

Tech OnTap

Insights for Simplifying Data Management



November 2005

WHAT'S INSIDE

- 10 Reasons IT Projects Fail
- Tom Mendoza's OpenWorld Keynote
- 5-Min. RPO, 15-Min. RTO
- Demo: FlexVol™ and FlexClone™
- NetApp's Secret Weapon: iSCSI

How to Talk to a CIO (If You Must)

Michael Friedenberg, President of CXO Media

While your CIO can't fix a downed server, he or she is awfully important when it's time to fund your next project. Win budget, resources, and maybe even a promotion by learning how to speak your CIO's language.

[More :-](#)



BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"One military customer told me if his data isn't available, people die. It didn't occur to me until later to ask whether he meant his own people or the vendor who sold him the equipment."

- Dave's Blog

TIPS FROM THE TRENCHES

Cookbook: Minimize DAS

Nick Triantos, Global SAN Systems Engineer, NetApp

"One size fits all" doesn't work for T-shirts, and it certainly doesn't fit your complicated IT environment. A top NetApp SE steps you through how to evaluate your existing infrastructure and application requirements (uptime, IOPS, bandwidth, and more).

[More :-](#)

Ten Reasons IT Projects Fail (and Ten Ways to Avoid Them)

Dale Defilippi, Supervisor of Hardware Engineering, IGT

In just six weeks, IGT's IT team rebuilt or updated nearly every engineering server, implemented disk-to-disk-to-tape backup, and proved 30-minute remote recoverability. Learn how the team did it and — more importantly — how you can make your next project equally successful.

[More :-](#)

Disaster Recovery, NetApp Style 5-Minute RPO, 15-Minute RTO

Devinder Singh,

Global Architect and Alpha Geek, NetApp IT

Learn why SNIA asked NetApp IT global architect Devinder Singh to describe the NetApp business continuity strategy at Storage Networking World.

[More :-](#)

NETAPP SECRET SAUCE


- Data ONTAP 7G White Paper

Data ONTAP™ 7G establishes a new industry benchmark for how data can be deployed and managed.

DRILL DOWN

▪ **Beaumont Hospitals Case Study**

Beaumont Hospitals increased storage utilization from 43% to 77% and enabled independent database backups.



NetApp iSCSI Storage Center


- See how NetApp stacks up versus the competition.

KICK THE TIRES

- **SQL Server™ 2005 Launch Tour**
Invitation code: Network Appliance.

- **Changing the Economics of Storage**
iSCSI Webcast series.

ORACLE OPEN WORLD RECAP



**Tom's Oracle OpenWorld Keynote :-
OpenWorld Demos and Presentations :-**

FREE STUFF ALERT

This month we're giving free t-shirts to Tech OnTap members. The catch ... you have to find the offer in this newsletter!



MICHAEL FRIEDENBERG

President and CEO of CXO Media

Michael Friedenber talks with CIOs every day in his role as the President and CEO of CXO Media, which produces magazines and Web sites focused on top executives and C-level management. In this Q&A Michael shares insights based on research from CIO magazine and the CIO.com Web site plus his more than 16 years of experience in high-tech media.

How to Talk to a CIO (If You Must)

Q&A with Michael Friedenber

Q: We often hear front-line IT staff complain that senior management doesn't understand their challenges and that it is difficult to get executive mind share and sponsorship for key initiatives. What do you think is the biggest disconnect between IT teams and their CIOs?

Michael: Too often, you have people trying to solve IT issues that may not be driving the business forward. They don't realize that it's no longer business & technology. The ampersand is gone — it's just business technology.

CIOs need to clearly understand the business strategy first and then create IT solutions that fulfill the business goals.

Q: How can IT teams get executive-level support for major projects that involve disruptive technology or significant upgrades? What type of information do CIOs need to make these kinds of decisions?

Michael: The margin of error that a CIO now has is basically zero. Any purchase has to go through a careful risk and opportunity assessment. On the risk assessment side, some of a CIO's questions are:

- If I invest these dollars, will the project not only keep the business going but also allow it to grow?
- Is the vendor going to be reliable, and will it be around next year? In five years?
- Is the project going to have a 5X or 10X multiplier on the service side?
- Is the SLA in line with expectations, and is everyone in full agreement with what is about to be signed?

The opportunity assessment is harder to get your arms around. CIOs need to understand:

- What will we be capable of doing that we weren't able to do before?
- If we don't do this project, what does that mean to us on a competitive basis?
- Will investing in this technology open up new markets for us?

RELATED INFORMATION

- [CIO Magazine](#)
- [How to Translate Geek Speak](#)
- [How to Do Customer Segmentation Right](#)
- [CIO.com Executive Podcast: John Baldoni's How to Ask for a Raise](#)

Get Noticed and Get Promoted

- Demonstrate clear understanding of the business goals and the ability to translate how technology helps achieve them
- Build great relationships — up, down, and out — within the organization and with your customers
- Think strategically and — even more important-show a track record of successfully executing upon that strategy

Best Practice: RBC Royal Bank

A perfect example of mapping IT to business needs is a

- Will this allow us to increase customer affinity and loyalty?

Providing these answers up front helps focus your CIO's attention on your project. Most organizations have a tremendous amount of backlog — a recent CIO.com tech poll suggests our readers have an average of 38 “top priority” projects currently in backlog — and one of the biggest frustrations a CIO faces is how to prioritize which project to work on first.

Q: What other areas must be addressed for CIOs to support major projects or initiatives?

Michael: CIOs need to know if the technology investment aligns with the business priorities and the business roadmap. This means that the IT organization needs to translate what the technology will ultimately do for the business. It's remarkable to me how few IT managers or staff spend any time with their actual customers to understand the needs of the people they ultimately serve. The IT team needs to communicate whether or not the project is going to make the organization a better company to do business with, if it's going to make the company more nimble, make it faster, allow increased margin on company products, or help lower costs. This is the mind set — or at least the language — you should be using with your CIO. Does it deliver business value and customer value?

Every organization right now is focused on growth, yet on average 80% to 85% of IT budgets are allocated toward items that just help keep the lights on. Balancing innovation and growth when only 15% of the budget is truly allocated toward these areas is one of the CIO's core pressures.

CIOs also care about the potential return on investment and whether or not it will integrate and scale in the proper way. They also want to be sure that everyone else within the purchase decision has bought into the solution. You need to have buy-in from each line of business (LOB) that will be touching and using this application or solution.

Q: Do any examples come to mind?

Michael: A perfect example is a CRM application. This is a customer-facing application that is not only driven by IT but also by LOB and corporate management. If you don't get all those camps bought into the solution you're proposing, along with its deliverables and expectations, I think you're destined for a very dark and unenjoyable process.

Q: Are there any instant turnoffs when CIOs are evaluating potential projects?

Michael: “Cool, hot technologies.” In today's market, there is tremendous pressure to be perfect on every project. We see a trend that CIOs of large-scale enterprises would rather be fast followers than early adopters, especially if being an early adopter has the potential to be painful and costly.

Q: After an IT project has been approved and is moving forward, should the IT staff continue to involve CIOs?

Michael: Absolutely. IT teams should report on small but quantifiable wins whenever possible. From a CIO perspective, there is nothing more effective than to immediately show the rest of the organization that the decision and the trust that were conveyed to the IT organization are being proven.

customer segmentation project undertaken by RBC Royal Bank. Typically banks view the 18-to-35 demographic as primarily unprofitable.

RBC, however, used technology to dig deeper into this age bracket to identify medical school students, for example, and determine which contacts have the most potential of becoming good customers when they are 45 or 50.

Within a year, RBC's market share among customers in the 18-to-35 market segment shot up from 2% to 18%, and the bank's revenue per client is now 3.7 times that of the average customer.


In board meetings CIOs may be expected to give daily, weekly, or monthly reports about how major IT projects are doing. Wherever a CIO can show when something is working (or, when it's not, what the solution is to turn it around), it helps build the CIO's reputation with the rest of the executive team. I believe that if IT organizations did more to promote their achievements, more companies would view IT as an enabler rather than a cost center.

Q: Is there anything else people at the lower levels of an IT organization don't understand about their CIOs?

Michael: Sometimes people within the ranks of IT don't understand that the CIO has officially crossed over into the CXO grouping. Three or four years ago CIOs were usually classified as part of the IT organization, but now they are seen as core members of the Executive Committee.

Today, the CIO is the multilinguist of the organization. That's what makes this role so unique and important to the business. Not only do CIOs need to talk to the CXO and the Board of Directors; they also need to translate the needs of their LOB, IT, and customers. Trying to connect the dots between these four constituents is what I believe keeps CIOs up at night.

© 2005 Network Appliance, Inc.



DALE DEFILIPPI

Supervisor of Systems Technology, Hardware Engineering, IGT

With over 15 years of IT experience Dale Defilippi has established a track record of successfully completing complex, multi-million dollar projects on time and on budget. Dale currently oversees IT for the hardware engineering department of IGT, a multibillion-dollar manufacturing company and the world's foremost supplier of gaming machines and gaming monitoring systems.

Ten Reasons IT Projects Fail (and 10 Ways to Avoid Them)

By Dale Defilippi

Pitfall 1: Lack of Commitment

If you aren't passionate about your project, why should anyone else be? At IGT, the senior management team required a personal guarantee before committing funds to such a large-scale project. I staked my reputation (and possibly my job) on this project's success.

Pitfall 2: A Blind Eye to Compromise

Look for areas to make compromises, especially when it comes to budget. For example, instead of replicating to a high-end expensive Fibre Channel array at our remote DR site, we recommended a less expensive architecture. We decided to reuse slower servers that had been production servers at the main site. Instead of Fibre Channel, we did SATA disk. Once management realized our DR implementation would be a third of the price of our main data center, it became an easier sell.

Pitfall 3: Overestimating Your Abilities

Think big, but don't be afraid to call on experts to help fulfill your vision. We knew it would be impossible to complete a project of this scale with my IT team, so we planned up front to enlist the help of our vendor's professional services organization. This enabled us to build a very skilled group with a broad range of knowledge — ranging from SAN management to backup with the VERITAS product — who assisted us with project planning, preparation, and the actual implementation.

Pitfall 4: Believing Vendor Promises

Let's face it — vendors aren't always 100% truthful. Always make them prove that they can deliver on their promises. After narrowing an initial list of 18 vendors down to a short list of three, we gave each vendor three weeks to perform an exhaustive proof of concept.

I personally visited each vendor's site and asked the vendor to show us

RELATED INFORMATION

- [IGT Protects Irreplaceable Intellectual Capital \(PDF\)](#)
- [Achieving Five-Minute RPO and 15-Minute RTO](#)
- [SAN Booting with Windows Server Operating Systems \(PDF\)](#)
- [Cookbook: Minimize DAS](#)
- [NetApp Data Protection Solutions](#)

IGT's Six-Week Implementation

- Installed 12TB SAN/NAS system
- Relocated community share and home drives to networked storage
- Reengineered and rebuilt backup strategy to enable disk-to-disk-to-tape approach
- Migrated local C: drive to SAN
- Implemented server boot off SAN functionality
- Updated all other servers connected to mission-critical systems
- Integrated multiple technologies (NSI Double Take, NetApp SnapMirror®, VERITAS® NetBackup)
- Built Las Vegas DR site,

a subset of our own data up and running in all our databases and file shares and prove that it could replicate the data and fail it over in our presence. One supplier opted out, while another completed roughly two-thirds of the requirements but was unable to boot off the SAN or demonstrate a clean recovery of an Oracle® Database at the DR site. NetApp was the only vendor that was able to do everything the sales team claimed it could within our deadline.

Pitfall 5: Unrealistic Timetable and Expectations

If you have a hard end date, try to build extra time into the schedule at the beginning or consider removing more complicated requirements from the first phase of your deployment. I wish we could have extended our timeline 30% or at least worked to get the PO cut earlier. As it was, we burned out a few people during the implementation. Part of the difficulty came from our efforts to make booting from the SAN work how we wanted. In 2003 we were one of the first people to boot off SAN using NetApp storage. In hindsight, knowing how complex that effort became, I would have probably postponed implementing bleeding edge functionality until later on.

“Think big, but don’t be afraid to call on experts to help fulfill your vision.”

Pitfall 6: Navigating without a Road Map

Before beginning a project like this, know the intricacies of your infrastructure in and out. Make sure you know how each of your applications interacts with the operating system and the disk storage subsystem, as well as how they handle network I/O. Learn about any possible latency issues on both sides before you begin a DR implementation. When you start moving drives and data from one point to another, you want to make sure the OS and the application know where everything is when it goes from your direct-attached storage to the new SAN. We had one somewhat unknown machine that had only been in-house for five or six months, and figuring out how to replicate it took up a lot of time we didn’t have during the final weeks of the project.

Pitfall 7: Inability to Expect the Unexpected

Remain creative and resourceful as the project unfolds: expect the unexpected. During the final week we realized that latencies and packet transfers were too slow to effectively complete the mirrors between the two sites. We had adequate bandwidth, but because of the physical cabling limitations of our service provider’s WAN routing it was actually a longer distance from Reno to Las Vegas than it would be from San Francisco to New York. An average latency of 28 milliseconds caused the mirroring process to take upward of three hours and consume more than half of our bandwidth.

NetApp took ownership of the problem with minimal finger-pointing. Within 24 hours we were testing eval units from a third-party vendor — Peribit Networks — helped us reduce latency over the WAN. This immediately resolved our mirroring issues. In my




The IGT team that implemented the consolidation and DR project (back to front): Barry Sethman, Dale Defilippi, Luke Mun, Jason Porterfield, Jason Hedrick, Brian Riffel, and Ruddi Kitchen.

including a 9TB nearline storage system

- Built and implemented VERITAS master server in Las Vegas

Results

- Completed on time, on budget
- Proven 30-minute remote recovery
- 450 hours of downtime eliminated



IGT Success Story

opinion, this experience underscores the benefit of a vendor who teams with other best-of-breed companies rather than trying to do everything on its own.

Pitfall 8: Failure to Adapt to Reality

Theorizing how a process should work and watching it work in live production environments are two different things. Don't be afraid to change your initial operating procedures once new facts come to light. When we originally set up our mirror between the two sites, we planned to perform the mirror once every 90 to 120 minutes. However, mirror performance was slower than expected because of how much data changed between Snapshot™ copies. Although we expected to see a delta of about 10% to 15%, we actually saw more in the range of 30% to 40% changed bytes.

We decided to experiment with the time between mirrors to see if that might improve performance. We first kicked it down to an hour, which gave us better performance at roughly the same bandwidth usage. Then we tried every 15 minutes, which got better. One weekend, we kicked it down to mirror every minute. This allowed the data transfer to take place in about seven minutes, with no increase in bandwidth. We now have about 7TB of data we continually keep synchronized between our active and replicated sites, with just a seven-minute lag time.

Pitfall 9: Failure to Test, Test, and Test Again

I cannot underestimate the importance of performing continuous testing, especially for a DR implementation. You have to know your data is accessible if you need it. After you get the initial mirror going, you want to make sure you can power the system up and make it operate at your DR site and that the data you might be in the middle of replicating doesn't become corrupted in the process. We run failover tests every quarter and do additional testing once a year with our auditors present so they can see the solution and that it's still functioning.

Pitfall 10: Buying Band-aids Instead of Solutions

The problem with spot solutions is that in a few months you're looking for another bandaid. After two years, most technology implemented is often due for replacement, but I'm confident that we implemented an architecture with the flexibility to support our organization for a while to come. We have a lot of ideas about where to take it from here. We've nearly tripled storage capacity since our initial deployment. We also recently upgraded to Data ONTAP™ 7G to enable instantaneous cloning of our production environment for development. That has been worth its weight in gold. It used to take us hours to create database clones, so this has saved a lot of time. Overall, we believe that NetApp has very impressive technology and we are just touching the surface of what it can do for us.

Interested in talking to a NetApp representative or partner about your upcoming storage projects?

Contact Me!

Oracle OpenWorld Recap

Nearly 30,000 people attended Oracle OpenWorld last month, making it the most popular OpenWorld ever.

Now, as a Tech OnTap member, you've been granted an all-access pass to the information that conference attendees traveled thousands (and tens of thousands!) of miles for.

- Hear Tom explain the vision that revolutionized the storage industry.
- See for yourself that it's possible to increase and decrease volumes on-the-fly with no performance hit.
- Learn what's top of mind for NetApp customers and top technical experts from NetApp Engineering, IT, and other areas.

▄▄ See Oracle OpenWorld presentations



Download:

- ▄▄ [Full keynote - Real \(45 min.\)](#)
- ▄▄ [Full keynote - Flash \(45 min./103MB\)](#)
- ▄▄ [Highlights - Flash \(4 min./10MB\)](#)


Excerpts:

- ▄▄ [Applying the Cisco "Magic" to Storage](#)
- ▄▄ [How Oracle IT Helped](#)
- ▄▄ [Define Data ONTAP 7G](#)

Oracle OpenWorld Keynote:

Simplifying Data Management: Doing More while Spending Less

-Tom Mendoza, NetApp President



- ▄▄ [Watch the demo](#)

DEMO

Advantages of FlexVol™ and FlexClone™ in Database Environments

This swing bench demo shows:

- Performance benefits of running OLTP database loads on aggregated storage
- Ability to increase and decrease volumes in seconds
- Ability to create database clones in under a minute for testing

Oracle® OpenWorld Presentations

Winning the War on Complexity: Building the Oracle Grid

Rich Clifton, VP and GM Enterprise Data Center and Applications, NetApp
Bill Weils, Sr. Director, Command Center and Oracle Global IT, Oracle

[Presentation1](#)
(pdf 4.88mb)

[Presentation2](#)
(pdf 952kb)

Migrating Your Business to the Oracle Grid

Eric Nielsen, Principal Architect, Thomson Legal and Regulatory

[Presentation](#)
(pdf 617kb)

Enabling the Oracle® Grid through the Storage Layer (with Speaker Notes)

Jeff Browning, Chief Grid Storage Architect, NetApp

[Presentation](#)
(pdf 4.92mb)

Improving Database and Application Performance and Management by Improving Storage Utilization

Radhakrishnan Manga and Blaine McFadden, Technical Marketing Engineers, NetApp

[Presentation](#)
(pdf 1.62mb)

Role of Storage System Architecture in Oracle Performance

Stephen Daniel, Director Database Platforms and Performance Technology, NetApp

[Presentation](#)
(pdf 1.34mb)

Best Practices: iSCSI, Oracle RAC, RH Linux® and NetApp Storage

Giovanni Brignolo, Director of Database Applications, Engineering, NetApp

Presentation

(pdf 1.15mb)

Better, Faster Oracle E-Business Suite Upgrades

Patel Tushar, Technical Engineer; Paul Turner, Director, Database Technical Solutions, NetApp

Presentation

(pdf 2.55mb)

For more information the show, check out Oracle OpenWorld 2005 Online.

© 2005 Network Appliance, Inc.

Oracle OpenWorld Keynote:

Simplifying Data Management: Doing More while Spending Less

Tom Mendoza, NetApp President

Excerpt: Applying the Cisco “Magic” to Storage

In the 1990s we took a step back and said, “What did Cisco do? Let’s do that to storage.” That was our business plan. In fact, I raised the money for our company about ten years ago with Mr. Don Valentine. Don is the most senior venture capitalist here in Silicon Valley. He founded Sequoia Capital. Don funded Cisco, he’s chairman of the board of Cisco, he helped fund Oracle, he funded Apple, he funded Yahoo!, and he currently owns 15% of Google.

So Don is a very knowledgeable guy. I had to sit with him and explain what I wanted to do with storage, and if I was wrong I probably would have found out pretty quickly. Don gave me \$7 million that day, which turned into \$250 million for Sequoia. We’re buddies now.

Let me tell you what I told him. There are three things that we believe were the magic of Cisco.

- Number one, no general-purpose operating system. UNIX® and the Microsoft® derivative in the early 1990s was eight million lines of code. Today it’s 40 million lines of code. If you put that outside your data center, you’re gonna have to manage it. So IOS was their OS, 100,000 lines of code, an embedded file system, or operating system, that you could not modify that was made for writes, made to move data in their case
- Second thing was off-the-shelf technology. They didn’t make any hardware at the beginning of Cisco. They made a conscious decision not to do ASIC. That was odd at the time. And so they rode the price curves down, and Motorola was their first processor. Just goes to show you that processor speed is not the issue, it’s IOPS. And Cisco moved on to other processors, but the fact that they could move and it didn’t affect their customers was one of the beauties. They could pass those savings on to their customers. The other thing was they could outsource their manufacturing. Like Dell, they had a very efficient manufacturing where they hold the inventory for a very short amount of time and can pass those savings on
- Lastly, Cisco was multiprotocol. If they would have stayed IP only you wouldn’t have cared about Cisco today. When they did SNA [IBM System Network Architecture protocol for mainframes], they became infrastructure

So NetApp said, “Let’s do those three things to storage.”

- **Excerpt: Oracle Contributions to Data ONTAP™ 7G**
- **Listen to the full keynote: Real (45 min) Flash (45 min)**
- **Oracle OpenWorld demo and presentations**



Watch Streaming Video or Download:

- [Full keynote - Real \(45 min.\)](#)
- [Full keynote - Flash \(45 min./103MB\)](#)
- [Highlights - Flash \(4 min./10MB\)](#)

Oracle OpenWorld Keynote:

Simplifying Data Management: Doing More while Spending Less

Tom Mendoza, NetApp President

Excerpt: How Oracle IT Helped Define Data ONTAP™ 7G

[Data ONTAP] 7G is a huge hit at Oracle; they actually helped us with some concepts on it. To keep it simple, here's two concepts I'll bring to you today.

First, storage utilization is too low. Gartner Group says the average company's storage utilization is 30% to 50%. Most people think they're 50%, but they're really 30%. What's the problem? A user asks for storage; you either allocate what they ask for or a little more because a fixed partition is too hard to change. You come back later and find they are at 30%. Are you going to go back and move that storage, which takes hours to do, so that person can shrink their volumes? Probably not, and therefore, you have an inefficient architecture.

With [Data ONTAP] 7G you allocate storage with one command, and you shrink or expand it with another command. So you don't have to worry about it. If it's too little you just take it down. Oracle themselves, down at Austin, Texas, has 3.2PB of NetApp storage and 14,000 Dell Linux® machines hitting it. They told us that they have about 30% better utilization. They said to us, you know, we're gonna buy less with this new technology, and I said, I know, but more people will buy more.

The second thing actually came directly from Oracle in Austin, and I'm sure this is true of your organization too. A lot of storage is used for tests. A lot of software vendors would say that up to 70% of the licenses they sell are in tests. If you're going to roll out enterprise apps, you're going to test them.

Now how does it normally work? Let's say I have a 100GB database, and I change 10%. I have to buy another 100GB, and then you have to keep on doing that to test those changes. So the underlying infrastructure is extraordinarily expensive, and you have to make a decision. Am I going to afford all of that, or am I going to test less?

What Oracle said to us was, NetApp, you have these very, very efficient Snapshot™ copies, which only take the deltas or the changed data. Why don't you let us write to them? When they first asked us that, we were like, "Write to them?" Well, that is a very smart idea, because now we allow them to clone their apps, and all you're doing is sending off the changes: you test your app, and you blow it away. You will find you can do way, way more testing, you can set up in minutes instead of hours, and you don't need the underlying infrastructure. This is huge for enterprise development. Ask the Oracle Austin people.

- **Excerpt: Applying the Cisco "Magic" to Storage**
- **Listen to the full keynote: Streaming MP3 (45 min./11MB)**
- **Oracle OpenWorld demo and presentations**



Watch Streaming Video or Download:

- [Full keynote - Real \(45 min.\)](#)
- [Full keynote - Flash \(45 min./103MB\)](#)
- [Highlights - Flash \(4 min./10MB\)](#)



DEVINDER SINGH

Storage Architect and Alpha Geek, IT Department, NetApp

Devinder Singh is the hottest ticket when customers visit the NetApp Executive Briefing Center. That's because he knows — and shares — the secrets of a bold plan to keep business operations up and running in the event of a complete site disaster.

Before NetApp, Devinder was senior director of Technology and Global Operations for Appshop, the largest independent Oracle® outsourcing company. While at Appshop, Devinder led the design of a high-availability architecture that supports over 450 separate Oracle Applications environments, including crown-jewel Oracle11i™ ERP apps.

Disaster Recovery, NetApp Style Five-Minute RPO, 15-Minute RTO

When Storage Networking Industry Association execs first heard the five-and-15-minute results of the NetApp IT business continuance implementation, they knew attendees of Storage Networking World (SNW) in Florida had to hear it too.

A Tech OnTap member perk: front-row seat access to a detailed implementation overview.

Business continuity SLA:

- Provide continuous application availability to *all* global locations
- No more than five minutes of data loss (five-minute recovery point objective)
- No more than 15 minutes of downtime (15-minute recovery time objective)
- Control and execute failover from any location

Key challenges:

- Deliver consistent SLAs for FC SAN, iSCSI, CIFS, and NFS environments
- Have minimal performance impact on applications

What made it work:

- Very smart people (and quite a bit of NetApp storage technology)
- VERITAS® Global Cluster software



See How NetApp Achieved
Five-Minute RPO and 15-Minute RTO

RELATED INFORMATION

- [Ten Reasons IT Projects Fail \(and Ten Ways to Avoid Them\)](#)
- [NetApp Data Protection Solutions](#)
- [Storage Networking World Fall 2005](#)





NICK TRIANTOS

Global SAN Systems Engineer, NetApp

A member of the elite Global Systems Engineering Group, Nick helps top NetApp global enterprise customers solve their toughest technical challenges. Nick has been in systems or support engineering roles for nearly 15 years, including positions at HP as an account support engineer (Server Group) and presales technical consultant (Storage Group).

Mapping Out a Strategy to Minimize DAS

Part I: Understand Your Options

by Nick Triantos

This is the first article in a two-part series to help you minimize reliance on local storage. Here's how to evaluate your current environment and explore which technologies might (or might not!) provide the best solution. Look for "Part II: A Roadmap for Networking Storage" in December.

If your company is like most organizations, some unlucky IT staffer (probably you!) spends many painful hours managing and backing up servers with direct-attached or internal storage.

Every day, I work with our top enterprise customers to help them build consolidated networked storage environments. As soon as we start talking about storage, I find most people already have an opinion about which protocols are the "right" ones.

However, "one-size-fits-all" doesn't work for T-shirts, and it certainly doesn't fit your complicated IT environment. Protocol decisions should be based on data criticality, availability, performance, existing infrastructure, IT expertise, and budget. Frankly, one's philosophical bent also plays an important role in the decision-making process.

Here are key considerations to keep in mind as you evaluate your requirements and consider different storage protocol options.

Do your applications require block or file I/O?

Application requirements should always be the first thing to dictate your storage decisions. Many of the enterprises I work with are focused on consolidating direct-attached storage for applications that require block-based I/O, so that's what I'll focus on here.

What levels of redundancy do your applications require?

Mission-critical environments demand full redundancy with no single point of failure. Multipathing capabilities and clustered host configurations are common.

RELATED INFORMATION

- [Microsoft Storage Devices Blueprint: Enterprise Design for SAN](#)
- [NetApp SAN and IP SAN Solutions \(Flash\)](#)
- [NetApp Fibre Channel SAN Solutions](#)
- [NetApp IP SAN \(iSCSI\) Solutions](#)
- ["Changing the Economics of Storage" Webcast Series](#)
- [NetApp iSCSI Storage Center](#)

Know This

1. Does your application require block or file I/O?
2. What levels of redundancy do your applications require?
3. Do your applications use small or large block I/O, how many IOPS are being pushed, and how much bandwidth do you need?
4. Is your current infrastructure pure Ethernet, or do you have Fibre Channel fabric?
5. How do requirements differ between your main data center, regional data centers, DR facilities, and

I used to automatically recommend Fibre Channel for full redundancy environments because all of the original iSCSI solutions offered NIC teaming rather than a true multipathing capability. Although NIC teaming is a proven and mature technique, it does not provide a fully redundant configuration because the aggregated interfaces must be connected to the same network, often to the same switch.

However, the technology has evolved.

Today you can achieve the same levels of redundancy for Windows® environments using the Microsoft® iSCSI software initiator 2.0, which provides a proven and stable approach to multipathing: MPIO (multipathed I/O). In fact, MPIO is not the only approach to achieving path redundancy in a Windows iSCSI environment. The iSCSI specification provides its own features for high availability and bandwidth aggregation, including multiple connections per session (MCS).

Both MCS and MPIO allow failover among multiple NICs, and both can support I/O load balancing as well. MCS offers multiple TCP/IP connections between the server and the storage device during the same iSCSI session. These can be on different physical links as well. MPIO lets the server log multiple sessions to the storage array.

Net:Net — Fibre Channel is currently the only option enabling full redundancy for block-based applications in UNIX® and Linux® environments but is no longer the default choice for Windows servers.

Do your applications use small or large block I/Os, how many I/Os per second (IOPS) are being pushed, and how much bandwidth do you need?

In my experience, more than half of our customers do not have a firm grasp of their performance requirements. For small block I/O applications, for example, the critical factor is IOPS-not bandwidth.

Depending on the application block size, the same number of IOPS may have significantly different bandwidth requirements. For example, an application requiring 5000 IOPS with a 4KB block size would result in a bandwidth requirement of about 20MB/sec. In contrast, an application that uses 16KB blocks with the same number of IOPS needs significantly higher bandwidth.

- 5000 IOPS x 4KB blocks = 20MB/sec
- 5000 IOPS x 16KB blocks = 80MB/sec

Fibre Channel provides much higher performance than iSCSI in terms of bandwidth, so it is the clear choice for streaming audio and video applications or any applications with sustained high bandwidth requirements.

The majority of Intel® servers and applications, however, cannot sustain 1GB/sec performance, so the bandwidth gains of Fibre Channel are usually irrelevant. IOPS capabilities are quite similar between the two protocols, for 70% to 80% of Windows and Linux servers, iSCSI performance is more than adequate.

The UNIX space is harder to generalize. For block-based UNIX applications with very high bandwidth requirements, Fibre Channel is your only option. Older midrange UNIX servers and low-end UNIX servers typically cannot sustain more than 1GB/sec performance, so bandwidth does not need to be

- remote offices?
6. Is getting the “lowest cost” solution a top priority?



Tech OnTap membership has its privileges. Today, we're offering members a FREE collector's edition NetApp Tech OnTap t-shirt.

Act fast; quantities are extremely limited. T-shirts are only available while supplies last.

➤ **Claim your shirt now!**



a key consideration.

Don't forget to take seasonality into account. Retailer billing applications, for example, may normally do 20MB/sec but see dramatic spikes during the holiday shopping season.

Is your current infrastructure pure Ethernet, or do you have Fibre Channel fabric?

Companies relying exclusively on TCP/IP often lean toward Ethernet-based iSCSI. Fibre Channel is certainly an option, but if you decide to go that route you'll need to invest in additional hardware and budget for training or more staff with FC expertise.

In companies with established FC fabrics, storage and system administrators often lean toward Fibre Channel as a known and trusted solution. However, depending on your infrastructure and application requirements, iSCSI might be a good option if you already have a Gigabit Ethernet infrastructure in place.

Every enterprise is different. One of the top U.S. telecom carriers recently approached us about iSCSI but ultimately decided to stick with their existing Fibre Channel network. Why? They were using 100 megabit Ethernet and didn't want to invest in gigabit switches. Instead, they bought a couple of host bus adapters (HBAs), implemented a NetApp Fibre Channel SAN, and used available ports on their FC switch.

How do requirements differ among your main data center, regional data centers, disaster recovery (DR) facilities, and remote offices?

Doing a site-by-site review of application requirements as well as the existing network infrastructure is a key element in planning your consolidation strategy.

Application criticality, performance requirements, infrastructure, and on-site staff resources can vary significantly among corporate data centers, regional data centers, remote offices, and DR sites. Your decision at one site shouldn't lock you into using the same technologies at all locations. It helps, of course, to work with a storage vendor that supports both protocols.

For example, one major enterprise relied almost exclusively on Fibre Channel at its main data center, but its 15 remote data centers used only direct-attached storage (DAS). In addition to being painful to manage, there were no centralized backups to protect data in the event of a site failure because the cost of buying host-based software to enable centralized replication and backup was too steep. After determining that expanding its Fibre Channel infrastructure would be too expensive, the customer deployed iSCSI to locally consolidate storage and replicate data back to the main data center.

It's a mistake to overlook smaller data centers because they house less critical data. In my experience, about 75% of an enterprise's data resides in regional and small remote data centers. In many cases, Fibre Channel has not penetrated these areas simply because of cost. That means the majority of servers are connected to DAS or operate with internal storage. If you are developing a plan to consolidate local storage, consider a technology like iSCSI.

Is getting the "lowest cost" solution a top priority?

Cost should not be your primary concern, but is an area where iSCSI has a

significant advantage. Even companies with existing Fibre Channel infrastructures may find it cost-effective to connect at least some servers with iSCSI.

	Fibre Channel (Low End)	iSCSI* (Low End)
HBA (each)	\$800	Free Microsoft initiator available for download
Cables (each)	\$150	~\$15
Switch (cost per port)	\$500* *	\$80
One extra GigE server card for full redundancy (iSCSI only)		\$100 (GigE NIC)
Cost to connect each server with full redundancy	\$2900	<\$300

* Assumes Windows environment with an existing Gigabit Ethernet infrastructure

* * Cost per port on Fibre Channel Director class switches is typically higher and in some cases double the cost stated

In a Nutshell

There is no one-size-fits-all solution that works for everyone. But answering the questions in this article will help you understand the necessary trade-offs and choose the best technology for your environment.

Next Month: A Roadmap for Networking Storage

© 2005 Network Appliance, Inc.