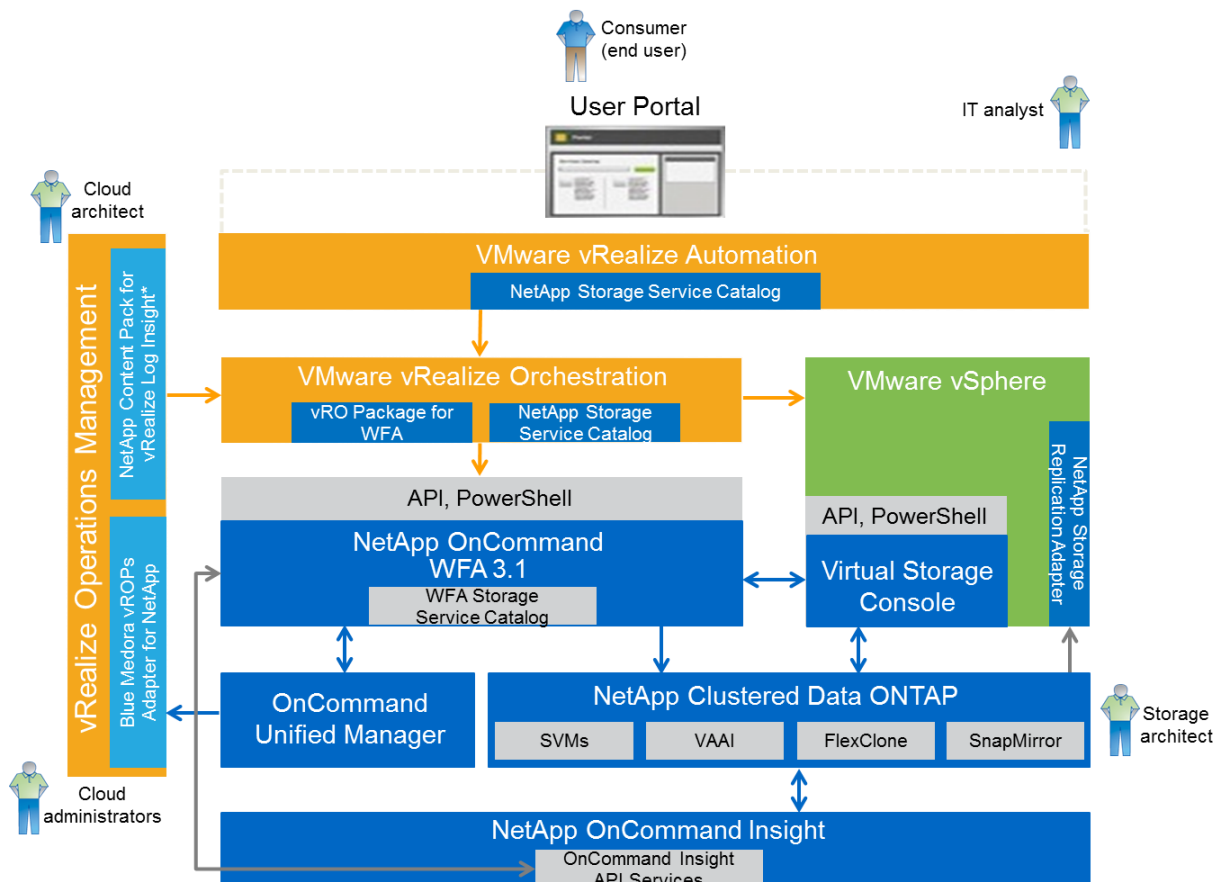
Terminology	Primary Technologies	Secondary Technologies	Definition
			<p>a VMware VM and acts as its own complete storage entity. This storage VM should have the capability to move from one storage controller to another, seamlessly and without downtime, allowing the possibility of continuous uptime in a cloud environment. This storage must also be integrated with the hypervisor layer to support the storage and VM creation, performance, and capability enhancement included with many hypervisors.</p> <p>Finally, the storage must have the ability to clone VMs at a rapid pace and then deploy those VMs to the hypervisor. At the very least it should support the VM cloning process that the hypervisor possesses. Each of the preceding components should have the capability to be orchestrated and in turn be placed into an automation framework.</p>
Storage automation and orchestration	OnCommand WFA	vRealize Orchestration, vRealize Automation	<p>This is the automation and orchestration framework that resides close to the storage and allows the creation, updating, deletion, and reading of the storage as needed.</p> <p>This framework allows storage administrators to orchestrate or automate specific functions within their storage environment. It also allows higher-level orchestration and automation engines to access it in order to create storage for other processes associated with virtualization.</p>
Hypervisor	VMware vSphere	None	This is the hypervisor and cloud environment for the SDDC. With it, there should be some type of integration from the hypervisor to the storage and the storage automation platform.
Hypervisor storage console	NetApp VSC	None	This software—as either a standalone application or a plug-in—performs storage administration services for a virtualization administrator from inside the hypervisor management layer. Common tasks of this software are datastore creation, datastore deletion, and rapid provisioning of VMs.
Orchestration	vRealize Orchestration	OnCommand WFA	This is an orchestration layer that hypervisors use in order to create and update hypervisor objects such as VMs, networking, and VM cloning. This layer calls the SDS orchestration layer to create or alter storage in conjunction with other creation and update procedures.
Problem remediation tool	OnCommand Insight, VMware vRealize Operations	OnCommand Unified Manager	This problem remediation tool automates operations management using analytics and an integrated approach to performance, capacity, and configuration management. This enables IT organizations to get greater visibility and actionable intelligence in order to proactively maintain service levels, optimum resource usage, and configuration compliance in dynamic virtual and cloud environments. This tool works in conjunction with

Terminology	Primary Technologies	Secondary Technologies	Definition
			the orchestration layer.
Service delivery and lifecycle management	vRealize Automation		This is the automation layer that calls other orchestrators. At this level, all processes are managed, and a user portal is available to users to request the resources they need. Also at this level, lifecycle management of the objects being automated is performed. This acts as the master control level, and all other orchestration is driven from this level.

## 5.1 NetApp and VMware SDS Overview

For multihypervisor environments, environments needing hardware orchestration, or enterprise-level environments, a combination of vRealize Orchestrator, vRealize Automation, and WFA can be used to automate and orchestrate SDS in an SDDC. With this combination, the automation of vSphere and software-defined networking through vRealize Automation and vRealize Orchestrator is combined with the SDS automation provided by WFA to create an end-to-end deployment process through all levels of a data center. Using vRealize Automation, a NetApp storage service catalog can be created that uses a combination of vRealize Orchestrator and WFA to create true SDS for the SDDC. Figure 23 is a high-level diagram of the SDS architecture. Figure 23 directly correlates to the generic SDS architecture shown in Figure 4, replacing the generic terminology with the corresponding NetApp and VMware technologies.

Figure 23) NetApp and VMware SDS overview.



## 5.2 Roles and Functionality

The SDS architecture includes the following roles and functions:

- **Consumer.** The ultimate consumer of services that the self-service portal offers can be internal or external. The consumer can be a human interacting with the portal user interface or an automation client interacting with the portal's RESTful interfaces.  
For more information, see [vRealize Automation 6.2.3 Programming Guide](#).
- **Cloud architect.** The cloud architect designs the services offered in the self-service portal for the cloud with which customers interact. The architect creates and defines the WFA, vRealize Orchestrator, and vRealize Automation service flows, combining storage resources offered by OnCommand WFA and compute systems, applications, and networking resources located in vCloud Director. The cloud solutions architect develops WFA, vRealize Orchestrator, and vRealize Automation workflows to create integration between different components.
- **Cloud administrator.** A cloud administrator manages day-to-day services. The cloud administrator has access to the defined workflows for vRealize Orchestrator, vRealize Automation, and WFA and uses them as needed, according to defined policies, to meet specific functions within the cloud environment. The cloud administrator acts as the gatekeeper for all compute and network resources in the architecture.
- **Storage architect.** The storage architect manages the storage infrastructure and can directly access the WFA workflows to perform storage administration tasks according to defined policies. These workflows are accessible to vRealize Orchestrator and vRealize Automation as storage resource operations. They act as the gatekeeper for all storage components in the architecture.
- **IT analyst.** The IT analyst manages the economics of which application is deployed to which cloud, as well as chargeback, wasted resources, and compliance with data security rules.

## 6 NetApp SDS

### 6.1 NetApp SDS Overview

SDS is part of the SDDC, an evolving architecture and set of technologies designed to speed delivery of IT services to application owners. SDS enables greater operational efficiency, higher availability, and lower cost in both SAN and NAS environments.

SDS does for storage what virtualization did for servers, breaking down the physical barriers that bind data to specific hardware. NetApp clustered Data ONTAP SVMs provide the following advantages:

- IT infrastructure teams become more responsive.
- Application owners become more autonomous.
- Business decision makers become more flexible.

NetApp delivers SDS with clustered Data ONTAP software, the industry's most comprehensive application integration, and programmable APIs for seamless workflow automation. SDS capabilities are woven into the NetApp products, not added later as a separate management layer. NetApp enable SDS through:

- Application self-service for fast user access to storage services, applications, and workflows
- Fully virtualized storage servers for dynamic resource provisioning based on service levels
- Deployment across a wide range of multivendor hardware and cloud service providers for platform flexibility

NetApp delivers SDS benefits through the following products:

- **Clustered Data ONTAP operating system.** Extends the innovation of Data ONTAP.
  - For more information, see the [NetApp Data ONTAP Operating System website](#).

- **NetApp OnCommand products.** Suite of management software and management integration tools.
  - For more information, see the [OnCommand Management Software and Management Integration Tools website](#).
- **NetApp FAS series.** Fabric-attached data storage systems.
  - For more information, see the [NetApp Data Storage Systems website](#).
- **NetApp FlexArray® virtualization software.** Enables management of both NetApp and third-party storage arrays using clustered Data ONTAP interfaces.
  - For more information, see the [NetApp FlexArray Storage Virtualization Software website](#).

## 6.2 SVMs

An SVM is a secure, logical storage partition through which data is accessed in clustered Data ONTAP. It is the core technology in SDS. A cluster serves data through at least one and possibly multiple SVMs to the cloud. An SVM is a logical abstraction that represents a set of physical resources of the cluster. The SVM abstracts the cluster hardware and is not tied to specific physical hardware. A variety of cloud-enabled applications can be deployed to SVMs because an SVM can support multiple data protocols concurrently.

An SVM is a secure entity; therefore, it is only aware of the resources that have been assigned to it and has no knowledge of other SVMs and their respective resources. Each SVM operates as a separate and distinct entity with its own security domain, which enables secure, separate storage to be granted to individual tenants within a cloud.

The cluster administrator can quantify the performance capabilities allocated to each SVM by using quality of service (QoS) policy groups for different service levels and for different storage needs in the cloud.

Clustered Data ONTAP is highly scalable. Additional storage controllers and disks can easily be added to existing clusters to scale capacity and performance to meet changing demands. New disk, cache, and network resources can be made available to an SVM to create new data volumes or migrate existing workloads to these new resources to balance performance and to create nondisruptive operations.

This scalability also enables the SVM to be highly resilient. SVMs are no longer tied to the lifecycle of a given storage controller. As new hardware is introduced, SVM resources can nondisruptively be moved from old controllers to new controllers, allowing the old controllers to be retired from service while the SVM is still online and available to serve data. Therefore, downtime for storage controllers—for maintenance, technology refreshes, or software upgrades—can easily be managed, allowing nondisruptive operations for cloud tenants created inside an SVM.

## 6.3 NetApp Storage Technologies for Cloud Storage

### Clustered Storage Architecture

All clustering technologies use the following guiding principles:

- **Nondisruptive operation.** The key to the efficiency and stability offered by clustering technologies is nondisruptive operation, which allows business-critical applications to stay online during routine upgrades and maintenance, add-ons, and even complete hardware refreshes.
- **Virtualized access as the managed entity.** Direct interaction with the nodes that make up the cluster violates the clustering concept. During the initial configuration of the cluster, direct access to the nodes is a necessity; however, steady-state operations are abstracted from the nodes because the user interacts with the cluster as a single entity.
- **Data mobility and container transparency.** The end result of clustering, which is the nondisruptive collection of independent nodes working together and presented as one holistic solution, is the ability to freely move data within the boundary of the cluster.

- **Delegated management and ubiquitous access.** In large, complex clusters, the ability to delegate or segment features and functions into containers that can be acted upon independently of the cluster means that the workload can be isolated. In addition, the cluster must not set conditions under which its contents can be accessed.

## Storage and Snapshot Efficiency

NetApp data deduplication and Snapshot copy technologies help customers manage data resources with greater efficiency and protect data with no performance impact and minimal consumption of storage space.

### Block-Level Deduplication

NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. This process is transparent to the application and the user, and it can be enabled and disabled easily and quickly. In a vCloud environment, deduplication can provide significant space savings because each VM can be an identical copy of the operating system, applications, and patches. Deduplication is also ideal for user and personal (profile) data stored in CIFS home directories. Also, deduplication should be enabled for datastores filled with VMs that have been cloned by using VMware vSphere vStorage APIs for VAAI full-copy offload. Patches, user data, and other similar blocks can be added to VMs inside the datastore after initial deployment. Enabling deduplication on these datastores allows space savings to be realized.

However, not all data within a vCloud environment is ideal for deduplication. Whenever possible, transient data should be placed in a separate volume with no deduplication operations that run on it. The advantages and disadvantages should be evaluated before performing deduplication on any data in a vCloud environment.

For more information about NetApp deduplication, see [Clustered Data ONTAP 8.3: Logical Storage Management Guide](#).

### Snapshot Technology

NetApp Snapshot technology enables data protection with no performance impact and minimal consumption of storage space. In a VMware cloud environment, the ability to create point-in-time copies of volumes or individual VMs (files) gives an administrator the ability to back up and instantly restore virtualized environments within the cloud. This process can then be automated in conjunction with an SDS backup and recovery service model.

### FlexClone Technology

A NetApp FlexClone volume is a clone of a volume, LUN, or file that is a writable, point-in-time Snapshot copy of a parent volume, LUN, or file, respectively. This copy, like a Snapshot copy, has minimal storage needs (essentially zero storage cost), and, because little data is replicated, it is highly time efficient. This feature gives a cloud administrator the ability to rapidly create clones of VMs either individually or as a group for almost zero cost in terms of space and performance to the storage controller. In an SDS environment, the ability to rapidly deploy VMs, volumes, or LUNs allows an administrator or architect to meet the constantly evolving needs of the SDDC by deploying individual VMs or entire tenants of preconfigured VMs with minimal lead time.

### NetApp QoS

Storage QoS workload management (available in clustered Data ONTAP 8.2 and later) allows administrators to define service levels by creating policies that control the resources that can be consumed by individual files, volumes, LUNs, or entire SVMs in order to manage performance spikes and improve customer satisfaction. QoS gives administrators the ability to set not-to-exceed performance



capacity limits (defined in terms of a maximum value for input/output operations per second [IOPS] or MBps) on a group of files, volumes, or LUNs within an SVM or on the entire SVM. This capability enables enterprise IT administrators and service providers to consolidate many workloads or tenants on a cluster without fear that the most important workloads will suffer or that activity in one tenant partition will affect another. As a result, administrators can push the storage infrastructure to higher levels of utilization.

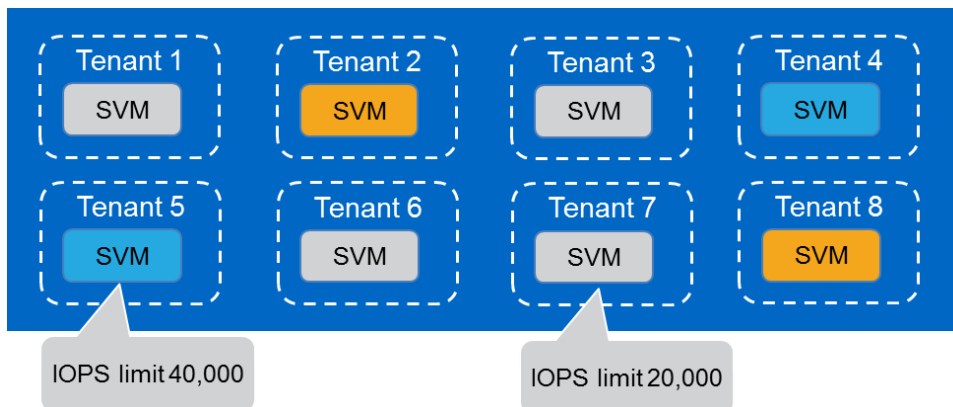
NetApp storage QoS works with both SAN and NAS storage, and it runs across the entire NetApp FAS product line, from entry to enterprise level. Administrators can incorporate third-party storage into their NetApp storage environment using NetApp FlexArray software. With FlexArray as a front end to SAN arrays from EMC, HP, HDS, and others, existing storage gets the full benefit of QoS limits and other NetApp capabilities. Storage QoS offers the following significant benefits for all types of storage environments:

- Achieve greater levels of consolidation.
- Set limits corresponding to SLAs.
- Add additional workloads with less risk of interference.
- Make sure that customers get what they pay for, but not more.

In a multitenant cloud environment, whether private or public, the first tenants on a particular resource might see a level of performance in excess of their SLA. This can create a perception of performance degradation as additional tenants are added and performance decreases. Storage QoS allows administrators to avoid this problem by assigning a performance capacity limit to each tenant in accordance with the tenant's SLA. Therefore, a tenant cannot exceed the set performance limit, even when resources are available, and the tenant is therefore less likely to notice changes in performance over time.

Another advantage of storage QoS is that it makes it simple to establish tiers of service based on SLAs. For instance, bronze service might correspond to a limit of 10,000 IOPS, silver to a limit of 20,000 IOPS, and gold to a limit of 40,000 IOPS. Figure 24 illustrates different SLAs for different tenants.

Figure 24) NetApp QoS with different SLAs for different tenants.



## 6.4 Accelerated Performance

The following sections summarize the integral components of the NetApp Virtual Storage Tier (VST)—NetApp Flash Cache™ and Flash Pool™ intelligent caching—which accelerates performance by providing higher speed and efficiency at lower costs.

### Flash Cache

Flash Cache is a hardware device that extends the native Data ONTAP VST capabilities. It improves read performance, increases throughput, and decreases latency. Flash Cache provides greater system

scalability by removing IOPS limitations caused by disk bottlenecks and lowers cost by providing equivalent performance with fewer disks. Leveraging Flash Cache in a dense (deduplicated) volume allows all of the shared blocks to be accessed directly from intelligent, faster Flash Cache instead of from disk. Flash Cache provides benefits in a VMware vCloud environment, especially during a boot storm, login storm, or virus storm, because only one copy of deduplicated data must be read from the disk (per volume). Each subsequent access of a shared block is read from Flash Cache and not from disk, which increases performance, decreases latency, and improves overall disk utilization. Flash Cache also increases the amount of available cache, which helps to reduce VM storm activities. For more information about Flash Cache technology, refer to [NetApp Flash Cache 2 Technical Specifications](#) on the NetApp website.

## Flash Pool

Flash Pool enables the addition of one or more RAID groups composed of solid-state drives (SSDs) to an aggregate that is otherwise composed of hard-disk drives (HDDs). The SSDs function as a high-performance cache for the working dataset, increasing the performance of the aggregate without incurring the expense of using SSDs for the entire aggregate. This hybrid-style aggregate combines the performance of SSDs with the capacity of serial-attached SCSI (SAS) drives, resulting in improved performance and lower costs.

## 7 SDS Orchestration and Automation Workflows

### 7.1 Where to Begin with Your First Workflow

Starting down the path of developing SDS workflows can be daunting. A top question is “Where do we begin?” Let’s take a look at common starting points for developing workflows.

The first place to look for developing workflows is business policies. These policies might already exist or might need to be created and then implemented: for example, a simple policy that all volumes created on storage must have a Snapshot policy set up for simple backup and recovery. This is a standard policy seen in storage environments and is an excellent example of an information technology organization’s storage setup being defined by a simple policy. This storage policy can be automated through a workflow that creates storage with a Snapshot policy with which it’s associated. The workflow (or script, depending on your implementation) then sets up the storage and always applies a Snapshot policy so that the business policy is always applied.

Another starting point for workflow development is repeated, routine storage tasks. Creating a volume is a basic task that all storage administrators must perform. Automating this task and including all functionality that could go with volume creation (for example, thin provisioning, deduplication, Snapshot policy, compression, autogrow, autoshrink, and others) make this repetitive, manual task into a “pushbutton” task. This task can then be given to junior-level administrators or pushed through as customer-initiated requests.

### 7.2 Benefits

Storage automation provides the following benefits to customers:

- When automation is used to create storage objects, all the desired properties are consistently set on the storage system each time. This brings efficiency to setting all the features that are needed for a storage object, such as deduplication, thin provisioning, and Snapshot copy retention policies.
- Automation of storage helps in the configuration of the storage features—such as the size and disk type of a volume or the IP address of a LIF—according to best practices and eliminates the human element, which can be error prone. Automation leads to reduced downtime, improved agility, and simplified operational processes.



Automating storage policies allows customers to leverage storage capabilities on demand, which ultimately allows the storage administrator to focus on design and architectural components and use less time on the operational aspects of operating a storage environment.

In order to achieve storage automation benefits, NetApp WFA is leveraged as a storage orchestration solution to automate and provision storage resources. With workflow creation, NetApp recommends using WFA to create simpler individual storage automation workflows for automating day-to-day tasks. Then these simple workflows are strung together at the vRealize Orchestrator level in a wrapper workflow to create one large, overarching workflow. This larger workflow can then be passed up to vRealize Automation from vRealize Orchestrator using an advanced services blueprint.

### 7.3 Architecture Capabilities

vRealize Automation, vRealize Orchestrator, and WFA combine to form a complete SDS architecture. The ability to automate and orchestrate the SDDC and SDS creates an agile and highly customizable data center at the push of a button. Correctly designed, this architecture can offload many manual processes from the administrator and allow the consumer of the SDDC to make the choices and deploy the environments on the back end, without the need for administrator intervention.

In today's fast-paced and constantly evolving world, organizations are looking to create their own private cloud, or at least use the cloud concepts in their data center. This NetApp and VMware joint architecture is a blueprint that allows SMBs and enterprise customers to easily manage and control their cloud environments.

At its core, this solution is designed to perform create, read, update, and delete operations on an as-needed basis for end users. When users, be they an end user, administrator, or architect, request a new environment, new storage, or an update to an additional environment or storage, the architecture has the capability to grant their request either automatically or after a change control process has been followed. The user is then granted the request, taking the storage administrator out of the day-to-day storage deployment equation.

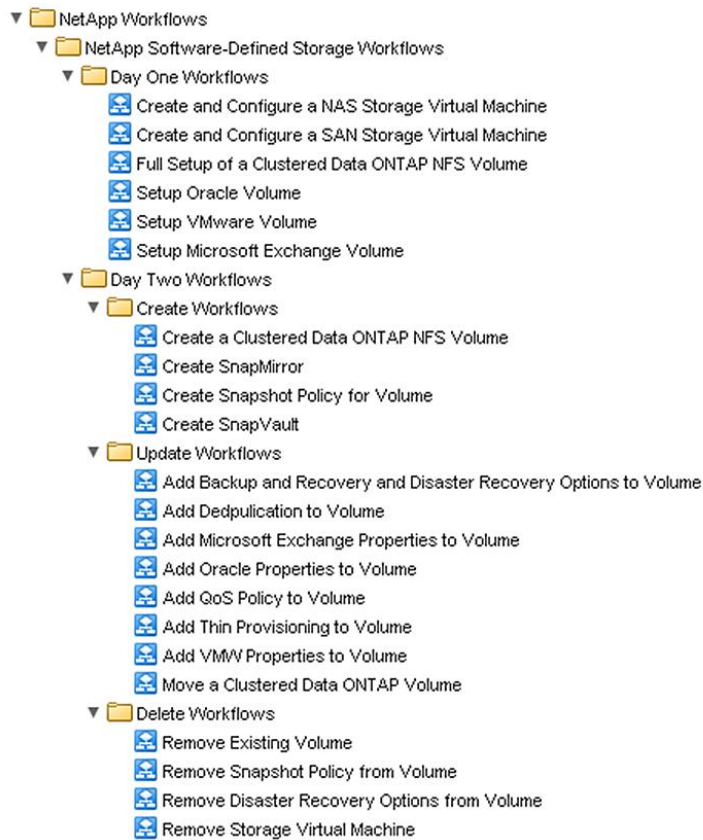
This solution is extensible to multiple implementation types and practices and has the following capabilities:

- Fast turnkey deployments
- Adherence to best practices
- Endless customization
- Low cost of management
- Securability, repeatability, and controllability
- Self-service storage management for tenants
- Automated provisioning, migration, and decommissioning of databases, applications, and virtualized infrastructure
- End-to-end orchestration processes

### 7.4 Understanding Workflow Concepts

Workflows can be used to accomplish multiple tasks within a storage automation and orchestration environment. As a best practice, NetApp recommends creating specific sets of workflows wrapped together to form a larger, more complete workflow. The workflow sets allow customization within a workflow environment, making the creation of more in-depth workflows possible. Figure 25 shows a series of workflows. The workflows in the Day One Workflows folder are large-scale workflows that are all crafted from workflows in the Day Two Workflows parent folder.

Figure 25) Series of workflows.



Using Figure 25, we will discuss the setup of VMware volume workflows. In this example, the end goal for the larger workflow is to have a datastore prepared for VMware with the following characteristics enabled:

- NFS
- 10GB
- Deduplication
- QoS of 13MBps
- Snapshot policy set and added to each host in the cluster

To accomplish this setup, the following day two workflows are wrapped together to form a larger workflow to create the overall workflow as follows:

1. Create a clustered Data ONTAP NFS volume. Creates a basic NFS volume of a size specified.
2. Add VMware properties to volume. Adds deduplication, thin provisioning, and QoS policy to a volume and makes it ready to become a VMware datastore.
3. Add backup, recovery, and DR options to folder. Adds NetApp Snapshot copy policies to the volume for the number of and time for the Snapshot copies to occur. It also adds SnapMirror replication and SnapVault® backup to the volume.
4. Add datastore on NFS volume. Standard workflow located in vRealize Orchestration under the Library > Storage folder. This adds the created volume and all of its properties to the vSphere environment.

Using the day two workflows, we created a day one workflow that creates a datastore with options specific to a VMware environment and sets up backup and recovery and DR as part of the process. When the datastore is set up this way, it is set up in the exact same way each time, removing human error from

the process. This process can be used to craft workflows for Oracle, Microsoft Exchange, or any other application layer.

Let's break down the day one workflows. Table 2 shows that each of the larger workflows (called wrapper workflows) is created by using smaller workflows from the Day Two Workflows folder.

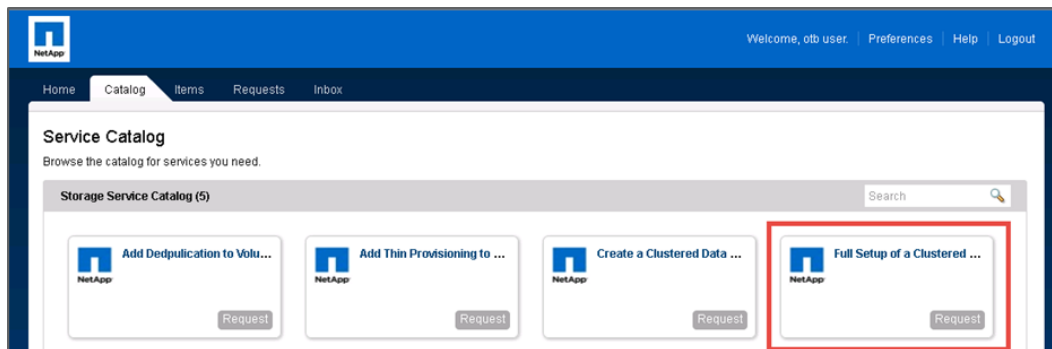
**Table 2) Workflow components.**

Workflow	Workflow for Volume	Workflow for Characteristics	Workflow for Backup and Recovery
Full setup of clustered Data ONTAP volume	Create a clustered Data ONTAP NFS volume	Add deduplication to volume, add QoS to volume, and add thin provisioning to volume	Add backup recovery and DR options to volume
Setup of VMware volume	Create a clustered Data ONTAP NFS volume	Add VMware properties to volume	Add backup recovery and DR options to volume
Setup of Oracle volume	Create a clustered Data ONTAP NFS volume	Add Oracle properties to volume	Add backup recovery and DR options to volume
Setup of Microsoft Exchange	Create a clustered Data ONTAP NFS volume	Add Microsoft Exchange properties to volume	Add backup recovery and DR options to volume

After you create the vRealize Orchestrator workflows, you can easily move them into vRealize Automation using advanced services blueprints. These workflows, after they are turned into blueprints in vRealize Automation, can then be made part of a storage service catalog.

Figure 26 shows the vRealize Automation blueprint created from vRealize Orchestrator workflows. This blueprint, when requested, calls both the vRealize Orchestrator and WFA workflows as needed to create a new volume.

**Figure 26) Storage service catalog.**



## 7.5 Security and Workflows

There are two specific things to think about concerning security in a vRealize Automation environment. First, when you create workflows that are moved into vRealize Automation (through advanced services), do not put credentials into the workflows. Instead, abstract credentials into the configurations of the environment. This creates a more secure environment. Second, be sure to use and document service accounts throughout the environment. This further heightens security.

## 8 Conclusion

With vRealize Automation, vRealize Orchestrator, and WFA, NetApp and VMware merged three technologies to form a single, cohesive orchestration and automation platform. If administrators follow the guides listed in the References section from start to finish, they can create an environment that handles the deployment of storage, applications, and virtualized infrastructure environments, including the VMs and storage required for each architecture.

Storage orchestration and automation are among the more difficult tasks to perform when creating environments and are often overlooked when architecting solutions. Using VMware and NetApp together to orchestrate and automate storage for virtualized environments provides the following benefits:

- When automation is used to create storage objects, all properties are set the same way each time. This process creates efficiency and provides the assurance that all features needed for a storage object—such as deduplication, thin provisioning, and a Snapshot retention policy—have been set.
- Automating storage establishes that all of the required features of a storage object—the size and disk type of a volume and the IP address of a LIF, for example—have been set. It also allows best practices to be adhered to. This process removes the human element, which can sometimes cause issues by users missing key settings, potentially resulting in downtime because of forgotten or overlooked storage object settings.
- Automating storage provisioning allows an organization to improve agility and time to market by having storage created without long, complicated processes (because those processes are automated) and without delay.
- Automating storage creation and deployment with best practices and having choices for storage properties, automation can be pushed to customers, making available storage-on-demand capabilities at the push of a button. Getting to this step takes time and effort, but, once there, day-to-day storage creation tasks are taken from the storage administrator and given to customers, allowing storage administrators to perform other tasks.

Getting to a state in which an SDDC and SDS environment is fully realized in a private cloud environment requires the understanding of a few key concepts:

- SDDC and SDS are about change. It is necessary to understand the roles and responsibilities in each private cloud environment. Furthermore, administrators must understand how those roles and responsibilities change after implementation of SDDC and SDS. There is a saying when discussing virtualization that “virtualization changes everything.” This is more so in an SDDC environment. Predicting and understanding the changes that occur facilitate a more successful deployment of an SDDC and SDS architecture.
- A successful implementation of SDDC and SDS takes planning. SDS and SDDC are a means to making an organization more agile and giving administrators the ability to rapidly meet the ever-changing demands of business. With this in mind, having the right tools and products that let the organization evolve into SDDC is key. VMware vRealize Automation and vRealize Orchestrator, coupled with NetApp WFA, allow businesses to gradually build their SDDC and SDS environments to fit their unique needs and schedules. In addition to having the right tools, make sure you have a plan that includes the required steps to achieve a true SDDC environment.

## References

### Supporting Documents

The following supporting documents were used for this document:

- Advanced Service Design: vRealize Automation 6.2  
<http://pubs.vmware.com/vra-62/topic/com.vmware.ICbase/PDF/vrealize-automation-62-advanced-service-design.pdf>

- Blue Medora vCenter Operations Management Pack for NetApp Storage Installation Guide  
<http://www.bluemedora.com/sites/default/files/resources/BM%20vCOps%20NetApp%20Storage%20Prod%20Sheet.pdf>
- Foundations and Concepts: vRealize Automation 6.2.3  
<http://pubs.vmware.com/vra-62/topic/com.vmware.ICbase/PDF/vrealize-automation-62-foundations-and-concepts.pdf>
- IaaS Configuration for Virtual Platforms: vRealize Automation 6.2.3  
<http://pubs.vmware.com/vra-62/topic/com.vmware.ICbase/PDF/vrealize-automation-62-iaas-configuration-for-virtual-platforms.pdf>
- Installation and Configuration: vRealize Automation 6.2.3  
<http://pubs.vmware.com/vra-62/topic/com.vmware.ICbase/PDF/vrealize-automation-62-installation-and-configuration.pdf>
- Installing the NetApp NFS Plug-In 1.0.21 for VMware VAAI  
[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP1389626](https://library.netapp.com/ecm/ecm_download_file/ECMP1389626)
- OnCommand Insight 7.0 Configuration and Administration Guide  
[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP1238056](https://library.netapp.com/ecm/ecm_download_file/ECMP1238056)
- OnCommand Unified Manager 6.3 Installation and Setup Guide  
[https://library.netapp.com/ecm/ecm\\_get\\_file/ECMP12469575](https://library.netapp.com/ecm/ecm_get_file/ECMP12469575)
- OnCommand Workflow Automation 3.1 Installation and Setup Guide  
[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP12470052](https://library.netapp.com/ecm/ecm_download_file/ECMP12470052)
- NetApp Storage Replication Adapter 2.1 for Clustered Data ONTAP Installation and Administration Guide  
[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP1388775](https://library.netapp.com/ecm/ecm_download_file/ECMP1388775)
- vCenter Site Recovery Manager 6.0 Site Recovery Manager Administration  
<http://pubs.vmware.com/srm-60/topic/com.vmware.ICbase/PDF/srm-admin-6-0.pdf>
- vCenter Site Recovery Manager 6.0 Site Recovery Manager Installation and Configuration  
<http://pubs.vmware.com/srm-60/topic/com.vmware.ICbase/PDF/srm-install-config-6-0.pdf>
- System Administration: vRealize Automation 6.2.1  
<http://pubs.vmware.com/vra-62/topic/com.vmware.ICbase/PDF/vrealize-automation-62-system-administration.pdf>
- vCenter Server and Host Management: ESXi 5.0 and vCenter Server 5.0  
<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-host-management-guide.pdf>
- Virtual Storage Console 6.1 for VMware vSphere Installation and Administration Guide  
[https://library.netapp.com/ecm/ecm\\_get\\_file/ECMP12412050](https://library.netapp.com/ecm/ecm_get_file/ECMP12412050)
- vRealize Operations Manager Customization and Administration Guide: vRealize Operations Manager 6.0.1  
<http://pubs.vmware.com/vrealizeoperationsmanager-6/topic/com.vmware.ICbase/PDF/vrealize-operations-manager-601-cust-admin-guide.pdf>

## Version History

Version	Date	Document Version History
Version 1.0	May 2014	Initial release
Version 2.0	March 2015	Added additional information
Version 3.0	January 2016	Support for WFA 3.1

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 © 1994–2016 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4308-0116