NetApp
Go further, faster

NETAPP ARCHITECTURAL PLANS FOR ENABLING A DYNAMIC DATA CENTER FOR TODAY AND IN THE FUTURE

As you strive to keep pace with ever-expanding data growth, your enterprise inevitably suffers higher costs and lower productivity due to an inflexible IT architecture. In today's resource-constrained world, success calls for streamlining the inefficiencies of your storage infrastructure and processes. Your underlying technology prohibits your IT organization from effectively managing costs and responding quickly to the demands of business.

This document outlines NetApp architectural plans for enabling a dynamic data center for today and in the future. The first paper, "Unified Storage Architecture Enabling today's dynamic data center", describes the elements of the NetApp® unified storage architecture in detail. This includes true multiprotocol support, a single management interface, integrated data protection, support for multiple tiers of storage (primary, secondary, and archive/compliance), quality of service, and the ability to act as a front end for legacy storage systems.

The second paper, "Scale-Out Storage and the Dynamic Data Center of the Future", describes the requirements for scale-out storage in the enterprise in terms of resiliency, scalability, manageability, and other factors. It also discusses the current state of NetApp® scale-out deployments as well as specific operational and business advantages that the NetApp approach delivers to the enterprise data center.

The final paper, "Building a Dynamic Data Center", describes the steps required to begin building your next-generation data center. It describes the elements of this new dynamic data center infrastructure as well as specific cost, time-to-market, and risk mitigation advantages that the NetApp approach delivers to the enterprise data center.

.

NETAPP WHITE PAPER

Unified Storage Architecture
Enabling today's dynamic data center

Paul Feresten and Ravi Parthasarathy, NetApp
October 2008 | WP-7054-1008

EXECUTIVE SUMMARY

The traditional model for enterprise storage requires a different storage system for each storage function. One storage architecture might be deployed for primary network-attached storage (NAS), another for storage area networks (SANs), with additional platforms for secondary storage, archive, and compliance. Because of the obvious complexity and cost of such an approach, many vendors in the storage industry have begun to talk about "unified" storage, co-opting the term from NetApp, which pioneered a unified storage architecture years ago. The solutions these vendors offer typically include the ability to accommodate both NAS and SAN protocols, but do so by combining NAS and SAN components with different architectures, management infrastructures, backup requirements, and so on, and rarely offer the scalability and performance necessary in today's enterprise environments.

This paper describes the elements of the NetApp® unified storage architecture in detail, including true multiprotocol support, a single management interface, integrated data protection, support for multiple tiers of storage (primary, secondary, and archive/compliance), quality of service, and the ability to act as a front end for legacy storage systems. NetApp is able to combine these features and more into a single platform capable of meeting your end-to-end storage needs, while demonstrating significant performance and cost-of-ownership advantages, setting the standard by which any storage solution that claims to be unified should be measured.

TABLE OF CONTENTS

1 INTRODUCTION

The enterprise data center is undergoing a radical transformation; server virtualization technology is changing the way both servers and applications are provisioned, while altering the workload on storage systems and increasing storage demand.

The trend toward server virtualization and consolidation is causing a transformation in the way that data centers are being designed, built, and managed. Data storage and data management are key elements of this transformation, because successful virtualization deployments depend upon a shared, networked storage infrastructure capable of eliminating the "silos" of storage associated with various application tiers. This data center transformation makes the idea of "unified storage" increasingly appealing. Unified storage is an important tool that can help address many existing data center challenges, including helping you get power, cooling, and space utilization under control. If you are in the midst of redesigning your data center or establishing a new infrastructure model, unified storage should be a key element of that effort.

But what is unified storage, and what are the requirements for a unified storage solution? Although the marketing literature for many solutions may claim a solution is unified, the reality is often something quite different.

NetApp believes that a true unified storage solution should be much more than a loosely coupled collection of components. In this paper, we'll examine a fundamentally simpler approach to storage, discuss unified storage building blocks, debunk some myths about unified storage performance, illustrate the solution's cost advantages, and describe what we see as the next steps in the evolution of NetApp unified storage.

DEFINITIONS

Let's begin with a few definitions so there is no confusion about what we mean by important terms.

- **Unified storage architecture.** A unified storage architecture creates a single, end-to-end foundation for dynamic data management. It goes beyond simple multiprotocol storage to provide integrated data management and data protection, support for all tiers of storage, quality of service, and other elements, all in a single platform. A unified storage architecture enables the creation of a common storage pool that can be networked and shared across a diverse set of applications with a common set of management processes. The ability to scale from small to very large storage capacity without compromising application performance or service levels is a key attribute.
- **Multiprotocol storage.** A multiprotocol storage system supports both block-based and file-based access in a single storage platform, incorporating Fibre Channel (FCP) and iSCSI access for blocks and NFS and CIFS access for files. Applications can utilize two or more of these protocols simultaneously. The ability to run Fibre Channel protocol over Ethernet (FCoE) is emerging as a key ingredient of modern multiprotocol infrastructures.
- **Dynamic data center.** The dynamic data center is one that leverages a utility computing model, centralized resource management, and rapid and flexible resource allocation to support a wide range of data center applications. This includes the ability to rapidly commission and decommission applications, provision (and reprovision) required resources, and nondisruptively migrate applications and data between resources to meet time-varying application service level requirements. It also includes the ability to seamlessly expand the underlying infrastructure and/or retire older components while maintaining continuous application operations through nondisruptive migration.
- **Unified fabric.** A unified fabric provides a high-bandwidth, single-wire approach in which all of a data center's I/O activity can be consolidated, regardless of underlying network protocols and data formats. From a server and storage perspective, this will reduce the required number of host bus adapters, NICs, switches, and cables. A unified fabric will allow increased data center scalability with improved performance and fault tolerance. The first step toward achieving a unified fabric is having the ability to run Fibre Channel over Ethernet.

ENTERPRISE CHALLENGES AND OBJECTIVES

When we look at typical enterprise challenges and objectives, the need for a unified storage architecture starts to come into focus. Many data centers have grown organically over time. New storage systems have been dedicated to each new application, often without enough consideration for what existed before. In many cases, this situation has been complicated by mergers and acquisitions that bring together disparate infrastructures.

Enterprise data centers often provide tiers of storage targeted to meet the needs of different applications. (Note that this type of storage tiering doesn't necessarily map directly to the technology tiers that storage vendors talk about—tier 1: high-speed FC disk, tier 2: SATA disk, tier 3: online archive, tier 4: tape). The highest storage tier in an enterprise data center offers the most robust features to meet the requirements of mission-critical applications. Lower tiers typically have decreasing levels of performance and availability (and cost) suited to different application categories. These enterprise tiers may be created using equipment from different vendors with different storage architectures.

The resulting silos of storage have increased overall cost because of complex management, low storage utilization, and direct data center costs for power, space, and cooling. Backup and recovery processes are complicated, and the time needed to accomplish them has increased to an unacceptable level. Implementing a consistent DR strategy is difficult or impossible, and the need to tailor each DR solution for each storage silo makes the cost and complexity prohibitive.

Table 1) Enterprise challenges that a unified storage architecture can help address.

Enterprise Challenge	How a Unified Storage Architecture Helps
Emergence of server virtualization makes backup, recovery, and DR more challenging.	Data protection is fully integrated at the storage level. Backup and recovery are designed into the underlying data structure and consistent across multiple classes of data.
Increased downtime and costs result from ad hoc processes across heterogeneous infrastructure.	Consistent processes enable faster response and fewer errors while enabling policy-based automation.
Increase the productivity of people, systems, and infrastructure across multiple applications.	Provides a single framework for data management with a common set of processes and training requirements.
Existing technology is difficult to repurpose and scale.	One storage platform adapts to the full range of data types, access methods, and capacities required across a diverse spectrum of enterprise applications.
Storage utilization is low resulting in poor data center efficiency.	Storage is pooled, allowing sharing of infrastructure and capacity with on-demand provisioning. Space saving technologies such as deduplication and thin provisioning reduce space, power, and cooling costs.
Infrastructure inherited through M&A and other means is difficult to integrate in a rational manner.	Common software architecture can accommodate and add value to storage devices/arrays from all major vendors.

THE NEED TO UNIFY STORAGE

In order to address these challenges and reduce storage acquisition and operational costs, enterprises are looking for ways they can simplify and unify storage infrastructure across both storage protocols and storage tiers. The goals of those seeking to unify storage include:

- Improve service levels.
- Reduce costs of redundancy and minimize the number of point solutions.
- Move toward an end-to-end data protection plan (backup/recovery, D/R, archiving) with a single management interface.
- Implement a recovery-centric approach to storage that minimizes downtime after an outage.
- Provide multisite infrastructure that protects as many applications as possible (cheaper network bandwidth contributes to making this practical).
- Add the ability to delegate more functions to application and system admins (this requires a common pool of storage coupled with policy-based tools).

When you begin to examine other vendors' "unified" storage offerings in light of these goals, you quickly realize that not all unified storage is created equal.

OTHER "UNIFIED" SOLUTIONS

For some storage vendors, unified storage just means the ability to support both NAS and SAN with a storage solution bundled under a single product family name. Vendors may make claims of multiprotocol support regardless of how integrated a solution is. Some solutions consist of a bundle of separate SAN and NAS components—such as a NAS gateway configured in front of SAN storage—each with its own user interface and data protection mechanism. At best, the commonality is the physical disks, and even those are dedicated for SAN or NAS without the flexibility you need to quickly reallocate storage for use with different protocols. This may be thought of as a "bolt–together" unified storage solution.

You may also find that these so-called unified storage products are bounded and nonscalable. Such solutions are often targeted at small and medium-sized businesses; they are composed of multiple components wired together in a way that won't scale beyond narrow limits.

If you evaluate such solutions against the enterprise goals for unified storage described in the previous section, you'll find that they do little to satisfy any of the goals. Section 7 of this paper provides five tests you can use to determine whether a particular solution has a truly unified storage architecture.

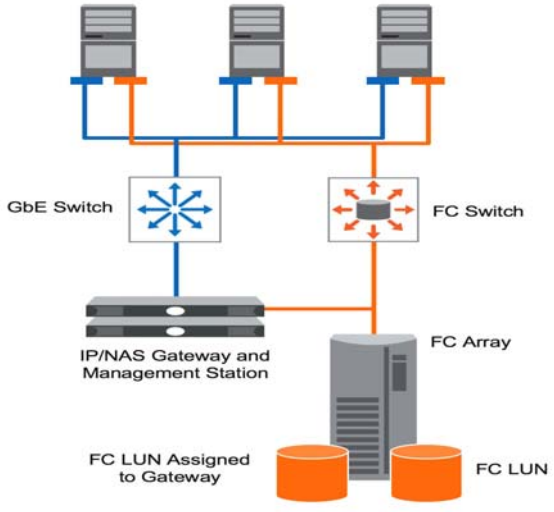

Figure 1) Bolt-together approximation of a unified storage solution.

2 A SIMPLER APPROACH TO UNIFIED STORAGE

In this section we examine the things that make NetApp storage—or any truly unified storage solution—simpler. In addition to having the ability to run multiple storage protocols on the same system, it comes down to how quickly and easily you can accomplish all the tasks associated with storage, from provisioning to data protection, and how broadly those management capabilities can be leveraged.

- **Single provisioning interface.** A single interface is needed to effectively pool storage, creating one model for defining data containers that can be allocated and dynamically managed (sized and resized) for use by a wide range of applications, whether they need block (SAN) or file (NAS) access.
- **Common management framework.** Having a single data model and toolset enables a consistent management framework across many applications and workloads. A set of common management services creates a hierarchy of value, from management of physical storage to application-level integration.
- **Policy-based automation.** Having a single provisioning interface and a common management framework in turn makes possible policy-based automation that allows storage administrators to delegate some or all of the responsibility for provisioning and management tasks. This extends a single toolset across the organization, allowing system administrators, application administrators, DBAs, and others to quickly accomplish storage tasks without storage administrator involvement.
- **Shared data protection at the storage level.** A single data protection architecture encompasses everything from a single file to full disaster recovery—all based on the same basic foundation. Starting with our Snapshot™ copy technology, NetApp has built a consistent set of tools that leverage that underlying capability, extending its use for a wide range of applications both in the data center and at remote offices. The result is near-instant data recovery and the ability to extend a consistent set of data protection tools to files, LUNs, databases, and applications.
- **Ability to incorporate legacy storage.** The reality in almost every data center is that you have substantial investments in storage from a variety of vendors. NetApp can extend many of the unified storage advantages described in preceding bullets to your existing storage. NetApp V-Series systems are able to virtualize storage from EMC, HP, HDS, IBM, and others.

As shown in Figure 2, this approach depends upon the ability to unify multiple technologies, protocols, and classes of data on top of a single software foundation, enabling much greater consolidation of storage infrastructure, including storage for primary, secondary, and archive data.

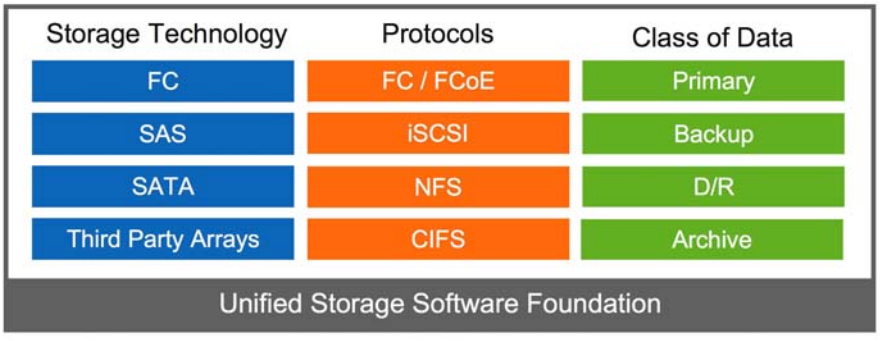

Figure 2) Integration of multiple technologies and classes of data.

3 NETAPP UNIFIED STORAGE ARCHITECTURE: BUILDING BLOCKS

This section describes the key technologies that make NetApp unified storage systems uniquely suited for enterprise storage needs. The combination of these building blocks, incorporating key technologies for thin provisioning, space-efficient scalable Snapshot copies, and deduplication, makes NetApp unified storage a fundamentally efficient architecture.

SINGLE OPERATING ENVIRONMENT

The NetApp Data ONTAP® operating environment is the software foundation that underlies every NetApp storage system and is therefore the core of the NetApp unified storage architecture. Data ONTAP runs on storage systems ranging from a few disk drives to over a petabyte of storage.

The result of years of continuous development, Data ONTAP is both stable and mature, and it contains a rich set of integrated features that other storage vendors usually charge for if they offer them at all:

- Space-efficient, scalable Snapshot technology, with negligible performance impact, which yields dramatic reductions in backup and recovery time and is the foundation of application-centric data protection
- Onboard deduplication of both primary and secondary data
- Dual-parity RAID that protects against double disk failures without significant performance impact; other vendors suggest using mirroring (doubling storage costs) for write-intensive workloads
- Thin provisioning, which allows multiple applications to share a single pool of on-demand storage so you don't have to provision more storage for one application while another application still has plenty of allocated but unused storage
- Quality of service (QoS) so that high- and low-priority workloads can share the same storage without impacting critical jobs

With a single software foundation, it is possible to meet the storage needs of the entire data center with one basic approach, which translates to less training, more staff productivity, and dramatically improved business agility.

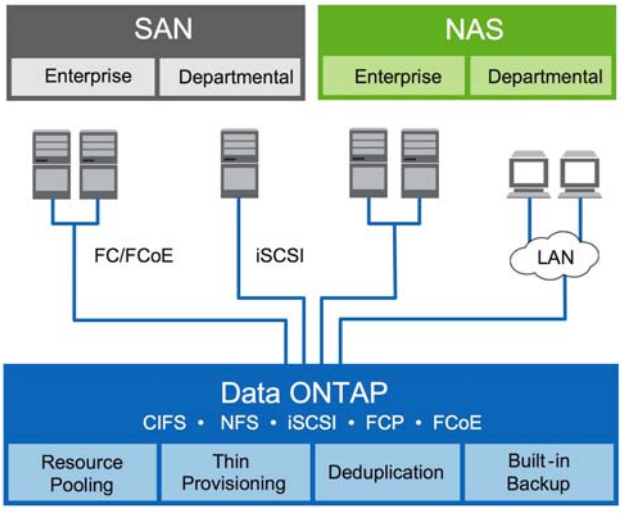

Figure 3) Unified storage foundation—Data ONTAP.

INTEGRATED, MULTIPROTOCOL ACCESS

While other multiprotocol storage solutions lack the level of integration necessary to make them truly unified, NetApp solutions integrate all the common file- and block-based storage protocols within Data ONTAP. Provisioning and management interfaces are the same, and data protection, management, and application-integration software work as expected, regardless of protocol.

You can easily run both file- and block-based storage protocols on a single storage system and apply core Data ONTAP features such as Snapshot, deduplication, and thin provisioning to all your data, regardless of the protocol you use. Of course, you can also dedicate individual storage systems for use by particular protocols if you prefer. You'll have a common management interface across all systems and the ability to repurpose a particular system and use it with a different storage protocol (or add a storage protocol) should your needs change. The NetApp unified storage architecture gives you flexibility to design the right solution for your needs today, without locking you into a limited solution that won't adapt to future needs.

SINGLE, SCALABLE ARCHITECTURE

Many of the "unified" solutions from other vendors lack the ability to scale. NetApp offers a single storage architecture across storage systems that range from workgroup or departmental solutions to full enterprise systems. That lets you choose the storage systems that are right for your requirements in various locations and to know that all will provide the multiprotocol access you need. A storage administrator who is trained on Data ONTAP will be able to leverage those skills to manage any NetApp storage system you have, whether it's a small departmental system in a remote office or an enterprise-class storage system with hundreds of terabytes of storage in your corporate data center.

NetApp storage systems also are able to scale both capacity and performance. You can start small with a single disk shelf and a few disks, and grow to many terabytes. Your investment in disk storage is protected in the event that you decide to upgrade your storage controller, so you don't have to go through a painful migration process. Other storage vendors often have different—and incompatible—operating environments for their low-end versus mid-range versus enterprise storage, so making the transition from one class of system to the next can require a painful data migration.

MULTIVENDOR ARRAY SUPPORT

Another significant dimension of NetApp unified storage is the ability to consolidate storage arrays from other vendors—in effect making your existing heterogeneous storage arrays part of a unified solution. With NetApp V-Series storage systems, you configure your existing storage arrays up front for use with the V-Series controller. After making that initial configuration, you can handle most of the ongoing management through NetApp management interfaces, effectively masking the heterogeneity of the underlying systems for more efficient management that leverages all the benefits of Data ONTAP, including thin provisioning, deduplication, and Snapshot. Your back-end systems may be Fibre Channel only, but using the NetApp V-Series you can repurpose that storage and use it for Fibre Channel, iSCSI, NAS, or any combination.

Some customers find that they can deploy NetApp in front of their existing storage, use deduplication to recover storage capacity, and actually delay making additional storage purchases.

FLEXIBLE STORAGE TIERS

The NetApp unified storage architecture has the ability to support different tiers of storage with the same architecture and on the same storage system. (In this case, we're talking about technology tiers rather than enterprise-defined tiers.) These include a high-performance tier (Fibre Channel disk); a low-cost, high-capacity tier (SATA disk); and online archive tiers using write once read-many (WORM) technology to provide data permanence (FC or SATA).

In most vendor environments, these tiers are implemented using separate, discrete storage systems—and at this point it should come as no surprise that the different tiers often have different, and incompatible, storage architectures. This once again raises complexity and management overhead and limits flexibility. Having the ability to have all these tiers in a single architecture with consistent management in which you can provision different storage tiers in a single storage system, whenever it makes sense to do so, gives you unprecedented flexibility and allows you to accomplish more with less overhead.

SINGLE, CONSISTENT DATA PROTECTION ARCHITECTURE

Unified storage from other vendors often consists of separate, discrete software stacks for NAS versus SAN. This means you are stuck with a set of different, and largely incompatible, tools for protecting different types of data. You may need two solutions for backup (one for NAS and one for SAN) and two more solutions for disaster recovery.

The NetApp approach to data protection is different. As you've already seen, all NetApp software tools work in essentially the same way regardless of the storage protocol you are using. NetApp data protection tools are built on top of our space-efficient Snapshot technology, which captures a point-in-time image of a data volume and consumes additional storage only as the volume changes. Snapshot copies serve as the foundation for standard tape (and virtual tape) backups. You can also vault your Snapshot copies to secondary storage for longer-term, online retention. Storage on standard Windows®, Linux®, and UNIX® servers can also be included in this vaulting process.

For disaster recovery, NetApp provides efficient asynchronous or synchronous replication. Snapshot once again serves as the foundation for efficient asynchronous data transfers that replicate only changed blocks (rather than whole files) for efficient use of network bandwidth. All NetApp data protection software can be configured, managed, and monitored across all storage systems and all storage protocols from a single management console.

For application backup, NetApp provides the SnapManager® suite of products that integrate with popular applications including Oracle®; SAP®; Microsoft® Exchange, SQL Server™, and SharePoint®; and VMware® VI3. These tools integrate directly with the application, allowing administrators to put the application in hot backup mode, capture a consistent Snapshot copy, and then resume normal operation in seconds—forming a foundation for both backup and disaster recovery. These tools allow storage administrators to delegate specific storage management capabilities directly to application administrators for improved efficiency.

FLEXIBLE, DYNAMIC PROVISIONING

One of the features of Data ONTAP that NetApp users consistently comment on is the ability to nondisruptively grow and shrink data volumes as needs change. For example, you can provision a data volume for use with either NAS or SAN protocols and grow it over time to meet changing needs. This contrasts sharply with the traditional approach to provisioning in which you guess how much storage you're going to need and provision it all upfront. If you guess high, the space is wasted (or painful to recover). If you guess low, growing the LUN or volume may be equally painful.

You also have the option of using thin provisioning, another feature built into Data ONTAP, which allows an underlying pool of free storage to be oversubscribed by multiple storage volumes. New storage capacity can be provisioned on a just-in-time basis rather than leaving large amounts of storage sitting idle. Rather than making capacity planning decisions and provisioning to meet the needs of each individual volume, you plan and provision for the needs of the entire storage system. This is easier, less prone to mistakes, and results in more efficient storage utilization so less storage is needed.

As a result of these flexible provisioning methods, typical NetApp storage systems have utilization rates of 60% or higher. The industry norm for storage utilization averages from 35% to 45%.

CONSOLIDATION OF MULTIPLE WORKLOADS AND QUALITY OF SERVICE

Another important feature of the NetApp unified storage architecture is the ability to consolidate multiple application workloads on a single storage system. Most storage administrators are hesitant to do this with critical applications because they don't want to risk having less essential work impact performance.

Data ONTAP includes quality of service (QoS) capability as a standard feature. This capability works at the volume level; by simply giving higher priority to the volumes used by more critical applications, you can safely support multiple applications with a single, consolidated storage system. A higher priority gives a volume a greater percentage of available resources when a system is fully loaded. If higher-priority applications aren't busy, lower-priority applications can use available resources without limitation.

For example, Leuven University Hospital (UZ Leuven) consolidated all its critical Sybase database storage along with storage used by less critical SQL Server applications on a single set of NetApp storage systems. UZ Leuven uses the built-in Data ONTAP QoS feature, FlexShare™ to prioritize Sybase volumes (and

hence Sybase workloads) over SQL Server volumes, successfully supporting over 120 database instances and 1,500 concurrent Sybase users on the consolidated infrastructure.[1]

4 PERFORMANCE IMPACT

One of the myths that have been perpetuated about the NetApp unified storage architecture is that performance naturally suffers when you combine all these capabilities in one system. Particularly for SAN performance, some storage vendors say that imposing any type of data layout overhead on the data volume reduces performance.

In fact, the reverse is true. The fundamental way that NetApp writes data to a volume, whether NAS or SAN, is what makes possible all the data management features that distinguish NetApp from the competition. NetApp unified storage performs as well or better than competing storage solutions, and benchmark results show that performance, particularly under real-world conditions, is superior.

SAN PERFORMANCE

NetApp recently compared the performance of a NetApp FAS3040 storage system with an EMC CLARiiON CX3-40 on the SPC-1 benchmark. (In the interest of full disclosure, the FAS3040 is a unified storage system capable of running multiple protocols simultaneously, but it was only configured for Fibre Channel for this benchmark. The EMC system is a SAN-only configuration not capable of serving file-based protocols as configured.)

SPC-1 is a standard benchmark that generates a workload with characteristics of typical business applications such as database and e-mail with random I/O, queries, and updates. In independently audited results, the NetApp system delivered higher baseline performance (nearly 20% more I/O operations per second), despite being substantially less expensive than the EMC configuration (see Table 2).

The NetApp performance advantage really becomes apparent, however, when the same test is repeated with snapshots turned on. When a feature becomes a regular part of daily operation across a wide base of users, it makes sense to test the performance impact of that feature. The performance of the NetApp FAS3040 dropped only 3%; it still achieved 97% of its baseline performance (performance level without snapshots). The EMC system didn't fair nearly as well. Performance dropped by 64%; it only delivered 36% of baseline performance with the equivalent snapshot capability enabled.

Table 2) SAN performance with and without snapshots enabled.

Performance/Pricing	NetApp FAS3040	EMC CLARiiON CX3-40
Baseline SPC-1 IOPS	30,985.90	24,997.48
Baseline Price Performance	$13.61/SPC-1 IOPS	$20.72/SPC-1 IOPS
SPC-1 IOPS with Snapshot	29,958.60	8,997.17
Perf Impact of Snapshot	3%	64%
Snapshot Price Performance	$14.89/SPC-1 IOPS	$59.49/SPC-1 IOPS

In a subsequent study of SAN performance, the NetApp FAS3170 achieved a result of 60,515.34 SPC-1 IOPS. (The FAS3100 series is the next-generation follow-on to the FAS3000 series, of which the FAS3040

[1] http://partners.netapp.com/go/techontap/matl/UZLeuven.pdf

is a member.) This result—approximately 2x the performance seen for the NetApp FAS3040 or the EMC CLARiiON CX3-40—corresponds to price/performance of $10.01 per SPC-1 IOPS versus $13.61 per SPC-1 IOPS for the earlier NetApp result and $20.72 per SPC-1 IOPS for the EMC configuration.

NAS PERFORMANCE

A comparison of benchmark results for the NetApp FAS3170 versus the EMC Celerra NS80G illustrates a similar advantage for NetApp in the area of NAS performance. (Once again, the NetApp FAS3170 system, while fully unified storage capable, was only configured for NFS for the purposes of the benchmark. The EMC system is a NAS-only configuration.)

Using the SPEC SFS97_R1.v3 benchmark, the NetApp system posted throughput of 137,306 operations per second (OPS) with an overall response time (ORT) of 0.94 milliseconds, while the EMC system achieved throughput of 86,372 OPS at an ORT of 1.49. The NetApp configuration outperformed the EMC configuration by 59% in throughput with an ORT that was 36% faster. A low response time (ORT) is just as critical as throughput for many applications.

OTHER PERFORMANCE MEASURES

While operations per second and response time are certainly important measures of performance, there are other real-world measures that storage administrators may find at least as compelling. In a recent comparison, VeriTest, an independent testing service now part of Lionbridge, compared provisioning and volume cloning time for a NetApp FAS3070 versus an EMC CLARiiON CX3-80[2]. Following the published best practices for each company, it took 2.5x longer to provision an equivalent volume on the EMC system and 233x longer to clone a 400GB LUN.

5 COST ADVANTAGES OF UNIFIED STORAGE

At this point, you're probably thinking, "There must be a catch." If NetApp unified storage delivers this much value, then it must cost more. The truth is just the opposite. Because of the efficiency of Data ONTAP with built-in features such as Snapshot, instantaneous cloning, deduplication, dual-parity RAID, and thin provisioning, unified storage actually saves you substantial storage costs up front; reduces your costs for space, power, and cooling; and reduces management costs for substantial reductions in total cost of ownership (TCO) for common data center applications.

In typical data center environments, NetApp storage achieves rates of utilization ranging from 65–75% versus 35–45% for traditional storage. All these space savings add up to substantial reductions in the amount of storage needed to achieve a given amount of usable space, as illustrated in Figure 4.

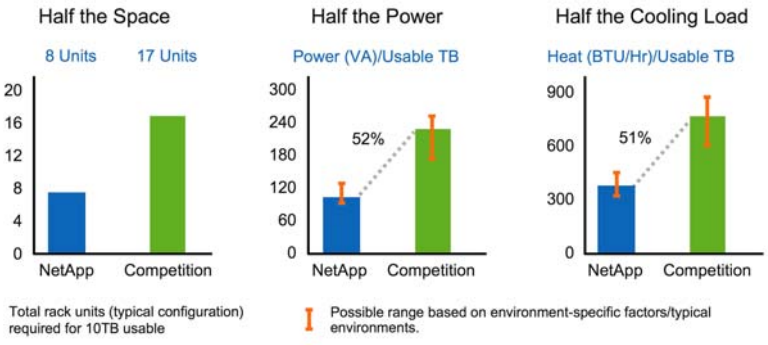

Figure 4) Data ONTAP enables significant cost reduction.

[2] http://www.lionbridge.com/competitive_analysis/reports/netapp/NetApp_FAS3070_vs_EMC_CX3-80_Executive_Summary.pdf

The final measure of efficiency is total cost of ownership in the real world running real applications. Oliver Wyman (formerly Mercer Management Consulting) undertook a series of TCO studies examining the cost of NetApp storage versus competing solutions for popular applications. The results speak for themselves, with TCO ranging from 30% to 55% lower.

Table 3) NetApp total cost of ownership vs. competing solutions (source: Oliver Wyman).

Application	TCO
Oracle	39% lower
SAP	55% lower
Exchange	30% lower
File Services	44% lower
Archive	35% lower
VMware	38% lower

6 THE FUTURE OF NETAPP UNIFIED STORAGE

Right now, there are a number of important trends in the industry as a whole, including growing interest in cloud computing and scale-out storage and the desire for a simplified network infrastructure. Ideally, your storage architecture should evolve to accommodate these trends without requiring disruptive changes to your entire operating environment.

The NetApp unified storage architecture has been proven to evolve to support new capabilities. In fact, since the Data ONTAP operating environment was first introduced in 1993, most of the enhancements that NetApp has added have been evolutionary—the company hasn't required users to throw out what they have and start again. NetApp has a deep commitment to that philosophy, and our unified storage architecture will deliver on the promise of the latest industry trends in a similar fashion.

SCALE OUT

The next step in the evolution of the NetApp unified storage architecture will be the addition of the scale-out storage technologies that NetApp pioneered. Your storage systems will become part of a global namespace, with the ability to transparently spread data across storage systems for greater performance, load balancing, and nondisruptive operation.

With all your storage under a common framework, the ability to scale out will become a logical extension. You'll be able to move to scale-out storage using existing storage; no forklift upgrade will be required. By joining together your existing storage with any future storage system purchases, you'll be able to achieve even greater levels of consolidation while scaling capacity, performance, and data availability far beyond today's limits.

You can find out more about scale-out storage in a companion white paper entitled *Scale-Out Storage and the Dynamic Data Center of the Future.*[3]

[3] http://media.netapp.com/documents/wp-7042.pdf

UNIFIED FABRIC

Another trend that enterprise data centers will be able to leverage in the near future is the unification of storage area networks and local area networks (LANs), by extensions to classic Ethernet. This "unified fabric" is made possible by the Data Center Bridging (DCB) standard. DCB technology allows 10GbE the flexibility to support transmission mechanisms beyond Internet protocol, including Fibre Channel over Ethernet. This is the culmination of a consolidation trend that began with server virtualization and that continues today with storage consolidation and unified storage. Most data centers today have a Fibre Channel network for SAN storage in addition to an Ethernet infrastructure. Maintaining multiple types of network equipment—along with the tools and expertise needed to effectively manage different network technologies—adds significantly to data center cost and complexity.

NetApp has been selling Ethernet-based storage for the past 16 years, with demonstrated leadership in both NAS and iSCSI protocols. Because NetApp has already unified the storage endpoint with support for both Ethernet and Fibre Channel, the company is logically in the best position to drive the conversion to a totally Ethernet-based fabric. Ultimately, server virtualization will unify the server, the intervening network fabric will be unified with FCoE, and NetApp will provide unified storage that can accommodate all back-end storage needs. NetApp has already forged close relationships with leading providers of server virtualization software so that our storage offers maximum benefit with VMware ESX, Microsoft Hyper-V™, Citrix, Oracle VM, and others.

7 FIVE TESTS FOR DETERMINING IF YOU'RE BUYING UNIFIED STORAGE

How do you know if a storage solution that is marketed as "unified storage" really delivers as advertised? Here are five simple tests to help you decide.

1. **Does one storage architecture address *all* your storage needs?/Does all the storage run under a single operating system?**
   Many vendors offer different systems for different needs: primary NAS, primary SAN, secondary storage, archiving, compliance, and so on. This forces you to train people for each system, resulting in inefficiency. This approach also creates islands of storage that make it impossible to provision available storage where it is needed. This increases cost and adds to data center inefficiencies. A true unified storage architecture allows you to train your staff on a single operating environment with a single set of tools.

2. **Are all protocols supported *natively* in the same box?**
   The ability to serve multiple protocols from the same storage system has many advantages. It makes consolidation more effective by combining block and file data on the same system. This helps you freely mix and match workloads to greatly improve utilization. It also helps make your infrastructure more flexible by giving you the right to change your mind. You can choose iSCSI today and change that to Fibre Channel without migrating data or changing your data layout. Native multiprotocol, as opposed to gateway approaches, results in noticeable performance improvements.

3. **Are all your business needs managed by a single storage architecture?**
   Meeting your business requirements may result in multiple copies of data for production, backup, test and development, disaster recovery, archival, and compliance. Having different architectures for each of these business needs adds significant management cost and increases the number of full copies required to support your business. Consolidating these classes of storage to a unified storage architecture with a common set of tools can reduce management overhead and lower training costs.

4. **Can you easily consolidate storage for space, power, and cooling efficiency?**
   Another key advantage of the superior consolidation made possible with a unified storage architecture is reduced space, power, and cooling cost. Effective pooling of storage across all business requirements with advanced features like deduplication, thin provisioning, and virtual copies helps you make the most of your consolidated storage by avoiding redundant copies, thereby reducing your space, power, and cooling requirements.

5. **Are you forced to move to a different platform to meet objectives for performance, availability, or scalability?**
   Life is full of choices, but you shouldn't be forced to choose between performance, availability, and scalability when choosing a storage platform. A unified storage architecture allows you to tailor your solution to cost-effectively meet your performance and availability objectives and scales seamlessly as you grow.

8 CONCLUSION

The traditional model in the storage industry is one with multiple, incompatible storage architectures with incompatible processes requiring lots of experts. Because integration between architectures is expensive, this model results in functional silos of storage in which available storage cannot be easily allocated to the applications that need it.

Recognizing the limitations of this approach, some vendors have begun to offer what they describe as "unified" storage, but these solutions typically fall short in functionality, performance, and scalability. Only NetApp offers a truly unified storage architecture capable of meeting all your storage needs—NAS and SAN: primary, secondary, and archive—from a single platform with a single architecture. With NetApp, you can implement a single set of processes for all data management functions, including backup and disaster recovery, and even delegate important storage tasks to other functions with policy-based management. The result is complete compatibility that allows you to do more and manage more data with fewer resources.

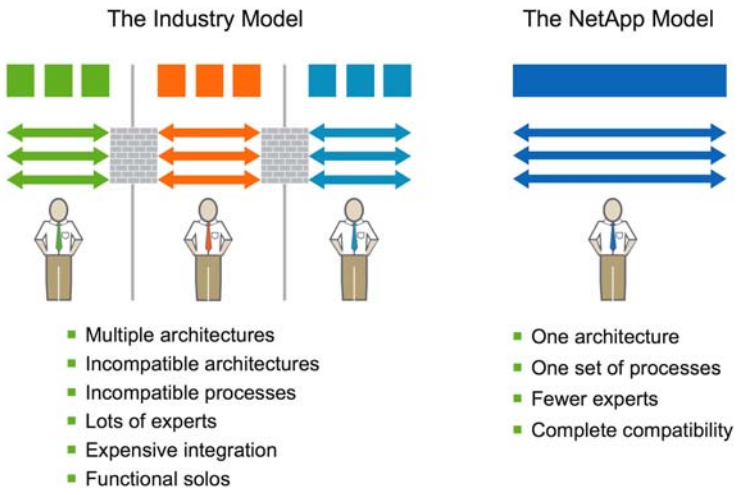

Figure 5) Unified storage architecture: summary of the NetApp advantage.

NETAPP WHITE PAPER

Scale-Out Storage and the Dynamic Data Center of the Future

Paul Feresten and Bruce Moxon, NetApp
May, 2008 | WP-7042-0508

EXECUTIVE SUMMARY

NetApp believes that the future of enterprise storage is a scale-out architecture that takes advantage of the superior price and performance of clustered components, facilitates nondisruptive operations, and employs policy-based management for improved efficiency and agility. We are working to make this vision a reality by delivering the proven feature set of Data ONTAP® on a scale-out foundation that uses the hardware building blocks that NetApp has successfully deployed in tens of thousands of existing storage installations.

Combining this scale-out architecture with emerging server virtualization capabilities is resulting in a new IT paradigm—a dynamic enterprise data center with the ability to rapidly commission and decommission applications, provision and reprovision resources, and nondisruptively migrate applications and data to adapt to changing service-level requirements.

This paper describes the requirements for scale-out storage in the enterprise in terms of resiliency, scalability, manageability, and other factors. It also discusses the current state of NetApp® scale-out deployments as well as specific operational and business advantages that the NetApp approach delivers to the enterprise data center. The pros and cons of alternative storage approaches are explored and some criteria are offered for evaluating scale-out solutions.

TABLE OF CONTENTS

1 INTRODUCTION

Enterprise data storage is at a crossroads. Massive data growth, tough economic conditions, physical data center limitations (power, heat, space), and new, agile software development methodologies are stressing IT organizations the world over. In the face of these pressures, traditional approaches to storage expose limitations in scalability, performance, availability, and—perhaps most importantly—manageability and flexibility. How will enterprise storage infrastructures evolve to meet growing demands as current trends continue? And how can enterprises store mountains of data in an efficient and cost-effective way while better leveraging both data assets and infrastructure investments?

This paper describes what NetApp sees as the future of enterprise storage—scale-out architectures that leverage the superior price and performance of commodity components; provide a dynamic, virtualized storage service infrastructure for nondisruptive operations; and employ policy-based management for improved efficiency and agility. The result is a storage and data management infrastructure that uniquely complements evolving server virtualization technologies to deliver on the promise of the dynamic enterprise data center.

Figure 1 shows the evolution of storage technologies, from direct-attached storage (DAS), to networked storage (SAN and NAS), and now to scale-out storage. Networked storage has allowed organizations to consolidate DAS and to realize the economies of scale of shared infrastructure and centralized management. In the same manner, scale-out storage consolidates networked storage devices, realizing further economies of scale, creating larger management domains, and enabling new resiliency models to support nonstop operations.

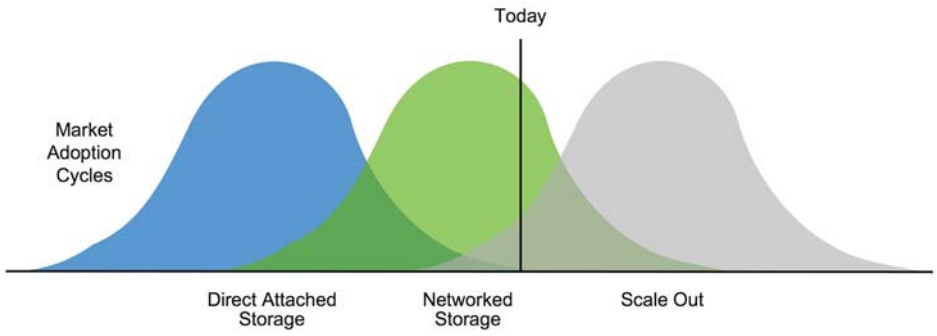

Figure 1) Storage evolution.

DEFINITIONS

Let's begin with a few definitions. to avoid any confusion about the meaning of important terms. (For more complete set of definitions, see the appendix, "Additional Definitions.")

- **Scale-out.** Scale-out is the strategy of using multiple smaller components in a collective manner to achieve the same performance as a single larger component. For example, today's scale-out computing approaches leverage racks of commodity 1U servers to achieve the same results that otherwise would require much more expensive supercomputer or symmetric multiprocessor architecture. Applied to storage, scale-out simply refers to the development of large storage systems based on many smaller storage subsystems; for example, modular midrange devices.
- **Dynamic data center.** A dynamic data center is one that leverages a utility computing model, centralized resource management, and rapid and flexible resource allocation to support a wide range of data center applications. This includes the ability to rapidly commission and decommission

applications, provision and reprovision required resources, and nondisruptively migrate applications and data between resources to meet time-varying application service-level requirements. It also includes the ability to seamlessly expand the underlying infrastructure and/or retire older components while maintaining continuous application operations through nondisruptive migration.

ORIGINS

The trend toward scale-out storage architectures follows on the success of the scale-out strategy in the high-performance computing world. In the past, many of the largest problems in scientific computing were run on monolithic supercomputers, developed in the 1970s by Cray Research and others. Today, such special-purpose machines have largely been supplanted by clusters of commodity servers with hundreds or thousands of nodes working cooperatively. In fact, 81% of the world's largest computing systems are now classified as "clusters."[1]

As these architectures have matured, they have increasingly found their way into production computing environments—ones that engineers rely on daily for circuit design and simulation, automotive design, development of digital effects, genome analysis, and large-scale financial analysis. Moreover, the same approaches have been applied in the deployment of large managed service infrastructures, such as those provided by large ISPs, and by software-as-a-service offerings such as Oracle On Demand.[2]

The main driver for this trend has been an economic one—the continued commoditization of standard server technology far outpaces any gains that could be accomplished with monolithic architectures. Of course, the ability to effectively use the physically distributed resources in a cluster, and to gracefully handle individual component outages, is critical to the success of this approach. Thus the development of software to manage and use these cluster computing architectures has been key.

The same economic driver is at the core of the current revolution in storage system architectures—the desire to leverage cost-effective components in a scale-out architecture that can meet increasing demands for capacity and performance while ensuring availability in the face of individual component failures.

ENTERPRISE PROBLEMS THAT SCALE-OUT STORAGE SHOULD ADDRESS

To date, scale-out storage architectures have predominantly been deployed in support of scientific and technical computing applications. However, the unique advantages of scale-out storage have begun to appeal to the broader enterprise storage market, which has its own unique set of requirements. Existing scale-out storage solutions may not be able to address all enterprise requirements. The ideal scale-out storage solution should address a number of key areas, as shown in Table 1.

---

[1] http://www.top500.org/stats/list/30/archtype

[2] http://www.oracle.com/ondemand/index.html

Table 1) Enterprise storage problems and requirements for a scale-out solution.

Enterprise Problem	Desired Scale-Out Solution
Capacity scaling. Capacity expansion in traditional storage systems may require downtime, either during physical installation or when redistributing existing data across the newly installed capacity.	Grow capacity incrementally, on demand, through the nondisruptive addition of storage shelves and growth of storage containers (pools, LUNs, file systems). Support nondisruptive redistribution of existing data to the newly provisioned capacity as needed.
Performance scaling. Standalone storage systems may lack the I/O/throughput to meet the needs of large-scale enterprise applications.	Grow performance incrementally, on demand, through the addition of storage controllers in small, economical (pay-as-you-grow) units without disrupting applications.
Availability. Traditional storage systems often have single points of failure that can affect data availability.	Leverage clustered controller configurations with appropriate failover support to ensure continuous data availability in the face of individual component faults.
Right-sized SLAs. Not all enterprise data requires the same level of service (performance, resiliency, and so on). Traditional storage systems support a single class of service, often resulting in poor utilization or unnecessary expense.	Support different levels of service and provide the ability to dynamically modify the service characteristics associated with stored data, by nondisruptively migrating data to slower, less costly disk, and/or by applying quality of service criteria.
Cost. With rapid growth, storage is consuming a larger and larger portion of shrinking IT budgets.	Control costs through the use of scale-out architectures that employ commodity components. Grow capacity and performance on an as-needed (pay-as-you-go) basis. Increase utilization through thin provisioning and data deduplication.
Leveraging the latest technologies. Traditional storage architectures require significant up-front investment and often lock customers into a particular "generation" of a system architecture.	Employ standard components to leverage rapid rates of improvement in storage density, performance, power consumption, cooling, and cost. Enable interoperability and compatibility across those components to incrementally leverage newer (more cost-effective) technology without the need for a forklift upgrade. Nondisruptively replace or remove older infrastructure to optimize price and performance, control maintenance costs, and ride technology commoditization curves.
Need for planned downtime. With traditional storage systems, software upgrades, data migrations, and hardware refreshes often require data to be taken offline	Eliminate the need to take data offline for such events by providing nondisruptive data and service migration throughout the storage cluster.
Complicated management. Discrete storage systems and their subsystems must be managed independently. Existing resource virtualization does not extend far enough in scope.	Provide a single point of management across the aggregate set of components. Leverage policy-based management to streamline configuration, provisioning, replication, and backup. Provide a flexible monitoring and reporting structure implementing an exception-based management model. Virtualize resources across multiple controllers so that volumes become simple-to-manage logical entities that span storage controllers for performance and dynamic redistribution of data.

2 CURRENT STATE OF NETAPP SCALE-OUT DEPLOYMENTS

NetApp began shipping clustered storage configurations with the release of Data ONTAP GX in early 2007. This marked the beginning of NetApp's commitment to a scale-out storage architecture, through a phased release plan in which the rich feature set of Data ONTAP 7G is delivered on a scale-out foundation. Existing NetApp customers continue to leverage the familiar and proven data management features of 7G, and they will be able to incrementally leverage scale-out features on their existing hardware, without the need to migrate data.

CURRENT SCALE-OUT DEPLOYMENTS

Most deployments of NetApp scale-out storage (Data ONTAP GX) to date have been in support of high-performance computing, aggressive digital media applications, and other applications that put significant demands on storage with file-based workloads, such as the following.

- **Back-end to large compute clusters.** These include scientific and engineering applications, seismic processing, automotive simulations, bioinformatics, and digital effects generation—applications in which large Linux® clusters are deployed to execute a single or limited set of applications, and where high aggregate data throughput rates are required.
- **Large-scale, active archives.** These environments may start small but grow to hundreds of terabytes or petabytes over time. Examples include large online photo and e-mail applications, studios digitizing videotape libraries for preservation and online distribution, and Picture Archiving and Communication System (PACS) installations that store medical imagery such as X-rays, MRIs, and CT scans. These applications typically have more modest performance requirements, the key being simplified management of vast quantities of data.
- **Dynamic enterprise infrastructure.** Many enterprise customers are looking to leverage server and storage virtualization to create an on-demand infrastructure for flexible application deployment and a more dynamic dev/test environment. Server virtualization, diskless boot facilities, and scale-out storage together provide an infrastructure where the applications (processes, OS/app images) and associated data can be nondisruptively migrated to alternative resources, ensuring adherence to service-level requirements in a nonstop operational model. Furthermore, if underlying resources are exhausted or if they fail, additional resources (servers, storage controllers, and capacity) can be brought online and application processes and data can be transparently redistributed to exploit newly deployed resources.

  These customers are leading the way toward the use of scale-out for the more heterogeneous application mix common in the enterprise. One NetApp customer, a major European hospital, has deployed Data ONTAP GX for use with multiple applications. They support two active archive applications: online, viewable versions of scanned paper patient records; and PACS. They also use Data ONTAP GX to store the active OS/application images used by VMware® servers. This allows VMware images to be redistributed across the storage cluster without disrupting running servers to balance load while meeting stringent application and data availability objectives.

OBSERVED BENEFITS FROM CURRENT NETAPP DEPLOYMENTS

Although the types of scale-out deployments have so far been limited, some significant benefits have been reported with immediate implications for enterprise data management.

- **Management simplicity.** Most Data ONTAP GX users cite as beneficial the ability to bring what otherwise would be a number of individually managed storage systems with significant capacity and I/O performance into a single system. With Data ONTAP GX, up to 12PB of data can be delivered from a single file system at an aggregate throughput in excess of 10GB/sec—all under a single point of management and the watchful "eye" of a single monitoring interface.
- **Ability to load balance.** In any storage installation, some storage systems are overloaded while others are underutilized. NetApp's scale-out storage architecture allows customers to transparently

migrate and redistribute data between storage controllers to meet performance objectives without any disruption to user access.

- **Elimination of planned disruptions.** For many applications, there's really no good time to take storage down for maintenance. And in the era of consolidation—where data from multiple applications is hosted on a single storage system—it's nearly impossible to plan maintenance periods. The NetApp scale-out storage model addresses these concerns in the following ways.
    - **Nondisruptive (rolling) software upgrades.** Individual controllers are upgraded one at a time. The inherent failover capabilities of the storage system ensure that clients have continuous access to data during this process, so data never goes offline.
    - **Elimination of forklift upgrades.** NetApp storage clusters support mixed hardware configurations, both controllers and disk subsystems. As a result, customers can roll in new controllers and disk subsystems and immediately initiate transparent data migration to distribute load and capacity to the newly added resources. Similarly, older hardware can be "drained" of data and decommissioned. This can be useful for removing and repurposing older controllers, and for transparently migrating data to newer-generation, denser disk drives.
    - **Transparent facilities moves.** Storage clusters can be geographically split across a suitable WAN link (20ms latency or less). This allows some Data ONTAP GX customers to do nondisruptive data center moves. In this process, data is migrated (again, nondisruptively) to a subset of the storage cluster nodes. The "empty" nodes are then transferred to the new facility and reintegrated into the cluster. Data is then (again, nondisruptively) migrated to the nodes in the new facility, completing a data center move with 100% data availability!
- **Suitability for virtualized server environments.** Virtual server environments create new challenges—and new opportunities—for storage. For example, VMware VMotion™[3] allows the entire state of a virtual machine to be migrated between ESX servers while the application is live. This is most typically used to balance virtual server load across a number of physical servers. The NetApp scale-out storage architecture provides the ability to nondisruptively migrate the *data* associated with an application (think of it as "DMotion"), in order to increase performance or to place less frequently used data on lower-cost SATA storage. It also supports scalable diskless boot implementations through the distribution of virtual machines across the storage cluster. Together, these capabilities provide the underpinnings of an end-to-end virtual application environment, as shown in Figure 2.

[3] http://www.vmware.com/products/vi/vc/vmotion.html

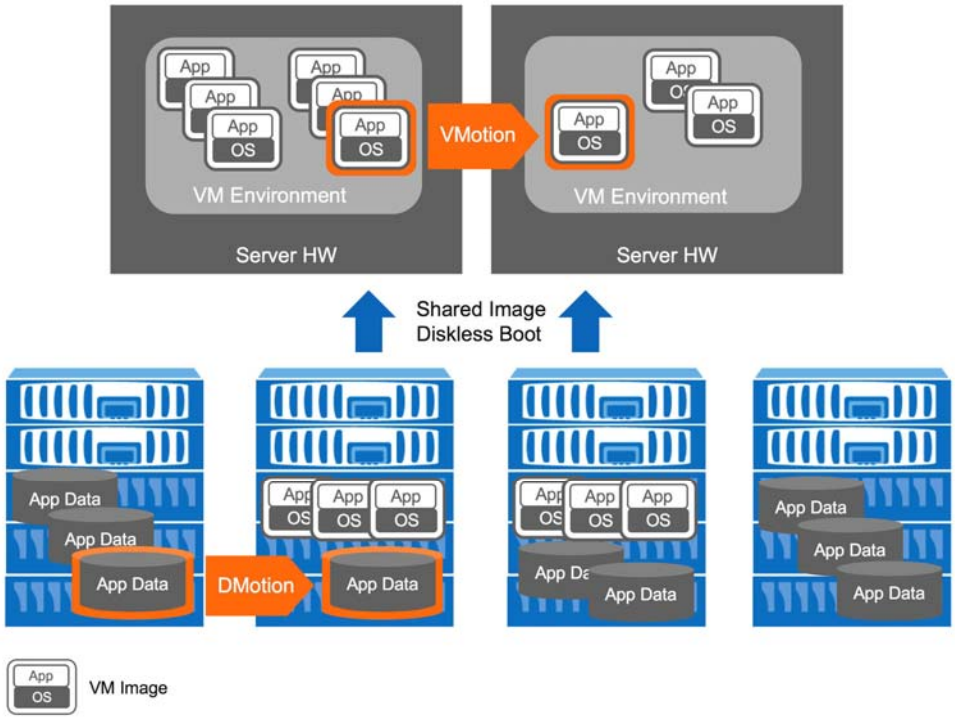

Figure 2) Combining server virtualization with scale-out storage to create a virtual application environment.

3 EXTENDING SCALE-OUT TO ENTERPRISE WORKLOADS

As scale-out technologies continue to mature, they will offer significant value for both next-generation business applications and traditional enterprise applications such as SAP, Oracle, and Exchange in terms of increased flexibility and improved data and resource management. This section describes the ways in which the NetApp scale-out storage approach can streamline enterprise environments and facilitate operational innovation.

THE PROMISE OF SCALE-OUT FOR ENTERPRISE APPLICATION ENVIRONMENTS

The NetApp scale-out architecture offers a number of benefits that enable enterprise data centers to make significant operational improvements. Most of these improvements can already be achieved with Data ONTAP GX. A few will become available in the future as Data ONTAP 7G and Data ONTAP GX converge.

TECHNOLOGY REFRESH

In terms of both disks and storage controllers, each new generation of storage technology is typically denser, faster, cheaper, and more environmentally friendly than its predecessor. In the dynamic enterprise, it can be important to migrate critical applications to the latest technology while repurposing and eventually decommissioning older technology. Because this process has traditionally been highly disruptive, however, companies have often chosen to make do with older technology for longer than they otherwise would, often incurring very high extended maintenance costs.

The NetApp scale-out storage architecture eliminates the disruption that results from technology refresh. When new storage is added, existing data can be migrated nondisruptively to leverage the new storage. Data can be migrated off older disks or controllers for decommissioning, or older technologies may be repurposed as part of a tiered storage strategy to extend its lifecycle.

TIERED STORAGE

The NetApp scale-out storage infrastructure can encompass different tiers of storage, including both high-performance disk (FC and SAS drives) and capacity-oriented disk (SATA drives). Drives in RAID groups are combined to form storage pools in which dynamically allocated, thin-provisioned volumes can be created. Additional semantics can be stamped on these volumes. For example, the SnapLock® feature allows the

creation of write once, read many (WORM) volumes with specified retention periods for compliance purposes. All of the volumes in the system are part of a global namespace; no matter where the data physically resides (that is, behind which controller), it remains accessible to user sessions bound to any of its controllers. The system manages this namespace, allowing data to be migrated nondisruptively from one storage pool (tier) to another or across controllers in order to meet changing service-level requirements.

RESOURCE AGGREGATION FOR APPLICATION PERFORMANCE SCALING

Another advantage of the NetApp scale-out storage approach is the ability to aggregate resources to provide performance exceeding what any individual resource can achieve. Analogous to the clustering of multiple database servers in Oracle® Real Application Clusters (Oracle RAC), storage resources across multiple storage controllers can be employed to deliver much greater I/O performance to an application than a single storage controller could achieve alone.

Storage pools can span different subsets of the storage system's resources, leveraging heterogeneous controller and disk configurations to support pools with different performance, cost, and resiliency characteristics. This in turn allows organizations to more effectively meet the unique needs of heterogeneous enterprise applications.

RIGHT-SIZE PROVISIONING AND REPROVISIONING

In the typical approach to application provisioning, you estimate what your peak demand will be and acquire servers and storage to meet that level of demand, much of it up front. The application is deployed in a test or pilot mode and then slowly rolled out until full production is reached. As the application ages, it enters a maintenance phase in which capacity utilization may remain high, but performance requirements typically diminish.

The best case under this scenario is that the initial estimate of peak load is correct, resulting in good performance during peak production but underutilization in the pilot and maintenance phases. The more likely outcome—because of the difficulty of estimating future application load—is that the application is either over- or underprovisioned. Overprovisioning means that you spend more than necessary on infrastructure. Underprovisioning means that the application won't meet service-level requirements and may need to be reprovisioned with new servers and/or storage. This process is almost always painful and disruptive.

The NetApp scale-out storage architecture gives you the ability to provision applications appropriately during all phases of operation. Starting with modest resources (capacity and I/O) during the pilot phase, you can transparently and dynamically reprovision as more resources are required to move into production and then scale back again when the application enters its maintenance phase. By combining a virtual server infrastructure with nondisruptive data migration among storage pools of varying service levels, it is possible to "right size" application resources throughout an application's lifecycle. This includes making resource adjustments to accommodate unforeseen growth or to meet cyclical changes in demand, such as quarter-end or year-end processing.

AGILITY

Enterprises need to respond quickly to business changes, resulting in new application requirements as well as changes in technology. NetApp allows storage resources to be dynamically adjusted to meet changing SLAs without affecting applications or users.

MAXIMUM SPACE EFFICIENCY

With the rate of data growth showing no sign of slackening, the best way to get storage costs under control is to maximize utilization of every disk deployed. The following NetApp technologies that increase utilization are already part of Data ONTAP GX:

- **Space-efficient Snapshot™ technology** avoids data copies and has no performance impact.
- **Disk-based backup** is inherently free of the data duplication found in most backup methods.
- **Thin provisioning** allows an underlying pool of free storage to be oversubscribed. New storage capacity can be provisioned on a just-in-time basis, rather than leaving large volumes of storage sitting idle.

The following additional technologies will become available as the NetApp technology convergence proceeds:

- **Space-efficient clones** will significantly reduce the footprint of the multiple copies of data that organizations generate to support development and test, QA, and analytic activities.
- **Deduplication** identifies and eliminates the duplicate blocks in a storage container. In virtual server environments, deduplication can recover 50% to 70% of the disk space by eliminating the duplication that results from having multiple virtual machines running the same OS and applications.

POLICY-BASED MANAGEMENT

Scale-out storage has the potential to bring a huge amount of storage capacity and I/O capability under a single point of management, but that still leaves a lot of individual entities to manage—the individual storage pools, volumes, LUNs, and any replicas required for backup or DR purposes. NetApp has made a significant investment in a new data management paradigm that employs policy-based administration and extensive delegation. These capabilities have been widely deployed in NetApp's existing Data ONTAP 7G installed base, and will be fully leveraged in the NetApp scale-out architecture in the near future.

With policy-based management, individual storage pool entities are easily aggregated into data sets. A policy, such as the replication type and frequency, can then be applied to a data set, ensuring that the policy is consistently applied to every member. Adding a new member to the data set results in the policy automatically being applied. If the policy changes, that change propagates to every entity in every data set to which that policy has been applied.

For example, suppose that you create a new set of data volumes and you want to create Snapshot copies of those volumes every 15 minutes. You create a data set to contain the volumes. Then, starting with a predefined template, you create a policy specifying the Snapshot interval and apply that policy to the data set. New volumes that are added to the data set automatically inherit the same policy.

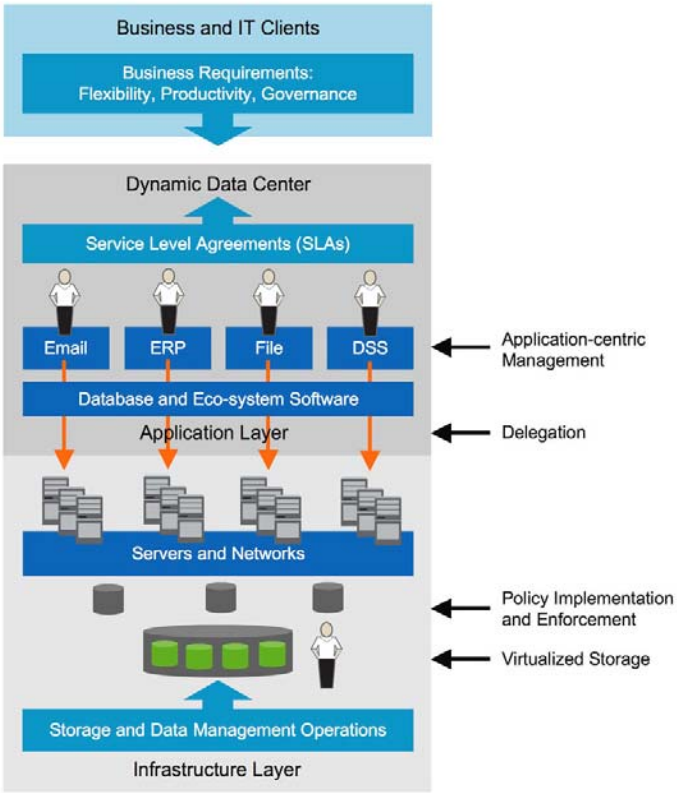

Figure 3) The NetApp management approach includes a policy-based strategy with delegation to applications.

Delegation extends the advantages of policy-based management, allowing members of the IT staff, such as application administrators and DBAs, to provision storage according to predefined policies and within limits set by the storage administrator. This effectively distributes many of the day-to-day storage provisioning tasks among members of the IT staff in a self-service model. Policy constraints and a rich set of reports provide the necessary controls and oversight to ensure consistency with operational guidelines.

This approach allows a large number of entities to be grouped and collectively managed efficiently with minimal time invested, allowing storage administrators to manage more storage, more effectively in less time. This, in turn, frees them up to pursue more strategic initiatives.

ENABLING BUSINESS BREAKTHROUGHS

The capabilities just described offer effective tactical approaches to streamlining data center operations. This section examines the way in which a scale-out storage architecture can provide a foundation for business breakthroughs—new operational approaches that can fundamentally change the business.

BETTER SUPPORT FOR SERVICE-ORIENTED ARCHITECTURES

Service-oriented architectures (SOAs) seek to create business applications by compositing a set of discrete underlying (Web) services. These services may be applied in different ways, and may be reused by different applications to provide a consistent look and feel, implement consistent data access techniques, and reduce the amount of software developed to support a diverse set of applications. By their very nature, SOA-based applications tend to create a more dynamic workload than traditional applications. They support rapid application development through service compositing, and they can often leverage multiple copies (instantiations) of a service on multiple hosts in response to peak activity or growing loads.

The NetApp scale-out storage architecture can adapt quickly to changing SOA application needs through improved agility in the storage system, allowing service levels to be changed to meet dynamic application requirements.

MANAGED SERVICES

An important trend in the IT industry has been the move toward managed services. Rather than host applications in-house on dedicated infrastructure, companies are increasingly turning to third-party hosted services such as Oracle On Demand, which can provide both production and development environments for a customer's Oracle applications.

Going forward, scale-out storage architectures promise to deliver more agility for large-scale managed service environments, enabling them to quickly deploy new application instances for new customers and to redeploy resources as necessary to enable optimal utilization of the infrastructure to meet customer SLAs.

VIRTUAL SERVER ENVIRONMENTS

Virtual server technology brings a number of benefits to application infrastructure in terms of the ability to rapidly and efficiently deploy whole applications and application stacks while leveraging the sweet spot of underlying hardware building blocks. It creates new challenges, however, in terms of data protection and disaster recovery.

Scale-out storage complements virtual servers by allowing data to be moved and managed in a manner that reflects the agility of virtual server environments. However, it may not intrinsically address the problems of data protection. For our part, NetApp has worked closely with VMware to integrate NetApp data protection, replication, and deduplication technologies in Data ONTAP 7G. In the future, virtual server environments will be able to leverage this mature set of data management capabilities in combination with the unique scale-out capabilities described here for load balancing, performance aggregation, and nondisruptive data migration.

4 NETAPP SCALE-OUT VERSUS ALTERNATIVES

NetApp's architectural vision combines its rich storage and data management software stack, Data ONTAP, with a switched storage architecture. The result is a storage architecture that delivers the feature set of Data

ONTAP on a scale-out foundation built from the same hardware components that NetApp has successfully deployed in tens of thousands of existing storage installations. This integrated approach builds on NetApp's core "DNA"[4] and uniquely positions NetApp to deliver scale-out storage as a critical element in a new IT paradigm, while leveraging customers' existing investment in hardware and software, training, and operational infrastructure.

In contrast with NetApp's unified approach to scale-out, alternative approaches typically attempt to bring together disparate architectural elements, including components obtained from a number of suppliers or through multiple acquisitions. Such components, by their very nature, have disjoint feature sets and incompatible management tools, often requiring customers to trade off features for desired service-level characteristics. Those tradeoffs can lock applications in to specific components and limit the ability to migrate applications in the face of changing service-level requirements.

More distributed architectural approaches to storage, as exemplified by the Google File System and other "cloud computing" approaches, have recently attracted much interest. Such systems offer a streamlined feature set and are best suited to certain classes of applications, including large-scale analytics and distributed fixed- content stores. They complement NetApp's architectural approach in supporting a new class of Web-scale applications, but they are less well suited to the general-purpose storage requirements of the enterprise data center.

5 SUCCESSFUL SCALE-OUT STORAGE DEPLOYMENT

When assessing scale-out solutions for possible deployment, you should keep the following questions in mind: there

- What should I buy now so that I can evolve smoothly to scale-out storage in the future?
- Can I leverage current storage investments, or will all existing infrastructure need to be replaced?
- What storage features will be required to support business applications now and in the future?
    - What protocols?
    - Is there specific support for application integration?
    - What degree of data protection and disaster recovery are needed? What resiliency options does the solution provide?
- Is any special hardware or software required on the client side?
- How space efficient is the solution? Does it provide options for thin provisioning, deduplication, and other space saving methods?
- How is the solution managed? Does it support policy-based management?
- How easy is it to add new resources? Do different generations of building blocks work together?
- Can the solution be upgraded nondisruptively? Can new software revisions be installed, new hardware added, and old hardware decommissioned without disrupting data access?

[4] "The Core NetApp DNA," Tech OnTap, February, 2007, http://partners.netapp.com/go/techontap/matl/NetApp_DNA.html

6 CONCLUSION

At NetApp, we strongly believe that the future of enterprise storage is a scale-out architecture that leverages the superior price and performance of clustered components; provides a dynamic, virtualized storage service infrastructure for nondisruptive operations; and employs policy-based management for improved efficiency and agility. We are working hard to make that vision a reality.

NetApp has been shipping scale-out storage systems with Data ONTAP GX for over a year—delivering on the promise of scale-out for customers who are running demanding file-based applications. As we complete the convergence of the Data ONTAP 7G lineage with this scale-out foundation, the full richness of our complete data management stack—including proven tools for application integration, data protection, disaster recovery, and advanced data management—will become available to enterprise scale-out deployments.

The result will be a new paradigm in enterprise IT: a dynamic enterprise data center with the ability to rapidly commission and decommission applications, provision or reprovision required resources, and nondisruptively migrate applications and data across pools of resources to fluidly adjust to changing service-level requirements.

.

APPENDIX: ADDITIONAL DEFINITIONS

- **Virtualization**. Virtualization is a broad and frequently used term in the computer industry. At the simplest level, it refers to the abstraction of a set of computer resources into a higher-level concept or entity.[5] It is used to hide the characteristics of the underlying physical components in order to present a simplified model of use. For example, a virtual server is an aggregation of a set of machine resources (processing, memory, I/O) that looks like a standalone server—able to be configured, have software installed, and run just as if it were a physical machine. In storage systems, virtualization occurs at many levels. For example, a *volume* is a virtual construct that aggregates a number of RAID groups, which in turn aggregates a number of disks.
- **Clustered storage system**. A clustered storage system is one that is composed of a number of storage controllers (in a scale-out architecture) that collectively manage the underlying storage pools and process the requests of connected clients. Clustering provides two significant benefits: 1) increased aggregate performance through the distribution of load across individual controllers; and 2) improved resiliency through service failover between controllers. A clustered storage system is managed and accessed as a single system, with requests transparently distributed among the individual elements. The clustered storage system virtualizes the underlying controllers and disk subsystems, providing a single point of management and access.
- **Utility computing**. Utility computing is an approach in which computational infrastructure (compute, network, storage) is deployed, managed, and typically billed in an on-demand model, similar to the way in which public utilities operate.
- **Cloud computing.** Cloud computing is a relatively new term that describes several recent system and application architecture trends. The first of these employs a highly distributed (hundreds to thousands of nodes) compute and storage architecture used for Web-scale processing and analytics. This is the GoogleFS[6]/MapReduce[7] approach to building massively parallel services for Web-scale data analysis. Yahoo has backed a similar activity, Hadoop[8] (and the Hadoop File system, HDFS[9]), an open source volunteer project under the Apache Software Foundation. Amazon's combined S3[10]/EC2[11]offering can also be used in this manner.

  The second trend uses a similar architecture to develop large, distributed, fixed-content stores. It emphasizes storage-oriented aspects of these projects and deemphasizes processing. It can be described as a storage-as-a-service architecture where clients can "get" and "put" objects in a multisite, intranet- or Internet-based "cloud." In some ways, this is a new twist on the content addressable storage (CAS) architecture[12]. Such services, including derivatives such as HBase, a Hadoop-based "databaselike" layer[13], may be the building blocks that support the evolution of higher-level (software-as-a-service) capabilities such as remote backup and recovery, online archive, e-discovery, and others.

---

[5] http://en.wikipedia.org/wiki/Virtualization

[6] http://labs.google.com/papers/gfs.html

[7] http://labs.google.com/papers/mapreduce.html

[8] http://hadoop.apache.org/core/

[9] http://hadoop.apache.org/core/docs/current/hdfs_design.html

[10] http://aws.amazon.com/s3

[11] http://aws.amazon.com/ec2

[12] http://en.wikipedia.org/wiki/Content-addressable_storage

[13] http://wiki.apache.org/hadoop/Hbase

NETAPP WHITE PAPER

Building a Dynamic Data Center:
Transform Your Environment to Cut Costs and Increase Efficiency

Hamish McGovern and Rob Pollock, NetApp
February 2009 | WP-7070-0209

EXECUTIVE SUMMARY

This paper describes the steps required to begin building your next-generation data center. It describes the elements of this new dynamic data center infrastructure as well as specific cost, time-to-market, and risk mitigation advantages that the NetApp approach delivers to the enterprise data center.

TABLE OF CONTENTS

1 INTRODUCTION

Data centers are at a crossroad. Massive data growth, tough economic conditions, and physical data center limitations (power, heat, space) have exerted extreme pressure on IT infrastructures. Finding ways to take cost, complexity, and associated business risk out of the data center without sacrificing service levels has become a major objective for almost every enterprise. In theory, consolidation and virtualization make it possible to change the way data centers are architected, built, and managed. But how do you start your improvement journey and avoid disruptions and added costs to your ongoing operations?

Those that have embarked on the journey to the next-generation data center have done so leveraging a dynamic computing model, centralized resource management, and rapid and flexible resource allocation to support a wide range of data center applications. This service-oriented infrastructure (SOI) approach is commonly referred to as a dynamic data center (DDC). NetApp has worked with many enterprises across the globe to create highly efficient dynamic data centers that reduce overall capital and management costs by 40% or more.

The NetApp® dynamic data center solution provides a proven project delivery methodology and IT services management framework that make it possible to rapidly commission and decommission applications, provision (and reprovision) required resources, and nondisruptively migrate applications and data between resources to meet changing service-level requirements. It also includes the ability to seamlessly expand the underlying infrastructure and/or retire older components while maintaining continuous operations.

This white paper discusses the key challenges that enterprises face and explains how the NetApp dynamic data center solution efficiently solves these challenges. It describes in detail the elements of the NetApp dynamic data center solution, including:

- An SOI that utilizes customer-defined standardized configurations to provide a strongly typed architecture
- A service management framework that clearly defines all the services necessary to run your data center
- A well-defined project delivery methodology that takes the risk and complexity out of transitioning your data center

We conclude by examining how IT as a service is offered through the form of a service catalog to business and application owners.

2 KEY DATA CENTER CHALLENGES

The challenges faced by enterprises can be grouped into three general areas:

- Reducing total cost of ownership
- Making sure of time to market
- Minimizing business risk

TOTAL COST OF OWNERSHIP

Total cost of ownership (TCO) includes not just the acquisition cost, but also ongoing management costs that are highly dependent on the operational efficiency of your data center. The deployment of point solutions to meet storage needs is a major factor that can negatively impact TCO. With point solutions, storage deployment is inconsistent, and each time you deploy a new point solution the less and less predictable.

A second factor that drives up storage TCO is inefficient storage management—a problem exacerbated by point solutions. If you have multiple types of storage systems with different interfaces requiring different management processes, managing that environment is going to be more complex and more costly.

A final factor affecting TCO is the necessity to cope with common data growth rates in most enterprises in the neighborhood of 50% to 100%. The time needed to procure and deploy this amount of storage every year has a huge destabilizing impact on an IT infrastructure.

The NetApp dynamic data center solution addresses these issues by standardizing your infrastructure and practices, reducing complexity and rationalizing procurement and deployment so you can deliver better results at less cost.

TIME TO MARKET

Latent time to market might result from inadequate capacity planning such that capacity is not available when it is needed, impeding business processes and reducing agility. The NetApp dynamic data center solution attacks these problems by driving the "white space" out of the procurement and deployment process so that you can deploy storage as it is needed so that you always have capacity on demand, without overprovisioning.

BUSINESS RISK

Heightened business risk results from the lack of coherent methodologies for backup and recovery as well as disaster recovery. A data center that consists of many diverse point solutions makes the application of standardized technologies and processes difficult or impossible. The NetApp dynamic data center solution provides a coherent architecture for backup and recovery and disaster recovery, so that processes are the same regardless of the underlying application.

Table 1) Enterprise challenges that the NetApp dynamic data center solution helps address.

Enterprise Challenge	How NetApp DDCS Helps
Substantially lower costs	Centralizes control and consolidates infrastructure to allow for pooled resources. Provides on-demand provisioning and activation at the company level rather than by project needs. Maximizes storage utilization, minimizing total raw capacity required and reducing associated costs for space, power, and cooling.
Faster time to market	Removes labor-intensive tasks and allows focus on policy administration. Dynamic infrastructure delivers rapid and reliable services through a single management framework with well-defined processes and a consistent delivery methodology.
Reduced operational risk	Data protection is fully integrated. Backup and recovery are designed into the underlying data structure and consistent across multiple classes of data.
Reduced deployment risk	Proven end-to-end delivery approach to take the guesswork out of moving to IT-as-a-service model.
Solid foundation for future cloud services	Modular approach toward vision that complements current projects and budgets. Dynamic infrastructure provides for a gold standard and default environment for reliable services delivered throughout the data center.

3 JOURNEY TO IT AS A SERVICE

As you can see, each step is taken with big picture design in mind and will generate significant benefits. The five-phase adoption to the dynamic data center involves:

1. **Centralize management** of IT to gain visibility of costs, take control of the IT offering, gain economies of scale, and begin the journey to service-oriented infrastructure.
2. **Standardize the offering** based on the key business requirements. Attempting to support bespoke solutions for each application is costly and slow to provision. Standardization of offering and creating repeatable processes are key to improving quality and provisioning times and reducing support cost and risk. Standardization is a prerequisite to successful consolidation and automation.
3. **Virtualize and consolidate** the physical infrastructure for virtualization consolidation, driving up asset utilization and with NetApp massive storage efficiency benefits. Virtualization happens at each level of the infrastructure stack: storage and data protection, network, and server layers to increase asset utilization and simplify asset lifecycle management through mobility of applications and data. With shared services and multitenancy come much faster time to market and considerably lower overall cost structures, and the responsibility of nonstop infrastructure. You've got the basis of SOI.

4. **Automate** the environment. Once the offering and processes are standardized and infrastructure virtualized, the step of automation is possible. Automation tools increase abstraction, providing simple controls for overall workflow management.
5. **Self-service** and APIs for delegated control. Handing back control is the crowning achievement. Allowing application administrators and owners the flexibility of scaling on demand, choosing different levels of performance and data protection as required, and automating recovery from applications errors are all possible through application integration and self-service capabilities making the jump from service-oriented infrastructure to enterprise cloud services.

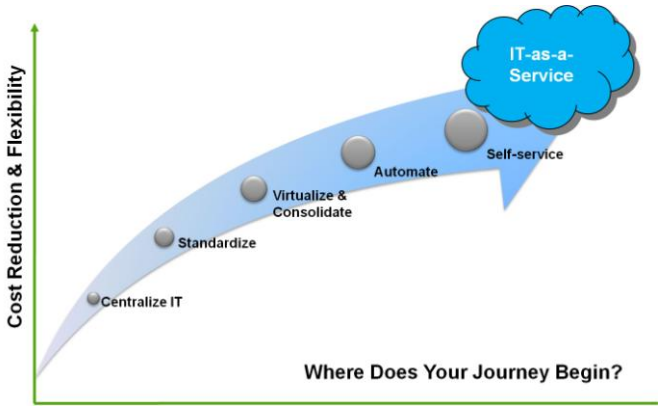

Figure 1) Conceptual view of the journey to IT as a service.

4 NETAPP SOLUTION ELEMENTS

The NetApp dynamic data center solution addresses enterprise storage challenges by taking the complexity and cost out of your data center. The combination of an SOI and a well-defined management framework takes the guesswork out of provisioning and management, while our defined project delivery methodology provides the means for you to transition to the solution with minimum risk to ongoing operations.

SERVICE-ORIENTED INFRASTRUCTURE (SOI)

Why do you need an SOI?

Simply put, for consistency. This strongly typed architecture allows you to consume and stand up resources in a repeatable manner. It is this ability to duplicate results that provides you with operational efficiencies and rapid provisioning. Now let's take a closer look at the logical overview of the NetApp dynamic data center solution.

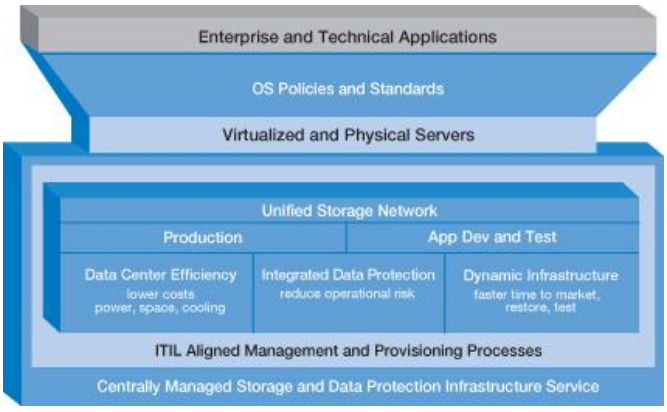

Figure 2) Logical overview of the NetApp dynamic data center solution.

The technology stack of the NetApp dynamic data center solution consists of four primary layers. The goal is to drive compliance through the stack in the data center when we want to either scale out the physical pool (provisioning) or consume storage from the pool (activate applications).

- **Application layer** has a number of application and OS standards attached to the environment that we call activation guides. These guides detail how we will consume the storage in the data center and get real process efficiency.
- **Server layer** comprises both physical and virtual servers. The goal is to provision, activate, protect, and recover your data quickly and cost effectively.
- **Network layer** design goal is reasonably simple: build a single data center storage solution that can simultaneously cater for Ethernet and Fibre Channel traffic. This provides the ability to deliver a consistent set of features and performance.
- **Storage and data protection layer** consists of a more agile and dynamic infrastructure that can change as the business demands change. This unified storage architecture gives you the flexibility to choose not only the right protocol, but also the right performance, and in the most cost-effective way to meet the most aggressive compound annual storage growth rates. This multitenancy service provider style solution integrates a broad range of products and capabilities.

Together these provide (1) high data center efficiency through thin provisioning, zero space cloning, multiple classes of service, and efficient data protection; (2) application-aware data protection for both system and site failures to minimize operational risk; and (3) dynamic infrastructure through rapid backup and recovery, instantaneous cloning for development, QA, data mining, and image cloning.

New capabilities can quickly be made available to all applications within the environment. The dynamic nature of the solution meets the changing demands of a wide range of business applications and flexibility to choose not only the right protocol, but the right performance, and in the most cost-effective way to meet the most aggressive compound annual storage growth rates.

In summary, this SOI lets you scale out and consume storage with minimum management overhead for maximum efficiency and is a requirement for automation of the solution.

SERVICE MANAGEMENT FRAMEWORK

The NetApp service management framework is closely aligned to the Information Technology Infrastructure Library (ITIL), which has been designed to increase the rigor of best practices used within the IT industry. A key goal of ITIL is to help IT teams better understand costs and reduce fixed costs wherever possible.

This service management framework defines all of the services necessary to operate a storage infrastructure within an enterprise. It includes seven services:

- **Lifecycle management** covers both configuration and asset management.
- **Capacity planning** defines how you determine what storage you need so that you can have capacity available on demand.
- The **solution architecture** service defines how you add new feature sets to enhance the capabilities of your storage environment. You might choose to add such features as compliance, encryption, thin provisioning, deduplication, and so on.
- **Procurement** is the act of buying storage to scale out your environment. The procurement service makes sure that you always buy the correct storage and deploy it in the proper way. The combination of a strongly typed architecture and this procurement service makes sure that the right purchasing decision is made every time.
- **Storage design and activation** define the way you deploy storage and include the design documents necessary to describe the environment as well as the methods for storage activation.
- **Change management** integrates your storage environment into your enterprise-wide change management processes, allowing you to effectively assess the risk of any changes that are made.
- **Maintenance and support** define the maintenance and support services that make sure of the health and availability of your storage environment.

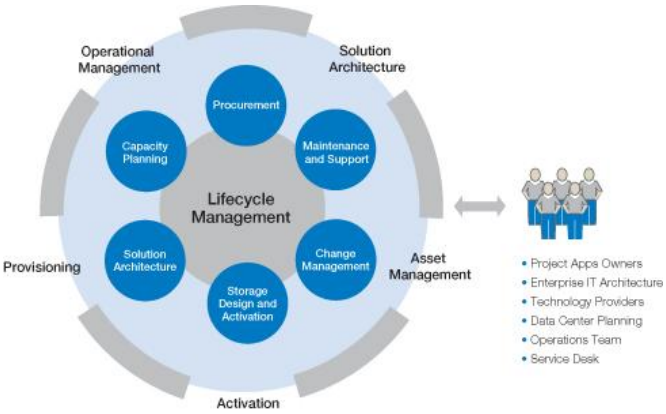

Figure 3) Conceptual view of storage management framework.

A service in and of itself doesn't do anything. We've designed five processes that overlie the seven services. These five services are all that are required to run your data center.

Provisioning Physical Capacity

The purpose of the provisioning process is to assess the project requirements and make sure that sufficient capacity is installed into the pool ready for activation as required. The process is primarily concerned with capacity planning, the deployment of the physical resources in accordance with the defined architecture standards.

A well-formed provisioning process can reduce the labor required to deploy new storage devices by using standard configurations and designs to remove the "white space" in the provisioning process.

Activation

The activation process uses virtualization to present storage and data protection capability to applications. Applications consume the preprovisioned storage in accordance with the activation standards. The process is primarily concerned with consumers of the storage and the configuration control of the deployed solutions.

A strong activation process can dramatically affect the time taken to activate resources in support of business initiatives by using standard activation guides. Consistent activation processes will also reduce the labor required, as activation is template driven. This removes the variability in the activation of technology and can drive efficiency right through the technology stack in the data center.

Solution Architecture

The purpose of the solution architecture process is to evolve the architecture in a controlled manner. This evolution may include software version upgrades, hardware models, or the introduction of new functionality.

The solution architecture process, in making considered decisions and maintaining the configuration of the environment, assists in reducing support costs and eliminating outages in the production environment.

Asset Management

The purpose of the asset management process is to manage the deployed assets. This includes leasing of hardware, software licenses, and versions. It also includes the replacement of aging assets.

Asset management makes sure of efficient use of assets within the infrastructure to maximize the return on investment.

Operational Management

The purpose of the operational management process is to maintain the deployed infrastructure. Interfacing to the customer's problem management processes in addition to the automated altering functionality provided by the NetApp toolsets, the process is the key to operational stability.

The strongly typed architecture used to provision and activate storage creates efficiency in the support process. Standard configurations reduce outages and ultimately reduce the support costs in the event an outage does occur because the configuration is known.

THE SERVICE CATALOG

The NetApp dynamic data center solution represents a paradigm shift for IT. No longer will very specific technology requirements be presented to IT on an ad-hoc basis. The data center of the future will require business and application owners to use a service catalog of standard configuration from which to build their applications. The service catalog is tailored to satisfying 90% of the business needs from a defined set of offerings that can be provisioned and integrated rapidly.

Typically the service catalog will provide:

- Classes of storage and integrated tiers to effectively match applications with their capacity and performance needs
- Different levels of data protection to meet the applications RPO and RTO requirements that cater to both on-site recovery points and those that are geographical dispersed
- Support for a standard set of physical and virtual OS and server types
- Activation guides to provide standards around application connectivity and data layout configurations, which have been tested for throughput and latency
- Long-term retention offering for archive or compliance needs and WORM-style compliance capabilities where needed

WELL-DEFINED PROJECT DELIVERY METHODOLOGY

The prospect of transforming your data center by implementing the NetApp dynamic data center solution might seem daunting, but NetApp has helped customers worldwide make the transition. Our dedicated Professional Services team and Authorized Professional Service Partners have the expertise to assess your situation, provide valuable insights, and create and execute a plan that will help you achieve your goals with the least cost, minimum risk, and the least disruption to your ongoing operations.

Our project delivery methodology is designed to mitigate the risks associated with the transition. We use a well-defined work breakdown structure; clear project management methodology; and a clear division of labor between NetApp, NetApp partners, and your IT team. Our phased approach allows us to rapidly deploy services in a predictable and repeatable fashion.

We recognize that everyone's needs are different and that not everyone is inclined or able to undertake a full data center transformation at one time. That's why we offer the Fast-Start Workshop to help illuminate the path to complete data center transformation, in a prescribed format that addresses your unique business needs.

The NetApp Fast-Start Workshop is a four-day program designed to:

- Identify application and infrastructure targets based on capacity, performance, and service-level needs
- Assess impacts in terms of cost savings, efficiency gains, and performance improvements
- Identify top five process improvements in terms of impact on cost and efficiency, agility and timeliness, and progress toward IT as a service

You choose the granularity of transformation that makes sense for your needs, from incremental steps to a single large effort. If you proceed in a stepwise fashion, all planning is done with the end goal in mind, so no effort is wasted. The savings that result from each step can help to fund the next step.

5 BENEFITS OF THE NETAPP SOLUTION

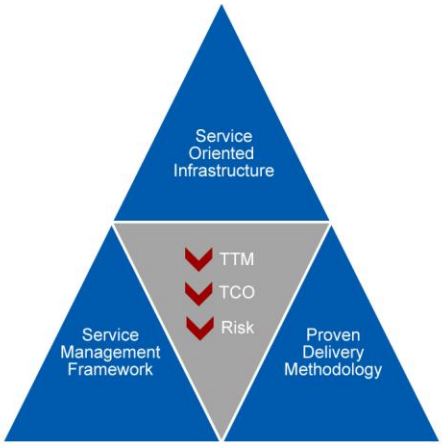

Figure 4) NetApp dynamic data center solution addresses the three biggest contributors to inefficiency in the data center.

TOTAL COST OF OWNERSHIP

In regard to total cost of ownership, the NetApp solution drives down acquisition costs by dramatically increasing storage efficiency. By eliminating point solutions and standardizing your storage architecture, you get the full benefit of storage consolidation. This is amplified by proven NetApp storage efficiency technologies, including unified storage with flexible volumes, space-efficient Snapshot™ copies and clones, thin provisioning, and deduplication. These technologies significantly increase your storage utilization and reduce the total amount of storage you need while at the same time reducing your requirements for floor space, power, and cooling. You can learn more about these technologies in a recent white paper in this series entitled “Unified Storage Architecture: Enabling Today’s Dynamic Data Center.”

A second factor driving TCO is operational efficiency—the number of people it takes to operate your storage environment. In one example of a NetApp dynamic data center solution Implementation the client was able to grow its storage environment by 10X (from 1.6PB to 16PB) without adding additional support staff.

If you are experiencing the typical 50 to 100% annual growth rates for storage, the NetApp solution—because of its strongly typed architecture and well-defined storage activation processes—can enable you to scale out your environment without adding to your support team. All the metrics that pertain to the TCO benefits of the NetApp dynamic data center solution are summarized in Table 2.

TIME TO MARKET

There are two important time-to-market benefits of the NetApp solution: a major reduction in the time it takes to deploy and activate production storage and dramatically reducing “standing up” an application development and test environment.

The traditional process of deploying a point solution for production storage includes:

- Determining what architecture to use
- Deciding on a storage system
- Finding budget and procuring the hardware
- Installing the hardware
- Laying out the storage volumes

This is obviously time consuming, and the end results are highly dependent on the quality of the decisions made at each point.

With the NetApp service management framework, all this is clearly defined in advance, so there are no delays, and the outcome is always as desired. We take the provisioning of production storage down from months to days—or hours if needed.

The second time-to-market benefit comes during provisioning for development and test. NetApp cloning technology allows you to provision test environments very rapidly without making space-consuming copies. Your storage team simply sets and enforces storage polices rather than constantly provisioning storage.

This cloning capability puts your development and test team in control of its own destiny. The team is free to clone all the test environments it needs without consuming the massive amounts of storage that would typically be required. The result is a significant speed up in development and test with improved quality.

5.1 BUSINESS RISK

A storage environment that consists of a number of disparate point solutions with a variety of backup and DR solutions that is growing at 50% to 100% annually will quickly become unsustainable. Such an environment is simply too labor intensive, and it ultimately becomes impossible to scale the environment further with reasonable resources.

The NetApp dynamic data center solution utilizes disk-based backup and the NetApp technology suite to define a single set of standard practices that encompass your entire storage environment to significantly reduce business risk.

Table 2) Benefits of the NetApp dynamic data center solution.

Area	Expected Benefit
Infrastructure operations	➢ Manage up to 2PB per full-time admin ➢ Reduced provisioning from weeks to minutes ➢ Increase utilization rates to at least 75% ➢ Cut raw storage in half
Data center operations	➢ Reduce power between 20% to 50% ➢ Consolidate and save up to 50% in floor space
Application development	➢ Complete frequent full-scale infrastructure testing with limited scale cost ➢ Test more and test more often to deliver a higher quality IT environment
Database operations	➢ Backup windows virtually disappear ➢ Reduced from hours to minutes

6 CONCLUSION

High storage growth rates combined with tough economic conditions are making traditional practices for data center management unsustainable. The NetApp dynamic data center solution addresses critical enterprise challenges to reduce costs, mitigate risk, and help IT teams succeed in the face of increasing budget constraints.

The NetApp solution combines a strongly typed architecture that simplifies scaling and activation of applications with a defined services management framework that takes the guesswork out of operating an enterprise environment for maximum operational efficiency. A fully defined project delivery methodology decreases the risks associated with transitioning to the NetApp solution.

APPENDIX: DEFINITIONS

- **Data center efficiency**. Efficiency achieved by looking holistically at the data center, including people, processes, and technology, to achieve higher application agility and flexibility. The net benefits of these activities are operational efficiency, reduced complexity, improved IT staff productivity, and improved storage efficiency.
- **Dynamic data center.** The dynamic data center (DDC) is one that leverages a utility computing model, centralized resource management, and rapid and flexible resource allocation to support a wide range of data center applications. This includes the ability to rapidly commission and decommission applications, provision (and reprovision) required resources, and nondisruptively migrate applications and data between resources to meet time-varying application service-level requirements. It also includes the ability to seamlessly expand the underlying infrastructure and/or retire older components while maintaining continuous application operations through nondisruptive migration.
- **Dynamic data services.** Dynamic data services (DDS) are a network and storage layer that forms the foundation of the DDC. DDS is grounded in an SOI architecture for data and storage layers, modular data services, and data management processes. These three components provide the foundation for the delivery of modular data services that improve agility, efficiency, and responsiveness.

  The DDS is based upon a common architecture and deployment methodology that is reused for virtualization, storage efficiency, storage management, and other data services. These services can be deployed as part of an end-to-end IT transformation or in a “bite-sized” modular fashion.
- **IT as a service.** A business model characterized by the delivery of data center solutions at an agreed level of service. Service characteristics may include performance and availability, visible capacity, processing consumed, RPO/RTO, security, and so on.
- **IT Information Library (ITIL)** is a globally recognized collection of best practices for information technology (IT) service management. ITIL provides businesses with a customizable framework of best practices to achieve quality service and overcome difficulties associated with the growth of IT systems.
- **Service-oriented architecture**. SOA is a business-centric IT approach that allows you to integrate your business using a set of linked, repeatable business tasks, or services. SOA helps users build composite applications, which are applications that draw upon functionality from multiple sources within and beyond the enterprise to support horizontal business processes.
- **Service-oriented infrastructure**. SOI includes all configurable infrastructure resources such as compute, storage, and networking hardware and software to support the running of applications. Consistent with the objectives for SOA, SOI facilitates the reuse and dynamic allocation of necessary infrastructure resources. The development of SOI solutions focuses around the service characteristics to be provided.
- **Storage efficiency**. The application of technology to solve the problem of minimizing the amount of storage used. The technologies include space efficient Snapshot copies, thin provisioning, RAID 6, thin cloning, thin replication, and deduplication. The main outcome is reduced power, space, and cooling load in the data center.
- **Unified fabric.** A unified fabric provides a high-bandwidth, single-wire approach in which all of a data center’s I/O activity can be consolidated, regardless of underlying network protocols and data formats. From a server and storage perspective, this will reduce the required number of host bus adapters, NICs, switches, and cables. A unified fabric will allow increased data center scalability with improved performance and fault tolerance. The first step toward a unified fabric is the ability to run Fibre Channel over Ethernet (FCoE).