



End-to-End Quality of Service:

Cisco, VMware, and NetApp Team to Enhance Multi-Tenant Environments



By Chris Naddeo, Cisco

Building shared infrastructure has always been something of a challenge. If you look at a typical corporate data center design, you find either that important applications have their own dedicated infrastructure or that shared elements have been overengineered to far exceed requirements. Either approach underutilizes resources and wastes your IT budget.

The problem is that no one really knows how infrastructure components such as servers, networks, and storage will behave as additional load is added. Will a resource become a bottleneck, decreasing the performance of an important application unexpectedly? If so, how can you quickly identify the source of such bottlenecks?

The current interest in cloud computing has made understanding all aspects of multi-tenant environments—infrastructures in which all resources are shared—even more critical. In fact, many companies hesitate to build cloud infrastructure or contract for cloud services because of fears about security and quality of service (QoS).

➤ [Continued on page 3](#)

Storage Infrastructure for the Cloud



With total IT spending on cloud computing forecast to grow at least three times by 2012, you've probably heard a lot about the potential benefits of cloud computing. Perhaps your company has already begun purchasing some cloud services rather than adding to its existing IT infrastructure every time a new requirement arises.

➤ [Continued on page 4](#)

Boosting Data Center Efficiency:

Innovate and Thrive in a Difficult Economy



Since 2005, the power required for every \$1,000 of IT equipment (servers, storage, and networks) has grown by four to six times, dramatically raising the "hidden" costs of operating a data center and setting the stage for significant problems. As a result, corporate data centers will have to become much more efficient in terms of both power density and asset utilization. If you deploy the latest technology without boosting utilization of those assets, you'll potentially waste a huge amount of power (at significant cost to your company). More important, if you fail to raise utilization levels, your data center will likely become constrained to the point where you can't deliver the power and cooling necessary to support the latest technology. In fact, 42% to 43% of all data centers in Europe and the United States are already power constrained.

➤ [Continued on page 6](#)

INSIDE

The NetApp Community

➤ [Page 2](#)

Deploying FCoE in Your Data Center

➤ [Page 8](#)

NetApp and Cisco: Virtualizing the Data Center

➤ [Page 10](#)

A Quantum Leap in Virtual Management

➤ [Page 11](#)

Virtualizing Microsoft Applications

➤ [Page 12](#)

Letter from the Editor

Welcome to the 2010 issue of the Tech OnTap magazine. Cloud (or IT as a service) is one of the hottest topics in IT right now, and here at Tech OnTap, we are all about it. Economic and business pressures are shifting storage decisions toward more cost-effective and responsive solutions such as IT as a service. NetApp has been laying the foundation for efficient and flexible IT every step of the way.

The storage decisions you make today will have a lasting impact on your company's future business performance. In this issue of Tech OnTap, we explore NetApp technologies that support the transition from silos to virtualization to shared storage infrastructures and even outsourced cloud services. Companies making this transition are gaining a huge competitive edge—reducing costs and increasing flexibility.

In our feature article, Chris Naddeo of Cisco explores the four pillars of success in cloud computing. Find out how our alliance with Cisco and VMware results in enhanced security and quality of service in multi-tenant environments and enables you to complete your journey to a fully virtualized, dynamic data center.

Efficient and flexible IT makes it possible to innovate and thrive in a difficult economy while laying a foundation for the future. Discover how you can buy less, use less, and waste less storage. And see how storage efficiency can translate into measureable business savings and value—not just today but also tomorrow. We also share tips for enhancing your data centers and IT practices to improve efficiency.

We have carefully selected the articles and information in Tech OnTap to give you useful, actionable information. Stop by our community online and let us know how we're doing: netapp.com/techontap/community. Are you getting the information you need when you need it? Are there any topics about which you'd like to hear more? I look forward to hearing from you. ■



Mina

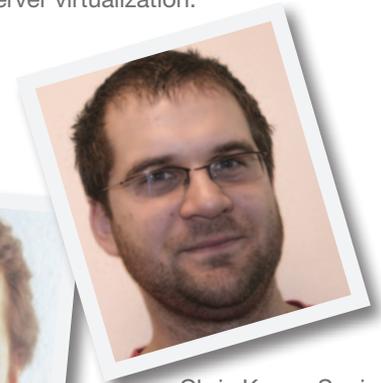
Mina Eng
Tech OnTap Editor
mina.eng@netapp.com

The NetApp Community

The NetApp Community recognizes members who show an active commitment and enthusiasm for assisting others, thereby helping to build a strong, thriving community. Meet our latest featured members, who have shown a willingness to help others and share insights in the community. ■



Andrew Miller, Systems Engineer. Loves the business and technical sides of storage and server virtualization.



Chris Kranz, Senior Technical Consultant. Specializes in storage, virtualization, and business processes.



Eric Barlier, Senior Storage Consultant. A proponent of using best practices and sharing knowledge with his peers.



Brendon Higgins, Technical Services Team Leader. Digital plumber on SAN and infrastructure.



Radek Kubka, Senior Technical Consultant. Focuses on IP SAN and data center virtualization.

Join Our Online Community



The NetApp Community consists of over 55,000 members to date from 150 countries. This powerful community of customers, partners, and employees helps connect professionals around the world to network and share insights with peers and learn how others deploy and manage NetApp® technology.

Community: netapp.com/community

- Tech OnTap: netapp.com/techontap/community
- User Groups: netapp.com/usergroup
- Community Events: netapp.com/communityevents
- Partner Community: netapp.com/partnercommunity

Sign up for the NetApp Community at netapp.com/community. ■

End-to-End Quality of Service:

Cisco, VMware, and NetApp Team to Enhance Multi-Tenant Environments

By Chris Naddeo, Cisco



Continued from page 1

Cisco has teamed with VMware and NetApp to design and test a secure, multi-tenant cloud architecture that can deliver on what we see as four pillars of secure multi-tenancy:

- **Secure separation.** One tenant must not be able to access another tenant's virtual machine (VM), network, or storage resources under any circumstance. Each tenant must be securely separated.
- **Service assurance.** Compute, network, and storage performance must be isolated and guaranteed during normal operations as well as when failures have occurred or certain tenants are generating abnormal loads.
- **Availability.** The infrastructure must make sure that required compute, network, and storage resources remain available in the face of possible failures.
- **Management.** The ability to rapidly provision, manage, and monitor all resources is essential.

In this article I describe the unique architecture the three companies have designed to address these pillars of multi-tenancy. I go on to discuss our efforts around the second pillar—service assurance—in more detail.

A recently released design guide (www.netapp.com/us/media/cisco-validated-design.html) provides full details of a Cisco® validated design that uses technology from all three companies to address all four pillars described above. An online Tech OnTap article (www.netapp.com/us/communities/tech-ontap/tot-secure-mobile-cloud-storage-1001.html) describes one element of the architecture, NetApp MultiStore®, in more detail.

ARCHITECTURE OVERVIEW

A block-level overview of the architecture is shown in Figure 1. At all layers, key software and hardware components are designed to provide security, quality of service, availability, and ease of management.

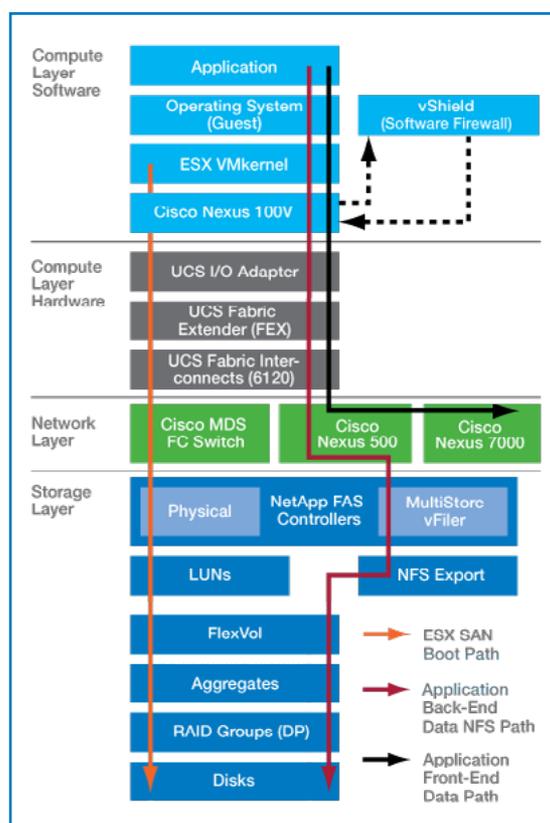


Figure 1) End-to-end block diagram.

Compute layer

At the compute layer, **VMware® vSphere™ and vCenter™** Server software provide a robust server virtualization environment that allows server resources to be dynamically allocated to multiple guest operating systems running within virtual machines.

VMware vShield Zones provides security within the compute layer. This is a centrally managed, stateful, distributed virtual firewall bundled with vSphere 4.0 that takes advantage of ESX host proximity and virtual network visibility to create security zones. vShield Zones integrates into VMware vCenter and leverages virtual inventory information, such as vNICs, port groups, clusters, and VLANs, to simplify firewall rule management and trust zone provisioning. This new way of creating security policies follows VMs with VMotion® and is completely transparent to IP address changes and network renumbering.

The Cisco Unified Computing System™

(UCS) is a next-generation data center platform that unites compute, server network access, storage access, and virtualization into a cohesive system. UCS integrates a low-latency, lossless 10-Gigabit Ethernet network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

Network layer

The network layer provides secure network connectivity between the compute layer and the storage layer as well as connections to external networks and clients. Key components include:

- Cisco Nexus™ 7000, which provides Ethernet (LAN) connectivity to external networks
- Cisco Nexus 5000, which interfaces with both FC storage and the Cisco 7000
- Cisco Nexus 1000V, a software switch that runs within the VMware kernel to deliver Cisco VN-Link services for tight integration between the server and network environment, allowing policies to move with a virtual machine during live migration

Continued on page 16



CHRIS NADDEO
TECHNICAL
MARKETING
ENGINEER FOR UCS
CISCO

Chris joined Cisco to focus on customer enablement and the design of optimal storage architectures for the Cisco Unified Computing System. He has an extensive storage background, including one year at NetApp as a consulting systems engineer for Oracle® and Data ONTAP® GX as well as nine years at Veritas, where he served as a product manager for Veritas™ storage software. ■

Storage Infrastructure for the Cloud

By Jeff O'Neal



► Continued from page 1

Right now there is economic pressure for companies to find new models of working, and technological maturity is making workable solutions both possible and affordable. You might already be considering ways to make your own data center more “cloud-like” to boost efficiency, cut capital costs, and provide the elastic scaling you need to adapt rapidly to changing business requirements. How to actually go about doing this—especially where storage is concerned—might still be a bit hazy.

In this article, I begin by defining what cloud computing is and then discuss the broad requirements for storage that meet the needs of cloud infrastructure—whether that infrastructure is a “private” cloud running within your data center to meet your company’s needs or a “public” cloud providing for-fee services to a broad community.

WHAT IS CLOUD COMPUTING?

In the broadest sense, when we talk about the “cloud,” we are referring to the ability to deliver IT as a service (ITaaS). Some people define the cloud from a technology standpoint, where exact architectural models and development protocols are specified, but at the outset it is simpler to think of cloud computing as a business model for delivering IT as a service. Cloud services are the end deliverable of cloud computing and can be broken down into four categories, as illustrated in Figure 1 on page 5. A cloud can be either private: limited to the ecosystem of end users, partners, and/or customers directly associated with a company; or public: available to more or less anyone with Internet access. (See Figure 2 on page 5.)

STORAGE INFRASTRUCTURE FOR CLOUD COMPUTING

Whether a cloud is public or private, the key to success is to create an appropriate infrastructure to deliver each cloud service, whatever it might be, as efficiently as possible. In a private cloud you might need to support a broad variety of applications,

so your goal is to create an infrastructure that can flexibly allocate resources to each application as required.

It’s probably already clear that, on the compute side, server virtualization provides an appropriate infrastructure for cloud services because it allows compute resources to be efficiently partitioned and quickly allocated, increased, decreased, or deallocated as needs change. A rapidly maturing set of management services also provides speed and flexibility as well as increased availability.

Much less has been written on how to create efficient and effective storage infrastructure for cloud computing. In fact, the first Storage Networking Industry Association (SNIA) work group focusing on cloud storage was announced at Storage Networking World in April 2009 (www.snia.org/about/news/newsroom/pr/view?item_key=b08b954cba2a94281c7dc86b19fdb6e46912dfb8), with the objective “to identify, develop and coordinate system standards and interfaces for cloud storage.”

Since there are no approved or de facto standards in place for cloud storage, there are a few questions you should consider when evaluating new or existing storage solutions for suitability:

- **Can your storage scale elastically?** Similar to what you do with virtualized servers, you need to be able to allocate, increase, decrease, and deallocate storage rapidly and with a minimum of overhead.
- **Can you automate storage management processes?** The more you can automate regular practices such as provisioning, backup, and replication, the better your environment will scale.
- **Can you meter and report on usage?** In order to implement cloud services, you must understand resources a user of your service needs; have the ability to report back on actual usage; and, now or in the future, be able to bill based on resource usage.
- **Can you move data freely?** If your data is tied to inflexible storage, efficiency and availability will suffer.

- **Can you establish multi-tenancy while guaranteeing that resources are sufficiently secure?** Allowing multiple business units or separate entities to share the same storage hardware is a necessity for efficient cloud storage.
- **Can you boost storage efficiency?** The first step is to increase utilization; beyond that, reducing overhead, thin provisioning, and eliminating redundancy are all opportunities to increase efficiency.
- **Can you efficiently protect your data?** A key to cloud success is integrating all your processes such that they are simple, repeatable, and efficient. Consistent data protection and disaster recovery processes that cover every service you provide, at the appropriate policy level, are essential.
- **Can you do everything on a single network fabric?** The arrival of Fibre Channel over Ethernet (FCoE) makes it possible to consolidate your SANs and LANs on a single Ethernet fabric for decreased cost and increased flexibility.
- **Does your storage environment support server virtualization?** Assuming that server virtualization will be a key component of your cloud infrastructure, you will want storage that integrates closely with any virtualization solutions you now use or are likely to adopt in the future.

WHAT SHOULD YOU DO NOW?

If you are working to evolve your existing data center toward a cloud model, you can get started by:

- Rethinking your data center design to accommodate the density and power requirements of the latest IT hardware; this topic is discussed in more detail in an onlineTech OnTap article (www.netapp.com/us/communities/tech-ontap/tot-datacenter-efficiency-0509.html).
- Rearchitecting racks, cabling plants, and network infrastructures so that applications can be moved around dynamically to better accommodate virtualization; in most existing data centers, key business applications are boxed in for security reasons.

- Virtualizing everything: servers (and possibly desktops), networks, and storage.

As you move toward a virtualized storage environment that offers the same advantages as server virtualization, make sure you evaluate storage with the guidelines from the previous section in mind.

NETAPP STORAGE FOR THE CLOUD

NetApp® data management solutions have already been deployed in a wide variety of public and private cloud environments, such as Telstra (<http://media.netapp.com/documents/telstra-0308.pdf>) and Sensis (<http://media.netapp.com/documents/sensis-storage-on-demand.pdf>).

Based on this experience, we continue to refine our product offerings to meet emerging cloud needs and offer a full feature set to meet the guidelines set out above.

Elastic scaling

NetApp flexible volume technology (FlexVol®) abstracts storage volumes from underlying disks. Any storage container (LUN or volume) regardless of size is automatically spread across a large number of disks for optimal performance and can grow or shrink nondisruptively. NetApp has already deployed tens of petabytes of storage in shared storage environments in which customers routinely scale volumes up and down as needs change.

Automated storage processes

For storage automation, NetApp provides a suite of management products that simplify storage operations, including data protection.

- NetApp Provisioning Manager allows you to create repeatable, automated provisioning processes based on policies you define. An at-a-glance dashboard lets you monitor a variety of metrics, including capacity utilization, policy compliance, and space management statistics.

- NetApp Protection Manager (http://partners.netapp.com/go/techontap/mat/backup_mgmt.html) provides capabilities similar to those of Provisioning Manager, but they are focused specifically on automation of data protection and replication.
- NetApp SANscreen® provides a real-time, multivendor, and multiprotocol service-level view of your entire storage environment.
- The NetApp SnapManager® suite of products extends the capabilities of NetApp storage to integrate with popular applications, including Oracle®, SAP®, Exchange, Microsoft® SQL Server®, and Microsoft SharePoint®.

Usage monitoring and chargeback

By automatically correlating your end-to-end storage infrastructure against business services, SANscreen allows you to monitor service-level agreements, meter usage for cost awareness and chargeback, and proactively manage capacity to maintain optimal utilization of all resources. Provisioning Manager also includes chargeback capabilities.

Moving data freely

NetApp has a long history of moving data efficiently during data protection operations by using NetApp SnapMirror® or SnapVault® software, NDMP, and so on. We continue to investigate more efficient ways to facilitate data movement in conjunction with VMware (www.netapp.com/us/communities/techontap/tot-vstorage-0309.html) and other virtualization partners wherever possible. Because data movement has been identified as a critical element for cloud infrastructure, we are focusing on this area.

Multi-tenancy

Traditionally, making sure of the highest level of storage isolation and security has meant independent hardware. Both private and public clouds need to be certain that security is as tight as it can be without sacrificing efficiency. NetApp MultiStore software lets

you create multiple, separate, and completely private logical partitions on a single storage system, so you can share storage without compromising privacy and security. (See Figure 3 on page 19.)

With MultiStore, storage can be provisioned, moved, and protected based on the boundaries that you define; virtual storage containers allow you to apply policies appropriate to each container that could correspond to a particular application or client of a cloud service.

Storage efficiency

By boosting storage efficiency, you can significantly reduce your total storage requirements, which translates into savings in power, cooling, and space. Because of the flexible provisioning enabled by NetApp FlexVol technology, NetApp customers routinely achieve storage utilization rates of 60% or higher, whereas the industry average is closer to 30%. NetApp offers a range of additional technologies that further boost storage utilization, as summarized in Table 1 on page 19.

➤ Continued on page 19



JEFF O'NEAL
SENIOR DIRECTOR OF DATA CENTER SOLUTIONS
NETAPP

Jeff and his team are responsible for market strategies and alliances for all NetApp data center solutions, including core systems, data center networking, manageability and data center automation software, and dynamic data center and cloud computing solutions. ■

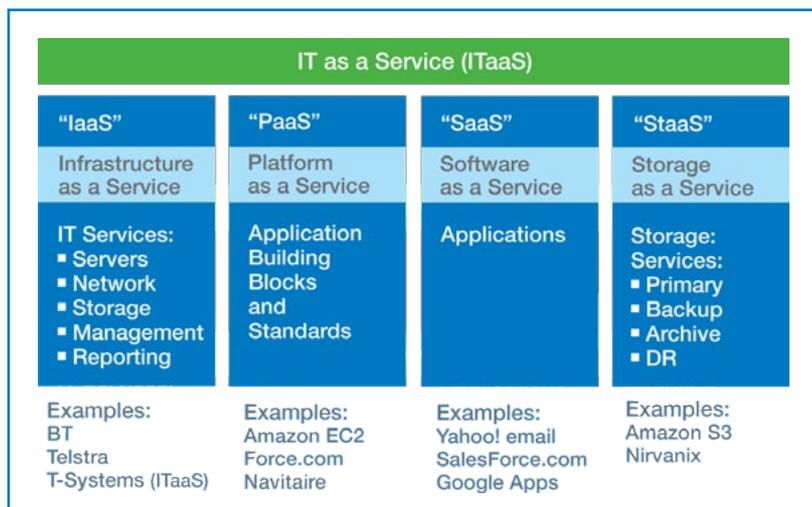


Figure 1) Different types of cloud services that are available.

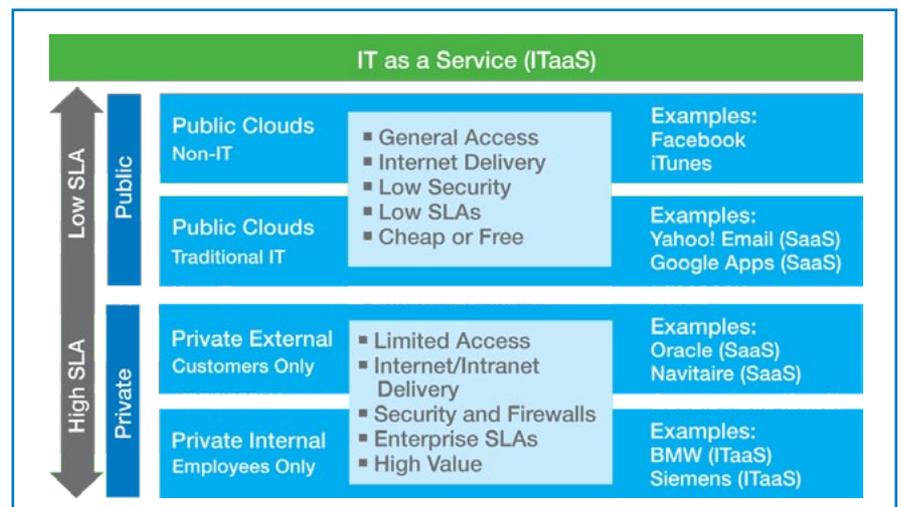


Figure 2) Private versus public clouds.

Boosting Data Center Efficiency:

Innovate and Thrive in a Difficult Economy

By Dave Robbins

➤ Continued from page 1

In this article, I explore this problem in a bit more detail, discuss how you can enhance your data centers and IT practices to improve efficiency, and describe the implications of these changes for storage administrators.

THE HIDDEN COST OF IT

If you deploy that low-cost server in an enterprise-class, tier-2 data center, the direct costs to power and cool it for a year could amount to as much as \$8,300. The associated operational expenses would be about \$1,320 annually, covering the depreciation of the capital investment, electricity, and site operations (security and so on).

Most IT departments do not include these costs in their business justifications for new server deployments. In fact, most IT organizations might never see a power bill until their data centers are out of power—so they have no idea these hidden costs even exist.

Since these costs are not seen, data center managers might not have the incentive to control them. At current growth rates we would need to build more than 10 new power plants in the United States alone, at a cost of \$2–\$6 billion each, to meet the growing power demands from data centers.

Of U.S. and EU companies surveyed:

- 47% do not monitor server utilization.
- 55% of data center managers are not aware of their monthly power costs.
- 43% of EU data centers are power constrained.
- 42% of U.S. data centers will run out of power in the next 18–24 months.

The key to getting your data center ready for the future as well as controlling hidden costs is to increase power efficiency and boost asset utilization to decrease the amount of equipment you need and slow the rate of future acquisition. You can accomplish this by:

- Designing facilities to support higher power density
- Rearchitecting core networks

- Virtualizing servers and storage so you can scale and move applications more fluidly

At NetApp, we've gained a lot of experience with all three of these items through efforts to improve the efficiency of our own data centers.

REDESIGNING DATA CENTER FACILITIES FOR EFFICIENCY AND DENSITY

To boost efficiency and density in NetApp data centers, we found that we needed to challenge a lot of old assumptions. Here are some of the things we changed on our way to increased efficiency:

- Facilities management and IT strategy are typically discrete functions, but we found that teaming up results in more energy-efficient solutions.
- Raised floors in a data center are artifacts of mainframe days. A cold aisle/hot aisle layout enables power and cooling delivery from overhead and is much more energy efficient.
- High-tech problems can sometimes be solved with low-tech solutions. A vinyl curtain can create a physical barrier between hot and cold aisles, saving more than 1 million KWh per year.
- Data centers don't have to be cold. We've gradually moved up the supply air temperature from 52° F to 70° F, and we let our hot aisle temperatures go as high as 95° F. This has resulted in less cooling and more hours of free cooling using outside air.

One important measure of data center efficiency is power usage effectiveness (PUE), the ratio of total facility power used to power used directly by IT equipment. Possible sources of electricity consumption are illustrated in Figure 1.

The current benchmark for good data center design is a PUE of 2.0: Everything else in the data center consumes as much power as the IT equipment. Other important measures of data center efficiency include:

- **How much power you can deliver per rack.** A typical value today is about 3kW. In other words, the equipment deployed in the rack can draw a maximum of 3kW of power.



- **How much space you need per rack.**

A rack in a typical data center can require as much as 28 square feet of floor space.

By applying techniques such as those mentioned above, NetApp has been able to lower the PUE in its data centers substantially below the 2.0 benchmark while delivering much more power per rack and consuming much less space.

For instance, our second-generation data center design uses variable air volume (VAV) control and a cold aisle damper to achieve a PUE of 1.40 with 801 racks drawing 3.5KW/rack (see Figure 2). This saves the company approximately \$1.7 million per year versus the same data center operating at a PUE of 2.0.

Our third-generation design resulted in a data center with 720 racks drawing 8KW/rack with a PUE of 1.30 that will save the company \$4.3 million per year versus a PUE of 2.0 (see Figure 3). We're also currently deploying a new combined IT and engineering data center that will include 1,800 engineering racks capable of delivering 12KW/rack with a predicted PUE of 1.2. Hot air from the data center is used to heat office space during the winter months.

You can read more about the strategies that we use to boost data center efficiency in a recent Tech OnTap article (http://partners.netapp.com/go/techontap/mat/DC_efficiency.html) as well as a recent white paper (<http://media.netapp.com/documents/wp-netapp-data-center-power-efficiency.pdf>).

Increasing the efficiency and density of your data center will only get you so far. You also need to substantially increase the utilization of the IT assets you deploy in your data centers by rearchitecting your core networks, consolidating as much as possible, and virtualizing both servers and storage. You can read about one approach to full data center virtualization using a combination of the Cisco Unified Computing System (UCS) and NetApp unified storage in a Tech OnTap article (www.netapp.com/us/communities/tech-ontap/tot-netappandcisco-datacenter-0509.html).

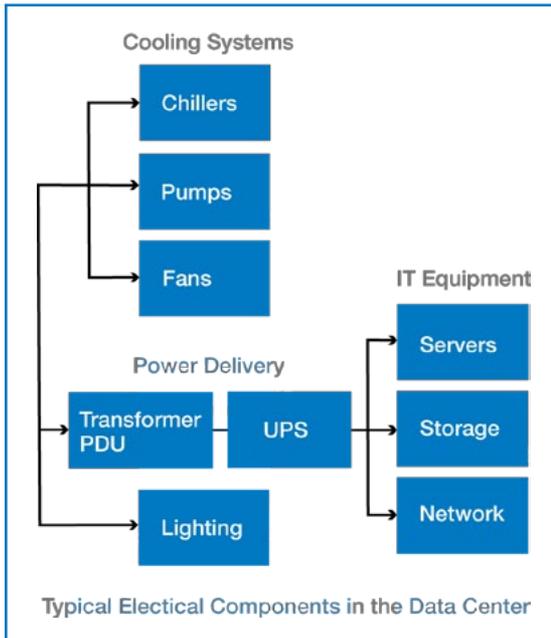


Figure 1) Sources of power consumption.

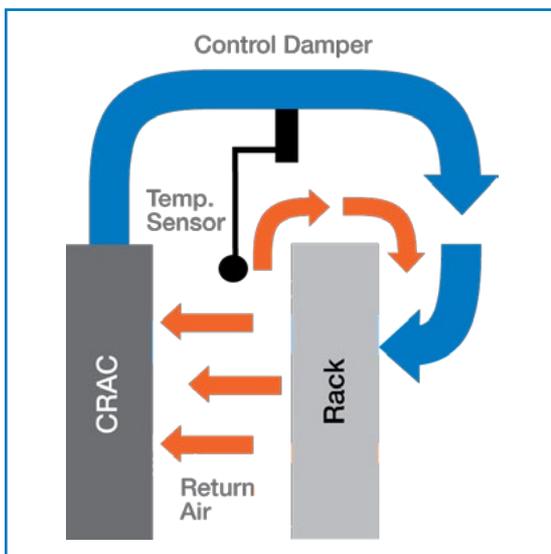


Figure 2) NetApp second-generation data center design.

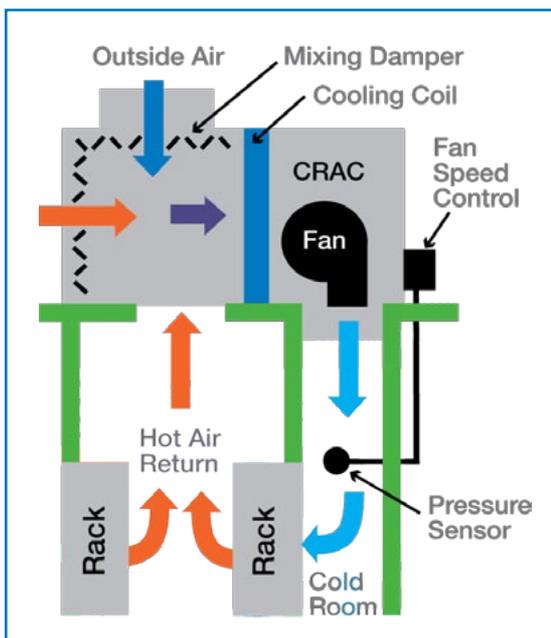


Figure 3) NetApp third-generation data center design.

REARCHITECTING CORE NETWORKS

In terms of core networks, there are two objectives:

- Standardize on a unified network fabric.
- Rearchitect your networks so that applications aren't "boxed in."

The arrival of converged network adapters (CNAs) and Fibre Channel over Ethernet (FCoE) makes it possible to incorporate your existing Fibre Channel devices on a single network fabric. Eliminating separate storage network infrastructures can further reduce your data center power consumption. By wiring once, you can connect to any network—SAN, LAN, or HPC—for faster rollout of new applications and services.

Most current networks are designed so that key business applications are isolated for security and performance. Unfortunately, this localizes IT resources and limits utilization. By opening up your networks, you make it possible for every host to mount any storage target. This drives storage consolidation, boosts utilization, and sets the stage for full virtualization of servers and storage.

VIRTUALIZE EVERYTHING

Rearchitecting your core networks makes it possible to take the greatest advantage of server and storage virtualization. Existing resources can be fully consolidated, and pools of resources can be efficiently brought to bear on any application. Because applications are free to move to any set of resources in the data center, or to alternate data centers, downtime can be significantly reduced.

NetApp has moved aggressively to consolidate older servers and storage. For example, a review of our engineering labs identified 4,600 x86 client servers that are candidates for virtualization. These systems result in about \$1.4 million in power and cooling costs and consume 192 racks of data center space. Full virtualization of these systems at 20:1 will result in a reduction in the total system count to 230 servers, resulting in about \$70,000 in annual power and cooling and occupying only 10 racks.

The initial phases of such a project in one engineering lab were described in a recent Tech OnTap article (<http://www.netapp.com/us/communities/tech-ontap/tot-ntp-p2vmigration-0409.html>).

Another Tech OnTap article focused specifically on NetApp's efforts to consolidate storage in one of our corporate data centers (http://partners.netapp.com/go/techontap/matl/netappit_case_study.html). This resulted in:

- An increase in storage utilization to an average of 60%
- A reduction in storage footprint from 24.83 racks to 5.48
- Replacement of 50 storage systems with just 10
- A decrease in direct power consumption of 41,184 kWh per month
- Substantial capacity and performance gains

The general approach is described in a white paper on reducing power consumption through storage efficiency (<http://partners.netapp.com/go/techontap/matl/reducing-datacenter-power.pdf>).

Many NetApp customers have achieved similar results in their own data centers. For instance, one large telecom customer consolidated 3,103 physical servers down to 134 (23:1) and increased storage utilization from 25% to 70% while saving \$2.25 million on power per year and freeing 660 racks and 8,500 ports. The project also resulted in a return on investment in just 8 months, enabled same-day server provisioning, and cut backup time from 96 hours to less than 30 minutes.

IMPLICATIONS FOR STORAGE ADMINISTRATION

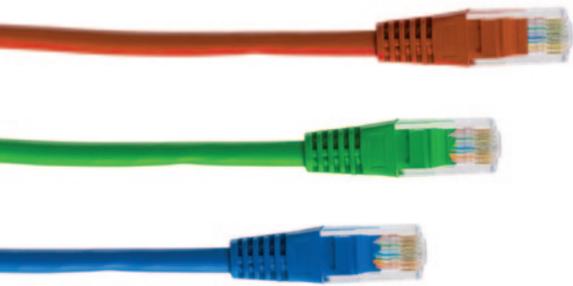
The move to a massively shared, virtualized ecosystem has some significant implications for storage administrators. Because servers and storage will now need to be accessible from everywhere—even from outside the data center—the way you do things will necessarily change.

➤ Continued on page 18



DAVE ROBBINS
CHIEF TECHNOLOGY OFFICER FOR IT
NETAPP

As an IT practitioner since 1979, Dave has seen and participated in many technology evolutions that have contributed to his keen understanding of data center efficiency and emerging cloud technologies. At NetApp, he is responsible for identifying and selecting new technologies and establishing adoption road maps for NetApp IT. He works hard to make sure of the use of NetApp technologies within NetApp IT, which serve as a model of best practices. ■



Network Convergence:

Deploying FCoE in Your Data Center

By Mike McNamara, NetApp, and Ahmad Zamer, Brocade

Many enterprise data centers use Ethernet networks for LAN and IP data traffic plus separate Fibre Channel (FC) networks for storage area network (SAN) traffic. The increased adoption of 10-Gigabit Ethernet (10GbE) in the data center, combined with the availability of Fibre Channel over Ethernet (FCoE) and new lossless 10GbE technologies, makes it possible to consolidate FC data flows with LAN and IP data traffic on the same Ethernet infrastructure. Network convergence lets you preserve your existing investments in FC storage, reduce data center costs and complexity, and simplify network management.

Although the benefits of using FCoE are compelling, many are still waiting to deploy the technology. This article addresses frequently asked questions about the technology and concludes with information on how you can make the move to FCoE using a gradual, phased approach.

IT CHALLENGES: MAINTAINING MULTIPLE NETWORKS

Most data centers maintain multiple networks for different purposes:

- Ethernet for local area networks (LANs) to transfer small amounts of information across short or long distances or in clustered computing environments. Ethernet provides a cost-effective and efficient way to support a variety of data types, including corporate LANs; voice-over-IP telephony; and storage with NFS, CIFS, and iSCSI.
- Fibre Channel for storage area networks (SANs) to provide block I/O for applications such as network booting; mail servers; and large, data-intensive databases. FC SANs are an excellent solution for storage consolidation, centralized storage management, high performance, reliability, and business continuance.

IP networks and FC SANs each play an essential role in the data center, but they differ in design and functionality. The two networks have their own security needs and traffic patterns, and use separate management toolsets. Each network is built and maintained on dedicated infrastructure, with separate

cabling and separate network interfaces on each server and storage system.

Managing two discrete networks increases the complexity and cost of your data center. Converging your Ethernet and FC networks can make your data center more efficient without sacrificing your investment in FC infrastructure.

FIBRE CHANNEL OVER ETHERNET

FCoE enables you to transmit IP and FC traffic on a single, unified Ethernet cable. In this way, the merged network can support LAN and SAN data types, reducing equipment and cabling in the data center while simultaneously lowering the power and cooling load associated with that equipment. There are also fewer support points when consolidating to a unified network, which helps reduce the management burden.

FCoE is enabled by an enhanced 10GbE technology commonly referred to as data center bridging (DCB) or Converged Enhanced Ethernet (CEE). Tunneling protocols, such as FCiP and iFCP, use IP to transmit FC traffic over long distances, but FCoE is a layer-2 encapsulation protocol that uses Ethernet physical transport to transmit FC data. Recent advances and upcoming additions to the Ethernet standard, such as TRILL (see sidebar) and the ability to provide lossless fabric characteristics over a 10-Gigabit link, are what enable FCoE.

FCoE delivers significant value to organizations that want to consolidate server I/O, network, and storage interconnects by converging onto a single network storage technology. For data centers with large investments, even the simplest reduction in the amount of equipment that has to be managed can reap significant benefits. And sharing the same network fabric—from server to switch to storage—removes the requirement of dedicated networks, significantly reducing TCO by preserving existing infrastructure investments and maintaining backward compatibility with familiar IT procedures and processes.

FCOE COMPONENTS

Some of the components needed to implement FCoE include:

- **Converged network adapters (CNAs).** These combine the functionality of Ethernet NICs and Fibre Channel host bus adapters (HBAs), reducing the number of server adapters you need to buy, cutting port count, and eliminating a healthy number of cables.
- **FCoE cables.** There are currently two options for FCoE cables: the optical cabling generally found in FC SANs and a new type of Twinax copper cabling. FCoE twin cables require less power and are less expensive, but, because their length is limited to less than 10 meters, you will likely need optical cabling to reach from top-of-rack switches to the LAN.
- **FCoE switches.** You need FCoE/DCB switches to connect servers to your storage arrays or native FCoE storage systems. For the early adopters, that means top-of-rack switches or end-of-row blades where possible.
- **FCoE/DCB storage systems.** These storage systems natively support FCoE and converged traffic. There are also storage systems that support Fibre Channel to an FCoE switch and FCoE from the switch to host servers.

IMPACT ON EXISTING SERVERS, NETWORKING, AND STORAGE

FCoE requires minimal changes to your existing IT infrastructure. It is a natural evolution of Fibre Channel technology, designed to carry data over Ethernet physical and data-link layers. Using Fibre Channel's upper layers simplifies FCoE deployment by allowing coexistence with deployed FC SANs and enables you to leverage enterprise-proven Fibre Channel software stacks, management tools, and existing training. Most importantly, you don't need to change your applications in order to benefit from the performance and potential cost benefits of FCoE.

ORGANIZATIONAL ISSUES

In traditional data center environments, the storage group owns and operates the FC SAN, while the networking group owns and operates the Ethernet LAN. Since the two groups have been historically separate, introducing FCoE into the data center might introduce beneficial changes to some IT practices.

Cultural, political, and behavioral concerns in data center and provisioning paradigms can present organizational obstacles to FCoE adoption. Some new business processes and procedures might need to be implemented so that proper control mechanisms are in place for FCoE networks. Purchasing patterns might have to be modified, and the reliability of Ethernet networks might have to be increased.

With the convergence of FC and IP created by FCoE, these two traditionally separate network realms overlap. Implementing FCoE requires little if any additional IT training. FCoE leverages the existing IT expertise and skill sets of your IP data and FC teams. Role-based management features in management applications allow your FC group to continue owning and operating the SAN and your IP networking group to continue owning and operating the data network.

WHERE TO DEPLOY

While the benefits of using FCoE are certainly compelling, you might still be waiting to deploy the technology. Fortunately, FCoE convergence is not a disruptive process and does not require a “rip and replace” upgrade. Moving to FCoE can be done gradually, using a phased approach. Most early FCoE deployments will likely be part of new server deployments in Windows® and Linux® environments in which virtualized tier-3 and some tier-2 applications are deployed.

Considering that FCoE is a relatively new technology, initial FCoE deployment is best suited for access-layer server I/O consolidation. As storage traffic requires the new lossless Ethernet, the 10GbE transport still requires Link Layer multipathing and multihop capabilities. Such features are currently under development, and should become available later in 2010. These capabilities will enable the deployment of larger FCoE networks, which will expand the reach of FCoE beyond access layer server connectivity and I/O consolidation.

Best practices for determining where to deploy FCoE include:

- Choose environments that already have a Fibre Channel skill base and Fibre Channel infrastructure
- “Green-field” deployments, in which new infrastructure is being introduced to accommodate data growth
- Consider beginning the transition to FCoE in your tier-3 or tier-2 applications; gain experience in labs or less mission-critical tier-3 environments and then use what you’ve learned to make the transition in tier-2 and, in some instances, tier-1 applications
- Start implementing FCoE on the top-of-rack access layer server I/O consolidation side—

that step may be combined with native FCoE storage deployment; extending FCoE beyond access layer servers should wait for multipathing and multihop standards (TRILL) to become practical

HOW TO BEGIN

Migration to FCoE can be accomplished with a gradual, phased approach, typically starting at the edge or switch, then moving to native FCoE storage, and eventually going deeper into the corporate network. Figure 1 depicts a typical data center architecture before network convergence begins. The FC SAN (illustrated by the orange line) is a parallel network requiring network ports and cabling over and above those required for the Ethernet IP LAN (illustrated by the blue line).

PHASE 1: MAKING THE TRANSITION TO DCB/FCOE AT THE EDGE OR SWITCH

Moving to a converged or unified Ethernet infrastructure can be done gradually and will likely begin at the edge (illustrated by the green lines in Figure 2), where the greatest return on investment can be realized. With FCoE convergence, port count at the servers and edge switches can be reduced by half,

➤ Continued on page 15



MIKE MCNAMARA
SR. MANAGER,
PRODUCT
MARKETING
NETAPP

Mike has over 20 years of computer industry marketing experience, 15 years of which have been specifically focused on storage. He worked at Adaptec, EMC, and Hewlett Packard before joining NetApp more than four years ago. Mike is also the marketing chairperson for the Fibre Channel Industry Association (FCIA).



AHMAD ZAMER
SR. PRODUCT
MARKETING
MANAGER
BROCADE

Ahmad has over 25 years of computer-industry experience, with special emphasis on networking and computer storage technologies. He worked at Philips and Intel before joining Brocade. Ahmad is a technical writer with more than 50 published articles to his credit. ■

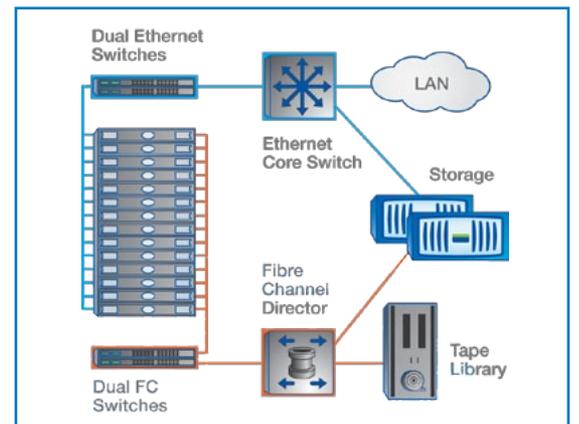


Figure 1) Layout of a typical data center before implementing DCB/FCoE.

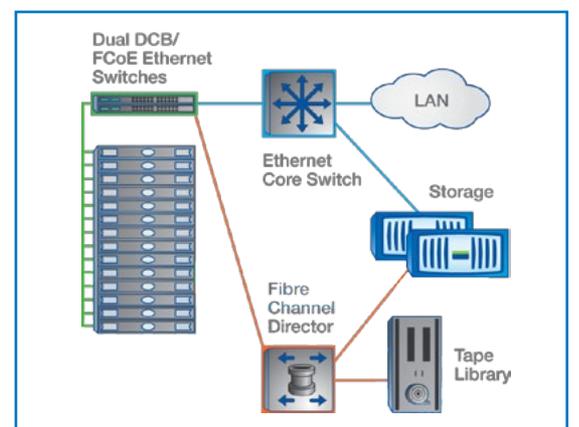


Figure 2) Phase 1: Making the transition to FCoE at the edge or switch.

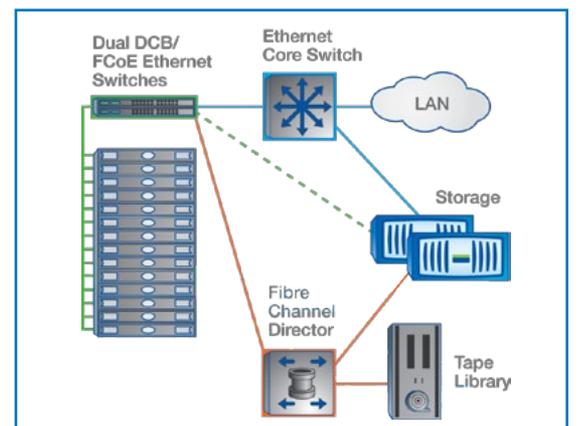


Figure 3) Phase 2: Making the transition to native FCoE storage.

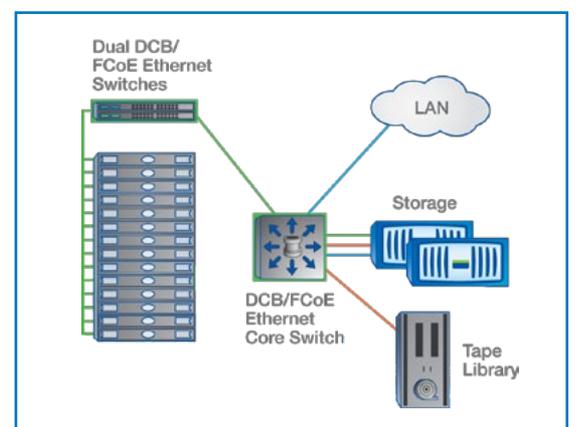


Figure 4) Phase 3: End-to-end FCoE, from edge to core to storage.

NetApp and Cisco: Virtualizing the Data Center

By Mike McNamara



The enterprise data center is undergoing a transformation: Server virtualization technology is changing the way servers and applications are provisioned while altering the workload and increasing storage demand. This data center transformation makes the idea of unified computing and storage appealing.

In March 2009, Cisco announced its Cisco Unified Computing System (UCS) (http://newsroom.cisco.com/dlls/2009/prod_031609.html), an architecture that integrates compute, networking, and virtualization in a single platform. In a

require multiple provisioning toolsets, data management models, and teams of people to manage. In addition, massive data growth, challenging economic conditions, and the physical limitations of power, heat, and space exert extreme pressure on IT staffs.

Cisco and NetApp have worked to create a holistic approach that allows the network and the applications it supports to work together in the most efficient way possible. The primary goals of the collaboration are to reduce costs, improve agility, and increase business within and between data centers.

Channel over Ethernet (FCoE) encapsulation. It eliminates the limitations of fixed I/O configurations with an I/O architecture that can be changed through software on a per-server basis to provide needed connectivity. NetApp's unified storage architecture supports multiple protocols—including native FCoE—providing investment protection, flexibility, and simplified configuration.

REDUCED MANAGEMENT OVERHEAD

A typical data center has to cope with far too many management interfaces. This joint solution allows you to manage your entire data center using just a few tools.

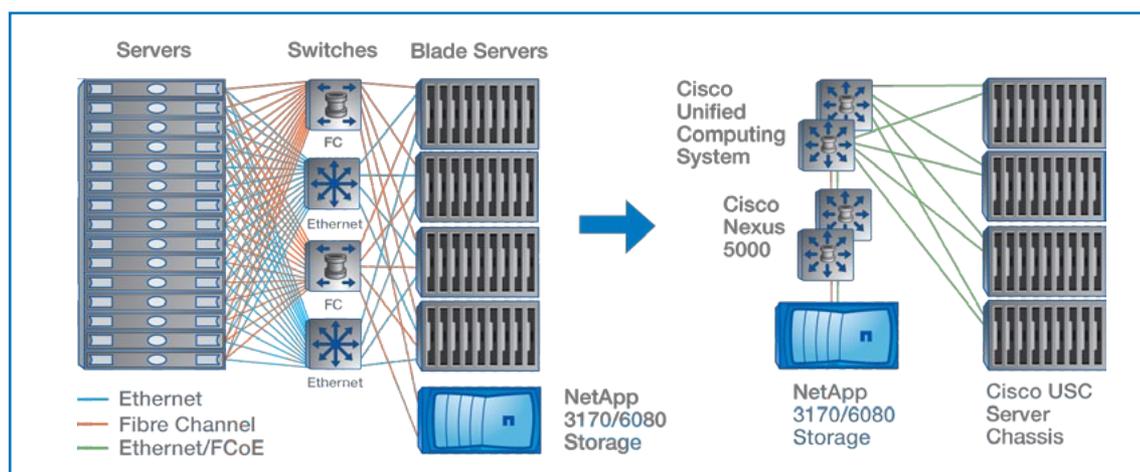


Figure 1) Simplified Cisco and NetApp data center architecture.

subsequent announcement soon after, Cisco announced a partnership with NetApp that pairs Cisco UCS with the advanced capabilities of the NetApp unified storage architecture to create a complete and easy-to-deploy virtualization solution for the dynamic data center (http://newsroom.cisco.com/dlls/2009/prod_041609c.html).

In this article, I discuss the technology highlights of this solution. You can learn more about the technical details and applications in a recent joint white paper (<http://media.netapp.com/documents/cisco-netapp-simplified-data-center-architecture.pdf>).

VIRTUALIZED DYNAMIC DATA CENTER

Current data center architectures are built around a complex, heterogeneous collection of servers and storage systems that result in poor utilization and captive resources and that

These goals are achieved by:

- Simplifying the overall architecture
- Reducing management complexity
- Enabling zero cost provisioning
- Leveraging NetApp storage efficiency, data protection, and availability technologies

SIMPLIFYING DATA CENTER ARCHITECTURE

With this joint solution your data center consists of just two elements: Cisco UCS provides compute, virtualization, and networking, while NetApp unified storage handles data storage and related storage functions such as data protection, replication, and so on.

Cisco UCS reduces the number of access-layer switches in the network and integrates compute resources around a unified I/O fabric that supports standard IP protocols and Fibre

Cisco UCS Manager lets you manage all UCS elements as a single, redundant, uniform pool of resources that can be configured on demand. The manager is embedded in the network fabric and automatically discovers resources as they are installed, adds them to inventory, and can automatically provision servers and I/O connectivity, putting new server instances into use in minutes rather than hours or days.

NetApp delivers unification of data storage and data management with SANscreen and Provisioning Manager. SANscreen extends data center automation to storage with real-time, multiprotocol, service-level views of your data center's storage environment.

> Continued on page 18



MIKE MCNAMARA
SR. MANAGER,
PRODUCT
MARKETING
NETAPP

Mike has 20 years of computer industry marketing experience, with 15 years specifically focused on storage. He worked at Adaptec, EMC, and Hewlett Packard before joining NetApp more than three years ago. Mike is also the marketing chairperson for the Fibre Channel Industry Association (FCIA). ■

A Quantum Leap in Virtual Management

By Trey Layton



Recently, I wrote a Tech OnTap article about a business process outsourcing company that uses the NetApp rapid cloning utility (RCU) to quickly deploy call center desktops (www.netapp.com/us/communities/tech-ontap/tot-vdi-0309.html). The approach I described has now become an integral part of that company's business process and a significant selling point for new business.

On a recent flight from Texas to New York, one of the company's virtualization engineers had a chance to demonstrate the power of this approach to a potential client with whom he was traveling. As they discussed the technical specifics, the client became curious about how the process worked.

Earlier that day, the engineer had noticed an advertisement for Gogo[®] Inflight Internet in the airport terminal and found a free code to try the service in the seatback pocket. Using Gogo, he was able to log in from the airplane through the company VPN to VMware[®] vCenter and launch the RCU plug-in to clone 300 virtual desktops. In the time it took the flight attendant to serve dinner, the entire process was completed, and the engineer booted a virtual desktop for a demo at 30,000 feet. The client was amazed that the process was so simple it could be done from an airplane. He was further impressed by the efficiency of the process—which would make it quick and easy to adapt to changing needs—and the 300 desktops only consumed a tiny amount of storage.

NetApp has taken significant steps to improve the management experience for users of VMware. In addition to the recent release of RCU 2.1, SnapManager for Virtual Infrastructure version 2.0 and the Virtual Storage Console version 1.0 have recently been released. All of these work with VMware vCenter, and they all generated a lot of interest during demonstrations at the VMworld 2009 conference in San Francisco.

RAPID CLONING UTILITY 2.1

RCU is a free management plug-in for VMware vCenter that works with NetApp FlexClone[®] to automate and accelerate virtual server and desktop provisioning in VMware ESX environments. RCU lets you:

- Automate VM-level and data store-level cloning within vCenter
- Use NetApp FlexClone to perform instantaneous cloning while consuming almost no additional capacity
- Automatically customize the guest operating system of each cloned VM and import into VMware View Manager

Several Tech OnTap articles have described the use of RCU: “vStorage Integration” (www.netapp.com/us/communities/tech-ontap/tot-vstorage-0309.html), “Consolidating 1,000 Physical Servers” (www.netapp.com/us/communities/tech-ontap/tot-ntp-p2vmigration-0409.html), and “A 9,000-Seat VDI Deployment” (www.netapp.com/us/communities/tech-ontap/tot-vdi-0309.html).

RCU 2.1 adds a number of new capabilities to the functionality of RCU 2.0, including:

- Deduplication management
- Data store provisioning
- Support for cloning on Fibre Channel, iSCSI, and NFS storage

RCU 2.1 is available on the NOW[™] site and requires Data ONTAP 7.3.1.1, FlexClone, and appropriate protocol licenses. You can watch a demo of RCU functionality (www.youtube.com/netapptv#play/user/FB8CE7ACF0FE6DEA/9/vqidYMcrOvI) or read the RCU v2.1 FAQ (<http://communities.netapp.com/docs/DOC-3870>).

SNAPMANAGER FOR VIRTUAL INFRASTRUCTURE 2.0

NetApp SnapManager for Virtual Infrastructure (SMVI) is designed to simplify the management of backup, restore, and disaster recovery operations in VMware environments. Like RCU, SMVI works with VMware vCenter for seamless operation. SMVI v2.0 adds a number of new features, including restore enhancements such as:

- **Single file restore.** Lets you restore one or more files from a guest VM disk (VMDK) without restoring the entire VM.
- **Self-service restore.** Users can see a list of backups for a VM, browse, and restore files.

You can watch a demo of SMVI v2.0 single file restore (www.youtube.com/netapptv#play/user/FB8CE7ACF0FE6DEA/6/fPg8FNaA_MY).

You can read more details about these features in a blogpost from frequent Tech OnTap contributor Nick Triantos (http://blogs.netapp.com/storage_nuts_n_bolts/2009/09/snapmanager-for-virtual-infrastructure-20.html). A series of blog posts from Kostadis Roussos in 2008 provides a great overview of SMVI (http://blogs.netapp.com/extensible_netapp/2008/08/smvi-part-i-the.html).

VIRTUAL STORAGE CONSOLE 1.0

The virtual storage console (VSC) is another vCenter plug-in that enables administrators to monitor and manage specific storage-side attributes of ESX hosts for hosts using both SAN and NAS protocols. VSC replaces the NetApp ESX Host Utilities Kit.

When you configure NetApp storage with VSC, it automatically makes sure that storage settings such as multipath settings for Fibre Channel, HBA timeouts, or NFS settings adhere to NetApp best practices. It also provides end-to-end troubleshooting should a connectivity problem arise. VSC incorporates the best practices for NetApp and VMware as

➤ Continued on page 17



TREY LAYTON
SYSTEMS
ENGINEER
NETAPP

Trey has been working at NetApp since 2006, specializing in the design of next-generation data centers using VMware. His wealth of experience with networking and virtualization makes him well suited to the current evolution to network storage. With over 18 years of IT experience, Trey began his career in the U.S. Army at USCENTCOM supporting U.S. Special Operations in the Middle East. He has held key network consulting and systems engineering positions at Eastman Kodak, GE, and Cisco. Read Trey's “Ethernet Storage Guy” blog (<http://communities.netapp.com/blogs/ethernetstorageguy>). ■

Virtualizing Microsoft Applications

By Abhinav Joshi



As you move toward the goal of 100% virtualization in your data center, careful attention to the virtualization of business-critical Microsoft applications—including Microsoft Exchange, Microsoft SQL Server, and Microsoft SharePoint Server—becomes essential.

To get from where you are today to an environment that delivers all the benefits of virtualization, including efficiency, improved availability, and decreased cost, you have to focus on virtualization of all layers of your infrastructure, including virtualization software, servers, networks, and storage (see Figure 1).

That's why NetApp joined forces with Cisco and VMware to create a complete solution for virtualizing Microsoft applications. This architecture combines the benefits of VMware vSphere 4 virtual infrastructure, Cisco Nexus unified fabric, and NetApp unified storage hardware and software.

This flexible architecture allows you to virtualize a mixed workload Microsoft application environment to deliver the full benefits of server, network, and storage virtualization. We've tested the performance of Microsoft applications on this solution to make sure there are no bottlenecks and that all performance metrics are well within Microsoft's published parameters.

This article briefly describes the reasons for virtualizing Microsoft applications and highlights the most important architecture and deployment considerations to help you get started. For full details on the joint solution, refer to the NetApp technical report "NetApp Solutions Guide: Microsoft Exchange Server, SQL Server, and SharePoint Server Mixed Workload on VMware vSphere 4, NetApp Unified Storage (FC, iSCSI, and NFS), and Cisco Nexus Unified Fabric" (www.netapp.com/us/library/technical-reports/tr-3785.html).

WHY VIRTUALIZE?

The reasons for virtualizing Microsoft applications are in large part the same reasons for virtualizing any application:

- **Reduced costs.** Upgrading to newer Microsoft server applications without virtualization can require even more server hardware to support an application that has already become excessively costly to run. VMware virtualization unlocks the full power of your existing hardware by running multiple workloads on each system. Increased utilization means less hardware and lower overall capital and management costs.
- **Advanced storage capabilities.** You can deploy Microsoft Exchange, SQL Server, and SharePoint on NetApp storage across any storage protocol, including FC, iSCSI, or NFS. NetApp FAS and V-Series storage arrays have been fully tested and certified for use in FC and IP-based VMware environments. By using NetApp storage efficiency and intelligent caching capabilities, you save significantly on storage costs. NetApp virtualizes storage by pooling available IOPS and capacity for on-demand use by multiple applications similar to how VMware virtualizes server resources.
- **High availability.** VMware can provide high availability (HA) for Microsoft server applications without the need for clustering at the virtual machine (VM) level. VMs are no longer tied to the underlying server hardware and can be moved across servers at any time with VMware VMotion. VMware HA provides server hardware fault tolerance for every VM and offers greater levels of availability over solutions designed to protect just the server. NetApp active-active storage configurations provide similar capabilities at the storage level.
- **Advanced backup/recovery and DR.** Backup and recovery for this solution are built using integrated VMware, Microsoft, and NetApp technologies for advanced, application-aware data protection. Deduplication-aware remote replication for disaster recovery with NetApp SnapMirror provides end-to-end data protection, and the addition of VMware Site Recovery Manager can automate the entire recovery process.
- **Enhanced mobility.** You have the option of easily and nondisruptively relocating

the virtual machines and/or storage used by your Microsoft applications for load balancing, upgrades, and maintenance or to meet other organizational goals.

Despite these obvious benefits, there are two persistent concerns about virtualizing critical Microsoft applications, but these concerns have been addressed:

- **Performance.** With the release of VMware vSphere 4.0, VMware has increased performance to the point where it is suitable for any mission-critical business application, as demonstrated in a recent VMware white paper that describes Exchange performance using various storage protocols (www.vmware.com/files/pdf/vsphere_perf_exchange-storage-protocols.pdf). We made performance validation a key part of the development of this solution to address this lingering concern.
- **Support.** There are still widespread concerns about support for virtualized Microsoft apps. The good news is that there are multiple ways to get necessary support. Microsoft fully supports virtualization through its Server Virtualization Validation Program—SVVP (www.windowsservercatalog.com/svvp.aspx?svvppage=svvp.htm). You also qualify for direct support of virtualized Microsoft applications if you contract for the Microsoft Services Premier Support Program. Support might also be available through your server OEM vendor, VMware Global Support Services (GSS), and Technical Support Alliance Network (TSANet).

KEY DESIGN CONSIDERATIONS

One of our key goals when architecting the joint solution was to provide clear design guidelines and at the same time provide enough flexibility so that you can create a solution that is tailored to meet the requirements of your environment. This section is structured around some of the key questions you will want to ask yourself as you move your Microsoft applications to a fully virtualized environment.

What storage protocol should I choose?

One of the great things about this solution—like all solutions that include NetApp storage—is that you have the flexibility to choose whatever storage protocol makes sense for your environment. We provide architecture guidelines for all protocols: FC, iSCSI, and NFS. A joint NetApp and VMware performance study demonstrates that all protocols perform within 10% of one another, so there is no reason based on performance to choose one protocol over another (www.netapp.com/us/library/technical-reports/tr-3808.html).

If you already have Fibre Channel (FC) infrastructure, you can continue to use it. If not, NFS and/or iSCSI can easily meet your storage needs. I advise you to look at each protocol in terms of cost to you (capital and operational), manageability, scalability, and flexibility and choose the one that fits your needs best. (A few more specific guidelines are forthcoming in the section on storage layout.)

What netapp software will I need?

We strongly recommend the use of a core set of four NetApp products:

- **Rapid Cloning Utility (RCU).** This free vCenter plug-in provides rapid, space-efficient provisioning of virtual servers and desktops leveraging NetApp FlexClone, data store deduplication management, data store provisioning, resize, and destroy operations.
- **Virtual Storage Console (VSC).** This free vCenter plug-in lets you manage and monitor NetApp specific storage-side attributes that pertain to VMware directly from within vCenter.
- **SnapManager for Virtual Infrastructure.** SMVI is an integrated data protection solution. It provides backup and recovery

for virtual machines and replication for DR. It uses the NetApp Snapshot™ capability with the option to invoke VMware snapshots for VM backups.

- **SANscreen VM Insight.** This vCenter plug-in provides monitoring and extensive reporting on the virtual to physical storage mapping (vmdk, data stores, LUN, storage fabric) to help in environment management and troubleshooting.

You can learn much more about the first three software tools shown above in a recent Tech OnTap article devoted to the subject (www.netapp.com/us/communities/tech-ontap/tot-vmware-management-0910.html).

In addition, you will want to install NetApp SnapDrive® and the application-specific SnapManager product inside guest VMs that hosts an Exchange Mailbox server, SQL Server, or SharePoint database and index server to provide application-consistent backup and granular restores of databases, logs, and so on. (Backup and DR are covered in more detail later.)

What storage layout should I use for different data components?

The storage layout you choose will depend in part on the storage protocol you have selected. Rather than trying to cover all possible storage layout and protocol options here, I'll simply focus on one of the most flexible IP-based storage layout options. If you are deploying from scratch or your infrastructure will support this approach, the layout shown in Figure 2 on page 14, combining NFS and iSCSI, is the one I would suggest. For FC or iSCSI layouts, refer to TR-3785 (www.netapp.com/us/library/technical-reports/tr-3785.html). The approach in all cases—and the logic behind them—is similar in most respects.

Here are the guidelines at a high level:

- Guest file system alignment is very important for optimal performance. NetApp TR-3747 provides best practices around file system alignment in virtualized environments (www.netapp.com/us/library/technical-reports/tr-3747.html).
- Create the VMs on an NFS data store using NetApp RCU.
- Host the virtual machine (VM) vswap and temp/page file on a separate NFS data store on a different volume on the NetApp storage system. (Separation of transient data allows faster completion of NetApp Snapshot copies and achieves higher storage efficiency.)
- Locate your application data (databases, logs, and so on) on iSCSI raw device mapping (RDM) LUNs, directly created and connected inside the guest VM using NetApp SnapDrive software (version 6.2 or higher must be installed on guest OS).
- Install application-specific SnapManager software inside the guest VM for consistent backup and granular restore.

This approach is recommended over guest-connected LUNs using the Microsoft iSCSI software initiator because if you want to implement VMware vCenter Site Recovery Manager for disaster recovery—now or at some point in the future—the failover/failback process is much simpler with application data on iSCSI RDMs, and you'll get better support from VMware. You also should put all data stores and RDM LUNs on the same storage system if you are going to use VMware vCenter Site Recovery Manager.

To leverage the benefits associated with SnapDrive and either iSCSI RDMs as recommended above or guest-connected

➤ **Continued on next page**

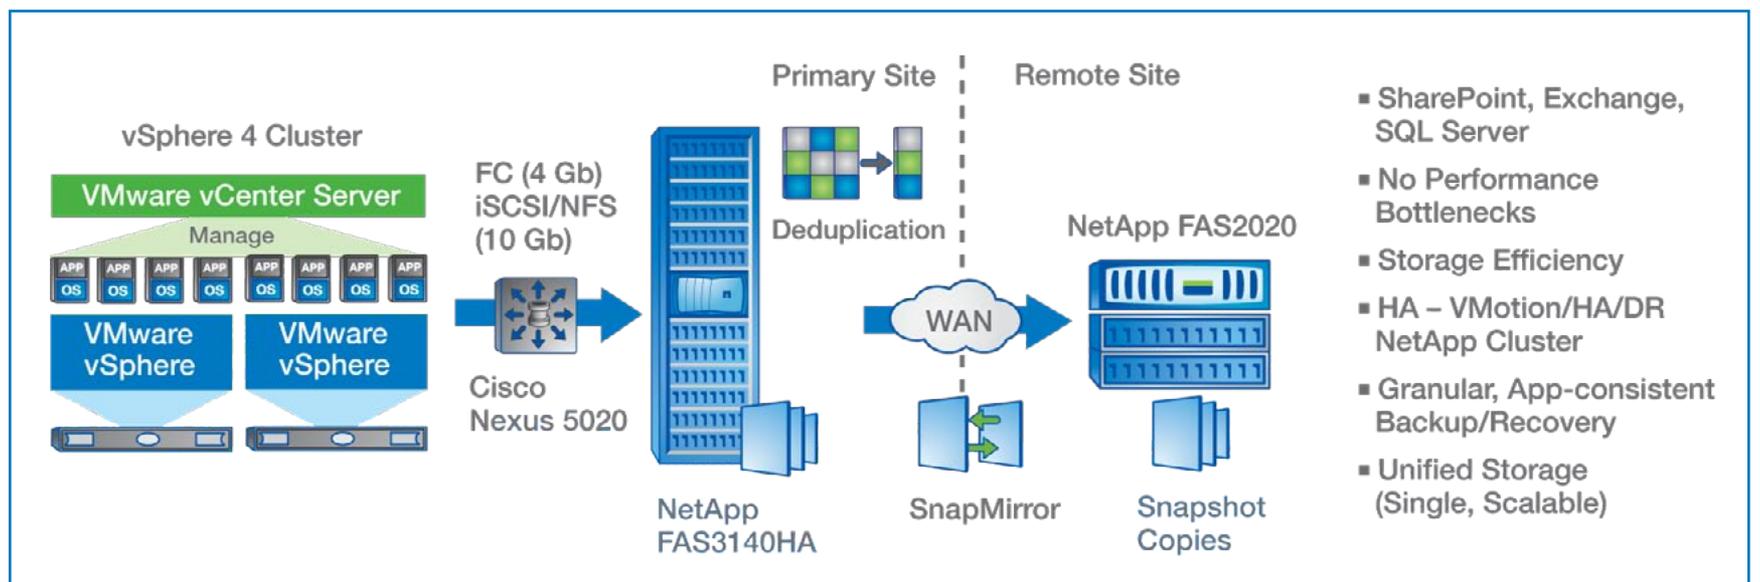


Figure 1) Key elements of the joint NetApp, VMware, and Cisco solution.

➤ Continued from page 13

Virtualizing Microsoft Applications

RDMs using the iSCSI S/W initiator, if you want to use application-specific SnapManager tools for backup of your Exchange, SQL Server, and/or SharePoint data, you must use RDMs (either FC RDMs, iSCSI RDMs as recommended above, or guest-connected LUNs using the Microsoft iSCSI S/W initiator).

If, for some reason, you must configure your environment using VMFS or NFS data stores for application data, your best backup option is SMVI. SMVI is capable of producing consistent backups for all three applications, but with some limitations. Currently, because of limitations in the VMware VSS Requestor (VMware uses copy enumeration for shadow copy), SMVI cannot provide automatic transaction log truncation or backup verification. Both have to be done manually. Also, the VMware VSS Requestor does not currently support application consistency for VMs running Windows® Server 2008. Therefore this solution is limited to scenarios where granular transaction-level restore is not required (for example, point-in-time restore for SQL Servers), manual backup verification can be performed after the backups, and alternate methods of transaction log truncation are possible, for example, with SQL Server databases in simple recovery model (SQL Server provides an automated method for log truncation).

How do I perform application-consistent backup and recovery?

The best way to achieve application-consistent backups for Microsoft applications is to install SnapDrive and the appropriate SnapManager product (SnapManager for Microsoft Exchange, SnapManager for Microsoft SQL

Server, SnapManager for Microsoft SharePoint Server) inside the guest OS for each VM as needed. These tools deliver the specific capabilities to provide application-consistent backups, automated backup verification, and granular restores. For example, SnapManager for Exchange provides single mailbox recovery capabilities. You can learn more about these SnapManager tools in a previous Tech OnTap article (http://partners.netapp.com/go/techontap/matl/app_integration.html).

What's the best way to implement DR?

NetApp SMVI and application-specific SnapManager products can provide replication and disaster recovery for VMs and hosted Microsoft apps. Fully automated disaster recovery can be achieved using VMware vCenter Site Recovery Manager in conjunction with these products. This solution provides complete failover workflow automation for complex environments as described in the Tech OnTap article "Using VMware Site Recovery Manager to Simplify DR" (www.netapp.com/us/communities/tech-ontap/srm-0708.html).

How do I implement multipathing?

If you want your environment to be robust, you must implement multipathing. For an FC-based architecture, I would recommend the Asymmetric Logical Unit Access (ALUA) protocol and round robin (RR) path selection policy. ALUA allows for the autonegotiation of paths between SCSI target devices and target ports, enabling dynamic reconfiguration. ALUA is enabled by default on ESX hosts. On NetApp storage arrays, ALUA should be enabled on the initiator groups, resulting in a more dynamic, or plug-and-play-like, SAN architecture. The RR path selection policy (PSP) provides path redundancy and bandwidth aggregation. Note that there is no need for a device-specific module (DSM) inside the guest VM.

For iSCSI, vSphere introduced support for multiple TCP sessions at the ESX host level for multipathing. You can have two vmkernel

ports and use round robin PSP to achieve plug-and-play multipathing. It provides multiple active paths, and no DSM is required inside the guest VM. Also, both the traditional and multiswitch trunking network designs can be used, as described in TR-3749 (www.netapp.com/us/library/technical-reports/tr-3749.html).

For NFS, multipathing can be achieved for both traditional and cross-stack switches. For details, see NetApp TR-3749 (www.netapp.com/us/library/technical-reports/tr-3749.html).

When using Cisco Nexus 10 Gigabit Ethernet (10GbE), only two 10GbE ports are required on the ESX host. Cisco virtual port channeling (vPC) provides redundancy, fault tolerance, and security.

Are there benefits to using deduplication and thin provisioning?

One of the benefits of this configuration is that no matter which protocol you choose, you can take advantage of NetApp storage efficiency capabilities (FlexClone, deduplication, and thin provisioning) to significantly reduce the amount of storage space you need.

Typical virtual environments have many copies of the same OS and application binaries in different VMs, consuming large amounts of space on expensive shared storage. By using NetApp storage efficiency capabilities, you can achieve more than 50% storage savings on primary storage. Figure 3 illustrates the 92% space savings we achieved while validating the joint solution.

How Do I Size My Environment?

Sizing your environment includes sizing both VMware data stores (containing the guest OS, application binaries, VM page file, and vswap file) and LUNs hosting application databases and logs. NetApp has developed sizing tools to properly size your environment. Your NetApp systems engineer or reseller can help you size your environment based on information gathered from your site:

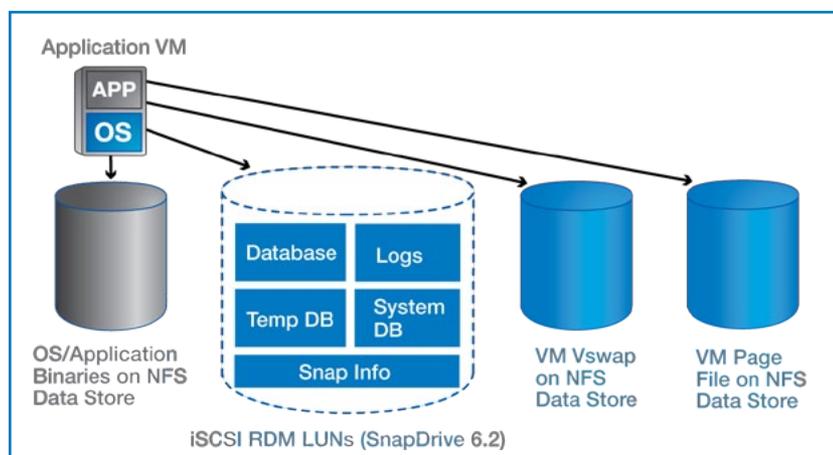


Figure 2) Storage layout using NFS data stores and iSCSI LUNs.

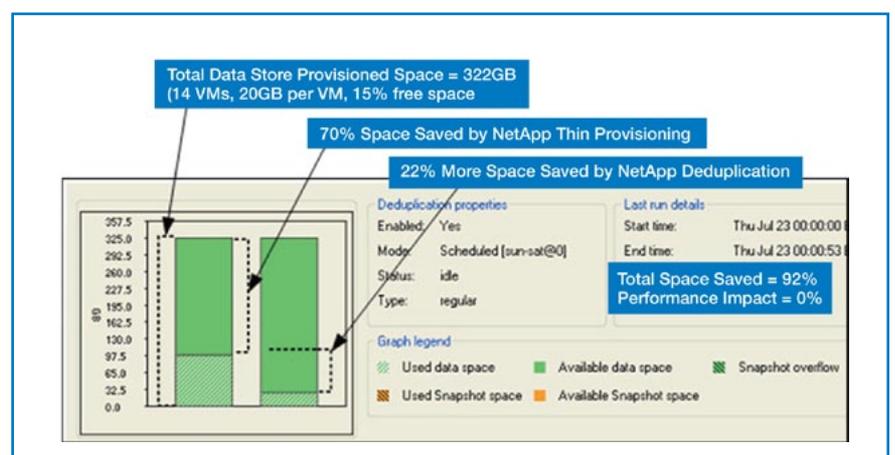


Figure 3) Space savings due to combining NetApp storage efficiency techniques.

- Number of application servers to be virtualized
- Number and type of Microsoft apps
- Capacity requirements for different data components, including expected growth rate
- Performance requirements, including read/write and random/sequential ratio
- For SQL Server databases, number and type of databases (OLTP, DSS, mixed)
- For Exchange server, number and size of mailboxes, user profiles
- For SharePoint server, number of users, space required per user, user concurrency percentage
- Backup/restore/DR requirements

How do I validate the performance of my virtualized microsoft application environment?

You can use the same set of performance validation tools available from Microsoft and third-party vendors that are used in physical environments. These tools can help you determine if performance is within Microsoft guidelines.

To test this joint solution, we used the Microsoft Exchange Load Generation Tool, Microsoft SQLIOSim utility, and AvePoint Sharepoint Test Environment Creator and Usage Simulator to validate performance. Several load tests were conducted for these applications, all running at the same time. Performance validation methods and success criteria for each application are described in TR-3785 (www.netapp.com/us/library/technical-reports/tr-3785.html). Our tests validated that:

- There are no CPU or memory bottlenecks within VMs or on ESX hosts
- No I/O, CPU, or disk bottlenecks exist on storage
- All read and write latencies were well within published Microsoft guidelines
- No network bottlenecks occurred

CONCLUSION

As you march toward your goals of a 100% virtualized data center, I hope the information provided in this article is helpful in understanding the process of virtualizing Microsoft applications. This article only covers the high points of the joint solution for Microsoft application virtualization. You can get all the information you need to deploy this solution in the detailed, 50-page solutions guide (www.netapp.com/us/library/technical-reports/tr-3785.html). The guide covers all the configuration details based on the careful work done by NetApp, VMware, and Cisco. The guide covers FC, iSCSI, and NFS implementations.

In addition to the various links provided in this article, other valuable resources include:

- TR-3749: “NetApp and VMware vSphere Storage Best Practices” (www.netapp.com/us/library/technical-reports/tr-3749.html). Best practices to implement VMware with NetApp storage.
- TR-3767: “Using the Performance Acceleration Module with Exchange 2007” (www.netapp.com/us/library/technical-reports/tr-3767.html). This technical report describes how PAM can boost the number of Exchange users you can support without adding spindles.
- SnapManager guides:
 - TR-3730: “SnapManager 5.0 for Microsoft Exchange Best Practices Guide” (www.netapp.com/us/library/technical-reports/tr-3730.html)
 - TR-3776: “SnapManager for MOSS: Best Practices Guide” (www.netapp.com/us/library/technical-reports/tr-3776.html)
 - TR-3598: “Protecting Exchange Server 2007 with NetApp SnapManager for Exchange” (www.netapp.com/us/library/technical-reports/tr-3598.html)
 - TR-3737: “SnapManager for Virtual Infrastructure Best Practices” (www.netapp.com/us/library/technical-reports/tr-3737.html)
- Site Recovery Manager:
 - TR-3671: “NetApp and VMware vCenter SRM Best Practices” (www.netapp.com/us/library/technical-reports/tr-3671.html) ■



ABHINAV JOSHI
REFERENCE
ARCHITECT
NETAPP

When Abhinav joined NetApp in 2008, he brought over nine years of experience with data center consolidation and virtualization to NetApp. His current responsibilities include developing scalable reference architectures and best practices for securely integrating NetApp virtualized storage and data protection solutions with VMware virtualization technologies and Cisco Unified Computing System and networking technologies—solving customer problems and helping save cost. Since joining NetApp, Abhinav has been an active author, having led and participated in the development of many of the solution guides referenced in this article. ■

➤ Continued from page 9

Deploying FCoE

driving significant capital and operational cost reductions as well as management improvements.

PHASE 2: MAKING THE TRANSITION TO NATIVE DCB/FCOE STORAGE SYSTEMS

Move to an end-to-end DCB/FCoE solution from the host to the network to native DCB/FCoE storage. As shown in Figure 3, the typical configuration has rack servers with CNAs connected to top-of-rack DCB/FCoE switches connected to unified storage that supports FCoE as well as other protocols. FCoE and converged traffic is supported throughout the infrastructure, providing optimal savings.

PHASE 3: MAKING THE TRANSITION TO DCB/FCOE AT THE CORE

After implementing FCoE at the edge or switch, enterprises can migrate to a comprehensive 10GbE-enhanced Ethernet network at the core (illustrated by the green lines in Figure 4) and then gradually move to storage that supports FCoE as well. The end goal is a 10GbE Ethernet infrastructure that supports multiple traffic types (FCoE, iSCSI, NFS, CIFS) from host to fabric to storage sharing the same Ethernet infrastructure.

CONCLUSION

FCoE brings together two leading technologies—the Fibre Channel protocol and an enhanced 10-Gigabit Ethernet physical transport—to provide a compelling option for SAN connectivity and networking. To simplify administration and protect FC SAN investments, FCoE enables you to use the same management tools and techniques you use today for managing both your IP and FC storage networks.

The benefits of converged networks will drive increased adoption of 10GbE in the data center. FCoE will fuel a new wave of data center consolidation as it lowers complexity; increases efficiency; improves utilization; and, ultimately, reduces power, space, and cooling requirements.

If you are planning new data centers or are upgrading your storage networks, you should seriously consider FCoE. By taking a phased approach to consolidating your data centers around Ethernet, you can build out your Ethernet infrastructure over time while protecting your existing fibre channel infrastructure investments. ■

► Continued from page 3

Cisco, VMware, and NetApp Team to Enhance Multi-Tenant Environments

- Cisco MDS 9124, a Fibre Channel switch that provides SAN connectivity to allow SAN boot for VMware ESX running on UCS

Storage layer

The storage layer consists of NetApp unified storage systems capable of simultaneously providing SAN connectivity (for SAN boot) and NFS connectivity for the running VMware environment. NetApp storage can also meet the specialized storage needs of any running application. Running the VMware environment over Ethernet provides a greatly simplified management environment that reduces costs.

NetApp MultiStore software provides a level of security and isolation for shared storage comparable to physically isolated storage arrays. MultiStore lets you create multiple completely isolated logical partitions on a single storage system, so you can share storage without compromising privacy. Individual storage containers can be migrated independently and transparently between storage systems.

Tenant provisioning

When a tenant is provisioned using this architecture, the resulting environment is equipped with:

- One or more virtual machines or vApps
- One or more virtual storage controllers (vFiler® units)
- One or more VLANs to interconnect and access these resources

Together, these entities form a logical partition. The tenant cannot violate the boundaries of this partition. In addition to security, we also want to be sure that activities happening in one tenant partition do not interfere indirectly with activities in another tenant partition.

END-TO-END QOS

Very few projects tackle end-to-end quality of service. In most cases, a QoS mechanism is enabled in one layer in the hope that downstream or upstream layers will also be throttled. Unfortunately, different applications have different characteristics—some might be

compute intensive, some network intensive, and others I/O intensive. Simply limiting I/O does little or nothing to control the CPU utilization of a CPU-intensive application. It's impossible to fully guarantee QoS without appropriate mechanisms at all three layers. Our team set out to design such a system.

Companies such as Amazon, Google, and others have built multi-tenant or “cloud” offerings using proprietary software that took years and hundreds of developers to create in house. Our approach was to use commercially available technology from Cisco, NetApp, and VMware to achieve similar results.

One design principle we applied in all layers is that when resources are not being utilized, high-value applications should be allowed to utilize those available resources if desired. This can allow an application to respond to an unforeseen circumstance. However, when contention occurs, all tenants must be guaranteed the level of service for which they have contracted.

Another design principle is to set the class of service as close to the application as possible, map that value into a policy definition, and make sure that the policy is applied uniformly across all layers in accordance with the unique qualities of each layer. We used three mechanisms in each layer to help deliver QoS, as shown in Table 1.

Table 1) QoS mechanisms.

Compute	Network	Storage
Expandable Reservations	QoS – Queuing	FlexShare®
Dynamic Resource Scheduler	QoS – Bandwidth Control	Storage Reservations
UCS QoS System Classes for Resource Reservation and Limit	QoS – Rate Limiting	Thin Provisioning

Compute layer

At the server-virtualization level, VMware vSphere provides many capabilities to make sure of fair use, especially of CPU and memory resources. A vSphere resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory. By correctly configuring resource pool attributes for reservations, limits, shares, and expandable reservations, you can achieve very fine-grained control and grant priority to one tenant over another in situations in which resources are in contention.

VMware Distributed Resource Scheduler (DRS) allows you to create clusters containing multiple VMware servers. It continuously monitors utilization across resource pools and intelligently allocates available resources among virtual machines. DRS can be fully automated at the cluster level so infrastructure and tenant virtual machine loads are evenly load balanced across all of the ESX servers in a cluster.

At the hardware level, Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS system. This industry-standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. System classes determine how the DCE bandwidth in these virtual lanes is allocated across the entire Cisco UCS system. Each system class reserves a specific segment of the bandwidth for a specific type of traffic. This provides a level of traffic management, even in an oversubscribed system.

Network layer

At the network level, traffic is segmented according to the class of service (CoS) already assigned by the Cisco Nexus 1000v and honored or policed by the UCS system. There are two distinct methods to provide steady-state performance protection:

- **Queuing** allows networking devices to schedule packet delivery based on classification criteria. The end effect of the ability to differentiate which packets should be preferentially delivered is providing differentiation in terms of response time for important applications when oversubscription occurs. Queuing only occurs when assigned bandwidth is fully utilized by all service classes.
- **Bandwidth control** allows network devices an appropriate amount of buffers per queue such that certain classes of traffic do not overutilize bandwidth. This allows other queues to have a fair chance to serve the needs of the rest of the classes. Bandwidth control goes hand in hand with queuing, since queuing determines which packets are delivered first, while bandwidth determines how much data can be sent per queue.

A set of policy controls can be enabled such that any unpredictable change in traffic pattern can be treated either softly, by allowing applications to burst/violate for some time above the service commitment, or by a hard policy, dropping the excess or capping the rate of transmission. This capability can also be used to define a service level such that noncritical services can be kept at a certain traffic level or the lowest service-level

traffic can be capped such that it cannot affect higher-end tenant services.

Policing as well as rate limiting is used to define such protection levels. These tools are applied as close to the edge of the network as possible to stop the traffic from entering the network. In this design, the Cisco Nexus 1000V is used for the policing and rate-limiting function for three types of traffic:

- **VMotion.** VMware traditionally recommends a dedicated Gigabit interface for VMotion traffic. In our design the VMotion traffic has been dedicated with a nonroutable VMkernel port. The traffic for VMotion from each blade server is kept at 1Gbps to reflect the traditional environment. This limit can be either raised or lowered based on requirements, but should not be configured such that the resulting traffic rate affects more critical traffic.
- **Differentiated transactional and storage services.** In a multi-tenant design, various methods are employed to generate differentiated services. For example, a “priority” queue is used for the most critical services, and “no-drop” is used for traffic that cannot be dropped but can sustain some delay. Rate limiting is used for fixed-rate services, in which each application class or service is capped at a certain level.
- **Management.** The management VLAN is enabled with rate limiting to cap the traffic at 1Gbps.

Storage layer

As described above, NetApp MultiStore software provides secure isolation for multi-tenant environments. MultiStore is described in more detail in a Tech OnTap article (www.netapp.com/us/communities/tech-ontap/tot-secure-mobile-cloud-storage-1001.html).

In the storage layer, delivering QoS is a function of controlling storage system cache and CPU utilization as well as making sure that workloads are spread across an adequate number of spindles. NetApp developed FlexShare to control workload prioritization. FlexShare allows you to tune three independent parameters for each storage volume or each vFiler unit in a MultiStore configuration so you can prioritize one tenant partition over another. FlexShare is described in more detail in more detail in an online Tech OnTap article (<http://partners.netapp.com/go/techontap/flexshare.html>). Both MultiStore and FlexShare have been available for the NetApp Data ONTAP operating environment for many years.

NetApp thin provisioning provides tenants with a level of “storage on demand.” Raw capacity is treated as a shared resource and

is only consumed as needed. When deploying thin-provisioned resources in a multi-tenant configuration, you should set the policies to volume autogrow, Snapshot™ autodelete, and fractional reserve. Volume autogrow allows a volume to grow in defined increments up to a predefined threshold. Snapshot autodelete is an automated method for deleting the oldest Snapshot copies when a volume is nearly full. Fractional reserve allows the percentage of space reservation to be modified based on the importance of the associated data.

When using these features concurrently, important tenants can be given priority to grow a volume as needed with space reserved from the shared pool. Conversely, lower level tenants require additional administrator intervention to accommodate requests for additional storage.

CONCLUSION

Cisco, VMware, and NetApp have teamed to define and test a secure, multi-tenant cloud architecture capable of delivering not just the necessary security, but also quality of service, availability, and advanced management.

This article introduced our end-to-end approach to QoS. You can read more about QoS or the other pillars of multi-tenancy in our recently released design guide (www.netapp.com/us/media/cisco-validated-design.html), which describes the elements of the architecture in detail along with recommendations for correct configuration. ■

“One of the biggest obstacles in deploying cloud infrastructures is making sure of the same level of security and isolation in shared virtual environments as you have in separate physical systems. Cisco, NetApp, and VMware have been collaborating to solve that challenge. Our new secure multi-tenancy solution, which is a Cisco Validated Design (CVD), uses virtualization across all layers of the architecture to provide the industry’s only architecture that lets you securely share IT resources in multi-tenant environments across all protocols.”

Parag Patel

VP of ISV and Storage Ecosystem Alliances
VMware

➤ Continued from page 11

Quantum Leap

laid out in TR-3428 (www.netapp.com/us/library/technical-reports/tr-3428.html). VSC includes the following capabilities:

- View status of storage controllers
- View status of physical hosts, including versions and overall status
- Check for the proper configuration of ESX settings as it applies to:
 - Multipathing settings
 - NFS timeouts
 - HBA driver timeouts
- Set the appropriate timeouts on multiple ESX hosts simultaneously with a single mouse click
- Gain access to mbrtools (mbrscan and mbralign) to identify and correct partition alignment issues
- Set credentials to access storage controllers
- Collect diagnostics from the ESX hosts, FC switches, and storage controllers
- View capacity reports on:
 - Data stores
 - LUNs
 - Volume
 - Aggregates
- View deduplication reports on:
 - Deduplication state
 - Deduplication status
 - Space savings
- View LUN status on:
 - Pathname
 - NAA ID
 - ALUA: enabled/disabled
 - Online/offline
 - Protocol

For more details, see the VSC demo (www.youtube.com/netapptv#play/user/FB8CE7ACF0FE6DEA/0/gYb68bVxmWg).

THE BUZZ ABOUT RCU

Since NetApp demonstrated these technologies at VMworld 2009, they’ve generated a lot of customer interest. RCU’s provisioning, deduplication management, and multiprotocol capabilities have been particularly popular. I recently demonstrated RCU provisioning for one customer team, and they were knocked out of their chairs.

Single file restore is perhaps the most popular new feature of SMVI. Most VMware managers can see immediate uses for that capability. For VSC, the ability for VM administrators to see the benefits of deduplication, even in block environments, is an important new feature, and the ability to configure storage according to NetApp best practices with just a few mouse clicks is seen as a big advantage. Taken together, these tools take the complexity out of managing and protecting your virtual storage environment. ■

› Continued from page 10

NetApp and Cisco

Provisioning Manager speeds the creation of new storage resources, helps improve capacity management, and enables policy-based automation to create repeatable, automated provisioning processes.

ZERO COST PROVISIONING

Cisco UCS can boot directly from NetApp storage and utilize a single storage system for both virtual environments and application data. With NetApp FlexClone, you can create operating system and application “golden images” and clone these for new deployments on the Cisco UCS platform. These clones do not require any additional space until they are modified, and only the modifications consume additional storage space. Cisco UCS profiles coupled with the instant nature of FlexClone volumes allow you to deploy fully validated application environments in minutes, without wasting any precious compute or storage resources. The fundamentals of this approach were described in a previous Tech OnTap article (<http://www.netapp.com/us/communities/tech-ontap/tot-vdi-0309.html>).

LEVERAGING NETAPP STORAGE EFFICIENCY, DATA PROTECTION, AND AVAILABILITY

The Cisco and NetApp solution is rounded out through the ability to leverage well-proven NetApp solutions for storage efficiency such as NetApp deduplication (www.netapp.com/us/communities/tech-ontap/tot-dedupe-unstructure-0409.html) now deployed on over 37,000 storage systems (www.netapp.com/us/company/news/news-rel-20090513.html), thin provisioning (www.netapp.com/us/communities/tech-ontap/tot-sddpc.html), NetApp Snapshot copies (http://partners.netapp.com/go/techontap/mat/vmware_snapshots.html), and FlexClone (www.netapp.com/us/communities/tech-ontap/tot-sddpc.html), as discussed above.

Data protection and availability options include Snapshot for instantaneous point-in-time copies, NetApp SnapVault for disk-based backup, and SnapMirror for remote replication.

FIND OUT MORE

If you want to learn more about the NetApp and Cisco virtualized dynamic data center—including benefits, user cases, and a detailed technical discussion of the building blocks—check out the full white paper (<http://media.netapp.com/documents/cisco-netapp-simplified-data-center-architecture.pdf>). ■

› Continued from page 7

Innovate and Thrive in a Difficult Economy

I asked Jessica Yu, NetApp IT manager for storage, to describe the advantages based on her direct experience.

Unlike dedicated environments, virtualized environments leverage all the resources that are available in the technological landscape. The concerns are usually around performance, storage capacity, scalability, fault tolerance, complexity in application dependency, and how to keep track of applications versus servers versus storage and information. Benefits include:

- **Capacity savings.** Flexibility in resources allocation, ability to leverage resources throughout the organization, and easier capacity planning all contribute to significant reductions in capacity requirements.
- **Ease of administration.** Fewer physical machines to manage and maintain also simplify the overall environment and lower maintenance costs. A simplified architecture leads to standardization and more and better automation.
- **Fault tolerance.** The shared storage environment has built-in redundancy when one of the hardware components fails.
- **Increased storage efficiency.** Virtualization of our storage environment makes it possible to utilize thin provisioning for all applications. Standardization also results in more duplicate data being colocated on the same storage, where it can be deduplicated to yield additional space savings.

With shared storage, everything becomes standardized. Therefore, it is much easier and faster to plan, procure, and implement projects requiring storage. We used to have multiple silos of storage. However, using this methodology, data center planning and growth become a nightmare, wasting both data center resources and administrator hours:

- We used to spend weeks talking about adding additional capacity; now we are able to handle capacity planning in a regular monthly meeting.
- Where we used to spend days discussing storage design with project teams, we now spend just a few hours to understand the requirements and review the design template.
- It used to take several hours to locate free capacity across all our storage systems to allocate space for new application implementations. Now we only have to

spend a few minutes looking at capacity availability to identify appropriate resources.

- Our virtualized environment makes tasks easier, so a task that would previously have been handled by level 3 support staff can now be delegated to level 2.
- By implementing SnapManager for Virtual Infrastructure (SMVI) for our VMware environment, we have been able to delegate many tasks, such as creating Snapshot copies and restoring files to system administrators instead of having storage administrators involved in every single task.

All these time savings add up to weeks of hours saved across multiple functional groups, and at the same time we are serving our user community better.

CONCLUSION

Data centers will need to become much more efficient to continue to meet business needs in the face of flat or decreasing IT budgets and increasingly power-hungry IT equipment. I've outlined some of the steps you can take to boost the power efficiency of data center infrastructure and maximize utilization of IT resources, as well as the implications of taking these steps.

A variety of additional initiatives are on the horizon that will also seriously affect IT. To learn more about these topics, see “Storage Infrastructure for the Cloud” on page 1 and “NetApp and Cisco: Virtualizing the Data Center” on page 10. ■

PARTNERSHIP VIDEOS

Cisco and NetApp are working together to deliver Cisco's Unified Compute System in conjunction with the NetApp unified storage architecture to virtualize the data center from end to end. Two recent six-minute videos explore the relationship in more detail:

- Cisco VP and CTO Ed Bugnion and NetApp VP Rich Clifton discuss the UCS architecture and joint product activities (www.youtube.com/watch?v=33x29xP6A1M).
- NetApp CMO Jay Kidd and Cisco VP Paul McNab describe how the companies deliver joint solutions (www.youtube.com/watch?v=C69I1vPBiBs&feature=related). ■

➤ Continued from page 5

Cloud Storage Infrastructure

Integrated data protection and DR

NetApp provides fully integrated and comprehensive data protection and disaster recovery based on the technologies discussed above, including Snapshot, SnapVault, SnapMirror, the SnapManager suite, and Protection Manager. These technologies offload the burden of data protection from servers to storage and provide a consistent approach to data protection across all your application services.

Unified fabric

The emergence of FCoE makes it possible to move your data centers toward a single Ethernet fabric for all your storage and networking needs. In reality, though, you'll probably be in transition for some time, supporting native Fibre Channel, FCoE, and Ethernet storage.

NetApp has been an Ethernet storage leader—first as a NAS pioneer and then as an early proponent of iSCSI. NetApp is the only vendor with native support for FCoE (www.netapp.com/us/communities/tech-ontap/tot-fcoe.html), which is a logical

progression for our unified storage approach that can simultaneously deliver Fibre Channel, iSCSI, and NAS on all our platforms. By supporting FCoE, we provide a simple evolutionary path for Fibre Channel SAN users to migrate to a unified fabric.

Server virtualization support

The items discussed in this section up to this point combine to make NetApp storage a compelling option in virtual server environments. NetApp adds to that close integration with the full range of virtualization solutions, including VMware (www.netapp.com/us/communities/tech-ontap/tot-vstorage-0309.html), Citrix XenServer (www.netapp.com/us/communities/tech-ontap/msa-citrix.html), and Microsoft Hyper-V™ (www.netapp.com/us/communities/tech-ontap/hyperv.html). You can use either NetApp deduplication or NetApp FlexClone technology to eliminate the storage burden that results from storing dozens of copies of the same operating software (www.netapp.com/us/communities/tech-ontap/tot-vstorage-0309.html).

CONCLUSION

Cloud computing is real, and it's happening now. Smart companies are already taking advantage of cloud services to meet many noncore IT functions. They are also evolving their internal IT infrastructures to become more cloud like and to focus on service

delivery to increase efficiency and flexibility while cutting costs.

Storage capable of meeting the needs of cloud computing infrastructures presents a number of unique requirements. To create an effective cloud infrastructure now, you should focus on storage scaling, automation, data movement, multi-tenancy, space efficiency, and support for virtualization. ■

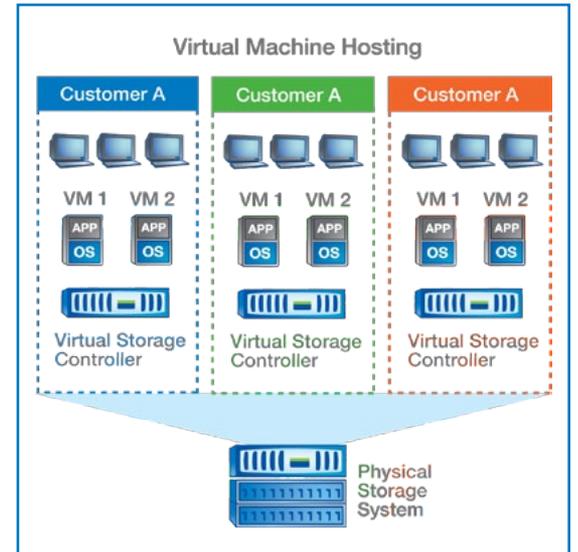


Figure 3) Multi-tenancy with NetApp MultiStore. Multiple customers are each allocated a “virtual storage controller” on a single physical storage system. Each virtual controller can be used to store and manage multiple virtual machines or to meet other storage needs just as if it were a physical storage system.

Table 1) NetApp efficiency technologies.

TECHNOLOGY	BENEFITS
Double-Parity RAID	The double-parity protection of RAID-DP® protects you in the event that two disks fail at once, saving 46% versus data mirroring.
NetApp Deduplication	NetApp deduplication (www.netapp.com/us/communities/tech-ontap/tot-dedupe-unstructure-0409.html) identifies and eliminates redundancy at the block level. Space savings range from 25% to 55% for most data sets, up to 95% for full backups stored on disk, from 70% to 95% for virtual server and desktop environments, and up to 70% for engineering environments.
Thin Provisioning	With NetApp thin provisioning, storage is treated as a shared resource, and capacity is consumed only as it is needed. Thin provisioning can reduce your storage capacity requirement by 20% to 30%. NetApp has been working closely with VMware to provide enhanced thin provisioning functionality for virtual server environments.
Snapshot Technology	NetApp provides space-efficient, nondisruptive Snapshot™ technology. NetApp Snapshot technology delivers up to 80% space savings over competing products.
Thin Replication	Replication is an effective way to make sure of business continuity. NetApp SnapMirror and SnapVault software perform only incremental block transfers—thin transfers—after an initial baseline, saving bandwidth and reducing the storage required for disk-based backups. If your source volume is deduplicated, the target volume will inherit the same space savings, and source and target storage systems need not be the same model and configuration.
Cloning	NetApp FlexClone technology lets you create a “virtual copy” of a data set in seconds and only consumes additional storage space as changes are made to the clone making it an ideal technology for virtual environments that maintain many identical copies of the same operating systems. Space savings can be as high as 80%.

ABOUT NETAPP

NetApp creates innovative storage and data management solutions that accelerate business breakthroughs and deliver outstanding cost efficiency. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, FlexShare, FlexVol, MultiStore, NOW, RAID-DP, SANscreen, SnapDrive, SnapManager, SnapMirror, Snapshot, SnapVault, Tech OnTap, and vFiler are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, Windows, SharePoint, and SQL Server are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Oracle is a registered trademark of Oracle Corporation. Veritas is a trademark of Symantec Corporation. VMware and VMotion are registered trademarks and vCenter and vSphere are trademarks of VMware, Inc. Cisco, Cisco Nexus, and Cisco Unified Computing System are trademarks or registered trademarks of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. The word “partner” or “partnership” does not denote a legal partnership between NetApp Inc. and any other company. NA-092-0810



Lower total cost

Double utilization

Back up frequently

Cut power consumption

Shrink your footprint

AND

AND

AND

AND

AND

AND

AND

AND

boost productivity.

enhance SLAs.

keep applications running.

manage data growth.

grow your business.

on.

on.

on.

Thanks to AND, you'll never have to settle for OR again.

With solutions from NetApp, you can both increase storage efficiency and improve business performance. To get storage efficiency without compromise, visit netapp.com/and.

