

» Tech OnTap Archive

November 2006

HIGHLIGHTS

- » One Way NetApp Simplifies SANs
- » Top Encryption Concerns Solved
- » NearStore® VTL Enhancements
- » Toasters: A Mailing List

Using Cutting Edge Technology Today

Q&A with Reinoud Reynders,
IT Manager, UZ Leuven

A critical database gets 0.5 ms response times via FlexShare while VMware ESX with iSCSI enables 10x faster database loads.

[More »](#)



BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"From the beginning I understood VMware had cool technology, but it took me longer to understand that VMware has the potential to revolutionize big enterprise data centers."

» [Dave's Blog](#)

DRILL DOWN

- » **NetApp Is Data Center Proven**
Check out the new FAS3070 and hear how First American Bank, Anadarko Petroleum and others use NetApp storage.
- » **NetApp FAS3070 vs. EMC CX380** (*pdf*)
See how the newest NetApp hardware platform stacks up.
- » **Storage Implications of 10GbE** (*pdf*)
As 10GbE becomes practical for storage interconnects, what's next for iSCSI?



ADMIN RESOURCES

- » **Free Tool: Toasters Discussion Group**
An independent, customer-managed mailing list "tuned for filer geeks."

TIPS FROM THE TRENCHES

Lost Keys, Performance, and More. Top Encryption Concerns Addressed

Shawna Meyer-Ravelli, Systems Engineering, Decru

Will encryption disrupt your environment? Will your database performance plummet? A field expert addresses these and other common questions.

[More »](#)

Applying Innovation to FC SANs: Focus on Dynamic Queue Management

Nick Triantos, Global SAN Systems Engineer, NetApp

FC SANs will never be simple, but NetApp helps make them less complex. See how NetApp reduces application shutdowns, host reboots, and host-side reconfigurations by managing queue depths from the array side.

[More »](#)



ENGINEERING TALK

No Compromises: How NetApp Built a More Efficient, Higher Performance VTL

Craig Johnston, Chief Product Architect, NetApp

Think it's impossible to get high performance with compression? See how the NetApp NearStore VTL engineering team implemented hardware compression and boosted performance with PCI-e, 4Gb FC, and more. Plus what's on the 2007 roadmap (hint: think deduplication).

[More »](#)





Reinoud Reynders

IT manager, Infrastructure and Operations, Leuven University Hospital

After joining Leuven University Hospital over 10 years ago, Reinoud worked his way up through the ranks to become an IT manager responsible for strategic planning and design as well as 43 staff members. In his spare time, Reinoud pursues outdoor activities that include hiking 'le Massif Central in France and Mount Kilimanjaro in Tanzania, sailing, and gliding.

Using Cutting-Edge Technology Today: How One IT Team Is Making It Work

Q&A with Reinoud Reynders

With more than 2,000 beds and 8,000 staff members, Leuven University Hospital (UZ Leuven) is one of the largest healthcare providers in Europe. The hospital's CIO, Professor Dr. Bart Van den Bosch, and his staff are also among the most innovative IT teams in the world. In 2004, UZ Leuven was recognized with a NetApp Innovation Award in the "Visionary" category.

One of the innovations UZ Leuven is most proud of involves a unique electronic patient records system that also functions as an order-entry system, operation room planner, appointment scheduler, and radiology system. Multiple hospitals outsource their IT environments to UZ Leuven to leverage this system, and the special database model is patented in both the U.S. and the European Union.



To maximize performance and uptime for this critical system, the IT team transitioned the Sybase database that supports it to NetApp storage. This was the final phase of a five-year plan to maximize performance, uptime, and efficiency by upgrading to a 100% NetApp storage environment. To achieve this vision, the team has embraced cutting-edge technology and was a beta test site for NetApp FlexShare™ technology and VMware Infrastructure.

The results are impressive:

- Even during peak periods, with 1,500 simultaneous users and 30 SQL Server databases accessing the same storage, Sybase database response times are a mere 0.5 millisecond. Systems availability exceeds 99.99%.
- In the event of a site disaster, 6,000 Exchange users won't lose a single e-mail due in part to NetApp MetroCluster software.
- By the end of the year, 75% of the hospital's Windows servers will be consolidated with VMware ESX 3.0, iSCSI, and NetApp storage, which is expected to provide a 10-fold performance boost over ESX 2.5 and local storage.
- Two people spend just 25% of their time managing over 227TB of storage.

Tech OnTap caught up with UZ Leuven IT manager Reinoud Reynders to learn more.

Q: You recently consolidated your Sybase database environment—the hospital's crown jewels—on NetApp storage. Could you tell us about this environment?

Reinoud: Our Sybase database environment is huge. All of the hospital's clinical records are stored on a Sybase system, and it supports over 1,500 simultaneous users across all areas of the hospital. Our entire workflow is stored in our clinical system. If

RELATED INFORMATION

- [Tech Case Study: UZ Leuven \(pdf\)](#)
- [UZ Leuven Email and Archive Strategy \(Symantec success story\)](#)
- [Workload Prioritization and NetApp](#)
- [Building Virtual Infrastructures with NetApp and ESX 3.0 \(pdf\)](#)

UZ Leuven Profile

Since 2000, UZ Leuven has used NetApp technology as a foundation for server and storage consolidation, improving availability and disaster recovery. Today, it runs its entire operation on NetApp storage.

the systems are down, we literally have to close our doors. The only thing we can keep open is the emergency room and intensive care to make sure no patient will die.

Q: What did your environment look like before the upgrade?

Reinoud: We used to host our Sybase applications on Sun™ servers with direct-attached storage in two separate data centers. We also have a lot of Microsoft® SQL Servers—more than 30—that are also important. For example, PeopleSoft Financials, a monitoring system for the intensive care and emergency room, and our billing software run on SQL Server.

Q: You must have some pretty tough criteria for storage.

Reinoud: Absolutely. Speed and performance are very important. The response time that the user sees depends on the write speed of your logs. So we set a response time goal for Sybase of 1.5 milliseconds in all circumstances. We also wanted to consolidate storage for our SQL Server databases, but only if it could be done without compromising the performance of our Sybase environment.

We looked at two vendors besides NetApp. We eliminated one right away, and the second vendor would have cost us substantially more because it couldn't do iSCSI and Fibre Channel on the same storage.

Q: What does your environment look like today?

Reinoud: Our Sybase database now uses Fibre Channel to access data on NetApp storage systems. Our 30 SQL Server databases access data on the same storage cluster, but they use iSCSI, not Fibre Channel. We use FlexShare to combine all of our major databases on the same storage cluster without compromising the performance of our Sybase database.

Q: So you use FlexShare to set the priority of Sybase applications over SQL Server applications?

Reinoud: That's right. For almost a year we tested the beta FlexShare software. It works perfectly. Sybase applications run with very high priority. SQL Server gets a lower priority so it doesn't compete for resources with Sybase even though its data is stored on the same storage cluster.



Case Study: UZ Leuven Consolidates Database Storage Using FlexShare

Q: What types of results did you see? Did you meet your 1.5 millisecond response goal for the Sybase environment?

Reinoud: It has been more impressive than we expected. Even when we do a load of a database on the system we still have a 0.5 millisecond response time for Sybase. That's three times faster than our original requirement and stays constant with 1,500 simultaneous users.

Queries take only half as long as they used to, and our end users commented on the improvement. Batch jobs that once took an hour now complete in 20 minutes.

Q: Did you have to sacrifice SQL Server performance to achieve these results?


Reinoud: Not at all. For our SQL Server environment we see 0.7 millisecond response times, which is at least twice as fast as the response times when we were using local storage.

Q: How do you ensure availability for your database environment?

Reinoud: We have the same storage and server configuration in two separate data centers. We use Sybase Replication Server to replicate data from one side to the other to keep the Sybase databases in sync. We use log shipping to do the same thing for SQL Server. We can switch an application between data centers in less than three minutes. In some cases we have to do a Sybase database load without disrupting other active applications running on the same server. When that happens, we use FlexShare to set the priority on the volumes being loaded to very low. That way, the load doesn't interfere with other database activity.

Q: Let's talk about Exchange for a minute. UZ Leuven has about 6,000 Exchange users who send and receive over a million e-mails a month. Your team uses NetApp MetroCluster software to protect this environment, which is something many European companies are familiar with but is less common in the U.S. and Asia. Could you explain your decision process?

Reinoud: E-mail is important to us, not because we are a hospital but because UZ



Leuven is a large company. For our databases, we always want to have a logical copy of the data, and we have methods to do that. In Exchange that's not possible. You can't do that with the software. We do tuning and use SME [NetApp SnapManager® for Exchange] to have an active copy that we can switch to if anything happens.

We use a Microsoft Exchange cluster so we have two nodes, one in each data center. One is active and one standby, just like with our databases. To keep the storage in sync between two locations we use MetroCluster. Each data center has one FAS3050 storage system, and NetApp SyncMirror® software, which is part of MetroCluster, duplicates all the data. The Exchange server connects to the FAS3050 via iSCSI.

Q: Why did you use MetroCluster software instead of NetApp SnapMirror®?

Reinoud: MetroCluster is completely synchronous. The data in each location is exactly the same. This means I have a fully redundant system. I can lose a physical server, I can lose a storage controller, or I can lose an entire data center and my operation keeps going.

Q: What other benefits has this provided?

Reinoud: The two Exchange nodes are identical, so it doesn't matter which one is active. We use the standby server for patching since Microsoft software needs to have patches added periodically. Once the standby is patched, we make it the active server. That takes just a few minutes. Then the server in the other data center becomes the standby and we patch it so the two are again the same. That way, everything stays current and we never take Exchange down to do the upgrades. Of course, if we have a problem or a failure, the standby is always ready to take over immediately.

Q: You also mentioned VMware. UZ Leuven was one of the first companies in the world to use VMware with iSCSI and is in the process of moving its environment to NetApp storage. Why take this direction?

Reinoud: I believe VMware is the future. You have to move toward server virtualization when you need to manage a lot of servers. On the storage side, NetApp has a lot of added value for VMware, especially when you use things like Snapshot™ and SnapVault®.

We've used VMware for almost two years now. We use 4-CPU machines and consolidate 20 to 25 Windows® servers on each physical server. It depends on the load on each individual server, of course. We originally started with version 2.5 with direct-attached storage because we didn't want the added expense of Fibre Channel SAN. We support our Sybase database with Fibre Channel, but as I have already said that's our most important environment. We want to keep the Sybase database as an island.

Instead, we've been waiting for iSCSI support in ESX3.0. Our team has years of experience with iSCSI, so we're very comfortable with it. We beta tested the VMware ESX software in April and May of 2006 with iSCSI using hardware initiator cards from Qlogic and NetApp storage.

Q: What type of read/write performance have you seen using VMware with iSCSI?

Reinoud: A database load is 10 times faster using VMware with iSCSI and NetApp storage than with ESX2.5 and local storage running a heavy load.

Q: What does your VMware environment look like today?

Reinoud: Right now we are in the midst of consolidating the rest of our Windows servers with VMware and migrating all of our VMware servers to 3.0. More than 50% of our Windows servers have already been consolidated on VMware, but I/O has been a bottleneck. All the applications that need a lot of I/O have been on direct-attached storage.

Using ESX 3.0 with iSCSI, we can migrate these applications very quickly and by the end of the year over 75% of our Windows servers will be consolidated.

Right now we have two servers running ESX 3.0. There are some limitations in the current release with support for hardware initiator cards, so at the moment we are



using the software initiator. We've been using the software initiator with Microsoft SQL Server and Exchange, so that's not really an issue, but we'd like to be able to use the hardware initiator to offload that work from the CPUs on our servers.

Q: So iSCSI is enabling you to migrate applications with high I/O requirements to virtual servers. What's driving the urgency to do this quickly?

Reinoud: We want to do the migration to ESX 3.0 very fast because of features such as VMotion, which lets you seamlessly migrate a virtual machine to another box. To do this in the old environment with the direct-attached storage, you have to shut down the server and copy the LUN. That takes something like 20 minutes. With VMotion, you can migrate a server and the server still keeps running. It's a very nice benefit.

Q: What will the back-end storage configuration for VMware look like?

Reinoud: Once the migration is complete our VMware environment will be served by the same NetApp FAS3050 MetroCluster that supports Exchange. The storage systems can easily handle both and MetroCluster gives me the same advantages for VMware as it does for Exchange. I can restart my Windows servers in the other data center if something affects the data center where an application is active.

Up to now, we've been using VMFS, the VMware file system, on the back end rather than using raw devices. This was because VMFS gave us the file system functionality we needed when we were running on local storage. Based on discussions during my NetApp visit, I see that there may be advantages to using raw devices as we migrate to iSCSI with NetApp storage.

The NetApp WAFL® file system gives us a lot of the same functionalities as VMFS, plus the ability to use Snapshot, FlexClone™, SnapMirror, etc. Backup is an issue right now, and I see the opportunity with raw devices to have a very sweet backup solution for VMware using SnapVault.

Q: What about provisioning new servers in VMware?

Reinoud: Today when I provision new servers with VMware tools, I need enough extra storage space to make a complete copy of my boot image for each new server. I have the functionality I need, but I use a lot of extra space there. Using raw devices and NetApp FlexClone, I can do the same thing much more space efficiently. I can just clone the boot image and only use additional space as things change. NetApp also gives me this same type of thin provisioning ability when it comes to allocating space for my application data.

Q: Do you see a lot of synergy between the functionality of VMware and NetApp?

Reinoud: Yes. The philosophy of both products is almost the same. One lets you consolidate and virtualize your servers and the other your storage. So far with NetApp I've done the consolidation part and a little bit of the virtualization. And it's just the same with VMware. With both you are so flexible. You can move your servers from one physical server to another; you can upgrade your hardware. When you want more resources you just add a new hardware box and you spread out your virtual servers. If one data center has a problem, you can just move all your critical servers to the other site automatically and very fast. VMware and NetApp work together very well in this regard.

Q: You've made the strategic decision to run your entire operation on NetApp storage. Is there anything that people who aren't familiar with NetApp would be surprised by?

Reinoud: I think the most important thing is the simplicity and flexibility, just as NetApp always says in its PowerPoint slides. It's really true. We run a 227TB shop with 17 NetApp storage systems with just half a person. Two people are assigned to managing NetApp, but each spends only a quarter of his total time on storage. This leaves them a lot more time to focus on other value-added tasks.

No other vendor enables you to be as efficient as Network Appliance.

Interested in learning more about UZ Leuven's deployment?

- [Technical Case Study: UZ Leuven Consolidates DB Storage Using FlexShare \(pdf\)](#)
- [UZ Leuven PACS Environment \(pdf\)](#)
- [UZ Leuven E-mail and Archival Strategy \(Symantec success story\)](#)



SHAWNA MEYER-RAVELLI

Director of Systems Engineering and Professional Services, Decru

Shawna Meyer-Ravelli oversees the strategic direction and delivery of sales engineering, professional services consulting, and training at Decru. As part of her job, she travels extensively meeting with customers to understand and address security problems. Shawna joined Decru in 2005 from Veritas Software (now Symantec). During her first year at Veritas, she worked in the CTO office exploring new storage technologies and platforms.

Lost Keys, Performance, and More:

Top Encryption Concerns Solved

By Shawna Meyer-Ravelli

With all the recent high-profile news stories about data loss, companies are waking up to the fact that perimeter security is not enough to protect sensitive and valuable data. This has heightened interest in encryption as a way to keep important data secure from prying eyes.

I've talked to hundreds of customers about implementing encryption solutions. One of the first things they typically want to know is whether or not encryption *really* addresses their business problems.

The answer for almost any compliance issue is, yes, encryption can help solve your business problem because it addresses the most rigorous requirements for protecting stored data.

For example, right now ensuring privacy for credit card transactions and information is a hot topic. Requirement 3.4 of the Payment Card Industry Data Security Standards (PCI DSS) *requires* that sensitive cardholder data be rendered unreadable *anywhere* it is stored. The standard goes on to establish stringent rules for key management.

Many regulations provide a "safe harbor" clause for companies that have implemented encryption that reduces or eliminates liability for companies acting in good faith when a breach occurs. Whether you need to comply with HIPAA regulations for healthcare, the Gramm-Leach-Bliley Act for financial services, the CASB1386 California Database Security Breach notification act, or DOD 5015.2 for federal government agencies, encryption is a generally accepted best practice for securing data.

Once an IT team is satisfied that its problem is one encryption can solve, they start asking questions about the impact of encrypting data. Four of the most common concerns include:

- How does encryption change my IT infrastructure?
- Will encryption degrade database or application performance?
- What if I lose my keys?
- Will encryption scale?

In this article, I'll tell you what I tell potential customers. I'll explain the trade-offs of various classes of encryption solutions, drawing from my field experience with Decru for examples of how encryption can be implemented successfully.

Concern 1: How does encryption change my IT infrastructure?

Everyone wants to know how implementing encryption will affect their existing infrastructure. Will it require radical changes? Will it be disruptive to ongoing



RELATED INFORMATION

- [Evaluating a Storage Security Solution \(pdf\)](#)
- [Considerations for Deploying Secure Backup \(pdf\)](#)
- [Webcast: Best Practices for Secure Backup](#)

Best Practices for Securing Backup Data

Customer privacy and corporate data security necessitate making backup and recovery data unconditionally secure. A recent TechTalk Online Event focused specifically on the use of encryption to protect backup data.

This Webcast covered a wide range of topics, including backup security trends, an evaluation of current encryption solutions, plus advanced backup encryption scenarios.

[Watch the recorded Webcast.](#)

Many questions were submitted via chat during the Webcast. We've compiled a complete transcript that covers a wide range of areas:

- Compliance
- Backup security
- Performance and interoperability
- Long-term security
- Key management
- Backup security deployment
- Solution architecture
- Purchasing

Sample questions include:

I'm concerned about getting locked into a single vendor for backup tapes.

operations? What processes and procedures will have to be revised? The answers depend in large part on how you approach encryption.

There are three basic options: host-based or application-based, appliance-based, and storage-based encryption.

- **Host-based or application-based solutions** are implemented on each server where encryption is desired, so little physical change to your hardware infrastructure is required to implement this type of solution. However, substantial changes may be required to your software infrastructure to implement encryption. This can take 12 to 18 months and substantial programming effort to implement.
- **Storage-based encryption solutions** implement encryption in special hardware integrated into the storage system. They work transparently with your existing servers, but you'll (obviously) be adding or replacing storage to gain encryption capabilities with this approach. Storage-based solutions have only recently become available, so maturity may be a concern.
- **An appliance-based approach** to encryption (such as that used by Decru) enables you to insert an appliance into your current infrastructure and begin encrypting right away with minimal disruption. Both new and existing applications and storage can be immediately accommodated. Decru is able to utilize all common types of back-end storage, including NAS, FC SAN, and IP SAN, which is a major advantage for heterogeneous storage infrastructures. (This may not be true of all appliance-based solutions.)

The appliance approach adds some complexity in the sense that you must manage a new, discrete device. In practice, implementing encryption in a separate appliance provides better role separation and delineation of responsibilities, and most of the customers I work with haven't found it to be a big factor.

Industry analysts such as Forrester Research recommend an appliance-based approach as the "best" encryption approach.

Concern 2: Will encryption degrade database or application performance?

Once again, this depends on your encryption approach. One of the potential disadvantages of host-based or application-based encryption is that it puts the burden on general-purpose host CPUs that must also perform many other tasks and is therefore typically slow. If you are writing data to tape drives and you want it compressed, you have to compress it before you encrypt it. This introduces a double hit since software-based compression is also CPU-intensive. Your servers may require additional CPUs to handle the added workload—the worst case scenario being a forklift upgrade of existing server hardware.

Storage-based encryption is a bit of an unknown quantity at this point since first-generation products are just starting to ship. The assumption would be that they perform at the line rate of the device.

Appliance-based encryption solutions implement encryption—and possibly compression—in hardware and can offer very good performance. For instance, Decru DataFort® appliances introduce only about a 100 microsecond port-to-port latency, which is the same as what you'd see with a network switch.

Decru also uses a number of optimizations to help ensure good performance. With early acknowledgment, Decru sends an acknowledgment to a writer while processing the write so the writer doesn't block; multiple threads are processed in parallel. Read ahead provides similar acceleration for reads, although there is a slight hit on read performance versus writes. In the case of tape writes, Decru actually improves performance by up to 30% in some cases because of early acknowledgment of writes plus full hardware-level compression prior to encryption. This allows Decru DataFort to stream compressed data at line rates to tape for improved performance versus writing directly to tape.

Concern 3: What if I lose my keys?

Let's say 10 years from now you need to restore data that you've encrypted today. If your keys are gone, your data might as well be. If you need to restore data for a regulatory or legal requirement, this loss could be catastrophic. Key management is one of the major features that differentiates one encryption solution from another.

How does Decru protect data from the administrator, especially when used to protect data in databases?

We're under the SOX, FERC, and NERC regulations. In IT, and especially here in IT Security, we know encryption is something that will be required, but it's still difficult to justify to the "powers that be." Do you have solution packs and/or strategies that can help with this piece?

[View the complete transcript \(pdf\)](#) to learn the answers to these and other questions.

The Tale of the Tape

What's the potential cost of not having encryption?

- Citigroup was recently fined \$30 million by the Federal Trade Commission due to an unencrypted tape that was lost by UPS. The tape was later recovered.
- Another top global bank was spending \$50,000 per month to rent Brinks trucks in order to securely transport backup tapes to a remote disaster recovery site. Total cost: \$600,000 a year.
- If a company loses 150,000 unencrypted customer records, it will cost the company approximately \$90 per record to deal with the problem. That's \$13.5 million in direct costs and doesn't include market capitalization damages, loss of customer loyalty, brand damage, or potential penalties and fines.

By comparison to these numbers, encryption starts to look pretty economical...and pretty smart.

Host-based	Appliance	Storage
• Complete software solution • Location at host • Software-based encryption	• Transparent to host applications • One-time setup and configuration • Strong logging and audit • Hardware-based encryption and key management	• Transparent to host applications • Standard with hardware
• CPU intensive • Host key management • Keys exposed in file • Complex to implement and manage • Poor coverage for mobile environments	• Requires additional device • Some vendors cannot handle encrypted hardware • Complete to implement and manage • Poor coverage for mobile environments	• Immediate key management • No support for heterogeneous environments • Look for the storage vendor • To work upgrade • Not backward compatible in many cases

A comparison of encryption approaches

Hardware failures or loss of encryption infrastructure due to disaster is a reality, which means that disaster recovery capabilities are a critical and essential part of key management no matter what solution you choose.

I'll describe the Decru® approach to key management as an example of one approach to protecting keys. Decru takes great care to make sure that keys are never lost or stolen, even in the face of failure or disaster. Each key is stored on each Decru DataFort appliance, which is certified to be secure and provides automated, self-contained key management. Keys are automatically and securely replicated to additional cluster nodes for protection against hardware failure.


A separate Lifetime Key Management™ (LKM) system provides an enterprise-wide secure key archive. If a Decru DataFort appliance should fail, a quorum of recovery smart cards (typically held by company officers) is presented to authorize the restoration of keys to the replacement hardware. Everything can be done through a Web-based interface, so the individuals holding the recovery smart cards don't have to be in one location. There is no single point of failure that puts your keys at risk of either loss or exposure, and the quorum requirement prevents spoofing, so your keys are protected against being stolen as well as being lost.

To ensure that keys cannot be stolen—the next worst thing that can happen after key loss—keys never leave Decru DataFort in clear text. This can be a particular problem in host-based solutions where the keys are often exposed in a general-purpose operating system and at risk of being stolen by intruders. Your data is only as secure as your keys, and your keys are only as secure as the OS storing them.

1. Each DataFort appliance provides automated, self-contained key management.

2. Keys are automatically and securely replicated to additional cluster nodes.

3. All DataFort appliances across the enterprise replicate keys to Decru Lifetime Key Management™ (LKM) system, providing automated, secure enterprise-wide key management. Recovery smart cards enforce quorum approval for sensitive operations.



Concern 4: Will my encryption environment scale?

Scalability for encryption infrastructure is really a question of two things: performance and key management.

Performance is a question of providing adequate encryption processing capability where it is needed. As we have already seen, this can be a problem for host-based encryption solutions, where you always run the risk of running out of horsepower.

Storage-based solutions are still an unknown entity, but are likely to provide encryption capability that matches the overall throughput capability of the storage they provide. It's possible that you might run into the opposite problem with storage-based solutions. For sites deploying storage infrastructures that require more capacity than throughput, you might end up paying for more encryption performance capability than you really need since each storage device has its own encryption hardware.

With appliance-based solutions, scaling performance is simply a matter of adding an additional appliance. This makes it simple to right-size your encryption capabilities to your needs. For instance, each Decru DataFort model has very well understood throughput capabilities, so sizing is simply a matter of determining the throughput you require at peak load and deploying the right number of appliances in the data path.

Scaling key management is just as important as scaling performance. Decru has made significant effort to make key management both scalable and simple to manage. For large enterprises, multiple LKMs can be deployed in a hierarchy to enable management of globally dispersed data centers from a central location. The Decru OpenKey Partner Program allows third-party encryption devices to also be managed

One of the top international banks uses this approach with Decru technology deployed across the entire enterprise. Decru DataFort clusters are used to encrypt tapes in branch offices around the world. Using Decru policies, this company is able to manage the operation and setup of every device around the world from its corporate headquarters in the United States.



Evaluating Potential Encryption Solutions

By this point, it's probably clear that there are some important criteria you should think about when evaluating encryption solutions:

- Does the proposed solution satisfy the regulations that govern your business?
- Will the solution fit well with your existing IT infrastructure?
- Can the solution deliver the performance you will need now and in the foreseeable future?
- Are the key management capabilities adequate to protect encryption keys from being permanently lost or stolen?
- Can both performance and key management capabilities scale meet your operational needs?

The good news is not only that encryption can help you meet security requirements, but also that you don't need to disrupt your entire infrastructure to implement it.

Comment on this article

**NICK TRIANTOS**

Global SAN Systems Engineer, NetApp

A member of the elite Global Systems Engineering group, Nick helps top enterprise companies solve their toughest technical challenges. Nick has been in systems or support engineering roles for nearly 16 years, including positions at HP as an account support engineer (Server group) and presales technical consultant (Storage group). Nick maintains a [blog](#) and has previously authored Tech OnTap articles on [iSCSI vs. FC SAN](#) and [SAN implementation tips](#).

Applying Innovation to FC SANs Focus on Dynamic Queue Management

By Nick Triantos

FC SANs will never be simple, but NetApp helps to make them less complex.

Instead of requiring separate storage islands for separate applications and storage needs, for example, NetApp enables you to easily support multiple applications using different protocols within a single unified storage environment. And that's only the beginning.

NetApp offers an unparalleled degree of integration and automation across the storage, server, and application layers via solutions, expertise, and processes that are consistent throughout the full line of NetApp entry-level, midrange, and high-performance storage systems. This has enabled NetApp to eliminate many manual storage-centric tasks in favor of application- and server-based management, which helps to keep the burden of routine day-to-day tasks from falling too heavily on the shoulders of the storage administrator.

Industry-leading NetApp management tools such as SnapDrive® and SnapManager® for Oracle®, Exchange, and Microsoft® SQL Server automate standard processes and enable server and application administrators to manage their own data. For example, [SnapManager](#) can be scheduled to create Snapshot™ copies, while [SnapDrive](#) for UNIX® or Windows® can be scheduled as a cron job or by Windows Task Scheduler to perform Snapshot copies, provisioning, LUN expansion, and so on. These tasks normally require multiple steps and typically cannot be done directly from a host.

Still not convinced? A perfect example of how NetApp brings its trademark simplicity to the Fibre Channel SAN space involves host-side and target-side queue depth management.

In a typical SAN environment, manually setting host queue depths can be a complicated process. Determining the appropriate setting requires specific knowledge of application throughput and response times, yet all of the vendor documentation I've seen explicitly specifies a standard host setting for host queue depth. This assumes that all initiators in the fabric will do the same type of I/O with the same I/O characteristics at the same time, which is a highly unlikely scenario.

Providing a standard queue depth setting is simply a "shot in the dark" approach that hopes the assigned queue depth will provide adequate application performance. But what if it doesn't? Then a host-side reconfiguration will be necessary, in addition to, depending on the platform, a subsequent reboot. This cumbersome procedure is the result of two things:

- Planning without a full understanding of application requirements
- Inadequate protocol monitoring and management features

To address this challenge, the NetApp engineering team implemented dynamic queue

RELATED INFORMATION

- [Tech Report: Thin Provisioning in a NetApp SAN Environment \(pdf\)](#)
- [Tech Report: Storage Consolidation with 4Gb FC SANs \(pdf\)](#)
- [More Technical Reports on SAN](#)

Thin Provisioning for NetApp SAN Environments

One of the big disadvantages of traditional disk arrays is that they force you to allocate dedicated storage space to a disk volume or LUN when you create it. Since it's often hard to gauge the amount of space you'll need up front, you end up over-provisioning and wasting valuable disk space.

In contrast, when you create a LUN on a NetApp system you don't have to dedicate specific disk blocks. Instead, blocks are allocated only as data is written. This way, multiple LUNs can flexibly share the same pool of free storage. You simply add more capacity when free storage gets low, and you can painlessly grow a LUN if more space is required.

Learn more. [Read the Technical Report.](#)

Redefining FC SAN Efficiency

There's more to deploying a SAN than just cabling together servers and storage.

A recent TechTalk Webcast describes how you can leverage NetApp storage systems and the virtual SAN (VSAN)

management, which shifts queue depth management from the host to the array target port.

What Is Dynamic Queue Management?

Simply put, dynamic queue management controls queue depths from the array side. By monitoring QFULL conditions, application response times on a per-LUN basis, and queue lengths, dynamic queue management automatically adjusts queue depth based on application I/O requirements. The goals of dynamic queue depth management are to:

- Provide a centralized location where queues can be monitored and reconfigured in real time
- Prevent host-side queue depth reconfiguration, which in some cases requires a host reboot
- Prevent initiators from flooding target ports while starving other initiators from LUN access
- Ensure that initiators have access to target queue resources on as-needed basis

With dynamic queue management, Data ONTAP® calculates the total amount of command blocks available and reserves the appropriate number for an initiator or a group of initiators based on the percentage you specify (0–99%). In the process we also specify a reserve queue pool, where an initiator or a group of initiators can borrow queues when they are needed by the application. On the host side, we simply set the queue depth to its maximum value.

The benefit of this practice is that it takes guesswork out of the picture and helps ensure that the application will perform at its maximum level without unnecessary host-side reconfigurations, application shutdowns, or host reboots.

Fibre Channel SANs Simplified

At NetApp, we believe that innovation is the key to minimizing the complexity of Fibre Channel SAN environments. While other vendors require administrators to statically configure (and reconfigure) queue depths—resulting in increased administrative hassles and potential application downtime—NetApp administrators don't even have to think about it. Best of all, dynamic queue management is just one example of the host of NetApp innovations designed to make your FC SAN simpler to deploy, monitor, and manage.

If you're not using NetApp storage in your data center, you don't know what you're missing.

Comment on this article

capabilities of Cisco switches to create a highly available and efficient SAN infrastructure, decrease storage provisioning time, and ensure application availability.

Ian Perdue of VECO describes the deployment of these technologies to streamline VECO operations. [View the Webcast](#) to find out more.



Consolidating Your FC SAN

In this TechTalk Webcast, Eric Tomasi, VP of Infrastructure for Folksamerica Reinsurance Company, explains how his team leveraged a NetApp FC SAN solution to achieved 80% storage growth without increasing IT headcount.

Today the company can complete a full remote recovery in less than 4 hours. Watch the Webcast to [learn more](#).





CRAIG JOHNSTON

Chief Product Architect—Heterogeneous Data Protection, NetApp

One of the original founders of Alacritus, Craig Johnston joined Network Appliance after the company was acquired in 2005. Craig has over 15 years of experience in the area of data protection and has been employed by many of the early innovators in the area of backup software, including Delta Microsystems (the developer of BudTool), PDC, IntelliGuard, and Legato.

No Compromises: How NetApp Built a More Efficient, Higher Performance VTL

Q&A with Craig Johnston

Virtual tape libraries (VTLs) enhance existing backup procedures without disruptive changes—simply drop a VTL system into a tape-based backup environment and reap the benefits of simplified management; shorter backup windows; and rapid, reliable restores without changes in processes, procedures, or your existing architecture.

In October of 2006, NetApp announced that it has extended the capabilities of its original VTL products to include hardware compression, greater performance, greater capacity, and a lower entry point. To learn more, Tech OnTap caught up with Chief Product Architect Craig Johnston.

Q: Hardware compression is a major component of the recent NetApp VTL announcement. Why is hardware compression an advantage?

Craig: First, it's important to understand that to date companies have been forced to choose between efficiency and performance. Compression makes great economic sense because it allows you to store more data—often twice as much—on the same amount of disk space. Keep in mind, however, that the primary reason that people implement VTL technologies is to increase backup speed and make their backup windows as short as possible. Compression methods that interfere with performance defeat the purpose.

Every other major VTL system on the market uses software compression, which places a significant burden on the processors and slows the overall speed of the VTL. In some cases, VTL throughput is cut in half when software compression is enabled. In the worst case scenario—enabling software compression with data that can't be easily compressed—throughput on competing systems drops over 85% when software compression is enabled.

By building compression into the hardware, the NearStore® VTL gives you double the usable capacity for the same dollars without compromising performance. Our VTL systems can achieve line speed compression with throughput upward of 850MB per second. In fact, our throughput actually goes *up* with hardware compression enabled since we're writing less to disk.

Q: NetApp is the first major storage vendor to offer hardware compression. What was involved with implementing this on the NearStore VTL platform?

Craig: Hardware compression is implemented on a PCI-Express card. Each card has four separate compression chips as well as an FPGA that coordinates data flow and balances load. As data comes in to be compressed, it is buffered in a circular "mover" buffer. When enough data accumulates, it triggers a compress cycle, which streams the data through the compression board and back into a separate mover buffer in



RELATED INFORMATION

- [NetApp VTL Product Overview](#)
- [VeriTest Performance Comparison](#)
- [And Then There Were Three \(TechWorld article\)](#)
- [Webcast: Reduce Backup Costs](#)

What Makes the NetApp NearStore VTL Unique?

- **Hardware compression** that dramatically increases capacity without degrading performance
- **Self-tuning** for optimal performance without manual intervention
- **Smart sizing** for more efficient utilization of physical tapes
- **Shadow tapes** reclaim the least recently used virtual tapes first for accelerated restores
- **Direct tape creation** writes directly to the physical tape in native format, which can be read without a VTL system
- **In-place upgradeability** from entry-level systems to higher performance systems
- **A dedicated engineering team** that owns every aspect of VTL development and delivers the capabilities that customers need


Reduce the Cost of D2D Backup

Not all VTLs are created equal.

In a recent Webcast Enterprise Strategy Group Senior Analyst Tony Asaro and NetApp General Manager of Heterogeneous Data Protection Krish Padmanabhan explained how to

system memory.

Compressed data is accumulated in memory until we have enough to do a full stripe write to disk. The compression driver coordinates closely with the VTL and RAID implementation so that we can write data very efficiently to get maximum I/O throughput to the disk subsystem.



The process is a little more complicated for decompressing the data. Since the user is likely to start reading somewhere other than at the start of the tape, we keep track of points on the tape, about every 10MB, where we can start decompressing. For hardware compression, we needed even more information. The engineer who wrote the compression driver, Jasbir (Jazz) Manotra, found that we needed to insert a small header into the compressed data to keep track of each compressed block.

Q: If hardware compression offers such obvious benefits for VTLs, why was NetApp, a relatively recent entrant into the VTL field, the first major storage vendor to introduce it?


Craig: I think it's natural to see hardware compression as a necessary part of a complete customer solution. You'll probably see most of the other major VTL vendors follow suit in the near future. It really comes down to how quickly you can innovate and implement new features. NetApp has a business unit dedicated to VTLs, with complete control over the technology. We're not dependent on OEM relationships the way some other VTL vendors are, which enables us to incorporate customer feedback and bring new features out quickly.

Experience is another key factor. We have a dedicated VTL engineering team with an impressive amount of experience in the physical tape drive arena.

leverage NetApp NearStore VTL solutions to:

- **Reduce the cost** of enterprise D2D backup by 50% to 67%
- **Increase the reliability** of the backup process
- **Increase D2D backup** performance up to 10 times
- **Simplify backup** and recovery management

[Watch the TechTalk Webcast.](#)



Q: Earlier you mentioned that throughput increases with hardware compression. What else have you done to enhance the performance of the new platform?


Craig: The NearStore VTL is based on the NetApp hardware platform, which means

that it's been tested and proven in thousands of deployments around the world. From a hardware standpoint, the new platform has a number of improvements over the previous models. These include more powerful processors, a faster memory bus, PCI-Express driven by a state-of-the-art chipset instead of PCI-X, and 4Gb Fibre Channel instead of 2Gb. The NetApp Platform Development team has produced a great new platform that we have the privilege of using for the VTL700.

From a software standpoint, we focus on using good practices to make sure we aren't squandering what the hardware engineers have given us. When not using hardware compression, data is DMA'd into a buffer in system memory from which it is eventually DMA'd out, avoiding unnecessary copies. Obviously with hardware compression the data has to pass through the compression board, but again we avoid time-consuming memory operations as much as possible.

Our data path is optimized for large sequential I/Os. We always do full stripe writes, so we only have to calculate parity once for each stripe. NetApp has a patent pending on the way we do RAID in the VTL.

The end result is that the NearStore VTL delivers up to three times the performance relative to products from competing vendors with compression enabled and provides compressed throughput performance on par with other VTL systems on the market that do not offer compression.



Comparison of NetApp VTL700 and EMC DL710 performance (source: NetApp performance testing).

Q: Another part of the VTL announcement involves the new VTL300 platform. Why did NetApp decide to introduce an entry-level platform?

Craig: It's critical that customers understand where the VTL300 fits. This is an entry-level enterprise platform. Although it has limited capacity and performance, the VTL300 includes all the enterprise features of other models in our VTL product line.

It's a great enterprise entry point at an attractive price for enterprise customers interested in seeing what a VTL system can do for them. And, this platform offers NetApp's trademark investment protection. As your requirements change, the VTL300 can be upgraded in place to a VTL700. I don't think any other VTL vendor offers this kind of upgradability.

Q: Does the new platform provide the same self-tuning, tape smart sizing, shadow tape, and other capabilities?

Craig: Absolutely. The new platform has all the features of the original, including self-tuning storage, tape smart sizing, and shadow tapes.

The self-tuning capability of NearStore VTL, for example, allows us to dynamically balance workloads. In other VTLs, virtual tapes reside at fixed locations on disk, and load balancing is static. You might have many backup streams going to one disk and few to others. Hot spots develop, and manual tuning is necessary to optimize performance. With NearStore VTL, virtual tapes don't reside at fixed locations. Backup streams are dynamically balanced across disk LUNs for optimal performance without

Other VTLs also don't understand physical tape drive compression. They use fixed capacity virtual tapes, and when those are written out to physical tapes, the result is a bunch of tapes that are only partially full. NearStore VTL tape smart sizing samples the backup data and adjusts the size of each virtual tape accordingly, based on how compressible the data stream is. The result is fully utilized physical tapes when virtual tapes are written out.

Tape smart sizing works exactly the same way with hardware compression enabled. When we write a physical tape, we decompress the data off disk before we send it to the tape library and let the tape drive do its own compression. This may seem like extra, unnecessary effort, but this way you end up with a standard tape that is in the native tape format. In a disaster, you don't have to have a VTL around to read your tapes.

Another thing we do is allow the user to save tape copies on the VTL disk as shadow tapes in a backup application-compatible manner even after the corresponding physical tapes have been created and off-sited. About 90% of restores happen in the first two weeks, so when you write a physical tape from your VTL we keep the virtual tape around and reclaim space only as needed. Then we reclaim the least recently used shadow virtual tapes first. As a result, if you need data from a tape, it's likely to still be available from the VTL—at fast VTL recovery speeds—saving you from having to go to the physical tape, which might be off-site already.

Q: Can you give us a sneak peek into what the VTL engineering team is working on now?

Craig: We're developing a number of new features. On the hardware side, we're focused on increasing usable disk capacity and increasing the total number of I/O ports available for connectivity. We're also adding ACSLS support so that the VTL can communicate directly with large Sun™ and STK tape libraries that use that protocol.

The most exciting and innovative developments involve a new deduplication capability that is planned for release in 2007. Deduplication will increase the amount of data that can be stored on a VTL by even greater amounts than we offer today, but will do so in a way that still delivers the enterprise-class backup performance required in the data center. The problem with existing VTL deduplication methods is that they are very slow. We're working on an algorithm that is optimized for large sequential I/Os. In ideal cases, this will mean that we can do deduplication at the same speed we could do nondeduplication writes.

Question: Any last comments?

Craig: In my opinion—and I'll freely admit that I'm biased—the NetApp VTL solution is easily the most innovative on the market. Every vendor offers its own flavor of VTL, with varying features and functionality, but NetApp is the first vendor to consistently invest in and deliver features that have a substantial impact on the economics of VTL without weakening the fundamental data center VTL performance value proposition. All of us on the VTL team are looking forward to an exciting 2007.

NearStore VTL Specifications	NearStore VTL300	NearStore VTL700	NearStore VTL1400
Max. Sustained Write Throughput (No Compression)	300MB/sec	600MB/sec	1200MB/sec
Max. Sustained Write Throughput (2:1 Compression)	500MB/sec	850MB/sec	1700MB/sec
Max. Raw Capacity	35TB	84TB	168TB
Max. Effective Capacity (2:1 Compression)	53TB	128TB	256TB

NetApp NearStore VTL Product Comparison (Source: [NetApp VTL Product Overview](#))

Comment on this article

TECH ONTAP HOME PAGE



Eyal Traitel
Storage Consultant

Eyal has over 9 years of experience working with NetApp storage, including several years as an admin at Motorola Semiconductors Israel (now Freescale). In his current role as a storage consultant, Eyal is involved in a variety of NetApp projects and activities. He is the creator of the [NetApp Filers Blog](#), an independent blog featuring real-world suggestions and tips for NetApp storage administrators.

"I am still an active Toasters user. After all these years, it seems it became part of my life."

Tool of the Month: Toasters Discussion Group

Every month, Tech OnTap showcases a free tool that just might make your life a little easier. [Recommend a tool](#) and get a FREE NetApp jacket.

Author: Tom Yates started the [Toasters](#) discussion list in 1997 because he "became very frustrated at my inability to find any other toaster owners out there in the big world."

What it is: Toasters is an independent mailing list tuned for "filer geeks" — a forum for system and storage administrators who are passionate about NetApp technology but who also need assistance and advice from each other about new products, new features, implementation tips, technical questions, and problems.

Several hundred users are registered to receive the Toasters mailing list. The load? About 18,000 messages have been sent over the period of close to 10 years, which averages out to about 5 messages a day.

How it works: Send an e-mail to lists@mathworks.com containing the body text subscribe toasters. More instructions can be found [here](#).

Why it's cool: Toasters is a stage where everyone can participate, from the most junior storage admins and NetApp newbies to legendary NetApp engineers. NetApp co-founder Dave Hitz, VP of engineering Brian Pawlawski, and former NetApp technical director Guy Harris are among those who have jumped in to answer Toasters questions over the years.

Caveats: Toasters has not been historically tied to NetApp in any formal way. It is not officially sponsored by NetApp, and contains much tribal knowledge that may not be endorsed by NetApp and may or may not be correct. Spam has been a topic of concern over the years as well, although recent efforts have dramatically improved the situation.

[Join the mailing list](#)
[Toasters archive \(hosted by Nabble\)](#)

Want to recommend a tool? If we feature it, you get a FREE NetApp jacket.

✱ **Nominate a tool**



RELATED INFORMATION

- ✱ [Toasters Archive](#)
- ✱ [The NetApp Filers Blog](#)
- ✱ **Previously Featured Tools:**
 - ✱ [Swatch and Kiwi](#)
 - ✱ [Firefox](#)
 - ✱ [Cacti](#)
 - ✱ [MRTG for Filers](#)
 - ✱ [Nagios](#)
 - ✱ [SIO](#)
 - ✱ [ToasterView](#)

Sample Toasters Query

