Technical Report

Secure Multi-Tenancy in Clustered Data ONTAP

Overview and Design Considerations

Julian Cates, NetApp
July 2013 | TR-4160

Abstract

This report discusses the implementation of secure multi-tenancy using Storage Virtual Machines (SVMs) in NetApp® clustered Data ONTAP®, covering design considerations and best practices. Throughout this paper, unless otherwise specified, the term “Data ONTAP” in isolation refers to clustered Data ONTAP. If a reference to 7-Mode or 7G is required, it will be specifically stated.

TABLE OF CONTENTS

1 Introduction 3

1.1 Purpose and Scope 3

1.2 Intended Audience 3

2 Overview 3

2.1 Secure Multi-Tenancy (SMT) 3

2.2 Storage Virtual Machines 4

3 SVM Design Considerations 6

3.1 SVM Layouts 6

3.2 SVM Networking 12

3.3 SVM Security 16

3.4 SVM Performance Monitoring and Isolation 23

3.5 Data Protection 25

3.6 Management Tools 28

LIST OF TABLES

Table 1) LIF maximums. 8

Table 2) Maximum NAS SVMs per cluster. 8

Table 3) Maximum SAN-enabled SVMs per cluster. 9

Table 4) Types of LIFs. 13

Table 5) Default cluster user roles. 17

Table 6) Default SVM user roles. 17

Table 7) SVM Language Recommendations 26

LIST OF FIGURES

Figure 1) Storage Virtual Machines. 4

Figure 2) Volumes junctioned into an SVM namespace. 11

Figure 3) Types of LIF configurations in clustered Data ONTAP. 12

Figure 4) LIF creation options as seen in OnCommand System Manager. 13

Figure 5) Single enterprise-wide IP network. 15

Figure 6) Multiple nonoverlapping IP networks within the same enterprise. 16

Table 7) Role-based access control definitions. 18

1 Introduction

1.1 Purpose and Scope

This technical report will describe the implementation of secure multi-tenancy using clustered Data ONTAP. It will describe common Storage Virtual Machine (SVM) deployment scenarios and will provide SVM best practice recommendations. This report is intended as a reference guide only and is not a replacement for product documentation, specific clustered Data ONTAP technical reports, or end-to-end clustered Data ONTAP operational recommendations. Where possible, these documents will be referenced.

1.2 Intended Audience

This document is intended for storage architects and storage administrators who want to understand secure multi-tenancy concepts in clustered Data ONTAP, different SVM deployment scenarios, and best practices. This document assumes that the reader has a fundamental knowledge of clustered Data ONTAP architecture, as covered in TR-3982.

2 Overview

2.1 Secure Multi-Tenancy (SMT)

What Is Secure Multi-Tenancy?

Traditionally, storage consumers who wanted to securely store their data did so by purchasing and deploying one or more physical storage arrays. The physical size of the array could vary based on the needs for capacity and throughput, but typically an entire array of some sort was required in order to provide the secure data and performance isolation that was required.

Secure multi-tenancy is the use of secure virtual partitions within a shared physical storage environment for the purpose of sharing the physical environment among multiple distinct tenants. For instance, a storage service provider might configure a storage array in such a way that each of three different customers is provisioned a certain portion of the array's disk capacity and network resources. In a secure multi-tenant environment, each customer would have access only to the resources explicitly provisioned to that customer. The customer would not have access to other customers' data or even be aware of the existence of the other customers or the fact that they share a common physical array. Secure multi-tenant environments should also provide a means to make sure that no single tenant consumes so much of the shared performance capability so as to affect the other tenants.

Benefits of Secure Multi-Tenancy

Traditional siloed models are inefficient. Deploying separate physical hardware stacks for each independent workload is costly and time consuming. Typically low resource utilization rates in these siloed environments translate directly to wasted resources. In contrast, multi-tenant environments enjoy the following benefits:

- **Reduced cost**. Dedicating a secure logical storage partition per tenant allows for economies of scale and is more cost effective than dedicated hardware.
- **Faster deployment time.** Logical partitions can be created in a fraction of the time required to rack, cable, install, and configure a separate physical array.
- **Improved resource utilization.** Sharing physical hardware increases resource utilization rates and eliminates the need to overprovision individual workloads in order to account for workload variability.
- **Secure isolation.** Secure multi-tenancy allows businesses to consolidate tenants onto shared resources, while being assured that tenants will not have access to resources not explicitly assigned

to them. Tenants sharing the same physical hardware can operate independently and with the expectation that no single tenant will consume resources unfairly. For example, production and dev/test can run on the same system without the risk of dev/test affecting production workloads.

SMT in Clustered Data ONTAP

Clustered Data ONTAP is an inherently multi-tenant storage operating system. Although Data ONTAP 7-Mode has traditionally contained optional features that enabled multi-tenancy, clustered Data ONTAP is architected in such a way that all data access is done through secure virtual storage partitions. It is possible to have a single partition that represents the resources of the entire cluster or multiple partitions that are assigned specific subsets of cluster resources. These secure virtual storage partitions are known as Storage Virtual Machines, or SVMs.

Figure 1) Storage Virtual Machines.

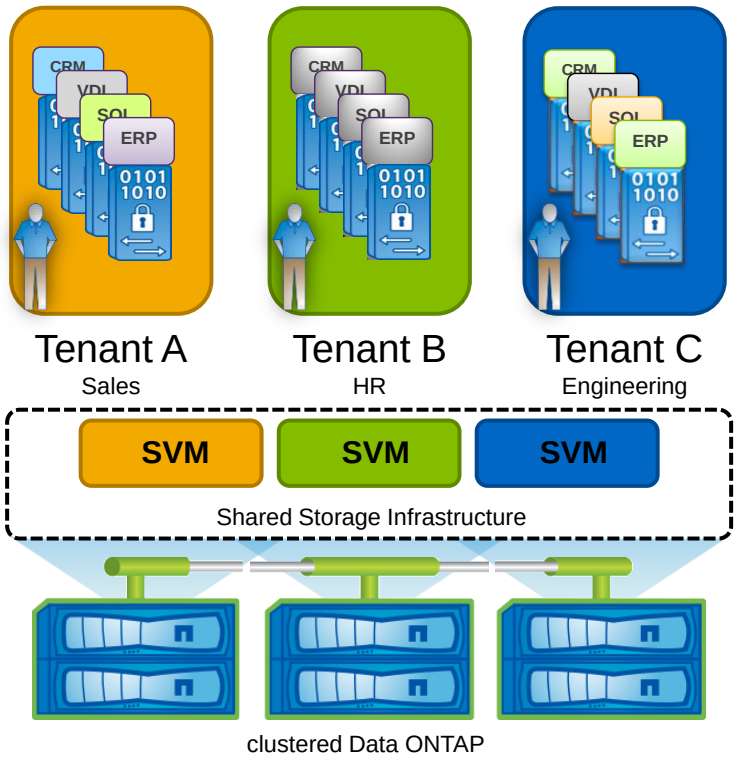

2.2 Storage Virtual Machines

Introduction to SVMs

The secure logical storage partition through which data is accessed in clustered Data ONTAP is known as a Storage Virtual Machine (SVM). A cluster serves data through at least one and possibly multiple SVMs. An SVM is a logical abstraction that represents a set of physical resources of the cluster. Data volumes and logical network interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved nondisruptively from one node to another. For example, a flexible volume may be nondisruptively moved to a new node and aggregate, or a data LIF could be transparently reassigned to a different physical network port. In this manner, the SVM abstracts the cluster hardware and is not tied to specific physical hardware.

An SVM is capable of supporting multiple data protocols concurrently. Volumes within the SVM can be junctioned together to form a single NAS namespace, which makes all of an SVM's data available through a single share or mount point to NFS and CIFS clients. To illustrate this example to the extreme, imagine a 24-node cluster licensed for UNIX® and Windows® file services with a single SVM configured with thousands of volumes and accessed from a single network interface on one of the nodes. SVMs also support block-based protocols, and LUNs can be created and exported using iSCSI, Fibre Channel, or Fibre Channel over Ethernet. Any or all of these data protocols may be configured for use within a given SVM.

Because it is a secure entity, an SVM 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. Tenants may manage the resources allocated to them through a delegated SVM administration account. Each SVM may connect to unique authentication zones such as Active Directory®, LDAP, or NIS.

An SVM is effectively isolated from other SVMs that share the same physical hardware.

From a performance perspective, maximums IOPS and throughput levels can be set per SVM using quality of service (QoS) policy groups. This allows the cluster administrator to quantify the performance capabilities allocated to each SVM.

Clustered Data ONTAP is highly scalable, and additional storage controllers and disks can be easily added to existing clusters in order to scale capacity and performance to meet rising demands. Because virtual storage servers within the cluster, SVMs are also highly scalable. As new nodes or aggregates are added to the cluster, the SVM can be nondisruptively configured to use them. In this way, new disk, cache, and network resources can be made available to the SVM to create new data volumes or migrate existing workloads to these new resources in order to balance performance.

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 to replace hardware that is to be retired, SVM resources can be nondisruptively moved from the old controllers to the new controllers. At this point the old controllers can be retired from service while the SVM is still online and available to serve data.

Note: Internal to Data ONTAP, an SVM is also referred to as a Vserver. In the command examples presented in this document, the `vserver` command is used to perform actions relating to SVMs.

Components of an SVM

Logical Interfaces

All SVM networking is done through logical interfaces (LIFs) that are created within the SVM. As logical constructs, LIFs are abstracted from the physical networking ports on which they reside. LIFs are described in further detail in section 3.2.

Flexible Volumes

A flexible volume is the basic unit of storage for an SVM. An SVM has a root volume and can have one or more data volumes. Data volumes can be created in any aggregate that has been delegated by the cluster administrator for use by the SVM. Depending on the data protocols used by the SVM, volumes can contain either LUNs for use with block protocols, files for use with NAS protocols, or both concurrently. For access using NAS protocols, the volume must be added to the SVM namespace through the creation of a client-visible directory called a junction.

Namespace

Each SVM has a distinct namespace through which all of the NAS data shared from that SVM can be accessed. This namespace can be thought of as a map to all of the junctioned volumes for the SVM, no matter on which node or aggregate they might physically reside. Volumes may be junctioned at the root of

the namespace or beneath other volumes that are part of the namespace hierarchy. Namespaces are discussed in section 3 and in further detail in TR-4129: Namespaces in Clustered Data ONTAP.

Infinite Volume

Instead of many flexible volumes, an SVM can be created that contains only a single volume that scales to multiple petabytes in size. SVMs with Infinite Volume can contain only this volume and have only one entry in their namespace, a top-level entry with a default value of /NS. No additional junctions can be created in an SVM with Infinite Volume. Rather, space is seamlessly added to the Infinite Volume as required. SVMs with Infinite Volume can span up to 10 nodes of a cluster. There can be multiple SVMs with Infinite Volume on a given cluster, and they can coexist on the same cluster that also contains SVMs with flexible volumes. For further information about Infinite Volume, refer to TR-4037: Introduction to NetApp Infinite Volume.

3 SVM Design Considerations

3.1 SVM Layouts

Single SVM Clusters

For many clusters, a single SVM that utilizes all of the physical resources available within the cluster could be the most logical and flexible option. This option is often the simplest to configure and maintain when there is a single tenant who will be consuming all of the available resources of the cluster. However, there are many instances, especially in multi-tenant environments, in which a cluster with multiple SVMs would be the most preferable. These use cases are discussed in the next section. A cluster initially configured with a single SVM can have additional SVMs defined as requirements or needs change.

Multiple SVM Clusters

Although it is possible to have only a single SVM for each cluster, there are also many scenarios where multiple SVMs would be required or offer advantages. Multiple SVMs could be created in any of the following scenarios.

Workload Separation

A benefit of creating an SVM for individual workloads on the cluster is that it enables the delegation of data management to the IT groups that are directly responsible for the data being stored. Application owners can be given the autonomy to control the datasets belonging to their specific application, while not having administrative rights to SVMs hosting other application workloads or the responsibility of overall cluster administration. When workloads are split off into isolated SVMs, QoS policies can be put into place to provide performance isolation at the SVM/workload level. For workloads with different data replication requirements, splitting into separate SVMs can be an effective way to logically group the volumes and LUNs that compose the workload.

NAS/SAN Separation

SVMs are, like clustered Data ONTAP itself, inherently multiprotocol and can serve NAS (CIFS, NFS) and SAN (iSCSI, FCP, FCoE) workloads concurrently. There is no technical requirement that NAS and SAN workloads run in different SVMs. However, there are a few reasons why separating NAS and SAN workloads into separate SVMs might be a good idea.

In many enterprises, there is a division of labor and responsibility when it comes to NAS and SAN. If there is a dedicated storage management team, it is common for that team to administer the storage hardware, including storage arrays and switches. The team also typically handles the creation and masking of LUNs along with creating the Fibre Channel zones required to securely assign LUNs to hosts. It is very likely

that this team would be responsible for overall management of the Data ONTAP cluster. It is a logical extension of the team's storage administration responsibilities to also manage the SVMs that will be providing block data services.

NAS, in contrast, tends to be very tightly coupled with the server teams that consume these storage services. In Windows environments, the server administrators have a vested interest in managing the CIFS shares. Likewise, UNIX administrators tend to be very familiar with managing NFS exports.

With this in mind, enterprises can create NAS SVMs and delegate administrative control to the respective server administrators, without granting them administrative access to the SAN SVMs.

File Server Consolidation

When consolidating multiple file servers or Data ONTAP 7-Mode systems into a single Data ONTAP cluster, there are some that will migrate easily into a single SVM, whereas others might require that additional SVMs be created. In general, if file servers to be consolidated share the same authentication services, belong to the same security zones, are managed by the same administrative team, and have no unresolvable overlap in share names, then a single SVM might suffice. Otherwise, creating additional SVMs will allow for file servers being transitioned to clustered Data ONTAP to be physically consolidated while maintaining the necessary secure logical separation. In general, the recommendation is to consolidate file servers into the smallest number of SVMs that is technically possible while still meeting authentication and performance requirements in order to streamline administration.

Infrastructure Service Providers

SVMs allow service providers to securely allocate storage resources to a tenant and delegate management of those resources without dedicating physical hardware to each tenant or exposing multiple tenants and their data to one another. Service providers can create tiers of service based on the types of cluster resources that will be made available to the tenant SVM, such as SSD storage, high-performance nodes with Flash Cache™, Gigabit Ethernet (GbE) vs. 10GbE interfaces, and so on. Volumes and LIFs can be nondisruptively reconfigured to use these resources, allowing service providers to maintain high availability for their customers. QoS also allows providers to control the data throughput allocated to each tenant. With QoS policies in place, tenants can share the same physical nodes of a cluster without one tenant consuming an unfair share of the node's resources.

SVM Limits

A minimum of one SVM is required to access data in clustered Data ONTAP; however, it is possible to create and use many more. The best practice for number of SVMs per cluster is arrived at by considering several factors along with the application and use case environments as discussed earlier:

- The number of nodes in the cluster
- The maximum number of IP and FCP LIFs per node
- The maximum number of IP, iSCSI, and FCP LIFs per port
- Keeping port capacity available to accommodate partner LIFs in the event of HA failover
- Whether SVM management will be done using a dedicated management LIF or a combined data/management LIF

Table 1 shows the per-node and per-port LIF maximums that must be taken into consideration. All of these factor into the maximum SVMs that are possible in a cluster, since each SVM requires a set of dedicated LIFs.

Table 1) LIF maximums.

Resource	Maximum Value
IP LIFs per node	256
FCP LIFs per node	512
IP LIFs per port	256
iSCSI LIFs per port	16
FCP LIFs per port	16

Note: An iSCSI LIF is defined as an IP LIF that has the iSCSI protocol enabled.

Note: Achieving the maximum iSCSI or FCP LIFs per node will require an adequate number of ports.

NAS SVMs

Table 2 shows the recommended number of NAS SVMs that can be created in clustered Data ONTAP 8.2. The choice to have separate management LIFs or to combine them with data LIFs will affect the number of SVMs it is possible to create. The following configuration options are recommended:

- **Combined data/management LIFs.** In this configuration each SVM requires one active IP LIF, which will be used for combined data and management access.
- **Dedicated management LIF.** In this configuration each SVM requires two IP LIFs: one active management LIF and one active data LIF.

Note: Care should be taken in an HA pair to make sure that enough LIFs are available on each node's partner so that HA failover is possible. Although the maximum IP LIFs per node is 256, for example, in practical terms this means 128 active LIFs and 128 that would only be instantiated in the case of failover.

Table 2) Maximum NAS SVMs per cluster.

Number of Nodes	Combined Data/Management LIF	Separate Data and Management LIFs
1	125	125
2	250	125
4	500	250
6	750	375
8	1,000	500
10–24	1,000	1,000

Note: It is possible to add more than the minimum required number of LIFs to an SVM. Doing so will reduce the maximum SVMs it is possible to create per cluster.

SAN-Enabled SVMs

Table 3 represents the recommended number of SAN-enabled SVMs that can be created in clustered Data ONTAP 8.2. The following configuration options are recommended:

- **FC/FCoE data LIFs with dedicated IP management LIFs.** In this configuration each SVM requires two FC/FCoE LIFs on each node of the cluster. For management, an IP LIF should be created and dedicated for management use.
- **iSCSI data LIF with dedicated IP management LIF.** In this configuration each SVM requires one iSCSI LIF on each node of the cluster, which will be used solely for iSCSI data traffic. It is recommended that the LIFs are created on top of interface groups, which combine multiple physical ports into a single virtual port for redundancy. An IP LIF is also needed for management per SVM.

Table 3) Maximum SAN-enabled SVMs per cluster.

Number of Nodes	FC/FCoE Data LIFs and Dedicated IP Management LIF	iSCSI Data LIF and Dedicated IP Management LIF
1	125	125
2	250	125
4	250	165
6	250	190
8	250	200

Note: It is possible to add more than the minimum required number of LIFs to an SVM. Doing so will reduce the maximum SVMs it is possible to create per cluster.

Mixed SAN and NAS SVMs

Clusters can contain a mixture of NAS-only and SAN-enabled SVMs. The exact number of SVMs that can exist in a single cluster will vary depending on how many of each type are created. The prime factor in determining exactly how many SVMs can be created in a cluster is the number of LIFs that the cluster nodes will support. In general, as long as the cluster has available capacity for the creation of required data and management LIFs while allowing for HA failover, an SVM can be created.

Tenant Tiering

Although SVMs have the potential to use any resource available within the cluster, cluster administrators also have the ability to control exactly to which resources—and which class of resources—a tenant would have access. This allows the cluster administrator to implement a tiering strategy whereby different business units, workloads, or customers could be assigned different classes of resources. A small cluster might have a small number of potential tiers; however, a large cluster with multiple controller and disk types can support many tiers.

Aggregates of various types can be created: SAS aggregates, SATA aggregates, SSD aggregates, and Flash Pool™ aggregates, for example. Tenant volumes can be provisioned in the appropriate aggregate based on requirements in place at the time of initial creation. If those needs or requirements change at a later time, cluster administrators can nondisruptively relocate the volumes to another aggregate of a different tier.

Aggregates can be located on nodes of differing capability as well, adding further potential to create a differentiation in tier. Workloads can be moved between nodes of differing memory and CPU potential, as well as differing amounts of flash-based cache.

For NAS workloads, logical interfaces can also be nondisruptively migrated to interfaces of differing capability. For instance, network access can be tiered by bandwidth (GbE vs. 10GbE) or degree of redundancy, such as a single physical port or an interface group composed of multiple physical interfaces.

SVM resources can be nondisruptively moved between tiers based on performance and capacity requirements or service-level agreements. Datasets with changing requirements throughout their lifecycle can be moved among tiers as required. For example, volumes can be created on high-performance tiers for intense data processing and later moved to cost-effective tiers for archival.

Note: Although volumes can be nondisruptively moved to new physical nodes and aggregates, they must remain in the same logical SVM. It is not possible to reassign a volume to a different SVM. If a volume must be migrated to a different SVM, a SnapMirror® data protection mirror can be used.

Language Considerations

Each SVM has a language setting that is specified at the time of SVM creation. This setting determines the default character set that will be used for the data stored in all volumes within the SVM. As of Data ONTAP 8.2, this SVM language setting can be changed at a later time if necessary. You cannot change the language for volumes that have been created within the SVM, however. Therefore, it is important to plan language choices ahead of time and be mindful to set them correctly.

One very important factor to consider when choosing an SVM language is whether or not the data will be replicated using SnapMirror data protection mirrors. When replicating volumes from one SVM to another, the source and destination volumes must have the same language setting. It is currently possible to replicate between SVMs whose language settings differ, as long as the volume language settings match. Where possible, it is advisable to choose an appropriate language setting based on technical requirements and standardize the use of this setting for every SVM and volume.

Best Practice
Unless there are technical requirements that mandate a specific SVM language, consider standardizing SVM deployments using the C.UTF-8 (POSIX with UTF-8) language.

SVM Namespace Considerations

Namespace

For NAS protocols, SVM data is accessed using a single hierarchical directory structure known as a namespace. Each SVM has a namespace that is separate and distinct from the namespaces of any other SVMs in the cluster. FlexVol® volumes created within a given SVM can only be mapped to that SVM's namespace.

Volumes are mapped into the namespace hierarchy through junctions. Each volume can have one junction path at a time, which designates where in the namespace the volume will be placed. A volume's junction path does not have to correspond with the volume name; however, doing so can make it easier to understand where in the namespace a particular volume is junctioned. Volumes can be junctioned at the root of the namespace or beneath another volume that is already junctioned. Figure 2 shows a sample namespace.

Figure 2) Volumes junctioned into an SVM namespace.

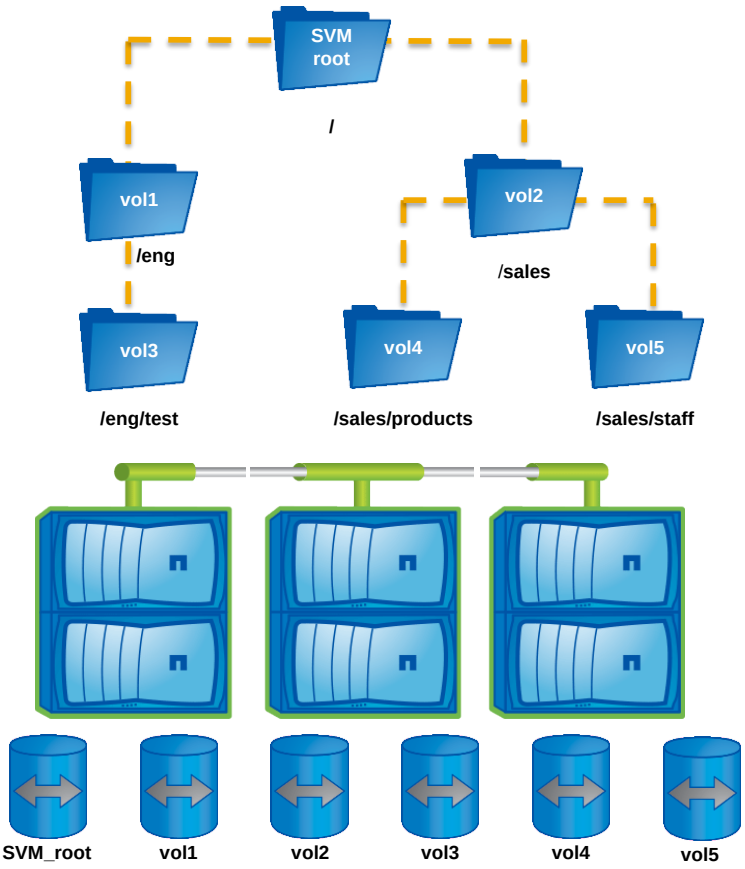

Note: Volumes are the only storage objects that can be junctioned into a namespace. Qtrees, volume subdirectories, and individual files cannot be. A volume can be junctioned directly off of another volume or off of a user-created subdirectory or qtree.

Export Policies

Access to each volume in the namespace is determined through the creation of an export policy. Export policies dictate which volumes are accessible from which hosts. A volume can have only one export policy, which applies to all data within the volume, including qtrees; however, the same export policy can be applied to many volumes. Each export policy can have multiple rules that specify a host or range of hosts and the type of access they are granted.

An example export policy is as follows:

```
             Policy          Rule    Access   Client                RO
Vserver      Name            Index   Protocol Match                 Rule
------------ --------------- ------  -------- --------------------- ---------
vs1          vsphere         1       any      192.168.1.1           any
vs1          vsphere         2       any      192.168.2.1           any
vs1          vsphere         3       any      192.168.3.0/24        any
```

Volumes that are assigned this export policy will be exported to the two hosts explicitly defined in the first two rules and to the entire subnet defined in the third rule.

Export policies are always created at the volume level. Although it is possible to have qtrees in clustered Data ONTAP, data in qtrees is exported by creating an export policy for the volume that contains the

qtree. A qtree inherits the export policy of its containing volume. For a FlexClone® volume, which is a space-efficient copy of an existing FlexVol volume, the export policy does not have to match the export policy of the parent FlexVol volume.

For a detailed discussion of namespaces and export policies, see TR-4129: Namespaces in Clustered Data ONTAP.

3.2 SVM Networking

Types of Network Objects in Clustered Data ONTAP

There are three types of network objects in clustered Data ONTAP: physical, virtual, and logical.

The first type is a simple physical port, such as a GbE port or 10GbE port.

The second type is a virtual network interface such as an interface group (ifgrp) or a VLAN. Interface groups aggregate multiple physical ports into a single virtual interface for redundancy. VLAN interfaces partition a single physical port or interface group into multiple isolated broadcast domains.

Lastly, logical interfaces (LIFs) are created as an abstraction on top of the physical or virtual interface layer. IP-based LIFs for NAS or iSCSI are assigned IP addresses, and FC-based LIFs are assigned WWPNs. LIFs are the only interface type that gets assigned directly to an SVM. Management LIFs and data LIFs used for NAS protocols are mobile and can be nondisruptively reassigned or failed over to a new physical port, ifgrp, or VLAN. Data LIFs used for SAN protocols do not require the same mobility and do not migrate or fail over to new physical ports. This is because MPIO and ALUA in clustered Data ONTAP automatically choose the most optimal port for host traffic. However, it is possible for SAN LIFs to be first brought offline and then reassigned to a new home port. This provides a means to accommodate new hardware being added to the cluster and the retirement of existing hardware. Note that bringing a SAN LIF offline does not mean that a host will lose access to mapped LUNs, as LUNs are exported through multiple SAN LIFs concurrently, and LUN pathing is managed through multipath I/O software on the host.

Figure 3) Types of LIF configurations in clustered Data ONTAP.

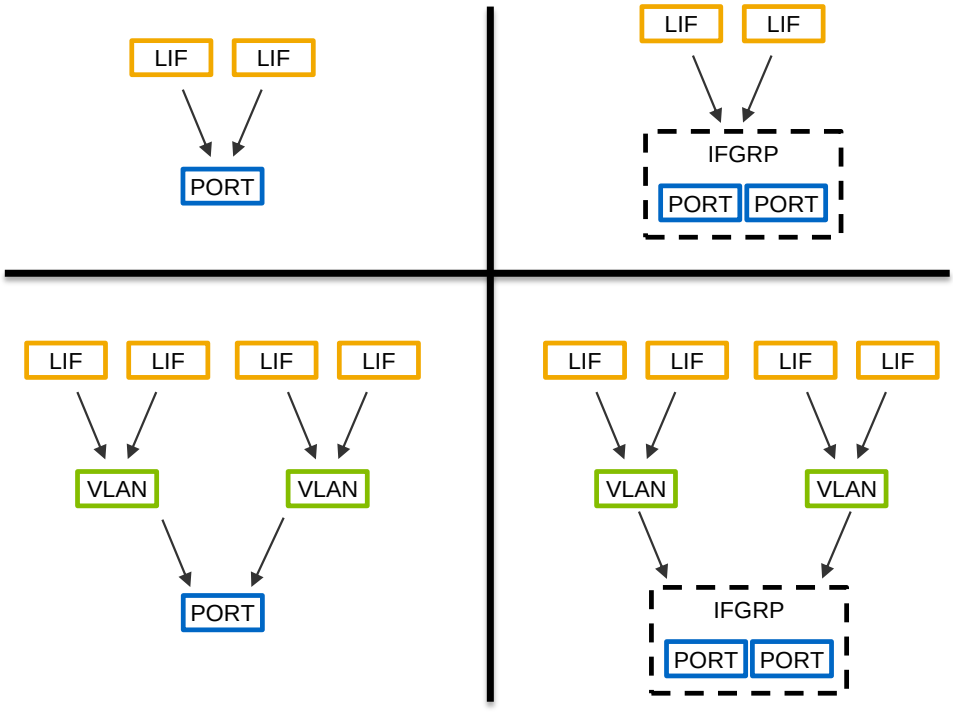

Logical Interfaces

There are several types of logical interfaces in Data ONTAP, and each is used for a distinct purpose. Table 4 lists the LIF types and their functions.

Table 4) Types of LIFs.

LIF Type	Function	Minimum Required	Maximum Allowed
Node management	Used for system maintenance of a specific node, SNMP, NTP, and ASUP™	1 per node	1 per port/subnet
Cluster management	Management interface for the entire cluster	1 per cluster	N/A
Cluster	Used for intracluster traffic	2 per node	2 per node
Data	Associated with an SVM and used for data protocols and protocol services (NIS, LDAP, AD, WINS, DNS)	1 per SVM	128 per node in HA config 256 per node in non-HA
Intercluster	Used for intercluster communication, such as setting up cluster peers and SnapMirror traffic	1 per node if cluster peering is enabled	N/A

Data LIFs can be used for data protocol access, SVM management access, or both, as shown in Figure 4.

Figure 4) LIF creation options as seen in OnCommand System Manager.

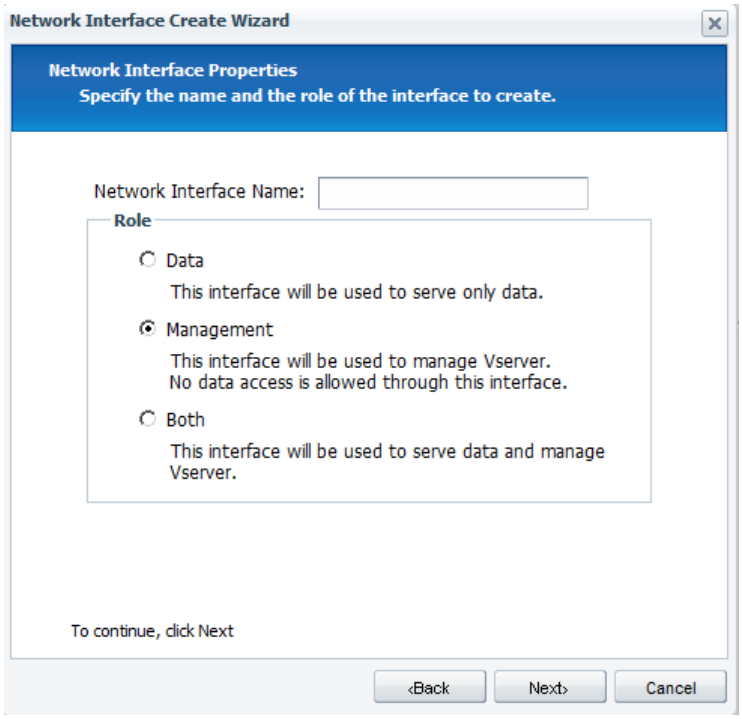

SVM Resource Placement

One of the key features of SVMs in clustered Data ONTAP is that each is a logical entity that exists on the cluster, not bound to any single controller or HA pair. Because of this, SVMs can contain resources from any node in the cluster and from multiple nodes concurrently. This empowers administrators with a great amount of flexibility. For example, data volumes for an SVM can reside on a single aggregate, or they can be distributed across multiple aggregates on multiple nodes. Using the data mobility features of Data ONTAP, these volumes can be relocated to different aggregates nondisruptively, even if the new aggregate is on a different node. Likewise, data LIFs are logical and can be moved nondisruptively to new physical ports, VLANs, or interface groups. These ports can theoretically be on any node of the cluster; however, care must be taken to make sure that the LIF gets moved to a physical port that is connected to an appropriate physical network. NAS clients can connect to shares or exports using an SVM's data LIF on any node and access all of the SVM's data volumes regardless of the nodes and aggregates in which those volumes are contained. This allows for unprecedented flexibility at the physical level to introduce new resources to the cluster, retire resources from the cluster, and balance workload and capacity across the cluster.

SAN LIF Considerations

Unlike LIFs used for NAS data access, LIFs that will be used for iSCSI, FCP, or FCoE access do not migrate from their assigned home physical port. Additionally, LIFs used for iSCSI cannot be used for any other protocol, such as NFS or CIFS, and must be dedicated for use with iSCSI. Consequently, a SAN LIF of the appropriate type is typically created on each node of the cluster.

It is expected that hosts will use multipath I/O drivers and determine the optimum LIFs to use for LUN access through the use of Asymmetric Logical Unit Access (ALUA).

For host operating systems that have upper limits on the number of paths per LUN, this could be a challenge in large clusters with many nodes and in smaller clusters with many available paths per node. To reduce the number of paths available to a host, consider using portsets. Portsets can limit the number of paths available per node or the number of nodes through which the LUN is available.

For further discussion of SAN best practices in clustered Data ONTAP, see TR-4080: Best Practices for Scalable SAN in Data ONTAP Cluster-Mode.

Best Practice

For SVMs using the iSCSI protocol, consider creating a LIF dedicated to SVM management with no data protocol access and placing that LIF in an appropriately configured failover group. Since iSCSI data LIFs do not migrate, this configuration promotes high availability for the management interface.

Network Isolation of Tenants

An important aspect of multi-tenancy is securing network traffic so that tenants can be securely isolated from one another. Although it might be desirable for SVMs in an enterprise context to share a common IP network, SVMs used by individual tenants within a shared infrastructure environment should not share IP networks and should remain separated. Both of these goals can be accomplished through the use of routing groups.

Routing Groups

Routing groups are used with SVMs to control outbound network traffic for the LIFs belonging to the SVM. Each routing group is a separate and distinct routing table. It is possible for an SVM to have more than one routing group, but routing groups are never shared between SVMs. Each LIF belonging to an SVM is associated with one and only one routing group. Multiple LIFs in the same SVM can share a common

routing group and must be on the same IP subnet. Routing groups provide secure, segregated traffic forwarding and SVM-scoped network administration and control.

Single IP Network

A common deployment scenario for the enterprise is to provide multi-tenancy at the workload, departmental, or administrative level. In this use case, a single enterprise-wide IP network is deployed. Here, routing groups scoped to a specific SVM provide the required control and separation for packet forwarding and network administration.

Figure 5) Single enterprise-wide IP network.

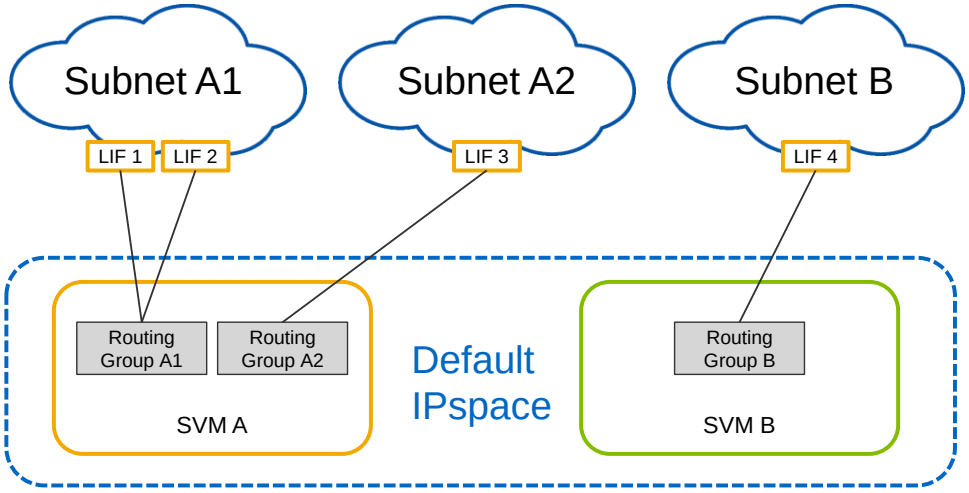

Multiple IP Networks

Another use case for secure multi-tenant network configuration is that of an enterprise or service provider whose tenants require segregated IP networks, each with a unique nonoverlapping IP address range. Examples of this could be an enterprise with test/dev SVMs on a separate network from production or that host data in a DMZ. In this instance, each SVM is confined to its own IP network, consisting of one or more subnets. The SVM LIFs would be created on top of a VLAN interface in order to securely scale the physical resources of the cluster and allow for many secure SVMs to be created. Because each routing group represents a distinct routing table, traffic is not routed between SVMs.

Note: The subnets depicted here could be connected to different networks entirely, such as a separate VLAN. It is possible through the use of multiple LIFs for the same SVM to be connected to multiple VLANs.

Figure 6) Multiple nonoverlapping IP networks within the same enterprise.

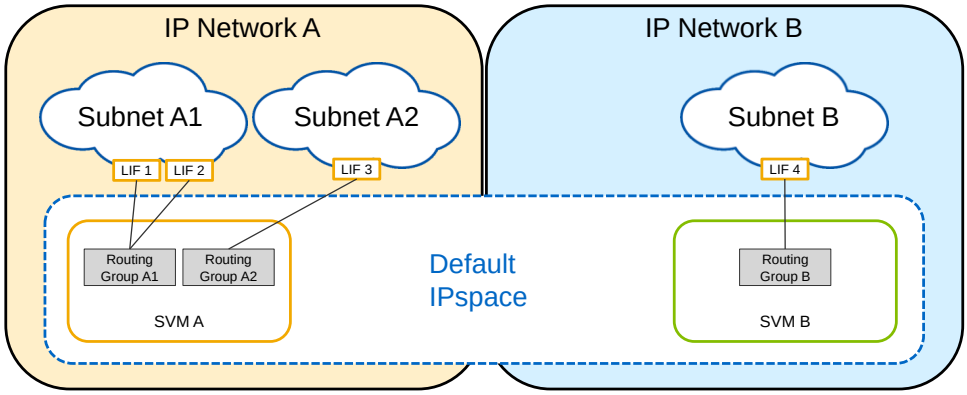

Configurations that require multiple IP networks and where identical IP address ranges are reused among tenants are not currently supported in clustered Data ONTAP.

Administrative Networking Control

Delegated SVM administrators using the predefined SVM administrative role have access only to a subset of the networking commands available to a cluster administrator. Only cluster administrators have the ability to move LIFs to different physical interfaces or VLANs. This is inherently secure from an SVM admin perspective, because the SVM admin does not have the potential to mistakenly reassign the LIF to an invalid port. It is also possible to create cluster administrator roles that would allow for administration of the cluster but restrict the ability to directly move LIFs or modify LIF failover and routing groups. The following is an example of such a custom role.

```
cluster::> security login role show -role NoNetworkAccess
  (security login role show)
           Role          Command/                                    Access
Vserver    Name          Directory                           Query   Level
---------- ------------- --------- ----------------------------------- --------
cluster  NoNetworkAccess
                         DEFAULT                                       all
cluster  NoNetworkAccess
                         network interface failover-groups             none
cluster  NoNetworkAccess
                         network interface migrate                     none
cluster  NoNetworkAccess
                         network interface migrate-all                 none
cluster  NoNetworkAccess
                         network routing-groups                        none
```

Roles are covered in further detail in the following section.

For an in-depth discussion of networking best practices, see TR-4847: Best Practices for Clustered Data ONTAP Network Configurations.

3.3 SVM Security

Users and Roles

Administrative Users

There are default administrative user accounts within clustered Data ONTAP as well as a robust means to create users with a customized set of privileges. The default cluster administrator is the `admin` user. Cluster administrators have the ability to administer the entire cluster and all of its resources. For SVMs, the default administrator is the `vsadmin` user. Although the `vsadmin` user is created with every SVM, it

must be explicitly enabled in order to delegate administration of the SVM. SVM administrators may only administer their respective SVMs.

User Roles

A role in clustered Data ONTAP is a collection of access control rules that specify what type of access a user will have to a given command directory, command subdirectory, or command.

There are predefined roles within Data ONTAP for both the cluster and SVM contexts.

Table 5) Default cluster user roles.

Role	Access Level	Capabilities
admin	All	All
readonly	Readonly	Read-only
none	None	None

Table 6) Default SVM user roles.

Role	Default Capabilities
vsadmin	• Manage owner user account local password and public key • Manage volumes, quotas, qtrees, Snapshot™ copies, FlexCache® files, and files • Manage LUNs • Manage data protection mirrors • Configure protocols • Configure services • Monitor jobs • Monitor network connections and network interface • Monitor the health of an SVM
vsadmin-volume	• Manage volumes, quotas, qtrees, Snapshot copies, FlexCache files, and files • Manage LUNs • Configure protocols • Configure services • Monitor network interfaces • Monitor the health of an SVM
vsadmin-protocol	• Configure protocols • Configure services • Manage LUNs • Monitor network interfaces • Monitor the health of an SVM
vsadmin-readonly	• Monitor the health of an SVM • Monitor network interface

Role	Default Capabilities
	• View volumes and LUNs • View services and protocols

Custom Users and Roles

In addition to the default users and roles, it is possible to create additional cluster and SVM users and to define custom roles that specify to what commands those users will have access.

Table 7) Role-based access control definitions.

Capability	Access	Description
Directory/subdirectory	All	Permits access to the directory and all subdirectories and commands contained within
	Readonly	Permits read-only access to the directory and subdirectories.
	None	Denies access to the directory and all subdirectories and commands contained within
Command	All	Permits execution of the command
	None	Denies execution of the command

Roles are defined in a hierarchical manner from general to specific. Rules defined for specific commands or subdirectories will override rules that are defined for their parent directory.

Creating Users with the CLI

To create a new user, use the `security login create` command. In the following example the user is created locally on the Data ONTAP cluster and logs in using a password using ssh. It is also possible to specify a domain or LDAP user or to specify the use of a public key in lieu of a password.

```
cluster::> security login create -username newuser -application ssh -authmethod password -role
customrole -vserver vs1

Please enter a password for user 'newuser':
Please enter it again:

cluster::> security login show -vserver vs1 -username newuser

Vserver: vs1
                             Authentication                  Acct
UserName         Application Method         Role Name        Locked
---------------- ----------- -------------- ---------------- ------
newuser          ssh         password       customrole       no
```

Creating Roles with the CLI

Roles are created in the CLI using the `security login role create` command. The role is created one rule at a time. The following series of commands creates the custom role discussed in the administrative networking control section earlier.

```
role create -vserver cluster -role NoNetworkAccess -cmddirname "network interface failover-
groups" -access none

role create -vserver cluster -role NoNetworkAccess -cmddirname "network interface migrate" -
access none

role create -vserver cluster -role NoNetworkAccess -cmddirname "network interface migrate-all" -
access none

role create -vserver cluster -role NoNetworkAccess -cmddirname "network routing-groups" -access
none

role modify -vserver cluster -role NoNetworkAccess -cmddirname DEFAULT -access all
```

Creating Users and Roles with OnCommand System Manager

To create and edit custom roles with OnCommand® System Manager, navigate to Configuration -> Security and select either Users or Roles.

Figure 7) Role Edit wizard in OnCommand System Manager.

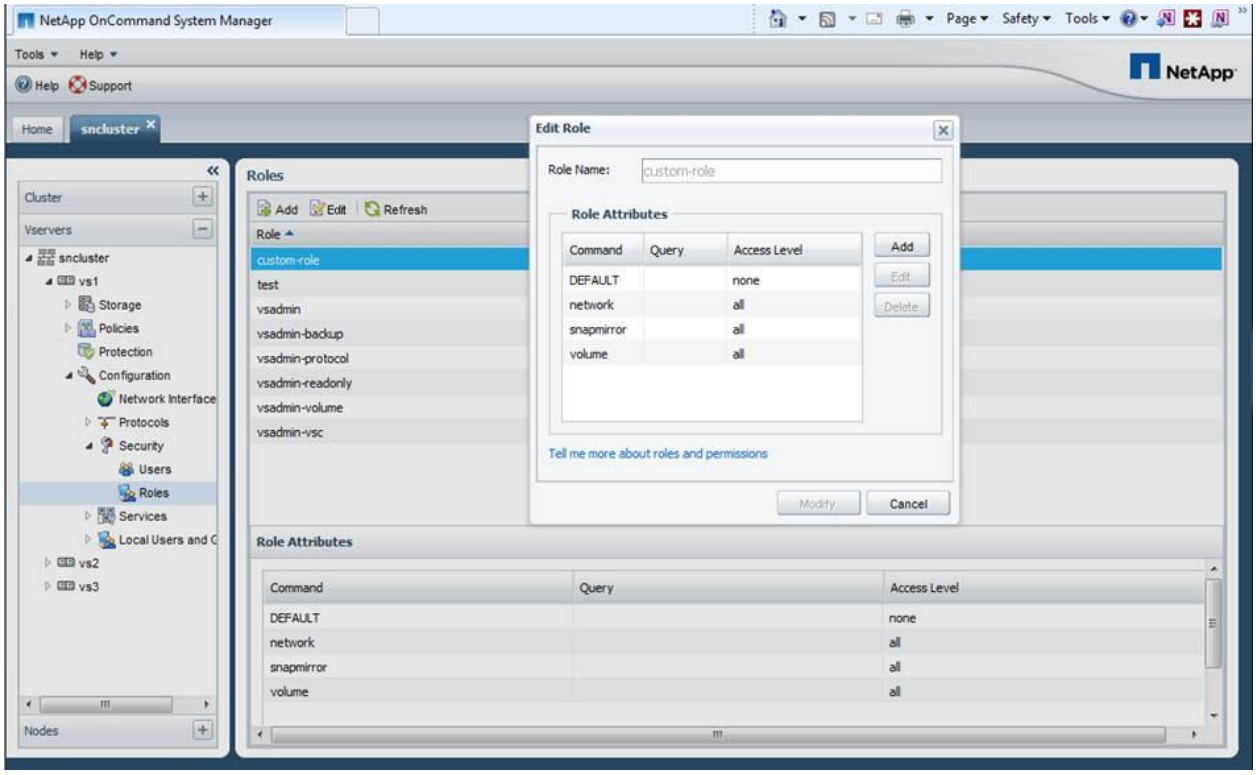

SVM Administrator Delegation

In addition to users who have administrative capabilities at the cluster level, it is possible for each SVM to have accounts enabled that have administrative rights only for that specific SVM. Each SVM has a `vsadmin` account created by default, but it must explicitly be enabled. To enable the `vsadmin` account, assign a password and then unlock the account.

```
cluster::> security login password -username vsadmin -vserver vs1

Enter a new password:
Enter it again:

cluster::> security login unlock -username vsadmin -vserver vs1

cluster::> security login show -username vsadmin -vserver vs1

Vserver: vs1
                             Authentication                  Acct
UserName         Application Method         Role Name        Locked
---------------- ----------- -------------- ---------------- ------
vsadmin          ontapi      password       vsadmin          no
vsadmin          ssh         password       vsadmin          no
2 entries were displayed.
```

It is also possible to enable the SVM administrator through the NetApp System Manager graphical user interface by choosing the Delegate Administration of this SVM option of the SVM setup wizard.

Figure 8) Delegating SVM administration.

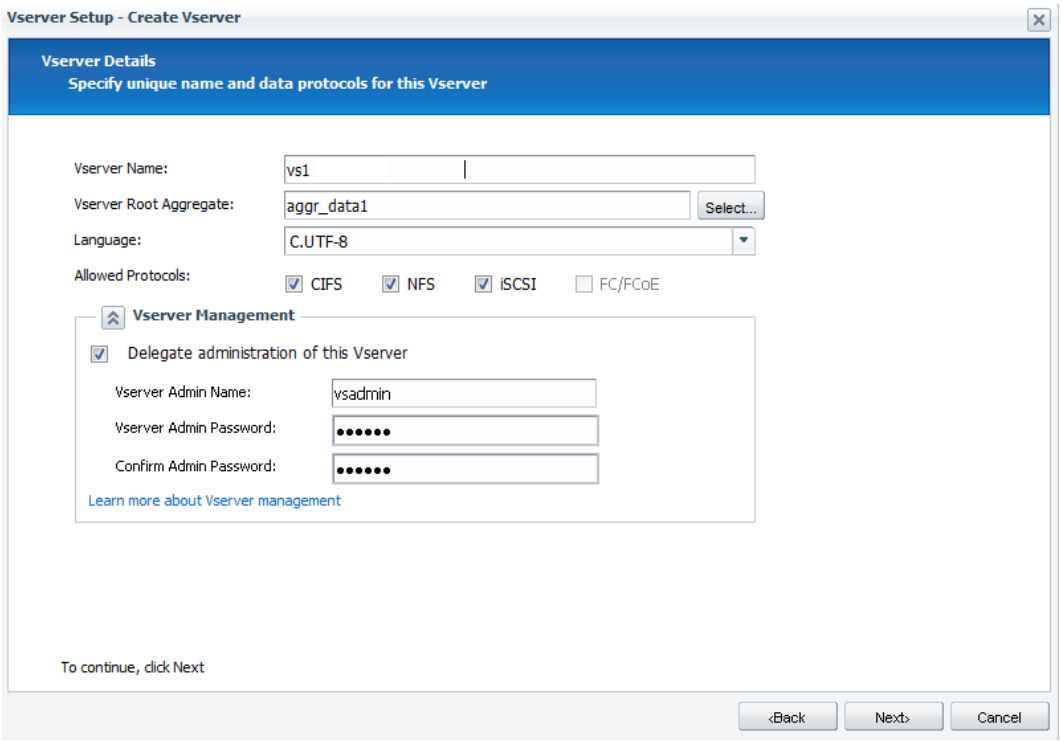

Aggregate Delegation

When delegating administrative access to an SVM admin, it is important to also delegate the aggregates that the SVM can use for provisioning new flexible volumes. A cluster administrator can use any aggregate in the cluster to create a new volume for an SVM or to relocate an existing flexible volume. An SVM admin, however, may only create new volumes in the aggregates that have been delegated to that SVM. The delegated aggregates are enumerated in the `aggr_list` option for the SVM.

In the CLI, a cluster admin can delegate aggregates using the `vserver modify` command.

```
cluster::> vserver modify -vserver vs1 -aggr-list aggr_data1,aggr_data2
```

The SVM show command can be used to check which aggregates have been assigned.

```
cluster::> vserver show -vserver vs1 -fields aggr-list
vserver aggr-list
------- --------------------
Vs1     aggr_data1,aggr_data2
```

NetApp System Manager can be used to delegate aggregates at creation time from within the SVM setup wizard, as shown in Figure 9.

Figure 9) Delegating aggregates for volume creation.

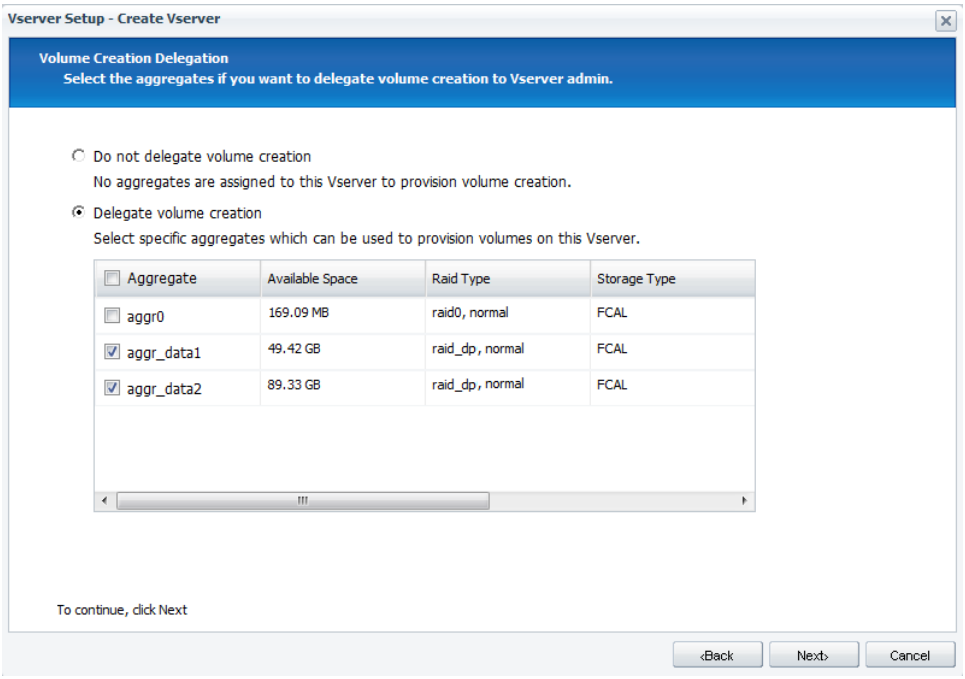

After initial setup, the SVM can later be modified with System Manager by editing the SVM and selecting the Resource Allocation tab.

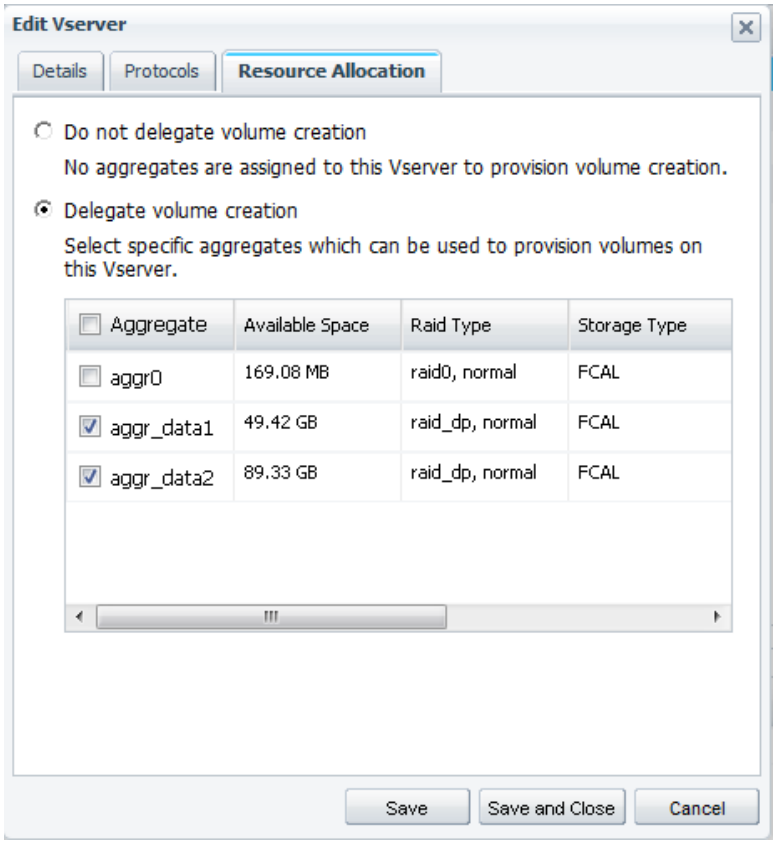

Firewall Rules

Data ONTAP includes functionality to restrict access to management service protocols available on the system. A number of system-defined firewall policies are included, and custom policies can be created using the `system services firewall policy` command directory.

The following example displays the firewall rules in place for the system-defined `mgmt` policy.

```
cluster::> system services firewall policy show -policy mgmt

Policy           Service    Action IP-List
---------------- ---------- ------ --------------------
mgmt
                 dns        allow  0.0.0.0/0
                 http       allow  0.0.0.0/0
                 https      allow  0.0.0.0/0
                 ndmp       allow  0.0.0.0/0
                 ntp        allow  0.0.0.0/0
                 rsh        deny   0.0.0.0/0
                 snmp       allow  0.0.0.0/0
                 ssh        allow  0.0.0.0/0
                 telnet     deny   0.0.0.0/0
9 entries were displayed.
```

To create a custom firewall policy using the `mgmt` policy as a base, the policy can be cloned using the `system services firewall policy clone` command.

```
cluster::> system services firewall policy clone -policy mgmt -new-policy-name mgmt-custom
```

To customize the newly cloned firewall policy, use the system services firewall policy modify command. For example, to allow ssh access only from the 192.168.1.0/24 subnet, use the following:

```
cluster::> system services firewall policy modify -policy mgmt-custom -service ssh -action allow
-ip-list 192.168.1.0/24

cluster::> system services firewall policy show mgmt-custom
Policy           Service    Action IP-List
---------------- ---------- ------ --------------------
mgmt-custom
                 dns        allow  0.0.0.0/0
                 http       allow  0.0.0.0/0
                 https      allow  0.0.0.0/0
                 ndmp       allow  0.0.0.0/0
                 ntp        allow  0.0.0.0/0
                 rsh        deny   0.0.0.0/0
                 snmp       allow  0.0.0.0/0
                 ssh        allow  192.168.1.0/24
                 telnet     deny   0.0.0.0/0
9 entries were displayed.
```

To assign the firewall policy to a LIF, modify the `-firewall-policy` attribute of the desired LIF.

```
cluster::> network interface modify -vserver vs1 -lif vs1_mgmt -firewall-policy mgmt-custom

cluster::> network interface show -vserver vs1 -lif vs1_mgmt

                    Vserver Name: vs1
          Logical Interface Name: vs1_mgmt
                            Role: data
                   Data Protocol: nfs
                       Home Node: cluster-01
                       Home Port: e0a
                    Current Node: cluster-01
                    Current Port: e0a
               Operational Status: up
                  Extended Status: -
                          Is Home: true
                  Network Address: 192.168.1.1
                          Netmask: 255.255.255.0
              Bits in the Netmask: 24
                  IPv4 Link Local: -
               Routing Group Name: d192.168.1.0/24
            Administrative Status: up
                  Failover Policy: nextavail
                  Firewall Policy: mgmt-custom
                      Auto Revert: false
   Fully Qualified DNS Zone Name: none
         DNS Query Listen Enable: false
             Failover Group Name: system-defined
                        FCP WWPN: -
                  Address family: ipv4
                         Comment: -
```

3.4 SVM Performance Monitoring and Isolation

Using Storage Quality of Service

Data ONTAP 8.2 includes a storage quality of service (QoS) feature that allows cluster administrators to manage system performance by setting maximum throughput thresholds for a variety of storage objects. By assigning storage QoS policies to a storage object, the administrator can make sure that the object does not consume more cluster resources than is expected.

Storage Objects

The following storage objects can have storage QoS policies applied to them:

- SVM
- FlexVol volume
- LUN
- File

Note: Only SVMs with FlexVol volumes can have storage QoS policies applied. SVMs with Infinite Volumes are currently not supported with storage QoS.

For purposes of this paper we will focus on applying storage QoS policies at the SVM layer.

Policy Groups

Policy groups are used to define throughput limits for a given SVM. It is also possible to create and assign a policy group that does not have a throughput limit set. This allows for monitoring of SVM throughput without enforcing limits. A maximum of 3,500 policy groups may be created per cluster, and up to 10,000 storage objects can be assigned to those groups.

Policy groups are created by using the `qos policy-group create` command. Existing groups can be modified with the `qos policy-group modify` command. Throughput limits can be defined using IOPS or MB/s as metrics.

```
cluster::> qos policy-group create -policy-group pg2 -vserver vs2 -max-throughput 1000iops

cluster::> qos policy-group modify -policy-group pg1 -max-throughput 1000MB/S
```

Policy groups can be displayed using the `qos policy-group show` command.

```
cluster::> qos policy-group show
Name             Vserver     Class        Wklds Throughput
---------------- ----------- ------------ ----- ------------
pg1              vs1         user-defined 0     0-1000MB/S
pg2              vs2         user-defined -     0-1000IOPS
2 entries were displayed.
```

Assigning Storage QoS Policy Groups to SVMs

After a policy group is defined, it can be assigned to an SVM by modifying the `-qos-policy-group` attribute of the SVM.

Per SVM, only one type of storage object can be assigned to a QoS policy group. If the entire SVM is added to a policy group, then one cannot add specific volumes and LUNs to policy groups as well.

Best Practice

Standardizing on the use of storage QoS policies at the SVM level allows cluster administrators to apply per-tenant throughput limits. Since SVM admins can create volumes and LUNs but cannot create or assign storage QoS policies to them, assigning policies to the SVM is a means for the cluster administrator to make sure that each storage object within a given SVM is covered under a storage QoS policy after being initially assigned.

The following is an example of applying a QoS policy at the SVM level:

```
cluster::> vserver modify -vserver vs1 -qos-policy-group pg1

cluster::> vserver show -vserver vs1 -fields qos-policy-group
vserver qos-policy-group
------- ----------------
vs1     pg1
```

Monitoring SVM Performance with QoS

When a storage object is assigned to a QoS policy group, this is said to define a QoS workload. QoS can provide a vast amount of detailed information regarding workload performance using the `qos statistics` command and its various options, including workload characteristics, latency, disk utilization, and CPU utilization. A complete discussion is beyond the scope of this document; however, further information can be found in the Clustered Data ONTAP 8.2 System Administration Guide for Cluster Administrators.

3.5 Data Protection

Clustered Data ONTAP provides a means to replicate data within the same cluster or to a peered cluster. Both load-sharing mirrors and data protection mirrors are supported. For detailed information on peering and mirroring, see TR-4015: SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode.

Cluster and SVM Peering

It is possible to create intracluster volume mirrors for load sharing and data protection. It is also possible to create intercluster mirrors if special intercluster LIFs are set up on both the source and destination cluster and the clusters are connected in what's known as a peer relationship. Replication traffic between the two clusters will occur over the intercluster LIFs.

Data ONTAP 8.2 introduces the concept of SVM peering in addition to cluster peering. For both intracluster mirroring and intercluster mirroring between peered clusters, the SVMs containing the source and destination volumes must also be in a peer relationship. This allows for a higher level of granularity in the control of data protection mirrors and provides a foundation for delegating the control of data protection mirrors to SVM administrators.

Note: SVMs can only be peered if they are within the same cluster or if their containing clusters are also peered.

SVM peers can be created and shown with the `vserver peer` command. The following example sets up an SVM peer relationship between two SVMs on the same cluster:

```
cluster::> vserver peer create -vserver vs1 -peer-vserver vs2 -applications snapmirror

Info: 'vserver peer create' command is successful.

cluster::> vserver peer show-all
            Peer        Peer                                  Peering
Vserver     Vserver     State        Peer Cluster             Applications
----------- ----------- ------------ ------------------------ ---------------
vs1         vs2         peered       cluster                  snapmirror
vs2         vs1         peered       cluster                  snapmirror
2 entries were displayed.
```

Because the two SVMs to be peered are in the same cluster, no further action is required after issuing the peer create command. If, however, the SVMs to be peered are in separate clusters that are in a cluster peer relationship, then the SVM peer relationship will be in a pending state until the cluster administrator of the remote cluster accepts the peer relationship.

```
cluster2::> vserver peer show

            Peer        Peer
Vserver     Vserver     State
----------- ----------- ------------
vs1          vs3         pending
```

To accept the SVM peer request, the remote cluster administrator should issue the `vserver peer accept` command.

```
cluster2::> vserver peer accept -vserver vs1 -peervserver vs3
```

Unique SVM Naming Requirements

In order to peer two SVMs, their names must be unique on both the source and destination clusters. For instance, for an SVM called vs1 on cluster1 to be peered with SVM vs2 on cluster2, cluster1 cannot have an SVM called vs2, and cluster2 cannot have an SVM called vs1.

Best Practice
Adopt a naming scheme that will make sure that SVM names are unique across peered clusters. One means to accomplish this is to name each SVM using a fully qualified domain name.

Language Considerations

In order to use SnapMirror to copy a volume from one SVM to another, both the source and destination volumes must have the same language setting. In versions of Data ONTAP prior to 8.2, all volumes contained within an SVM shared the same language setting as the SVM. The SVM language could only be set when the SVM was created and could not be changed. As a result, SnapMirror copying could only occur between two SVMs of the same language.

Starting in Data ONTAP 8.2, volumes can have a different language than the SVM in which they are contained. By default, a volume will inherit the language of the containing SVM if no language is specified when the volume is created. It is possible to specify a different language for the volume at creation time. It is possible to change the language of an SVM; however, volume language cannot be changed. In Data ONTAP 8.2, SnapMirror copying between SVMs of different language types can occur, but the source and destination volumes must have the same language.

The default language for SVMs in Data ONTAP 8.2 is C.UTF-8. This setting provides a neutral, non–country specific encoding. Unless technical requirements dictate the use of a different language encoding, consider using C.UTF-8.

UTF-8 encoding is variable length. The presence or absence of certain bits in a byte of UTF-8 indicates the number of bytes that make up the encoded character. Therefore, the use of certain special characters in *[language]* will cause errors if *[language].UTF-8* parsing is used.

Use these qualifying questions to determine the best option for the SVM language encoding:

- Are there any older CIFS clients (Windows® 95/95/ME)? If yes, match the SVM language encoding to the *[language]* client locale.
- Are there any NFSv2/3 clients not using UTF-8? If yes, match the SVM language encoding to the *[language]* client locale.
- Are all the CIFS clients post Windows 95/98/ME and all the NFS clients using a UTF-8 locale? If so, set the SVM language encoding to C.UTF-8. If some clients are not using UTF-8, set the SVM language encoding to the *[language]* client locale.

Table 7 describes how to specify an SVM's language. en_US is used in each example, but can be substituted with the appropriate client locale.

Table 7) SVM Language Recommendations

Client Protocol	Client Encoding Type	SVM Language Recommendation
CIFS (Windows 95/98/ME)	ISO 8859-1	Use *[language]*. Example: "en_US."

Client Protocol	Client Encoding Type	SVM Language Recommendation
CIFS (Windows NT® 3.1+)	UCS-2	• If all clients use UTF-8, use C.UTF-8. • If any client does not use UTF-8, use *[language]*. Example: “en_US.”
NFSv2/3	Non-UTF-8 client locale	Use *[language]*. Example: “en_US.”
NFSv2/3	UTF-8 client locale	Use C.UTF-8.
NFSv4	UTF-8	Use C.UTF-8.
FC or iSCSI	N/A	C.UTF-8 preferred, C/POSIX is acceptable.

In deployments with mixed non-Unicode client language encodings, the SVM language encoding should match the highest priority client language encoding.

Incorrectly displayed characters might not be remediable in situations where different clients use the same file access protocol (for example, NFSv3), but with different language encodings. The same is true when clients have a mix of client locales.

The following best practices should also be considered when selecting an SVM’s language setting:

- For environments using both CIFS and NFS clients, match the NFS client language encoding.
- Do not use non-ASCII characters in your file names, including “smart quotes” and currency symbols (for example, € for euro, £ for UK pound).
- UTF-8 language encoding is preferred over non-UTF-8 encodings. However, there are situations in which this is not recommended. For example, when NFS clients are using non-UTF-8 encodings, the SVM’s language should be set to the same non-UTF-8 encoding.
- When sharing files between CIFS and NFS, only use characters that are legal for both and are in the NFS character set.
- For customers planning to move from NFSv3 to NFSv4, the SVM’s language should match the NFSv3 client language.
- If a volume inside an SVM will be a SnapMirror destination, choose the same language used on the volume that will be the SnapMirror source.

Load-Sharing Mirrors

Load-sharing (LS) mirrors are a special type of SnapMirror that can increase performance and availability of volumes accessed using CIFS or NFSv3. LS mirrors can be used to distribute reads across multiple nodes for datasets that are read-only. LS mirrors can also be used to increase the availability of an SVM’s namespace when used to create multiple instances of the SVM root volume.

Creating LS mirrors of an SVM root volume on each node of the cluster provides multiple redundant copies that can be used should the primary volume become unavailable. If the root volume should become temporarily unavailable, read access to the volume is provided through the LS mirrors, and the namespace remains available. If the root volume is permanently unavailable, one of the mirrors can be promoted to make write access available.

There are some trade-offs that must be made in order to achieve this degree of availability. Because the SVM root volume is the location of the SVM namespace root, any new volumes that are junctioned into the namespace at the root are not visible until the LS mirror set has been updated. This update can only

be done by a cluster administrator, because the command required (`snapmirror update-ls-set`) is not available at the SVM level. The cluster administrator can schedule these updates to occur on a regular basis. SVM admins should be mindful that updates to the namespace will only be visible after the mirror set is updated. Any automated workflows that operate under cluster credentials and update the namespace should also update the LS mirror set.

Best Practice
For added resiliency and namespace availability, consider creating a load-sharing mirror set for SVM root volumes.

Data Protection Mirrors

Data ONTAP provides the ability to asynchronously mirror volumes from one SVM to another SVM on the same or a different cluster by creating a SnapMirror data protection mirror. These data protection mirrors are used to maintain local backup copies or to provide a remotely replicated copy, which can be used for disaster recovery and business continuance.

Vaulting

With the 8.2 release, SnapVault® technology is introduced to clustered Data ONTAP. SnapVault allows for asymmetric Snapshot retention between source and destination volumes. Local Snapshot copies could be retained for short-term backup and recovery, while SnapVault copies are retained for long-term archival storage.

3.6 Management Tools

CLI

The Data ONTAP command line interface is available to cluster administrators as well as SVM administrators. Cluster administrators can use ssh to connect to the cluster management LIF in order to access the CLI in a cluster context. From this context, the admin can manage items related to the cluster as a whole, as well as each SVM on the cluster.

SVM administrators can use ssh to connect directly to the CLI in an SVM context. From this context they are able to manage items directly pertaining to the SVM to which they have been delegated access. SVM admins are able to manage data protocols and services; storage objects such as volumes, LUNs, and qtrees; and Snapshot copies of volumes and monitor the overall health of the SVM.

Cluster admins, once logged in to the CLI, can switch to an SVM context using the `vserver context` command.

The command line interface is currently the only interactive management tool available to SVM administrators.

Note: There is a maximum of 64 concurrent SSH sessions in Data ONTAP.

NetApp Manageability SDK

Data ONTAP includes a rich API set available through the NetApp Manageability SDK. The SDK provides APIs and sample code to develop custom management tools using C, C++, Java®, Perl, C#, VB.NET, Windows PowerShell™, Python, and Ruby. These APIs can be used to integrate Data ONTAP management with existing orchestration tools or create custom management portals.

Data ONTAP PowerShell Toolkit

For those who prefer to manage and automate through the use of Windows PowerShell, NetApp offers a powerful set of cmdlets for use with Data ONTAP. Available through the NetApp communities portal, the Data ONTAP PowerShell Toolkit is a powerful and versatile means to manage clusters and SVMs alike. For more information, see Making the Most of Data ONTAP PowerShell Toolkit.

OnCommand System Manager

NetApp OnCommand System Manager is a graphical management tool available for Windows and Linux®. It provides GUI-based management of storage systems and storage objects, including the ability to configure data protocols, provision storage objects, and create and manage SVMs.

Note: OnCommand System Manager login is only available to cluster administrators.

Figure 10) OnCommand System Manager.

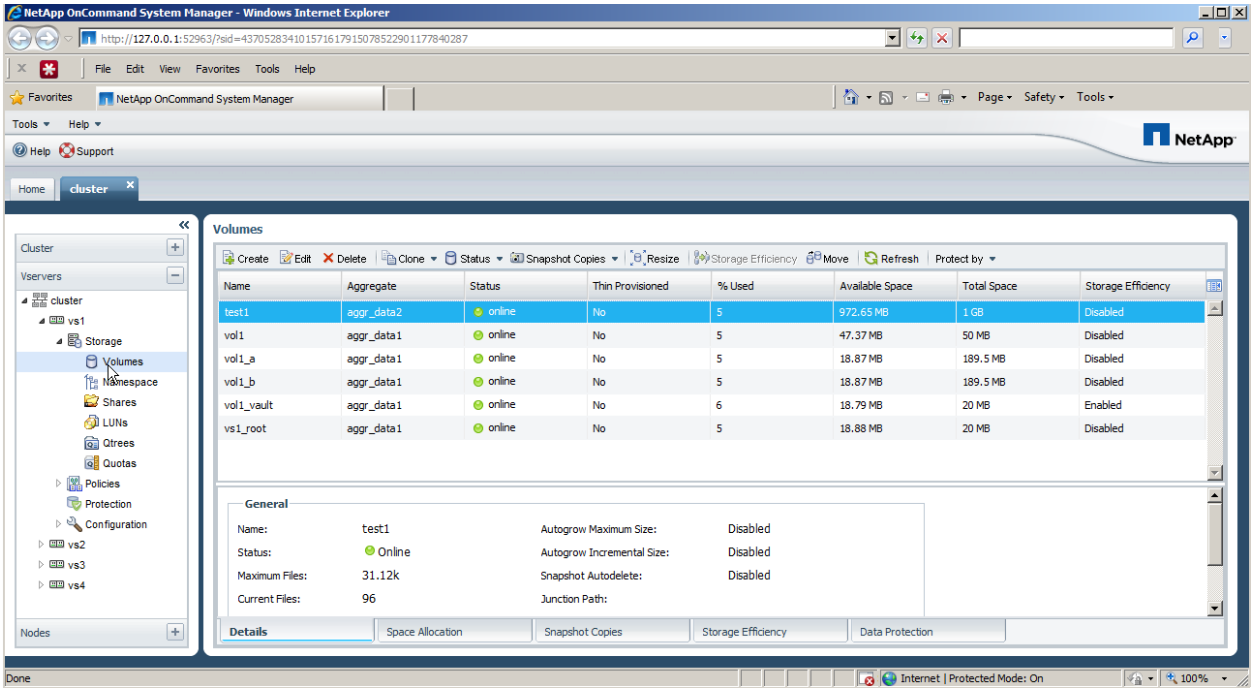

OnCommand Unified Manager

OnCommand Unified Manager is a centralized interface for managing multiple clusters across multiple environments. It includes capabilities for monitoring, alerting, and reporting on storage infrastructure at scale.. Further information about OnCommand Unified Manager can be found at the NetApp OnCommand landing page: http://www.netapp.com/oncommand.

Figure 11) OnCommand Unified Manager.

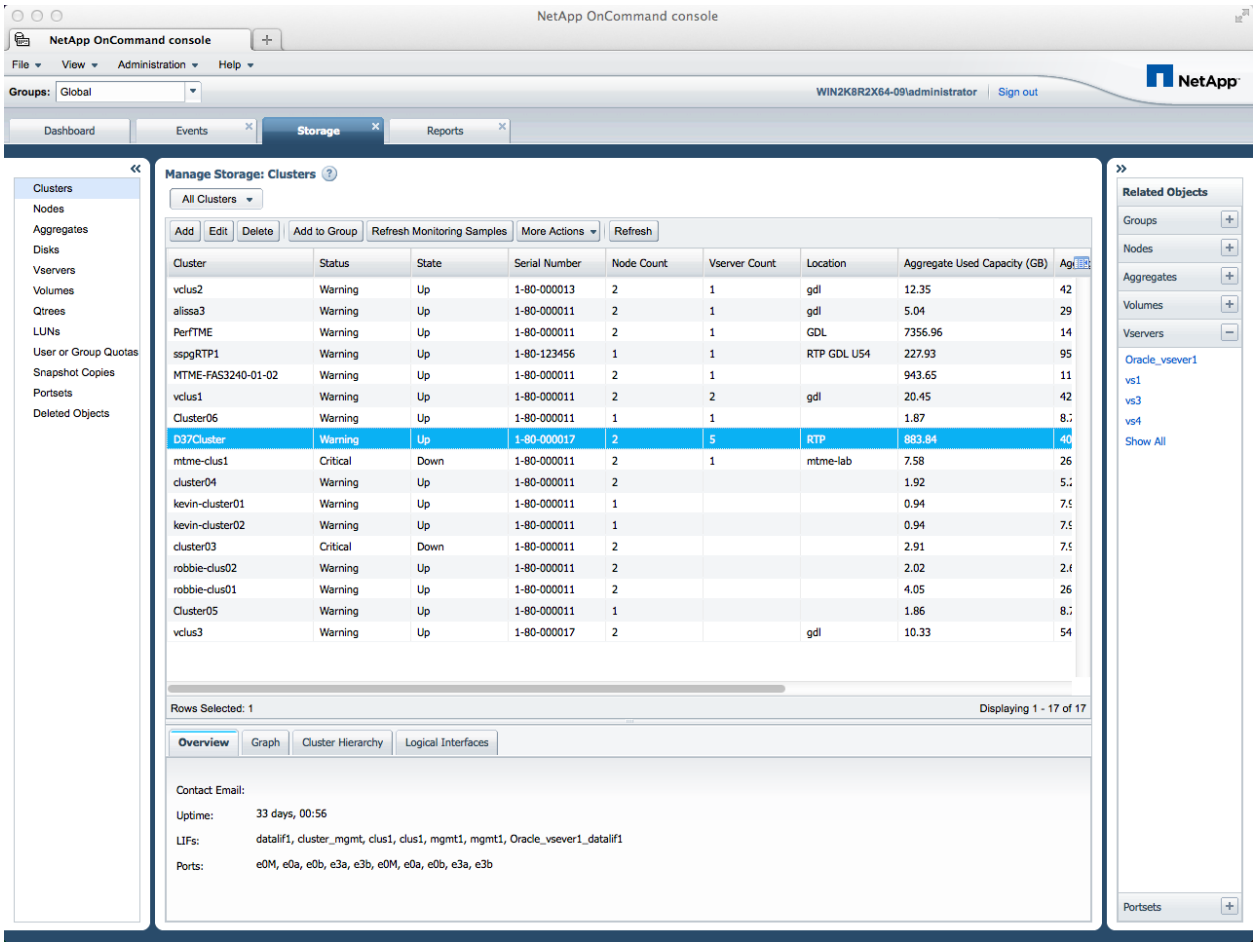

OnCommand Workflow Automation

OnCommand workflow automation (WFA) is a powerful tool for storage automation. WFA allows storage administrators to create, test, and publish custom workflows for an endless variety of storage functions and tasks. For further information on WFA, see the OnCommand landing page on the NetApp Communities at www.netapp.com/oncommand_community.

Figure 12) OnCommand workflow automation portal.

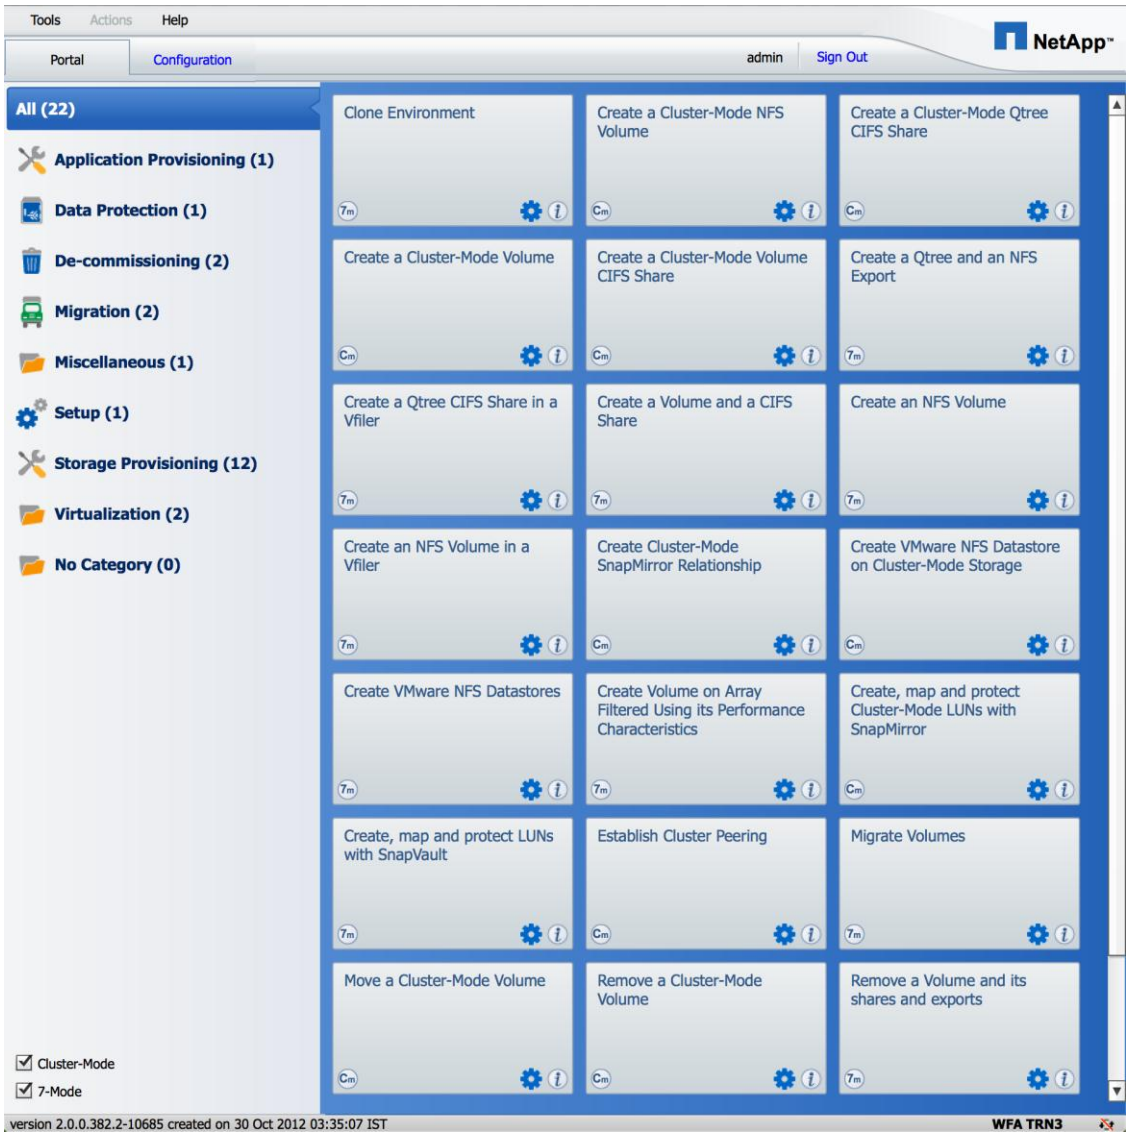

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. 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.

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.

Go further, faster®

www.netapp.com