Technical Report

ITaaS Solution Guide
VMware vCloud Director and NetApp Unified Storage

Abhinav Joshi, NetApp
Wen Yu, VMware

September 2010 | TR-3866

TABLE OF CONTENTS

1 INTRODUCTION

In today's ultra-competitive environment, companies have to innovate rapidly to stay ahead of their competition. The increased business pressure has a direct impact on the IT organization to innovate and meet growing business needs rapidly. At the same time, CIOs are struggling with legacy IT infrastructure and processes that come with high operational costs and which are incapable of meeting the ever growing requirements. The entire IT organization is continually being asked to deliver more while facing immense pressure to reduce costs.

IT organizations require a new, innovative approach to efficiently meet the ever evolving business requirements. This new approach requires technological innovations that can both reduce the cost of delivering services and provide the agility to meet rapidly changing requirements.

Cloud computing is a style of IT that improves efficiency by enabling pooling and on-demand consumption of IT resources so that businesses can meet their needs in a more agile and cost-effective way. Sometimes known as utility computing, cloud computing provides a set of virtualized services from where the users can consume compute, storage, and network resources only when needed, often paying only upon usage. Cloud computing promotes agility, efficiency, and affordability for end users, and can be delivered in several different ways. Cloud infrastructure built inside of the enterprise firewall is called an internal cloud and cloud infrastructure hosted externally by a service provider for enterprise access is called an external cloud.

VMware® vCloud Director is VMware's cloud computing solution that provides the automation and management required for building both internal and external clouds. It leverages the power and flexibility of virtualization through VMware vSphere, and allows the abstraction and pooling of infrastructure resources into virtual datacenters.

The value of VMware vCloud Director, including the ability to provide cost-efficiency, agility and scalability made possible by cloud computing, can only be fully realized if the right type of storage is used as part of the solution. It requires a storage system that supports all of the attributes necessary for cloud computing without requiring sacrifices or trade-offs. The NetApp® unified storage extends all its efficiencies and flexibility already available for VMware vSphere virtualized environments to VMware vCloud Director, enabling companies to more efficiently realize the full benefits of cloud computing.

VMware vCloud Director along with NetApp unified storage enables service providers and enterprise customers to build very cost effective IT-as-a-Service (ITaaS) cloud solutions that deliver the capabilities required to meet business needs: scalability, secure multi-tenancy, efficiency, performance, ease of manageability, chargeback, high availability, and integrated data protection.

1.1 SCOPE

The solution described in this document is applicable to either a service provider building an external cloud or an enterprise building a internal cloud solution with VMware vCloud Director and NetApp unified storage. There are multiple tenants that require secure, isolated, self-service access to a virtual environment with compute, storage, and network resources. The solution allows tenants to self-provision and manage their virtual machines on demand.

This document primarily focuses on the solution design best practices. Detailed step-by-step instructions on deploying the solution is outside the scope of this document.

For detailed information on a joint NetApp, VMware, and Cisco Reference Architecture primarily focused on enterprise internal cloud use cases, refer to the NetApp Secure Multi-Tenancy (SMT) solution.

1.2 KEY TERMINOLOGIES

This section describes the VMware vCloud Director terminology used throughout this document. VMware vCloud Director has two types of resources: vSphere resources and cloud resources. For more information, see the VMware vCloud Director documentation on http://www.vmware.com/products/vcloud-director/.

- **vSphere resources**
  - vSphere resources are the vCenter™ Servers, ESX hosts, resource pools, datastores, vNetwork Distributed Switches, and port groups that are used to provision cloud resources in VMware vCloud Director.
- **Cloud resources**
  - **Cloud cells:** Cloud cells are the Red Hat Enterprise Linux® 5 (RHEL5) servers that run the VMware vCloud Director software. Multiple cloud cells form the VMware vCloud Director cluster.
  - **Provider vDC**: A provider vDC is a group of compute, memory, and storage resources from one vCenter. You can allocate portions of a provider vDC to your organizations using VMware vCloud Director .
  - **External network:** An external network uses a network in vSphere to connect to a network outside of your cloud. The network can be a public network such as the Internet or an external VPN network that connects to a given organization.
  - **Organization:** An organization is the fundamental grouping in VMware vCloud Director. An organization contains users, the vApps they create, and the resources the vApps use. An organization can be a department in your own company or an external customer to which you are providing cloud resources.
  - **Organization vDC**: An organization vDC provides an organization with the compute, memory, storage, and network resources required to create vApps.
  - **Network pool:** A network pool is a collection of VM networks that are available to be consumed by vDCs to create vApp networks and by organizations to create organization networks. Network traffic on each network in a pool is isolated at layer 2 from all other networks.
  - **vApp:** A vApp is a virtual application that contains one or more VMs.
  - **Catalog**: A catalog allows you to share vApp templates and media images with other users in your organization or with other organizations in VMware vCloud Director.

2 VMWARE VCLOUD DIRECTOR AND NETAPP SOLUTION HIGHLIGHTS

VMware vCloud Director and NetApp offer a cost-effective and scalable cloud solution with the following key benefits:

- **Scalability and secure multi-tenancy** with VMware vSphere, VMware vCloud Director, and NetApp unified storage
- **Unified storage architecture** to meet all the VMware vCloud Director tenant and infrastructure storage requirements from the same storage array
- **Operational Efficiency** with rapid provisioning of tenant vApps leveraging NetApp FlexClone®**,** storage efficiency with NetApp deduplication and thin provisioning, and being able to meet performance requirements with NetApp Flash Cache and FlexShare® capabilities
- **Ease of manageability and chargeback** with NetApp Virtual Storage Console (VSC), SANscreen® VMInsight, Operations Manager, and VMware vCenter Chargeback
- **High availability and integrated data protection** with NetApp Snapshot™ for VMware vCloud Director tenant data and infrastructure VMs

2.1 SCALABILITY AND SECURE MULTI-TENANCY

VMware vCloud Director and NetApp unified storage provide the capability to build highly scalable, elastic, and secure multi-tenancy solution.

VMware vCloud Director provides a rich set of built-in role-based access control (RBAC), enabling the cloud administrators to have a holistic view of the cloud environment and also allowing individual tenants to view and control their resources, clearly separate from those of other tenants in a multi-tenant model. The VMware vSphere compute, storage, and network resources can be pooled to serve multiple tenants, with vSphere resources dynamically assigned and reassigned according to changing customer demands. VMware handles this through the use of the ESX/ESXi host hypervisor, ESX host clusters, vSphere resource pools, vNetwork distributed switches, and VMware vCloud Director network pool. The compute cluster can easily scale up and out by adjusting the resource pool allocation and adding additional ESX servers to compute clusters, respectively.

NetApp unified storage with MultiStore® capability offers a very dynamic and scalable virtual storage infrastructure. MultiStore extends the value of VMware vCloud Director by providing the capability to divide a physical storage array into multiple, isolated virtual storage partitions called vFiler® units. IP spaces and VLAN segmentation ensure secure isolation between multiple vFiler units. Individual vFiler units can be used to provide:

- Dedicated and isolated storage access for VMs belonging to individual or multiple tenants and vCloud Director infrastructure VMs, depending on the requirements and SLAs.
- Secure, storage as a service (CIFS, NFS, and iSCSI) to different tenants in order to meet their application storage needs.

NetApp Data Motion™ allows seamless migration of an entire vFiler unit from one physical storage array to another without disrupting ongoing tenant activity. NetApp Data Motion is complementary to VMware vMotion™, making it simple to migrate data on a large scale. This provides seamless data mobility at every layer of your VMware vCloud Director infrastructure for load balancing, nondisruptive upgrades, or to satisfy other data center needs. For detailed information on MultiStore, refer to the article NetApp MultiStore: Security and Mobility for Cloud Storage.

NetApp aggregates and flexible volumes with the auto grow feature provides the capability to build a very large, scalable, virtualized pool of capacity and IOPS resources that can be shared by multiple tenants and used on demand by individual VMs. This capability can be compared to the ESX/ESXi host hypervisor that provides the capability to pool compute and memory resources to be used on demand by different VMs.

Through VMware vCloud Director and NetApp Virtual Storage Console (VSC), compute and storage resources can be rapidly provisioned and elastically grown or shrunk to provide dynamic scale-out and scale-in.

2.2 UNIFIED STORAGE ARCHITECTURE

The NetApp unified storage architecture enables customers to deploy an agile and scalable shared storage infrastructure that can meet all the VMware vCloud Director storage requirements from a single storage array.

- VMware vCloud Director tenant data (vApps, VMs, catalogs with vApp templates and media files) hosted on NFS or VMFS (FC, FCoE, or iSCSI) datastores
- Storage as a service (guest-connected storage) for tenant vApps and VMs (for example, shared NFS mounts or iSCSI LUN directly connected inside the VMs in a vApp)
- VMware vCloud Director infrastructure VMs hosted on NFS or VMFS (FC, FCoE, or iSCSI) datastores
- NFS shared storage mounted on all VMware vCloud Director server hosts as the transfer area required for uploading/downloading vApp templates and media files to/from local computers

All NetApp storage systems utilize the Data ONTAP® operating system to provide SAN (FC, FCoE, and iSCSI), NAS (CIFS and NFS), and HTTP capabilities from the same storage array. This provides a significant cost savings for building a scalable VMware vCloud Director environment on a scalable storage array.

2.3 OPERATIONAL EFFICIENCY WITH RAPID PROVISIONING, STORAGE EFFICIENCY, AND PERFORMANCE

RAPID PROVISIONING OF TENANT VMS

Cloud tenants demand a very elastic and scalable environment that allows multiple VMs to be rapidly provisioned on demand and also allows for individual tenants to do self-serviced VM provisioning as needed.

VMware vCloud Director and NetApp provide the capability to meet both these requirements with agility and high levels of storage efficiency. The solution offers the following two VM provisioning methods:

- **VMware vCloud Director administrator-directed VM provisioning**: In this scenario, the VMware vCloud Director administrator provisions the desired number and types of VMs as requested by the tenant. The VMs are created in VMware vCenter and imported as vApps into the tenant organization vDCs.
- **Tenant self-service VM provisioning**: In this scenario, the tenant administrator self-provisions the VMs and vApps in their organization vDCs.

In a VMware vCloud Director environment built on traditional storage arrays, both these operations result in a VMware full clone operation. However, because of the nature of VMware full cloning, each VM/vApp-provisioned method consumes time and requires distinct amount of storage for each VM that essentially have the same or a similar type of data (as shown in Figure 1).

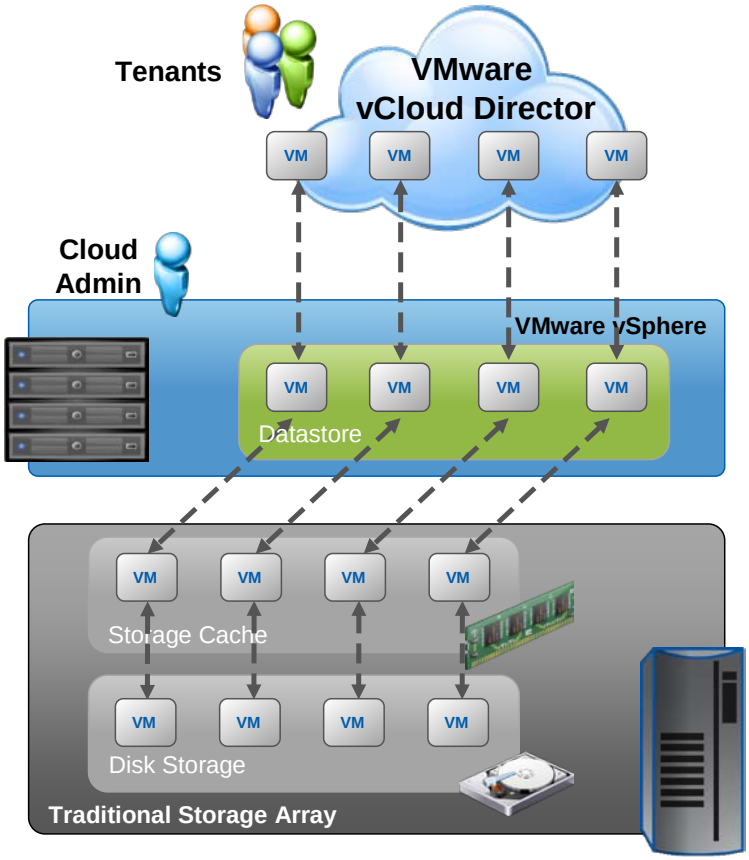

Figure 1) VM Provisioning and storage consumption on traditional storage.

In the VMware vCloud Director administrator-directed VM provisioning use case, the VMware vCloud Director administrator can efficiently meet the tenant demands by rapidly provisioning hundreds of space-efficient VMs in minutes using NetApp VSC. These VMs can then be imported into the tenant organization's vDCs. VSC leverages the NetApp FlexClone capability to provision multiple VMs for the cost of one and does not require unique storage for each cloned VMs (as shown in Figure 2). Only the new writes to the cloned VMs consume additional storage.

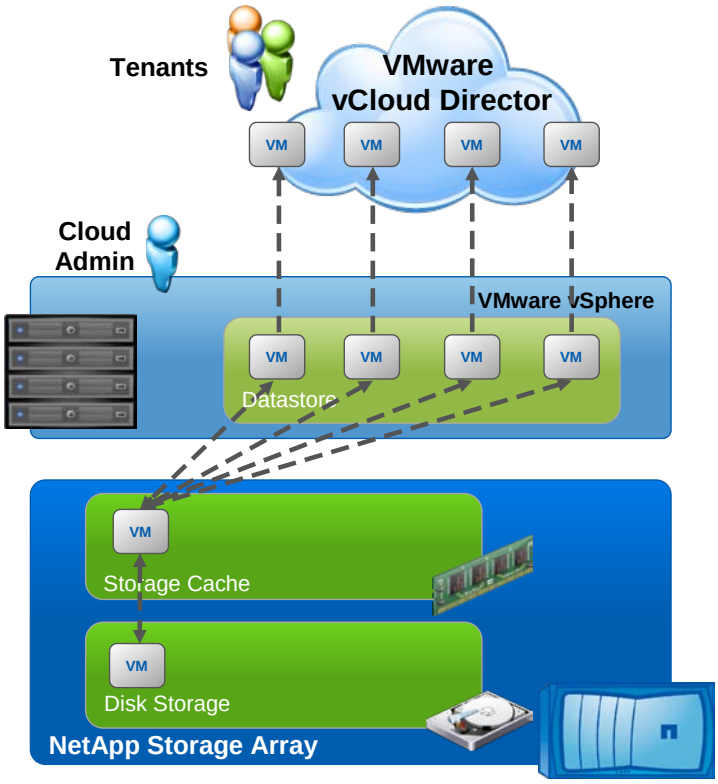

Figure 2) VM Provisioning and storage consumption on NetApp storage.

Before provisioning the VMs, make sure that the template VM or the source VM is properly aligned. For details on guest file system alignment, see NetApp TR-3747: Best Practices for File System Alignment in Virtual Environments.

STORAGE EFFICIENCY AND PERFORMANCE

NetApp and VMware vCloud Director provide multiple levels of storage efficiency (50% guaranteed[1]) with NetApp FlexClone, primary storage deduplication, and VMware and NetApp thin provisioning. These savings are achievable for both VMware vCloud Director tenant data (VMs, vApps, catalogs with vApp templates and media files) and VMware vCloud Director infrastructure VMs deployed on any of the storage protocols (FC, FCoE, iSCSI, and NFS).

As highlighted earlier in this section, NetApp FlexClone capability can be leveraged for the VMware vCloud Director administrator directed VM provisioning operation to achieve high levels of storage efficiency. However, the new writes to the VMs consume additional storage. In addition, in the tenant self-serviced VM provisioning operation leveraging VMware full clone operation, each VM consumes unique storage.

For both these scenarios, NetApp primary storage deduplication can be leveraged to achieve high levels of storage efficiency. NetApp Deduplication provides a lot of flexibility with scheduling - daily, weekly, or depending on the change rate.

Performance SLAs can be achieved with the NetApp deduplication aware Flash Cache (PAM II) and FlexShare capabilities. These capabilities strongly complement the NetApp storage efficiency capabilities by accelerating the performance of deduped data and also providing performance QoS on a per datastore basis. NetApp FlexShare provides up to five priority levels to control workload prioritization, making it possible to create a VMware vCloud Director environment with multiple “tiers of storage services”. For more information on FlexShare, see the NetApp TR-3459: FlexShare™ Design and Implementation Guide.

2.4 MANAGEABILITY AND CHARGEBACK

Customers can more efficiently manage and monitor their VMware vCloud Director infrastructure resources by leveraging NetApp Virtual Storage Console (VSC), Operations Manager, and SANscreen VMInsight. VMware vCenter Chargeback and NetApp SANscreen VMInsight solution helps provide chargeback and metering for tenant virtual infrastructure and storage resources, providing granularity up to the individual vApp and VMs level.

NETAPP VIRTUAL STORAGE CONSOLE

NetApp virtual storage console (VSC) is a VMware vCenter plug-in that enables cloud administrators to centrally perform storage operations such as monitoring, provisioning datastores, VM cloning, deduplication, backup and recovery, and data replication without requiring coordinated support from storage administrators. In addition, the plug-in includes real-time discovery, health monitoring, and capacity management for a more efficient environment. This combination of unique capabilities provides VMware administrators with the necessary tools to improve server and storage visibility and efficiencies while still enabling storage administrators to own and control storage policies. For further details on NetApp VSC, see http://www.netapp.com/us/products/management-software/vsc/virtual-storage-console.html.

NETAPP OPERATIONS MANAGER

NetApp Operations Manager provides a comprehensive monitoring and management solution for NetApp unified storage. It provides comprehensive reports of utilization and trends for capacity planning and space usage. It also monitors system performance, storage capacity, and health to resolve potential problems. For more information about Operations Manager, see www.netapp.com/us/products/management-software/operations-manager.html.

VMWARE VCENTER CHARGEBACK

User behavior is properly affected by providing transparency into actual resource usage and quantifying the metrics in financial terms. Metering involves measuring and reporting of consumption units appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, reported, and accounted for appropriately, providing transparency for both the

[1] NetApp Virtualization Guarantee* Program: http://www.netapp.com/us/solutions/infrastructure/virtualization/guarantee.html.

provider and consumer of the utilized service. VMware handles this through several technologies, including VMware vCloud Director and vCenter Chargeback.

VMware vCenter Chargeback provides the capability to account, monitor, and report on costs associated with virtual infrastructure. The key benefits of the solution are:

- **Improved resource utilization**: Associating costs to VMs eliminates many of the "unused" VMs, which frees resources for higher priorities.
- **Optimized budgets**: Business units can understand how much they are paying for resources and how much goes to unused, allowing them to optimize resource consumption and costs.

The key features of the solution are:

- Fixed, allocation, and utilization-based costing
- Charge different amounts for tiers of infrastructure
- Schedule reports and email results

For more information, refer to the documentation on www.vmware.com.

SANSCREEN VM INSIGHT

NetApp SANscreen VMInsight is a VMware vCenter plug-in that provides cross-domain visibility from the VMs to the shared storage, allowing both storage and cloud administration teams to more easily manage their VMs and infrastructure resources. VMInsight provides service-level information for VMs, physical servers, and storage devices, as well as VM volume, allocated capacity, and datastore information directly from within VMware vCenter. The enterprise-class reporting provides cloud administrators the ability to access, query, and analyze VM data. When deployed with SANscreen Capacity Manager, it enables capacity planning and storage chargeback for VMs. For more information about SANscreen VM Insight, see www.netapp.com/us/products/management-software/sanscreen/sanscreen-vm-insight.html.

2.5 HIGH AVAILABILITY AND DATA PROTECTION

VMWARE CLOUD DIRECTOR AND VSPHERE AVAILABILITY

VMware vCloud Director utilizes stateless cells that provide all of the functionality required to provide the higher order resource management and structure. For a highly available self-service architecture, a minimum of two VMware vCloud Director cells are required.

The underlying VMware vSphere infrastructure enabled by VMware HA and DRS cluster provides high availability for tenant vApps and infrastructure VMs without the need for clustering at the VM level. In the event of ESX host failure or isolation, the VMs on the failed ESX host can be automatically restarted on the other ESX hosts in the cluster. VMware DRS continuously monitors ESX host utilization and intelligently allocates available resources among VMs according to business needs. VMs are no longer tied to the underlying server hardware and can be moved across servers at any time with VMware vMotion. For more information, refer to the VMware vCloud Director and vSphere documentation on www.vmware.com.

NETAPP ACTIVE-ACTIVE STORAGE CONTROLLERS

NetApp active-active storage controllers provide the capability to achieve high levels of availability on the storage arrays. Both storage controllers in a cluster pair are active and serving data during normal production hours. In the event of one controller failing, the surviving controller takes over the personality of the failed controller, thereby providing continuous data availability.

RAID-DP

With any VMware vCloud Director deployment, data protection is critical because any storage array RAID failure could result in hundreds to thousands of tenants with downtime, resulting in lost productivity. NetApp storage arrays provide a unique level of data protection with RAID-DP®. RAID-DP provides performance that is comparable to that of RAID 10, with much higher resiliency than RAID-5. It provides protection against double-disk failure as compared to RAID 5, which can only protect against a single disk failure. For more information on RAID-DP, see NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection.

NETAPP SNAPSHOT AND SNAPMIRROR

NetApp storage array–based Snapshot and replication with NetApp SnapMirror provides a very scalable data protection solution for VMware vCloud Director tenant data and infrastructure VMs. The NetApp Snapshot backups are very fast, only consume block-level changes to each VM, and can provide multiple recovery points throughout the day. The backups happen directly at the storage array level without consuming any ESX host resources or requiring separate backup infrastructure resources.

NetApp Snapshot technology also integrates with the NetApp SnapMirror replication technology to preserve the deduplicated storage savings from the backup source to the destination storage arrays. Deduplication is not required to be rerun on the destination storage array. NetApp SnapMirror also helps achieve efficient WAN acceleration with deduplication-aware built-in compression.

3 SOLUTION DESIGN

The scalable cloud solution showcased in this document represents a flexible use case with four tenants that require efficient hosting of virtualized applications with ability to self-provision resources on demand. The design showcases two sample deployment scenarios:

- **Use case 1**: Two tenants, Org1 and Org2, represent small customers that consume resources from a shared pool of compute, network, and storage resources. These are the tenants with flexible resource-sharing requirements.
- **Use case 2**: Two tenants, Org3 and Org4, are larger tenants that require end-to-end isolation of logical resources to meet their scalability needs. These are the tenants with dedicated logical resource assignment requirements. Org4 also has the additional storage as a service requirement (guest-connected storage) for some individual VMs.

All the four tenants require a set of VMs to be pre-provisioned by the VMware vCloud Director administrator. All the tenants also require access to a shared catalog with sample vApp/VM templates and media files to be used to deploy additional VMs as needed. Also, each tenant requires a percentage of its vApps to have high levels of compute, memory, performance, and storage resources.

3.1 HIGH-LEVEL ARCHITECTURE

Figure 3 depicts the high-level VMware vCloud Director solution architecture with the four tenants (Org1, Org2, Org3, and Org4) mentioned previously. For tenants to self-provision new vApps on demand, the cloud administrator has created an organization called “Published Catalogs.” This organization contains one org vDC with catalogs containing VM/vApp templates and media files that are sharable by all the tenants. The tenants can provision vApps on demand by leveraging the vApp templates and media files in the published catalog.

To meet the compute, memory, and storage requirements for different VMs, the cloud administrator has created tiers of service (virtual datacenters) in VMware vCloud Director, called Provider vDCs. The different provider vDCs contain compute, memory, and storage resources with varying SLAs.

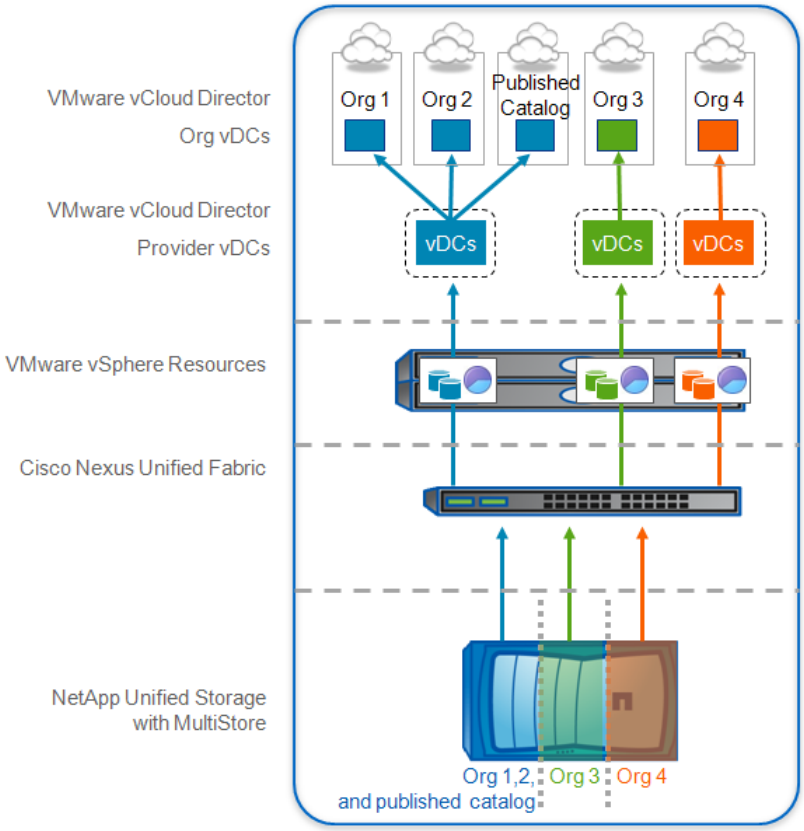

Figure 3) High-level solution architecture.

3.2 DETAILED SOLUTION ARCHITECTURE

Figure 4 represents the detailed end-to-end solution architecture.

The solution architecture is based on the concept of VMware vCloud Director tenant resources (called VMware vCloud Director resource groups) and VMware vCloud Director provider infrastructure resources (called VMware vCloud Director Management cluster). This separation of tenant and provider resources provides a greater level of security and repeatable, independent scaling of tenant and management resources.

VMWARE VCLOUD DIRECTOR RESOURCE GROUP

A VMware vCloud Director resource group contains clusters of ESXi (ESX) hosts under management by a vCenter server that are dedicated to providing vSphere resources to tenants. Resource pools/sub resource pools are created in the vSphere clusters to provide compute resources to the provider vDCs in VMware vCloud Director.

VMWARE VCLOUD DIRECTOR MANAGEMENT CLUSTER

The VMware vCloud Director management cluster is anchored by the vCenter server and includes the VMware vCloud Director cell(s), vCenter server, vCenter Chargeback, vShield manager, and Operations Manager. The VMware vCloud Director management cluster, for governance/security requirements, is placed within a dedicated management cluster of ESX/ESXi hosts managed by a vCenter server. The management cluster is used to house the infrastructure VMs for the environment. This includes:

- VMware vCenter Server
- VMware vCenter SQL Server®
- VMware vCenter Chargeback server
- VMware vCenter Chargeback SQL Server
- Two or more VMware vCloud Director cells
- NetApp Operations Manager
- NetApp SANScreen
- AD/DNS servers

TENANT ORG LAYOUT

- As discussed previously, Org1 and Org2 are relatively smaller tenants and require only a few vApps, and without any requirement of end-to-end data isolation. Therefore, the org vDCs containing vApp data for Org1 and Org2 are provisioned from the same set of gold, silver, and bronze provider vDCs (blue set of provider vDCs) as shown in Figure 4**.** The datastores attached to these provider vDCs are provisioned from a dedicated vFiler unit on the NetApp storage array using NetApp VSC.
- Org3 and Org4 require more vApps when compared to Org1 and Org2 and also require end-to-end data isolation. For ease of manageability, data mobility, and end-to-end isolation with large tenants, the cloud administrator has decided to host the datastores for Org3 and Org4 on separate vFiler units on the NetApp storage array, exposed through gold and silver provider vDCs (“green” and “orange” sets) as shown in Figure 6.. This allows end-to-end isolation and better manageability of the environment for large tenants, while offering tiers of storage services with predefined SLAs from the same unified storage array.
- In addition, Org4 also has multiple RHEL VMs that require shared NFS storage accessed by all the VMs in a vApp. Org4 also has a SQL Server vApp that requires advanced transaction log-level recovery. Therefore, the Org 4 is also using the storage as a service consumption model to obtain some NFS exports for RHEL5 VMs, and guest-connected LUN using Microsoft® iSCSI software initiator, for SQL server databases, being backed up by NetApp SnapManager® for SQL Server (SMSQL). The NFS exports and iSCSI guest connected LUNs are provisioned on the vFiler dedicated to Org 4.

The VMware vCloud Director infrastructure VMs are hosted on datastores provisioned and managed on a separate vFiler unit in the NetApp storage array (as shown in Figure 5). This allows secure isolation of infrastructure VMs and no downtime with storage migrations and non-disruptive upgrades. For details on the

NetApp best practices for VMware vSphere deployments, see [TR-3749: NetApp and VMware vSphere 4 Storage Best Practices](#).

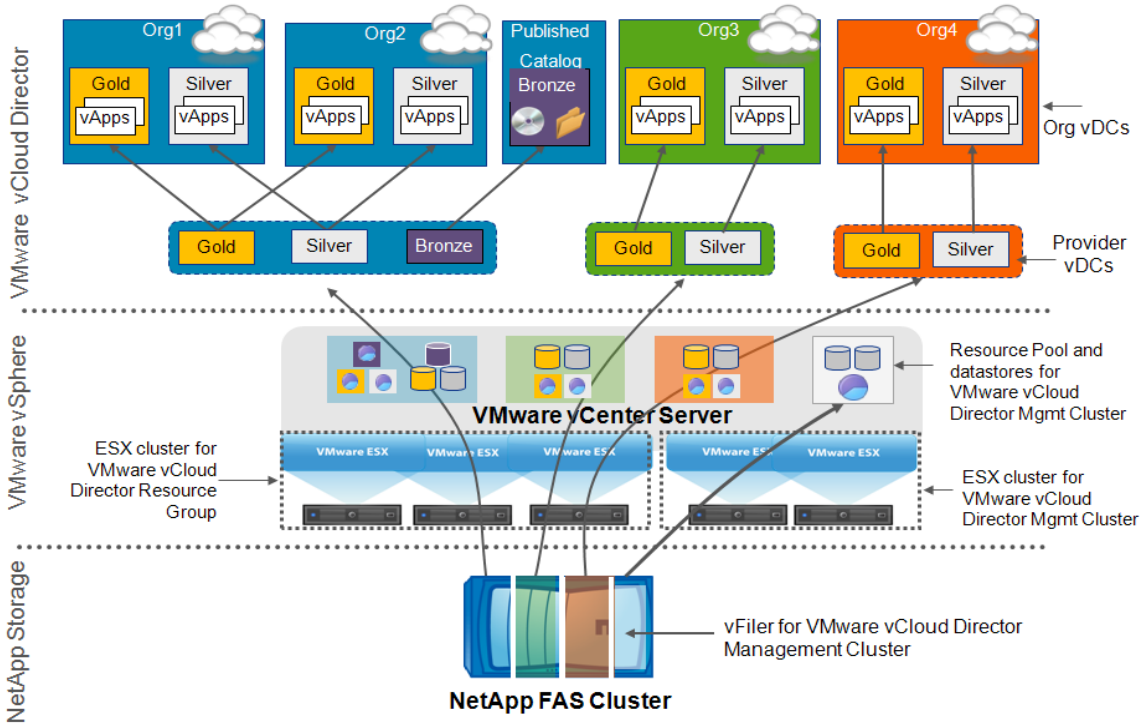

Figure 4) Detailed solution architecture.

3.3 SOLUTION TECHNOLOGY COMPONENTS

This section describes each of the physical and virtual building blocks that make up this solution.

Table 1 lists the software components used in this solution.

Table 1) Software components.

Solution Component
VMware vCloud Director Abstracts and coordinates underlying resources. This includes: • VMware vCloud Director Server (minimum of two instances, each installed on a Linux VM and referred to as a "VMware vCloud Director cell") • VMware vCloud Director Database (one instance per clustered set of VMware vCloud Director cells; Oracle® RAC is recommended for high availability) **VMware vSphere 4.1** Foundation of underlying cloud resources. This includes: • VMware ESX hosts (two for VMware vCloud Director management cluster, and three for VMware vCloud Director resource group) • vCenter Server • vCenter Server Database (one instance per vCenter Server) • VMware vShield, providing network security services, including NAT and firewall. This includes vShield Manager (one instance per vCenter server) and vShield Edge (deployed automatically by VMware vCloud Director as virtual appliances on ESX hosts)

Solution Component
VMware vCenter Chargeback 1.5 Provides resource allocation, metering, and chargeback models. It includes: • vCenter Chargeback Server (one instance) • Chargeback Data Collector (one instance; more instances can be installed as environment scales) • vCloud Data Collector (one instance; more instances can be installed as environment scales) • VSM Data Collector (one instance; more instances can be installed as environment scales)
Storage Protocol • vApps OS/application binaries/data: NFS datastores (Note that iSCSI, FC, or FCoE VMFS datastores can also be used for the solution) • Linux vApp requiring access to shared storage: NFS exports mounted directly inside the vApps • SQL Server vApp requiring advanced transaction-level backup and recovery: iSCSI LUN directly mounted inside the VMs in the vApp
NetApp • Data ONTAP 7.3.4 • FCP or NFS, iSCSI, deduplication, MultiStore, FlexClone, FlexScale®, SMSQL, SnapMirror, SnapRestore®, SnapDrive® for Windows®, and NearStore® licenses
NetApp Management Software NetApp Virtual Storage Console (VSC) 2.0.1, SnapManager for SQL Server 5.0, SnapDrive for Windows 6.3, SANscreen 6.0, Operations Manager 4.0, Systems Manager 1.01

Table 2 lists the hardware components used in this solution validation.

Table 2) Hardware components.

Solution Component
(5) ESX hosts with 16GB RAM and (1) quad core CPU per host; (2) 10Gb Ethernet ports
(1) NetApp FAS3140HA with (2) DS4243 shelves; (24) 300GB 15K RPM SAS drives per DS4243 shelf; (1) 256GB Flash Cache (PAM II) card per controller
(2) Cisco Nexus 5020 switches with multiple 10Gb Ethernet ports

3.4 VMWARE VSPHERE ARCHITECTURE

Figure 5 shows the details of the VMware vSphere architecture for hosting the VMware vCloud Director tenant and infrastructure VMs. Separate ESX clusters were used to host the tenant VMs (VMware vCloud Director resource group) and the VMware vCloud Director management VMs (VMware vCloud Director management cluster).

- **VMware vCloud Director management cluster**: A single management ESX cluster containing all core components and services needed to run the cloud.
- **VMware vCloud Director resource group**: A resource group that provides dedicated resources for consumption by tenants. The VMware vCloud Director resource group is a cluster of ESX hosts managed by the vCenter Server and is under the control of VMware vCloud Director.

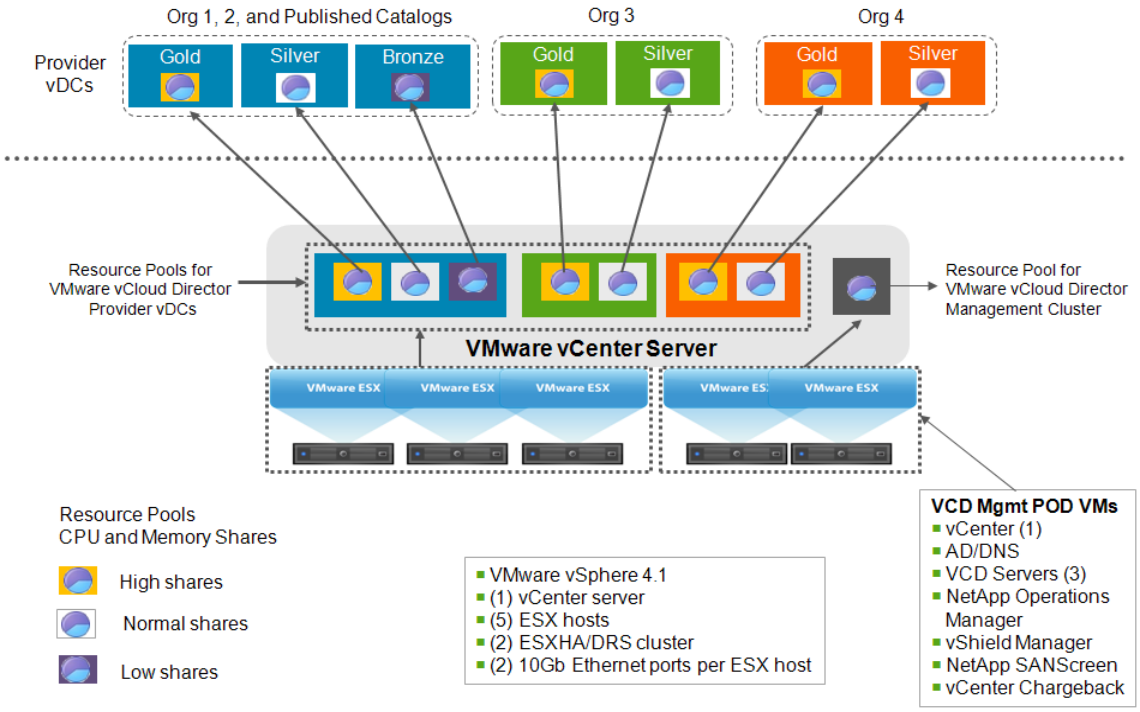

Figure 5) VMware vSphere architecture.

For the VMware vCloud Director resource group, separate resource pools were created for different provider vDCs with different levels of CPU and memory shares based on the VMware vCloud Director service offerings (for example, gold, silver, and so on).

- Resource pools with “high” CPU and memory shares were created for the “gold” provider vDCs.
- Resource pools with “normal” CPU and memory shares were created for the “silver” provider vDCs.
- Resource pools with “low” CPU and memory shares were created for the “bronze” provider vDCs.

3.5 STORAGE ARCHITECTURE

This configuration used one NetApp FAS3140HA with active-active controllers, two DS4243 shelves (300GB 15K RPM SAS drives), and one 256GB Flash Cache (PAM II) card per controller.

Figure 6 shows the logical storage layout for hosting the tenant data and VMware vCloud Director management VMs. In this deployment, all tenant data (vApps and catalogs) was hosted on provider vDCs on NFS datastores provisioned on the relevant vFiler units on a NetApp storage array. Each NetApp FAS3140 storage controller has one 22-disk aggregate with RAID-DP to provide protection against double-disk failure. The 22-disk aggregate provided the capability to pool IOPS and storage capacity to be used on demand by different tenant workloads. This is a unique capability that is a significant value add for dynamic and elastic cloud deployments.

NetApp VSC 2.0.1 was used to provision NFS datastores (500GB each) from the different vFiler units and connect to the ESX clusters. All the volumes on the NetApp storage array were thin provisioned with primary storage deduplication and volume autogrow enabled to allow the datastores to grow automatically as capacity demands increase. The datastores were named so that it is easy to identify the provider vDC to which they are mapped: for example, PvDC_Gold_Storage1. NetApp also allows provisioning datastores by leveraging the storage service catalog capability in Provisioning Manager. For details on how to leverage this capability, consult your NetApp representatives.

While provisioning the provider vDCs, we also made sure that the same NFS datastores were not shared by multiple provider vDCs. This allowed better end-to-end data manageability and ease of troubleshooting in case of performance or scalability issues.

- The NFS datastores belonging to Org1, Org2, Org3, and published catalogs were provisioned and mapped through the vFiler units on controller A.
- The NFS datastores belonging to Org4 and management VMs are provisioned and mapped through the vFiler units on controller B. Also, the NFS volume mounted as NFS export to multiple RHEL VMs in a vApp and the GOS-connected LUNs inside one of the SQL Server vApps are also provisioned on the vFiler unit for Org4.

One 256GB Flash Cache (PAM II) card was used in each storage controller to accelerate the performance of the heavily deduped vApp data.

FlexShare was enabled on all volumes to provide performance QoS at the individual datastore level.

- All NFS datastores mounted to the gold provider vDCs had FlexShare priority of “very high.”
- All NFS datastores mounted to the silver provider vDCs had FlexShare priority of “medium.”
- All NFS datastores mounted to the bronze provider vDCs had FlexShare priority of “very low.”
- The following NetApp volumes had FlexShare priority “very low”:
    - Shared NFS storage mounted directly to the VMs in a vApp for Org4
    - Shared NFS storage mounted to all the VMware vCloud Director server hosts as the transfer area required for uploading/downloading vApp templates and media files

All the volumes had the FlexShare system and cache setting as “default.” For more information on FlexShare, see http://partners.netapp.com/go/techontap/flexshare.html.

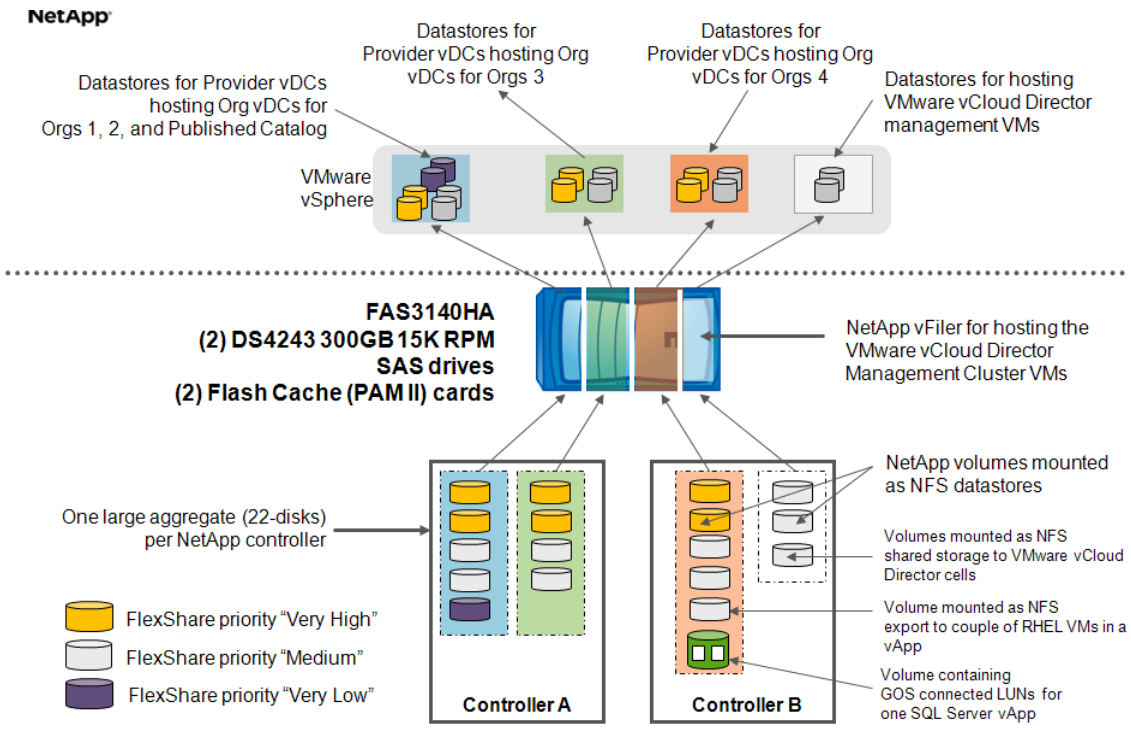

Figure 6) NetApp storage architecture.

Note: The exact sizing (ESX hosts, network, storage, and so on) is not within the scope of this document. The exact sizing for your environment will be based on your custom requirements. When sizing for your environment, consult your NetApp and VMware representatives for the exact solution configuration.

3.6 NETWORK ARCHITECTURE

Each ESX host had two 10Gb Ethernet ports configured with multiple port groups for vApps, VM networks, org networks, vMotion and storage VMkernels, and service console traffic. A VMware vNetwork Distributed Switch was used for the storage, VM networks, VMware vCloud Director external network, org networks, and vApp networks.

VMWARE VCLOUD DIRECTOR STORAGE NETWORK ARCHITECTURE

Figure 7 shows the storage network architecture for this solution. The VMware vNetwork distributed switch used in this configuration provides ease of manageability, scalability, and automatic provisioning of port groups by VMware vCloud Director for creating secure organization networks. For secure data isolation, separate VMkernel ports (on different private VLANs) were configured for each of the following entities:

- NFS datastores for Org1, Org2, and published catalog were mounted on a VMkernel port vmk1 with private VLAN 460
- NFS datastores for Org3 were mounted over a VMkernel port vmk2 with private VLAN 461
- NFS datastores for Org4 were mounted over a VMkernel port vmk3 with private VLAN 462
- NFS datastores for hosting VMware vCloud Director management VMs were mounted over a VMkernel port vmk4 with private VLAN 463VMware vCloud Director

A separate port group was created (on private VLAN 464) for connecting to the external network for NFS mounts insides the VM and also for iSCSI guest-connected LUNs.

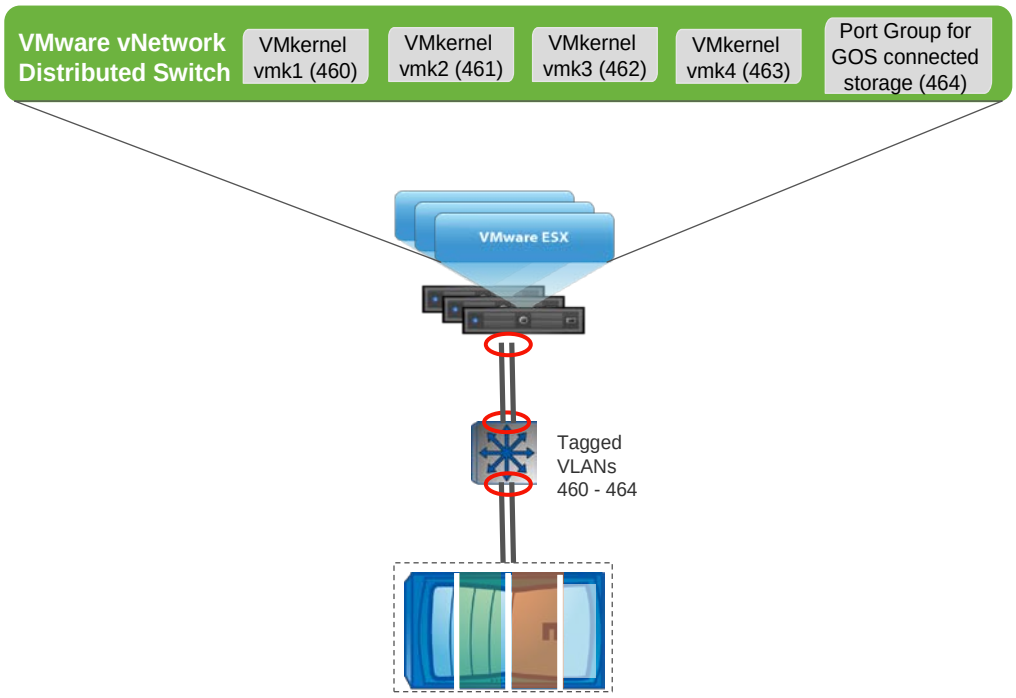

Figure 7) VMware vCloud Director storage network architecture.

VMWARE VCLOUD DIRECTOR ORG AND VAPP NETWORK ARCHITECTURE

Org1, Org2, and Org3 had a requirement for an internal organization network with vApps not requiring connectivity to the external network. Therefore, an internal network was created for each of these orgs leveraging the "VMware vCloud Director network isolation-backed" network pool capability in VMware vCloud Director. The creation of an internal org network automatically created a vShield edge virtual appliance for each network. Org4 had additional external network requirements for Web servers requiring access to an outside network. Therefore, an additional external organization network (NAT-routed) was created for this org and connected to the external network. This also automatically provisioned a vShield Edge virtual appliance in vSphere.

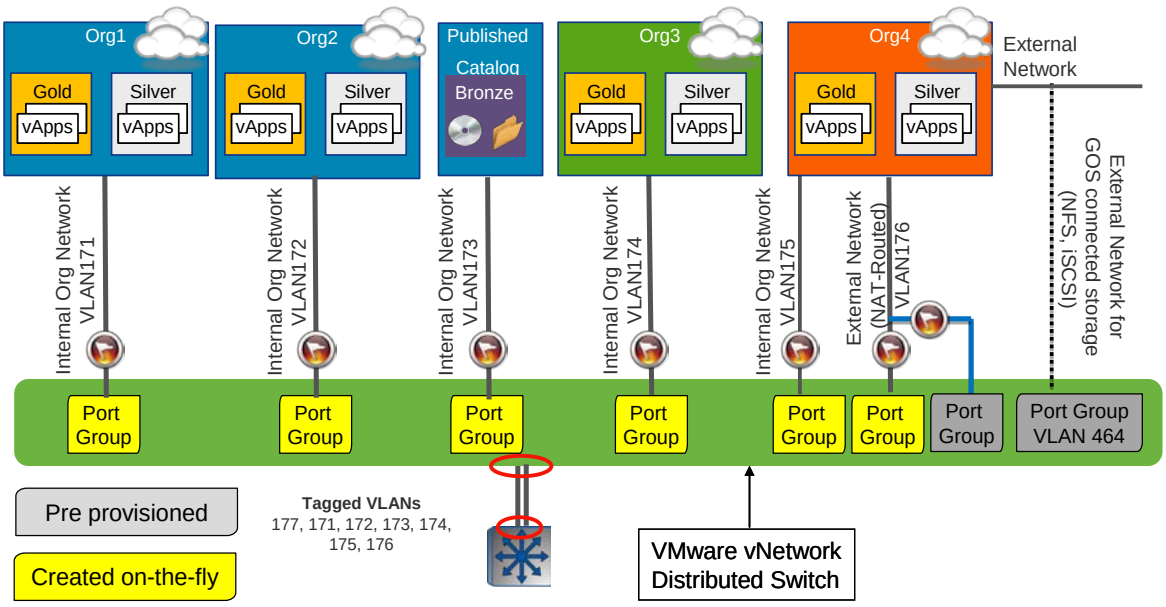

Figure 8) VMware vCloud Director org network architecture.

VAPP NETWORK

Shown in Figure 9 is the sample vApp network for the vApps in Org4 that require different types of networks.

- For network connectivity between different vApps in the organization, VMs in the vApps are connected to the internal organization network. VMs in Org1, Org2, and Org3 also have a requirement for this type of vApp internal networking.
- For VMs requiring NAT-routed connectivity to the outside world (for example, Web servers), VMs in certain vApps are connected to the external org network (NAT routed).
- For VMs that required storage as a service (NFS mounts and iSCSI LUNs directly connected inside the guest VMs), one additional network - “external org network (direct connect)”, was created and connected to the public network.

For information on other networking options available in VMware vCloud Director, see the VMware vCloud Director documentation on www.vmware.com.

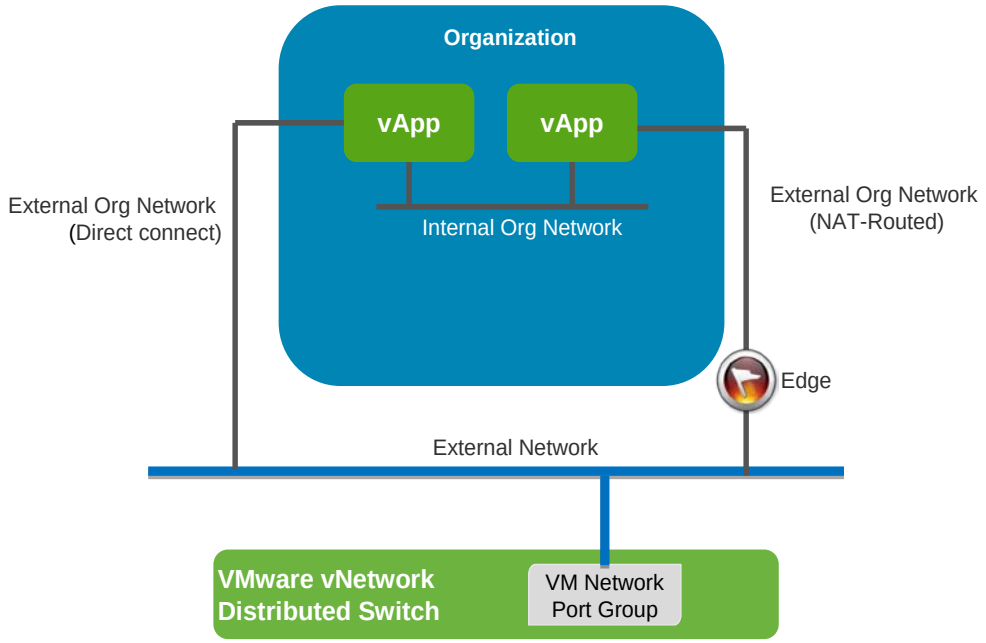

Figure 9) vApp networking.

3.7 BACKUP AND RECOVERY

NetApp storage array Snapshot and SnapMirror technologies were leveraged to provide point-in-time backup, recovery, and remote replication for the VMware vCloud Director environment. Snapshot backups and remote replication were set up at the datastore level with the capability to restore individual VMs. NetApp solution also allows setting backup schedules and retention policies. Deduplication savings were preserved with NetApp SnapMirror replication from the backup source to the destination storage array. NetApp SnapMirror also helped achieve efficient WAN acceleration with deduplication-aware built-in compression. For details on how to leverage the NetApp Snapshot backup solution for the VMware vCloud Director environment, contact your NetApp representative.

4 SOLUTION VALIDATION

Table 3) Solution validation.

Solution Test	Results
Scalable and secure multi-tenancy with VMware vCloud Director and NetApp unified storage	The solution provided secure, end-to-end isolation for different VMware vCloud Director tenants. The tenants were able to securely access only their data with end-to-end QoS.
Unified storage architecture	All the VMware vCloud Director storage requirements were successfully met from the same NetApp FAS3140HA controller pair. • VMware vCloud Director tenant VMs • Tenant VMs 'storage as a service' requirement for guest connected storage (e.g. NFS exports, CIFS shares, iSCSI guest connected LUNs) • VMware vCloud Director management/infrastructure VMs • VMware vCloud Director shared storage required for upload/download of vApp templates and media files
Operational Efficiency	Rapid provisioning of tenant vApps leveraging NetApp FlexClone® • Successfully provisioned multiple space efficient VMs in minutes using NetApp VSC. The VMs were then imported into VMware vCloud Director tenant Org vDCs. Storage efficiency and Performance • Invoked and managed deduplication of datastores directly from within vCenter using NetApp VSC. • Achieved more than 50% storage efficiency across the different datastores, with up to 90% in some cases. • Flash Cache (PAM II) provided the capability to accelerate performance of the heavily deduped data, and FlexShare provided the capability to build storage service tiers.
Manageability and Chargeback with NetApp SANscreen, VMInsight, Virtual Storage Console (VSC), Operations Manager, Systems Manager, and VMware vCenter Chargeback	• Successfully created thin-provisioned NFS datastores on different vFiler units on the NetApp storage array directly from VMware vCenter (with autogrow enabled). • Successfully provisioned space-efficient VMware vCloud Director infrastructure and tenant VMs using NetApp VSC, leveraging NetApp FlexClone capabilities on the storage array. • Monitored the end-to-end virtual and physical infrastructure with reporting and alerting using NetApp SANscreen. • VMware vCenter Chargeback and SANscreen VM Insight were able to successfully perform metering and Chargeback for individual VMs, storage, and networks in VMware vCloud Director. • Successfully managed storage and provided detailed alerting with NetApp Operations Manager.
High availability and data protection for VMware vCloud Director tenant data and infrastructure VMs	• VMware HA and DRS provided the desired high availability and workload balancing for the vSphere infrastructure. • Successfully backed up and restored VMware vCloud Director infrastructure and tenant VMs. • Successfully sustained multiple disk failure and controller failure on the NetApp storage array.

5 SUMMARY

To summarize, VMware vCloud Director provides the automation and management required for building internal and external clouds. The value of VMware vCloud Director, including the cost efficiency and the business justification to move to clouds, can only be fully realized if the right type of storage is used in the solution. It requires the storage solution to be intelligent, agile, scalable, and cost efficient without any negative tradeoffs. This is where NetApp unified storage extends the value of VMware vCloud Director, as highlighted in this paper.

VMware vCloud Director along with NetApp unified storage enables service providers and enterprise customers to build a very cost-effective IT-as-a-service (ITaaS) cloud solution that has the desired scalability and secure multi-tenancy, efficiency and performance, ease of manageability and chargeback, high availability, and efficient data protection.

The key benefits of the VMware vCloud Director and NetApp solution are:

- **Scalability and secure multi-tenancy** with VMware vSphere, VMware vCloud Director, and NetApp unified storage
- **Unified storage architecture** to meet all the VMware vCloud Director tenant and infrastructure storage requirements from the same storage array
- **Operational Efficiency** with rapid provisioning of tenant vApps leveraging NetApp FlexClone®**,** storage efficiency with NetApp deduplication and thin provisioning, and being able to meet performance requirements with NetApp Flash Cache and FlexShare® capabilities
- **Ease of manageability and chargeback** with NetApp Virtual Storage Console (VSC), SANscreen VMInsight, Operations Manager, and VMware vCenter Chargeback
- **High availability and integrated data protection** for VMware vCloud Director tenant data and infrastructure VMs

6 ACKNOWLEDGEMENTS

The authors of the document would like to thank the following for their contribution to the design, validation, and creation of this solution guide:

VMware: Pang Chen, John Arrasjid, Mike Dipetrillo

NetApp: Eric Forgette, Sudhir Srinivasan, Lee Dorrier, Radhika Krishnan, Esther Smitha, Greg Kleiman, Graham Smith, Jim Lyons, Mike Zimmerman, Spencer Sells

7 FEEDBACK

If you have questions or comments about this document, contact xdl-vgibutmevmtr@netapp.com.

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.

www.netapp.com