Technical Report

# Desktone DaaS on NetApp FAS Storage with Clustered Data ONTAP

Ganesh Kamath, NetApp
Tony Alvino, VMware
January 2014 | TR-4258

TABLE OF CONTENTS

## LIST OF TABLES



## LIST OF FIGURES

# 1 Desktone DaaS on Clustered Data ONTAP

## 1.1 Overview

The enterprise desktop world is changing quickly, responding to forces of IT consumerization and end-user mobility, platform updates, high costs, and security concerns. An increasing number of businesses are looking to centralize the management and provisioning of desktop environments. In addition, they are stumped by either technical complexity issues or the upfront costs required to make the transition.

Service providers who want to offer virtual desktops as a service (DaaS) as a means to address these needs are increasingly turning to Desktone on NetApp. DaaS overcomes the challenges of deploying traditional on-premises VDI solutions by eliminating the need for large upfront capital expenditures (capex) and the necessity of in-house technical expertise to deploy and manage the VDI infrastructure. DaaS transforms desktops from the capex outlay inherent in enterprise on-site VDI and physical desktop hardware refreshes into a predictable, easy to budget opex item.

NetApp complements Desktone's capabilities by delivering innovative, scalable multi-tenant storage and data management solutions that boost efficiency and flexibility, accelerating and optimizing the Desktone DaaS delivery process.

The Desktone on NetApp® solution provides the essential components for a successful cloud-hosted desktop service:

- A comprehensive blueprint that enables service providers to dramatically simplify and accelerate time to market
- A highly cost-effective and scalable multi-tenant platform built on proven, patented technology
- Unparalleled expertise that enables our partners to quickly onboard, successfully service, and continually grow revenue from customers

## 1.2 Target Audience

This document describes the Desktone DaaS on NetApp solution. The target audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy this solution.

## 1.3 About Desktone

Desktone by VMware is changing the way people buy and manage desktops. Going forward in this document, Desktone by VMware will just be referred to as Desktone. The Desktone Cloud, Desktone's DaaS offering, provides all of the benefits of virtualized desktops without any of the hassles. By delivering virtual desktops in the cloud, Desktone enables businesses to rapidly provision desktops to users on any device, anywhere, without the upfront costs and complexity of traditional desktop virtualization, transforming desktops from a capex to an opex item. For more information about the Desktone Cloud, visit www.desktone.com.

## 1.4 Business Benefits

Desktone DaaS on NetApp enhances service provider capabilities in several ways. This section describes the benefits from a high-level perspective.

### On-Demand Self-Service

- Tenant enterprise admins can provision and remove virtual desktops as needed automatically without involvement of the service provider.

- End users can configure session settings, default desktops, power on/power off, and reboot and reset the protocol on their own desktops without involvement of the enterprise admin or the service provider.
- The Desktone resource manager appliances leverage NetApp FlexClone® technology to rapidly deploy new virtual desktops on thin-provisioned storage without affecting the CPU and memory utilization of the host platform.

## Broad Network Access

Virtual desktops hosted on the Desktone DaaS on NetApp solution can be accessed by a broad range of devices, including thin clients, desktops, laptops, smartphones, and tablets, using the following protocols.

- RDP
- VNC
- PCoIP
- NX
- RGS

## Rapid Elasticity

- The Desktone Enterprise Center web console enables tenant administrators to quickly and easily provision virtual desktops for end users based on desktop images.
- The platform allows you to add capacity by simply adding additional servers and NetApp storage, tenants, and data centers as needed.
- The virtual desktops are hosted on the VMware® ESXi™ hypervisor and NetApp storage platforms, both of which allow thin provisioning and oversubscription.
- NetApp FlexVol® volumes are used to store data in the environment. Because a FlexVol volume is abstracted from the underlying disk, you can create a volume to meet your capacity needs without regard for physical layout. Furthermore, service providers can grow or shrink a FlexVol volume easily at any time.

## Resource Pooling

The Desktone DaaS on NetApp solution's network, computing, and storage resources are pooled to serve numerous consumers using a secure multi-tenant architecture.

From a storage perspective, clustered Data ONTAP® is an inherently multi-tenant storage operating system. 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.

For infrastructure service providers, multiple SVMs should be deployed to securely allocate storage resources to different tenants and delegate management of those resources (resource pooling illustrated in Figure 5.) without dedicating physical hardware to each tenant or exposing multiple tenants and their data to one another. Service providers can create tiers of service or resource pooling 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™ or Flash Pool™, Gigabit Ethernet (GbE) versus 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.

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. 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. Volumes can be easily moved when tenant storage needs change. Workloads can be moved between nodes of differing memory and CPU potential, as well as differing amounts of flash-based cache, by using Flash Pool and/or Flash Cache as performance tiers.

Another 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 VLANs and routing groups.

## Measured Service

- Service providers define quotas for each of the tenants to control the number of desktops they can provision. Tenant administrators are free to provision and delete desktops up to their defined quota.
- The Desktone DaaS platform enables service providers and tenants to report desktop usage for each tenant. For each data center associated with a tenant, the platform captures the desktop model quota and desktops in use for that model.
- The NetApp OnCommand® suite of tools can be leveraged to manage, monitor, and report on host and storage utilization.
- Resource pooling that is inherent in Data ONTAP can be used to provide different levels of “measured service” for each tenant based on that tenant’s requirements.

## Massive Scale

- The Desktone DaaS platform on NetApp scales to hundreds of thousands of desktops across thousands of tenants across dozens of data centers. This scale is achieved with a minimal set of Desktone management and infrastructure components, thereby reducing the operational expense associated with operating VDI at scale: that is, DaaS.
- The ability of NetApp clustered Data ONTAP to deliver seamless infinite scalability along with nondisruptive operations and proven efficiency is complementary to the Desktone platform’s ability to scale, and the integration of both technologies provides for a very compelling solution.
- The NetApp storage platform can easily scale up and out to support the virtual desktop workload as well as any application workload the tenants might choose to host in the environment.
- NetApp Flash Cache, Flash Pool, and data deduplication greatly reduce the number of disks required to support the environment, along with the requisite power and cooling to support them.

## Geographic Distribution

- The Desktone platform is natively multi-tenant, multi–data center capable. Therefore, the service provider administrators can centrally manage data centers worldwide, allowing the service provider to scale out and provide better performance with less network latency between customers and their virtual environments.
- Tenant environments can span multiple data centers and be centrally managed.
- NetApp SnapMirror® technology can be leveraged to replicate data between data centers. SnapMirror is also dedupe aware, which provides for thin replication: that is, only changed and deduplicated blocks are replicated to the secondary site, which provides significant cost and bandwidth savings.

### Virtualization

- The platform is deployed on ESX®, and virtualization is leveraged throughout the solution.
- The Desktone service provider appliances, resource manager appliances, and tenant appliances are hosted on VMware ESXi hypervisors.
- The storage controllers are virtualized on the NetApp FAS systems and presented as storage virtual machines that provide a level of abstraction between the physical storage controllers, allowing storage portability and security.
- Desktone also integrates with NetApp VSC, which is a VMware vCenter™ plug-in for automated storage and virtual desktops provisioning along with storage monitoring, backup, and restore capabilities.

### Low-Cost Software

- The principal barrier to offering a cost-effective desktop-as-a-service solution is software licensing. One of the key benefits of the Desktone solution is that no additional SQL Server® or Microsoft® Server licenses are required beyond the hypervisor.
- 
- One of the key features of Desktone is multi-tenancy and hence the ability to leverage common infrastructure and benefit hugely from the economies of scale.
- Additionally, competing solutions require costly per-user licenses that drive the cost per desktop out of consumers' reach. By offering a more competitive licensing model, NetApp and Desktone can drive the cost down to a level that allows widespread adoption of this solution.

# 2 Desktone DaaS Design Elements

To provide a best-in-class, multi-tenant, cloud-based desktop-as-a-service offering for service providers, NetApp has partnered with Desktone.

## 2.1 Clustered Data ONTAP

Desktone DaaS on clustered Data ONTAP leverages several technologies that are unique to NetApp, resulting in higher performance, lower total cost of ownership (TCO), and tighter integration with the Desktone environment:

- Clustered Data ONTAP addresses the challenges facing your growing and dynamic business by extending the innovation of NetApp Data ONTAP, the world's number-one branded storage operating system. Our unified cluster architecture scales and adapts to your changing needs, reducing risk and cost. Clustered Data ONTAP is designed to eliminate downtime, allowing you to service your infrastructure without disrupting access to user data and applications, even during regular business hours.
- Proven operational efficiency helps you simplify your overall storage environment and manage storage infrastructure at scale by automating important processes and increasing productivity. You can add capacity as you grow across both SAN and NAS environments without reconfiguring running applications. We let you start small and grow big without the disruptive hardware upgrades required by other storage vendors.
- Clustered Data ONTAP provides up to 24 storage controllers—or nodes—managed as a single logical pool so your operations scale more easily. NetApp supports the broadest set of storage protocols and is the only provider to deliver both SAN and NAS data access from a single, unified scale-out platform.

Some of the key benefits NetApp technologies provide to this solution are:

- **Unified architecture**. Significant operational cost savings with unified scale-up/scale-out capabilities to support cloud-scale deployments of Desktone-powered DaaS platforms.

- **Storage efficiency**. Significant cost reductions are realized through the use of RAID-DP® technology, thin provisioning, deduplication, thin replication, and FlexVol volumes.
- **Seamless scalability** (performance and capacity) as and when requirements and demands change. This includes vertical as well as horizontal scaling. A given cluster can be scaled up, down, or out based on performance and capacity requirements, physical limitations, best practices, data center power and cooling availability, and so on.
- **Performance and virtual storage tiering**. Enhanced user experience with NetApp's unique Virtual Storage Tiering and write I/O optimization, which strongly complements NetApp's storage efficiency capabilities.
- **Flash Cache** is part of the NetApp Virtual Storage Tier and in a virtual desktop environment allows the storage controllers to cache the most frequently accessed blocks of the desktop virtual machines into flash memory, dramatically reducing disk reads and increasing both read and write performance.
- **Operational agility, nondisruptive operations and data mobility**. Enhanced Desktone solution management with tight platform integration and storage mobility capabilities.
- **Flash Pools**. NetApp Flash Pool is a storage cache option within the NetApp Virtual Storage Tier (VST) product family. A Flash Pool aggregate configures solid-state drives (SSDs) and hard disk drives (HDDs) into a single storage pool (aggregate), with the SSDs providing a fast response time cache for volumes that are provisioned on the Flash Pool aggregate.
- **FlexClone** is leveraged by the Desktone platform to efficiently clone virtual desktop images in the storage tier, offloading this from the typically CPU- and memory-intensive process that is required to handle the cloning process in the virtualization compute tier.
- **Deduplication** reduces storage utilization by rates often over 90% for virtual desktop environments; it also enhances related caching and replication capabilities by only having to store/replicate a single instance of each replica block.
- **OnCommand suite** provides centralized storage management and monitoring.
- **Data protection**. Enhanced local and remote protection of the Desktone management platform, virtual desktop OS data, and user data, with very low overhead for both cost and operations.
- **Snapshot™** copies provide the capability to protect and retrieve deleted files easily with the capability to instantaneously revert file systems to a previous point in time. This allows quick recovery from application corruption issues.
- **SnapMirror** provides the capability to replicate virtual desktop and user content between storage systems in a thin and deduplicated format.

Figure 1) Clustered Data ONTAP customer benefits.

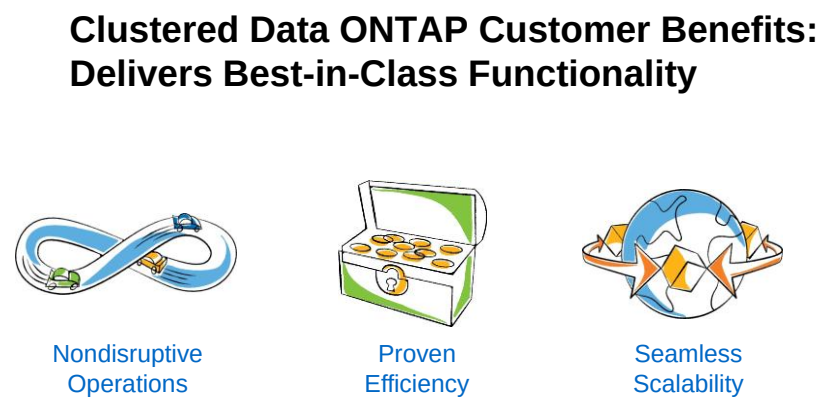


## Nondisruptive Operations

- Performs lifecycle operations without interrupting business operations.
- Clustered Data ONTAP improves service levels over the lifecycle of an application with the ability to dynamically assign, promote, and retire storage resources.
- Clustered Data ONTAP provides continuous data access when upgrading controllers and shelves, taking less time (minutes versus hours/days) and fewer resources than competitor data copy approaches.

Figure 2) Clustered Data ONTAP NDO.

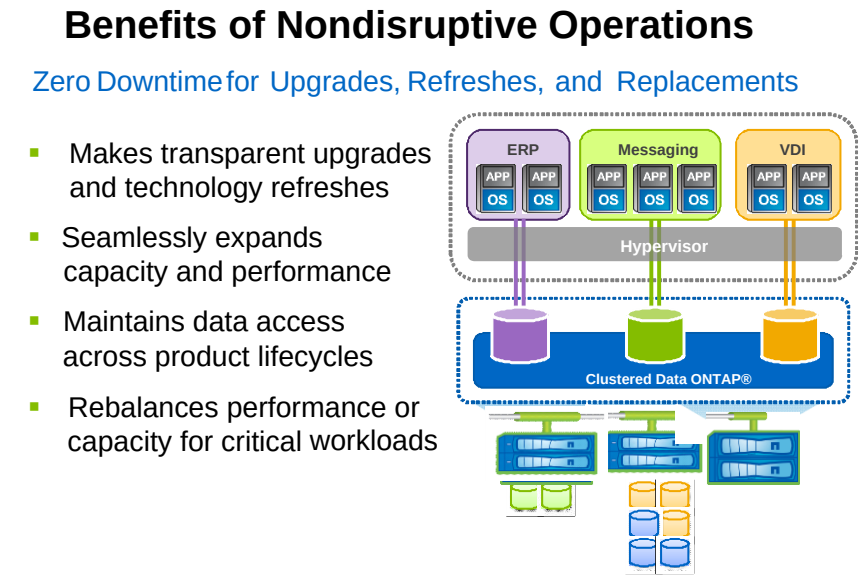


## Proven Efficiency

- Simplifies, automates, and increases productivity while lowering IT costs.

- Clustered Data ONTAP enables pervasive cost reductions with the industry's most comprehensive storage efficiency offering.
- NetApp is the only storage provider to deliver proven efficiencies for both SAN and NAS on entry-level, midtier, enterprise, software-based, and virtualized third-party arrays.

**Figure 3) Clustered Data ONTAP efficiency features.**

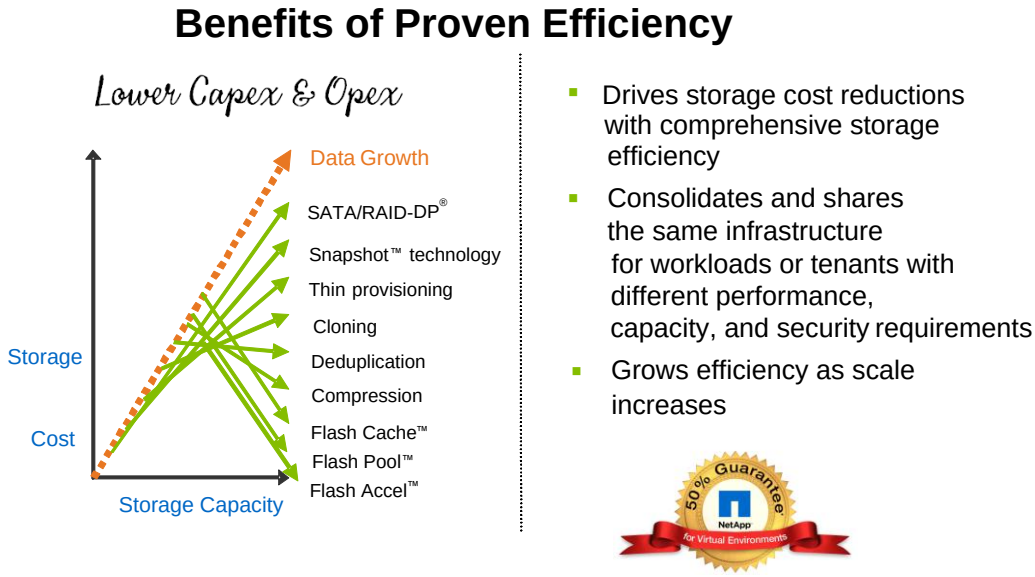


## Seamless Scalability

- Scales in both SAN and NAS environments: start small and grow big. Storage systems running clustered Data ONTAP can scale SAN and NAS from terabytes to tens of petabytes transparently and without reconfiguring running applications.
- Clustered Data ONTAP makes it possible to scale capacity, performance, and operations without compromise, regardless of application. Seamlessly expand capacity and performance.
- Scales capacity, performance, and operations without compromise.

Figure 4) Clustered Data ONTAP seamless scalability.

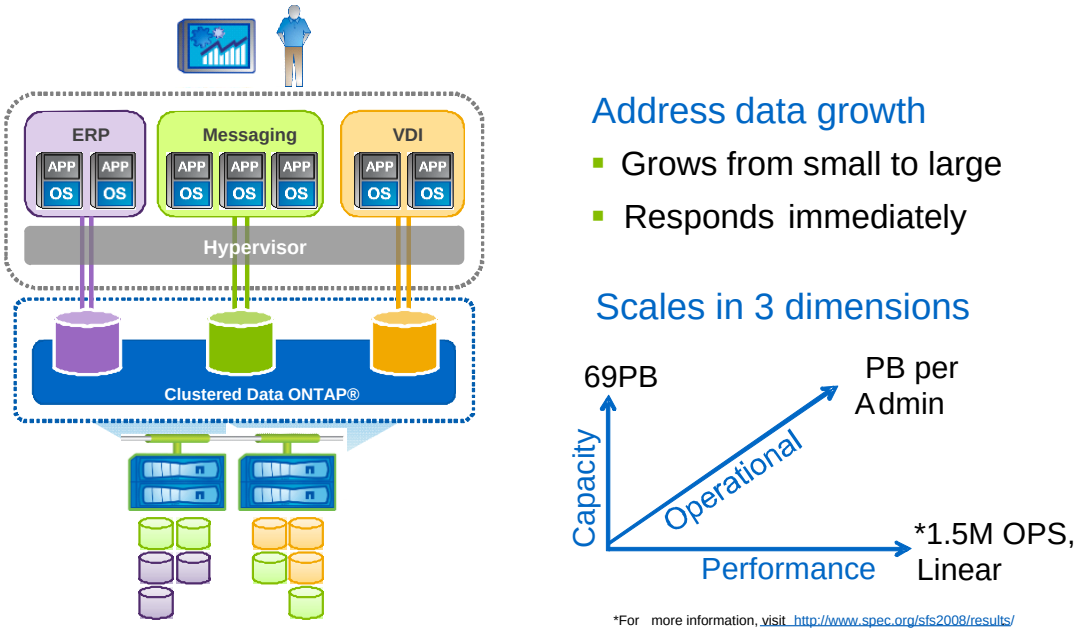


## Multi-Tenancy

Clustered Data ONTAP is an inherently multi-tenant storage operating system. 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 5) Clustered Data ONTAP multi-tenancy.

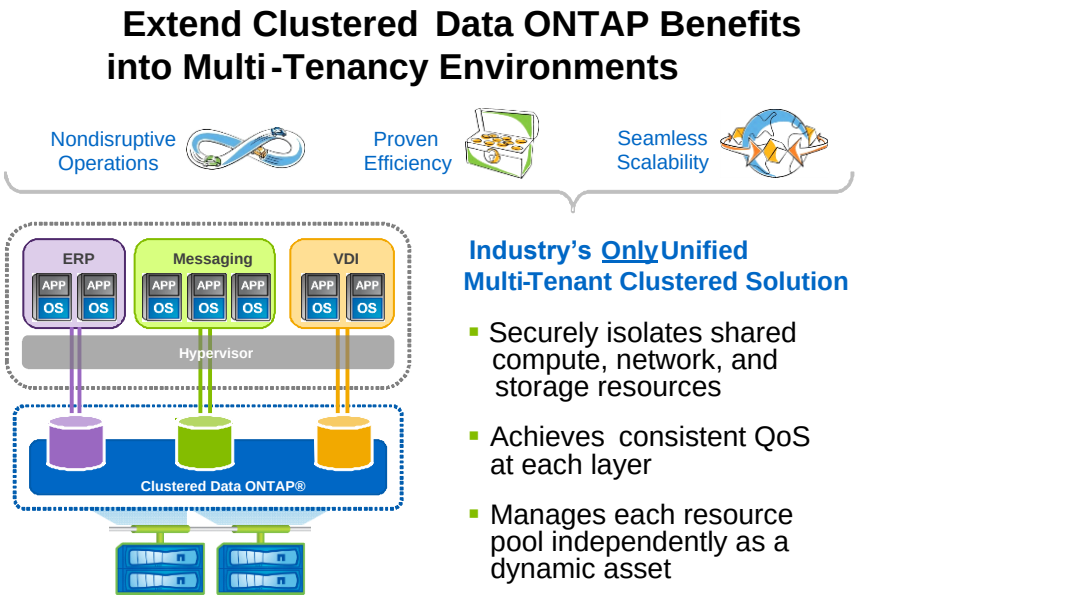

## 2.2 Desktone Platform Overview

Desktone's DaaS platform is a virtual desktop solution built from the ground up for service providers, with many features specifically for operating as a service in the cloud.

**Figure 6) Desktone DaaS platform (graphic supplied by Desktone).**

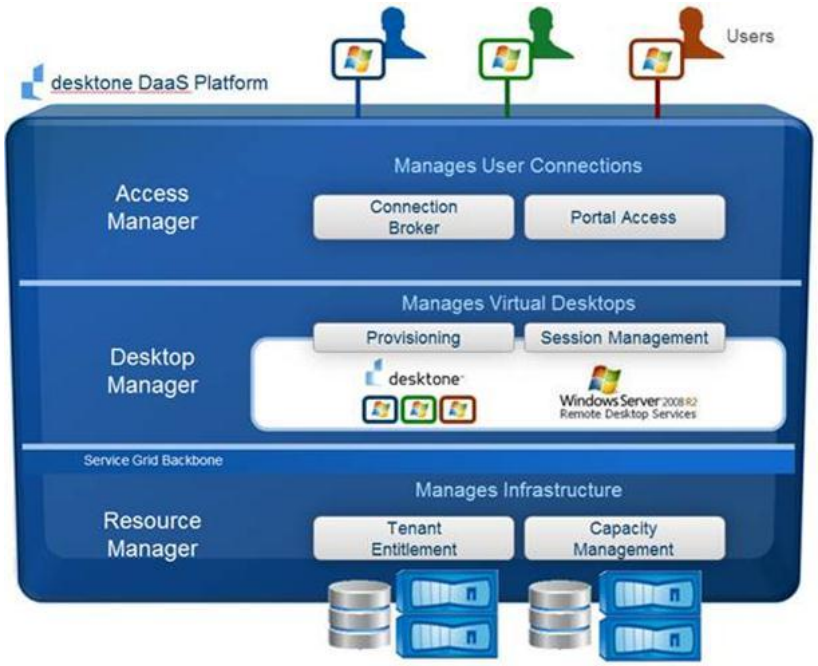


The list of features is long, but the top five that set Desktone apart are:

- **Multi-tenancy.** The Desktone DaaS platforms inbuilt multi-tenancy feature allows service providers to manage numerous tenants across multiple data centers from a single management platform, eliminating the need for the per-tenant management platforms of traditional VDI solutions.

**Figure 7) Desktone multi-tenancy (graphic supplied by Desktone).**

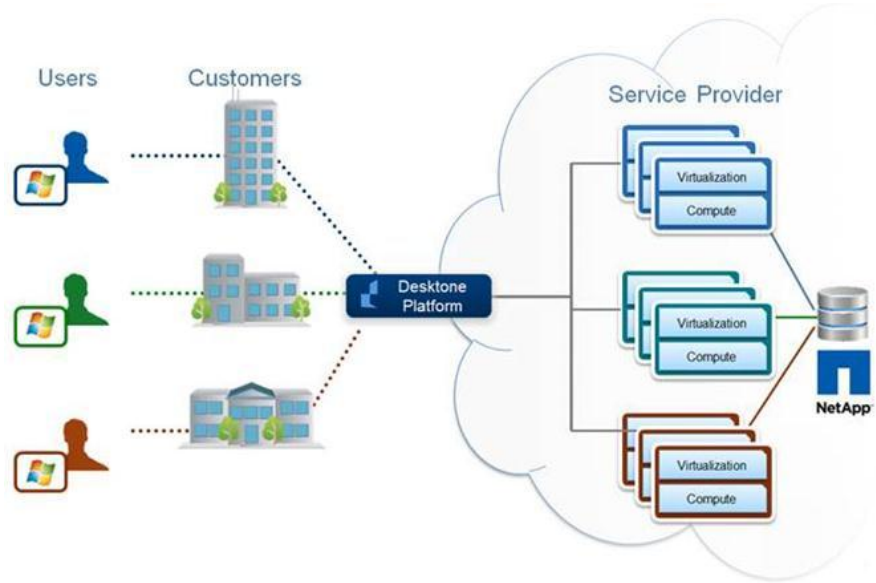

- **Elastic scalability.** The Desktone Enterprise Center web console enables tenant administrators to quickly and easily provision virtual desktops for end users based on new or existing desktop images. Desktone's tight integration with NetApp FlexClone technology offloads the virtual desktop provisioning process from the tenant host to the storage controllers, significantly reducing deployment time and impact on tenant host CPU/memory resources.

**Figure 8) Desktone DaaS elastic scalability (graphic supplied by Desktone).**

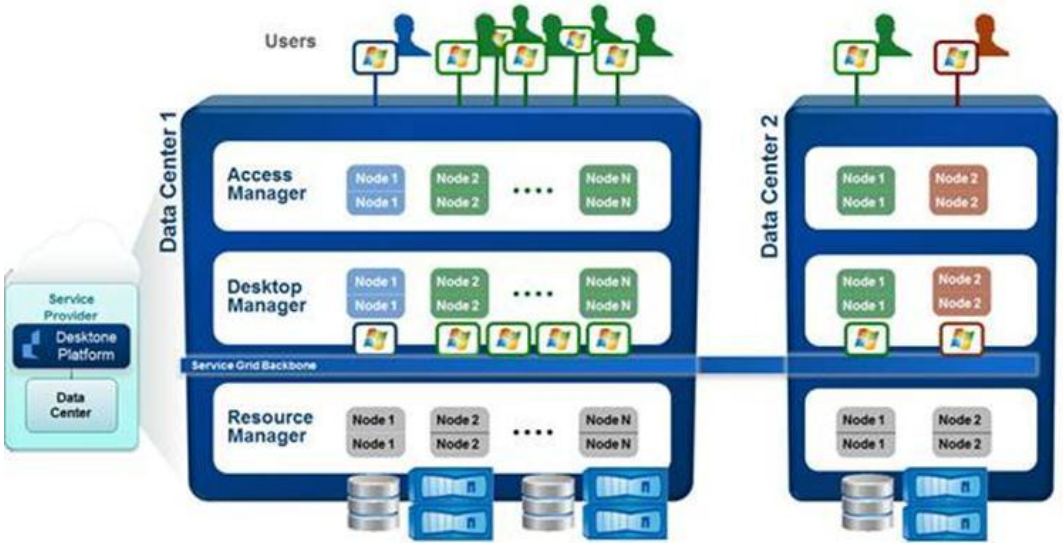


- **Enterprise integration.** The Desktone DaaS platform integrates with the tenant's internal IT infrastructure and leverages the tenant's Active Directory® service to control access to the virtual desktops. In addition, it enforces group policies, which allows administrators to control the behavior of the virtual desktops.

**Figure 9) Desktone DaaS enterprise integration (graphic supplied by Desktone).**

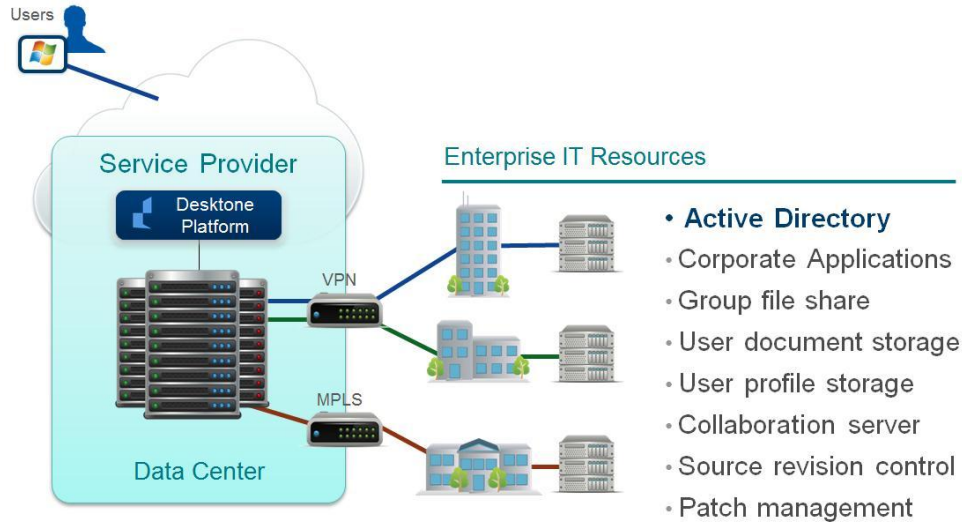


- **Cost-effective model.** The scale of the solution allows service providers to share resources across multiple tenants, reducing the financial burden and distributing the costs across tenants. Integration with NetApp technologies such as FlexClone, Flash Cache intelligent caching, Flash Pool, deduplication, SVMs, Snapshot, and SnapMirror further drives down costs while increasing

performance, efficiency, and flexibility. Implementing cloud-based virtual desktops for numerous tenants on traditional VDI platforms is highly cost prohibitive. Desktone's licensing is more competitive than the traditional VDI licensing offered by traditional enterprise VDI solutions, most of which require a separate siloed deployment per tenant.

- **Tiered role separation.** There are clearly defined administrative boundaries for service provider administrators, tenant administrators, and end users that are separated by the web portal for each role.

**Figure 10) Desktone DaaS tiered role separation (graphic supplied by Desktone).**

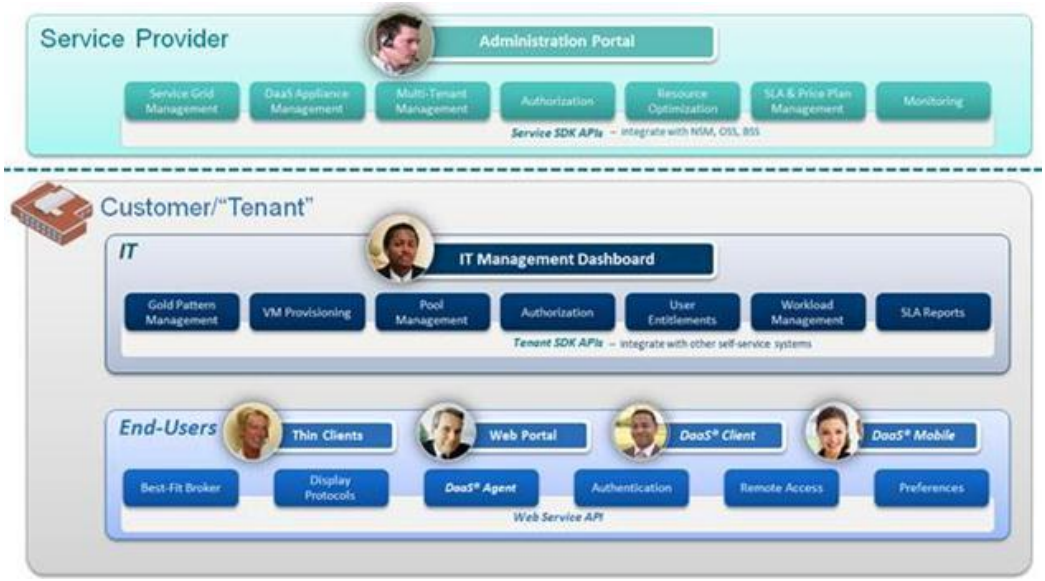


The Desktone DaaS platform consists of two distinct tiers: enterprise and service provider. Separating the enterprise and service provider tiers is the key to enabling multi-tenancy VDI for DaaS. It allows enterprises to manage virtual desktop environments supported by physical resources owned and maintained by a third party. Service providers benefit from optimizing use of their data center assets while offering a new, highly scalable, multi-tenant, value-added service to customers.

Organizations consuming the Desktone DaaS on NetApp solution benefit from several key advantages over the traditional VDI/terminal server approach, including:

- No upfront capex costs
- No infrastructure to manage, no server hardware, and no data center/power/cooling costs
- Fully provisioned Windows® virtual desktops, rather than shared sessions on a single server
- Access to virtual desktops in the cloud from any device, rather than shared terminal service sessions that lead to additional application compatibility issues
- Lower-cost model

The Desktone DaaS platform enables enterprises to quickly realize the full potential and benefits associated with VDI environments—reduced deployment complexity; improved management, security, and compliance; and superior end-user productivity—without the capital expense and complicated systems integration of building and deploying a customized internal solution. It offers enterprises the flexibility to centrally deploy, manage, and update virtual desktops on an elastic environment for a geographically dispersed workforce.

The Desktone platform has been designed to tightly integrate with NetApp storage features to provide a highly manageable, efficient, and cost-effective cloud solution. Desktone leverages key NetApp

technologies, such as NetApp FlexClone, deduplication, Flash Cache, and more, in order to provide the performance and efficiency criteria mandated by today's VDI-based infrastructure deployments.

Desktone's technical differentiators, including its patents for cloud-hosted desktop, provisioned virtual computing, multi-tenant data center design, and virtual computing services network, among others, come to life in the Desktone DaaS platform.

They can be seen in its four major components:

- **Resource manager**. The resource manager is used by service providers to manage both infrastructure capacity and tenant entitlement. Providers can define the desktop options available to tenants for purchase, such as Windows 7, Windows Server® (Hosted Personal Desktop), Windows XP, or Linux®; how much memory and CPUs a particular virtual desktop has; and whether it is persistent or dynamic (and whether the customer will be allowed to choose between those options). You then manage the capacity, including the ratio of virtual to physical machines and the amount of memory overallocation, for each physical host.
  - This is also where you choose whether to share a host or dedicate compute to the virtual desktop. This unique feature is what makes it possible for you to offer Windows 7 virtual desktops while remaining in compliance with Microsoft. Only Desktone supports both models: sharing hosts for Windows Server (Hosted Personal Desktop) and Linux desktops or dedicating a host for a cloud-hosted Windows 7 desktop offering.
  - Based on the desktops that a tenant chooses, the Desktone platform also calculates how much of a resource, such as storage and compute, you need to meet the tenant's needs. You will see whether you have enough compute assigned to the tenant and enough capacity in the service grid overall to meet tenant needs.
  - The resource manager leverages NetApp FlexClone technology to rapidly deploy new virtual desktops on thin-provisioned storage without affecting the CPU and memory utilization of the host platform, making the desktop provisioning process extremely fast and very efficient.
  - Using the resource manager, you provision the tenant: register the tenant, provide a unique password, upload a custom look and feel, and assign a unique VLAN. The Desktone platform then auto configures an access manager and desktop manager specific to the particular tenant.
- **Access manager**. The access manager is the front end that is used by the tenants to manage the connection broker for mapping users to desktops. It provides the portal access by which the customer's users connect to their desktops.
- **Desktop manager**. Each customer uses its own desktop manager to provision pools of desktops and to continuously monitor and manage the desktop session state.
- **Service grid backbone**. This resides between the service provider (that is, resource manager) and tenant (that is, desktop manager and access manager) layers. This technology, which is unique to the Desktone DaaS platform and critical to service provider success, has the resource manager living in both the service provider network and the service grid backbone. Likewise, the access and desktop managers live in both the tenant network and the service grid backbone. Network packets cannot cross the network, but the address space used by the service provider and the tenant can be identical without any issues. Desktone built this as a nonroutable network so that both the service provider and the tenant can run the same subnet IP ranges without conflict.

  The service grid backbone also contributes to the Desktone platform's comprehensive security. Service providers and tenants cannot bridge into each other's networks, but they can still communicate with each other because the tenant appliances (access manager/desktop manager) and the service provider appliance (resource manager) are dual-homed on the nonroutable network.

**Figure 11) Desktone DaaS platform (graphic supplied by Desktone).**

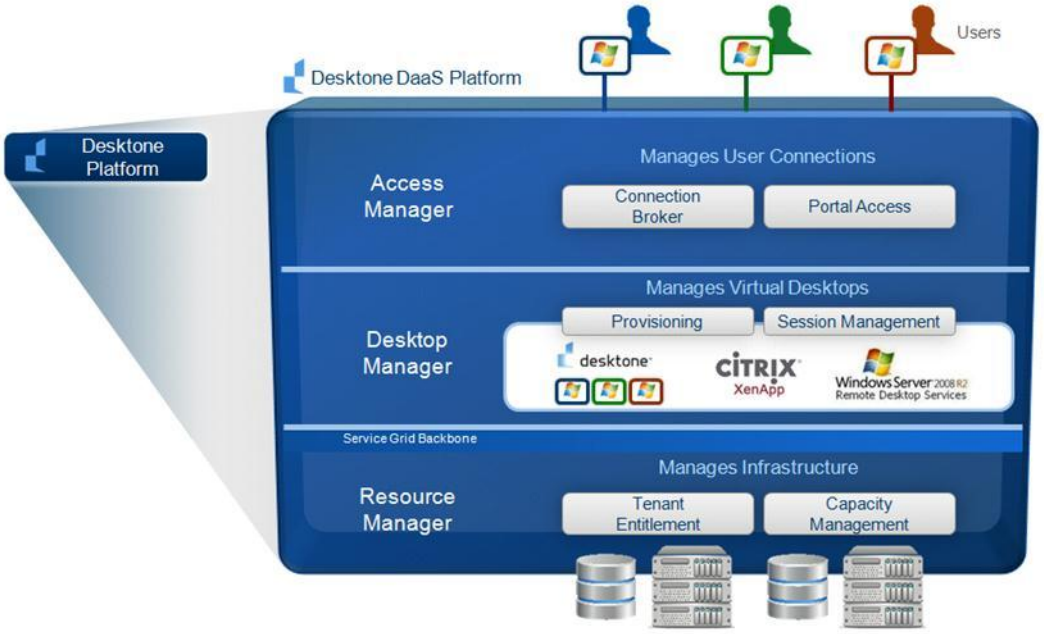


The DaaS environment can be broken down into four key elements: compute, storage, network, and the Desktone platform. Through our patented assembly of these resources managed by the Desktone platform, your DaaS solution can scale to hundreds of thousands of virtual desktops across hundreds of tenants in multiple data centers around the world:

- **Desktone management hosts**. An HA pair of physical machines that run a hypervisor and host multiple Desktone management virtual appliances for the service provider and tenants.
- **Virtual desktop hosts.** Physical machines that run a hypervisor and host tenant desktop VMs. The Desktone platform allows for sharing of a virtual desktop host between multiple tenants; however, Microsoft licensing for desktop operating systems prohibits this configuration. If tenants are running a Linux OS or using Windows Server (for which there is SPLA licensing available), then sharing of the virtual desktop host across tenants is permitted.
- **Storage system.** Clustered Data ONTAP with block- or file-based storage access.
- **Networking.** The network must support VLAN tagging, or alternatively distributed virtual networking (also referred to as DVS) can be used in conjunction with VMware vCenter. A unique network should be defined for the management network (containing the hosts and storage systems), the service provider network (literally, an extension of the service provider's network into the Desktone data center), the Desktone management network (referred to as the backbone link local network), and one or more isolated networks for each tenant.
- **Desktone appliances.** Virtual servers that live on the Desktone management hosts and support the Desktone DaaS on NetApp design elements.

  This section describes the different elements of the Desktone DaaS on NetApp solution. Desktone DaaS on NetApp is not a rigid configuration; the customer can build an infrastructure that includes best-in-class technologies from any qualified hardware platform on both the VMware and NetApp hardware compatibility list.

# 3 Key Components of the Desktone Platform

There are logical, virtual, and physical constructs in this solution. In this section we define the logical and virtual constructs in the context of this solution, including data centers, Desktone management hosts, tenant virtual desktop hosts, Desktone virtual appliances, tenants, gold images, virtual desktops, desktop pools, Desktone portals, compute platform, networking components, and NetApp storage technologies.

Figure 12 illustrates the Desktone DaaS logical constructs.

**Figure 12) Desktone DaaS logical view (graphic supplied by Desktone).**

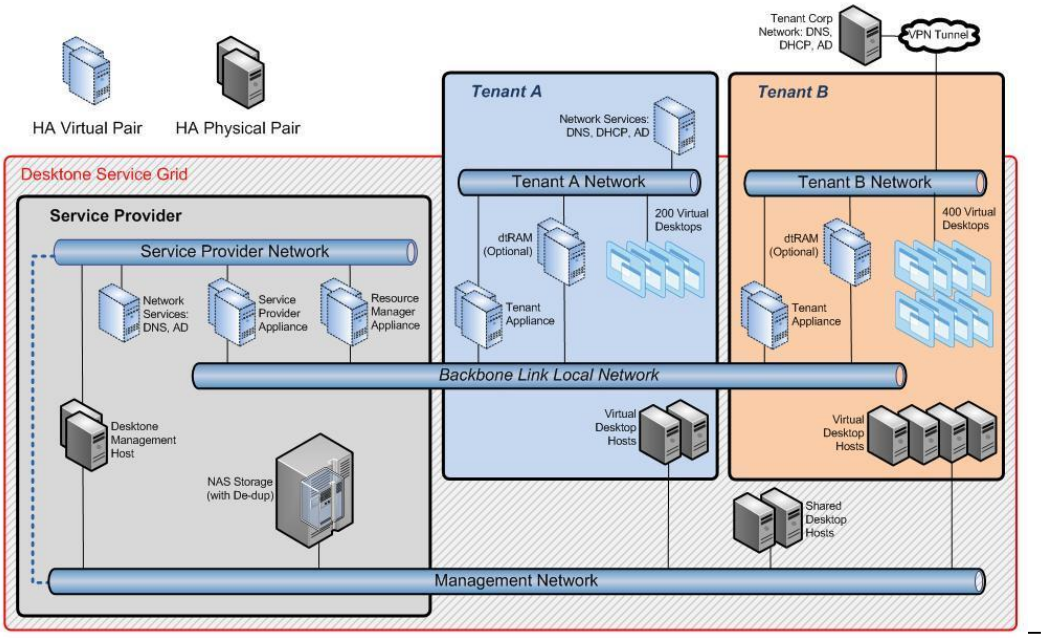


## 3.1 Desktone Components

### Data Centers

In the Desktone DaaS environment, data centers are a logical representation of the physical data center that houses tenant compute resources.

### Cross–Data Center HA

If you configure a tenant in two data centers, there needs to be a plan for failover to a backup data center in the event of the failure of one data center. The simplest solution is to redirect users to the backup data center by changing the DNS record for the portals. More advanced solutions are available, such as configuring the F5 BIG-IP global traffic manager to perform the failover automatically. The requirements for HA have to be determined before choosing the best solution.

### Desktone Hosts

Desktone management hosts are physical servers required to support the Desktone platform. These physical hosts are typically deployed in pairs and host the Desktone appliances. These hosts are used by both the service providers (service provider appliances) and tenants (tenant appliances) to provision and manage the entire DaaS environment.

Desktone virtual desktop hosts are physical servers required to support virtual desktops for the DaaS tenants:

- **Desktone management host (service provider and tenant).** HA pair of physical machines that run VMware ESXi and host the Desktone management virtual appliances that constitute the Desktone platform.
- **Virtual desktop host (tenant).** Physical machines that run VMware ESXi and host tenant desktop VMs.

  **Note:** The Desktone platform allows sharing of a virtual desktop host between multiple tenants; however, Microsoft licensing for desktop operating systems prohibits this configuration for tenants using Windows XP, Windows 7, or Windows 8. If tenants are running a Linux OS or Windows Server desktop, then sharing the virtual desktop host across tenants is permitted.

Although both types of hosts support virtualized servers or desktops, the optimization of each of these hosts is somewhat different.

## Desktone Appliances

A Desktone appliance is a virtual machine combined with a functional unit of Desktone software within the Desktone platform. The appliances are separated into two tiers: service provider and tenant.

Figure 13 shows the relationship between Desktone management hosts and Desktone appliances in a multi-tenant environment.

**Figure 13) Relationship between Desktone management hosts and Desktone appliances in a multi-tenant environment (graphic supplied by Desktone).**

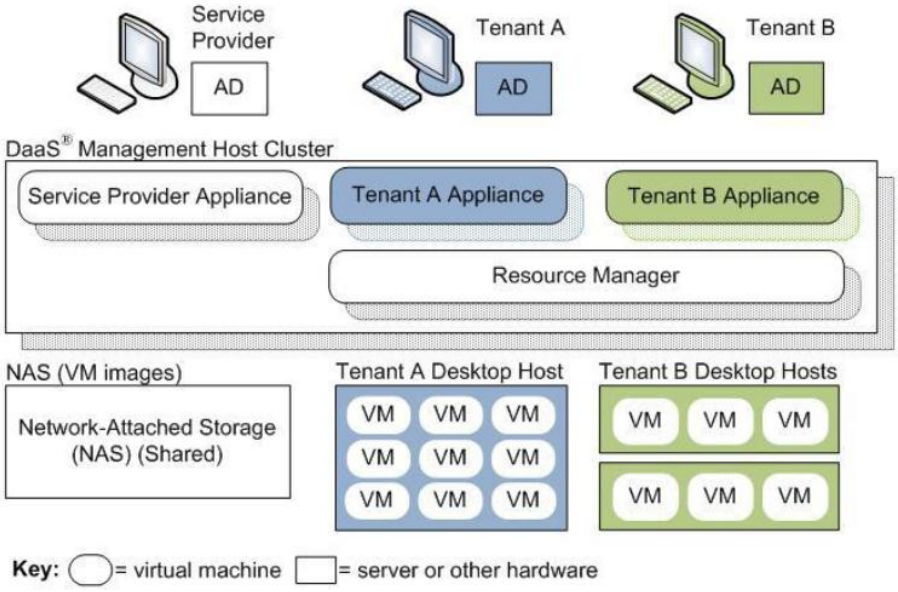


The following Desktone management appliances are virtual machines that are used to control and run the Desktone platform:

- **Service provider appliance**. Provides two types of access to the system: through the service center web-based UI and as a transit point for enabling SSH access to all the management appliances in the data center. The service provider appliance is the first appliance installed in a data center and, once bootstrapped, provides the foundation to install the remainder of the Desktone application.
- **Resource manager appliance**. A resource manager appliance integrates with the hypervisor and storage infrastructure in a given data center. A single resource manager appliance can be shared across multiple tenants. The resource manager abstracts the specifics about the desktop infrastructure from the tenant appliances.

- **Tenant appliance**. Provides the tenant with both end-user and administrative access to the tenant's virtual desktops. End users can access and manage their individual virtual desktops using the Desktone portal. Administrators have the ability to create and manage their virtual desktops through the enterprise center. All Desktone appliances are derived from a common template deployed from an OVA file that has 2GB RAM and a 20GB HDD. Desktone recommends one pair of HA tenant appliances per DC/tenant/5,000 users. Additional tenants are added to the data center by adding additional tenant appliances. For specific sizing guidelines for tenant hosts and management hosts, refer to the Desktone sizing module.

## Tenants

The service provider provisions tenants and allocates hosts to the tenants in one or more data centers. Tenants have administrative control of the gold images that are used to provision virtual desktops that are within pools for consumption by users. Multi-tenancy and secure separation enable tenant administrative domains and resources to be independent of each other. There is also a new feature called a supertenant, or shared tenant, where an MSP could partition off users/desktops within one tenant from each other, allowing the MSP to sell desktops and services to very small organizations in a cost effective manner.

## DaaS Agent

The Desktone DaaS agent runs on virtual desktops within the tenant environment. It facilitates the desktop login process, virtual desktop cloning, and virtual desktop power operations.

## Gold Pattern

A gold pattern is a nonexecutable desktop pattern that includes the operating system, applications, and default resources for the desktop. This type of pattern is used as a master template when creating pools of virtual desktops. Tenants create gold patterns and assign them to pools. Every tenant environment must contain at least one gold pattern, but many will contain several, depending on the requirements of the tenant's user community.

## Virtual Desktop

A virtual desktop is a virtual machine with one of the following operating systems: Linux, Windows XP, Windows 7 32-bit, Windows 7 64-bit, or Windows Server 2008. The desktop is accessed through the desktop portal. Tenants provision virtual desktops from gold patterns when they create or expand pools.

## Desktop Models

Desktop models are used to define the specs and protocols assigned to a desktop. The specs include the amount of RAM; the number of CPUs; and, sometimes, the amount of storage space. The protocols available are RGS, PCoIP, NX, and HTML5. All of the desktops in a pool have the same model. Gold pattern configuration is modified to the hardware specification of the model as part of the provisioning process.

Table 1 is an example or a rough guideline for what a desktop model library might look like for gold, silver, and bronze service levels.

**Table 1) Example template library.**

| Desktop Model | Specifications | IOPS |
|---|---|---|
| Gold model | 2 vCPUs, 4GB RAM, 60GB disk with 12GB unique data | 20 |
| Silver model | 1 vCPU, 2GB RAM, 30GB disk with 6GB unique data | 15 |

| Desktop Model | Specifications | IOPS |
|---|---|---|
| Bronze model | 1 vCPU, 1.5GB RAM, 20GB disk with 2GB unique data | 10 |

## Pools

A pool is a logical unit within the context of a tenant that can span multiple geographic locations. A pool is normally created based on criteria such as worker type, department or business unit (for example, finance, HR, engineering, help desk, call centers, senior management, and so on), or notable attribute groupings such as a service level (for example, platinum, gold, silver). As an example, the platinum pool could consist of static desktops targeted at senior management, whose expectations are high; the desktop pools could be configured with hourly Snapshot copies and regular off-site replication. Consequently, a bronze pool could simply consist of dynamic desktops with no Snapshot copies or replication defined; this pool would generally be targeted at a less business critical use case such as temporary workers.

As stated earlier, most businesses have a mix of worker types (for example, finance, HR, engineering, help desk, call centers, senior management, and so on). Each use case might require one or more ways in which the virtual desktops can be provisioned, managed, and protected. Based on the requirements and mix of user profiles, companies might choose to implement several functional (desktop) pools.

The following sections provide details about the pool types and user assignment types and validate how the NetApp solution strongly complements them to achieve the service provider's ultimate goals and objectives. Note that data such as guest OS paging, temp files, and vSwap should not be deduplicated because these files contain a considerable amount of transient data that is created and destroyed frequently.

Three pool types are available for Desktone tenants: static pools, dynamic pools, and session-based pools. A pool is defined to manage a collection of virtual desktops; a typical data center will have a mix of static and dynamic pools deployed to support individual tenant desktop requirements. The definitions of the different pools are as follows:

- **Session (static) pool.** This consists of virtual desktops that are assigned to individuals. The first time a persistent user logs in, that user is allocated an available virtual desktop from the pool. After that time, that virtual desktop is dedicated to that user for future use. The number of virtual desktops in the pool should equal the number of users assigned to the pool if all users have one-to-one mappings. It is worth noting that a single static virtual desktop can be assigned to multiple individuals if needed on a shared desktop assignment basis.
- **Dynamic pool.** This consists of virtual desktops that are assigned on an as-needed basis. A user is never assigned to a specific virtual desktop; a desktop is essentially floating, and the end user is assigned to the dynamic pool, logging in to any available virtual desktop from that pool. The state of the virtual desktop in a dynamic pool can be recycled to a predefined state between sessions. The dynamic, nonpersistent pool defines the number of users who can be connected concurrently. Therefore, the number of users assigned to the pool can exceed the number of virtual desktops available in the pool.
- **Session-based pool.** Session-based pools allow users to establish a session within a Windows Server instance that is acting as an RDS server. This allows multiple/many users to connect to the same Windows Server instance at the same time as in traditional RDS deployments.

  The static and dynamic pool types offer two options in which virtual desktops can be provisioned by using hardware-assisted cloning, that is, NetApp FlexClone using Virtual Storage Console (VSC); and/or by using hypervisor-based cloning capabilities such as VMware ESX clones (in the case of FC or iSCSI datastores). The option for cloning is not user selectable; the Desktone storage manager determines the storage system in use. If a NetApp system is detected with FlexClone licensed, Desktone automatically utilizes the NetApp VSC rapid cloning capability for provisioning. If a NetApp system is detected but the FlexClone license is not installed, then the basic hypervisor cloning

capability is used, exactly as it would be if generic storage other than NetApp is in use. Leveraging FlexClone to provision and clone Desktone appliances and tenant desktop images offloads this workload to storage, eliminating the CPU and memory hit on the compute platform and the network traffic that is required during traditional cloning processes.

**Note:** Desktone still does not support NetApp VSC rapid cloning on FC and iSCSI datastores; native VMware cloning is used on FC and iSCSI datastores.

Both the static and the dynamic pool types require an element of preprovisioning. This means that for a user to log in to a virtual desktop, the actual desktop must already exist in the virtual desktop pool. Because virtual desktops are created through cloning operations on a deduplicated volume on NetApp storage, this preprovisioning causes minimal overhead on the storage system because the desktop image essentially consumes no additional space or I/O until used.

For a static pool, virtual desktops are precreated and assigned to individuals as the users log in. With a predefined number of virtual desktops available in a static pool, up to that number of desktop users have access to a static virtual desktop assuming a one-to-one user desktop assignment.

With dynamic pools, virtual desktops also need to be preprovisioned. However, they can be powered on demand as new users log in. A dynamic desktop pool has the ability to define a desktop overhead in which, as an example, an overhead of 10 powered-on desktops may be defined in the pool for fast user logon. This overhead of powered-on virtual desktops is managed dynamically by the Desktone platform. In addition to having the capability to dynamically power on virtual desktops, the capability exists to power down desktops should demand for virtual desktops in the dynamic pool decrease.

## Desktone Portals

The Desktone platform presents three browser-based graphical user interface portals:

- Service Center
- Enterprise Center
- Desktop portal

### Service Center

The Service Center is used by the service provider administrators to manage the data center resources, such as hosts, storage, and the Desktone management appliances. The Service Center also enables the management of tenant contracts defining tenant models and quotas as well as the configuration of tenant appliances and networks.

The Service Center supports creating and assigning additional roles and permissions among the service provider administrators to securely distribute management tasks among larger organizations.

### Enterprise Center

The Enterprise Center is used by the enterprise administrators (tenant administrators) to manage their virtual infrastructure. Each enterprise has its own customizable Enterprise Center portal. Enterprise administrators can provision both static and dynamic pools of desktops based on templates they have customized or new templates they may upload. Enterprise administrators can also add additional domains and map groups or individuals to either specific virtual desktops or pools.

The Enterprise Center supports creating and assigning additional roles and permissions among the service provider administrators to securely distribute management tasks among larger organizations.

### Desktop Portal

The desktop portal enables individual users to connect to their virtual desktops. Every tenant has its own customizable portal. Users log in to the portal and have the option of being directly connected to a

desktop they have defined as their default or presented with a list of available desktops and enabled to choose to which virtual desktop to connect. Users can also set default protocols per VM and additional protocol customizations. Users can connect to the desktop portal from a variety of clients, including thin clients (both WTOS-based WYSE clients and any thin clients running Windows Embedded), thick clients (such as PCs running Windows, Mac® OS, or Linux), as well as iOS- and Android-based mobile devices.

The desktop portal facilitates connections using a wide collection of remote protocols:

- **RDP (Microsoft).** Microsoft's Remote Desktop Protocol is a very strong protocol with broad support. The protocol supports a good multimedia experience with less than 20ms of latency and a good user experience when using office productivity apps when latency is under 50ms.
- **PCoIP (VMware).** The PCoIP experience provides a very good multimedia user experience in situations with both high latency and constrained bandwidth.
- **RGS (HP).** Remote graphics software developed by HP primarily for WAN deployments provides good multimedia support with latency as high as 100ms when provided with ample bandwidth. RGS is particularly popular for graphics-intense use cases such as CAD.
- **HTML5 (Ericom).** This client allows users to access their desktop through any HTML5-compatible web browser. It does not require any additional plug-ins, add-ons, or installation of any kind on the end-user device, which makes it suitable for devices such as Chrome OS netbooks.
- **NX (Linux).** This protocol enables access to Linux desktops.

# 4 Compute Resources

Compute resources refers to the physical servers necessary to support the Desktone platform and the software required on those hosts. Desktone management appliances and desktop virtual machines cannot reside on the same physical server. Separate servers must be used for the following:

- Desktone management host (service provider)
- Virtual desktop host (tenant)

Although both types of hosts support virtualized servers or desktops, the optimization of each of these hosts is slightly different. Therefore the process for sizing each server is defined separately in the following sections.

## 4.1 Hardware Requirements for Desktone Management Hosts

The management hosts are a pair of physical machines that contain the Desktone management appliances (both service provide and tenant appliances). Several sample profiles are defined here; after a server is full, you can simply add additional management hosts to the platform.

### Desktone Appliance Sizing Requirements

Table 2 shows the prescribed sizing requirements for Desktone appliances.

**Table 2) Desktone appliance sizing requirements.**

| Appliance | Template (Memory/Disk Space) | Sizing |
|---|---|---|
| Service provider appliance | Standard (2GB/20GB) | 1 pair/dc |
| Resource manager appliance | Standard (2GB/20GB) | 1 pair/dc/20,000 VMs |
| Tenant appliance | Standard (2GB/20GB) | 1 pair/dc/tenant/5,000 users |
| dtRAM appliance | FreeBSD (512MB/8GB) | 1 pair/dc/tenant |

The smallest environment begins with two management hosts, each with one service provider appliance, one resource manager, and one tenant appliance. From there, additional tenants are added to the data center by adding an additional tenant appliance to each management host. The size of the management host is generally referred to by the number of tenants it can support.

### Sizing a Management Host for a Specific Number of Tenants

There are three variables to consider when determining the hardware configuration for a Desktone management host:

- **CPU.** Each core supports 10 tenants. For example, four cores are required to support 40 tenants.
- **Memory.** Each tenant requires 2.5GB of RAM on each Desktone management host (2GB for each of the tenant appliances and 0.5GB for each of the dtRAM appliances). For example, 125GB of RAM on each host is required to support 50 tenants.
- **Storage.** Each tenant requires 56GB of storage, 28GB allocated to each Desktone management host (20GB for each of the tenant appliances and 8GB for each of the dtRAM appliances). For example, a pair of management hosts that can scale to 50 tenants requires 1400GB (1.4TB) of storage each (2.8TB total).

Table 3 defines the server hardware used for two sample Desktone management hosts.

**Table 3) Sample hardware requirements for each Desktone management host.**

| Component | Trial Environment | Production Recommendation |
|---|---|---|
| CPU | 1 CPU | 1 CPU |
| CPU architecture | 2 cores | 6 cores |
| Minimum RAM | 48GB | 128GB |
| Data disk configuration (see note) | 560GB | 1.4TB |
| Supported tenants | 20 | 50 |

**Note:** For virtual desktops (not including user home directories), we typically see over 80% to 90% savings in capacity after the NetApp storage efficiency technologies are applied. All the management appliances have similar footprints, so identical data will get deduplicated, saving large amounts of space.

## 4.2 Hardware Requirements for Virtual Desktop Hosts

Virtual desktop hosts are sized based on the number of CPU cores and amount of memory installed in the server. All tenant virtual desktops reside on shared storage. Therefore, desktop hosts only require a pair of small disks for the hypervisor OS installation.

### Sizing a Desktop Host for a Specific Number of Desktops

The guidelines for sizing the memory and CPU for a desktop host are:

- **CPU.** Desktone recommends setting a 10x ratio for CPU; that is one physical core for every 10 virtual CPUs. Virtual to physical CPU ratios will vary based on use case.
- **Memory.** Desktone recommends setting a 1.5x overcommit ratio for memory; that is you need 32GB of physical memory for each 48GB of virtual memory allocated for desktops. Memory overcommit ratios can vary based on the workload on a host. The more similar the workload, the higher the memory overcommit ratio can be. For example, if a host has all Windows 7 SP1 VMs, then page sharing will be optimal, and therefore the memory overcommit can be high.

### Host Sizing Calculations

Two formulas are helpful when sizing your hosts. There are four variables for each formula, and you can solve for any of the four variables, so long as you know three. The variables are the number of VMs on the host, the amount of memory or CPU cores assigned to each virtual desktop, the amount of physical memory or CPU cores installed in the host, and the overcommit ratios for either the memory or CPU cores.

### Memory

[Number of VMs] x [virtual memory per VM] <= [physical memory] x [memory overcommit ratio]

### CPU

[Number of VMs] x [virtual CPUs per VM] <= [number of physical cores] x [CPU overcommit ratio]

Table 4 defines the server hardware used for the sample virtual desktop hosts. The table assumes a virtual desktop that consists of 2GB of RAM and one virtual CPU. This example uses a 10 x CPU overcommit ratio and a 1.5 x memory overcommit ratio.

**Table 4) Minimum hardware requirements for virtual desktop host.**

| Number of Desktops | Cores Required | RAM Required (GB) |
|---|---|---|
| 20 | 2 | 32 |
| 40 | 4 | 64 |
| 60 | 6 | 96 |
| 80 | 8 | 128 |
| 120 | 12 | 192 |

### CPU Speed Considerations

As the diversity of server processors continues to expand, a simple core ratio might not be appropriate for all situations, particularly in the case of very fast or very slow processors. The Desktone 10x recommendation applies best to CPUs running around 2.4 GHz. Some hardware vendors have begun suggesting megahertz-based sizing. To size desktops to a server using megahertz-based sizing, simply multiply the clock speed of the processor by the number of cores, then divide by a per user allocation. Desktone recommends allocating at least 250 to 300 megahertz per desktop.

As with host sizing equations, the equation has four variables; you can solve for any of the four variables, so long as you know three.

[Number of VMs] x [MHz allocated per desktop] <= [number of physical cores] x [CPU clock speed]
Example: How many desktops can I host with a single 1.9Ghz six core CPU?

Number of VMs = 6 x 1900 / 250 = ~45

## 4.3 Compute Resource Software Requirements

The Desktone platform manages all compute resources virtually. Therefore, a hypervisor must be running on the management and desktop hosts. The Desktone platform supports the VMware vSphere® hypervisor. There are two options in how vSphere is managed: either directly or through VMware vCenter. There are several considerations when determining how to manage your vSphere hosts. Some of these are:

- **Ease of management**. Using vCenter simplifies the management of vSphere because many vSphere hosts can be aggregated under a single vCenter instance.
- **Scale and performance**. Introducing an additional management layer could negatively affect scaling and performance. There is additional work being done by vCenter, and it is not as responsive as the hypervisor directly even when given significant resources (CPU and memory). The limits on host and VM counts could potentially affect the scalability of an individual tenant. At a minimum, using vCenter might require more desktop managers because only a single vCenter instance can be assigned to a given desktop manager, and the scalability of a desktop manager can be greater than a vCenter instance.
- **Cost**. vCenter introduces additional licensing considerations. At a minimum there is a license for vCenter itself and the Windows Server license on the system where vCenter is installed. There are also CPU and memory resources required to run vCenter.
- **Monitoring**. Many service providers already have existing investments in vCenter for the purposes of monitoring, which they would prefer to leverage with the Desktone platform.
- **Reporting**. vCenter provides additional reporting and retention over and above what can be obtained from the hypervisor directly.
- **High availability**. There are pros and cons regarding high availability when it comes to using vCenter. On the positive side, a key feature of vCenter is host clustering, which enables automatic failover of VMs in case of a failure (see the following section regarding vCenter integration for information about vCenter host clustering). This is more relevant to virtual desktops than management appliances because HA is built into the Desktone platform independent of the hypervisor manager. On the negative side, vCenter introduces an additional layer between the Desktone platform and the hypervisor. Should the vCenter go down, the platform would be unable to provision or perform other important functions, such as power operations.
- **Storage**. There are additional storage options available with vCenter. Specifically, using iSCSI and FC storage is possible. Additionally, the VM cloning speed can vary depending on the type of storage being used. There are specific integrations with several NFS storage systems when using ESXi directly that significantly accelerate clone times. On the vCenter side there is integration with the NetApp VSC, which leverages the Rapid Cloning Utility (RCU) for fast cloning.
- **Networking**. If you want to use distributed virtual networking such as the vSphere Distributed Switch or Cisco Nexus® 1000V, you must use vCenter with the Desktone platform. When using distributed virtual networking, vCenter is the network controller. Therefore, any API calls made by the Desktone platform to configure the virtual network on a VM or otherwise must be done to vCenter.

The decision of whether to use vCenter or manage vSphere hosts directly is one that should be made prior to installation of the platform. As a general rule, management of the vSphere hosts should be consistent with your operational practices. This means that if you are managing your vSphere hosts with vCenter, then vCenter should be used by the Desktone platform as well. There are known issues related to discovering vSphere hosts directly while they are connected to vCenter. For example, it is not possible to mark a virtual disk nonpersistent on a vSphere host directly when that host is under vCenter management.

## Native ESXi Integration

The Desktone platform has been optimized to communicate directly to the ESXi hosts and not use vCenter. Due to the lack of a cloning API on ESXi hosts, the Desktone platform takes ownership of VM manufacturing. This necessitates access to the file system of the datastore and is why NFS storage is required. When managing ESXi hosts directly, the Desktone platform officially supports three major versions of ESXi at any one time. When a new version is supported, it means that the oldest version is no longer supported. The currently supported versions of vSphere are ESXi 4.1, 5.0, and 5.1.

### vCenter Integration

When managing ESXi hosts through vCenter, the Desktone platform officially supports two major versions of ESXi at any one time. The currently supported versions of vSphere are vCenter 5.1 and ESXi 5.0 and 5.1. There are several advanced features of vCenter that are currently untested and, as a result, are unsupported. These features include, but are not limited to, vMotion®, DRS, and host clustering. Although these features might appear to work, they can cause problems with things such as host quotas that are tracked by the Desktone platform. These features are being considered for future Desktone platform releases. When in doubt about a vCenter feature, contact Desktone Support.

**Note:** Desktone does not support VMware clustering. The Desktone ESXi hosts should not be clustered; otherwise, VM provisioning will fail.

**Note:** Desktone also does not support direct ESXi connectivity in a NetApp configuration. The Desktone hosts must be managed by vCenter, and the NetApp VSC plug-in for VMware must be installed and configured.

# 5 Network Resources

There are two key components to the network to assure tenant separation when assembling the Desktone platform: VLAN tagging and VRF support. In the Desktone environment, a tenant network is **not** a subnet within the SP network; each tenant network is a logical extension of the tenant network existing in the SP data center.

## 5.1 Desktone Networks

Figure 14 shows the distinct networks within the data center:

- The link local backbone network is fully controlled by the SP. This network is a link local nonroutable subnet (169.254.0.0/16) that is logically separated from all tenant networks. The backbone network connects all Desktone management appliances. For example, the TenantA appliance connects to the SP resource manager through the link local backbone network.
- The management network is used to segregate all the physical hosts and storage systems.
- The SP network is an extension of the service provider's network into the data center. The Service Center is accessed through the service provider appliances on the service provider network. The SP VLAN must have access to the management VLAN for access to virtual desktop hosts and the storage systems.
- The tenant networks are fully controlled by the tenants, again as a discrete VLAN that is separate from the SP and other tenant networks. The tenant network connects the tenant appliances to a tenant's virtual desktops. The tenant VLAN is not accessible to the SP.

Figure 14 emphasizes the clean separation of SP and tenant networks. The area labeled Desktone service grid provides the core of the Desktone platform. The portion of the diagram directly connected to a tenant network represents the components of the system that are duplicated for multiple tenants.

Note the following network architecture:

- The tenant networks are not a subnet of the SP network. It is a logical extension of the Tenant A or Tenant B network existing in the SP data center.
- High availability indicates redundant pairs to make sure failover integrity. Most of the computing components are pairs of virtual machines (shown in light blue).
- The number of physical servers in a tenant network largely depends on the size and number of virtual desktops the tenant is hosting.

**Figure 14) Desktone logical view (additional tenants).**

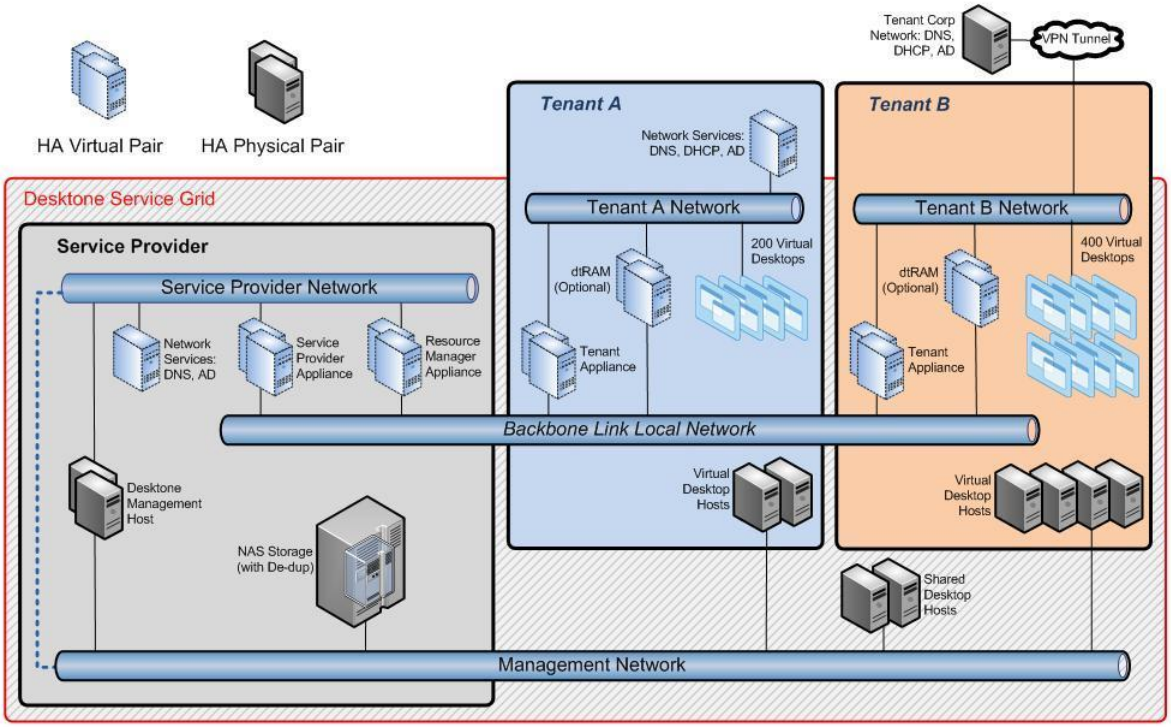


## 5.2 Virtual LANs in the Desktone Platform

A VLAN is an emulation of a standard LAN that allows data transfer to take place without the traditional physical restraints placed on a network. An understanding of the use of VLANs is important to planning and implementing the Desktone platform due to their role in making sure of separation of tenants and SP, optimizing the performance of data and management information flows, and in increasing the scalability of the Desktone platform.

### Distributed Virtual Networking

If you are using VMware vCenter, an alternate network configuration option is available by using distributed virtual networking as opposed to using the local vSwitch with VLANs. Distributed virtual networking allows the network configuration to be done in a software appliance-based controller rather than in the physical hardware device (that is, the router). Two distributed virtual networking options are available with the Desktone platform. They are VMware vSphere Distributed Switch (VDS) and Cisco Nexus 1000V. Refer to the product documentation for further information.

### Layer 2 Segregation Using VLANs

VLANs provide segregation of traffic at layer 2 to prevent two tenants from seeing each other's traffic while still sharing the same physical network path. However, in order to reach the end user, this traffic must pass through a router. Without additional segregation at layer 3, customers would be able to route to each other or could have a nonresolvable conflict of IP addresses.

Figure 15 illustrates layer 2 segregation with two tenants. In a typical environment, a multilayer switch is required. Note that although four physical switches and two routers are shown in Figure 15, only one of each might be needed; virtual switches and routers can be used. The VLANs are trunked. In addition, each of the virtual desktop hosts and the Desktone management host supports multiple virtual machines.

Figure 15) Layer 2 segregation using VLANs.

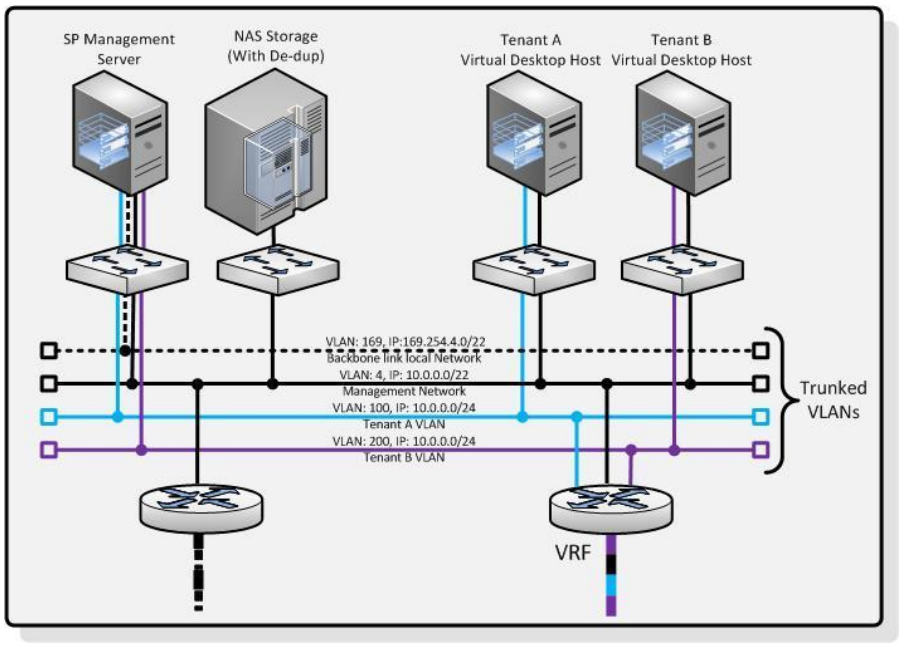


## Layer 3 Segregation Using VRFs

Virtual Routing and Forwarding (VRF) allows multiple instances of a routing table to coexist within the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other.

VLANs in a Desktone installation are aggregated into VRFs. After being aggregated into a VRF, segregation is handled at layer 3, and the VLAN IDs are effectively discarded. This means that VLANs require uniqueness only under one VRF, and the same VLAN IDs can be reused under multiple VRFs within a single data center.

## 5.3 Networking Resources

Switches, routers, load balancers, and gateways play a role in any Desktone installation. This section lists the characteristics required of each of these components. Desktone is hardware agnostic, only requiring that equipment supports the characteristics listed in this document.

For example, any layer 3 network devices, including firewalls or routers, installed between the customer site and the customer's virtual desktops at the SP's data center must meet one of the following requirements, in order of preference:

- Supports multiple independent routing tables (VRFs)
- Be dedicated to the customer and support out-of-band management
- Be dedicated to the customer and managed in band from the customer's network

Not meeting at least one of these requirements can result in IP address conflicts between the SP and the tenant.

## Switches

Switches must support trunking of VLANs. Here are some guidelines regarding the connectivity requirements for the switches:

- Connectivity between the desktop hosts and the storage should be a 10GB Ethernet network. For additional bandwidth, you could aggregate the two links to achieve 20GB.
- Connectivity for tenant/protocol traffic can be 1GB connections.

### Routers

Tenant networks are VLAN tagged; for uniformity of management, SP networks are also VLAN tagged. Routers must support VRFs. If there will be customers with VPN access back into their corporate network for access to network services such as DNS/AD/DHCP or for access to other applications, then the router must have the ability to tie that VPN tunnel to the VRF for a tenant.

### Load Balancing

Load balancing is only required in front of the tenant appliances for a tenant with 25,000 or more desktops in a single data center. Load balancing is required in front of the dtRAM appliances for a tenant with 5,000 or more concurrent external desktop connections in a single data center.

### Gateways

Gateways must support VRFs.

### Cross–Data Center HA

If you configure a tenant in two data centers, there should be a plan for failover to a backup data center in the event of a failure of one data center. The simplest solution is to redirect users to the backup data center by changing the DNS record for the portals. There are more advanced solutions available, such as configuring the F5 BIG-IP global traffic manager to perform the failover automatically. The requirements for HA have to be determined before choosing the best solution.

## 5.4 Managing External IPs and Global DNS Entries

In order for a customer to access its cloud-hosted desktops from anywhere, the service provider must allocate external IP addresses to each tenant. The edge router must be able to direct traffic destined for the tenant portals to the tenant appliance shared IP and dtRAM-enabled desktop connections to the dtRAM VIP. This can most easily be accomplished by using two external IPs per tenant: the first is NATed to the tenant appliance shared IP, and the second is NATed to the dtRAM. Depending on the desired DNS name, a global DNS entry should be created for the portal IP on the service provider's DNS (such as tenant.SvcProsDesktops.com) or in the tenant's global DNS (such as desktops.tenant.com). The domain where the name is hosted matters because this also defines who is responsible for supplying the SSL certificates for the tenant.

Many routers support port-based NATing. If your router is capable of redirecting traffic based on the incoming port, you can condense the number of external IPs needed to one per tenant. In such a configuration, configure traffic destined for port 80 or 443 to redirect to the tenant appliance shared IP and dtRAM ports (typically 8001-8050) to redirect to the dtRAM VIP.

**Figure 16) External port-based NAT.**

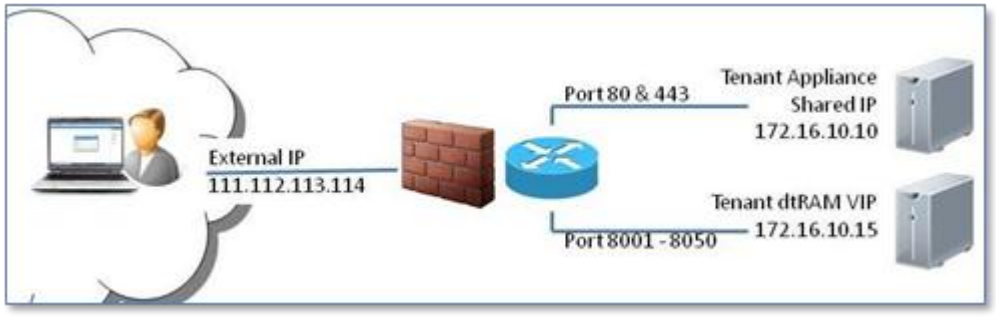


## 5.5 Wiring the Data Center

Figure 17 represents a sample data center wiring diagram. For clarity, only a sample of connections is represented to demonstrate connectivity from, for example, the desktop hosts to the top of rack switch, to the core switch, then to the NAS storage.

**Figure 17) Sample data center wiring diagram.**

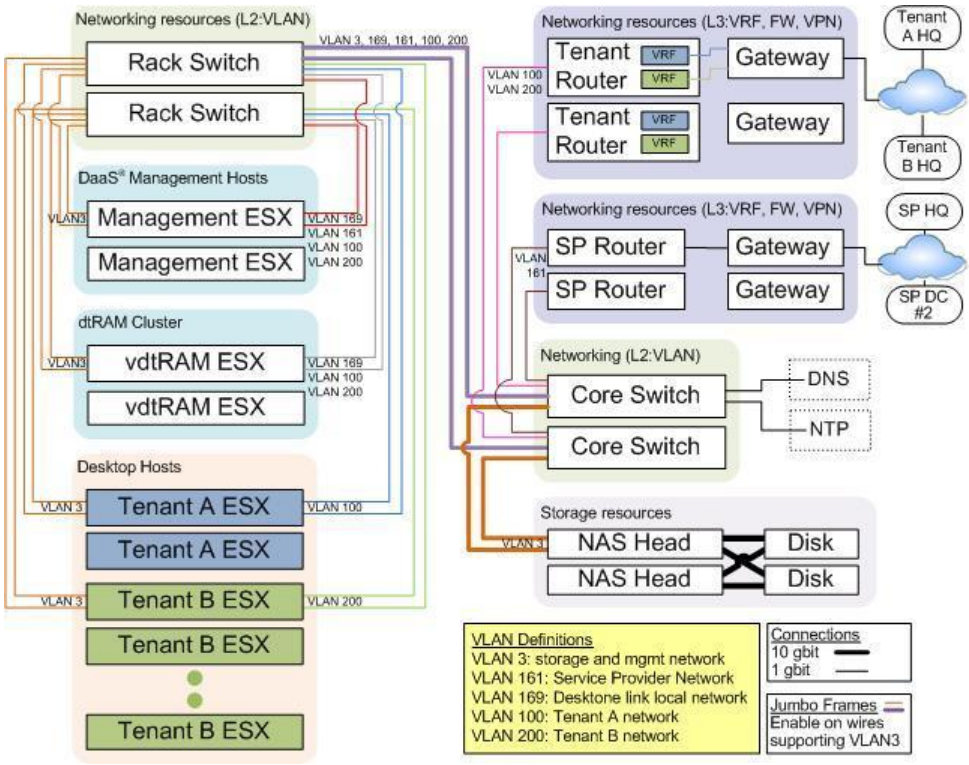


Figure 18 demonstrates an example of how to configure networking on an ESXi host for use with the Desktone platform. Your management hosts will need to have the default tenant VLANs available because the tenant appliances reside there. Your desktop hosts need to have only the specific tenant VLAN that is assigned to that host. These need to be configured manually on the ESXi hosts before you add the host to the Desktone platform.

**Figure 18) Sample ESX network configuration for both tenant and management host.**

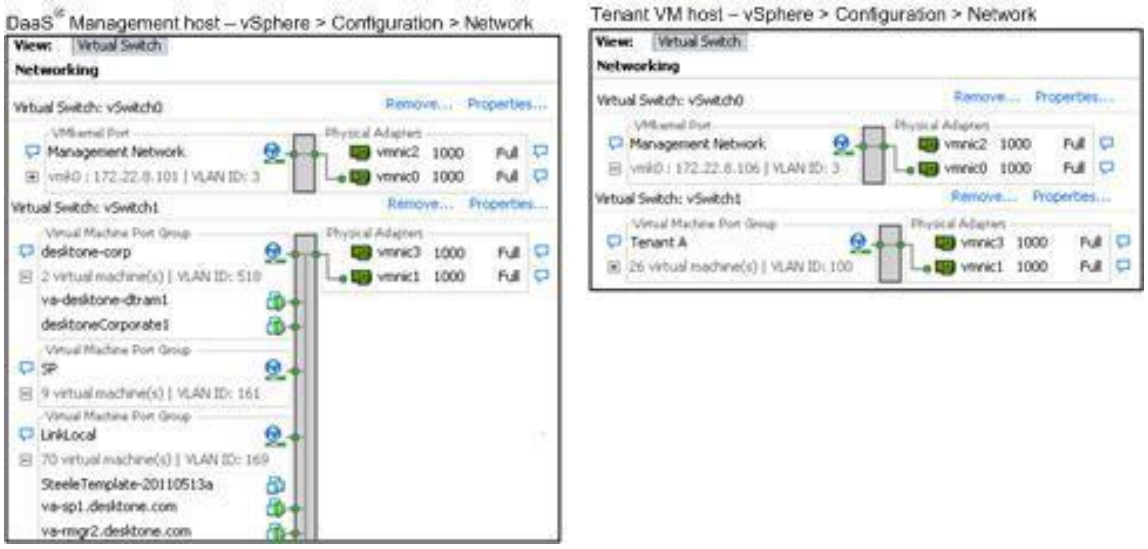


## 5.6 Networking Between Data Centers

As noted earlier, a tenant network can be separated into functional pools rather than only into geographic locations. For example, Tenant C in the example can be divided into sales, manufacturing, and finance through the creation of pools because a pool is a logical unit that can span multiple geographic locations. Virtual desktops within a pool have no awareness of where their VM is located; end users should only be concerned with their user experience. Alternatively, pools could be aligned to a specific SLA or use case (that is, stateless pool versus stateful pool).

Figure 19 shows a segment of the example network in which Tenant C is divided into organizational functions by pools rather than geographic locations. Finance and development each have staff located in both New York and London.

**Figure 19) Multiple data centers.**

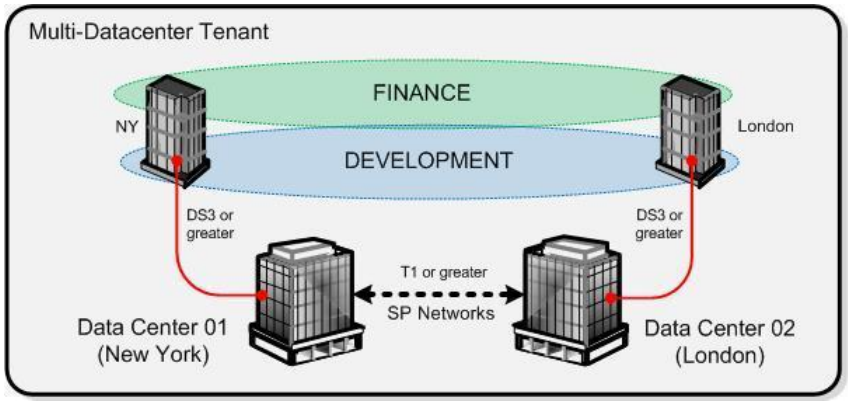


### Traffic Between Data Centers

SPs that maintain multiple data centers require network connectivity between the data centers for data sync operations and application traffic. The SP nodes in each data center need to be able to communicate with each other using their IP addresses (NAT will not work). This can be accomplished through a VPN tunnel.

The traffic between data centers tends to be bursty in nature. Experience suggests that a T1 connection between data centers is sufficient to handle all cases without a disruption of service.

Note that the Desktone management resource overhead scales independently of the number of tenant virtual desktops. Increasing the number of virtual desktops reduces the percentage of the cost of an installation that is required by Desktone management alone. The Desktone management overhead is an increasingly smaller portion of the cost of operation as data centers are scaled out with the addition of virtual desktops.

It is important to realize that the example used in this document is a set of calculation methods based on a hypothetical case in order to demonstrate the methods. Optimization of tenant virtual desktops could reduce the number of required desktops, resulting in decreases in the number of physical servers.

# 6 Storage Resources

This section of the solution guide provides a high-level overview of the components and features that should be considered when deploying Desktone infrastructure on NetApp storage.

NetApp hardware is capable of fast cloning, thin provisioning, and deduplication; typical figures for space savings with deduplication using a NetApp configuration:

- >70% typical storage reduction
- >90% with virtual desktop infrastructure (VDI)
- Saves up to 95% for full backups; 25% to 55% for most datasets

For further reading on deduplication in clustered Data ONTAP, refer to TR-3966: NetApp Data Compression and Deduplication Deployment and Implementation Guide.

For additional information on designing, deploying, configuring, and best practices for vSphere SAN and IP networks, refer to TR-4068: VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1.

## 6.1 Storage in the Desktone Solution

Storage and storage planning are critical parts of any virtualization environment. In the Desktone solution on clustered Data ONTAP, you must plan for four types of storage: FC boot LUNs for ESXi hosts and FC or iSCSI storage for tenant virtual machines, NFS storage for Desktone appliances, NFS storage for tenant virtual machines, and CIFS shares for user data.

- FC Boot LUNs are used by Desktone hosts to boot VMware ESXi and have a one-to-one relationship with hosts. The storage requirements are fairly static, as is the LUN data.
- NFS volumes for Desktone appliances are mounted to Desktone management hosts and used to house the service provider and tenant appliances.
- NFS volumes for tenant virtual desktops are mounted to each of the tenant's virtual desktop hosts in a given data center. NFS is used for storage of virtual desktops and virtual desktop patterns. Make sure your NFS datastore has sufficient I/O capacity. For planning purposes, assume an average of 12–20 IOPS per virtual desktop.
- CIFS shares are accessed by tenant virtual desktops to store user data.

  **Note:** Desktone does NOT support VSC for SAN (FC or iSCSI) storage. The Desktone platform will use the native vCenter cloning APIs and NOT VSC to provision and clone management appliances when using FC or iSCSI datastores. It is recommended to use NFS datastores in the Desktone solution on clustered Data ONTAP.

Fibre Channel boot LUNs are created through the Data ONTAP CLI or OnCommand System Manager and presented to Desktone hosts through FCoE or FC SAN. These LUNs are provisioned on a volume with data deduplication to maximize storage efficiency. FC boot LUNs are typically created as you scale out your environment, when adding new physical hosts.

The NFS and CIFS storage (detailed in the preceding section) are provisioned using the Data ONTAP CLI, OnCommand System Manager, and VSC (for NFS only). The NFS and CIFS storage is routinely provisioned as tenants are onboarded into the environment. The SVMs hosting the NFS and CIFS shares may also be provisioned through the Data ONTAP CLI or OnCommand System Manager. An alternative is to automate the entire process by using the OnCommand Workflow Automation (WFA) tool.

## 6.2 NFS Storage

Notable characteristics of this architecture are:

- NFS shares are presented and mounted on corresponding host servers.
- Isolation is accomplished by presentation of NFS shares to specific hosts.
- Multi-tenant architecture. Separate routing tables and VLANs, and different SVMs if necessary.
- Failover and VM restart done at the tenant appliance layer.
- Deduplication done at NFS storage layer with up to 95% space saving for the tenant VMs.
- VSC for VMware integration in the Desktone solution.

Desktone requires NAS for primary VM storage when discovering ESXi hosts directly. Some of the reasons for choosing NAS over SAN are:

- Each NAS mount point can be shared to the entire data center.
- There is no built-in limit to the number of hosts that can mount a single NAS mount point.
- Capacity management can be done at data center granularity.
- Failure of any individual head can be handled automatically by software. No manual intervention is required.
- Single NAS head can host multiple tenants using NFS export controls to provide firm tenant separation. You can also use different SVMs as zones of separation across tenants.

## 6.3 SAN Storage

Fibre Channel and iSCSI storage can only be used with vCenter. Local storage is not an option for storing desktop VMs. When using vCenter, all storage configuration changes are made outside of the Desktone platform. Typically, this would be done using the vSphere client connected to vCenter. One or more LUNs must be mapped to all ESXi hosts assigned to a desktop manager. The datastores associated to the LUNs must be created on all hosts for that desktop manager and have the same exact name (case sensitive). The Desktone platform will clone desktop VMs to the same datastore on which the gold pattern resides using the native vCenter cloning API. For additional information, refer to VMware documentation.

**Note:** Desktone does NOT support VSC for SAN (FC or iSCSI) storage. The Desktone platform will use the native vCenter cloning APIs and NOT VSC to provision and clone management appliances when using FC or iSCSI datastores. It is recommended to use NFS datastores in the Desktone solution on clustered Data ONTAP.

## 6.4 Tenant Data Storage

Desktone recommends not storing data in the virtual desktop. Instead, user data directories, such as My Documents folder, should be redirected to a separate storage location. The reasons for not storing the data on the virtual desktop are:

- The storage used for desktop images is highly optimized for performance: high-speed FC or SAS disk is required; other optimizations are potentially in place to improve I/O. The storage requirements and performance profile for user data are significantly different from the desktop images, and thus a different class of storage can be used.

- Protecting data that is stored in a virtual desktop requires that each of the virtual desktops is backed up individually. If user data is redirected to a CIFS share, the external storage can be better protected and maintained independently of the desktop image.

There are several options for tenant data storage. A tenant's user data can be:

- **Colocated with the desktops in the service provider's data center**. The service provider can offer an add-on service that provides storage space for user data. Typically, this would be CIFS storage as a service where the service provider offers secure tenant-specific user data containers based on CIFS shares with integration into the tenant's own Active Directory.
- **In the tenant's own data center**. Access to file shares and other data from the virtual desktops would be over the site-to-site VPN or MPLS connection between the tenant and service provider data centers. This connection is often referred to as the backhaul connection. Performance of the backhaul connection depends on the latency between the tenant and service provider data centers. In certain cases, the backhaul connection can be optimized with a WAN accelerator.
- **In the cloud**. The tenant can use a cloud storage service for storing user data.

In the NetApp solution, all the CIFS user data is redirected to a dedicated tenant SVM that is in the tenant AD domain.

# 7 NetApp Integration

## 7.1 NetApp and Desktone Validated Configuration

**Table 5) Validated configuration.**

| Component | Description |
|---|---|
| Desktone platform | The Desktone version qualified for the solution is version 5.4. |
| Supported hardware | NetApp FAS series, NetApp V-Series, and IBM N-Series. |
| NFS permissions | Desktone integrates with VSC; therefore, it is recommended to use VSC to manage and provision the datastores. VSC automatically configures the export permissions; alternatively, you can also choose to provision the exports manually and only use VSC to mount the provisioned datastores. Desktone requires that the hypervisors have root access to the appropriate NFS exports. |
| Access credentials | Desktone uses VSC to talk to NetApp storage. Configure VSC to manage the storage clusters before configuring Desktone. Specify appropriate users with the required credentials that can perform all the capabilities in VSC on NetApp storage. Refer to NetApp Virtual Storage Console for VMware vSphere Install and Admin Guide. |
| Data ONTAP | The Desktone platform is qualified to work with the following version of Data ONTAP:<br>• Clustered Data ONTAP 8.1.2. |
| VSC | The NetApp VSC is a plug-in for VMware vCenter leveraged by the Desktone platform for rapid cloning of virtual machines when vSphere hosts are discovered through vCenter. The following version of VSC is supported:<br>• VSC 4.1 |
| VMware ESXi | VMware ESXi 5.1 and vCenter 5.1. |
| Support Matrix | Refer to the Interoperability Matrix on the NetApp Support site for more details for exact driver/firmware versions and supported host platforms. |

| Component | Description |
|---|---|
| | NetApp Interoperability Support Matrix |
| Licenses | • **FlexClone**. FlexClone provides the ability for the Desktone platform to clone gold images in the matter of seconds (optional for Desktone, but recommended).<br>• **NFS**. This is the base license required to use the NetApp storage system for NFS (required by Desktone).<br>• **CIFS**. This is the base license required to use the SVM for CIFS user data.<br>• **SnapMirror**. A SnapMirror license is required for thin replication (optional for Desktone, but if the design incorporates DR, then SnapMirror license is recommended).<br>• **SnapRestore®**. To rapidly restore from backup. An alternative would also be to clone and mount the volume and copy the VMDK disk manually. |

## 7.2 Virtual Storage Console Integration

Desktone integrates with NetApp VSC, which leverages the RCU for fast cloning. NetApp VSC for VMware vSphere provides a provisioning and cloning API, which is leveraged along with the VI SDK by Desktone for quick and efficient provisioning of VMs.

As a vCenter Server plug-in, VSC for VMware vSphere is available to all vSphere Clients that connect to the vCenter Server. Unlike a client-side plug-in that must be installed on every vSphere Client, you can install the VSC for VMware vSphere software on a Windows Server instance in your data center.

The plug-in provides different capabilities to perform the following functions:

- Storage configuration and monitoring using the monitoring and host configuration capability.
- It enables you to manage ESX and ESXi servers connected to NetApp storage. You can set values for host timeouts, NAS, and multipathing as well as view storage details and collect diagnostic information.

  **Note:** It is very important to set the timeout values, NFS recommended settings, and multipath configuration on all of the ESXi hosts used in the Desktone configuration using VSC.

- Datastore provisioning and virtual machine cloning using the provisioning and cloning capability.
  - The provisioning and cloning capability uses FlexClone technology to let you efficiently create, deploy, and manage the lifecycle of virtual machines.
- Online alignments and single and group migrations of virtual machines into new or existing datastores using the optimization and migration capability.
  - The optimization and migration capability enables you to quickly check the alignment status of virtual machines. If alignment issues are encountered with virtual machines in VMFS datastores, you can in most cases resolve those issues without having to power down the virtual machines.
- Backup and recovery of virtual machines and datastores using the backup and recovery capability.
  - The backup and recovery capability allows you to rapidly back up and recover multihost configurations on NetApp storage.

Refer to the following for installing and configuring NetApp Virtual Storage Console for VMware vSphere.

## 7.3 Tenancy Models

The tenancy model for Desktone is shown Figure 20. Each tenant will have two SVMs associated, out of which, one SVM will be for the tenant VDI VMDK data, and the second SVM will be for CIFS user data. The SVM hosting the tenant VDI VMDK data can be shared or dedicated depending on customer needs. Secure separation is by way of export-level security and hypervisor-level security if the SVM is shared for VDI VMDK data.

Basically, each tenant is mapped to two SVMs:

- One shared SVM with other tenants containing the respective tenants' VDI VMs in separate volumes. The SVM hosting the tenant VDI VMs can also be dedicated depending on the architectural needs.
- One dedicated SVM containing the tenant CIFS user data.

For provisioning new tenants, the Data ONTAP components will have to be duplicated for deploying the new tenants. New environments can be deployed seamlessly and efficiently in a matter of minutes with clustered Data ONTAP.

When a shared SVM is used for tenant VDI VMs, the volumes that host the tenant VMDKs are exported out to specific tenant hosts only (export level security) - volumeA > TenantA and volumeX > tenantX and so on. All the user data would be stored on the CIFS shares in a dedicated SVM so the tenants would still have secure separation from a user data perspective.

**Figure 20) Tenancy model.**

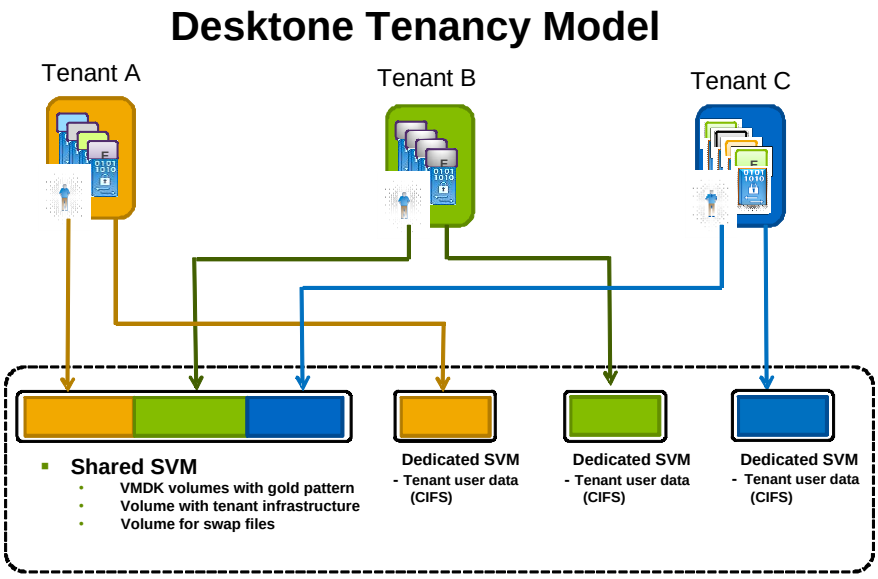


This following section describes the NetApp clustered Data ONTAP storage layout from aggregate, volume, and networking perspective along with the logical layout for different tenants.

## 7.4 Aggregate Layout

Figure 21 shows an example disk layout for production data on the NetApp storage controllers for a Desktone deployment. The example assumes a two-node cluster with NodeA and NodeB.

**On NodeA**

- **Root aggregate**. Aggr0 is used only for the root file system and is typically a three-drive RAID-DP configuration.
- A "management" aggregate, which will hold the boot LUNs for the ESXi hosts and the NFS volumes that host the Desktone management appliances, such as SP appliance, resource management appliance, and tenant appliances.
- A VDI aggregate hosting the VDI VMDKs and CIFS user data for tenants. This is the data aggregate hosting all the production VDI VMs with the required number of spindles. Refer to the System Performance Modeler (SPM) Sizing Tool for Sizing.

**On NodeB**

- **Root aggregate.** Aggr0 is used only for the root file system and is typically a three-drive RAID-DP configuration.
- A VDI aggregate hosting the VDI VMDKs and CIFS user data for tenants. This is the data aggregate hosting all the production VDI VMs with the required number of spindles. [Refer to the System Performance Modeler (SPM) Sizing Tool for Sizing.]

Spread the tenant load across both the nodes. For instance, if there are two tenants, have TenantA use volumes on NodeA and have TenantB use volumes on NodeB. You can also similarly choose to spread the boot LUNs across the nodes to distribute workload and increase resiliency.

**Figure 21) Storage aggregate layout.**

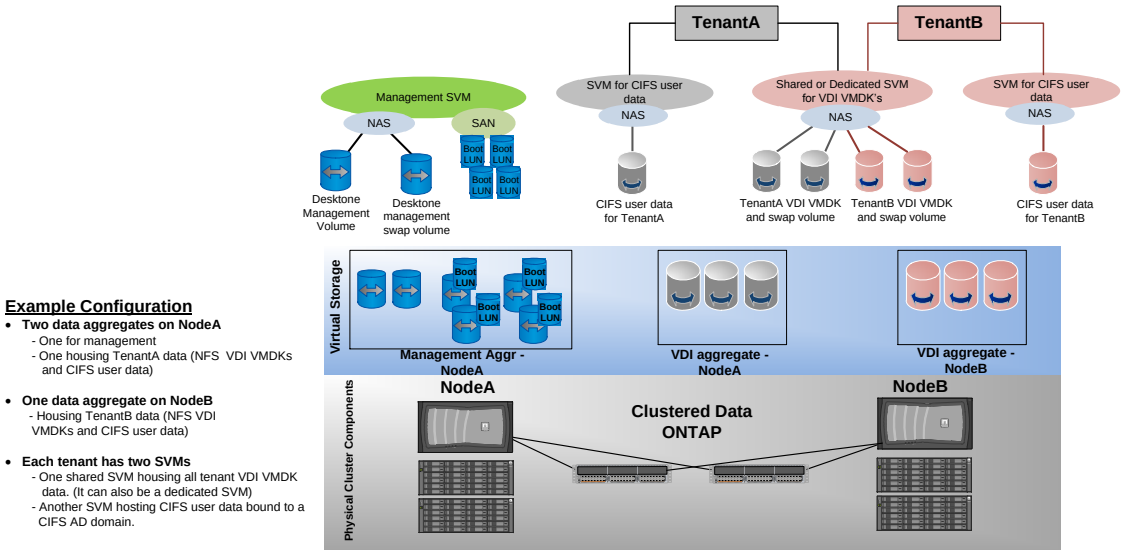


**Example Configuration**

- **Two data aggregates on NodeA**
  - One for management
  - One housing TenantA data (NFS VDI VMDKs and CIFS user data)
- **One data aggregate on NodeB**
  - Housing TenantB data (NFS VDI VMDKs and CIFS user data)
- **Each tenant has two SVMs**
  - One shared SVM housing all tenant VDI VMDK data. (It can also be a dedicated SVM)
  - Another SVM hosting CIFS user data bound to a CIFS AD domain.

## 7.5 Network Link Aggregation (Physical)

To achieve optimal performance and resiliency, maximize the number of Ethernet links for both controllers in the NetApp active-active controller configuration. An interface group is a port aggregate containing two or more physical ports that acts as a single trunk port. Expanded capabilities include increased resiliency, increased availability, and load sharing.

NetApp recommends dynamic multimode ifgroup configuration if the switches support LACP. Dynamic multimode interface groups implement Link Aggregation Control Protocol (LACP) to communicate group membership to the directly attached switch.

Figure 22 shows an example LACP/VIF configuration on clustered Data ONTAP.

**Figure 22) LACP VIF configuration (physical network configuration).**

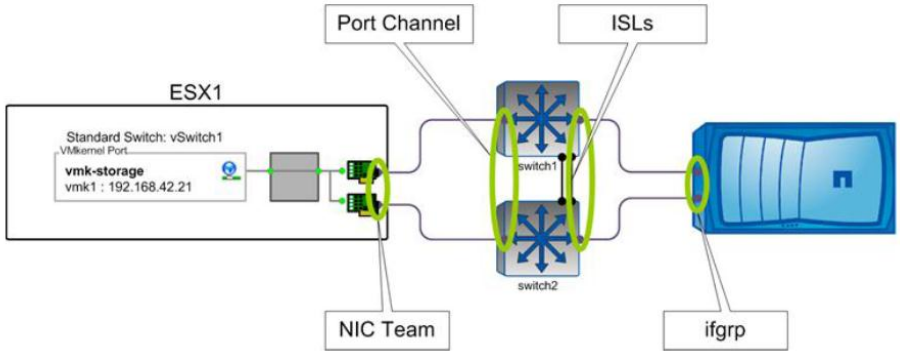

**Best Practices**

- Use switches that support link aggregation of ports on both switches.
- Disable LACP for switch ports connected to ESXi and enable LACP for switch ports that are connected to NetApp storage nodes.
- Use IP hash on ESXi and use dynamic multimode (LACP) with IP hash on NetApp storage nodes.

Refer to the following links for best practices and configuration guidelines:

- vSphere on clustered Data ONTAP 8.1: http://www.netapp.com/us/media/tr-4068.pdf
- Data ONTAP 8.1 Network Management Guide for Cluster-Mode: https://library.netapp.com/ecmdocs/ECMP1141762/html/frameset.html

## 7.6 Network VLAN (Logical)

Figure 23 shows an example of the logical network configuration. VLANs have been configured on the LACP ifgroup, and logical interfaces (LIFs) have been configured on the VLAN interfaces to achieve secure separation. The tenants have their own LIFs and routing groups/routes for secure separation.

**Figure 23) Example network VLAN configuration on NetApp nodes.**

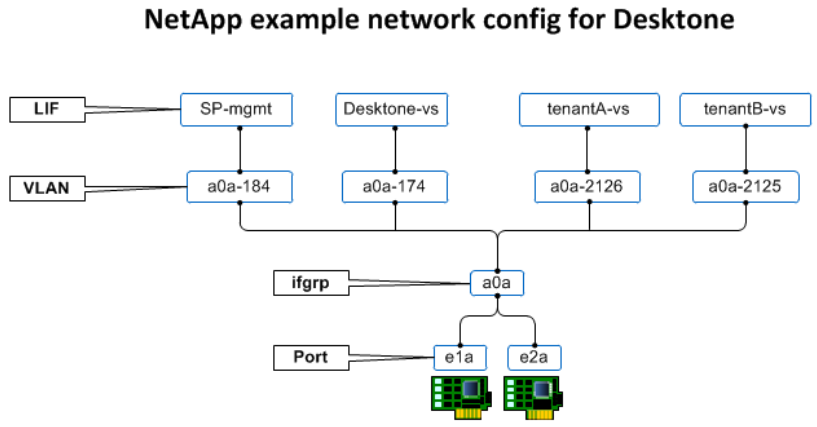


## 7.7 Storage Volume Layout

As discussed in the earlier section, from a Desktone perspective, NFS storage will host the management appliances and the VMs and CIFS storage hosting the tenants' CIFS user data.

Figure 24 shows the volume layout for the Desktone management base infrastructure. A single dedicated SVM used for Desktone management purposes will host all the Desktone infrastructure volumes as well as the boot LUNs for all the virtual ESXi hosts.

The Desktone management VMs consist of service provider appliances, resource manager appliances, and tenant appliances. These VMs will be hosted in the Desktone management volume, and their swap files will be on a separate volume.

**Figure 24) Desktone management infrastructure volume layout.**

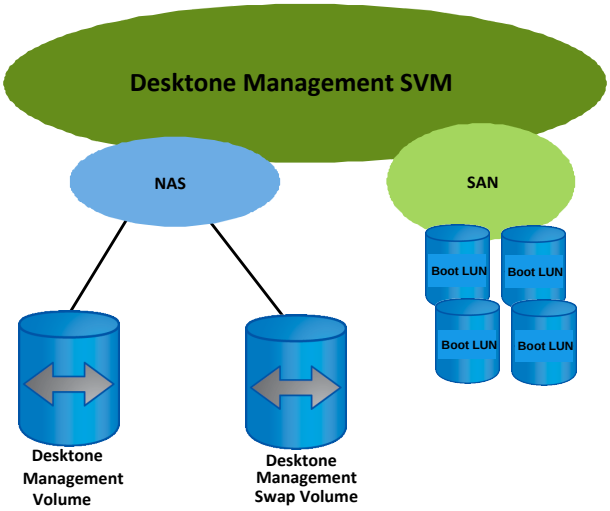


Figure 25 shows the volume layout for tenants. Each tenant has SVMs that host volumes with the base VMDK files of the Desktone VDI VMs, the swap files for the VMs, the tenant infrastructure (AD, DNS, DHCP, and so on), and the redirected CIFS user data.

Each tenant volume layout has:

- One shared SVM
  - Tenant VDI VMDK volumes
  - Tenant VDI swap file volumes
  - Tenant infrastructure volumes
- One dedicated SVM
  - Tenant CIFS user data volumes

Figure 25 shows each tenant having two SVMs, with one SVM hosting the VDI VMs and the other storage system hosting the CIFS user data, which is accessed directly from tenant virtual desktop systems. The tenant components illustrated in Figure 25 are to be duplicated for deploying or standing up multiple tenants. If the design dictates, then the SVM hosting the tenant VDI VM data can also be a dedicated SVM instead of shared.

The volumes that host the tenant VMDKs are exported out only to specific tenant hosts (export-level security) - volumeA > TenantA and volumeX > TenantX and so on. All the user data would be stored on the CIFS shares in a dedicated SVM, so the tenants would still have secure separation from a user data perspective.

Figure 25) Tenant data storage layout.

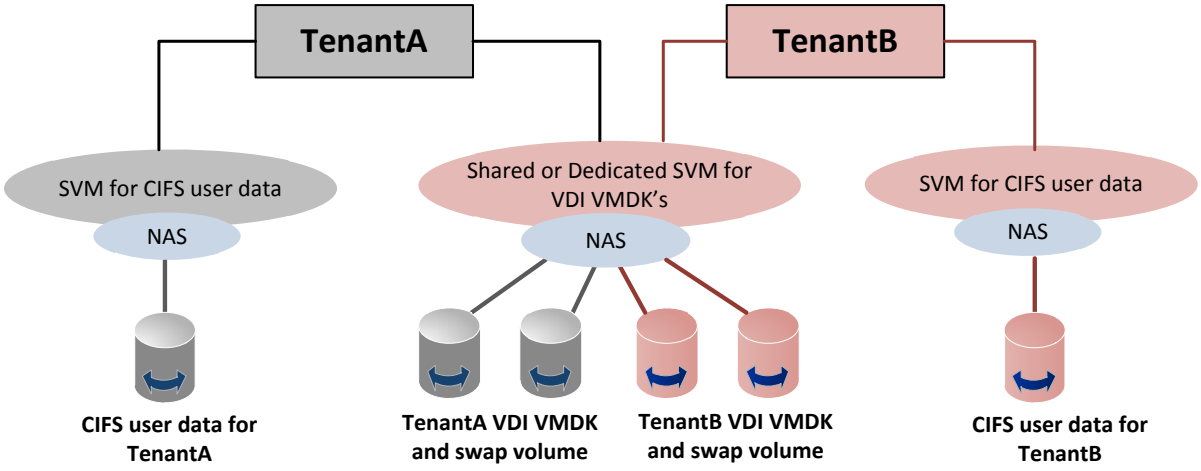


We can also provide different levels of "service" for different user types within a single tenant or between tenants. There can be different FlexVol volumes providing different levels of service within a single tenant for power users versus knowledge workers. It could very well be that all of the VMs will be housed on a single FlexVol volume to maximize the storage efficiency features. There are different combinations for volume deployment inside the SVMs for tenants, and that will depend on customer environments and their specific needs.

Different VM templates (gold patterns) will be hosted on different volumes because Desktone clones only the pattern on the same volumes on which the pattern resides. A VM pattern with a 4GB RAM, 2 vCPUs, and 50GB hard drive can be on a "gold volume" (thick provisioned, local Snapshot copies, and mirrored), and a VM pattern with a 2GB RAM, 1 vCPU, and a 30GB hard drive can be on a "silver volume" (thin provisioned and local Snapshot copies), and a VM pattern with a 1.5GB RAM, 1 vCPU, and 20GB hard drive can be on a bronze volume (thin provisioned and no Snapshot copies).

**Note:** Desktone clones gold images to the same volume on which the image resides, so make sure that the gold pattern already exists in the volume to which you want to host the VDI VMs. More details in the use case section.

Desktone recommends that users do not store data inside of the virtual desktop. Instead, user data directories, such as the My Documents folder, should be redirected to a separate storage location. There are several reasons for this:

- The storage used for desktop images is highly optimized for performance: high-speed FC or SAS disk is required; other optimizations are potentially in place to improve I/O. The storage requirements and performance profile for user data are significantly different from the desktop images, and thus a different class of storage can be used.
- Protecting data that is stored inside of a virtual desktop requires that each of the virtual desktops is backed up individually. If user data is redirected to a CIFS share, the external storage can be better protected and maintained independently to the desktop image.

## 7.8 Datastore Attributes

**Table 6) Datastore attributes.**

| Datastore | Attributes | Deduplication | Backup and Recovery | Replication |
|---|---|---|---|---|
| VM datastore | NFS volumes for Desktone management (appliances) volumes and for Desktone tenant virtual desktops (VMDK data of all the tenant VMs). | Yes | Yes, if tenant VMs are dedicated and static. | Yes, if tenant VMs are dedicated and static. |
| VM swap | NFS volumes for swap files for the Desktone management (appliances) volumes and for Desktone tenant virtual desktops (VMDK data of all the tenant VMs). | No | No. | No. |
| CIFS user data | Volumes hosting the tenant CIFS user data. | Yes | Yes. | Yes. |
| Template datastore | NFS volumes hosting the gold pattern. | Yes | Yes. | Yes. |

Some best practices are:

- Remove all transient data from the template VM before deploying virtual desktops.
- When using Network File System (NFS) with VSC, perform a space-reclamation process on the template to make the VM as small as possible.
- After the virtual desktops are provisioned, NetApp recommends running deduplication on the system to reduce the amount of consumed storage. It might be necessary to provision a portion of the desktops, run deduplication, and then provision additional desktops so that not all of the storage is consumed in the datastore. This process should be repeated as many times as necessary.

  **Note:** Depending on the number of virtual desktops and the commonality of the data in the VM datastores, the initial deduplication process might take a considerable amount of time to complete.

## 7.9 How to Create a Windows 7 Gold Template

Before defining a VM as your gold template, you need to create your template. We strongly recommend against using a physical-to-virtual (P2V) conversion tool. Instead a new OS install should be customized to VDI best practices. The template VM must exist on the same host, storage, and VLAN mapped to the tenant.

There are numerous online publications on Windows 7 VDI best practices, such as VMware-View-OptimizationGuideWindows7-EN.pdf.

The following steps must be part of the VM preparation:

1. Create a template VM.
2. Install VMware tools and verify NIC settings (NIC type should be VMXNET3).
3. Enable administrator account and confirm RDP access.
   a. On the **Member Of** tab, confirm **Administrator** is a member of **Remote Desktop Users**.
4. Install the Desktone DaaS agent by copying DaaSAgent_5.4.0.msi onto your VM and running the install.

5. The DaaS agent must be configured to point at the tenant appliances. This can be done using one of these methods:
   a. **DaaS agent discovery**. The tenant appliance addresses can be automatically discovered by the DaaS agent through DHCP by using the option code 74. Configure DHCP option code 74 (IRC chat) to point to the two IPs allocated for the tenant appliances. For more information, refer to Desktone Tenant Installation Guide. (You will need to contact Desktone for the username/password to access the portal.)
   b. **Update of DaaS agent configuration file**. The tenant appliance addresses can be manually updated in the DaaS agent configuration file. Open the file C:\Program Files (x86)\DaaS Agent\service\MonitorAgent.ini with a text editor such as the Notepad.

   **Note:** On 32-bit systems, the path will exclude the "(x86)."

   Remove the semicolon in the line containing the parameter "standby_address" and enter a comma-separated list of the tenant appliance IP addresses. A restart of the DaaS agent Windows service is required after making this change.
1. Install the PCoIP protocol (VMware View® agent and the PCoIP GPO).
2. Add the VM to the domain and add the appropriate domain groups to **Remote Desktop Users** group.
3. Set the guest OS timeout value.
   a. It is important to set the Windows guest OS timeout value to help improve recoverability after a storage failover event. The NetApp Knowledge Base Article KB3013622 describes the recommended settings and the rationale for deciding which timeout value to use.
   b. Run the guest OS timeout script for Windows as described in the Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide.
4. Set power options for the VM to HIGH Performance.
5. Confirm Windows firewall is disabled, or at least the necessary ports are configured.
6. Confirm Windows updates are current.
7. Optional: Log in as administrator and remove all other accounts on the VM.
8. Optional if users connect through other browsers: Confirm remote settings are not using NLA. NLA might interfere with IE users trying to connect to their desktops: "Allow connections from computers running any version of Remote Desktop."
9. Optional: Disable Ctrl+Alt+Del Secure logon. Some protocols (and users) struggle with entering Ctrl+Alt+Del to log into their VM.
10. Install the appropriate protocol drivers if you choose additional protocols beyond RDP.
11. After your VM has been configured with the appropriate software, you can begin the conversion process.
12. Confirm that the VM has been imported into the **Enterprise Center** and appears in **Imported Desktops**.

# 8 NetApp Use Cases

Following are the main use cases validated with clustered Data ONTAP and Desktone:

- SMT
- VSC integration for provisioning and cloning
- Backup and recovery

## 8.1 Secure Multi-Tenancy (SMT)

Clustered Data ONTAP is an inherently multi-tenant storage operating system. Clustered Data ONTAP is architected in such a way that all data access is done through secure virtual storage partitions known as storage virtual machines (SVMs).

Some of the many benefits of running clustered Data ONTAP are seamless data mobility, scalability, and nondisruptive operations (immortality). This scalability also enables SVMs to be highly resilient. SVMs are no longer tied to the lifecycle of a given storage controller. This means that tenants and tenant resources can be nondisruptively moved anytime during lifecycle operations or maintenance operations. In this way, new disk, cache, and network resources can be made available to the tenants to create new data volumes or migrate/move existing workloads to these new resources in order to balance performance.

For example, a tenant VDI VM 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 tenant SVM abstracts the cluster hardware and is not tied to specific physical hardware.

Because Desktone is a solution specifically catering to service providers, clustered Data ONTAP is a great fit because 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. It also provides service providers the ability to easily 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) versus 10GbE interfaces, degree of redundancy such as single physical port or an interface group with multiple physical interfaces, and so on.

For instance, a tenant might prefer to segregate power users on high-performance storage and knowledge workers on SATA. It is also a very simple one-click activity to move between different “tiers of service” should the administrator desire to move the users from one service level to another. Depending on needs or requirements, cluster administrators can nondisruptively relocate the volumes to another aggregate of a different tier.

**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.

### SVM Limits

Because the SVMs hosting the tenant CIFS user data will be dedicated to each tenant, it is important to know SVM limits per node and for the entire cluster.

For example, the choice of the service provider administrator to have separate management LIFs for each tenant SVM or to combine them with data LIFs will affect the number of total SVMs it is possible to create.

Keep the following points in mind:

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

  **Note:** It is possible to add more than the minimum required number of LIFs to an SVM. This will reduce the maximum SVMs it is possible to create per cluster. For example, a tenant might

have volumes that need a higher level of “service” that have one-to-one mappings with LIFs. In this situation, more LIFs are consumed, which also affects SVM limits.

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

For specific guidelines and recommendations, refer to the SMT Best Practices Guide.

## Network Isolation of Tenants

Each tenant SVM is only aware of the resources that have been assigned to it and has no knowledge of other tenant 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. Every tenant SVM serving CIFS user data connects to a unique AD domain.

An important aspect of multi-tenancy is securing network traffic so that tenants can be securely isolated from one another. 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.

In the NetApp architecture for Desktone, the use case is for many secure multi-tenant networks whose tenants require segregated IP networks, each with a unique nonoverlapping IP address range. Each tenant 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.

Figure 26 shows an example of routing group configuration in a Desktone deployment. Different routing groups for the tenants make sure of secure separation. The SP management subnet and the storage subnet are also in different routing groups through different VLAN LIFs, though hosted by the same SVM.

Figure 26) Example routing groups in a Desktone deployment.

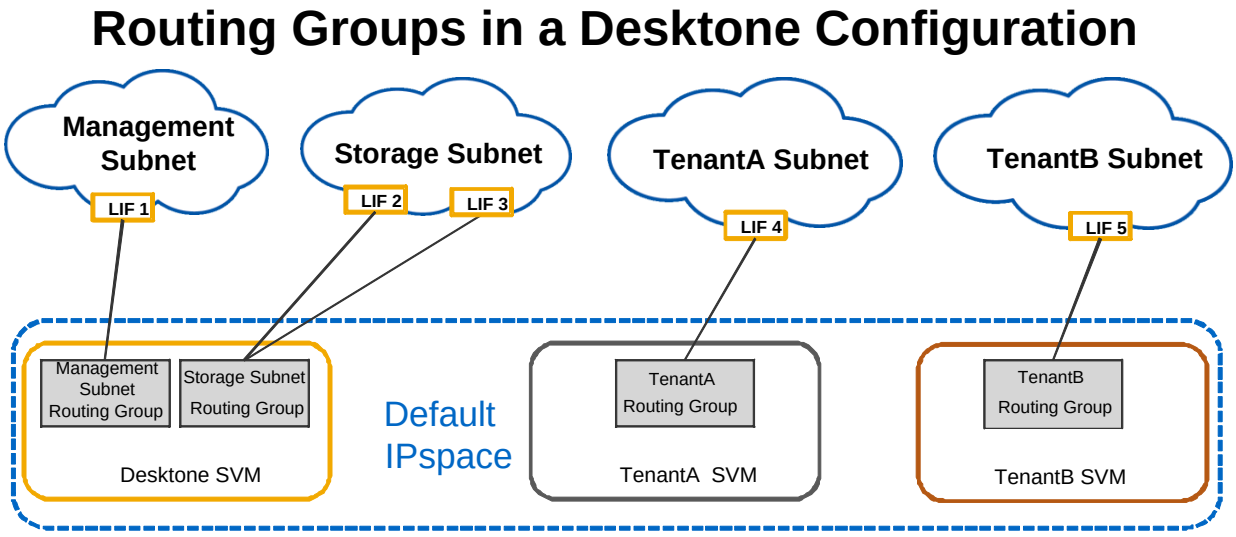


## Failover Groups

In the event of a failure, LIFs need to be migrated in a coordinated manner to other surviving nodes/ports to make sure of seamless operations. This is done through failover groups. A failover group contains a set of network ports on one or more nodes, and it is a best practice recommendation that LIFs capable of utilizing failover groups should be assigned to appropriate failover groups with the correct network ports.

When a LIF is created, it is assigned to a system-defined failover group by default. However, the behavior of the default failover group might not be sufficient for every different type of environment. Failure to determine that ports are in the same subnet and failure to assign LIFs to appropriate failover groups result in loss of connectivity to data. It is recommended to created user-defined failover groups so that LIFs do not fail over to a VLAN port that resides in a different subnet, which could lead to a failure.

Because multiple subnets are likely to exist in a Desktone deployment; the default system-defined groups could cause LIFs to fail over to VLAN ports that reside in a different subnet. Basically, you have to be certain that LIFs that are configured on top of VLAN ports move to ports that can communicate with the other devices that are members of that same VLAN configuration.

You can also define failover groups if you want to logically group a certain type of interface (for example, 10GbE-based LIFs only fail over to 10GbE ports).

Because the Desktone deployment has LIFs in different VLANs or broadcast domains, failover groups must be configured for each VLAN or broadcast domain. You must then configure the LIFs hosted on a particular VLAN or broadcast domain to subscribe to the corresponding failover group.

**Note:** Failover groups do not apply in a SAN iSCSI or FC environment.

Figure 27 shows an example failover group configuration in a Desktone deployment for illustrative purposes. For example, with the default system-defined failover group configuration, the storage LIF hosted on a0a-174 on NodeC is likely to fail over to a VLAN on NodeD, which is not the desired outcome because there is no a0a-174 VLAN on NodeD. Similarly, the a0a-2125 for TenantA could possibly fail over to a VLAN port on NodeA, which is also not desired. Hence, configure failover groups on every node with the relevant VLAN ports included and add the LIFs to the failover groups.

Figure 27 is not an actual recommendation for failover group configuration because it is highly possible that in a production environment all of the nodes might have VLANs for all the tenants, unlike what is illustrated in Figure 27 as an example.

Figure 27) Example failover groups in a Desktone deployment.

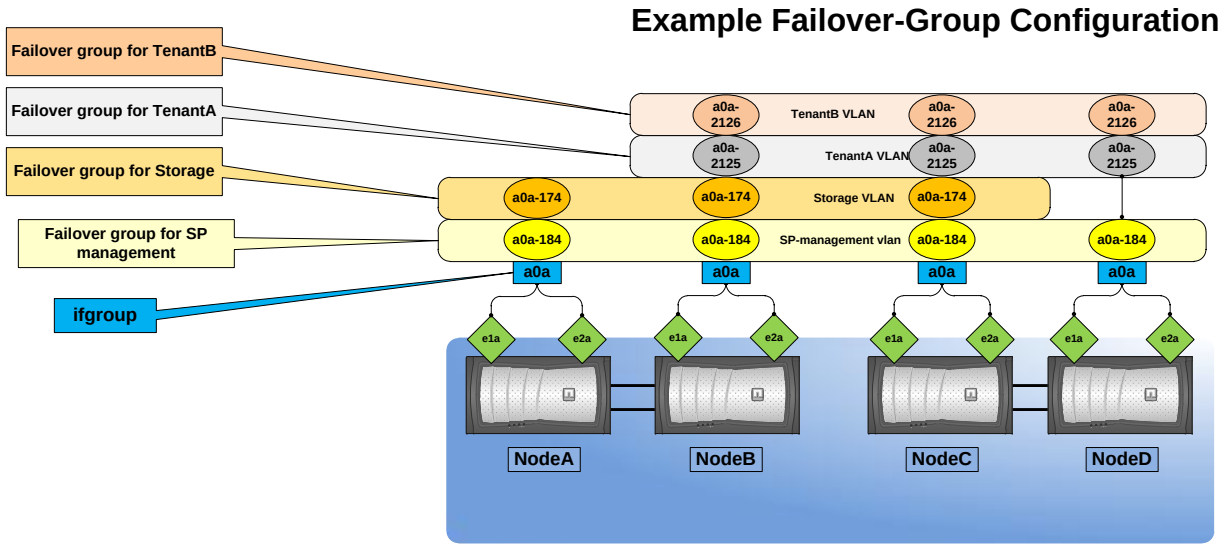


## 8.2 Provisioning and Cloning (VSC Integration)

Desktone integrates with the NetApp VSC, which leverages the RCU for fast cloning. VSC for VMware vSphere supports the provisioning and cloning API, which is leveraged with the VI SDK by Desktone for quick and efficient provisioning of VMs.

VSC must be installed and configured for the vCenter data center, which is to be used for the Desktone deployment. Make sure that the Data ONTAP cluster or the SVM has been added to VSC with the appropriate credentials that have the authorization to perform all the operations on the cluster or the SVM.

For installing and configuring, refer to NetApp Virtual Storage Console for VMware vSphere.

Earlier versions of Desktone did not have vCenter support, but Desktone 5.4 has vCenter support, and it integrates with the VSC APIs, which enable Desktone to clone and provision hundreds of VMs very quickly based on the gold templates in a particular volume.

### Shared Storage Requirements (NFS, FC, and iSCSI) for Desktone Management Appliances

The first stage in the Desktone deployment is to install and configure all the management appliances (service provider appliances, resource manager appliances, and finally the tenant appliances). After the first management appliance is installed, the remainder of the management appliances are cloned out of the template on the NFS datastore.

The default behavior of the Desktone platform is to clone out management appliances on the local disk (through a local datastore), but when using shared NetApp NFS storage, the default behavior must be changed within the Desktone platform to allow cloning and provisioning on the NFS datastores on clustered Data ONTAP. The following section provides the steps for turning off local disk provisioning.

Some guidelines for shared storage:

- Any shared storage (NFS, iSCSI, or FC) can be used. Make sure that the appropriate configuration has been done (for example, LUN masking, zoning, and so on) to allow one or more LUNs to be mapped to all of the management hosts. Desktone does NOT support VSC for FC or iSCSI block storage. The Desktone platform uses the native vCenter cloning APIs to provision management appliances using FC or iSCSI. For this to work properly, datastores must be created and mapped to the same LUNs on all the management hosts with the exact names (case sensitive).

- Datastores must be manually created on each of the management hosts.
- The datastore name must be identical (case sensitive) on each management host.

It is recommended to use VSC to provision the NFS datastores and mount them with the same datastore names on all the management hosts designated for Desktone appliances.

**Note:** Desktone does NOT support VSC for FC or iSCSI block storage. The Desktone platform uses the native vCenter cloning APIs and NOT VSC to provision and clone management appliances when using FC or iSCSI datastores.

## Turn Off Local Disk Provisioning

As mentioned earlier, the default behavior of the Desktone platform is to clone out management appliances on local disk (through a local datastore) using native vCenter cloning APIs, but because we are using NFS datastores hosted on clustered Data ONTAP, the default behavior must be changed so that Desktone clones and provisions the appliances on shared storage (NFS) instead of local storage.

The first step after installing the initial Desktone management appliance and getting access to the Service Center portal is to change the default behavior of the application to allow for installation on shared storage.

To access the installation and configuration guide for Desktone in a vCenter environment, click here. For the username and password for accessing the portal, contact Desktone.

1. After the first management appliance is installed and configured, to start the **Service Center,** enter the URL or IP address in a browser. For example:
   a. https://<IP for eth0 (SP)>/service

**Figure 28) Desktone Service Center portal.**

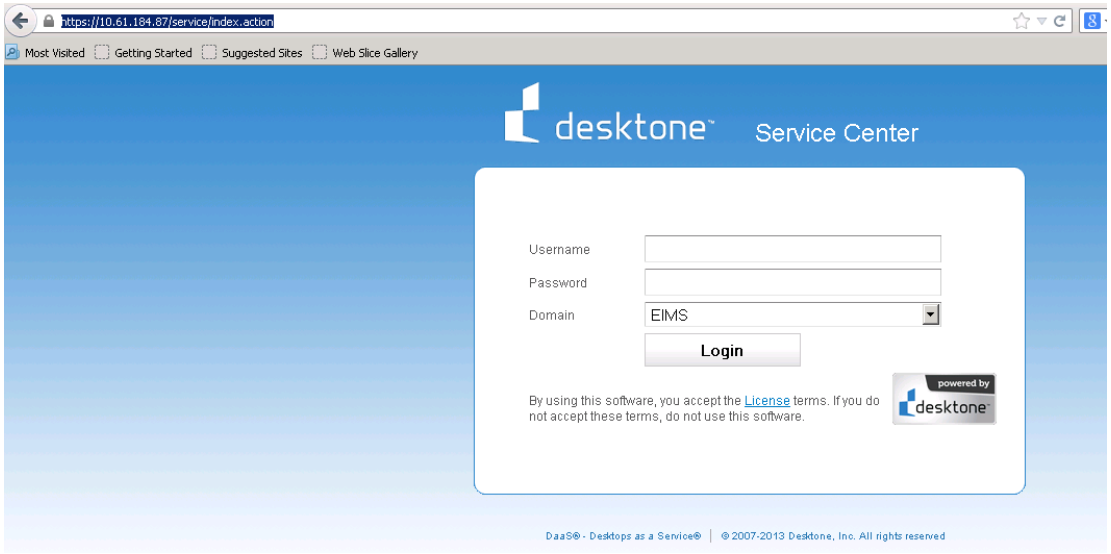


2. Log in to the **Service Center** and select the **tenants** > **policy**.

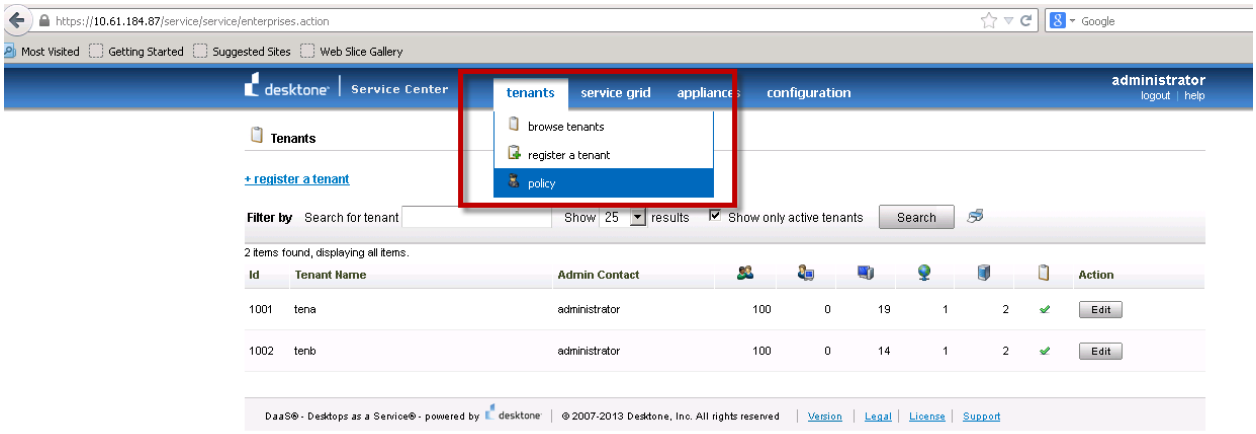


3. The **Policy configuration** page is displayed. Select **Service Provider** from the **Tenant Name** drop-down menu.


4. Scroll through the list of policies and set the vmgr.appliance.local.disk policy to **false**.

| | |
|---|---|
| tenant.authentication.override | false |
| upgrade.snapshot.keep.seconds | 259200 |
| user.event.report.allow | true |
| user.event.report.interval | 5 |
| userportal.session.timeout | 60 |
| vmgr.appliance.local.disk | false |
| vmgr.isilon.concurentCopy | 5 |
| vmgr.state.sync.interval | 86400000 |
| vmgr.vcloud.discovery.enable | false |

## VSC Credential Configuration and Fetching It from vCenter

Because Desktone has been enabled to use shared storage for appliance deployment, the next step is to configure Desktone to use NetApp VSC for provisioning and cloning.

Some details are required for configuring storage on the Desktone portal.

1. Note the credentials you have used for the NetApp cluster in VSC. Navigate to VSC from vCenter, right-click the cluster, and click the **Modify Credentials**.
2. To verify that the given credentials have all the necessary privileges to perform the operations on storage, click **Show Privileges**.

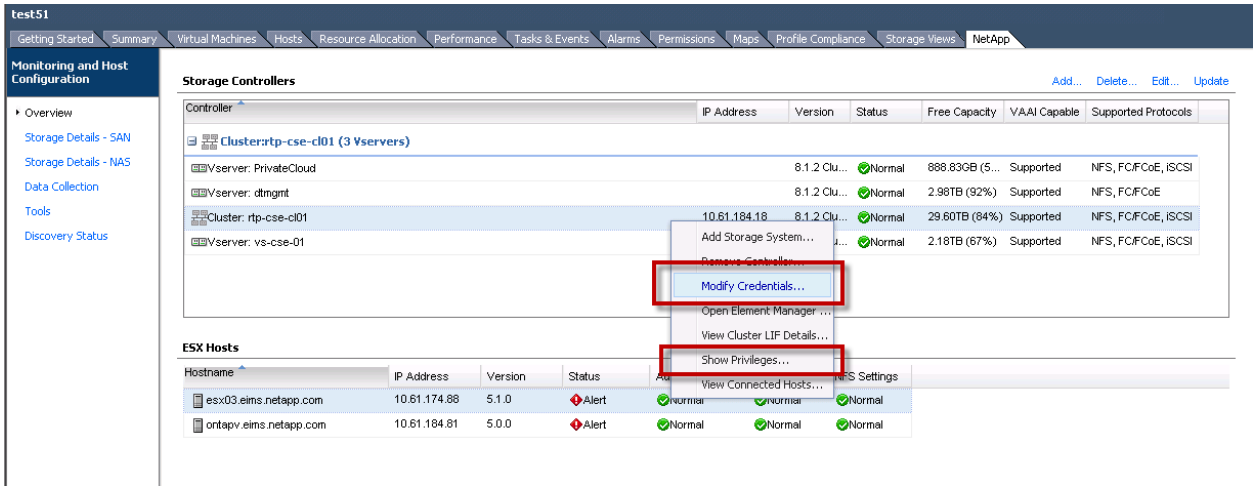


The **Modify Credentials** pops up the credentials box, where you can see the **User name** that has the privileges to perform all the VSC operations on the NetApp cluster. Following are the exact details that are required later for configuring storage on the Desktone portal:

- IP address of the cluster (you can also specify the delegated **Management IP address** of the SVM).
- **User name**.
- **Password**.

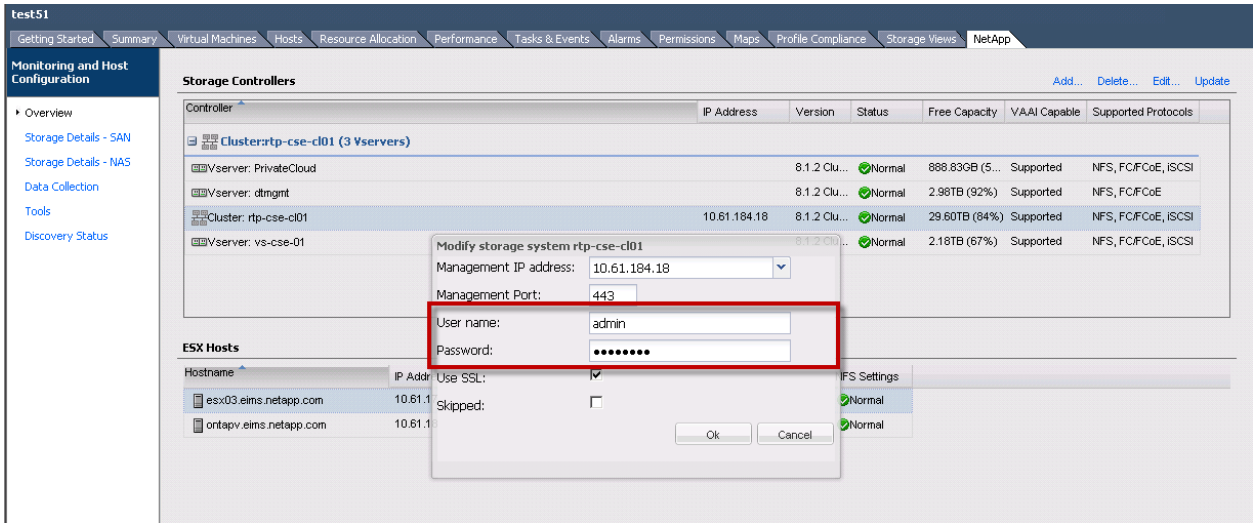


## NetApp Storage Configuration for Desktone Management Appliances

1. Navigate to the **Service Center** portal and click **service grid** > **resources**.

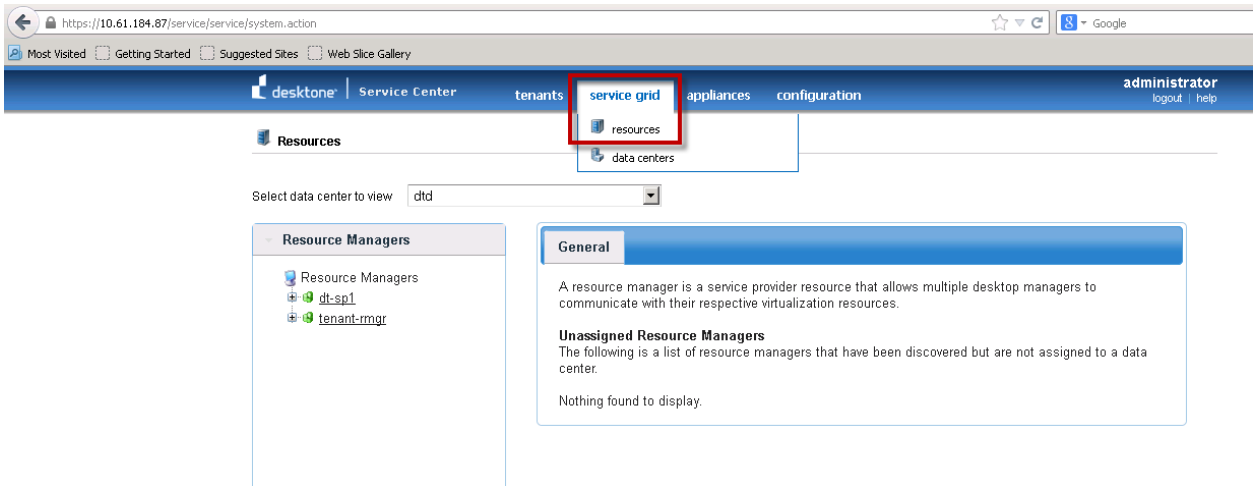


2. Click **Service provider** rmgr in the left pane and click **Storage Systems**.

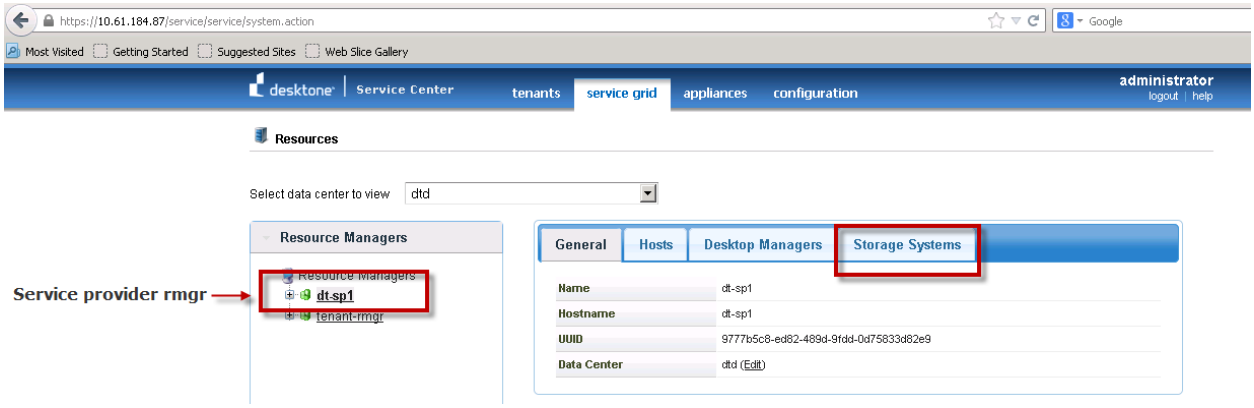


3. On the Storage Systems tab, click Add Storage System.

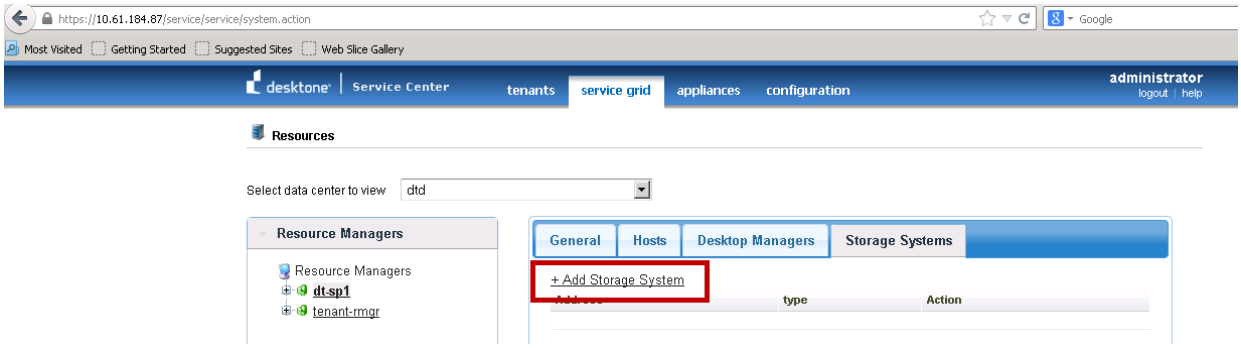


4. The values for the fields shown in the following screenshot are the same values that were fetched from the VSC credential page. The **Address** field must match how the VSC plug-in discovered the NetApp cluster. If VSC discovered the cluster using an IP address, then enter the IP address. If it was discovered as an FQDN, then you must enter the complete domain name in the **Address** field. Enter the exact details as shown in the VSC credential page.

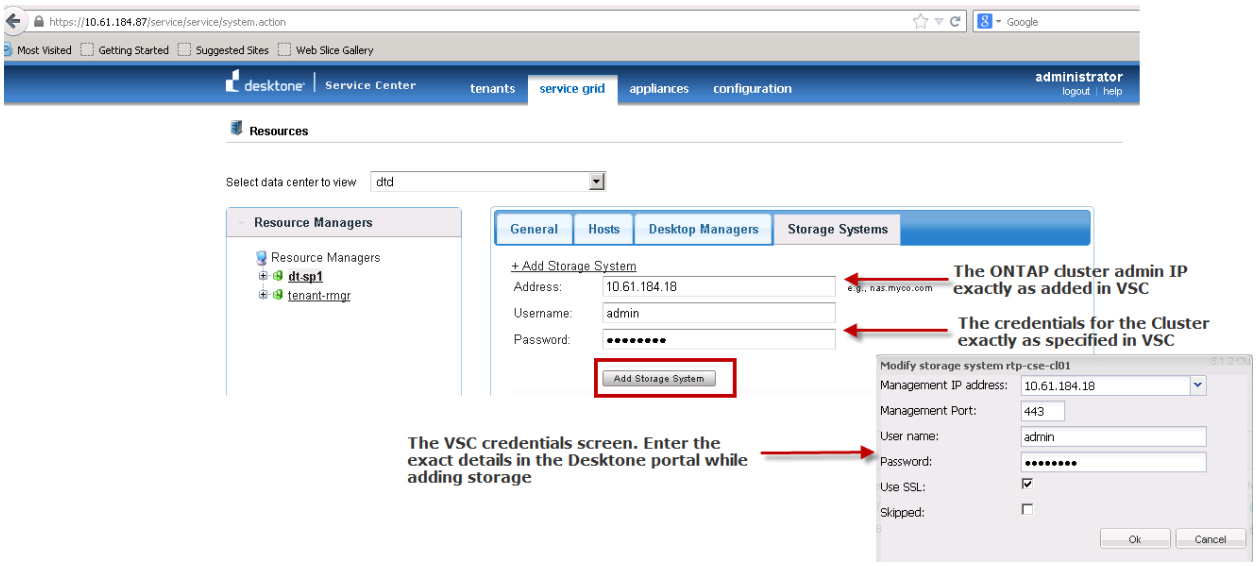


5. Click **Add Storage System**. The system adds the name of the storage system to the **Storage Systems** tab.

**Note:** It is very important to make sure that the storage system shows up as **NetApp** under the **type** column, as shown in the following screenshot. When **NetApp** is displayed on the screen, it confirms that VSC rapid cloning will be used while provisioning the VMs. Sometimes, if the configuration details entered are incorrect, the storage will be added as **NFS** instead of **NetApp** under the **type** column. In that case, remove the storage system and enter the correct configuration details and verify that the storage system shows **type** as **NetApp**.

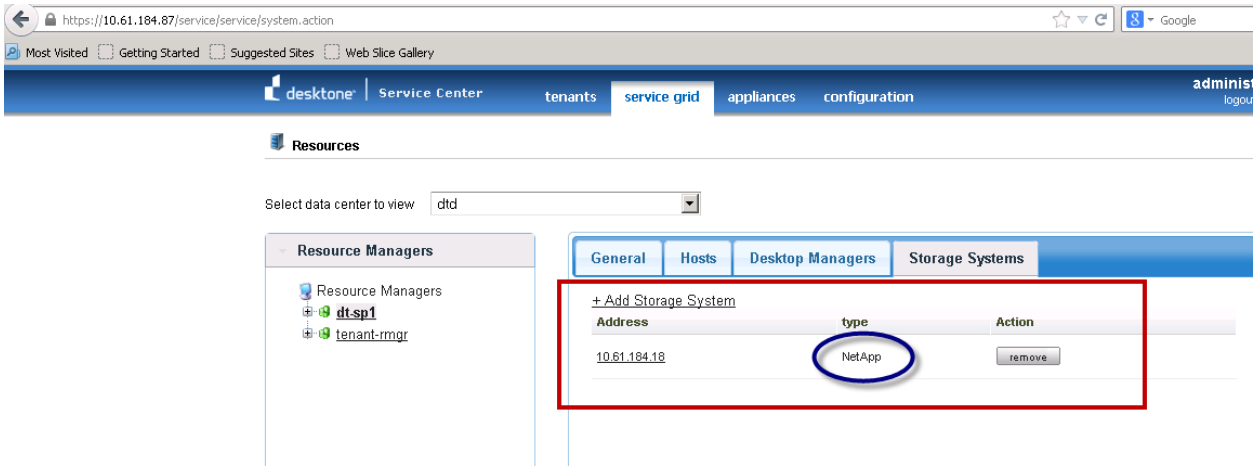


6. After the storage configuration is completed, the remainder of the management appliances can now be deployed. This will be the secondary service provider appliance and the resource manager HA pair appliances. The vCenter event details will display that the NetApp RCU is being used for the appliance deployment. Additionally, you can also verify that the cloning process is using NetApp VSC by looking at the service provider appliance logs that detail the provisioning and cloning process. The following screenshot illustrates that VSC rapid cloning is being used for the appliance deployment.

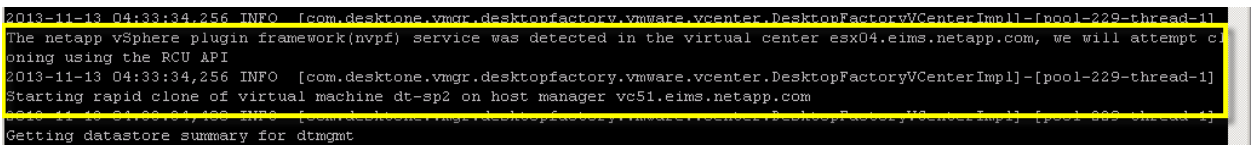

**Note:** Desktone does NOT support VMware clusters. If the management or the tenant VDI hosts are in a VMware cluster (even in an empty cluster with HA and DRS disabled), then the Desktone cloning process will fail. Make sure that the Desktone hosts are not in a VMware cluster before starting the deployment.

## NetApp Storage Configuration for Desktone Tenants

The procedure for configuring NetApp storage for Desktone tenants is similar to the procedure detailed in the preceding section. The only difference for tenant storage is that the storage systems must be configured in **Tenant resource managers** and not in the **Service provider** tab, as done in the earlier section.

1. Navigate to the **Service Center** portal and select **service grid** > **resources**.
2. Click **Service provider** rmgr in the left pane and click **Storage Systems**. On the **Storage Systems** tab, click **Add Storage System**.

   **Note:** In the earlier section, the storage configuration was done for the service provider. This storage will be used ONLY for the Desktone management appliances (the service provider appliances, resource manager appliances, and tenant appliances themselves when tenants are instantiated). In this section, steps to add storage to the tenant resource managers are provided. This is the storage that will be used for tenant VDI VMs. The configuration details are the same because we are entering the NetApp cluster details in both the places.

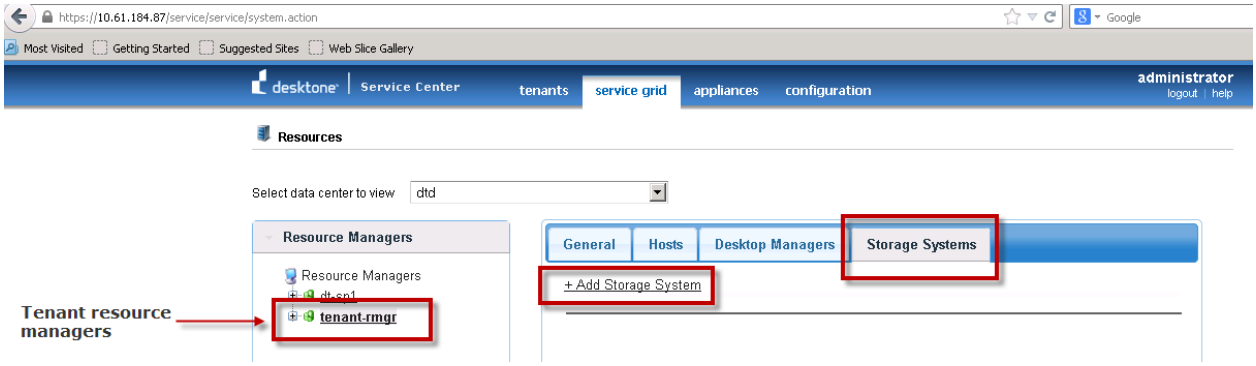


3. The values for the fields shown in the following screenshot are the same values that were fetched from the VSC credential page. The **Address** field must match how the VSC plug-in discovered the NetApp cluster. If VSC discovered the cluster using an IP, then enter the IP. If it was discovered as an FQDN, then you must enter the complete domain name in the **Address** field. Enter the exact details as shown in the VSC credential page.

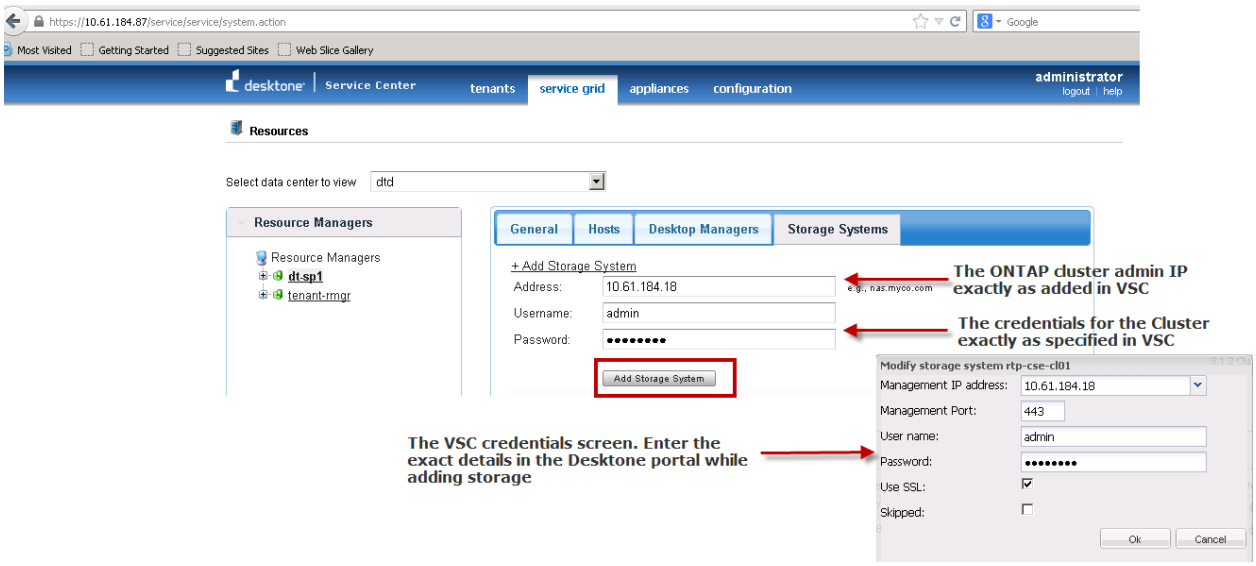


**Note:** As noted in the earlier section, verify that the storage type that is added shows up as **NetApp** and not **NFS**.

## 8.3 Backup and Recovery (VSC Integration)

Integrated data protection in accordance with the type of data, and the requisite customer requirements will have to be implemented so that data is automatically protected from the time it is provisioned. For instance, the volume hosting the Desktone management appliances will have local Snapshot copies and volumes housing the CIFS User data can be Snapshot and SnapMirror protected, so different strategies can be used depending on customer requirements.

NetApp VSC should be used to protect the datastores used by the Desktone hosts. The backup and recovery capability of VSC for VMware vSphere enables you to rapidly back up and recover multihost configurations running on the NetApp storage systems.

You can use the VSC backup and recovery capability to perform the following tasks:

- Create on-demand backups of individual virtual machines, datastores, or a data center
- Schedule automated backups of individual virtual machines, datastores, or a data center
- Support virtual machines and datastores that are located on either NFS directories or VMFS file systems
- Mount a backup to verify its content prior to restoring it
- Restore datastores to their original locations
- Restore virtual machines to their current locations
- Restore virtual machine disks (VMDKs) to their current locations or alternate locations

For details on configuring the backup and restore capability of VSC, refer to the VSC 4.1 Documentation.

The tenant VDI VM data protection depends on the architecture that is implemented. Only static VMs might need protection depending on the user profile. For example, protection might be desired for a few VMs that are statically assigned to power users. In general, tenant VDI VM volumes might not need protection because all user data is expected to be redirected to CIFS shares, which will be protected. VSC, being a vCenter plug-in, cannot be used to protect tenant CIFS user data.

### Desktone Appliance Backup

With the exception of the database embedded in the Desktone appliance, the appliance is disposable: it is not required that an entire appliance be backed up. Desktone provides backup and restore scripts for the embedded databases with the platform. The backup scripts create a tar ball backup of the embedded databases, and the certificates restore scripts for the embedded databases with the platform. The backup scripts should be executed depending on the amount of changes that occur in the Desktone environment (provisioning new tenants, VMs and others). For accounts experiencing a high rate of daily delta change in data, it is recommended to schedule com jobs to run the script once or twice daily. These scripts need to be run on the service provider appliances and the tenant appliances (not required to be run on the resource manager appliances). Additionally, a rotation should be identified to send backups off site.

NetApp VSC can be used to create on-demand or automated backups of the appliances or the entire datastore itself. The backups can be scheduled to occur at preset intervals to complement the Desktone native backup strategy. The appliances can be restored to the point-in-time Snapshot copy for recovery.

### Desktone Backup

- **Backups and restores for gold images.** In the cases where user data is redirected to external file shares, backup of individual desktop is not required because the desktop can easily be recreated. However, the gold images used to create the desktops should be backed up regularly. This is because a gold image is a powered-off virtual machine; the backup should be performed by the service provider. The service provider could leverage Snapshot copies and replication from the storage system as a strategy to protect the gold images. This has the benefit of being able to restore quickly and the ability to provide disaster recovery. In the NetApp solution, a single volume can house all the different gold patterns, and VSC can be configured to protect the datastore.
- **Backups and restores for desktop VMs.** The static desktops can leverage NetApp VSC for backing up and restoring the VM's and files in the VMs. Use the administrator-assisted or the limited self-service recovery option in VSC to recover files in the tenant VMs. With dynamic desktops and CIFS user data redirection to NetApp, it might be easier to provision new VMs instead of implementing protection and recovery of dynamic desktops.

### User Data Backup Replication

The user data redirected from the desktop should be backed up regularly. Using the storage system, Snapshot copies and replication can be an effective strategy with the ability to restore files to a specific point in time.

If user data is stored at either the service provider or tenant data center, file shares should be replicated if high availability and/or disaster recovery is desired. If user data is stored in the cloud, coordination will need to take place with the cloud storage provider in order to understand how to implement HA.

In the NetApp solution, all the volumes housing the CIFS user data should be protected using local Snapshot copies and can also be mirrored, depending on architectural needs.

## 8.4 Nondisruptive Operations

Today's business environments require 24/7 data availability. Constant data availability begins with architecting storage systems that facilitate nondisruptive operations (NDO). Nondisruptive operations have three main objectives: hardware resiliency, hardware and software lifecycle operations, and hardware and software maintenance operations.

NetApp clustered Data ONTAP enables data mobility and nondisruptive operations, allowing data movement within the customer environment without incurring an outage.

Refer to the following link for DataMotion for Volumes Overview - Data ONTAP 8.1 Operating in Cluster-Mode.

# 9 Storage Best Practices

In general, following these guidelines makes sure that all the best practices are in place for a Desktone VDI deployment:

- **Aggregate creation**. The recommendation is to provision fewer large aggregates over more, smaller aggregates. The advantages to larger aggregates are that the I/O has more disks across which to write, therefore increasing the performance of all volumes contained within the aggregate. Also, disable scheduled aggregate Snapshot copies and set the aggregate snap reserve to zero.
- **Recommended settings for ESXi using VSC**. Monitoring and host configuration capability set ESX or ESXi host timeouts (FC) and other settings (NFS) to make sure of the best performance and successful failover. Refer to the Virtual Storage Console Installation and Administration Guide for configuring the recommended settings.
- **File system alignment**. File system misalignment is a known issue in virtual environments and can cause performance issues for virtual machines (VMs) and therefore could affect the performance of a Desktone virtual desktop deployment. It is therefore critically important that the NetApp file system alignment practices are followed as Best Practices for File System Alignment in Virtual Environments. Note that this issue is not unique to NetApp storage arrays and can occur with any storage array from any vendor.
- **Maximize aggregate size for desktops**. It is recommended to create an aggregate at the maximum size supported for the particular model of NetApp controllers. The reason for this is to include as many spindles as possible across which to spread data to have optimum I/O performance.
- **Separate volume and NFS exports per tenant**. For maintainability, separation, protection, and portability purposes, it is recommended that a separate volume hosting NFS exports is created for each tenant virtual desktops. Note: For smaller tenant deployments, it might not be wholly appropriate to follow the separate volume rule due to reduced efficiencies.
- **Separate volume to host the transient data (swap and so on)**. This is required to separate redundant data from the volumes hosting the base desktop image. This approach reduces the redundant data that would otherwise be locked into Snapshot copies and also replicated for DR. These volumes should have ASIS (deduplication) disabled.
- **Use NetApp FlexClone**. It is highly recommended that NetApp FlexClone is used for desktop provisioning operations; the ability to provision with FlexClone is built into the Desktone platform. FlexClone technology is rapid creation of space-efficient, writable, point-in-time images of individual files. FlexClone provides the ability to clone hundreds and possibly thousands of desktop images from a base desktop image, providing significant cost, space, and time savings.
- **Enable deduplication**. It is recommended to enable deduplication for the virtual desktop and any user data host volumes. Note that this needs to be enabled on a per volume basis (disabled by default).
- **Antivirus operations**. Optimize the AV operation policies for the VDI deployment. Optimizing the traditional AV policies means better planning the scheduled AV scan and virus definition update so that not all the virtual desktops run AV scan or virus definition updates at the same time, creating CPU contention within the environment. By staggering the scheduled AV operations and distributing the load at different points in time, you can avoid a large percentage of this contention. In addition to modifying the schedules, it is important to verify that these schedules do not interfere with other scheduled events such as backup or replication. For more details, refer to TR-3107: NetApp Antivirus Scanning Best Practices Guide.
- **Use Flash Cache**. It is highly recommended that Flash Cache is used to provide a large front-end read cache. This helps lighten the I/O load on the disks and thus helps deal with high read I/O load such as boot storms. The size of the Flash Cache storage varies depending on the FAS model. With clustered NetApp, each storage system head should have its own Flash Cache card.
- **Use Snapshot copies**. Snapshot copies can provide an almost instantaneous way to protect and recover appliance and user data and gold images. Snapshot copies do not affect performance and provide the ability for very fast restores.

- **SnapMirror for DR**. SnapMirror can be used to replicate desktop and/or user data off site in the case of a site failure. SnapMirror complements NetApp FlexClone and deduplication, providing a thin replication capability where the data reduction persists during the replication process.

## References

- VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1
- Nondisruptive Operations - DataMotion for Volumes Overview - Data ONTAP 8.1 Operating in Cluster-Mode
- SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP
- NetApp Data Compression and Deduplication Deployment and Implementation Guide Clustered Data ONTAP
- TR-3949: NetApp and VMware View 5,000-Seat Performance
- NetApp Antivirus Scanning Best Practices Guide
- Storage Subsystem Resiliency Guide
- File System Alignment in Virtual Environments
- VMware Horizon View 5 Solutions Guide
- Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide
- Storage Design Guidelines for Sizing Storage for Desktop Virtualization

## Version History

| Version | Date | Document Version History |
|---|---|---|
| Version 1.0 | January 2014 | Initial release |

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.

Go further, faster®




www.netapp.com