

How to Deliver a Home Run

FOR VIRTUALIZED SERVER ENVIRONMENTS

WHITEPAPER | FEBRUARY 2010




BEDROCK[™]
TECHNOLOGY PARTNERS



Table of Contents

- A. Executive Summary:
THE GAME PLAN
- B. Cover Your Bases with
STORAGE SOLUTIONS FOR VIRTUALIZED
SERVER ENVIRONMENTS
- C. First Base:
DELIVER EFFECTIVE, EFFICIENT DATA
PROTECTION IN A SNAP
- D. Second Base:
IMPROVE STORAGE UTILIZATION WITH
FLEXIBLE, THIN PROVISIONING
- E. Third Base:
MANAGE REDUNDANCY WITH
DEDUPLICATION
- F. Make Your Virtualization Investment
A HOME RUN

ACCORDING TO BOTH GARTNER AND IDC, YOU NEED TO CONSIDER STORAGE AS PART OF THE SOLUTION IF YOU WANT TO **REALIZE THE PROMISE OF A TRULY VIRTUALIZED ENVIRONMENT.**



A. Executive Summary: **THE GAME PLAN**

Virtualization has Changed the Game Plan for the Corporate Data Center.

Is your Storage Line-Up Supporting a Win?

There have been few technology movements that can rival the impact of virtualization. Not only has it changed the economics of information technology forever, but virtualization has enabled data centers to quickly evolve into dynamic environments capable of easily meeting increasing business demands.

While server virtualization is the first step to creating a more modern data center, the only way to realize the full effect and benefits of your investment is to virtualize the whole stack—servers, desktops, network, and storage. With the increasing demands virtualization places on storage resources, addressing your storage infrastructure now can make a dramatic difference in the overall return on investment from your virtualization initiatives. In fact, savings from your server virtualization efforts will be eroded by increasing storage costs if you do not adapt your storage infrastructure to accommodate a virtualized environment. Industry experts now predict storage moving from a supportive and peripheral role to being an essential and fundamental pillar on which IT infrastructures will be architected.

This paper explores storage technologies for server virtualization offered by Bedrock Technology Partners using NetApp solutions, and why you will want to include storage as a key consideration in your virtualization initiatives. After reading this paper, you will have a solid understanding of the impact of storage in your virtualized data center and the benefits to be gained by evolving your infrastructure sooner rather than later.

B. Cover Your Bases with **STORAGE SOLUTIONS FOR VIRTUALIZED SERVER ENVIRONMENTS**

Virtual server environments inherently generate larger quantities of data. Simply put, more servers—including virtual machines—create more data and require greater storage resources to support applications, backup and recovery, testing operations and more. Typical storage solutions do not offer the flexibility, efficiency and capacity requirements that are needed to meet the shared storage demands of virtualized servers including:

- >> Online scalability
- >> Changing or increasing demands for virtual server deployments
- >> Ability to support concurrent server virtual workloads simultaneously
- >> Ability to match storage to meet the different virtual server demands
- >> Support virtual data centers locally and in remote locations

As a critical element to realizing the full benefit of your virtualization investment, your storage solution needs to match the flexibility and cost efficiency of your virtualized servers, while meeting the increased data protection and availability requirements of a virtualized environment. Bedrock Technology Partners offers NetApp storage solutions that accommodate data growth and deliver key technologies—including snapshot, thin provisioning and deduplication technologies—to help you cover your bases and realize up to 50% less storage in your virtualized environment, guaranteed.

C. First Base: **DELIVER EFFECTIVE, EFFICIENT DATA PROTECTION IN A SNAP**

There are many unique challenges to protect data and applications in a virtual server environment. In non-virtualized server environments, the loss of a physical server would impact only the applications running on that server. In a highly aggregated or consolidated environment, however, the loss of a physical server supporting many virtual machines would have a much more significant impact by affecting all the applications supported by the virtual servers. Additionally, the growing number of virtual machines puts a strain on backup and recovery windows and creates CPU and IO bottlenecks. Even with enough CPU to push a backup, limitations in the number of shared network or storage adapters available to virtual machines may degrade backup performance—resulting in backup operations that could extend beyond the organization's backup window. This makes architecting a data protection strategy of critical importance.

Another challenge is protecting the growing amount of structured and unstructured data—both in primary data centers and offsite locations, including branch and field offices. Multiple-location enterprises need to protect and quickly recover data in the event of natural and man-made disasters, user errors, and technology and application failures. They also need a space-efficient method of distributing data to and from remote locations. Without an effective data protection and distribution strategy, business operations can be brought to a standstill.

Snapshot is an industry term representing the ability to record the state of a storage device at any point in time and preserve that 'snapshot' as a guide for restoring the storage device in the

BEDROCK AND NETAPP CAN HELP TO **REDUCE YOUR STORAGE REQUIREMENTS IN VIRTUALIZED ENVIRONMENTS BY 50%**, COMPARED TO TRADITIONAL APPROACHES.



event that it fails. Typically, a snapshot copy is done instantly and made available for use by other applications—such as data protection, data analysis and reporting, and data replication applications. The original copy of the data continues to be available to the applications without interruption, while the snapshot copy is used to perform other functions on the data. Snapshots provide an excellent means of data protection and enable better application availability, faster recovery, easier backup management of large volumes of data, reduced exposure to data loss, virtual elimination of backup windows and lower total cost of ownership (TCO).

NetApp Snapshot technology is a feature of the WAFL® (Write Anywhere File Layout) storage virtualization technology that is a part of Data ONTAP™—the microkernel that ships with every NetApp storage system. Created in less than one second, this ‘frozen’ read-only view of a WAFL volume provides easy access to old versions of files, directory hierarchies, and/or LUNs (logical unit numbers). Once a NetApp Snapshot has been taken, changes to data objects are reflected in updates to the current version of the objects, as if the Snapshot copy did not exist. Meanwhile, the Snapshot versions of the data remain completely stable. Performance is not impacted and you get more frequent recovery points with less space, continuous availability with transparent failover, and fast and efficient data replication for disaster recovery. Incurring no performance overhead, users can comfortably store up to 255 Snapshot copies per WAFL volume, all of which are accessible as read-only and online versions of the data. NetApp solutions leveraging snapshot technology include NetApp FlexClone, NetApp SnapMirror and NetApp SnapVault.

Using a technique sometimes referred to as “writable” Snapshot copies, **NetApp FlexClone**® instantly replicates data volumes and data sets as transparent, virtual copies—without requiring additional storage space—increasing productivity and saving storage space without compromising performance. This is accomplished by creating a virtual copy of the primary data set and storing only the data changes between a parent volume and a clone. With FlexClone, all unchanged data remains on primary storage and is used by both the primary application and the secondary clone copy. Multiple clone copies can be created from a single primary data set, enabling users to perform multiple simulations and to compare the characteristics of each data set after the simulations are complete.


By replicating data at high speeds over a LAN or a WAN, **NetApp SnapMirror**® software not only provides disaster recovery protection for your business-critical data, but also enables your

DR site for other business activities—so you can turn your disaster recovery solution into a business accelerator. SnapMirror can meet a full range of recovery point objectives from zero seconds to hours, and the recovery process supports low recovery time objectives, so you can get back to business faster. At the core of SnapMirror software is its storage-efficient design. First, a baseline mirror is performed between the SnapMirror source and SnapMirror destination systems. Next, at user-defined intervals, Snapshot copies are taken of the source system and only the new and changed blocks are sent incrementally over the network to the destination system. When the data is received at the destination, the changed blocks are merged with the existing data blocks, resulting in a full mirror copy of the source system. By replicating only the data that has changed since the last Snapshot copy, SnapMirror significantly reduces network bandwidth requirements by as much as 70%. Data deduplication facilitates further space reduction at the source system, at the destination system, and during the data transfer between the two systems.

With NetApp SnapVault® software, a complete copy of the primary data is stored on the SnapVault backup system. This initial, or baseline, transfer is much like a level-zero backup to tape. Each subsequent backup, however, occurs at the block level—copying only the data blocks that have changed since the last backup, not entire files. This enables backups to run more frequently and utilize less capacity because no redundant data must be moved or stored. To the user, however, each backup session is virtualized to appear as though it were a complete, level-zero backup copy, greatly simplifying the process of data restoration.

For example, if SnapVault backed up a 100GB primary data set for the first time, it would consume 100GB of disk space on the SnapVault system. Over the course of several hours, suppose that users change 10GB of data on the primary file system. When the next SnapVault backup occurs, SnapVault writes the 10GB of changes to the SnapVault system and creates a new Snapshot copy. At this point, the SnapVault system contains two Snapshot copies; one contains an image of the file system as it appeared when the baseline backup occurred and the other contains an image of the file system as it appeared when the incremental backup occurred. The copies consume a combined total of 110GB of space on the SnapVault system, but they are the equivalent of two complete 100GB backup copies. System administrators can refer to any of the backup instances to retrieve their files and can easily store dozens of backup images in a reduced amount of space.

BY SOME ESTIMATES, AN **AVERAGE 60%**
OF PRIMARY DISK STORAGE REMAINS
UNUSED SIMPLY BECAUSE OF OVERPROVISIONING.



SnapVault can improve backup and recovery speeds by up to 92% and, like SnapMirror, data deduplication can be combined with it for additional savings. The result is a dramatic reduction in the physical storage requirement for disk-to-disk backups. Benefits include fast and easy restoration of files from disk and the option to retain backups on disk for longer periods of time because each subsequent backup requires very little disk space.

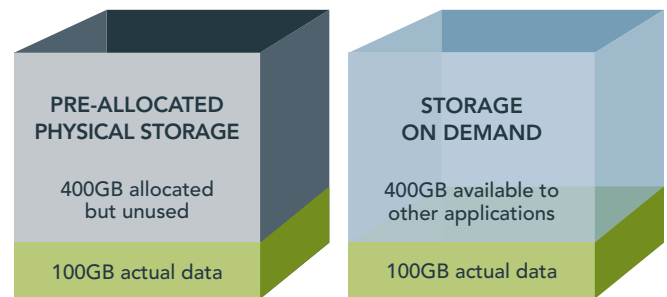
D. Second Base: **IMPROVE STORAGE UTILIZATION WITH FLEXIBLE, THIN PROVISIONING**

Traditional storage provisioning is slow and consumes excessive amounts of space. By some estimates, an average 60% of primary disk storage remains unused simply because of this type of overprovisioning. In a virtualized server environment, provisioning needs to support the speed of deployed virtual machines and also needs to be flexible to meet continuously changing storage demands. Additionally, skyrocketing storage requirements are placing an increased focus on cost control, increased utilization and simplified management in an effort to eliminate storage waste.

Thin Provisioning is an approach to managing storage where instead of allocating a fixed amount of storage, capacity is removed from the available pool only when data is written to disk. With thin provisioning, the storage administrator creates volumes for different users/applications to use, but does not actually allocate any space upfront to those volumes. The unused aggregate space is available for the existing thinly provisioned volumes to expand or for use in creation of new volumes. Once the aggregate free space declines to a predefined threshold, the storage administrator

can attempt to reclaim space from volumes that have obsolete data; can delete volumes that may be no longer in use; or can purchase more storage. By allowing as-needed provisioning and space reclamation, thin provisioning can result in storage utilization rates that are 2 to 3 times higher than with traditional provisioning, and efficiency can be automatically driven up towards 100% with very little administrative overhead. Organizations can purchase less storage capacity up front, defer storage capacity upgrades in line with actual business usage, and save the operating costs associated with keeping unused disk capacity available.

THIN PROVISIONING



TRADITIONAL PROVISIONING

THIN PROVISIONING

Thin provisioning is really “just-in-time” storage, because physical storage is fully allocated only when it is truly needed by the application. The rest of the storage remains in the “free pool,” to be used on an as-needed basis. Storage administrators can set thresholds that trigger alerts when more disks should be added to the pool.

With **NetApp FlexVol®** thin provisioning, users can create flexible volumes that appear to be a certain size but are in reality much smaller physically and can be managed and moved independently from physical storage—enabling you to respond to changing storage needs fast, with little effort and no disruptions. FlexVol pools physical storage, so you can create and resize virtual volumes as your application needs change—improving, and even doubling, the utilization of your existing storage and saving the expense of acquiring more disk space. You can also improve I/O performance and reduce bottlenecks by distributing volumes across all available disk drives.

NetApp storage systems make it possible to oversubscribe free space and adapt rapidly to the changing needs of the enterprise by leveraging thin provisioning. This is essentially an allocate-on-demand model that allows distribution of storage to application needs, whether or not all of the provisioned capacity is actually installed up front. The ability to present more storage space to the hosts or servers connecting to the storage system than is actually available, makes it possible to defer storage purchases until real application capacity thresholds are realized, resulting in dramatic cost savings.

E. Third Base: **MANAGE REDUNDANCY WITH DEDUPLICATION**


System administrators often feel helpless as they watch their data volume continuously grow larger. Once data is created, it seems to multiply faster than it can be controlled. Redundancy is a root cause of data proliferation and, in a virtualized environment, server data takes up substantial amounts of primary storage space. The average enterprise disk volume contains thousands or even millions of duplicate data objects. As these objects are modified, distributed, backed up, and archived, duplicate data objects begin to proliferate at an alarming rate. The result is inefficient use of storage resources. Server virtualization projects are routinely jeopardized by inefficiencies in storage.

Data deduplication is a proven way to eliminate redundancy to reduce the amount of data that has to be backed up or replicated. Duplicates are removed and new files are stored at each backup, resulting in faster backups by cutting down the amount of information processed, transferred and stored. Only one unique instance of the data is actually retained on storage media and redundant data is replaced with a pointer to the unique data copy. For example, an email system might contain 100 instances of the

same one megabyte (MB) file attachment. If the email platform is backed up or archived, all 100 instances are saved—requiring 100 MB storage space. With data deduplication, each subsequent instance is just referenced back to the one saved copy. In this example, a 100 MB storage requirement could easily be reduced to only one MB. Highly redundant virtual servers can be deduped freeing up storage capacity, allowing you to add additional virtual servers for greater scalability.

NetApp Deduplication® has taken the powerful benefits of deduplication benefits even further by enabling it to be implemented across a wide variety of applications and file types, including data backup, data archiving, and unstructured data volumes. NetApp deduplication divides newly stored data objects into small blocks. Each block of data has a digital “signature,” which is compared to all other signatures in the data volume. If an exact block match exists, the duplicate block is discarded and its disk space is reclaimed—resulting in a dramatic savings in storage requirements. A 2009 study by Oliver Wyman reported that NetApp deduplication customers have decreased their upfront storage requirements by 54% and achieved space savings of anywhere from 40% to 70% at no additional cost.

DEDUPLICATION



PERVASIVE DEDUPLICATION

- + 20:1 reduction in secondary
- + 30-70% reduction in primary
- + Capability, not a platform

With NetApp’s unique deduplication technology, customers are able to eliminate duplicate data across all tiers of storage, allowing them to combat data proliferation while creating a more efficient storage infrastructure.

F. Make your Virtualization Investment

A HOME RUN

As an experienced solution provider specializing in architecting, implementing and supporting IT solutions, Bedrock Technology Partners has helped Fortune 1000 corporations, large and small enterprises, government agencies and educational institutions achieve more from their technology investments. Our broad technical expertise, combined with comprehensive services and in-house technology labs, enable us to pinpoint the right solution for any requirement.

Bedrock Technology Partners simplifies storage challenges by providing proven, centralized and automated data management solutions featuring the industry's most innovative technologies. We offer storage solutions with a full range of efficiency capabilities—including, deduplication in primary and secondary environments, thin provisioning, and snapshot technology—that help you deliver more value from your virtualization initiatives. Our solutions can be used individually or together to help you boost utilization of existing storage and reduce your storage requirements by at least 50% guaranteed.

When planning your data management objectives, partner with the dependable team at Bedrock for solutions that deliver the highest return on your investment and create a solid foundation for your continuing success.

CONTACT US TODAY for a
FREE storage evaluation and learn how Bedrock
can make your virtualization project a home run!



TOLL FREE

1.877.9BEDROCK

EMAIL

sales@bedrock-tech.com

WEB

www.bedrock-tech.com

