White Paper

Delivering IT as a Service for Virtualized Data Centers

NetApp
August 2009 | WP-7083-0809

TABLE OF CONTENTS

1 INTRODUCTION

As the IT landscape rapidly changes, cost reduction pressures, focus on time to market, and employee empowerment are compelling enterprises and IT service providers to develop innovative strategies to address these challenges. The cloud computing approach to IT service delivery has become a key area requiring attention.

The abstraction of compute, network, and storage infrastructure is the foundation of cloud computing. The infrastructure is a service, and its components must be readily accessible and available to the immediate needs of the application stacks it supports. Cloud computing removes the traditional application silos within the data center and introduces a new level of flexibility and scalability to the IT organization. This flexibility helps address challenges facing enterprises and IT service providers that include rapidly changing IT landscapes, cost reduction pressures, and focus on time to market.

Cloud computing is generally described as IT delivered as a service (ITaaS). Three primary ITaaS cloud-computing models have emerged to address particular customer use cases. These models offer a new level of automation and security to successfully achieve application stack autonomy and customer business objectives. The ITaaS cloud models include:

- Infrastructure as a service offers compute resources, operating systems, networking, and storage to client environments. Clients are typically responsible for all operations in, administration of, and configuration of their environment beyond the baseline deployment.
- Storage as a service (StaaS) offers storage provisioning from pooled resources for use as application datasets, backup, and content repositories.
- Software as a service (SaaS) delivers a turnkey, managed application offering for direct use, typically hosted in the service provider's data center.

Many enterprises and IT service providers are developing cloud service offerings for public and internal customer applications. Regardless of whether the focus is on public or private cloud services, these pursuits share many objectives:

- Increase operational efficiency through cost-effective use of expensive infrastructure.
- Drive up economies of scale through shared resourcing.
- Rapid and agile deployment of customer environments or applications.
- Improve service quality and accelerate delivery through standardization.
- Promote green computing by maximizing efficient use of shared resources, lowering energy consumption.

Achieving these goals can have a profound, positive impact on profitability, productivity, and product quality. However, leveraging shared infrastructure and resources in a cloud-services architecture introduces additional challenges, hindering widespread adoption by IT service providers who demand securely isolated customer or application environments but require highly flexible management.

As enterprise IT environments have dramatically grown in scale, complexity, and diversity of services, they have typically deployed application and customer environments in silos of dedicated infrastructure. These silos are built around specific applications, customer environments, business organizations, operational requirements, and regulatory compliance (Sarbanes-Oxley, HIPAA, PCI), or to address specific proprietary data confidentiality. For example:

- Large enterprises need to isolate HR records, finance, customer credit card details, etc.
- Resources externally exposed for outsourced projects require separation from internal corporate environments.
- Healthcare organizations must insure patient record confidentiality.
- Universities need to partition student user services from business operations, student administrative systems, and commercial or sensitive research projects.
- Telcos and service providers must separate billing, CRM, payment systems, reseller portals, and hosted environments.
- Financial organizations need to securely isolate client records and investment, wholesale, and retail banking services.
- Government agencies must partition revenue records, judicial data, social services, operational systems, etc.

Enabling service providers to migrate such environments to a cloud architecture demands the capability to provide secure isolation while still delivering the management and flexibility benefits of shared resources. External cloud service providers must enable all customer data, communication, and application environments to be securely isolated from other tenants. With external clouds, the separation must be so complete and secure that the tenants can have no visibility of each other. Internal cloud service providers must deliver the secure separation required by their organizational structure, application requirements, or regulatory compliance.

However, lack of confidence that such secure isolation can be delivered with resilient resource management flexibility is a major obstacle to the widespread adoption of cloud service models. NetApp, Cisco, and VMware have partnered to create a unique infrastructure solution that incorporates comprehensive storage, network, and computing technologies that facilitate dynamic, shared-resource management while maintaining a secured and isolated environment. VMware® vSphere, Cisco Nexus, and NetApp® MultiStore® with NetApp Data Motion™ deliver a powerful solution to fulfill the demanding requirements for secure isolation and flexibility in cloud deployments of all scales.

One of the main differences between traditional shared hosting (internal or external) and a typical IaaS cloud service is the level of control available to the user. Traditional hosting services provide users with general application or platform administrative control, whereas IaaS deployments typically provide the user with broader control over the compute resources. The secure cloud architecture further extends user control end to end throughout the environment: the compute platform, the network connectivity, storage resources, and data management. This architecture enables service providers and enterprises to securely offer their users unprecedented control over their entire application environment. Unique isolation technologies combined with extensive management flexibility delivers all of the benefits of cloud computing for IT providers to confidently provide high levels of security and service for multi-tenant customer and consolidated application environments.

This overview details the features and technologies that NetApp, Cisco, and VMware provide to deliver a fully featured cloud infrastructure solution that is as secure, reliable, and powerful as traditional silo deployments, while achieving the flexibility, efficiency, and OPEX/CAPEX reduction benefits of a cloud architecture.

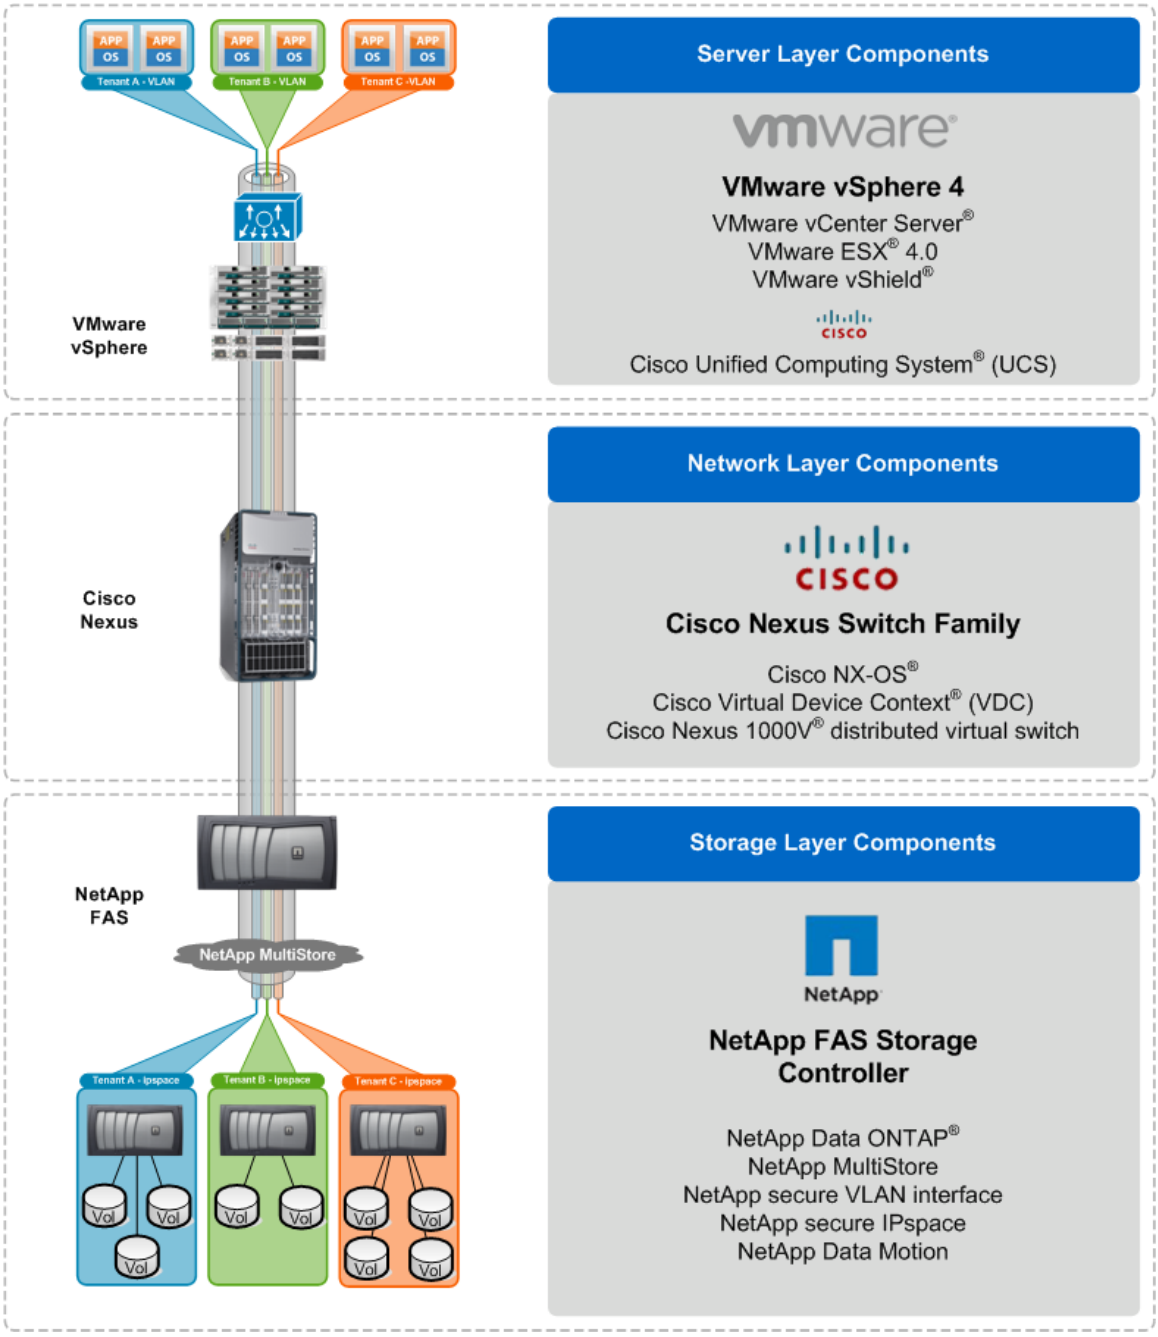

Figure 1) Secure cloud architecture overview.

2 ISOLATION AND SECURITY

In legacy deployments, environments that needed additional security and isolation required a separate set of equipment and infrastructure. While this provided the required separation, efficiency dropped significantly due to lower utilization, while costs increased due to higher management and capital expenditures. The secure cloud architecture addresses these concerns by developing a secure, isolated multi-tenant solution at every layer: storage, network, and server.

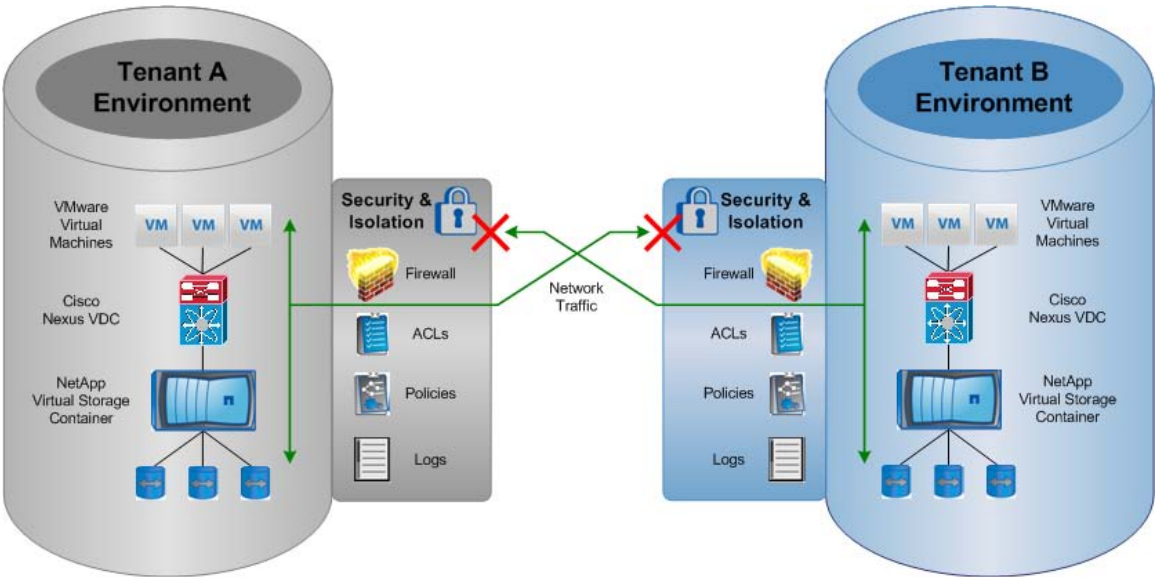

Figure 2) Isolation within the secure cloud architecture.

The goals of this isolation and security focus are confidentiality and integrity of data throughout the environment. Each component in the secure cloud architecture offers unique features to address these needs. The following sections detail these for storage, networking, and compute resources.

CONFIDENTIALITY

In order to provide confidentiality, communication and data inside a virtual environment must be isolated and secure. The industry has accepted the ability to virtualize storage resources through NetApp MultiStore, networking through Virtual LANs (VLANs) and Virtual Device Contexts (VDCs), and server resources through hypervisors within vSphere.

NetApp **MultiStore** allows cloud providers to quickly and easily create separate and completely private logical partitions on a single NetApp storage system as discrete administrative domains called virtual storage containers. These virtual storage containers have the effect of making a single physical storage controller appear to be many logical controllers. Each virtual storage container can be individually managed with different sets of performance and policy characteristics. Providers can leverage NetApp MultiStore to enable multiple customers to share the same storage resources with minimal compromise in privacy or security, and even delegate administrative control of the virtual storage container directly to the customer.

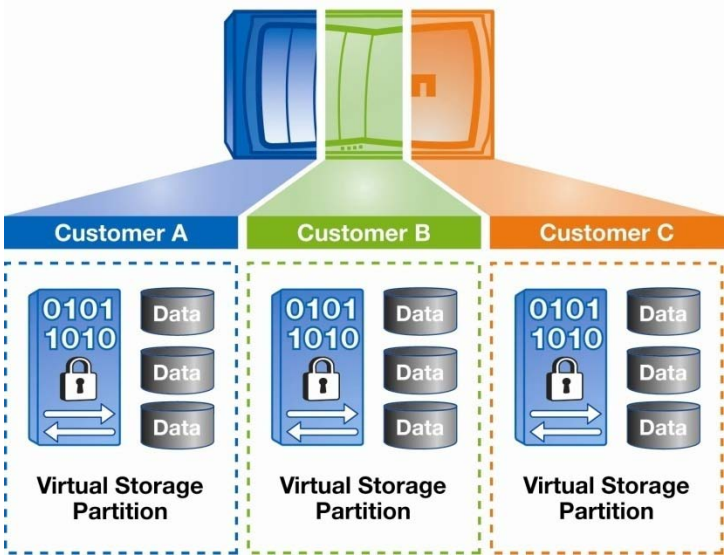

Figure 3) NetApp MultiStore.

Multi-tenant service providers commonly use **Virtual LANs (VLANs)** to logically segment and secure a switched network according on an organizational basis, function, project team, or application. Using VLANs in such cases enables organization-based networks to be individually secured regardless of their physical connections to the network or the various arrangements of those connections. VLAN segmentation allows the network to be reconfigured according to the needs of individual tenant groups through logical methods instead of physically moving network hardware and connections. VLAN technology logically segments the switched infrastructure into separate Layer-2 broadcast domains, enabling service providers to assign one or more discrete network address spaces to each tenant and manage them as an organizational unit. This also offers the ability to provide differing services on each logical segment.

NetApp VLAN interfaces create private networking segments in which each interface can be bound to an **IPspace**. An IPspace is a discrete, private, and secure logical networking partition within the NetApp array that represents a private routing domain. MultiStore virtual storage containers are created and bound to one or more IPspaces. As each new customer is added to the multi-tenant infrastructure, IPspaces allow easy consolidation of multiple separate customers onto one shared storage platform, even if those customer environments contain duplicate IP addresses or hostnames. There is also no need to build or manage static routes at the network or storage layer.

Security between individual MultiStore containers is provided by the NetApp Data ONTAP® operating system. In 2008, NetApp commissioned a security audit of MultiStore by Matasano Security validating the security and isolation capabilities of MultiStore. Access to this analysis is available at the link below:

http://www.netapp.com/us/library/white-papers/wp-multistore-analysis.html

The Cisco Nexus series of switches offers security features tailored to virtual environments, allowing consistent visibility, control, and isolation of the application stack within a multi-tenant compute cloud.

The Cisco Nexus 7000 Series switches running NX-OS have introduced the capability to divide a single physical switch into up to four virtual switches, referred to as **virtual device contexts (VDCs)**, for a more secure method of isolation**.** Each VDC operates similar to a standalone switch with a distinct configuration file, complement of physical ports, and separate instances of Layer-2 and Layer-3 services available. This feature provides the option to use a single physical switch pair to serve multiple roles and/or tenants within the environment. Different VDC design options can use this feature for service integration, enhanced security, administrative boundaries, and flexibility of hardware deployment during changing business needs.

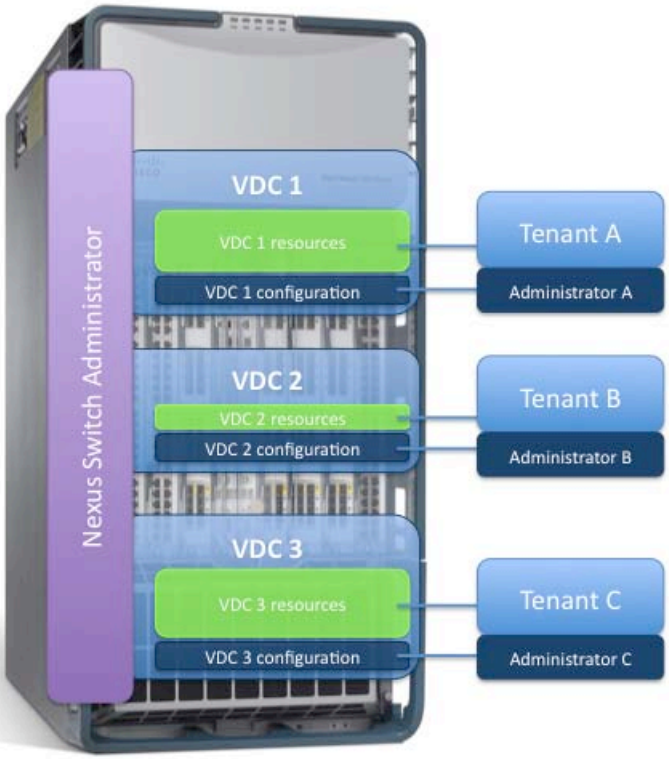

Figure 4) Cisco virtual device contexts.

Also supported are multiple virtual routing instances per VDC, further enhancing the abstraction of Layer-3 forwarding and routing services in the cloud. The deployment of **virtual routing and forwarding (VRF)** provides the capability to direct traffic independently through multiple virtual contexts (VDCs), providing more granular Layer-3 isolation within a VDC for the cloud tenants and their independent application stacks.

IP and MAC-based access controls (ACLs) are filters that provide stateless traffic filtering based on IP address or the MAC (Media Access Control) physical hardware address. This provides the ability to control and monitor virtual machine traffic flows within the virtual and physical cloud infrastructure.

Control Plane Policing (CoPP) is a Cisco feature designed to allow users to manage the flow of traffic handled by the route processor of their network devices. CoPP is designed to prevent unnecessary traffic from overwhelming the route processor that, if left unabated, could affect system performance. This can protect all tenants of the cloud environment by maintaining the integrity of the cloud infrastructure.

The necessity to secure the application stack between virtual machine nodes is critical and possible in a Nexus-enabled cloud. The Cisco **Nexus 1000V** Series switches provide security services to the individual virtual machine as well as policy portability, so the proper network and security policies follow every virtual machine as it moves around the data center. This isolation allows individual access restrictions to be placed on each virtual machine, as if they were completely separate physical machines.

Cisco **VN-Link** bridges the server, storage, and network management domains, so changes in one environment are communicated to the others. Cisco VN-Link helps enable new capabilities and features, simplifies management and operations, and allows scaling for server virtualization solutions. Specific benefits include:

- Real-time policy-based configuration

- Mobile security and network policy, moving policy with the virtual machine during VMware VMotion live migration for persistent network, security, and storage compliance
- Non-disruptive management model, aligning management and the operations environment for virtual machines and physical server connectivity in the data center

Cisco VN-Link network services are available in VMware vSphere environments with the Cisco Nexus 1000V Series Switches and through the Cisco Unified Computing System. Cisco VN-Link storage services are available as part of the intelligent fabric delivered by the Cisco MDS 9000 Family of multilayer SAN switches.

In deployments using the vNetwork Distributed Switch, VMware **vShield Zones** provide data center administrators with greater visibility and enforcement of network activity to provide the necessary level of isolation and means for compliance for each tenant. Previously, this level of policy enforcement required the use of external appliances, segmenting resources into smaller, disconnected clusters, disrupting the flexibility and efficiency of a shared computing pool or cloud. VMware vShield Zones allow the creation of logical zones that span all the physical resources of the virtual server layer, so that distinct levels of trust, privacy, and confidentiality can be maintained between tenants' virtual server resources.

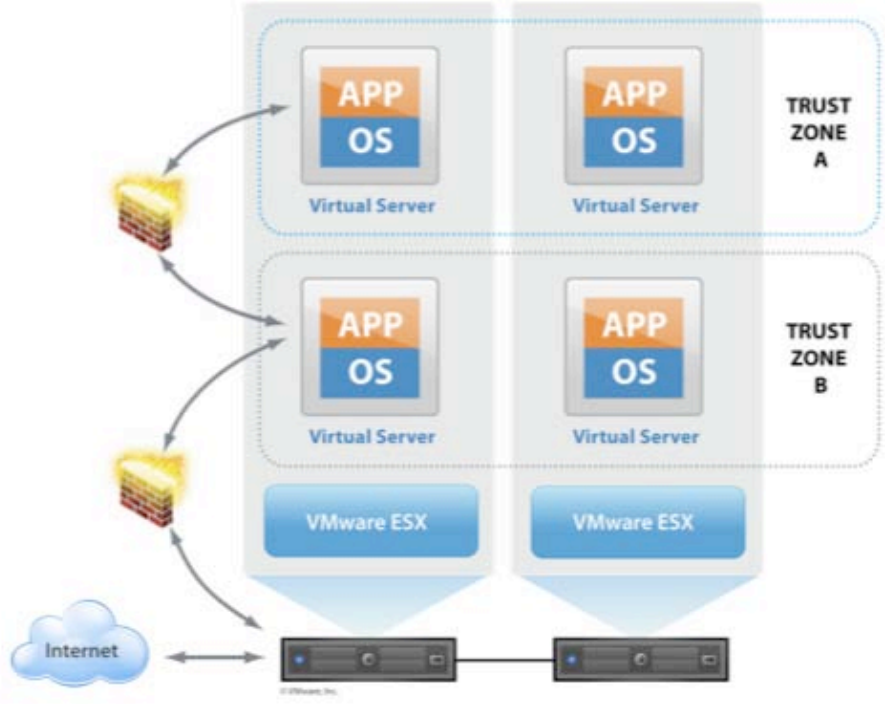

Figure 5) VMware VShield Zones.

VMware **VMsafe** is a new security technology for virtualized environments that helps protect the virtual infrastructure in ways previously not possible with physical machines. VMsafe provides a unique capability for virtualized environments through an application program interface (API)-sharing program that enables partners to develop security products for VMware environments. This allows direct integration of firewalls, intrusion detection, data leakage protection, and other security capabilities.

Confidentiality is a key aspect of any cloud service, and the combination of isolation and security features offered by NetApp, Cisco, and VMware allows administrators to create complete, secure, and isolate IaaS "containers."

INTEGRITY

Confidence that data is safe from unauthorized tampering, leakage, and corruption is a key component in any shared infrastructure. Typically this access to data has been safeguarded at the application level, but, with a shared infrastructure, all points of the environment must provide safeguards.

NetApp **DataFort** systems combine secure access controls, authentication, hardware-based encryption, and secure logging to help protect stored data. The appliance can be transparently deployed inline to gain security advantages without impacting user workflow. NetApp DataFort fits seamlessly into existing storage environments, supporting wire-speed encryption and compression for CIFS, NFS, IP SAN, FC Disk, FC Tape, and SCSI protocols. NetApp DataFort protects data both at rest and in flight with encryption certified by the National Institute for Standards and Technology. It allows control over sensitive data by providing access only to properly authorized individuals to help keep trade secrets and other intellectual property from falling into the wrong hands.

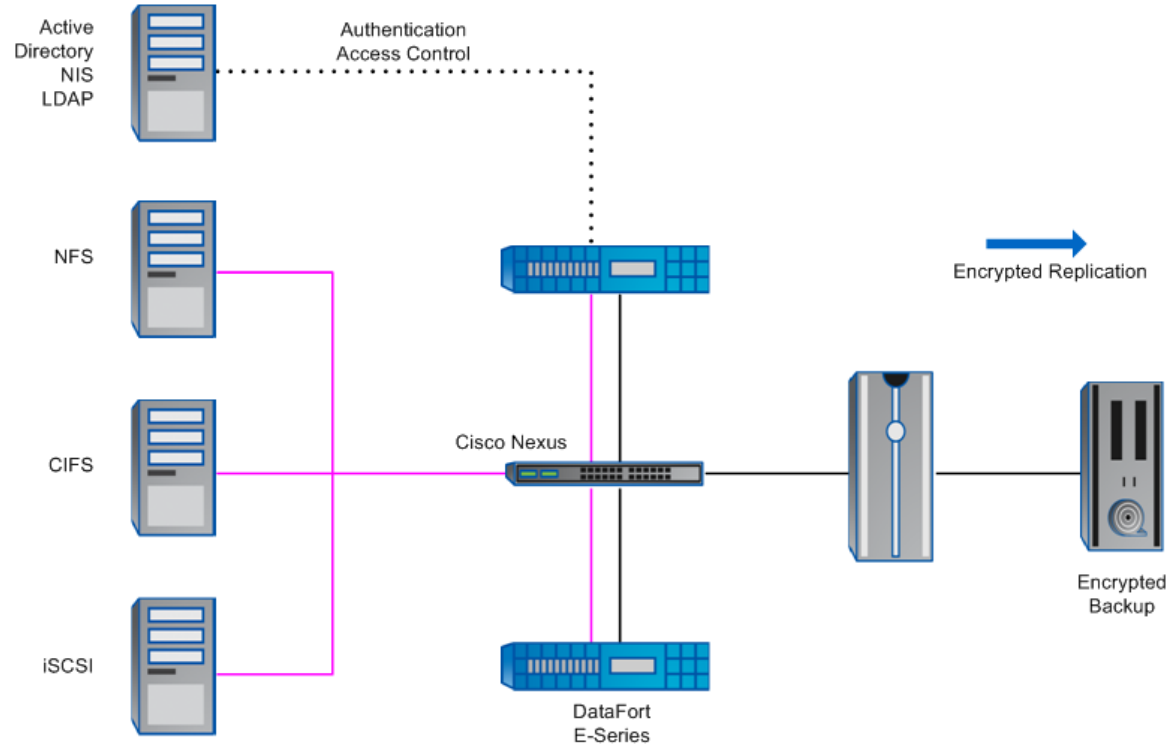

Figure 6) NetApp DataFort.

NetApp **SnapLock**® is a flexible data permanence solution that allows organizations to meet the strictest data retention regulations or internal IT governance policies. SnapLock aids in compliance with records retention regulations that require archival of e-mails, documents, audit information, and other data in an unalterable state for years. It is especially beneficial when faced with the dilemma of needing to retrieve unregulated yet crucial reference data that is not changed or deleted but must be accessed fast. SnapLock creates nonrewritable, nonerasable volumes to prevent files from being altered or deleted until a set retention date. NetApp allows an administrator to back up this Write Once Read Many (WORM) data to disk or tape for an additional level of data protection.

Transport security with Cisco **TrustSec** and hardware-based **SAN fabric encryption** is a new security architecture providing security within the cloud by creating a cloud of trusted fabric devices. The components of the Cisco cloud architecture authenticate themselves to one another, allowing encrypted transport between devices, packet classification, and access control among other services. Each of these capabilities supports the security requirements of cloud computing.

Combined, these capabilities can be leveraged across the secure cloud architecture to provide an extremely robust and tailored security solution safeguarding the assets of each tenant.

3 FLEXIBILITY

Flexibility within the secure cloud architecture enables administrators to maintain an adequate balance between current business needs and future requirements. Within a multi-tenant environment, this requirement becomes more important as both cloud and tenant administrators look to optimize existing deployments and efficiently provide IT-as-a-service offerings. To provide this flexibility, resources within both internal and external clouds must be capable of efficient provisioning, must be scalable, and must be mobile.

STORAGE PROVISIONING

Provisioning resources efficiently is critical to providing a flexible multi-tenant environment. These technologies allow data center administrators to quickly provision resources and help create more value with fewer physical resources. The ability to satisfy business needs faster with fewer resources to enable organizations to efficiently scale is key to a cloud offering.

NetApp **FlexVol®** technology enables the multi-tenant service provider to pool physical storage resources into large aggregates and then allocate discrete, virtual volumes of storage that can each be flexibly managed to suit the customer's requirements. Within each pooled aggregate of physical storage, providers can dynamically create, delete, resize, and reconfigure many NetApp FlexVol volumes with minimal disruption to storage services. Each FlexVol volume in turn can be presented to customer environments. NetApp FlexVol volumes enable multi-tenant service providers to deploy and manage logical storage containers in a highly efficient and adjustable manner, providing all of the benefits of virtualization to storage resource management and administration.

NetApp **MultiStore** software provides rapid provisioning capabilities through storage controller virtualization. By separating the enterprise storage needs from the physical hardware, administrators manage logical units called virtual storage containers, which are deployed based on workload requirements. These virtual storage containers, delivered on demand, have no impact on the current configuration and facilitate dynamic workloads and enterprise requirements.

Thin provisioning with NetApp storage enables administrators to allocate "just-in-time" storage resources yielding substantial storage efficiency benefits. A traditional approach to storage provisioning dictates that a request for 100GB of storage would be satisfied by an equal amount of raw space. Often this space is not fully used for several months if ever, thereby creating storage inefficiencies. With NetApp's thin provisioning, storage is allocated to allow for future growth while minimizing impact to currently available storage. This is transparent to the end user and applications. As the application demands increase, additional storage is allocated from the thinly provisioned data pool.

NetApp **Provisioning Manager** enables storage administrators to utilize policy-based automation to create repeatable, automated provisioning processes to improve the availability of data and enable provisioned storage to comply with individual tenant policies. These processes are faster than manually provisioning storage, easier to maintain than scripts, and help minimize the risk of data loss due to misconfigured storage.

SERVER PROVISIONING

Traditionally, server provisioning has been a process that is measured in days or weeks. This is not viable for any sort of cloud infrastructure in which the flexibility or on-demand nature of adding compute resources requires servers to be provisioned or reprovisioned in minutes. Cisco's UCS provides the ability to achieve this by making the physical hardware completely stateless, meaning it has no specific identify by which an operating system or application can "tie" itself to this particular blade. UCS uses a construct called a service profile to accomplish this. A service profile abstracts unique hardware identifiers such as server BIOS, adapter firmware settings, UUID, MAC address, HBA WWNs, and several others. A provisioning process is a simple as associating a service profile with an available anonymous blade. When that blade reboots it has the exact settings and attributes defined in the service profile. This entire process takes from three to five minutes to accomplish.

The **Cisco Unified Computing System** bridges the silos in the data center, enabling better utilization of IT infrastructure in a fully virtualized environment. It creates a unified architecture using industry-standard technologies that provide interoperability and investment protection. The architecture unites computing, network, storage access, and virtualization resources in a scalable, modular design that is managed as a single energy-efficient system through the Cisco UCS Manager.

The **Cisco UCS** is architected to meet the provisioning requirements of cloud computing. With a "wire once" deployment model, this robust compute platform readily supports the repurposing of its uniform infrastructure for virtual or native systems, allowing administrators to realign resources based on immediate business needs. The agility afforded by UCS means the cloud "transforms" in minutes, not days.

VMware Host Profiles creates a profile that encapsulates the host configuration and helps to manage the host configuration, especially in multi-tenant environments where an administrator manages many hosts in vCenter Server. Host profiles eliminate per-host, manual, or UI-based host configuration and maintain configuration consistency and correctness across the datacenter. These policies capture the blueprint of a known, validated reference host configuration and use this to configure networking, storage, security, and other settings on multiple hosts or clusters. The host or cluster can then be validated against a profile's configuration for any deviations.

VMware Lifecycle Manager provides a method to track and control virtual machines using standardized processes and best practices, and automates time-intensive manual tasks to reduce costs and optimize resources. This allows administrators to implement a consistent, automated process for managing the lifecycle of virtual machines in the data center, from provisioning to operation to retirement. It automates the virtualization workflow to improve efficiency and productivity and enables compliance with company policies, critical processes within a multi-tenant environment.

SCALABILITY

As demand for cloud resources grows, multi-tenant environments must have the ability to scale to meet tenant needs. Individual physical resources have capacity limitations, but with the addition of complementary software features, the result is a scalable infrastructure able to meet growing demand.

Another challenge as organizations scale is to enable workloads to be prioritized according to business needs. This can be accomplished through **quality of service (QoS)** mechanisms on each layer of the infrastructure. Quality of service is the ability to recognize types of traffic to control and prioritize them differently. Workload prioritization allows systems that are bandwidth constrained or under heavy load to prioritize the most critical tasks to be completed first, with subsequent, lower-priority jobs following. This allows cloud administrators to maintain different service levels for each of the application stacks within the fabric.

The NetApp storage system can leverage **MultiStore** software to easily scale with customer demand in a multi-tenant environment. As physical resources become more heavily utilized, additional physical controllers could be added to support further scale. The ability to add logical resources first, then virtual storage containers, enables cost-effective scalability both horizontally and vertically.

Through the use of NetApp's **data deduplication**, tenant's applications, data sets, and operating systems can be quickly provisioned, requiring less additional storage space. This is because within any given set of data there exist duplicate files and blocks. Data deduplication eliminates those common blocks on secondary and primary storage. This technology helps to enable data to be stored as efficiently as possible, effectively freeing up valuable storage space that is then available to other resources as the environment scales.

As tenants grow within a shared environment, additional resources will inevitably be required, even if they are less than in traditional approaches. **FlexClone® volumes** enable data center administrators to use cloning to efficiently provide these additional resources. FlexClone volumes can be split from the original active file system and used as live instances with their own Snapshot™ schedule. This feature is especially useful for test and development environments, in patch upgrade analysis, and in scaling of resources to additional data centers. In a netbooted environment, FlexClone volumes aid in scalability since instantaneous copies of master golden images can be created yielding deployment times of seconds for dozens or even hundreds of servers or virtual machines.

The NetApp **rapid cloning utility**, a user-friendly plug-in within VMware's vCenter, can greatly increase the efficiencies of virtualization deployments, allowing organizations to scale quickly. The rapid cloning utility automatically creates datastores, populates naming conventions, and allows the insertion of customized settings for each virtual server or desktop instance. The rapid cloning utility utilizes NetApp's file level cloning, which is similar to traditional FlexClone cloning. Users can clone VMs individually or in batches with no impact or dependencies placed on the original VM.

As service providers scale shared infrastructure to support additional customer and application environments, NetApp **FlexShare®** helps drive up their return on storage investments by maximizing the use of existing resources. FlexShare, an embedded feature of NetApp storage systems, allows storage administrators to increase resource utilization without impacting the performance of business-critical workloads. FlexShare enables service providers to consolidate multiple customer environments and applications and prioritize them based upon service-level requirements within a flexible NetApp shared storage infrastructure.

Often the data center is a mix of workloads in which equipment is designated for NAS traffic, SAN traffic, and emerging technologies such as Fibre Channel over Ethernet (FCoE). With NetApp **unified storage**, NAS, SAN, and FCoE protocols can be leveraged on the same storage system.

All servers must have a consistent and ubiquitous set of network and storage capabilities for cloud computing. One of the simplest and most efficient ways to deliver these capabilities is to deploy a **unified fabric**. The shift to a unified fabric gives all servers, physical and virtual, direct access to the SAN and LAN, allowing more storage to be consolidated in the customer's environment for greater efficiency and cost savings. This unified approach enables cloud administrators to provision resources dynamically, knowing that workloads can be shifted without physical reconfiguration or installation of additional hardware.

To consolidate server I/O, the server access layer must be adapted to support a unified fabric. Unified fabric is the consolidation of Ethernet and Fibre Channel networks across the same physical cable. Fibre Channel over Ethernet (FCoE) is the protocol that combines the two separate networks together. Cisco Nexus Family products enable this support, since they support FCoE, Small Computer System Interface over IP (iSCSI), and Fibre Channel. On the network side, supporting FCoE is simply a matter of enabling FCoE features in the Cisco Nexus 5000 Series switches and installing either the Fibre Channel or Fibre Channel and Data Center Bridging (DCB) uplink modules. The Cisco Nexus Series is designed to support Unified Fabric today and tomorrow.

The Cisco Nexus 5000 family delivers on the promise of a unified fabric by consolidating LAN, storage area network (SAN), and server cluster networks at the access layer. It enables I/O consolidation through Fibre Channel over Ethernet, providing benefits of lower power consumption, simplified cabling, reduced cost, and increased performance. The Cisco Nexus 2000 is a fabric extender that connects to the Nexus 5000 and provides an economical and manageable single platform that supports 1GE to 10GE connectivity. The Nexus 2000 assists in the transition from devices that only support Gigabit Ethernet while providing 10-Gigabit uplinks to the Nexus 5000, providing flexibility in the wiring plant.

The Nexus 7000 is built with modularity, virtualization, and resiliency at its core to deliver operational benefits in mission-critical environments at the core and aggregation layers. Other virtualized devices have been available from Cisco for some time, including the Firewall Services Module (FWSM), Application Control Engine (ACE) module, and the Cisco ASA 5580 platform. Virtual device contexts provide multiple virtual switching instances on a physical chassis with distinct sets of interfaces, control plane, and management. These features allow significant flexibility in network design in addition to more security as discussed previously.

Cisco **QoS** is a fundamental capability to provide prioritized processing to particular network communication in order to deliver a guaranteed level of bandwidth or performance. Service providers can leverage Cisco QoS to keep the network resources consumed by one customer from adversely affecting other customers sharing the same network infrastructure. The service provider could also grant a higher network service priority to those customers who pay a premium for enhanced performance or bandwidth beyond the baseline service level. From a resource management perspective, Cisco QoS allows service providers to granularly control the distribution of their shared network infrastructure capacity in order to maximize the efficient utilization of resources while complying with the service-level requirements of their customers.

VMware **vCenter Server** provides unified management of all hosts and virtual machines in a data center from a single console. This allows administrators to improve control, simplify day-to-day tasks, and reduce the complexity and cost of managing a scalable virtualized multi-tenant environment.

VMware **Distributed Resource Scheduler (DRS)** is a set of tools for managing resources and allocating them to virtual machines. DRS provides a hierarchical organization of VMs into resource pools, allowing resources to be allocated more easily and consistently to groups of VMs using shares, limits, and reservations. DRS can also manage the placement of VMs using VMotion under the following scenarios:

- Initial placement - When a VM is powered on, DRS can determine the best server in the DRS cluster on which to power on the VM based on current resource utilization.
- Dynamic load balancing - DRS constantly monitors resource utilization and can invoke VMotion to move VMs to better balance VMs across resources in the DRS cluster. The migration threshold of the DRS cluster can be set between conservative and aggressive.
- Maintenance mode - When an ESX server requires maintenance such as upgrades or replacing failed components, it can be placed into maintenance mode. In fully automated mode, DRS will automatically evacuate running VMs to other servers in the cluster using VMotion, balancing the load across the other servers.

DRS also provides affinity and anti-affinity rules, to require or prohibit VMs from running on the same ESX server. VMs that are part of a tiered application may benefit from affinity rules, keeping them on the same ESX server and allowing them to communicate at memory speeds instead of network speeds. Redundant VMs such as domain controllers should be separated using anti-affinity rules so that they do not all suffer an outage if a single ESX server fails. These rules can have network security implications since VMs with anti-affinity rules will not normally run on the same ESX server, meaning they must communicate over a physical network.

MOBILITY

Resource needs often change as project and customer requirements evolve. The ability to dynamically shift workloads and configurations between storage, compute, and network allows the infrastructure to grow and adapt to changing requirements. Routine data center maintenance such as software upgrades, deployment of new hardware, or troubleshooting often requires downtime. A properly designed cloud infrastructure utilizing the following technologies enables an administrator to securely migrate workloads across physical hardware, with no disruption to the services.

NetApp **MultiStore** software further aids flexibility because data provisioned into a virtual storage container can be migrated to other resources nondisruptively. Virtual storage containers can be easily migrated to different physical storage in the event of disaster, to aid in nondisruptive upgrades and also to solve scalability challenges. If a disaster occurs at the primary data center, a redundant copy housed on storage at a remote site can easily be accessed on the destination storage since no client configuration changes are necessary. Additionally, upgrades can be made on primary storage with minimal downtime to users because virtual storage containers can be migrated to a redundant controller and returned following upgrade completion.

The seamless live migration of virtual storage containers among physical resources is called **data motion**. This migration balances performance requirements in secure virtual partitions across the enterprise. This flexibility is valuable in a secure cloud architecture because all virtual machine data or application data within a virtual storage server is moved as an entire container. Data motion solves challenges with respect to mobility and availability and is critical to the live migration of virtual storage server instances. Data motion is primarily targeted at the storage infrastructure manager to manage the storage infrastructure pool regarding capacity and performance utilization and for asset lifecycle and maintenance. It allows migration of one or more datastores with hundreds of VMs.

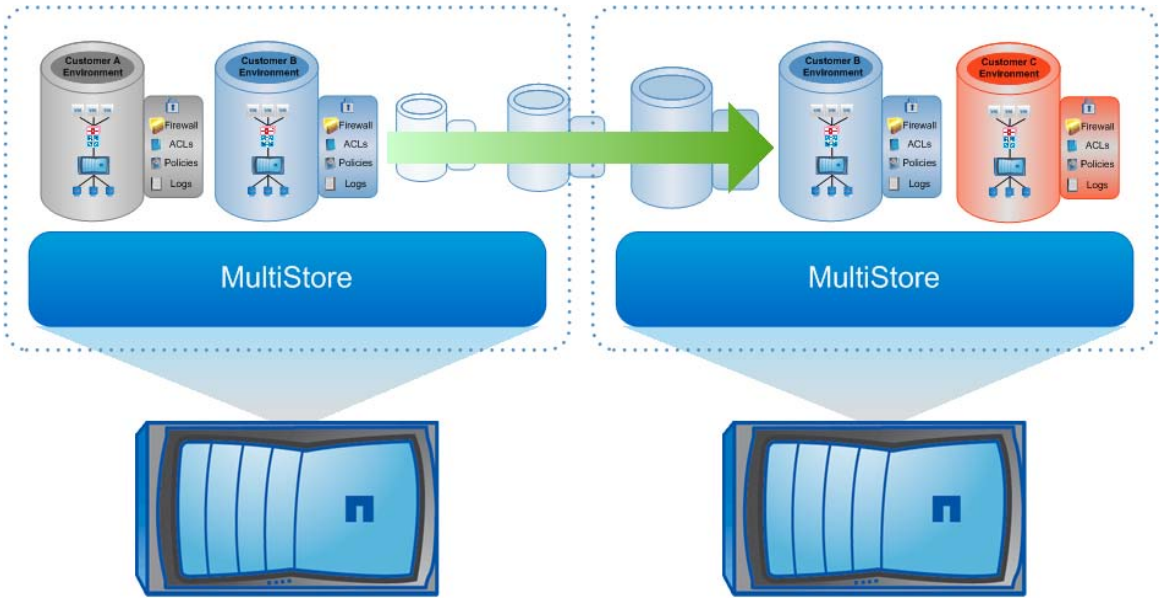

Figure 7) NetApp Data Motion.

Tenants within a cloud environment may leverage resources from both a shared cloud infrastructure and a traditional environment. NetApp **SnapMirror**® provides the ability to move applications, data sets, and operating systems between different NetApp controllers throughout multiple infrastructures. This feature allows tenants to expand and contract and leverage resources as needed between multiple environments. NetApp FlexVol volumes can also be moved among physical storage aggregates using SnapMirror, enabling service providers to dynamically and flexibly manage the underlying resources.

VMotion is a VMware feature that allows the live migration of running VMs from server to server during normal operations with no disruption of users or applications. VMotion is used by many VMware customers to balance VM and application load across the physical ESX servers and to evacuate VMs from a server in preparation for maintenance on the server while VMs are still running on other servers.

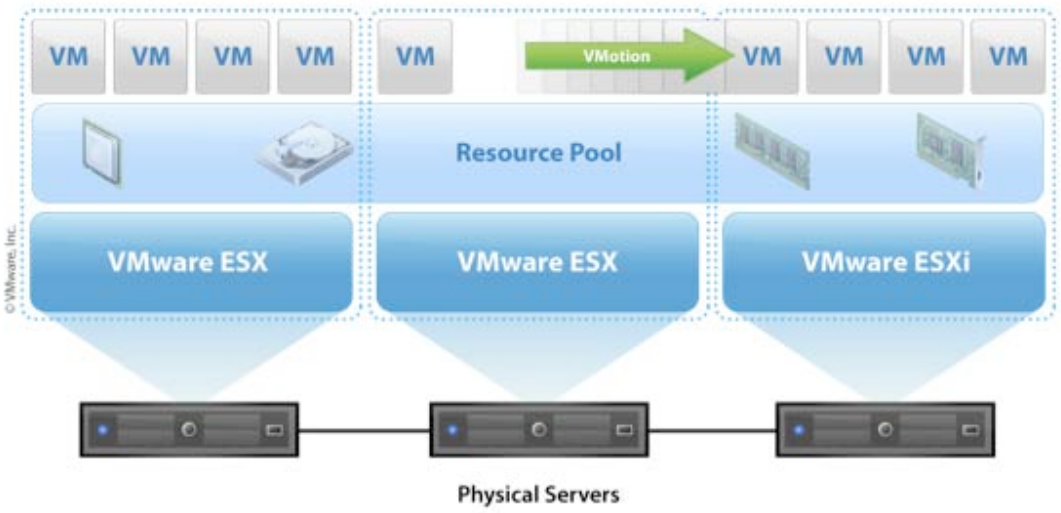

Figure 8) VMware VMotion.

As VMs are moved with VMotion, the virtual network connections associated with these VMs are tracked by vShield Zones and network security is maintained as VMs migrate. This enables stateless virtual machines to be migrated anywhere within the deployment.

Much like VMotion, VMware **Storage VMotion** allows administrators to migrate VM storage from datastore to datastore, even across storage devices and protocols as necessary, while the VM is running. Storage vMotion is fine-grain control for the VI admin to manage the performance of individual VMs, load balance capacity and load on individual datastores, and aid in populating the datastores.

The Cisco Nexus 1000v is a pure software-based implementation of the Nexus product line that provides enhanced virtual distributed switching functionality to VMware vSphere environments. The Cisco Nexus 1000v provides the following advantages:

- Policy-based virtual machine connectivity
- Mobile VM security and network policy
- A nondisruptive operation model

Each of these Nexus 1000v benefits enables and/or enhances the cloud environment, but the ability to manage the fluid atmosphere introduced with cloud computing may be its greatest advantage. The Nexus 1000v leverages port profiles to consistently apply network and security policies to the mobile VMs within the compute cloud.

The Cisco UCS platform details every provisioned server by a service profile. The service profile is a software definition of a server and its I/O adapters. Service profiles define the server hardware requirements, identity information, firmware version, and adapter configuration. Service profiles are maintained within the UCS fabric and managed by the Unified Compute System Manager (see Management section below) assigning identity to stateless UCS server platforms. The service profile construct allows server, network, and storage administrators to maintain configuration policies reliably and repeatedly across the cloud.

UCS and service profiles make it a trivial operation to move existing workloads among physical blades. One simply reassociates a given service profile from, say, blade Z to blade A and, when blade A reboots, it appears to the operating system and all higher-up applications that the server originally ran have rebooted.

However, now it has additional or different resources, such as more memory, additional I/O adapters, or a faster CPU clock speed.

The flexibility offered by the joint solution is unique in the industry. The features delivered by each vendor allow easy and efficient provisioning, dynamic scale-up or -out capabilities, and motion to and from cloud and noncloud offerings.

4 RESILIENCE

The ability of multi-tenant infrastructures to adapt to unforeseen events is imperative. The shared nature of such environments reinforces the importance of both data protection and high availability to provide business continuity for each tenant. The secure cloud architecture offers several technologies to enable an otherwise commodity infrastructure to cope with unplanned circumstances at any layer. The result is a resilient solution that enables uninterrupted operation in the face of the inevitable real-world problems encountered in the data center.

DATA PROTECTION

The foundation of a resilient shared infrastructure is data protection: The solution must provide a robust mechanism for protecting the integrity of data as it is moved and stored throughout the environment. This need for data protection is especially important when higher utilization and greater density are the main goals of cloud computing. Real-world equipment is prone to failure. The use of sophisticated software to prevent hardware failures from impacting data integrity is essential. NetApp's suite of data protection technologies can provide these robust measures so that information is protected.

Data protection falls into three main categories:

- Protection from physical failures
- Protection from user and application error or corruption
- Protection from facility loss

Hard disk failure is a common cause of data loss. RAID ("Redundant Array of Independent Disks") is a familiar technique developed in order to protect valuable data against disk failure and has been widely used in all types of enterprise environments. In many deployments, RAID 5 ("single parity") groups are used to provide protection against a single disk drive failure. Unfortunately, as disk capacity grows, so does the probability of a double disk drive failure, which would lead to total data loss and downtime as backups are restored. However, NetApp **RAID-DP® ("double parity")** technology can sustain two simultaneous disk drive failures. Statistically, this technique is over 10,000 times more reliable than a single-parity solution, and is more reliable than even RAID 1 mirroring. Leveraging NetApp RAID-DP within the secure cloud architecture enables the stack to be immune to downtime or data loss as a result of disk failure. Each tenant can utilize the infrastructure with the confidence that all application and underlying operating system data remains consistent and protected from physical storage failures.

While RAID-DP protects against disk failure, it does not prevent human error. If data is mistakenly deleted or corrupted, it is important that it be recovered. NetApp **Snapshot copies** offer a solution to the issue of accidental deletion or corruption: A consistent view of the file system can be captured instantaneously with virtually no storage overhead. These Snapshot copies can be created manually or automatically on a regular schedule set by the administrator. Because multiple Snapshot copies can be taken and saved per volume, tenants can schedule hourly, daily, weekly, and monthly Snapshot copies to create a schedule that best suits their individual data protection needs. Files can be recovered by accessing the snapshot directly without burdening the storage administrator or by invoking a NetApp **SnapRestore®** operation to roll back part or all of the active file system. NetApp Snapshot copies provide a highly efficient means to recover from data loss due to human error or software problems.

There is no substitute for a true backup to deal with hardware failure or a site-wide disaster. If an event destroys the local copy, having multiple full copies keeps valuable data safe at an alternate location. It is also often necessary to provide an off-site disaster recovery infrastructure to maintain continuity in the event of a site-wide failure. NetApp **SnapMirror** is a data replication feature that efficiently mirrors data,

applications, and their underlying operating systems between two NetApp storage controllers. The source and destination controllers can live in the same data center or be geographically separated, as both LAN and WAN links are supported. SnapMirror runs at continuous intervals to keep the destination controller up to date, so the RPO (recovery point objective: the amount of data not yet backed up) is dramatically lower than with legacy, periodic backup routines. A SnapMirror relationship may also be placed into synchronous mode, which effectively sets the RPO to zero. A NetApp destination appliance is more than a simple backup depot: It can also take over for the source appliance in a disaster recovery scenario. This means that the RTO (recovery time objective: the time from disaster to being back online) is minimal.

NetApp **SnapDrive®** technology extends control of NetApp Snapshot and SnapMirror capabilities to the individual tenants' host administrators. SnapDrive is a host-side storage management utility that allows server administrators of high-priority applications greater visibility and control of their underlying storage infrastructure. The utility allows the server administrator to create, schedule, and manage Snapshot copies of particular data sets as well as restore data from Snapshot copies using SnapRestore. All of these operations are performed from the host-side SnapManager® GUI, requiring no direct access to the storage itself.

Long-term backup storage is also a major concern in cloud environments. NetApp **Virtual Tape Library (VTL)** allows storage administrators to tightly integrate tape backup for off-site, long-term storage. By employing deduplication, compression, and a disk-to-disk-to-tape workflow, VTL technology can dramatically improves both the speed and space efficiency of tape backup. Further, the VTL's virtual tapes allow recovery of recent backups directly from disk, allowing the most common restore operations to proceed at 10 times the speed of traditional systems. In the context of the secure cloud architecture, this means that tape backups of tenants' data can be created and restored faster, with less tape, and at greater frequency than would otherwise be possible.

Data protection solutions are critical to the success of a high-density cloud environment, and NetApp offers the features and functionality to maintain and help protect the large volumes of data created.

HIGH AVAILABILITY

The second key aspect of a resilient shared infrastructure is high availability. High availability means that when resources, whether network bandwidth, CPU cycles, or data storage, are requested, those resources are online and always available to users. Any disruption in availability can ultimately result in a loss of productivity and revenue. The following secure cloud architecture features enable a highly available environment to meet the demands of every tenant.

The consolidated and shared multi-tenant environment introduces greater risk for storage availability. Business-critical applications and operations now depend on fewer physical storage resources. The NetApp **high-availability system configuration** (an "active-active" cluster) can be leveraged to enhance the availability of the environment, minimizing physical disruption and reducing operator errors. In an active-active controller configuration, each NetApp storage controller services its own workload while monitoring the other controller. If a controller fails, the other seamlessly takes over the additional workload. This failover is transparent to users and applications. A manual failover can also be triggered to perform scheduled maintenance and upgrades. This allows a cloud provider to stay up to date and compliant while continuing to provide superior storage availability.

The network platform in a multi-tenant infrastructure must be highly available to provide a reliable conduit for moving data between layers. The Nexus family has the ability to handle unanticipated hardware and network failures without affecting the end user. The Cisco Nexus family of products provides features that provide availability of data paths throughout the infrastructure. The Cisco NX-OS Software platform provides a technology called **virtual PortChannel (vPC)**. Although a pair of switches acting as a vPC peer endpoint looks like a single logical entity to PortChannel-attached devices, the two devices that act as the logical PortChannel endpoint are two separate devices. This environment combines the benefits of hardware redundancy with the benefits of PortChannel loop management.

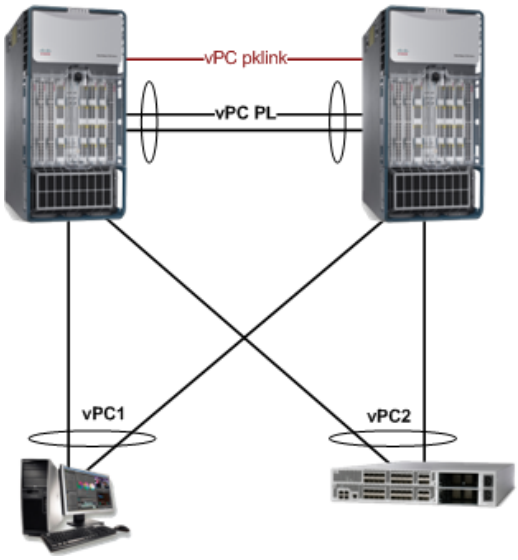

Figure 9) Cisco Virtual Port Channels (vPCs).

Redundant server connectivity to separate Nexus 5000 access switches provides redundancy not only for network connectivity but also for storage access.

- Two Nexus 5000s per rack provide access to SAN fabrics A and B, one fabric on each Nexus 5000 switch, maintaining the best of both Ethernet and SAN architectures.
- Scaling Layer 2 in a highly reliable fashion can be accomplished with virtual port channels that span separate chassis and add not only a flexible but a resilient network design.
- The Cisco Nexus 7000 supports core and aggregation layers in the network with redundant connectivity.
- The Cisco Nexus 7000 also provides lossless nondisruptive upgrades for zero-downtime service through no single point of failure in the system hardware and a modular operating system.

The Cisco Nexus family of products provides the cloud with a strong and flexible network fabric for critical services to run across. Unification of the Ethernet and storage traffic provides for one infrastructure instead of the usual two and reduces cost while improving scalability and flexibility. LAN and SAN can be managed as one unified fabric or split across separated teams on the same hardware with role-based access control as required. The result of combining LAN and SAN is a single, very manageable unified fabric with less equipment, less cabling, and lower power requirements and that requires less configuration complexity.

Cisco UCS provides added differentiation to the classic HA architectures in three different ways:

- UCS can enable customers to deploy N:M clustering topologies by reducing the risk of the dreaded second failure scenario. Once a server does fail, UCS can quickly reassociate the service profile with a spare blade, thus restoring the original level of redundancy in the cluster topology. This reduces risk and enables customers to move to more aggressive and cost-effective clustering deployments.
- For those applications that don't qualify for third-party clustering frameworks due to application tiering (gold and silver apps get clustering, bronze and below do not, for example), UCS can provide a manual method of rapidly reprovisioning a server.
- Scale-out applications benefit from the ability of UCS to provision a new server in minutes and quickly add it to an existing compute farm.

The VMware vSphere virtualization suite contains multiple features to provide high availability within the virtual server layer. Providing continuous operation of tenants' virtual machines is essential, and the VMware **High Availability** feature addresses this concern. With HA enabled, the virtual environment is monitored

continuously to detect OS and hardware failures. If a problem is detected, the affected virtual machines are automatically restarted on another physical host. This functionality is built directly into the hypervisor, sidestepping the complex configuration and dependencies of an application-centered solution. This feature means that downtime due to software error or hardware failure is bounded, enabling cloud providers to provide strong uptime service level agreements.

While VMware's HA feature limits downtime, VMware **Fault Tolerance (FT)** goes a step further by allowing hardware failures to have zero effect on the virtual machine. This is achieved by running two instances of the virtual machine on separate physical hosts: a primary VM and a shadow VM. The two systems are synchronized, so if the primary VM's host fails, the shadow VM seamlessly and instantly takes over. While this does require twice the number of computing resources, the advantage offered is staggering for the most important VMs. A virtual machine protected by FT will not suffer downtime as a result of hardware failure. By eliminating a major source of downtime, cloud providers can provide tenants with an even stronger uptime service level agreement.

The combination of VMware **Site Recovery Manager (SRM)** and NetApp storage creates the ability to simplify and automate processes associated with typical disaster recovery operations. SRM detects failures at the primary site and alerts administrators of the problem. With the administrator's authority, SRM will facilitate the migration of operations for a predefined grouping of virtual machines from the primary site to the secondary site. The predefined groups can be based on tenant priority or service-level tiers that can correspond to specific service actions in a disaster recovery scenario. The higher priority or tier groupings receive quicker actions and failover before the lower priority or tier groupings. Virtual machine identities and network configurations are preserved during the failover process using SRM to provide the same secure isolation at the secondary site for each tenant's environment. Because SRM is augmented with the NetApp Site Recovery Adapter, it can leverage NetApp SnapMirror and Snapshot technologies to provide low-RTO/RPO recovery in the virtual infrastructure.

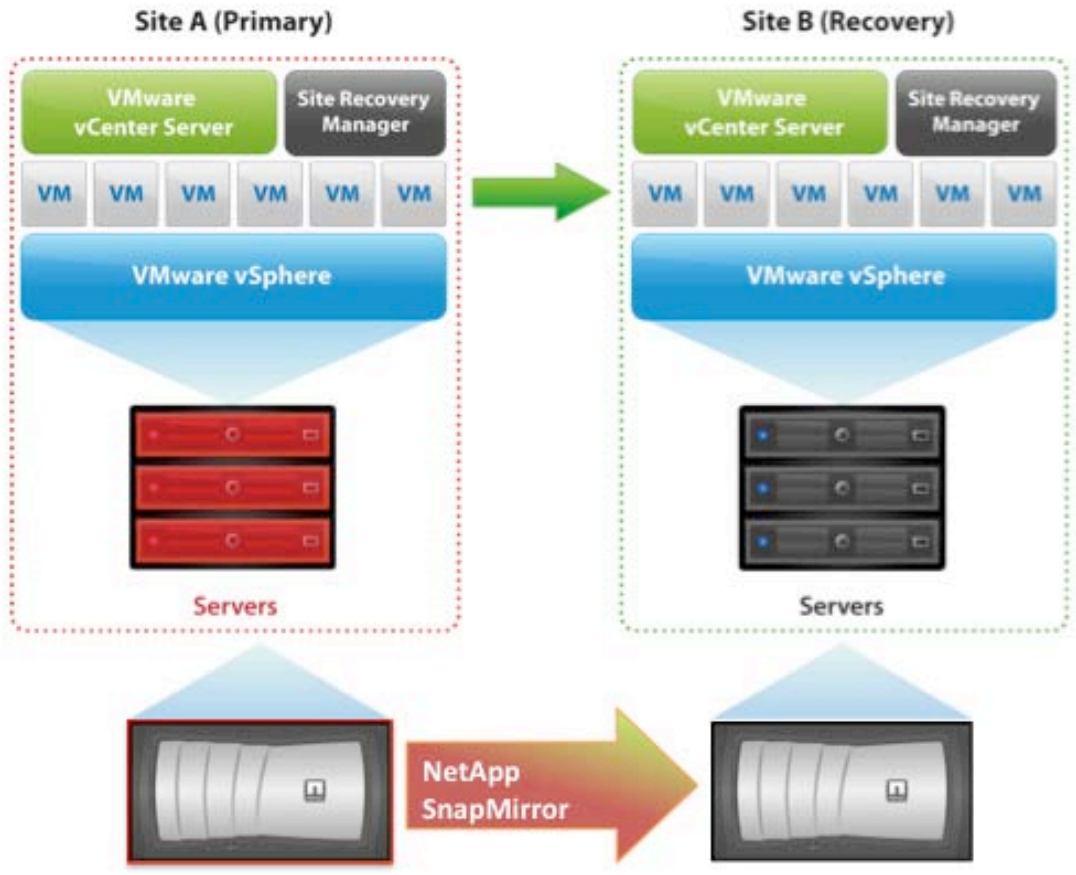

Figure 10) VMware SRM

The full environment resilience offered by this joint architecture offers both cloud providers and tenants the peace of mind that their critical applications and data will be available for business whenever and wherever needed.

5 MANAGEMENT

The ability to effectively manage and monitor resources within the multi-tenant cloud environment is key to a successful implementation. Cloud management tools must address both the physical nature of the cloud in addition to its storage, compute, and network layers of abstraction. Management platforms simultaneously achieving these objectives will allow IT organizations to realize a new level of operational efficiency, flexibility, and agility promised by cloud computing.

For the service provider there is an increasing need to provide faster response times, more flexibility, and continuous business operation to effectively meet more dynamic and varied customer requirements. Cloud architecture offers the benefit of resource and infrastructure consolidation, but introduces unique challenges with deploying and coordinating multiple tiers of resources as cohesive environments while maintaining customer access and service levels. Effectively managing shared storage, network, and compute resources

in a cloud architecture requires the ability to consolidate not only resources, but also provisioning and operational tasks according to organizational units. Multi-tenant service providers require comprehensive control and extensive visibility of their shared infrastructure to provide security and service levels for their customers. The secure cloud architecture delivers end-to-end management, monitoring, and accounting solutions that enable higher efficiency, utilization, and availability of cloud service delivery.

PARTITIONING

NetApp **MultiStore** provides secure partitioning of network and storage resources into virtual domains that can be consolidated and managed on a common storage system. NetApp MultiStore enables the multi-tenant service provider to provision, move, and protect customer data based on user and application boundaries. Since customer resources are maintained in virtual storage containers, providers can establish operational and business continuity policies that are customer or application specific.

The Cisco Nexus 7000 network switch **virtual device contexts**, in addition to security and flexibility, also allow efficient management in a multi-tenant design**.** Multi-tenant service providers can configure and deploy multiple VDCs on each physical Nexus switch, each running as a discrete entity with its own configuration, network administrator, and set of running processes. Cisco Nexus VDC technology delivers to multi-tenant service providers the powerful ability to extend logical partitioning of customer environments into the network device layer, realizing the management and efficiency benefits of resource consolidation while leveraging a common physical network infrastructure.

The Cisco Nexus line of products meets these multi-tenant cloud management requirements via command line and GUI user interfaces, in addition to industry standard application programmatic interfaces (APIs) for integration with third-party management suites. The Cisco cloud management offering allows IT organizations to readily manage a unified fabric or distinct LAN and SAN deployments that have been virtualized to support the multiple application stacks of the cloud. The fundamental pieces to this end are Cisco's Data Center Network Manager (DCNM), Fabric Manager (FM), and Device Manager (DM) and their associated Web service exposed APIs, in addition to the Nexus 1000v management features.

The Cisco UCS provides not only a unified fabric, but also a unified point of management. The Cisco UCSM that is embedded in the network fabric has an intuitive GUI, a command-line interface (CLI), and a robust API for managing all system configurations and operations. The UCSM management software is deployed in a clustered active-passive configuration so that the management plane remains intact even if failure in the fabric occurs. As part of the fabric, the UCS Manager automatically discovers resources as they are installed in the system, adds them to inventory, and can automatically provision every aspect of the UCS servers and their I/O connectivity. Leveraging service profiles, the UCSM helps automate provisioning along with allowing data center managers to deploy applications in minutes instead of days. The Cisco UCS Manager helps increase IT staff productivity by expediting the use of storage, network, and compute resources for tenants within the cloud by providing a consistent fabric and a uniform platform.

VMware **vCenter** simplifies the management of multi-tenant environments through the use of resource pools and clusters. Service providers can create resource pools to hierarchically partition the shared CPU and memory resources of an individual host or cluster, then delegate control over the pools to their customers. VMware vCenter allows service providers to add, remove, and reorganize resource pools as needed in response to the changing workloads or management requirements of their customers. Defining resource pools across VMware vSphere clusters creates a separation of resources from the underlying hardware, which allows the provider to manage resource allocation from a pooled, aggregate perspective, regardless of the actual host machines that contribute the resources. This abstract management of resources enables service providers to modify or upgrade the underlying hardware capacity without affecting the resource allocations of individual or groups of VMs assigned to their customers. Service providers can link together multiple VMware vCenter deployments to leverage a single, consolidated management interface for very large-scale virtual environments.

VMware vShield allows service providers to deploy distributed **vShield Zones** virtual appliances on each vSphere host, which provides the provider with visibility and enforcement of virtual network traffic across their customer virtual server environments. The distributed vShield Zones appliances are administered by VMware **vShield Manager**, which integrates seamlessly with the provider's VMware vCenter deployment to present policies and events in the context of the existing virtual machines, networks, host, and clusters used to service their customer deployments. Service providers can use the VMware vShield Manager unified

dashboard overview to manage and deploy policies for the entire VMware vCenter environment, leveraging their existing virtual infrastructure containers as organizational zones across physical hosts, virtual switches, and networks. VMware vShield enables multi-tenant service providers to leverage common interfaces and organizational units to deploy and manage their customers' vSphere virtual network environments cohesively with their virtual server resources.

ACCOUNTING AND MONITORING

NetApp provides cohesive accounting and monitoring solutions that enable cloud providers to achieve improved efficiency, utilization, and availability with insightful, end-to-end control of the shared storage infrastructure, customer resources, and service-level delivery.

NetApp **FilerView®** is the primary, element-level graphical management interface available on every NetApp storage system. NetApp FilerView is an intuitive, browser-based tool that can be used for monitoring and managing individual NetApp storage systems. NetApp FilerView provides control over administrative and user access. Storage providers can use NetApp FilerView to inspect the health and status of NetApp storage systems, as well as configure notification and alerting services for resource monitoring.

The varied and dynamic requirements of accommodating multiple customers on a shared infrastructure drive cloud providers toward storage management solutions that are more responsive and comprehensive while minimizing operational complexity. NetApp **Operations Manager** delivers centralized management, monitoring, and reporting tools to enable service providers to consolidate and streamline management of their NetApp storage infrastructure. Service providers can reduce costs by leveraging comprehensive dashboard views into optimizing storage utilization and minimizing the IT resources needed to manage their shared storage infrastructure. Operations Manager delivers to service providers comprehensive visibility into their storage infrastructure, providing continuous monitoring of resources and analysis of utilization and capacity management and insight into the growth trends and resource impact of their customers. NetApp Operations Manager also addresses the business requirements of multi-tenant service providers, enabling charge-back accounting through customized reporting and workflow process interfaces.

NetApp **SANscreen®** enables service providers to further improve the quality and efficiency of their storage management with real-time, multivendor, and multiprotocol service-level views of their storage environment. NetApp SANscreen is a suite of integrated products that delivers global visibility into the service provider's networked storage infrastructure. NetApp SANscreen delivers end-to-end visibility, flexible and proactive management, and service-level assurance for multi-tenant service providers.

The Cisco Nexus family of products is managed by **Data Center Network Manager.** Focused on the management requirements of the data center network, DCNM delivers service providers a robust framework and rich feature set that fulfill the routing, switching, and storage networking needs of present and future data centers. In particular, DCNM automates the provisioning process, proactively monitors the storage area network (SAN) and local area network (LAN) by detecting performance degradation, secures the network, and streamlines the diagnosis of dysfunctional network elements. It is purpose built for interoperability with the NX-OS operating system of the Nexus product family, allowing the consolidation of administration instrumentation across the LAN, SAN. and/or unified fabric.

Cisco's **Fabric Manager** and **Device Manager** tools allow IT organizations to support Cisco MDS and Nexus 5000 platforms leveraged in the storage fabric of the cloud. Historical performance monitoring, centralized management services, and integrated traffic analysis enable configuration and frame-level visibility within a traditional or unified SAN fabric. Cloud topology views and granular configuration capabilities allow these two management platforms to effectively and efficiently support multi-tenant cloud requirements.

The Cisco Nexus switches provide several features that directly address the challenges of service providers to support, and monitor the network services delivered to customers in a multi-tenant environment. Some of these Cisco Nexus features include Role-Based Access Control (RBAC) and Nexus Flexible NetFlow.

Role-Based Access Control allows service providers to define hierarchical control policies to delegate administrative access according to customer environments and reflect organizational roles and responsibilities. Service providers create roles that determine a privileged scope of access or use of network devices or services, then associate administrators or customers with those roles according to organizational or operational function.

As service providers host more customer environments on shared infrastructure, understanding the network's detailed behavior is critical for network availability, performance, and troubleshooting. Cisco Nexus **Flexible NetFlow** enables multi-tenant service providers to optimize utilization of their network infrastructure, improve capacity planning, reduce operational costs, and improve security incident detection with increased flexibility and scalability. Cisco Flexible NetFlow delivers network service metering and accounting information to facilitate usage-based billing processes for multi-tenant service providers.

VMware vSphere offers multi-tenant service providers centralized tools to monitor and manage access to their VMware virtual infrastructure. Administrative and operational permissions are assigned to user-defined roles, enabling service providers to customize access and privileges across the virtual infrastructure. VMware **vCenter** provides an audit trail of configuration changes and reports for event tracking, allowing service providers to maintain a record of changes to their customer environments. Providers can leverage VMware vCenter to continuously monitor the physical hosts and VM availability within a vSphere deployment. Using VMware vCenter, service providers can obtain real-time performance monitoring of VMs, resource pools, and physical host utilization at customizable granularity and collect this data for historical reporting, trending, and analysis of their customer environments.

To consolidate and simplify VMware virtual network security management, VMware **vShield Manager** provides centralized management of monitoring and access policies across an entire VMware vCenter deployment. Service providers can customize access to VMware vShield privileges according to existing VMware vCenter roles or create distinct access roles for vShield administration according to customer or organizational requirements. VMware vShield Manager consolidates logging of network activity between VMware **vShield Zones** and to the outside networks, delivering service providers a single interface to obtain real-time and historical reporting of event data. It also provides log archival facilities for compliance with organizational and regulatory policies.

VMware **vCenter CapacityIQ** helps provider administrators monitor and manage the capacity of virtualized multi-tenant shared services infrastructure . CapacityIQ is installed as a VirtualCenter plug-in. CapacityIQ ensures that virtualized infrastructure capacity is efficient and predictable. The what-if scenario feature of CapacityIQ allows administrators who are managing capacity to model changes to the virtualized environment and assess the effect of a change without implementation costs or time consuming trial and error.

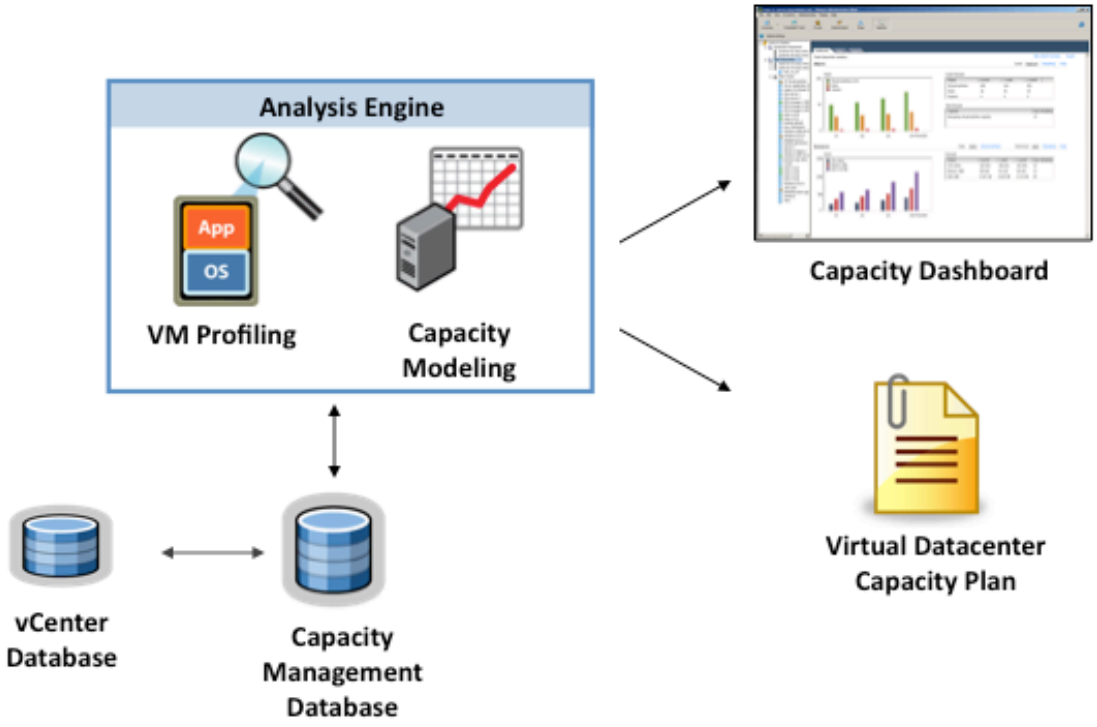

Figure 11) VMware Capacity IQ

The secure cloud architecture delivers a unique, end-to-end infrastructure solution combining complementary technologies from NetApp, Cisco, and VMware that enable service providers to efficiently manage consolidated customer and application environments on a shared-resource cloud architecture. Providers can cohesively manage resources, access, and control according to each customer environment or application, delivering reliable standardized services while leveraging shared resources and infrastructure. The secure cloud architecture offers providers and their customers extensive control and visibility into the storage, network, and compute resources of each deployment to enable appropriate access and service levels while maximizing efficiency and utilization.

6 CONCLUSION

Architecting a successful cloud computing infrastructure can be a daunting task. No single vendor has the full product or feature set across the entire environment to do it alone. The secure cloud architecture, the result of years of planning, offers the best products and features to create a cloud service that can realize the efficiencies of shared infrastructure while still offering the security, flexibility, resilience, and ease of management of a traditional IT deployment.

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 must be used solely in connection with the NetApp products discussed in this document.