



White Paper

Enabling the Path to Private Cloud: Automation

Kristina Brand, NetApp
August 2014 | WP-7130

Abstract

A critical step in moving from a shared virtualized environment toward a private cloud is the transition from infrastructure optimization to service optimization. NetApp's design approach, technologies, and ecosystem integration build upon our storage efficiencies to deliver service-level efficiencies using centralized, policy-based management and automation. With NetApp® OnCommand® Workflow Automation, an efficient service-oriented approach can be implemented that enables storage provisioning, data protection, and operational processes to be integrated and automated. As a result, enterprises can meet workload service-level requirements while greatly reducing management costs and improving business agility.

TABLE OF CONTENTS

1	Private Cloud White Paper Series	3
2	Why Automation Is an Essential Element of the Cloud	4
2.1	Moving from Infrastructure Optimization to Service Optimization.....	4
2.2	How Automation Drives Service Efficiency	6
3	NetApp Approach Toward Automation	6
3.1	Delivering Storage and Service Efficiency	6
3.2	Automating and Integrating Storage Services.....	6
3.3	Centralized, Integrated Provisioning and Data Protection.....	15
4	Automation Use Cases	16
4.1	Use Case 1: Automatic Provisioning of Shared Storage with Data Protection	16
4.2	Use Case 2: Cloning in a Private Cloud.....	17
4.3	Use Case 3: Backup and Restore of Virtualized Infrastructures	17
4.4	Use Case 4: Multi-Tenancy.....	18
5	Ecosystem Integration	20
6	Automation: A Key Enabler for the Private Cloud	22

LIST OF TABLES

Table 1)	NetApp private cloud white papers.....	3
Table 2)	Workflow components in the Storage Service Catalog Package.....	9
Table 3)	The OnCommand management portfolio.	15

LIST OF FIGURES

Figure 1)	Organizations are evolving from silos to virtualization to cloud deployments.....	4
Figure 2)	Mapping technical attributes to service levels.	8
Figure 3)	Data ONTAP cluster credentials in WFA.....	10
Figure 4)	Storage service catalog.....	10
Figure 5)	Executing the workflow.	11
Figure 6)	Creating a storage domain (aggregates).....	11
Figure 7)	Creating a provisioning policy.	12
Figure 8)	Creating a schedule.	12
Figure 9)	Creating a local protection policy.	13
Figure 10)	Creating a remote protection policy (mirror).....	13
Figure 11)	Creating a storage service.	14
Figure 12)	Create automated backup of critical enterprise applications using OnCommand Workflow Automation and tools such as Snap Creator framework.....	18
Figure 13)	Provisioning SVMs for secure multi-tenancy.....	19
Figure 14)	NetApp open management interfaces.....	20

1 Private Cloud White Paper Series

Enterprise IT departments are under extreme pressure to reduce capital and operating expenses, driving them to virtualize infrastructures to improve hardware utilization and scalability and move toward the enhanced operational efficiency and flexibility of cloud computing. The cloud landscape includes private, public, and hybrid clouds. A private cloud is a shared virtualized infrastructure that remains within the control of the enterprise's IT organization behind a firewall. IT departments in a private cloud essentially take on the role of brokers of services in delivering applications, storage, and server resources to internal customers as services. A hybrid cloud is when a company uses a combination of private and public clouds.

Many organizations have virtualized portions of their infrastructures but are not sure how to navigate the next steps toward a fully automated, service-driven model that enables them to further reduce costs, improve efficiency, and deliver IT as a service. The transition to cloud computing is a multiyear and multiphase effort, and most enterprises are still in the early stages of data center transformation. This transformation involves a fundamental shift in focus from the infrastructure optimization provided by virtualization to the service optimization necessary for the cloud.

NetApp has helped many industry-leading firms deliver data and applications as an on-demand service delivery model built on clustered Data ONTAP®. This model can evolve from a private cloud to a hybrid cloud on a single platform. Based on NetApp's experience with countless IT environments, we have identified some fundamental elements that organizations of all types and sizes should include as they move to a private cloud. These elements are captured and explained in the NetApp private cloud white paper series shown in Table 1. These white papers explain how NetApp helps enterprises transition from a shared virtualized infrastructure to a private cloud. Each paper describes the design, deployment, and benefits of one of the key elements as it relates to a service-oriented infrastructure.

An important point: these papers are not focused on NetApp hardware. Instead, they explore the NetApp management software that enables policy-driven service efficiency as well as many advanced storage efficiency capabilities. They also describe how NetApp APIs integrate with third-party or customized orchestration solutions at each step, enabling organizations not only to deliver comprehensive storage management, but also to select other cloud resources tailored to their needs as part of their automated, end-to-end service fulfillment capabilities.

The NetApp private cloud white papers do not necessarily have to be read in sequence. In addition, some elements described in the documents overlap and can be deployed together.

Table 1) NetApp private cloud white papers.

Service Analytics	Optimize your services with centralized monitoring, metering, and chargeback to enhance visibility into both costs and service-level agreement (SLA) management.
Automation	Rapidly deploy new services by automating storage processes and integrating with third-party orchestrators for automated end-to-end service delivery.
Self-Service	Empower IT and users by allowing service requests to be fulfilled through a self-service portal.

2 Why Automation Is an Essential Element of the Cloud

Many organizations today find themselves in a struggle to stay competitive in the face of high data growth, infrastructure complexities, and shrinking IT budgets. They are looking for ways to reduce the overall cost of IT operations and increase business agility by addressing the following challenges:

- Inflexibility and poor responsiveness in meeting changing business demands, with the constant pressure to roll out systems more quickly and more cost-effectively
- Inefficient use of IT resources, including both underutilized hardware and overworked staff members who spend an inordinate amount of time on manual, repetitive, low-value activities
- High costs of management, including multiple tools and disparate service provisioning processes

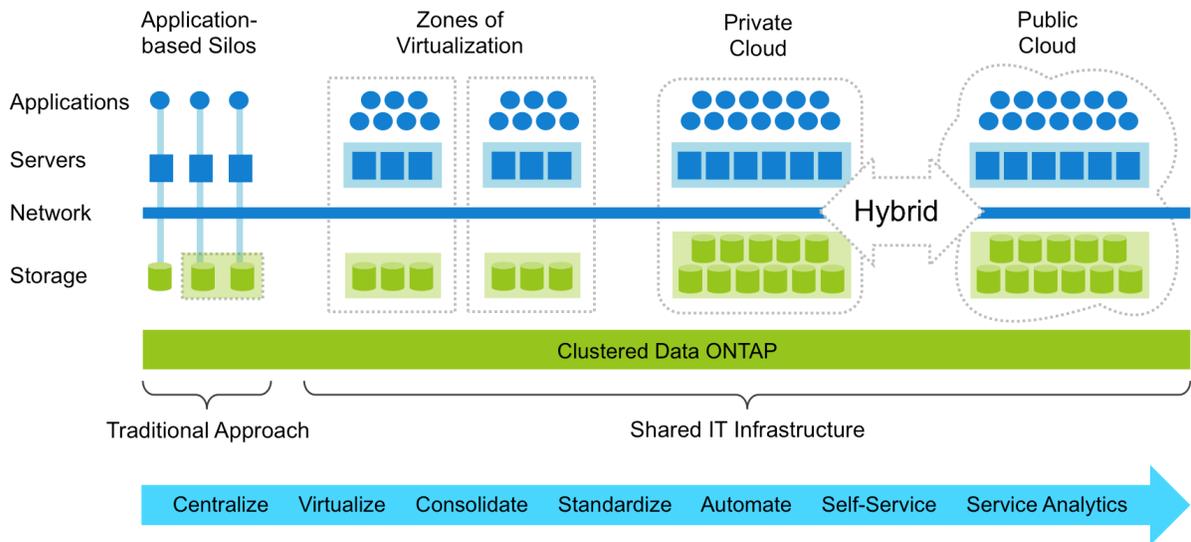
Cloud environments promise to deliver the solution through standardization, virtualization, and automation. The concept of delivering infrastructure as a utility service that can provide high resource utilization, multi-tenancy, and a pay-as-you-go model is very compelling to enterprises as they evolve toward an efficient private cloud environment. A recent ESG survey of IT decision makers revealed that approximately half of the respondent organizations already have a private cloud infrastructure in place, while an additional 20% plan to deploy one in the future.*

Automation is a key component of the cloud infrastructure, providing end users with almost immediate access to hardware resources and enabling enterprises to dynamically scale their services in accordance with client demand. Automation can also be used to integrate multiple disparate technologies and processes together to provision a higher value service to internal customers using a self-service interface.

2.1 Moving from Infrastructure Optimization to Service Optimization

In a critical step toward cloud computing, the majority of enterprise IT departments either have fully embraced virtualization as a concept or have begun to virtualize large portions of their data centers. In doing so, they begin to transition from the inefficiency and inflexibility of application silos to the enhanced agility and resource utilization of a shared virtualized infrastructure based on pooled hardware resources (see Figure 1).

Figure 1) Organizations are evolving from silos to virtualization to cloud deployments.



*Trends in Private Cloud Infrastructure, Enterprise Strategy Group, April 2014.

IDC expects that by 2017, over 80% of storage revenue and capacity shipped for x86 platforms will go toward virtualized environments. However, virtualized environments, while highly dynamic and efficient at the hardware resource level, are also typically very complex. If the environment isn't managed efficiently to enable service optimization, the cost savings that result from infrastructure optimization can quickly begin to unravel.

Today, the amount that enterprises spend worldwide managing virtualized infrastructures is already significantly greater than the amount spent implementing the virtualized infrastructure, and this gap is rapidly increasing. IDC has also discovered that flexibility and responsiveness are even more important drivers than cost for businesses that plan to move toward a private cloud.

The continued consolidation of virtual machines (VMs) in enterprise data centers increases the burden on storage systems, which must accommodate data from an increasing number of VMs. Therefore it's not surprising to find that a large share of the management complexity and cost in virtualized environments is associated with storage. Storage management processes tend to be largely manual intensive and poorly integrated, driving down productivity and responsiveness and driving up administrative costs. Without a more efficient, automated means of delivering storage capabilities to end users, it becomes nearly impossible to implement a service-driven model.

Storage Management Challenges That Prevent a Service-Driven Model

Enterprises typically experience storage service inefficiencies in the following areas:

- **Provisioning.** Provisioning storage in virtualized environments is a cumbersome process, requiring time-consuming, manual tasks and custom scripts, which can result in inappropriately allocated storage and poor utilization. In addition, many storage management tools are not integrated with virtualization solutions, so they are unable to rapidly provision storage in coordination with virtual machines to meet fluid business demands. The combination of NAS and SAN storage within the virtualized environment also leads to multiple tiers of physical storage with varying attributes that further complicate the provisioning process.
- **Protection.** As the number of managed devices and volume of data grow, it becomes increasingly difficult to prevent, diagnose, and resolve failures. This is especially true in virtualized environments where storage systems protect data for large numbers of virtual machines that are often migrated across different hosts. The impact of operator errors, component failures, and unexpected outages on users and the overall business is also much greater in highly consolidated virtual infrastructures. Without the right policies in place, storage might be manually configured without the required levels of protection to meet SLAs and RPOs/RTOs.
- **Management integration:** In addition to vendor-specific storage management tools, virtualized environments typically rely on a variety of device managers, network management tools, and orchestration solutions. Unfortunately, these disparate tools are not always integrated. Often they do not function across all storage platforms or provide centralized cross-domain management. This adversely affects administrative productivity and responsiveness.

Note: Storage inefficiencies related to poor infrastructure visibility and analysis, reactive capacity decisions, cost, and SLA management are also important topics. These factors are addressed in detail in the service analytics white paper "Enabling the Path to Private Cloud: Service Analytics" (WP-7133), part of this private cloud white paper series.

Organizations that are moving toward full data center transformation and cloud computing have realized that beyond infrastructure optimization is the realm of service optimization: the delivery and consumption of IT as a service. A service model transforms a virtualized environment into a service-oriented infrastructure that leverages automation to address management inefficiencies and deliver IT as a service to end users. Service automation represents a major inflection point on the path to a private cloud.

2.2 How Automation Drives Service Efficiency

Enterprises that have moved to a private cloud understand that automation is a key infrastructure attribute. It enables them to provide consistent, on-demand computing and storage resources to a variety of applications and users using well-defined policies. Automated policy-based management not only reduces administrative overhead, but also helps establish consistent, predictable service and reduces wasted resources. A service-driven model also improves IT efficiencies in other ways. For example, if organizations are able to associate the resources that users consume with services, it becomes easier for them to assign a cost to those services and thereby make departments or business units accountable for using them.

3 NetApp Approach Toward Automation

As organizations move from shared, virtualized infrastructures to private clouds, they are measuring their success by the efficiency of their infrastructure and services. NetApp's objective is to help enterprises manage the complexity of this transition and gain higher levels of storage efficiencies and translate these storage efficiencies into a form of higher service efficiencies. Our automation technologies enable this process to occur simultaneously by putting NetApp's shared storage differentiators to use. Enterprises can leverage these advanced storage automation capabilities through NetApp product interfaces or through their existing orchestration solutions without requiring manual intervention or extensive knowledge of the underlying technologies.

3.1 Delivering Storage and Service Efficiency

NetApp provides a number of technologies that enhance *storage efficiencies*, including virtualization and thin provisioning of volumes; deduplication of redundant data; and space-efficient backup, cloning, and replication of data based on Snapshot™ technology, just to name a few. In fact, we even guarantee that you will use 50% less storage by deploying NetApp solutions in a virtualized environment:

www.netapp.com/us/solutions/infrastructure/virtualization/guarantee.html.

NetApp delivers *service efficiency* by simplifying the use storage efficiency technologies with policy-based automation to create storage services based on those technologies. Policy-based automation helps eliminate manual processes, reduce errors, and increase operational efficiency. In a storage context, automation sets the stage for faster provisioning of storage and virtual machines, better data protection, and reduced management complexity. Our integrated storage automation capabilities have enabled many companies to execute large-scale provisioning, cloning, and backup processes in minutes instead of days or weeks.

3.2 Automating and Integrating Storage Services

NetApp OnCommand Workflow Automation contains categories of standard storage services and over 50 workflows right out of the box. Organizations can automate and standardize common storage processes such as provisioning, cloning, migrating, and decommissioning storage. They can integrate with orchestrators and third-party IT systems to form a single service offering when provisioning storage. OnCommand Workflow Automation enables ease of integration with orchestrators to provide end-to-end automated service delivery. It employs REST APIs that can be used to invoke automated workflows from the data center orchestration software or an external self-service portal.

This ability to integrate and automate is especially important for organizations looking to deploy cloud services. OnCommand Workflow Automation can play a vital role, enabling integration between storage and orchestration layers for cloud service delivery. It is a flexible framework that allows users to build an automated solution based on your existing processes and integrate it with your existing infrastructure. You can also use building blocks to create custom workflows, so as business needs change, the organization can easily update them to meet new requirements. For example, if the company is deploying a private or a hybrid cloud, the existing provisioning workflows can be executed in these new environments.

Although the NetApp Unified Storage Architecture represents an ideal platform for standardizing storage *infrastructure*, OnCommand Workflow Automation enables the organization to standardize the *utilization* of the storage infrastructure for improved performance, availability, and compliance with company- or industry-wide requirements. To optimize service efficiency in a cloud environment, IT management needs to make sure that the storage operations are as stable and predictable as possible when resources are added or replaced by standardizing the services deployed and automating as many processes as possible.

It might be useful to think of the storage service catalog as a means of simplifying the transaction between the provider and consumer of storage services:

- The provider (for example, a storage architect or storage administrator) creates the storage services using NetApp technologies that are accessible using OnCommand Workflow Automation and packs such as the Storage Service Catalog Package.
- The storage requester (for example, an end user or application administrator) consumes these storage services and isn't really concerned with how the services are delivered. Although OnCommand Workflow Automation is used to automate the delivery of the service, the requester can consume the storage using the OnCommand Workflow Automation user interface, a self-service portal, a NetApp API, an orchestrator, or a partner-enabled user interface (which, in turn, requests storage using the NetApp API).

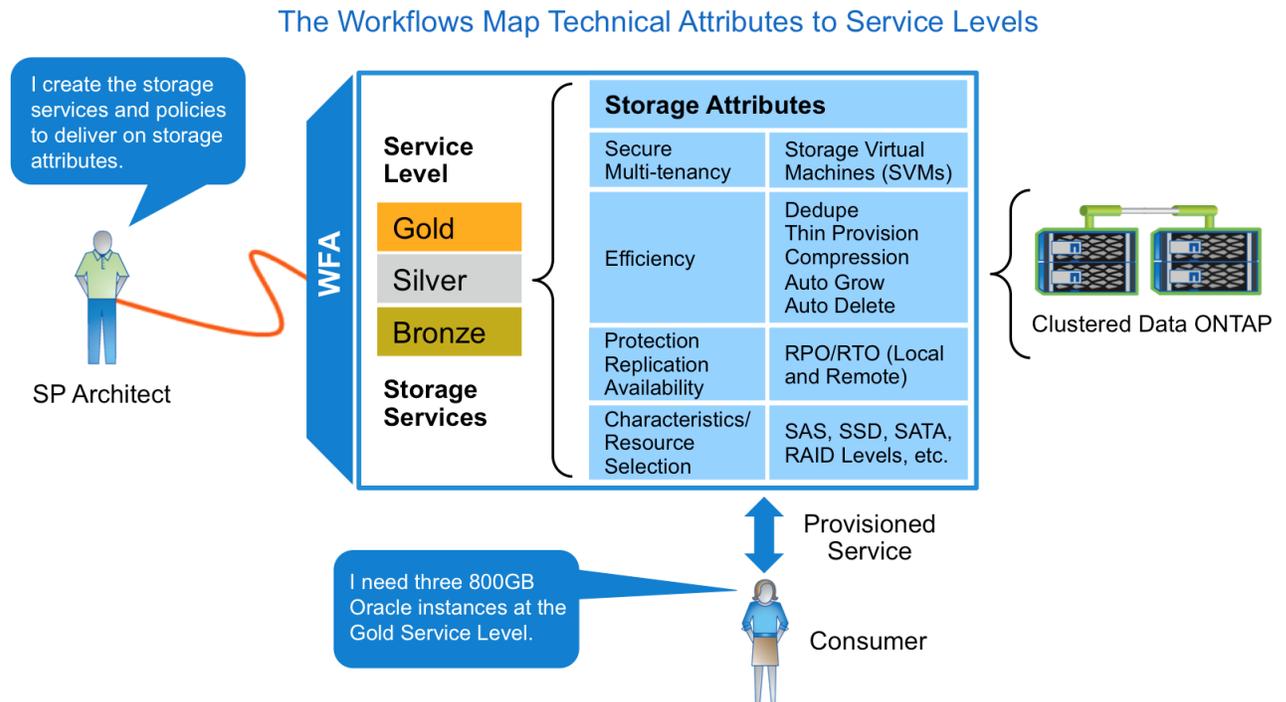
Service-Based Tiers Versus Physical Storage Tiers

Before discussing how the storage service catalog is created and used, it is important to understand NetApp's differentiated approach to service-level management. The traditional approach maps SLAs directly to static storage configurations, often designated as tier 1, tier 2, tier 3, and so on. The problem with this approach is that the service level requested is mapped to one or more storage systems that might not have all of the performance, data protection, availability, or capacity attributes necessary to fulfill the SLA over time. Conversely, the physical storage attributes might exceed the requirements of the SLA, leading to wasted high-end resources.

The service catalog instead maps SLAs to specific service levels, which are each associated with a set of storage attributes that determine performance, efficiency, and data protection capabilities, as shown in Figure 2. (The service levels in the diagram are designated as Gold, Silver, and Bronze as an example only. OnCommand Workflow Automation users can choose names that appropriately describe the service-level offerings for the specific environment.)

The underlying storage is managed as an optimized resource pool based on storage service catalog attributes. To further enhance utilization, the same resource pool can serve more than one service level. The policies created can thereby combine performance, provisioning, and protection attributes into a single workload-specific service level that leverages a dynamic pool of storage resources, rather than the static configurations of specific storage systems.

Figure 2) Mapping technical attributes to service levels.



Setting Up the Storage Service Catalog for Automation

The storage service catalog is composed of a set of defined service-level offerings (for example, SLA 1, SLA 2, SLA 3 or Gold, Silver, Bronze, and so on). Each of these service levels maps to a specific provisioning policy, a specific protection policy, a resource pool of one or more NetApp arrays, and an optional storage virtual machine (SVM) template (discussed in section 4).

There is a Storage Service Catalog Package for OnCommand Workflow Automation that contains table definitions, workflows, commands, filters, finders, and functions. All these objects can be viewed under the “Designer” tab of OnCommand Workflow Automation. The workflows in the Storage Service Catalog Package consist of provisioning and unprovisioning exports, LUNs, and shares, as well as workflows to create and remove the storage service catalog objects. The Storage Service Catalog Package is imported to OnCommand Workflow Automation to begin defining storage service policies according to SLAs. Table 2 lists the workflow components of the Storage Service Catalog Package and their uses.

Table 2) Workflow components in the Storage Service Catalog Package.

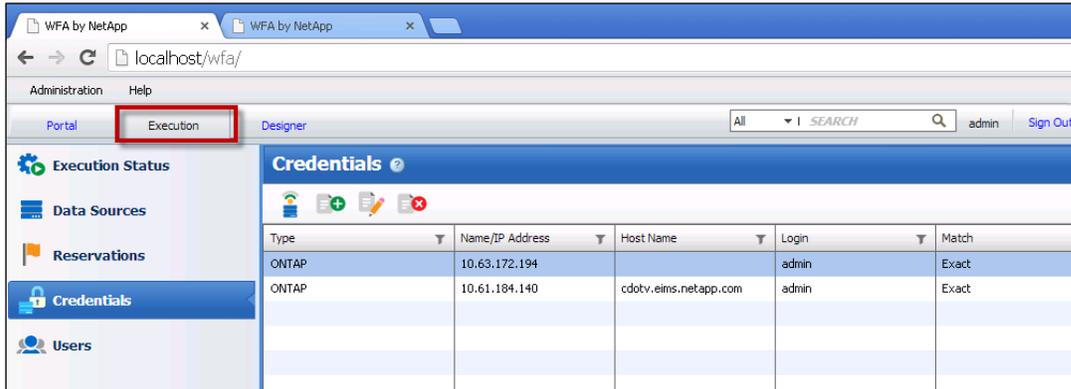
Workflow	Use
Create a set of predefined schedules	The set of predefined schedules can be used as examples to create new schedules or used just as they are
Create schedule	Creates a new schedule
Create storage domain	Groups aggregates for provisioning
Create provisioning policy	Groups storage object attributes and storage domains and creates a provisioning policy
Create local protection policy	Creates local protection attributes
Create remote protection policy	Creates remote protection attributes
Create storage service	Associates provisioning and protection policies to create a storage service
Create consumer	Creates a consumer and assigns storage services, clusters, and SVMs (primary and/or secondary)
View consumer objects	Views the consumer's associated objects
View provisioning policy	Views the provisioning policy's associated objects
View storage domains	Views the storage domains and associated members
View storage services	Views the storage service's associated objects
View storage objects by consumer	Views the storage objects that are associated to a consumer
Destroy schedule	Removes a schedule
Destroy storage domain	Removes the association of aggregates to a storage domain
Destroy provisioning policy	Remove a provisioning policy
Destroy local protection policy	Removes a local protection policy and its associated schedules
Destroy remote protection policy	Removes a remote protection policy and its associated schedules
Destroy storage service	Removes the storage service and the associated provisioning and protection policies
Destroy consumer	Removes the consumer
Provision export	Provisions an export
Provision LUN	Provisions one or more LUNs
Provision share	Provisions a share
Unprovision export	Removes an export
Unprovision LUN	Removes one or more LUNs starting with the LUN prefix
Unprovision share	Removes a share

Creating Automated Storage Services

The following is a step-by-step overview of how to create a new storage service, including provisioning, scheduling, and protection, using OnCommand Workflow Automation and the Storage Service Catalog Package.

1. Confirm that the Data ONTAP cluster credentials are configured in WFA, as shown in Figure 3.

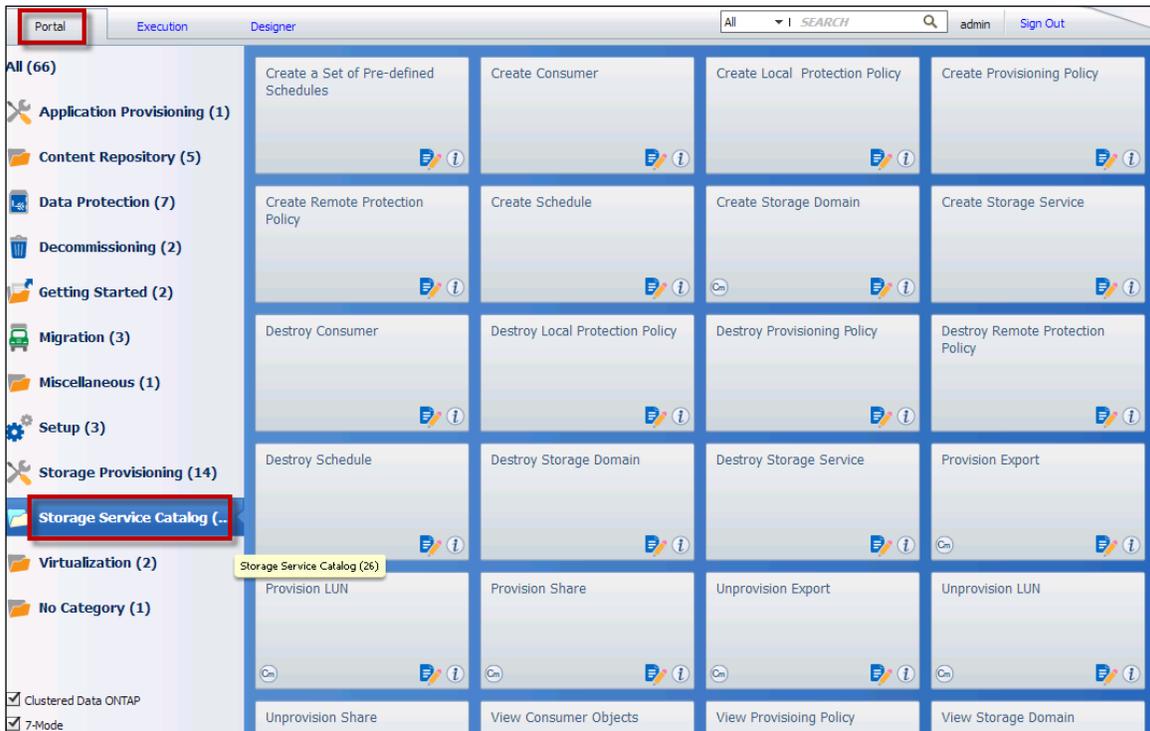
Figure 3) Data ONTAP cluster credentials in WFA.



Type	Name/IP Address	Host Name	Login	Match
ONTAP	10.63.172.194		admin	Exact
ONTAP	10.61.184.140	cdotv.eims.netapp.com	admin	Exact

2. The workflows can be filtered by clicking the storage service catalog category under the Portal tab to list all the custom workflows that have been imported.

Figure 4) Storage service catalog.



Create a Set of Pre-defined Schedules	Create Consumer	Create Local Protection Policy	Create Provisioning Policy
Create Remote Protection Policy	Create Schedule	Create Storage Domain	Create Storage Service
Destroy Consumer	Destroy Local Protection Policy	Destroy Provisioning Policy	Destroy Remote Protection Policy
Destroy Schedule	Destroy Storage Domain	Destroy Storage Service	Provision Export
Provision LUN	Provision Share	Unprovision Export	Unprovision LUN
Unprovision Share	View Consumer Objects	View Provisioning Policy	View Storage Domain

3. Execute the “Create a set of predefined schedules” workflow by clicking the workflow. This workflow creates a set of predefined schedules that should suit most protection needs. If a custom schedule is desired, it can be created using the “Create local protection policy” workflow.

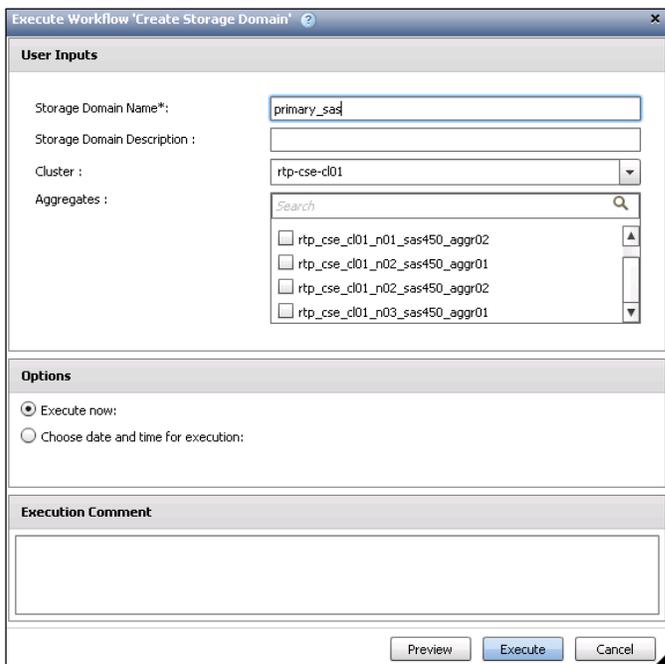
Figure 5) Executing the workflow.



Building Storage Service

1. The first step in building a storage service is to start with creation of storage domains. Execute the “Create Storage Domain” workflow.

Figure 6) Creating a storage domain (aggregates).



2. Create a provisioning policy and associate it with a storage domain. Storage efficiency parameters (thin provisioning, deduplication, compression, and so on) are specified here.

Figure 7) Creating a provisioning policy.

Note: The workflow also supports Data ONTAP Edge. If you want to create a storage domain with Data ONTAP Edge, select the RAID type to be Raid_0 in the drop-down menu for RAID type. One of the use cases would be to use Data ONTAP Edge as a mirror destination.

3. Create a schedule if the preexisting schedules do not match the requirements.

Figure 8) Creating a schedule.

4. Create a local protection policy and associate the desired schedule.

Figure 9) Creating a local protection policy.

The screenshot shows a dialog box titled "Execute Workflow 'Create Local Protection Policy'". It contains several sections:

- User Inputs:** Includes "Local Protection Policy*" with a dropdown menu set to "local", and "Local Protection Policy Description" with an empty dropdown menu.
- Schedule:** A section with a minus sign icon and "Schedule Name*" dropdown menu set to "Daily at 8 pm".
- Schedule Attributes:** A section with a minus sign icon containing "Retention Count*" (text input with "4"), "Prefix" (empty text input), and "Remote Protection Label" (empty text input).
- Options:** Contains two radio buttons: "Execute now:" (selected) and "Choose date and time for execution:".
- Execution Comment:** An empty text area.
- Buttons: "Preview", "Execute", and "Cancel" at the bottom right.

5. Create a remote protection policy that defines the replication characteristics.

Figure 10) Creating a remote protection policy (mirror).

The screenshot shows a dialog box titled "Execute Workflow 'Create Remote Protection Policy'". It contains several sections:

- User Inputs:** Includes "Remote Protection Policy*" (text input with "mirror every 2 hours"), "Remote Protection Policy Description" (empty text input), "Relationship Type*" (dropdown menu with "mirror"), "Transfer Priority" (dropdown menu with "normal"), "Schedule" (dropdown menu with "Every 2 hours"), "Tries" (text input with "8"), "Ignore Access Time" (checkbox), and "Restart" (dropdown menu with "default").
- Vault Rule:** A section with a plus sign icon and a "Vault Rule" label.
- Options:** Contains two radio buttons: "Execute now:" (selected) and "Choose date and time for execution:".
- Execution Comment:** An empty text area.
- Buttons: "Preview", "Execute", and "Cancel" at the bottom right.

6. Create a storage service and associate the previously created components to build the desired service. Associate the primary provisioning policy and the local protection policy to the storage service. If the storage service needs to have a secondary, then associate the secondary provisioning policy and the remote protection policy to the service as well. This will also determine what the primary and the secondary SVMs/clusters are when creating a “consumer” in the subsequent steps.

Figure 11) Creating a storage service.

Execute Workflow 'Create Storage Service' ?

User Inputs

Storage Service Name*: gold_primary

Storage Service Description : gold service with mirroring

Primary Information

Primary Provisioning Policy : aggr1node1

Local Protection Policy : local

Secondary Information

Secondary Provisioning Policy : cdotv

First Remote Protection Policy : remote

Options

Execute now:

Choose date and time for execution:

Execution Comment

Preview Execute Cancel

3.3 Centralized, Integrated Provisioning and Data Protection

How do you leverage NetApp storage and service efficiency to reduce management complexity while also improving storage provisioning and protection processes? The NetApp OnCommand portfolio includes a number of management software products designed to meet specific business needs such as simple device management, centralized monitoring and management, automation and integration, and enterprise storage resource management. The NetApp Unified Storage Architecture supports all NAS and SAN data, and all of our storage platforms across a shared storage infrastructure can be centrally managed from a single console. To simplify storage management, we provide flexibility and choices for automating the storage process for efficient storage service delivery.

Table 3) The OnCommand management portfolio.

System Manager	Simple Device Management
Unified Manager	Monitor and manage at scale
Performance Manager	Monitor performance
Workflow Automation	Automate and integrate storage processes
Balance	Optimize performance and prevent problems
Insight	Provides configuration and change management, capacity planning, enterprise reporting for chargeback, and showback
Snap Creator™ framework SnapDrive® SnapManager® SDK/API	Integrate management and data protection

OnCommand Workflow Automation facilitates a service-based architecture by capturing the key features and functionality of NetApp storage and associating them with services that can be delivered without complex scripts, spreadsheets, or other specialized understanding of NetApp technologies. There are numerous capabilities within OnCommand Workflow Automation, and many are beyond the scope of this paper. The OnCommand Workflow Automation provisioning and protection capabilities are essential to service automation. These integrated, automated provisioning and protection capabilities will be elaborated upon in the use cases described in section 4.

Open API Integration with the Broader Ecosystem

Another key NetApp advantage is our ability to integrate with many higher level management frameworks and third-party orchestrators. OnCommand offers plug-ins and other means of integrating with VMware® and Microsoft® virtualized infrastructures and can also be easily used within your existing orchestration solutions. By leveraging OnCommand Workflow Automation or the NetApp Manageability Software Development Kit (SDK), you can create a self-service portal and use third-party IT service management (ITSM) platforms, virtualization management solutions, or your own custom management tools to gain the same level of access to NetApp storage automation technologies and the storage service catalog as you can directly using NetApp OnCommand. The NetApp Manageability SDK will be discussed further in section 5.

4 Automation Use Cases

4.1 Use Case 1: Automatic Provisioning of Shared Storage with Data Protection

Business agility is a major driver for organizations moving to private clouds, and the ability to rapidly provision storage is an essential requirement. NetApp has taken a major step toward storage service efficiency by enabling resource pools, provisioning policies, and protection policies to be encapsulated to create an automated service. By easily selecting a service from OnCommand Workflow Automation and the storage service catalog offerings, administrators can significantly reduce the amount of time and energy they spend on configuration, provisioning, backup, recovery, and replication processes.

Simplify Storage Provisioning

OnCommand Workflow Automation automatically provisions NAS and SAN storage from resource pools established in the storage service catalog. You can set up repeatable processes with policies that consistently select resources for each provisioning activity. This eliminates the need for administrators to painstakingly search for available storage resources with the appropriate performance, availability, and efficiency attributes. For example, a high-end service level could be selected to automatically provision storage with a combination of attributes from the storage service catalog:

- **Fast performance.** NetApp FAS8000 series with all-flash FAS configuration
- **High efficiency.** Thin provisioning and deduplication for disk space optimization
- **Enhanced data protection.** RAID SyncMirror® (synchronized mirroring) and a combination of Snapshot, SnapVault®, and SnapMirror® technologies for fast recovery at local and remote locations
- **Secure isolation.** Provisioning storage virtual machines as part of a secure multi-tenancy environment

In contrast, you might use a midrange service level to provision storage for less critical data. This service might be associated with less expensive storage and slower disks and could exclude replication based on Snapshot.

Improve Storage and Service Efficiencies

Policies defined in OnCommand Workflow Automation and the storage service catalog can be used to automatically assign, add, and reclaim storage across the entire storage infrastructure. For example, a policy that includes thin provisioning and deduplication capabilities can greatly simplify capacity planning, optimize utilization, and improve service efficiency. Thin provisioning and deduplication can also be automatically applied to secondary or backup storage resources. These processes are much faster than manual administration, often requiring only minutes. They are easier to maintain than scripts and help minimize the risk of data loss or out-of-capacity conditions due to incorrectly configured storage. In the event that any policies don't conform to established service levels, OnCommand Workflow Automation can automatically detect and correct them.

SLA-Driven Data Protection

OnCommand automates end-to-end data protection activities, ranging from point-in-time copies for instant backups (Snapshot) to data replication for disaster recovery (SnapMirror). It applies consistent, SLA-driven data protection policies across all storage systems and automatically provisions secondary storage, enabling you to meet your RPOs and RTOs. For example, you can set up a policy that determines how often data is backed up, how long to retain backup copies (Snapshot copies), and when a backup copy should be replicated to another system for disaster recovery purposes.

4.2 Use Case 2: Cloning in a Private Cloud

Sophisticated IT professionals understand the importance of dynamically deploying private cloud resources to enable a flexible data center infrastructure: one that lets you anticipate and quickly respond to changing business needs. You might have investigated cloning, an increasingly effective and popular means of accelerating scalability and resource allocation in VMware environments. However, typical VM cloning methods can be expensive, inefficient, and slow, requiring full copies of VMs and time-consuming scripts to integrate cloning processes with storage solutions. NetApp offers a straightforward step-by-step method for cloning quickly and easily.

Virtual Appliance Cloning Example

The following is a description of the process for performing virtual appliance (vApp) cloning using OnCommand Workflow Automaton and VMware vCloud® Director. Begin by logging into vCloud Director and reviewing the vApp. Then select the vApp you want to clone. Next, open your browser and log onto OnCommand Workflow Automation. Then review the available workflows in the vCloud category listed and select the workflow for cloning a new vApp. Next, you will be shown a limited set of user inputs, including the source vApp name. Give the new vApp the name you would like to use and then input the host IP for vCloud Director and vCenter™. Then click execute to test execution for vApp cloning. After testing is passed, just click execute to process this workflow. The workflow will plan to see if it will succeed, and then it will initiate the execution step. Finally, switch over to vCloud Director and click refresh to confirm that the new vApp has been created.

4.3 Use Case 3: Backup and Restore of Virtualized Infrastructures

Virtualized environments introduce an entirely new set of data protection and business continuity requirements compared to physical infrastructures. Reduced physical server footprints provide higher efficiencies but often lead to a simultaneous increase in storage growth and complexity. Because a virtual server is primarily a set of storage files, the traditional boundaries between storage and servers have also become less clear. As a result, virtualized infrastructures are expanding the role of VI administrators, causing them to spend a greater amount of time on storage operations such as backup and restore.

Integration with VMware vCenter and Microsoft Hyper-V

To simplify and automate storage management activities performed by VI administrators, OnCommand Workflow Automation integrates with VMware vSphere® and Microsoft Hyper-V® virtualized infrastructures. VI administrators can access and apply automated data protection capabilities to their virtualized infrastructures directly from OnCommand Workflow Automation, without requiring storage administrator assistance. Our integrated workflows automatically group virtual machines into datasets and apply standard backup and replication processes to create consistent point-in-time copies of VMs, entire datastores (in VMware ESX® environments), or entire physical disks (in Hyper-V environments). They also enable fast restores at multiple levels of granularity: datastore/physical disk, VM, VMDK (ESX), or VHD (Hyper-V). Because the majority of these tasks are executed at the storage rather than the server layer, your server CPU cycles can be devoted to running applications, not backups.

Nondisruptive Backup

By leveraging OnCommand Workflow Automation policies, VMware and Hyper-V administrators can adopt a “set-and-forget” methodology:

1. Set a backup retention period based on time or number of copies based on Snapshot.
2. Replicate your data after every backup to make sure you are DR ready.
3. Schedule backups at the datastore or physical disk level so that all VMs provisioned with that datastore or physical disk are automatically protected.

NetApp Snapshot technology can be used while applications are running, enabling hundreds of VMs to be backed up almost instantaneously. And because backups are completed with almost no impact on host or network performance, they can be scheduled as frequently as needed. Because only incremental changes are stored, Snapshot copies are very space efficient, allowing more backups to be cost-effectively kept on disk for rapid VM recovery from up to hundreds of recovery points.

Your VI administrators can recover VMs within minutes by simply selecting the desired recovery point stored on disk. You can also streamline remote data replication for rapid disaster recovery. Replication from the primary system to your DR site can be automatically triggered immediately following a backup. In the event of a disaster, the Snapshot copy at your DR site can be rapidly promoted to a production copy.

These backup and recovery capabilities offered to VI administrators can also be made available to subscribers of these services using a self-service model. Enterprises can create a web-based self-service portal that provides access to the service levels in the storage service catalog and the protection policies described earlier through their higher level orchestration solutions. By integrating the orchestration framework with both the virtualization solution and the storage service catalog, users of the self-service portal are able to automate backup and recovery processes without being exposed to the underlying technology and workflow complexity.

Figure 12) Create automated backup of critical enterprise applications using OnCommand Workflow Automation and tools such as Snap Creator framework.

Execute Workflow 'Add Snap Creator Oracle Backup capability to Volumes'

User Inputs

Storage

Enter the Application Name*: orade_demo

Enter the Oracle SID*: eng1

Vserver Management Interface*: oracle01_mgmt

Snap Creator

Select the SnapCreator Server*: 192.168.0.5

Enable Agent? :

Agent

Backup

What policy type*: daily

How many snapshots?*: 1

Delete snapshots by age? : no

Preview Execute Cancel

4.4 Use Case 4: Multi-Tenancy

Multi-tenancy represents an ideal storage solution for private cloud deployment because it combines the flexibility and efficiencies of shared, virtualized storage with the security and dedicated resources of a physical array. It enables data for multiple tenants (applications, departments, or business units) to reside on a single storage system and simultaneously provides the ability to establish secure partitions and unique service levels for each tenant. Because each partition is treated as a separate virtual storage system, multi-tenancy can be used to improve both storage utilization and quality of service (QoS) management for each tenant residing on the array. OnCommand Workflow Automation lets you set up and automatically provision storage for multi-tenant environments by coordinating the service-level offerings of the storage service catalog with NetApp MultiStore® software capabilities.

NetApp MultiStore

Using NetApp MultiStore software, a single NetApp storage array can be partitioned into multiple SVMs that each appear as a storage system on the network. Each SVM uses assigned storage resources to deliver policy-based services to its tenant, just as a physical storage system does. Because each SVM is securely isolated from all other SVMs, one tenant cannot access storage resources and data from another. This would make sure, for example, that users from an organization's marketing department would not be able to access data from the HR department. Storage resources assigned to each tenant can also be dynamically resized to meet changing business demands without compromising security and SLAs.

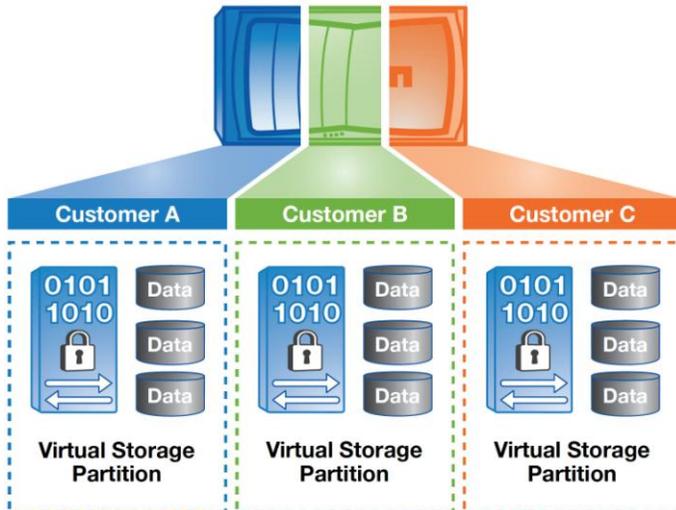
End-to-End Secure Multi-Tenancy

To apply the benefits of multi-tenancy to a cloud infrastructure stack, NetApp collaborated with VMware and Cisco to develop the industry's first secure multi-tenancy solution, enabling shared storage, server, and network resources to be securely partitioned and isolated for each tenant. With secure multi-tenancy, the security, flexibility, and service efficiencies provided for the storage layer are extended to all layers of the infrastructure, facilitating a dynamic shared environment that serves as a foundation for cloud services.

Provisioning SVMs

How does the provisioning of SVM fit into the context of cloud computing? The storage service catalog provides a template for the creation and provisioning of virtual NetApp arrays (SVMs) for systems partitioned with NetApp MultiStore and enables service-level selection for each SVM. For example, you can support several applications on a single NetApp array by provisioning an SVM for each application. Each SVM would be provisioned with the corresponding service level that provides the unique storage attributes, performance, and data protection needs of each application (SVM tenant). You can therefore optimize array utilization by partitioning physical resources *and* optimize service efficiency by automatically delivering unique service levels for each tenant.

Figure 13) Provisioning SVMs for secure multi-tenancy.



5 Ecosystem Integration

NetApp has adopted a differentiated strategy when it comes to the management of a private cloud infrastructure. Storage is our area of concentration and expertise. We help you optimize storage resource and service efficiencies and integrate with best-in-class partners that can provide end-to-end IT service management.

Flexibility to Choose the Right Solution

We integrate with partners that fall into these broad categories:

- ITSM or orchestration platforms
- Virtualization management solutions provided by VMware, Microsoft, and Citrix

Integration with Orchestration and ITSM Platforms

Large enterprises and service providers with complex IT environments are challenged to manage multiple moving parts from different domains and to provide the required service levels as prescribed by the various businesses they support. Perhaps you currently use or are considering IT service management or orchestration platforms from vendors such as BMC, CA, DynamicOps, Fujitsu, Gale Technologies, IBM, or NewScale. These management platforms consolidate the management of various data center elements and give enterprises and service providers the ability to orchestrate their entire infrastructure operations from a single management console.

Many storage vendors provide access to third-party orchestration platforms. However, they also typically offer their own infrastructure management solutions, which might not provide the same level of IT service management. We've taken a different approach. Instead of creating our own infrastructure management platform, we tightly integrate with these higher level management platforms. Our open APIs offer third-party management solutions or your own customized tools the same level of access to NetApp technologies as our own management tools.

Your third-party orchestration and ITSM platforms can integrate with NetApp technologies at a much higher level of abstraction, so you can deliver and enforce policy-based automation tasks using your existing tools. You can also use your orchestration platforms to provide detailed tenant and infrastructure workflows and end-to-end management, monitoring, and troubleshooting using self-service portals.

Figure 14) NetApp open management interfaces.



Open Access

Our open management solutions make it easier for you to provision new services across servers, storage, and network resources from a variety of (or combination of) management interfaces. By integrating with leading virtualization and orchestration solution vendors, you can empower your administrators to deliver end-to-end service automation, from service request to deployment, management, and deactivation. Integration with these tools facilitates automatic enforcement of storage provisioning, protection, and efficiency strategies, such as tiering or multi-tenancy, using your existing tools. On-demand provisioning lets you spend less time managing your storage environment and more time innovating.

Integration

The REST APIs in OnCommand Workflow Automation and the APIs in the NetApp Management SDK enable you to integrate with NetApp technologies at a higher storage abstraction layer. For example, you can use OnCommand Workflow Automation to quickly create a custom workflow for automated storage provisioning to be driven by your orchestration software without exposing the complexity of the underlying operations. As a result, you can consistently meet service levels and scale your environment with minimal effort as you add new cloud-based applications and services.

NetApp has taken extraordinary measures to make sure that our partners can easily leverage our APIs. We support all industry standard operating systems and offer web services, API documentation, solution guides, sample code, and developer tools for all NetApp technologies related to automated storage provisioning and data protection.

Enabling a Self-Service Model

We have discussed the automation capabilities provided by OnCommand Workflow Automation from the perspective of both storage and VI administrators. All of the use cases covered earlier can also be made available to service subscribers (end users) directly through a web-based self-service portal that provides direct access to the storage service catalog offerings. The subscriber can simply select one of the service-level offerings from a pull-down menu and enter the size of the storage share to be provisioned. The rest of the provisioning process would be automated in the same manner that it is when directly initiated through the OnCommand Workflow Automation operator portal. Creating a self-service environment in which your internal customers can request and receive appropriately configured IT resources with little or no intervention is an important step in cloud deployment. Details regarding the development of a self-service model for delivering cloud services are available in the self-service white paper “Enabling the Path to Private Cloud: Self-Service” (WP-7139), part of this NetApp private cloud white paper series.

6 Automation: A Key Enabler for the Private Cloud

NetApp has taken a holistic approach to cloud computing, where automation is one of several essential elements that are required to help organizations move from shared virtualized infrastructures to a private cloud and further evolve to a hybrid environment to fit business requirements. The role of automation in this process cannot be overemphasized because it enables organizations to move beyond the hardware efficiencies provided by infrastructure optimization toward the management efficiencies provided by service optimization. Storage automation is particularly critical to this transition because much of the administrative inefficiency in virtualized data centers is associated with complex storage management and exponential data growth.

In order to help enterprises address rapidly increasing capital and operating expenses while delivering resources on demand for improved business agility, NetApp delivers both advanced technology-based storage efficiency and policy-based service efficiency. We enhance service efficiency by encapsulating storage resources with provisioning and data protection policies into automated services that can be easily accessed and consumed. As reflected in the preceding use cases, storage service automation using OnCommand Workflow Automation and its Storage Service Catalog Package improves efficiencies across a variety of environments that are geared toward cloud computing, from rapid provisioning and cloning of virtualized infrastructures to secure multi-tenancy.

Clustered Data ONTAP and other NetApp storage technologies are designed with open APIs to extend these automated services beyond storage environments into the broader ecosystem using several methods, including self-service portals. As a result, organizations using virtualization solutions or higher level orchestration frameworks can integrate the benefits of automated storage services into their end-to-end infrastructure service delivery models.

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

© 2014 NetApp. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Data ONTAP, MultiStore, OnCommand, Snap Creator, SnapDrive, SnapManager, SnapMirror, Snapshot, SnapVault, and SyncMirror are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft and Hyper-V are registered trademarks of Microsoft Corporation. VMware, ESX, VMware vCloud, and VMware vSphere are registered trademarks and vCenter is a trademark of VMware, Inc. Oracle is a registered trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7130-0814