NetApp

Tech OnTap

Highlights October 2009

A Quantum Leap in Virtual Infrastructure Management

Trey Layton, Systems Engineer

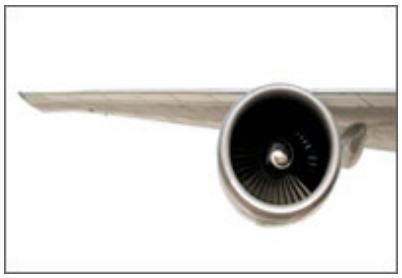

Managing VMware® servers and storage has become so simple that you can now perform complex provisioning tasks even while flying coast to coast.

Learn what was hot at the NetApp booth at VMworld 2009 with the latest on the NetApp® Rapid Cloning Utility (RCU) v2.1, SnapManager® for Virtual Infrastructure (SMVI) v2.0, and Virtual Storage Console v1.0.

More

Tips from the Trenches

Case Study: Consolidating Windows Storage and Servers for Big Savings

Keith Alioto, IT Director

Osiris Therapeutics is growing fast and thinking big. A recent round of Windows® storage consolidation and server virtualization really boosted IT productivity. Learn how.

More

Blog of the Month

The Upcoming Hypervisor Wars

Since storage is becoming increasingly important in virtualized data centers, it's probably worth asking your storage vendor how they are going to help you with the new options and potential future considerations.

Read the Blog

Engineering Talk

Accelerate Replication on Low-Bandwidth Links

Srinath Alapati, Technical Marketing Engineer

For many, bandwidth constraints limit the use of replication. SnapMirror® network compression changes the game by lowering bandwidth utilization up to 70%. Find out how.

More

Community Spotlight

Are You an Innovator?

Nominate your company for a NetApp Innovation Award for achieving positive business results through innovative use of NetApp.

Hottest Virtualization Community Topics

See the latest discussions about VDI, RCU, SMVI, Snapshot™ best practices, and more.

A Quantum Leap in Virtual Management

By Trey Layton, NetApp

Back in March, I wrote a Tech OnTap article about a business process outsourcing company that uses the NetApp® rapid cloning utility (RCU) to quickly deploy call center desktops. The approach I described has now become an integral part of that company's business process and a significant selling point for new business.

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

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

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

Rapid Cloning Utility Version 2.1

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

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

Several recent Tech OnTap articles have described the use of RCU: vStorage Integration, Consolidating 1,000 Physical Servers, and A 9,000-Seat VDI Deployment.
RCU 2.1 adds a number of new capabilities to the functionality of RCU 2.0, including:

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

RCU 2.1 is available on the NOW™ site and requires Data ONTAP® 7.3.1.1 along with FlexClone and appropriate protocol licenses. For more information, watch a demo of RCU functionality or read the RCU v2.1 FAQ.

SnapManager for Virtual Infrastructure Version 2.0

NetApp SnapManager for Virtual Infrastructure (SMVI) is designed to simplify the management of backup, restore, and disaster recovery operations in VMware environments. Like RCU, SMVI works with VMware vCenter for seamless operation.

SMVI v2.0 adds a number of new features, including important restore enhancements such as:

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

Explore

A 9,000-Seat Virtual Desktop Deployment

Author Trey Layton recently worked with a company that used VMware and NetApp technology to deploy 9,000 virtual desktop seats for call center agents working anywhere in the world.

The article describes the decision process and details the architecture used.

More

Consolidating Up to 1,000 Physical Servers on VMware

The NetApp Bangalore Engineering Lab is tasked with meeting the compute and storage needs of roughly 700 busy engineers. By August 2008, the lab had deployed over 1,000 physical servers and was continuing to deploy 40 new servers a month, creating unreasonable demands in terms of power, space, and cooling.

NetApp has initiated an internal project to consolidate these physical servers down to as few as 20 VMware servers.

More

Watch a demo of SMVI v2.0 single file restore.

You can read more details about these features and more in a recent blogpost from frequent Tech OnTap contributor Nick Triantos. A series of blog posts from Kostadis Roussos in 2008 provides a great overview of SMVI.

Virtual Storage Console Version 1.0

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

When you configure NetApp storage with VSC, it automatically makes sure that storage settings such as multipath settings for Fibre Channel, HBA timeouts, or NFS settings adhere to NetApp best practices. It also provides end-to-end troubleshooting should a connectivity problem arise. VSC incorporates the best practices for NetApp and VMware as laid out in TR-3428.
VSC includes the following capabilities:

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

For more details, see the VSC demo.

The Buzz from VMworld 2009

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

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

Taken together, these tools take the complexity out of managing and protecting your virtual storage environment.

Got opinions about VMware and NetApp Management?

Ask questions, exchange ideas, and share your thoughts online in NetApp communities.

Trey Layton
Systems Engineer
NetApp

Trey has been working at NetApp since 2006, specializing in the design of next-generation data centers using VMware. His wealth of experience with networking and virtualization makes him uniquely well suited to the current evolution to network storage. With over 18 years of IT experience, Trey began his career in the U.S. Army at USCENTCOM supporting U.S. Special Operations groups operating in the Middle East. He has also held key network consulting and systems engineering positions at Eastman Kodak, GE, and Cisco. Read Trey's "Ethernet Storage Guy" blog.

Case Study: Consolidating Windows Storage and Servers Yields Big Benefits

By Keith Alioto, Osiris Therapeutics, Inc.

Taking a new therapeutic agent through all phases of the clinical trial process is an expensive, time-consuming, and data-intensive process that can tax the resources—especially the IT resources—of a small company.

When I joined Osiris Therapeutics two and a half years ago, we were in the midst of clinical trials to bring the world's first stem cell therapy to market, and our data storage needs were growing exponentially. Since then, we've evolved the infrastructure to fully virtualize our Windows® servers with all Windows storage and file serving consolidated on a single, clustered NetApp® storage system for improved availability, scalability, and ease of management. This infrastructure is designed not only to meet our needs now, but also to grow with us over at least the next 12 to 24 months, so we're no longer constantly struggling to maintain enough available storage.

In this article, I'll examine the evolution of our infrastructure and describe where we are now and where we're headed. I'll also talk about why we made the decisions we have and the importance of choosing the right partners.

Struggling with Exponential Data Growth

When I joined Osiris, we had in place a pretty standard Windows environment with a combination of physical and VMware® virtual servers running Microsoft® Exchange, Microsoft SQL Server®, and Symantec™ Enterprise Vault™, plus several Windows file servers. All servers had local storage.

SAN Fails to Address Growth Needs

We quickly realized that this approach could not continue and implemented a 6TB iSCSI SAN using an HP storage array. We used the SAN for VMware virtual machines, archived e-mail storage and other needs, and quickly outgrew it. We ended up buying a second HP array to address virtual machine performance issues and again had to add an additional disk shelf almost immediately because we were still running out of space.

Windows File Server Woes

On the Windows file server side, we had similar problems. The mountains of file data coming in from clinical trials was stored on Windows file servers with local disk storage. We quickly expanded from one file server to three file servers. We were continuously adding more disks and resorted to compressing files and folders in some cases to meet our space needs.

Time to Reassess

At this point we took a step back to ask what solution could carry us through the next two years, rather than having to apply a band-aid every six months or less to meet our storage growth. We started looking at NetApp plus a variety of other offerings from HP, Dell, EMC, and others.

Ultimately we settled on NetApp because I want the partner that I trust with my infrastructure to be more interested in building a relationship toward the future than just selling me a box. With our system integrator, CTI, and NetApp that was what we got. I also really liked the versatility of the NetApp unified storage architecture, which gave us the ability to support all the protocols we needed—NFS, CIFS, and iSCSI—on a single platform. A single storage system with a single management interface and data protection strategy would replace all our existing storage, including SAN storage, local disks, and Windows file servers. We also saw that deduplication was going to be key for us, and we liked the phone home capabilities that the NetApp system offered.

Consolidating with NetApp

We worked with CTI to develop a full plan to consolidate onto a NetApp 2050c with 17TB of raw storage and the ability to meet our iSCSI, CIFS, and NFS needs. This plan contained a number of elements:

- Consolidate VMware data stores using NFS
- Consolidate existing Microsoft SQL Server databases and logs using iSCSI
- Virtualize Symantec Enterprise Vault (P2V) and centralize storage using NFS
- Perform a physical-to-virtual (P2V) Exchange migration
- Migrate Windows file shares to NetApp

Explore

About Osiris Therapeutics

Osiris Therapeutics, Inc. has been working since 1992 to develop and commercialize cellular therapies based on stem cells isolated from readily available and noncontroversial adult bone marrow. These stem cells offer the opportunity to provide revolutionary treatments for a variety of ailments such as osteoarthritis, heart attack, type 1 diabetes, and Crohn's disease.

A treatment for graft versus host disease (GvHD) just completed Phase III clinical trials.

- More

About CTI

CTI—Osiris Therapeutics technology partner—is a systems integrator serving small-, mid-, and large-sized businesses in Maryland; Washington, DC; and central Pennsylvania. CTI provides a full spectrum of IT services, including:

- Storage
- Virtualization
- E-mail and document archiving
- Network security
- Business continuity and DR
- Platform upgrades

CTI is a NetApp Gold Partner.

- More

- Implement appropriate elements from the NetApp SnapManager® Suite for Virtual Infrastructure, Exchange, and SQL Server

VMware on NFS
We chose to move our VMware data stores from iSCSI to NFS for performance and ease of backup based on CTI's recommendation and have been happy with that decision. SnapManager for Virtual Infrastructure (SMVI) was implemented to facilitate data management and backup for the VMware environment, replacing our previous vRanger implementation. We do daily Snapshot™ copies combined with weekly backups to tape to protect this environment.

We've implemented deduplication in our VMware environment to eliminate the storage duplication that arises from having multiple virtual machines running the same operating system.

SQL Server and Enterprise Vault
We use Microsoft SQL Server databases primarily to support our Blackberry server and Symantec Enterprise Vault. Because the SQL Server units were already virtualized, it was simply a matter of migrating storage from the SAN to NetApp. NetApp SnapManager for SQL Server gives us backup and other functionality analogous to SMVI for this environment.

Enterprise Vault was still running on a physical server. As part of the project, we did a P2V migration of the server and moved the vault to an NFS volume on the NetApp storage.

Exchange
Because of the ongoing clinical trials, communication within the company, with partners, and with hospitals participating in the trials is critical, and protecting that information is vital. The Food and Drug Administration (FDA) has several guidelines for data supporting clinical trials. The system or systems have to be dependable with both physical and logical security. We need to provide audit trails of file and folder access as well as accurate date and time stamps. We can't accomplish this using traditional servers based on Windows with local storage and making changes every other year.

We previously kept Exchange on a physical server because of concerns about performance, but the P2V migration and the transition to NetApp for back-end storage has created no issues. As with the other application environments, we've implemented SnapManager for Exchange (SME) for fast, consistent backups. Using SME, we make hourly Snapshot copies of our Exchange environment to protect this critical resource.

Windows File Shares
Largely because of the clinical trials, the data in our file shares has grown more than sixfold over the past 2.5 years. Migrating everything from our three Windows file servers to the central NetApp storage system eliminates three servers and allows us to make Snapshot copies of this critical data every hour. We're also in the process of converting paper documentation to electronic form to accelerate access to data that might previously have required us to request paper records retrieval from Iron Mountain.

Results So Far

All phases of this transition were implemented in June 2009, and we've been extremely happy with how it went. As a result of these changes we retired:

- Five physical servers and associated direct-attached storage (DAS)
- Two HP SANs with 6TB of SATA and 1TB of SAS storage

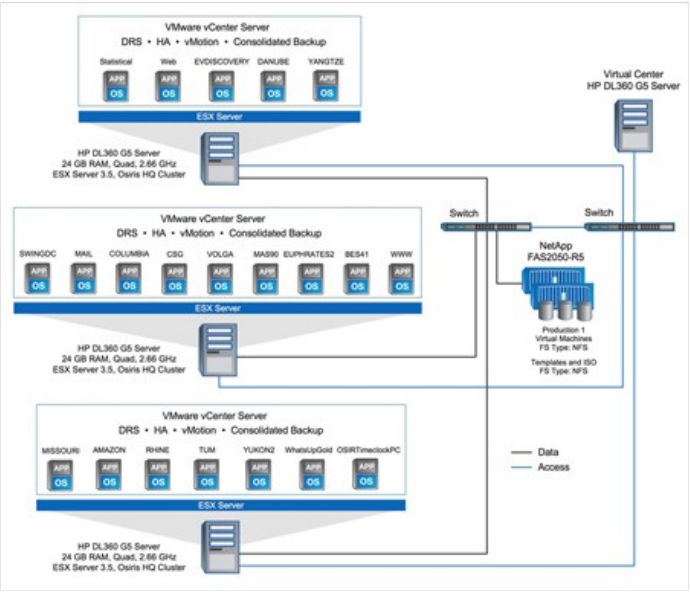

Figure 1) Osiris Therapeutics' current Windows environment.

While I haven't made any direct efforts to quantify the savings, we did drop from four ceiling-based air-conditioning units to two floor units because of the substantial reduction in heat output. Obviously, that corresponds to less electricity used by the infrastructure and less electricity needed for cooling.

Deduplication
Additional savings result from the use of NetApp deduplication on our primary storage. Deduplication on our CIFS and NFS volumes results in significant savings:

- CIFS: 35%
- NFS: 22% to 76%

In addition, we can now add additional virtual servers with almost no additional storage consumed because deduplication eliminates redundancy in virtual server environments.

Rationalized Backup
From a management standpoint, the biggest savings result from rationalizing our approach to backup. Previously, we struggled with four different methods of backup. Now we make hourly Snapshot copies for the most critical applications (Exchange and file storage) and daily Snapshot copies of the rest (including SQL Server and Enterprise Vault). Instead of hours to do backups it takes seconds; restores are just as quick, and we've noticed no issues or conflicts with VMware features such as VMotion®. In addition, backups always succeed, and we get regular notifications to tell us so. With our previous methods, backup failures were common.

NetApp SnapManager tools automate the process and make sure that backups are consistent for our VMware environment, Exchange, and SQL Server. We used to suffer noticeable database drops from our Blackberry server when running SQL Server backups, but these have been eliminated.

Simple, Fast Management
Overall, I've been surprised how user-friendly the NetApp storage system is to manage compared to other storage I've worked with such as EMC and HP. Recently, we had a problem with Enterprise Vault when we changed our retention policy and had to do a complete dump back to Exchange. Even though both the vault and the Exchange database were on NetApp storage, the process proceeded very quickly. When Exchange ran out of space, I was able to expand the volume on the fly so the dump could continue. This was a real lifesaver, and the process completed quickly after that.

Performance
From a performance standpoint, we've seen the same or better performance across the board as a result of these changes. Enterprise Vault performance has actually gotten better since the server has been virtualized because both the database and vault are now spread across far more disk spindles than they had been previously. (NetApp flexible volumes—FlexVol® volumes—automatically spread I/O for each volume across the maximum number of spindles possible.) The same goes for Exchange. We reduced the amount of RAM allocated to Exchange from 16GB on the physical server down to 6GB on the virtual server, and performance is still better because of improved I/O.

Ready for the Future

All the improvements I've mentioned above are important, but for Osiris the biggest benefit is that our scientists can now do their jobs without concern for possible IT or data storage limitations. We used to sometimes ask external vendors to house data for us; we can now bring that data in house and analyze it quickly and efficiently, and—because our data is better protected—we all rest easier. This solution lets us work more efficiently now and prepares us to meet our growth needs for the foreseeable future.

The next step for us will be to implement SnapMirror® to allow us to replicate critical data off-site. We expect to implement that solution as soon as we reach the point where off-site storage is needed. We'll also be looking at adding Single Mailbox Recovery for Exchange.

The new technology in this deployment was easy to implement and set up. Osiris Therapeutics is a small organization with just 70 employees and two dedicated IT staff. Having reliable partners such as NetApp and CTI was essential to making this project a success. The technical experts at CTI helped us make sure that all phases of the plan were implemented quickly and with as little disruption as possible.

Got opinions about Storage Consolidation for WIndows?

Ask questions, exchange ideas, and share your thoughts online in NetApp communities.

Keith Alioto
IT Director
Osiris Therapeutics, Inc.

Keith has been IT director at Osiris for the past 2.5 years. His primary objective is to support the company's goal of achieving FDA approval by making sure that IT is fast and efficient. Prior to Osiris, he spent 1.5 years as an IT project manager for LCG Technologies and four years as IT director for a venture capital firm. .

Accelerate Replication on Low-Bandwidth Links: SnapMirror Network Compression

By Srinath Alapati, NetApp

For many companies, available bandwidth is a major factor limiting the use of replication. In a recent survey of companies with greater than 1,000 employees, 44% reported that bandwidth limitations were the number one issue affecting their replication and data protection plans for branch and remote offices.

NetApp® SnapMirror® thin replication software has become popular with NetApp customers in large part because of its efficient use of network bandwidth. Unlike other replication products that copy entire files across the network any time a single block is changed, SnapMirror replicates only changed blocks, significantly reducing bandwidth requirements. A 50% reduction in management overhead versus competing solutions and the ability to use replicated data sets for dev/test, data mining, or other purposes add to the appeal of SnapMirror.

Despite the proven efficiency of SnapMirror, however, situations still arise where available network bandwidth is not sufficient for replication. To better address cases where bandwidth is the limiting factor, NetApp recently announced SnapMirror network compression, which can lower volume SnapMirror bandwidth utilization by up to 70% or more (depending on the compressibility of your data set).

SnapMirror network compression has been added to Data ONTAP® 7.3.2 and is available to SnapMirror users at no additional cost. This article explores how network compression works, talks about when you should (and shouldn't) use it, and describes some observed results for common data sets.

How SnapMirror Network Compression Works

With SnapMirror network compression data is compressed only while it traverses the network; data on source and destination systems remains uncompressed. Enabling compression results in two additional steps:

- Compression on the source system
- Decompression on the destination system

On the source system, data blocks that need to be replicated are handed off to a compression engine, which compresses them. The compression engine creates multiple threads corresponding to the number of CPUs on the storage system. The multiple compression threads compress data in parallel. Compressed blocks are then transmitted over the network. On the destination system, compressed blocks are received and decompressed using a similar multithreaded approach. Decompressed data is then written to the appropriate volume.

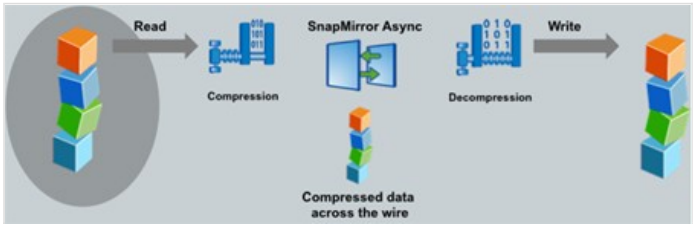

Figure 1) Functional diagram of SnapMirror network compression.

The compression and decompression engines can either be configured to conserve network bandwidth or complete a transfer in the shortest time possible, depending on user preference. SnapMirror network compression is supported on all NetApp storage platforms (including V-Series virtualization systems and the IBM N-series) in the asynchronous mode of operation only. The semi-synchronous and synchronous modes of SnapMirror operation are not currently supported with network compression enabled.

When to Use Network Compression

SnapMirror network compression is useful in the following scenarios:

When network bandwidth would otherwise be a limitation. Depending on the compressibility of your data set (see following section for information on data sets), enabling network compression can make it possible to use replication on links where it would otherwise not be possible to meet (or continue to meet) your goals.

In general, the suitability of a replication solution for disaster recovery is dictated by your recovery point objective (RPO), which defines the point in time (relative to the time a failure occurs) to which you want to be able to recover. You must be able to replicate data quickly enough to meet this objective. For example, if your RPO is one hour, you must be able to replicate all changed data every hour, even during peak periods of activity. If your network bandwidth is limited, network compression can help you:

- Maintain the same RPO level in the face of data growth and/or increased change rates
- Improve your RPO without buying additional bandwidth

Similarly, if you are using replication for data distribution, network compression can help you continue to meet your existing time objectives or to meet more stringent time requirements without adding bandwidth.

To save precious network bandwidth. Network compression can help you continue to perform the same replication schedule while conserving network bandwidth for other functions.

To perform baseline transfers without saturating network links. To use SnapMirror (or any replication product) you must first perform a baseline transfer in which all data is replicated from source to destination. Once the baseline is in place, subsequent transfers only require that changed or new data be replicated.

Creating a baseline can be very bandwidth intensive. Enabling compression can allow baseline transfers to complete more quickly while conserving network bandwidth. (SnapMirror provides a throttling feature that can prevent the software from saturating the network during baseline transfers or other operations.)

When Not to Use Network Compression

The major factor in determining when to enable SnapMirror network compression is available CPU capacity. The compression and decompression engines create additional load on CPU cores on both source and destination systems and should factor this load into your decision to enable network compression. You might want to avoid using network compression under the following conditions:

- **When CPUs are already heavily loaded** on either the source or destination storage system. Because compression and decompression result in increased CPU utilization, it is not generally recommended to enable compression under conditions where available CPU capacity is likely to become a bottleneck and affect other workloads.
- **When network bandwidth is not a limitation**. Network compression necessarily creates processing overhead and might limit throughput relative to SnapMirror without compression enabled. You should not turn on compression in situations where the network is not a bottleneck and you don't need to conserve bandwidth for other traffic.
- **When using "fan-in" configurations.** Decompression creates only about 60% of the overhead of compression, so in most situations, performance of the destination system is unlikely to be an issue. However, if you have a configuration where you have multiple source systems all replicating to a single destination system at the same time, enabling compression for all transfers could overwhelm the processing capabilities of the destination system.

The following section provides more details on use of network compression with real workloads and effects on CPU resources.

How Well Does Network Compression Work with Typical Data Sets?

We chose three different data sets to measure the performance of SnapMirror network compression under laboratory conditions: Exchange database, Oracle® Database, and home directory data. We performed baseline transfers of eight 50GB volumes in each case and looked at both compression ratio and CPU utilization.

Oracle data was the most compressible, achieving a ratio of 3.5:1, followed by home directories at 2.7:1 and Exchange at 1.5:1. The reduction in bandwidth required (holding transfer time constant) is roughly proportional to the compression ratio as shown in Table 1.

Table 1) Compression ratios and bandwidth savings for common workloads.

Data Set	Compression Ratio	Bandwidth Savings
Exchange	1.5:1	34%
Home directory	2.7:1	60%
Oracle	3.5:1	70%

The effect on CPU overhead for each data set depends on whether you hold transfer time constant or accelerate transfer time.

- **Transfer time held constant**. If transfer time is kept the same as for the transfer without compression, overhead from compression processing is typically not significant. For example, assume a data set that yields 2:1 compression and that without network compression the SnapMirror update takes one hour using a bandwidth of 100Mb/sec. With network compression enabled only 50Mb/sec bandwidth is needed to achieve the same transfer time. Because over-the-wire bandwidth is lower, CPU utilization due to network processing is decreased, compensating somewhat for the increased CPU used by compression. In the case of Oracle Database, CPU overhead of just 14% reduces the needed bandwidth by over 70%.
- **Reduced transfer time**. If a transfer is allowed to use all available bandwidth with SnapMirror compression enabled, CPU overhead will be higher. For example, consider the same data set as above with 2:1 compression in which an update without compression takes one hour using a bandwidth of 100Mb/sec. With network compression enabled, the transfer completes in 30 minutes. Because the work is completed faster by using the entire bandwidth, network processing overhead is higher, and the compression processing must also be completed in half the time. If CPU utilization is too high, you can use SnapMirror throttling (either per transfer or global) to adjust the throughput so that CPU utilization does not go too high. The following figure summarizes the effects of compression on total transfer time for our three data sets.

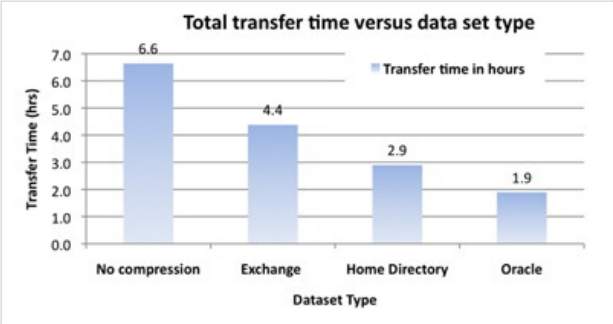

Figure 2) Reduction in transfer time for different data sets. CPU overhead ranges from 38% for Exchange to 55% for Oracle. Total amount of data transferred was the same in all cases. Storage system was a FAS3070 using a single 155Mb/sec network connection.

- Another factor that affects CPU overhead is the compression ratio achievable for a data set. If a data set yields a high compression ratio, the total amount of effort to compress and send it over the wire is less compared to a data set that is hard to compress. In contrast, if you don't throttle bandwidth for a transfer, a data set with a high compression ratio could require a lot of compression work to fill the pipe, raising CPU utilization. For example, to fill a 100Mb/sec pipe with Oracle data that compresses at 3.5:1, your storage system would have to compress data at a rate of 350Mb/sec, while filling the same pipe with Exchange data that compresses at 1.5:1 would only require a compression/decompression rate of 150Mb/sec.

Conclusion

In situations where network bandwidth is limited, SnapMirror network compression can make a big difference, making replication possible in situations where bandwidth would previously have been the limiting factor. If you're already a SnapMirror user, you can take advantage of network compression simply by upgrading to Data ONTAP 7.3.2 (or later).

With SnapMirror network compression, there is a strong interdependence between data transfer time, network bandwidth consumption, and CPU overhead. Decreasing transfer time necessarily consumes more network bandwidth and system CPU resources, as illustrated in Figure 3.

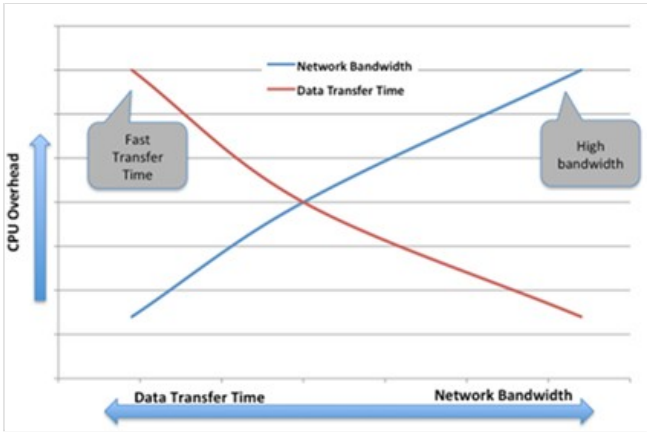

Figure 3) Relationship between transfer time, bandwidth utilization, and storage system CPU overhead.

To find the sweetspot for your environment, pay attention to the guidelines outlined in this article. Plan to do a few trial transfers to assess the compressibility of your data sets and the resulting CPU overhead on source and destination systems before enabling network compression in production.

Got opinions about SnapMirror Network Compression?

Ask questions, exchange ideas, and share your thoughts online in NetApp communities.

Srinath Alapati
Technical Marketing Engineer
NetApp

Srinath joined NetApp in 2004 and has been a member of the Data Protection group for over two years. He has 10+ years of experience in IT, managing servers and storage infrastructure. Srinath has authored or coauthored multiple technical reports on SnapMirror, MetroCluster, VMware®, and Exchange and speaks at various technical conferences. He is also a core team member involved in NetApp IT's disaster recovery implementation.