

HIGHLIGHTS

- » A 1500-Node Diskless Server Farm
- » NetApp's Approach to Deduplication
- » Case Study: Streamline Oracle® Dev
- » A Simple Tool for Remote Backup
- » Resources:
 - Demos: Data Protection and E-mail
 - Webcast: E-mail Archive/Compliance

Achieving 99.99% Uptime via Storage and Server Virtualization

Dale Smith, Director of Technical Services, Sentara

Sometimes availability really is a matter of life and death. How a recent Innovation Award winner uses VMware, NetApp controllers, and HP storage arrays.

[More »](#)



BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"Let me explain how our recent data deduplication announcement fits into our long-term strategy..."

- » Dave's Blog

DRILL DOWN

- » Webcast: NetApp E-mail Archive and Compliance Solutions
Eliminate the need for separate storage silos without sacrificing performance.
- » Interactive Demos: Exchange and Data Protection Manageability Tools
Preview the new NetApp Manageability Family demo and see how NetApp solves challenges you encounter every day.
- » Sneak Preview: A Practical Approach for Grid Computing
An in-depth look at the first viable example of a large-scale grid for commercial apps.



ADMIN RESOURCES

- » [visionapp Remote Desktop](#)
One user's approach to managing 150 remote Windows® servers, an Exchange 2003 infrastructure, and backups to NetApp storage.

TIPS FROM THE TRENCHES

A Case Study: Architecting Storage for Application Development and Test

Bikash R. Choudhury, NFS Product Engineer, NetApp

This major application developer streamlined the creation and management of thousands of dev/test environments using Data ONTAP® 7G and FlexClone® technology.



[More »](#)

How a 1500-Node Diskless Server Farm Evolved into a Fully Virtual Ecosystem

Gregg Ferguson, Manager, NetApp Engineering

NetApp Engineering clones LUNs in milliseconds and changes OSs in the time it takes to reboot a server, plus over the past two years its test lab has evolved to include:

- Booting over iSCSI, FC, and NFS
- Provisioning of virtual environments
- A unique cooling design

[More »](#)

ENGINEERING TALK

NetApp A-SIS: Deduplication Comes of Age

Blake Lewis, Technical Director, NetApp

A-SIS deduplication can help you reduce storage consumption up to 95%. A behind the scenes look at how this free technology leverages the NetApp WAFL® file system to conserve disk space while keeping system overhead low.

[More »](#)

FEEDBACK

TECH ONTAP ARCHIVE



Dale Smith

Director of Technical Services, Sentara

Dale Smith is responsible for computer technical services at Sentara, a healthcare organization which is nationally recognized for its use of innovative and leading edge technology and has been named to the '100 Most Wired' Healthcare systems for multiple years. Dale leverages over 22 years of Information Technology experience to lead teams focused on computer infrastructure, Operations, Help Desk, and Enterprise architect functions. His responsibilities also include implementing long-term system strategies in conjunction with corporate goals and objectives.



Leveraging Technology to Deliver Superior Care: How Sentara Achieved 99.99% Availability via Storage and Server Virtualization

By Dale Smith, Sentara Healthcare

Sentara serves approximately two million patients across Virginia and North Carolina. In most cases, we are focusing on eliminating paper records and keeping everything from medical histories to patient exams in and outside the hospital online. This dramatically improves care efficiency, but has brought about a paradigm shift in how the hospital—and everyone in our organization—views IT.

At Sentara, the cost of downtime is no longer measured in dollars per minute but in potential lives lost. Even though downtime procedures are a part of the clinical workflow, as patient care providers become more dependent on technology the requirements for non-disruptive technology are imperative. Our internal threshold used to be that no patient care system could be down for more than four hours. In reality, we are finding that threshold really needs to be minutes (or much less!) because every bit of patient information is electronic. The customer mindset has changed from some downtime being acceptable to any downtime, either planned or unplanned, being completely unacceptable.



Our internal threshold used to be that no patient care system could be down for more than four hours. In reality, we are finding that threshold really needs to be minutes (or much less!) because every bit of patient information is electronic. The customer mindset has changed from some downtime being acceptable to any downtime, either planned or unplanned, being completely unacceptable.

Overview: eCare Initiative Expansion Drives Storage Evolution

In 2005, we launched the Sentara eCare™ Health Network, a comprehensive electronic medical record (EMR) system linking clinical information with scheduling, billing, and registration data over a secure network. This technology allows the secure sharing of patient information between hospitals, physician offices, diagnostic centers, and patients' homes. When fully implemented, the system will enable patients to book their own appointments from home and hospitals to transfer full medical records to other states or countries as needed. It has also facilitated smart prescription centers that automatically check medication allergies and conflicts and can even make recommendations about potential treatments.

As part of the eCare initiative, our team originally implemented a pair of HP StorageWorks XP12000 disk arrays to support UNIX®-based hosts. However, as the implementation expanded we needed to deliver the same class of availability to Windows® hosts while simplifying provisioning and other storage-management processes.

RELATED INFORMATION

- [NetApp VMware solutions](#)
- [Five Ways to Use NetApp Snapshot Copies in VMware Environments](#)
- [Simplifying Data Management for EMC and Other Storage Arrays](#)

About Sentara

Sentara Healthcare is a not-for-profit integrated healthcare provider serving more than two million residents in southeastern Virginia and northeastern North Carolina.

Sentara is comprised of 7 acute care hospitals; health plans with 300,000 covered lives, 7 nursing centers, 3 assisted living centers, and 235 physician medical group. Sentara was ranked the number three integrated health care network in the United States in 2005 by Modern Healthcare magazine and is the only health care system in the nation to be named in the top 10 for seven consecutive years.

Sentara has a total of 4PB of data.

V-Series: Simplifying Data Management for Existing Storage Arrays

The NetApp V-Series product line is unique in its ability to unify block and file access to storage array products from Hitachi, HP, IBM, Fujitsu, and—most recently—EMC.

V-Series systems take the back-end

After a thorough review of options, our team chose NetApp storage virtualization technology. This ended up being only the first in a three stage deployment that has helped revolutionize our storage environment:

- Stage #1: Virtualize existing storage arrays for multi-OS and multi-protocol access
- Stage #2: Maximize data availability
- Stage #3: Consolidate servers

As a result, we've not only streamlined backups and significantly expanded our environment without increasing costs but—most important—we have achieved our 99.99% uptime goal.

Stage #1: Virtualize Existing Storage with NetApp V-Series

Every dollar spent on IT is a dollar that can't be spent on new types of medical training or technology. Any IT investment is therefore under extreme scrutiny and must provide a maximum return. We chose XP to support mission-critical solutions on HP Superdome Servers, but required a more flexible solution for the Windows environments. In 2005, we were in a situation in which our XP arrays were supporting our UNIX environment, but we had no budget to build a parallel infrastructure and no simple path to expand this environment to support Windows systems.

The Sentara SAN team initiated an evaluation of the NetApp V-Series system. (see sidebar) Basically V-Series provides virtualization for pre-existing storage arrays by taking functions that would happen either on the host or the array and pushing those functions to the virtualization layer.

We were blown away by the early results, and extended our original 90-day proof-of-concept operation to a 120-day in-depth evaluation of the NetApp storage virtualization solution.

Ultimately, there were three key reasons we chose the NetApp solution:

1. Ability to present (via FCP and iSCSI) the existing XP12000 arrays as storage to our Windows hosts while simplifying provisioning and enhancing data protection
2. Ability to streamline backups and provide rapid restores via integrated tools
3. Flexible ways to replicate data via existing IP infrastructure

We deployed clustered V3050 controllers and connected them to our existing HP StorageWorks XP12000 disk arrays. The benefits of this deployment have been considerable.

First, storage virtualization gives us the full range of NetApp capabilities—including Snapshot™, FlexVol®, and FlexClone®—in conjunction with our existing storage arrays. With a relatively modest additional investment, we were able to leverage our existing infrastructure to not only support our Windows environment but also gain substantial flexibility.

Multi-protocol support has been a substantial advantage. While our primary Microsoft® SQL Server™ applications are on a Fibre Channel SAN, the same NetApp V-Series clustered system also provides NAS connectivity using the CIFS protocol to a repository of some 30 million TIFF images in our OnBase document management system from Hyland Software.

Additionally, if we connect a host running Microsoft SQL Server directly to a storage array, we would typically have to take a database offline to do backup. With our V-Series solution, we use NetApp SnapManager® for SQL Server, which coordinates activities between the host and storage to create consistent, point-in-time Snapshot copies of a database that can then be backed up without significantly interfering with ongoing database activity.

We also utilize a capability called NetApp FlexClone to clone existing databases for test and development and a variety of other purposes. FlexClone clones can be created in seconds and only consume additional storage space as changes are made to the cloned volumes, so they are both time and space efficient compared to traditional cloning methods.

The FlexVol volumes that the V-Series presents to hosts are spread across all the

LUNs presented by existing storage arrays and group them together so they can be more flexibly managed and utilized.

Essentially, the V-Series creates a single large storage pool (analogous to a standard NetApp aggregate) out of a collection of storage array LUNs. The back-end storage continues to provide RAID protection while the V-Series virtualizes back-end LUNs so you can create flexible volumes (FlexVol volumes) that span multiple LUNs without having to worry about underlying storage layout.

V-Series systems are identical to NetApp FAS storage in supporting:

- All common front-end protocols, including CIFS, NFS, iSCSI, and FCP
- NetApp Snapshot, SnapMirror®, SnapVault®, SyncMirror®, and NDMP software for data protection and DR
- All host-side software, including SnapDrive®, and the SnapManager suite of products for application support

Learn more. Read a recent [Engineering Talk Article on the V-Series](#).

NetApp Snapshot Copies and VMware

NetApp Snapshot technology is ideally suited for use with VMware. A recent Tech OnTap article explored five uses of Snapshot and other NetApp technologies derived from it in VMware environments:

- Near-instantaneous VM backup
- Fast and flexible VM recovery
- Accelerated data management through cloning
- Disaster recovery
- Application backup and management

Learn more. Read [Five Ways to Use NetApp Snapshot Copies in VMware Environments](#).

disks in the back-end storage pool for greater performance. On one very large database we have to do a re-index on a particular table monthly or quarterly. This took 24 hours on the old storage infrastructure. When we moved it to NetApp we were surprised (and pleased) to find that the same procedure took a much more manageable four hours to complete—a 6X improvement.

While we haven't analyzed all possible reasons for the speed-up, it is likely that it results from spreading out the activity across more disks in the back-end storage arrays. Our use of another NetApp product, NetApp SyncMirror®, to ensure availability probably also contributes to this performance improvement. Which leads us to...

Stage #2: Ensure Availability with SyncMirror

Sentara's availability goal is 99.99% uptime. By combining a virtualized storage infrastructure with NetApp SyncMirror synchronous mirroring software we successfully achieved this goal for the past 12 months. With SyncMirror, our clustered V3050 controllers maintain two fully consistent copies of all critical data, one on each XP12000 array. This protects us against all types of hardware outages including multiple disk failures or the failure of an entire array.

SyncMirror improves the performance of random disk read operations using simultaneous round-robin reads from both mirrored copies of data. This results in up to 80% improved random read performance, which is particularly beneficial for database environments. The substantial re-indexing speed-up discussed above probably resulted in part from SyncMirror's improved read performance.

Stage #3: Consolidate Servers with VMware VI3 and FAS Storage

About four months after the initial implementation we purchased a NetApp FAS3020c to provide back-end storage for our server consolidation efforts using VMware. Our initial experiences with NetApp had been extremely positive, plus we knew that we wanted to use iSCSI for connectivity within this environment.

iSCSI has been a perfect fit—it offers relatively simple management and low-cost Ethernet networking. (We used Fibre Channel rather than iSCSI with the V-Series deployment because we were inserting it as a virtualization layer within an existing FC SAN.)


Because the FAS3020c is clustered, we can perform rolling upgrades on each head to avoid downtime during upgrades. Through the combination of NetApp FAS storage systems, VMotion, and VMware High Availability (VMware HA) we have been able to maintain service levels above four nines (>99.99% availability) for the past 12 months.

So far, we've consolidated about 30% of our server environment. We've replaced about 190 physical servers with direct attached storage (DAS), with virtual servers using NetApp storage. That's resulted in about \$700K in savings over the past four months. That number includes savings from a wide variety of factors including reductions in air conditioning, power, administration costs, etc. It does not include cost reductions due to reduced downtime, which is a significant benefit we get from NetApp. As a result of this effort, we've cut the deployment time for new servers from four weeks down to two weeks.

Availability and Flexibility on the Critical Path to Success

The digital revolution in the healthcare industry has put IT infrastructure on the critical path for patient care. Application and data availability have now become critical to healthcare delivery. Those of us tasked with creating the IT infrastructure to support leading-edge healthcare facilities are focused on delivering the highest possible availability with the flexibility to adapt to rapid change.

For Sentara, NetApp has become a key partner in achieving these goals. It may have been possible to do what we've done without using NetApp technology, but NetApp has made it simple and relatively painless. Within two days of installing the V-Series, we were presenting storage for use by hosts, and the availability we've achieved speaks for itself.



The NetApp V-Series has allowed us to better utilize the storage we already have, protecting that investment. Both the V-Series and our NetApp FAS systems give us the ability to support more platforms using more protocols while providing new capabilities such as FlexVol, FlexClone, SnapManager, and SyncMirror. That flexibility will be critical going forward.

Sentara's core mission is to improve health every day and lead the community by providing the best healthcare possible. This requires technology that is always available—obviously no small feat, but a goal NetApp is helping us achieve.

TECH ONTAP ARCHIVE



Bikash R. Choudhury

NFS product partner engineer, NetApp

Bikash started at NetApp seven years ago as a technical support engineer, and later became the technical global advisor (TGA) for one of NetApp's largest customers. As a TGA he focused on building strong customer relationships as well as technical solutions. In his current role as an NFS product partner engineer, Bikash spends his time doing functional testing and certifications, documenting best practices, and providing architectural and configuration recommendations to customers.

A Case Study: Architecting Storage for Application Development and Testing

By Bikash Choudhury

Whether you're creating commercial applications or doing in-house development work, ensuring the effectiveness of your development and testing efforts is critical. As applications continue to grow in size and complexity, the simple act of provisioning development and test environments can consume vast amounts of storage and become a major bottleneck.

In this article, I'll look at the methods that one customer uses to create fast and reliable replicas of the master Oracle® Databases and Applications in its development environment.

Overview: Customer Environment

This customer is a major developer of business applications running on Oracle Databases. Its development services organization supports thousands of engineers all over the world. To enhance productivity and efficiency, the IT team has centralized shared services. Having all development resources and tools in a central location accessible from anywhere in the world 24 x 7 has proven to be cost effective, scalable, and easier to manage than a distributed environment.

A master environment consisting of the required version of the Oracle Database plus Applications has been created to serve as a template. A special tool manages and configures replicas of the original template for specific development and testing requirements. The method for creating these replicas has evolved over time to accommodate growth and increase efficiency.

Challenge: Traditional Replication Approach Lacks Scalability

Historically, the master development environment lived either on a Fibre Channel SAN or on direct attached storage (DAS) systems and was 150–200GB in size. Developers used replicas of the master environment to validate their work. These replicas were created using basic UNIX® copy commands such as "rsync" and "cp." Custom scripts could be run against the replicas to further configure them to each developer's requirements.

As the size of the replicas continued to grow, however, existing storage solutions were not scaling to meet developer requirements.

Backup and retention of configured replicas




Figure 1) Traditional replication approach: the master volume is completely copied for each individual test and development

RELATED INFORMATION

- [Data ONTAP 7G for Database Applications \(pdf\)](#)
- [Optimizing Oracle on NFS\(pdf\)](#)
- [Best Practices for Oracle \(pdf\)](#)
- [Webcast: Data Center Solutions for Oracle](#)

Empowering DBAs with SnapManager® for Oracle

SnapManager for Oracle (SMO) integrates closely with the Oracle Database, providing features that allow database administrators (DBAs) to perform a variety of storage-related tasks and busy storage admins to save time and effort.

In a recent article, NetApp database and grid architect Alvin Richards took an in-depth look at the inner workings of SMO, including:

- Snapshot copies
- Full and partial restores
- Cloning
- ASM integration

Read the article to find out more. You can also [watch a demo](#) of SnapManager for Oracle in action.

Data Center Solutions for Oracle

Managing critical tasks such as backups, restores, and provisioning in Oracle Database environments is becoming more time consuming and costly than ever. A recent TechTalk Webcast highlights how to maximize

were also becoming difficult. Because multiple storage products were deployed and multiple backup applications were in use, the process was difficult to manage, and tape-based solutions weren't keeping up with data growth. Backups were taking too long and restores were unreliable.

environment.

Intermediate Solution: Move to NFS

To address the limitations of this traditional approach, the development services organization decided to move to NFS with a network-attached storage (NAS) back end.

This decision was primarily made to reduce the complexity of deploying and managing storage. Using NFS allowed the IT team to speed deployment and simplify support for a large number of replicated development environments.

With data growth exceeding 67% per year, NFS in a NAS environment proved highly cost effective. Ethernet-based NAS was much easier to understand and implement, and it was already compatible with most of the existing infrastructure. NetApp NAS, which has been proven in hundreds of Oracle Database deployments, met the team's performance requirements of response times under 15 milliseconds, plus gave the IT team multi-host file system access, file-level data sharing, and the necessary level of security.

Clustered NetApp storage systems were chosen for this deployment to ensure high availability. Existing data was migrated to partitions called qtrees residing in large storage volumes on the NetApp systems. The master development environment was staged inside a single qtree. Copies of the master environment could then be made into other qtrees in the same volume, to other volumes on the same storage system, or to another storage system. A tool was created to copy the master environment and custom configure the replica. This tool uses NDMPCOPY in the background to replicate the master environment based on developers' requests.

Backing up the master and test databases in hot backup mode was made possible using NetApp Snapshot™ copies, which copy only incremental data and make backup and recovery fast and seamless. Using this approach, it is possible to recover a 4TB database in a matter of minutes and keep several hundred Snapshot copies (and therefore hundreds of recovery points) online. Check [The Core NetApp DNA](#) for more information about the unique NetApp approach to snapshots.

Despite the improvements, this infrastructure still had significant caveats:

- **The copy process took hours to complete**, and sometimes the IT team would run out of space while replicating the master environment because automatic space provisioning was not intuitive.
- **Ensuring that each developer had the right setup at the right time remained a challenge**, especially since the company had hundreds of developers worldwide and this number had doubled over the past 5 years.
- **Performance issues** (mostly disk bottlenecks) were introduced by maintaining replicas of the master in different partitions in the same volume.
- **Developers had to work through the IT team** if they needed to roll back to a previous configuration and start again. If a developer had to refresh a test environment, another NDMPCOPY process had to be initiated.
- **Deleting partitions that were no longer needed was a long and cumbersome process**, and the IT team was completing multiple refresh cycles every month as tests were completed.

Final Solution: Space-Efficient Cloning

With the release of the Data ONTAP® 7G operating system, NetApp introduced the concept of aggregates with flexible volumes (FlexVol® volumes) and flexible copies (FlexClone® copies). If you're unfamiliar with NetApp technology, this basically allows automatic space provisioning and volume resizing on the fly.

FlexVol volumes and FlexClone copies are contained inside an aggregate, which consists of a large number of physical disks grouped into RAID groups using the NetApp dual-parity RAID implementation: RAID-DP. RAID-DP is an advanced, cost-effective failure/error protection solution that protects against double disk failure within a single RAID group.

Using FlexVol, storage admins can grow and shrink volumes at will as long as the aggregate has enough space. Because aggregates contain a large number of disk spindles, performance bottlenecks are greatly reduced; even the smallest FlexVol volume within an aggregate is spread across all spindles. Finally, Data ONTAP 7G provides ZAPIs that help administrators build tools to automate the creation of

the return on investment in your database environment by:

- Managing initial purchase price
- Maximizing operational efficiency
- Reducing downtime

[View the Webcast](#) to learn more.

The Core NetApp DNA

In the early 1990s, Network Appliance revolutionized storage networking with a simple architecture that relied on NVRAM, integrated consistency points, and a unique file system to do things that the file servers of the time could not.

These basic building blocks of NetApp technology support:

- Multiprotocol environments (NFS, CIFS, FC, iSCSI, and so on)
- Clustered failover, mirroring, and disk-to-disk backup
- RAID-DP™ and other software-based resiliency features
- The near-instantaneous creation of writable clones
- Block deduplication using A-SIS

Learn about [The Core NetApp DNA](#).

development and test environments using these basic features. (See [TR 3373](#) for more information on using Data ONTAP 7G with database applications.)

As shown in the table below, these capabilities offered an elegant solution to the issues the IT team was still experiencing.

Customer Challenge	Data ONTAP 7G Advantage
- Time-consuming copies - Ensuring developer has right setup - Performance issues - Developers can't save point-in-time images - Time-consuming refresh cycles	- FlexClone makes copying unnecessary - ZAPIs integrate storage with custom tools; FlexClone accelerates up-replication - FlexVol spreads all volumes across a large number of spindles - Developers can make Snapshot copies of individual FlexClone copies for easy rollback - FlexVol volumes can be expanded, contracted, or deleted instantaneously; space can be provisioned on demand

To implement the new solution, the master environment was staged in a FlexVol inside an aggregate with a large pool of physical disks. A master Snapshot copy of the FlexVol supports the creation of copies or clones.

Upon request the Snapshot copy of the master environment can be cloned instantaneously using FlexClone. The cloned volumes that are created only consume space incrementally as changes are made to the original image.

This saves a significant amount of storage space. The overall architecture offers the following additional advantages:

- Test environments are now refreshed quickly. When a user is done with an environment, he simply destroys the clone and requests another copy of the master and is back to work in minutes.
- IT queries have been reduced more than 50% now that developers can create Snapshot copies of their assigned clone volumes to retain point-in-time images of the test environment as they make changes and revert to any point if necessary.
- Disk-to-disk backup provides a centralized, simple-to-manage solution that ensures backup, disaster recovery, and retention. File recovery is fast, and data is accessible from backup if primary storage is out. Overall, downtime has been reduced 80% – since upgrading to Data ONTAP 7.2.2, the team has had zero downtime except for scheduled maintenance – as well as significant cost savings.
- If any files are accidentally deleted in a cloned volume, the developer uses Single File SnapRestore® to retrieve the file without wasting a lot of time. This feature has eliminated about five helpdesk requests per developer per month, and has helped improve time-to-market.




Figure 2) Using FlexClone, the master environment can be instantaneously cloned as often as necessary. Additional storage is only consumed as data is changed.



Figure 3) Typical testing environment with FlexClone.

Impact

Moving from traditional storage to NetApp helped this customer scale to better meet its business needs while consolidating storage to reduce its data center footprint. With Data ONTAP 7G and NetApp NAS storage running NFS the customer can quickly configure test, development, maintenance, and staging environments and accelerate test cycles for improved quality and faster time-to-market.

TECH ONTAP ARCHIVE



Gregg Ferguson

Kilo-Client manager, NetApp Engineering Support

A 22-year industry veteran, Gregg held a variety of engineering, systems administrator, and IT manager positions before joining NetApp in 2000. After helping found the original NetApp North Carolina facility, Gregg spent over four years in the field as a systems engineer before returning to the NetApp Engineering IT team. He leveraged his years of experience with customer environments to conceive and architect the Kilo-Client test lab. Today, he manages a team of five that supports an average of 12 to 14 virtual environments and manages 1,498 blades, 7102 fabric ports, and 87 storage controllers representing 1038TB of storage.

How a 1,500-Node Diskless Server Farm Evolved into a Fully Virtual Ecosystem

By Gregg Ferguson

When we first got the idea for our engineering test lab in Research Triangle Park, it was in response to a growing need within NetApp to be able to test our products against large grids or server farms and to quickly reproduce any customer problems that might occur in such environments. Our original plan was to use server blades in which each blade booted from a local disk. However, as the project progressed, it became clear that the time and administrative overhead required to copy a boot image to a thousand local disks would result in more time spent configuring and managing the cluster than running actual tests.

Instead, we designed our test lab to include 1,120 server blades booting over iSCSI. We dubbed the lab Kilo-Client, and believe it was the largest iSCSI-based diskless server farm in the world when it was launched in 2005 (and may still be!). We later added an additional 98 blades with iSCSI HBAs and 280 blades capable of booting over Fibre Channel. Check out the sidebar for specific hardware and software components.


The result: a 1,500-node server farm that packs massive performance and flexibility into a footprint of just over 389 square feet.

While the motivation for Kilo-Client today remains more or less the same as it was in the beginning, the lab has evolved to keep up with emerging technologies. Plus, over the past two years we've learned a lot about operating and maintaining a big environment. This article focuses on aspects of the current test lab design that customers and partners tend to find most interesting, including:

- Rapid server provisioning
- Provisioning virtual environments
- Booting over iSCSI, FC, and NFS
- iSCSI over 10 gigabit Ethernet
- Management automation
- Thin provisioning
- Data center cooling

Before I dive in, however, I want to make it clear that the Kilo-Client builds on just five or six different technologies, each of which is currently in use at hundreds of NetApp customer sites. Creating the architecture was mostly a matter of pulling together all those elements—each of which I had been exposed to during my years as a NetApp

Sneak Peek: A Practical Approach for Grid Computing



RELATED INFORMATION

- [Sneak Preview: A Practical Application for Grid Computing \(pdf\)](#)
- [Kilo-Client Case Study \(pdf\)](#)

A Practical Application for Grid Computing

Enterprise Management Associates, Inc. (EMA), recently took a fresh look at the Kilo-Client from an outsider's perspective. Because it provides technology guidance and research to Fortune 1000 companies, EMA is intimately familiar with large-scale IT, and describes the Kilo-Client as "probably the world's largest commercial grid."

This white paper focuses on a number of aspects of the Kilo-Client design, including:

- Use of standard components
- "Podule" architecture
- Automated provisioning
- Centralized monitoring and management
- "Greening" the grid

A final perspective measures the Kilo-Client architecture against modern IT requirements and concludes that "Any data center staggering under the load of servicing multiple shifting workloads... on multiple disconnected platforms will have to pay attention."

Read [A Practical Application for Grid Computing](#).

SE—in a single infrastructure.

In short: there is absolutely nothing in our test lab that can't be leveraged by any NetApp customer.

Rapid Server Provisioning

One of our early goals was to quickly provision a compute grid capable of meeting specific test characteristics. This meant that servers had to be quickly booted with any OS/application environment. We solved this problem using NetApp FlexClone® technology to enable the rapid creation of system images without making full physical copies of those images.

A set of "golden" boot images is created (as iSCSI and Fibre Channel SAN LUNs) for each operating system and application stack required in the server farm. Using SnapMirror® and FlexClone, we can quickly reproduce hundreds of clones (a FlexClone clone for each server being configured for a test); only host-specific "personalization" needs to be added to the core image for each provisioned server. This unique approach affords near-instantaneous image provisioning with a near-zero footprint (only the blocks of the images that differ need to be added to the storage system, which keeps track of the individual images), enabling us to configure and boot all or a subset of our nearly 1,500 blades in minutes.

Virtual Environments

What we ultimately discovered was that simply provisioning the server environment—although our method works extremely well—is not enough. What NetApp engineers wanted, and what we really needed to be able to do, was to rapidly provision complete virtual environments, including compute grids, interconnect fabrics, and storage grids.

That's exactly what we do today. We can automatically configure a compute grid running almost any OS (including VMware) and connect it via a vLAN (IP), vSAN (Fibre Channel), NFS, or even CIFS (we don't have the capability to boot over CIFS but can test CIFS functionality) to any of five possible storage grids. A typical virtual environment—which might include 100 servers, multiple OSs, and five to six storage controllers—is usually up and running in an hour or less. The most complex environment we ever created took about 10 hours to get up and running, and involved 500 servers, 30 NetApp FAS 6070s, 72 shelves of 300GB FC drives (~500TB), and the Data ONTAP® GX operating system.



Figure 1) A true virtual environment.

At any given time, our lab is running 12–15 virtual environments that are used for everything from product and interoperability testing to troubleshooting to proof-of-concept testing. Tests can be preempted by halting a server and creating a space-efficient, derived clone of that system (using FlexClone). Test configurations of any environment can be preserved or shared with other users and re-run months or years later, perhaps even on an alternative (albeit with the same architecture) system. Once we've built an environment, we never have to rebuild it. For example, say we build a Red Hat Linux® environment and the team requesting this environment loads Oracle 10g™. After they're done with the test they can create a clone, and this pre-configured environment can be reused as necessary in the future.

A final important point is that these virtual environments can be accessed and managed from anywhere in the world. A NetApp engineer working in any of six global facilities or a systems engineer at any location worldwide can schedule resources and run tests remotely.

Booting over iSCSI, FC, or NFS One of the unique differentiators of NetApp storage is the ability to support iSCSI, FC, and network storage (NFS and CIFS) all from a single storage platform. Lots of customers find it most efficient to deploy iSCSI for some applications and Fibre Channel SAN for others, and to support additional applications using network-attached storage. As a result, we face new challenges on a daily basis in our test lab, and having a highly flexible environment that can support just about any

Kilo-Client Overview

The NetApp engineering test lab dubbed Kilo-Client is divided into six "pmodules," groups of nodes sharing a common infrastructure consisting of an Ethernet switch, a NetApp 6070c (used as a boot appliance), and 16–20 server chassis. Currently, each podule accommodates from 224–280 nodes, although a podule may be any size an application requires.

Components of the Kilo-Client architecture include:

1,498 Blade Servers

- 1,218 blades with QLogic iSCSI hardware initiators
- 280 blades with QLogic Fibre Channel HBAs capable of SAN booting via FCP or iBoot

Operating Environments

- Images available for Windows®, SUSE Linux, Red Hat Linux, VMware ESX, Solaris™ 10 (other operating systems supported upon request)

Network Infrastructure

- 17 Cisco 4948 switches (boot infrastructure)
- 10 Cisco 7609 switches (test infrastructure)
- Multiple 10GbE switches for 7609s(test environments)

SAN Boot Storage

- 5 NetApp FAS980s and 1 FAS6070 cluster (active boot images; 1 per 224–252 blades)
- NetApp NearStore® (master boot images)

Test Virtual Environments

- 24 FAS6070s—Data ONTAP GX
- 20 FAS3050s—Data ONTAP GX
- 6 FAS3070s—Data ONTAP GX
- 4 FAS3050s—Data ONTAP GX
- 2 FAS6070s—Data ONTAP 7G
- 5 FAS3050s—Data ONTAP 7G
- 7 FAS980s—Data ONTAP 7G
- 2 FAS270 clusters (4 actual heads) for Data ONTAP 7G Testing

The test environment infrastructure includes additional F980s NetApp storage systems.

NetApp Shatters Industry Record for NFS

protocol we throw at it is a huge advantage.

The original Kilo-Client design allowed us to boot our server blades over iSCSI using hardware initiators (iSCSI HBAs). Today we can boot servers using any of four approaches:

- Over iSCSI using the hardware initiator (1,218 blades)
- Over iSCSI using the software initiator (entire environment)
- Over Fibre Channel using an FC HBA (280 blades)
- Over NFS (entire environment)

This allows us to test and compare various environments and booting methods. If we aren't specifically testing booting methods, we tailor our approach based on test requirements. For instance, if someone wants to perform Fibre Channel testing with fault injection, we would typically boot the servers being used for other tests over iSCSI or NFS to leave the Fibre Channel free for the testing.

iSCSI over 10GBE

A while back I was asked to do a presentation about the Kilo-Client design at an event sponsored by blade.org. After my talk, virtually every vendor at the show wanted to sell me his or her new technology for use with the Kilo-Client. I was even approached by one overenthusiastic salesman in the bathroom!

When I got home, I went through all the cards that had been forced on me and discovered that several were from vendors of 10 gigabit Ethernet gear. I called them up and ultimately we created a test kit using the IBM Blade Center with NetXen controllers connected to a NetApp cluster also outfitted with 10 gigabit Ethernet cards. The result was a configuration with 10 gigabit Ethernet from end to end that was capable of diskless booting using iSCSI. We took that configuration to an event in New Orleans where it generated a lot of interest and the hardware went on to shows in Paris and Singapore (although I did not).

So far, we've done mostly functional testing, but this architecture gives us the ability to do large-scale performance comparisons involving 10 gigabit Ethernet versus Fibre Channel—as well as anything else we might want to test.

Automated Configuration Management

When the Kilo-Client was created, we had a few scripts to help with configuration and that was about it. As we freely admitted at the time, that was the weakest element. Today, our work follows a predictable pattern of schedule→provision→monitor→adjust resources based on load→de-provision→re-schedule and so on.

We've got an automation framework in place to do all of those things that is about 70% of the way there—a big improvement. Customers struggling with scalability issues are interested in our management approach, because it shows how a very limited staff can effectively manage a dynamic, high-volume, high-request environment.

Thin Provisioning

I never actually used the term "thin provisioning" in association with Kilo-Client until a Gartner analyst pointed out that this is one of the best large-scale, real-world examples of it. He's right—our lab is highly space efficient because cloned images (LUNs) only consume additional disk space as boot images change, providing over 1,500-fold efficiency of capacity.

For example, let's say we wanted to boot all 1,498 servers with Red Hat Linux. The total storage requirement in our test lab would be 7.63TB (assumes 20GB for each of seven boot storage systems and 5GB per blade). In a traditional server farm—or even with traditional diskless booting—we would need a full 20GB per server so our total storage requirement would be 30TB. Ouch! As I said upfront, we'd spend more time configuring and managing the cluster than running the tests.

Cooling Design for Dense Configurations

One of the questions I get asked most often is, "How in the world do you cool this beast?" Part of this ties into the point I made about thin provisioning: there just isn't as much to cool as there would be in a traditional environment.

Still, 1,500 blades, 7102 fabric ports, and 87 storage controllers are an awful lot of equipment packed into a dense area. In our original data center, we used a hot aisle/cold aisle approach. We created a cold aisle on the front side of the equipment (where air is drawn in) by adding extra cooling equipment there. That gave us as much as a 30-degree delta from front to back.

Performance

In early 2006, NetApp delivered 1,032,461 SPECsfs97_R1.v3 operations per second—more than triple the previous performance record—using the Kilo-Client lab, the Data ONTAP GX operating system, 24 FAS6070 nodes, and a single namespace.

Additional details are available in a [press release on these results](#).

We recently moved to a new data center, and in our new lab we took a different approach—we created a cold room. We purchased new floor-to-ceiling cabinets and made sure that all openings were completely sealed from front to back, creating an air conditioning plenum. The only place for the cooled air in front of the equipment to go is through the equipment, and it never mixes with heated air coming out the back. Air pressure is also slightly higher on the cold side to ensure the flow is only in one direction. Using this approach, we get about 8 kilowatts of cooling in the lab versus 4 kilowatts with the previous design.

Incidentally, some visitors have asked if we use controllable power strips that power down clients that are not used. In truth we didn't even consider it since our goal has been 100% utilization since day 1. The servers are 100% reserved and automated tests run overnight so there are no shut-down periods.

Summary

Over the past two years we've learned a lot about managing a large-scale environment. We've also discovered from customers and analysts that this architecture is impacting how they think about technology and data center design. Key benefits include:

- **Huge reductions in server provisioning time**
- **Massive flexibility in the infrastructure for quick reconfiguration**
- **Ability to manage the infrastructure with a small team**
- **Ability to save and re-use an environment that has been configured**

The ultimate promise of this architecture is scalability. How can a company grow at 30% without increasing hardware at the same 30% rate? Many companies can no longer build out their data centers fast enough to accommodate growth, and the types of technologies we're using have the potential to bend that acquisition curve.

Find Out More

I've touched on a lot of different topics here, and space limitations have kept me from going into much detail on any of them. Although the Kilo-Client architecture is a moving target, there are additional resources available. Most recently, EMA released a white paper that provides an [outside assessment of the Kilo-Client architecture](#). A [technical case study](#) from January 2007 also provides more details.

TECH ONTAP ARCHIVE



Blake Lewis

Technical Director, Data Retention, NetApp

Blake Lewis joined NetApp in 1996 and has contributed to many areas of the Data ONTAP® operating system. For several years, he had architectural responsibility for the NetApp WAFL® file system. Currently, he is a technical director in the Data Retention group, where his focus is on making secondary storage more useful and less expensive.

A-SIS: Deduplication Comes of Age

By Blake Lewis

Everyone knows that the capacity of storage systems is going up at a breathtaking pace. In the last 10 years, NetApp has gone from shipping storage systems with tens of gigabytes to hundreds of terabytes, an astonishing 10,000-fold increase. Most businesses, however, find that their appetite for storage has grown even faster and – in addition to the costs of disk or tape to store all this data – data center space and power are increasingly expensive. Using storage as efficiently as possible is therefore a critical objective.

NetApp has long been an industry leader in efficient storage utilization, from its unique incremental-only Snapshot™ technology, which requires minimal disk space to store hundreds of Snapshot copies, to FlexVol® technology, which enables sysadmins to expand and contract volumes on the fly.

In May, NetApp announced a new deduplication technology that can significantly increase the amount of data stored in a set amount of disk space: Advanced Single Instance Storage (A-SIS) deduplication. This technology is available (at no charge!) for NetApp NearStore® R200 and NearStore on FAS systems.

Deduplication improves efficiency by finding identical blocks of data and replacing them with references to a single shared block. The same block of data can belong to several different files or LUNs, or it can appear repeatedly within the same file. A-SIS deduplication is an integral part of the NetApp WAFL file system, which manages all storage on NetApp FAS systems. As a result, deduplication works "behind the scenes," regardless of what applications you run or how you access the data, and its overhead is low.

How much space can you save? It depends on the data set and the amount of duplication it contains. Here are a couple of examples of the savings that NetApp customers have seen:

- A global oil and gas company achieved a 35% space savings for its home directory storage.
- An investment management company reduced backups copies of their VMware images by 90%.
- A test and measurements manufacturer realized a 98% space savings on daily database backups.

How A-SIS Deduplication Works

At its heart, A-SIS deduplication relies on the time-honored computer science technique of reference counting. Previously, WAFL kept track only of whether a block was free or in use. With A-SIS deduplication, it also keeps track of how many uses there are. In the current implementation, a single WAFL block can be referenced up to 256 times in different files or within the same file. Files don't "know" that they are

**OnDemand
TechTalk Webcast:**

Cut Data Storage
Use by As Much As
95%: Using
Deduplication to
Drive Down Storage
Costs

RELATED INFORMATION

- [Demo: NetApp A-SIS Technology](#)
- [Webcast: Cut Data Storage Use by As Much As 95%](#)
- [Dave's Blog: How Data De-Duplication Fits into Our Master Plan](#)

87% Space Savings in 60 Days

The County of Sacramento has a backup environment that includes approximately 2.1TB of data. The agency saves daily incremental Snapshot copies for two weeks, and keeps weekly full backups for 60 days. Every day, approximately 5% of files and 0.5% of blocks change.

By deploying a NetApp storage solution that includes A-SIS deduplication technology, the agency saw results that included:



- ~ 87% space savings at 60 days
- Backup 16x faster
- Restores 6x faster: 5 min vs. 30 min

Listen to Keith Scott, County of Sacramento IT Analyst, describe the deployment.

How Data Deduplication Fits into the NetApp Master Plan

Following is an excerpt from a recent post on NetApp Founder Dave's Blog:

sharing their data – bookkeeping within WAFL takes care of the details invisibly.

How does WAFL decide that two blocks can be shared? The answer is that for each block, it computes a "fingerprint," which is a hash of the block's data. Two blocks that have the same fingerprint are candidates for sharing.

When A-SIS deduplication is enabled on a volume, it computes a database of fingerprints for all of the in-use blocks in the volume (a process known as "gathering"). Once this initial setup is finished, the volume is ready for deduplication.

To avoid slowing down ordinary file operations, the search for duplicates is done as a separate batch process. As the file system gets updated during normal use, WAFL creates a log describing the changes to its data blocks. This log accumulates until one of the following occurs:

- The administrator issues a sis start command
- The next time specified in the sis config schedule occurs
- The changes to the log exceed a predetermined threshold

Any of these events will trigger the deduplication process. Once the deduplication process is started, A-SIS sorts the log using the fingerprints of the changed blocks as a key, and then merges the sorted list with the fingerprint database file. Whenever the same fingerprint appears in both lists, there are possibly identical blocks that can be collapsed into one. In this case, WAFL can discard one of the blocks and replace it with a reference to the other block. Since the file system is changing all the time, we of course can take this step only if both blocks are really still in use and contain the same data.

The implementation of A-SIS deduplication takes advantage of some special features of WAFL to minimize the cost of deduplication. NetApp discovered a long time ago that to ensure the integrity of data stored on disk, a belt-and-suspenders approach is warranted. (In fact, several pairs of suspenders is best.) Accordingly, every block of data on disk is protected with a checksum.

A-SIS uses this checksum as its fingerprint. Since we were going to compute it anyway, we get it "for free" – there is no additional load on the system. And since WAFL never overwrites a block of data that is in use, fingerprints remain valid until the block gets freed. The tight integration of A-SIS deduplication with WAFL also means that change logging is an efficient operation. The upshot is that A-SIS deduplication can be used with a wide range of workloads, not just for backups, as has been the case with other deduplication implementations.

What Sorts of Environments Are Good Candidates for A-SIS?

In the first place, your data should be fairly long-lived. There isn't much point in working hard to find duplicates if you are going to be changing the data soon. The system should also have some CPU headroom. Change logging and fingerprint matching are designed for efficiency, but nothing is free. If your system spends long periods at high CPU utilization, the extra load that deduplication brings could be the last straw.


Other Approaches for Saving Disk Space

NetApp offers a variety of other alternatives to use disk space more efficiently, each with its pluses and minuses. It isn't necessary to pick just one; for the most part, they can all be used in conjunction.

Snapshot Copies

From the beginning, WAFL has allowed block sharing through Snapshot technology. As a file changes over time, you can capture several versions of it using Snapshot copies, and the storage cost is just equal to the amount of change between versions.

Snapshot copies have proven their value both as a feature in their own right and as the



Buying less storage is the small picture. The big picture is that we want to help customers create a disk-based copy for all of their primary storage. ...

Interesting things start to happen when you create a disk-based copy of everything. Instead of doing searches on primary storage, which could hurt performance, why not search the secondary copy? If the people running decision support systems want their own copy of a critical database, why not clone the secondary instead of paying for a whole new copy? Why not create lots of cloned copies for the test and development team preparing to upgrade to the next version of Oracle or SAP?

When you create a copy of everything, and add functionality like Snapshot copies and clones, what you end up with is a smart copy infrastructure that can completely change the way you think about data management.

This won't happen overnight. We understand that. But anything that helps people reduce the cost of creating copies helps us achieve our vision more quickly. In the short run, data deduplication helps customers save space and save money, but what's more important is that by reducing the cost of copies, it helps us achieve our master plan.

Read the full [Dave's Blog post on Deduplication](#)

The Core NetApp DNA

In the early 1990s, Network Appliance revolutionized storage networking with a simple architecture that relied on NVRAM, integrated consistency points, and a unique WAFL file system to do things that the file servers of the time could not.

These basic building blocks of NetApp technology support:

- Multiprotocol environments (NFS, CIFS, FC, iSCSI, and so on)
- Clustered failover, mirroring, and disk-to-disk backup
- RAID-DP and other software-based resiliency features
- The near-instantaneous creation of writable clones
- Block deduplication using A-SIS

Learn about [The Core NetApp DNA](#).

basis for applications such as SnapVault® and SnapMirror®. In WAFL, they come for free as far as performance is concerned. Their main limitation is that they can provide block sharing only between different versions of the same file, unlike A-SIS, which shares duplicate blocks between different files.

Incidentally, if you haven't used NetApp storage before, the NetApp "incremental-only" approach to Snapshot copies is unique among all major storage vendors and is the fundamental technology behind our SnapVault and SnapMirror products, and the main reason for their success.

Compression

Compressing data before it is written to disk is a good way to save space. Algorithms such as gzip can cut the size of a file in half or more, and it works even if there is no duplicated data for sharing. The drawbacks of compression are that it is CPU-intensive. Also, some types of data such as images are already compressed and get no benefit. Because A-SIS deduplication can collapse hundreds of copies of the data into one, it has the potential for much greater savings than compression in environments with lots of duplication.

NetApp currently offers compression in its Decru® and VTL products.

Content-Addressable Storage (CAS)

Although the implementation is usually quite different, content-addressable storage is conceptually similar to A-SIS deduplication. A "blob" of data gets hashed, and the hash value is used to identify it. Only one copy of data with a given hash value is stored. A file can consist of a number of blobs.

In one way, CAS is more flexible than A-SIS deduplication, since CAS blobs do not need to be whole file system blocks. However, in a very important way, CAS is less flexible. With A-SIS deduplication, WAFL can share blocks using fingerprints as keys, but its basic data structures remain unchanged and the sharing is invisible (and of course, you can always turn A-SIS deduplication off). By contrast, in most CAS implementations, blobs are always found through their hash keys. This makes it hard to get good performance, with the result that CAS is generally used for write-mostly archival applications and not for applications that require a quick response to bursts of reads, such as e-discovery and data recovery.

One aspect of CAS that sometimes sparks controversy is that it considers two blobs to be identical if they have the same hash key. If two different blobs happen to hash to the same value, data is lost. This is known as a "hash collision" or a "false positive." There are good statistical arguments for why such an event is highly unlikely, but many people still feel uneasy. A-SIS deduplication takes a conservative approach in this regard, and shares blocks only if their contents (and not just their fingerprints) are identical. Before deleting a block as a duplicate, A-SIS does a byte-by-byte comparison to make sure that the data is indeed the same.

Conclusion

A-SIS deduplication leverages the unique characteristics of WAFL to conserve disk space while keeping system overhead low. In many environments, the space savings can be substantial. Even in primary storage applications, such as a home directory environment, A-SIS deduplication can often produce significant savings.

Just as with NetApp Snapshot technology, the A-SIS deduplication machinery will almost certainly provide the basis for interesting new applications in the future (cloning a file, for instance). It's an exciting development in the ongoing evolution of WAFL.



William Morlett

Enterprise Systems Engineer

William is a Systems Engineer at a rapidly expanding entertainment company with 10 years of IT experience. Over the past two and a half years his responsibilities have come to include implementation and management of NetApp storage in a globally distributed environment. He is also directly responsible for a 25-node Microsoft Exchange environment and replication of remote site data for disaster recovery purposes.

Featured Tool: visionapp Remote Desktop (vRD)

Author: vRD is a freeware tool developed by [visionapp](#).

What it is: The visionapp Remote Desktop (vRD) is a simple, graphical tool that allows you to remotely connect to Windows® boxes, maintaining simultaneous connections to multiple servers. Appropriate settings and login credentials for each connection can be logically organized in folders for fast and easy access to important servers. Console logins are also supported.

How it works: The vRD tool uses the Remote Desktop Protocol (RDP), which allows you to establish a remote desktop connection to any system running Microsoft® Terminal Services. Client software that implements RDP can make it simple to connect to and manage remote servers or any Windows system. However, the Microsoft Remote Desktop client only provides a single long list of machines for which connections have been configured. By comparison, vRD lets you sort your various connections into folders and apply credentials at the folder level. Credentials can be inherited within a folder so you don't have to manually type in the information for each connection. (The Microsoft tool also lacks this inheritance feature.)

How do you use Remote Desktop? The company I work for publishes and develops video games for all popular gaming platforms. The company currently has 14 NetApp storage systems – with more on the way – to meet its storage needs and about 150 servers distributed across 25+ sites and data centers.

I'm primarily responsible for managing the company's Exchange 2003 infrastructure and replication of remote site data for disaster recovery purposes. My environment includes about 25 different locations across the globe. Most all of the Exchange servers have now been configured on NetApp storage via iSCSI and run SnapManager for Exchange (SME).

To facilitate managing my environment with vRD, I've organized the servers into folders which reflect sites or NetApp function. Using vRD, I can quickly locate and connect to servers to perform any number of tasks needing to be performed with SME, such as tune backup schedules, change Snapshot™ copy retentions, etc. I also have Open Systems SnapVault® (OSSV) nodes in remote locations to manage, so I have a separate folder which I've organized these into. At times, it's also necessary to manage the remote office backup servers on which I run NetApp SnapDrive and NetApp Single Mailbox Recovery (SMBR) to recover email items from snapshot. I've also organized these servers into a separate folder.

The company also uses NetApp Operations Manager hosted on a dedicated server for direct management of its NetApp storage systems. Occasionally, I need to make configuration changes on the server that hosts Operations Manager, so I have a separate folder for that connection as well.

As soon as I launch vRD, I'm presented with an organized folder list. Because the vRD tool allows me to organize my connections, it's easy to find the node I want to open a session to. At that point, I can access the features of SnapDrive®, SME, or Single

RELATED INFORMATION

✦ [vRD download](#)

✦ **Previously Featured Tools:**

✦ [Hobbit](#)

✦ [NTOP](#)

✦ [Toasters](#)

✦ [Swatch and Kiwi](#)

✦ [Firefox](#)

✦ [Cacti](#)

✦ [MRTG for Filers](#)

✦ [Nagios](#)

✦ [SIO](#)

Mailbox Recovery to perform whatever task I need to perform.

This may not seem like a big deal at first glance since Microsoft Remote Desktop has existed for some time, but for admins that manage a large number of servers, having everything in one window, being able to control how your view of available connections is organized, and applying the right credentials for each connection automatically can make RDP connections easier to manage. While Microsoft's version of Remote Desktops offers much of this same functionality they don't allow admins to organize their Remote Desktop connections into folders in order to organize them. The Microsoft version requires multiple MMCs to achieve somewhat the same functionality and credentials cannot be shared across MMCs or even groups of servers.

Why do you like vRD? It's a simple, free tool that helps me better organize all my remote desktop sessions, and it allows me to manage all my remote desktop connections into a single window. When a session is opened to a system, the icon for that system is changed to include a green arrow so open sessions can be identified, unlike the Microsoft version that only changes the icon for the connection you are actively using. In an organization that has many Windows servers to manage, this tool helps to more effectively and efficiently manage the servers.

Caveats:

None.

✱ [vRD download](#)