

# Nanoseconds, Stubborn SAS, and Other Takeaways from the Flash Memory Summit 2019

Every year at the Flash Memory Summit held in Santa Clara, CA, attendees get a firsthand look at the technologies that will impact the next generation of storage. This year many of the innovations centered on forthcoming interconnects that will better deliver on the performance that flash offers today. Here are DCIG's main takeaways from this year's event.

## ***Takeaway #1 – Nanosecond Response Times Demonstrated***



PCI Express (PCIe) fabrics can deliver nanosecond response times using resources (CPU, memory, storage) situated on different physical enclosures. In meeting with PCIe provider, [Dolphin Interconnect Solutions](#), it demonstrated how an application could access resources (CPU, flash storage & memory) on different devices across a PCIe fabric in nanoseconds. Separately, [GigaIO announced](#) 500 nanosecond end-to-end latency using its [PCIe FabreX](#) switches. While everyone else at the show was boasting about microsecond response times, Dolphin and GigaIO introduced nanoseconds into the conversation. Both these companies ship their solutions now.

## ***Takeaway #2 – Impact of NVMe/TCP Standard Confirmed***

Ever since we heard the industry planned to port NVMe-oF to TCP, DCIG [thought](#) this would accelerate the overall adoption of NVMe-oF. Toshiba confirmed our suspicions. In discussing

its [Kumoscale](#) product with DCIG, it shared that it has seen a 10x jump in sales since the industry ratified the NVMe/TCP standard. This stems from all the reasons DCIG stated in a previous [blog entry](#) such as TCP being well understood, Ethernet being widely deployed, its low cost, and its use of existing infrastructure in organizations.

### ***Takeaway #3 – Fibre Channel Market Healthy, Driven by Enterprise All-flash Array***

According to FCIA leaders, the Fibre Channel (FC) market is healthy. FC vendors are selling 8 million ports per year. The enterprise all-flash array market is driving FC infrastructure sales, and 32 Gb FC is shipping in volume. Indeed, DCIG's research revealed 37 all-flash arrays that support 32 Gb FC connectivity.

Front-end connectivity is often the bottleneck in all-flash array performance, so doubling the speed of those connections can double the performance of the array. Beyond 32 Gb FC, the FCIA has already ratified the 64 Gb standard and is working on the 128 Gb FC. Consequently, FC has a long future in enterprise data centers.

FC-NVMe brings the benefits of NVMe-oF to Fibre Channel networks. FC-NVMe reduces protocol overhead, enabling GEN 5 (16 Gb FC) infrastructure to accomplish the same amount of work while consuming about half the CPU of standard FC.

### ***Takeaway #4 – PCIe Will Not be Denied***

All resources (CPU, memory and flash storage) can connect with one another and communicate over PCIe. Further, using PCIe eliminates the need for introducing the overhead associated with storage protocols (FC, InfiniBand, iSCSI, SCSI). All these resources talk the PCIe protocol. With the PCIe 5.0 standard formally ratified in May 2019 and discussions about PCIe 6.0 occurring, the future seems bright for the growing

adoption of this protocol. Further, AMD and Intel having both thrown their support behind it.

## ***Takeaway #5 – SAS Will Stubbornly Hang On***

DCIG's research finds that over 75% of AFAs support 12Gb/second SAS now. This predominance makes the introduction of 24G a logical next step for these arrays. A proven, mature, and economical interconnect, few applications can yet drive the performance limits of 12Gb, much less the forthcoming 24G standard. Adding to the likelihood that 24G moves forward, the SCSI Trade Association (STA) reported that the recent 24G plug fest went well.

*Editor's Note: This blog entry was updated on August 9, 2019, to correct grammatical mistakes and add some links.*

---

# **DCIG 2014-15 Flash Memory Storage Array Buyer's Guide Now Available**



DCIG is pleased to announce the March 30 release of the **DCIG 2014-15 Flash Memory Storage Array Buyer's Guide** that weights,

scores and ranks more than 130 features of thirty-nine (39) different storage arrays from twenty (20) different storage providers.

### **Flash Memory Array Performance and Sales Rising Rapidly**

Flash Memory Storage Arrays promise to deliver the dramatic performance benefits of flash memory including hundreds of thousands to millions of IOPS with sub-millisecond latencies while using as little as 1/10<sup>th</sup> the rack space, power and cooling of traditional enterprise storage arrays. The most recent generation of flash memory storage arrays generally deliver twice the IOPS of the prior generation and deliver a more complete set of features that enables them to address a broader set of use cases. Although enterprise storage professionals are traditionally cautious about adopting new technologies, many all-flash array vendors report that sales are growing in excess of 100% per year.

### **Flash Memory Storage Arrays Now Replacing Traditional Enterprise Arrays**

Multiple vendors we spoke with indicated that prospective customers are now looking to do a complete tier 1 storage refresh; transitioning to an all-flash environment for their critical business applications. Reflecting this trend, International Data Corporation (IDC) [forecasts](#) “*capacity shipped in 2016 will increase to 611PB [petabyte] with a 2012–2016 CAGR [compound annual growth rate] of 110.8%*”.

### **Driver #1: IT Budget Savings**

For those with a responsibility for the technology budget, a *flash-storage-enabled rethinking of the data center can generally achieve hard cost savings of over 30% in data center hardware and software, and realize an ROI of less than 11 months.*<sup>2</sup> In some cases, the Flash Memory Storage Array may prove less expensive than just the maintenance cost of the

former SAN.

## **Driver #2: Accelerating the Enterprise**

Flash-based storage systems typically create a seven-fold improvement in application performance. This accelerated performance is enabling progress on initiatives that were hampered by storage that could not keep up with business requirements. Savvy business people are finding many ways to generate business returns that make the flash storage investment easy to justify. In one [case study](#) *the installation of a flash memory storage system was directly attributed with avoiding the need to hire between 10 and 40 additional employees. Flash storage enabled the company to grow their business without growing their head count.*

## **About the DCIG 2014-15 Flash Memory Storage Array Buyer's Guide**

The plethora of vendors and products in the all-flash array marketplace—combined with a lack of readily available comparative data—can make product research and selection a daunting task. DCIG creates Buyer's Guides in order to help end users accelerate the product research and selection process—driving cost out of the research process while simultaneously increasing confidence in the results.

The DCIG 2014-15 Flash Memory Storage Array Buyer's Guide achieves the following objectives:

- Provides an objective, third party evaluation of hybrid storage arrays that evaluates and scores their features from an end user's perspective
- Ranks each array in each scoring category and then presenting these results in an easy to understand table
- Provides a standardized data sheet for each of the arrays so users may do quick side-by-side comparisons of products
- Provides insights into what features the arrays offer to

- optimize integration into VMware environments, as well as support for other hypervisors and operating systems
- Provides insight into which features will result in improved performance
- Gives any organization a solid foundation for getting competitive bids from different providers that are based on “apples-to-apples” comparisons

The DCIG 2014-15 Flash Memory Storage Array Buyer’s Guide Top 10 solutions include (in alphabetical order):

Hitachi Data Systems HUS VM  
HP 3PAR StoreServ 7450 Storage  
NetApp FAS3250 Series AFA  
NetApp FAS6290 Series AFA  
Nimbus Data Gemini F400  
Nimbus Data Gemini F600  
PureStorage FA-400 Series Controller  
SolidFire SF6010  
SolidFire SF9010  
Tegile Zebi HA2800

The **Nimbus Data Gemini F600** earned the *Best-in-Class* ranking among all Flash Memory Storage Arrays in this buyer’s guide. The [Gemini F600](#) multiprotocol unified all-flash storage system stood out in the following ways:

- Captured the highest score in the Management & Software as well as Hardware categories, and scored near the top in VMware Integration. In addition to excellent VMware integration, the Gemini F600 is one of just a handful of arrays supporting Microsoft System Center (SCVMM SMAPI), Microsoft Offloaded Data Transfer (ODX), and SMB 3.0.
- Raw flash storage density of 24 TB per rack unit, scalable to 1PB in a 42U cluster. This raw flash density was exceeded by only two other arrays in this guide.
- The already high flash storage density is further enhanced through comprehensive flash optimization

support, including lossless in-line deduplication and compression.

- One of the few arrays supporting 56 Gb FDR InfiniBand and 40 Gigabit Ethernet , the fastest and lowest latency interfaces for host connectivity.

[Nimbus Data](#) claims that the Gemini F600 can deliver up to 1 million 4K write IOPS and 2 million 4K read IOPS with latencies of approximately 50 microseconds; at a 35% lower cost compared to the prior generation. The Gemini F600 starts under US\$80,000, and typically sells for between \$150,000 and \$250,000 as configured by customers.

## **Flash Memory Storage Arrays Are a Systemic Business Opportunity**

The purchase of a Flash Memory Storage Array will be most easily justified and have the greatest benefit in businesses that think this through as a systemic opportunity. Many who do so will discover that “Flash is free.” That is, the return on investment within the IT budget will be rapid, and the business benefits of accelerating all enterprise applications could truly present an opportunity to accelerate the enterprise.

The DCIG 2014-15 Flash Memory Storage Array Buyer’s Guide is immediately available through the DCIG analyst portal for subscribing users by following this [link](#).

---

# **Gemini X-series Delivers**

# Painless Scale-out to Enterprises and Cloud Service Providers; Interview with Thomas Isakovich, Nimbus Data Systems, Inc. Chief Executive Officer and Founder, Part 3

In this final blog entry from our interview with [Nimbus Data](#) CEO and Founder Thomas Isakovich, we discuss his company's latest product, the Gemini X-series. We explore the role of the Flash Director and how it Gemini X-series appeals to enterprises as well as cloud service providers.

**DCIG:** *Tell us more about the role of the Flash Director. Does it have a role in how deduplication functions?*

**Thomas:** The Flash Director does not have a role in deduplication. The IOs come into the switchboard on the Flash Director, and it's just giving them out based on how many nodes are active behind it—and it's doing so in a balanced kind of striping manner. But the Flash Director also has another cool feature, which is the ability to stripe with parity, in either single or dual parity. This is for customers who want to be able to survive a complete node failure in the cluster.

Our approach to this is very cool because you don't have to mirror. You can do it with a RAID5 algorithm, as opposed to having it mirror the data. Again, I'm not sure how many customers are going to want to do that because every node is inherently HA; but if they want to, that ability is there. The Flash Director would handle that.

***DCIG: That's a RAID-style technology or erasure codes***

**Thomas:** It's RAID. If we went with erasure codes, we'd have to potentially have more capacity, or you'd have more overhead.

***DCIG: So deduplication is offloaded to the controllers on the Flash Nodes?***

**Thomas:** Yes. Behind this RDMA port is a CPU and everything that's dedicated to doing those hardware offloading functions.

***DCIG: The reason for choosing RDMA was low latency and high data rates?***

**Thomas:** Exactly that, yes. The latency is key. RDMA allows you to keep the CPU out of the data transfer process. And that keeps latency down.

***DCIG: Are you looking at FICON at all?***

**Thomas:** The direction that we're going in is more the RDMA protocols. We currently have SRP, but we're also going to be introducing iSER which is the iSCSI RDMA transport. The other big thing we're working on is SMB 3. You'll see that in the HALO 2014 announcement. And that allows us to kind of target the Hyper-V space where Microsoft wants to realign everyone away from iSCSI back to SMB. There are very few native SMB 3 implementations out there.

***DCIG: Could you talk about what developments were tailored to the enterprise versus the cloud service provider?***

**Thomas:** I think the key was to provide the multiprotocol services, because an enterprise customer is going to have a diversity of storage requirements and we didn't want to diminish that. If you try to name the multiprotocol scale-out storage systems in the market, there really aren't any.

So, I think the fact that we can deliver very painless scale-out with full multiprotocol services and data services, so it

plugs right in and it looks and smells just like the storage you have now from a connectivity and functionality perspective, but of course is way faster and way more efficient – I think that's what's going to be appealing to the enterprise customer.

Cloud customers may look at this and talk about their desire to use NFS. Because there are a lot of cloud providers that really like NFS. And NFS is high enough performance, especially if you combine it with flash and you're using it in kind of a cloud context.

And so here by providing a single name space, near-petabyte flash repository that is VMware certified and fully offloaded, with all the VAAI integration and so on, I think that's going to be another real big win. There isn't anyone else that's got that—at least that I'm aware of—because all the other flash arrays are pretty much block. There are maybe a couple here or there that are file, but certainly not scale-out file, so that's another big advantage of the Nimbus Data [Gemini X-series](#).

*In the [Part 1](#) of this interview series, Thomas Isakovich guided us through the development of the Nimbus Data Gemini X-series, and where he sees it fitting into the current market.*

*In the [Part 2](#) of this interview series, Thomas Isakovich described how the Gemini X architecture delivers highly available inline deduplication and consistent low latencies at scale.*

---

# Gemini X All Flash Scale-Out Storage Ready to Replace HDD as Enterprise Tier One; Interview with Thomas Isakovich, Nimbus Data Systems, Inc. Chief Executive Officer and Founder, Part 1

Recognized as an innovator in storage system technology, Thomas Isakovich sat down with [DCIG](#) to discuss the development, capabilities, and innovation in [Nimbus Data's](#) latest release: the Gemini X. In this first blog entry, he guides us through the development of the X-series, and where he sees it fitting into the current market.

**DCIG:** *Can you tell us what is so different about the X-series?*

**Thomas:** In terms of availability, this is probably the most advanced product – well, it *is* the most advanced product we've ever made, because it builds on everything that we've been improving. It takes that [Gemini](#) technology and then amps it up with true scale-out capability that is managed by our all-new Flash Director device. We've been working on it for the past two years. It's been a real challenge and also a pleasure developing it.

And, really, for us it completes the story. We believe we have the most competitive all-flash system currently in the market with the Gemini F. The only caveat being how do we scale to huge, huge sets of capacity? We now provide that with the Gemini X, and we've done it in a way that keeps the software and the hardware building blocks about 90 percent shared

between the two platforms.

Customers can start with the F series and go to the X series later. There's a lot of commonality between the two. From a manageability perspective, that familiarity will be a big plus. I think from an all-flash array portfolio perspective, we've got customers covered from three terabytes to a petabyte now – from a \$50,000 entry point to multimillion dollar solutions – all on the same Nimbus Data technology.

The timing of this product from our perspective is pretty perfect because our sales force is increasingly encountering customers that want to do wholesale refreshes of their entire tier-1 infrastructure. Not just flash for individual applications like databases and VDI, but really observing all-flash as a potential contender for the entirety of the tier-1 infrastructure. So having the ability to scale is well-timed and we're excited to be putting it out there now.

***DCIG: Can you talk more about the deduplication and compression of Gemini X?***

***Thomas:*** The deduplication and compression is really a sneak preview to an important feature of our forthcoming HALO 2014 operating system that we'll announce later this spring.

One of the challenges in scaling an array that uses inline deduplication is managing the vast metadata hash table that is the result of that, and keeping it in a manner where it's very rapidly accessible. And, as you know, a lot of solutions consume inordinate amounts of RAM to hold all this. But it's actually the RAM constraints of the controllers that may be limiting the ability for these inline deduplicating storage arrays to scale. So, many of those guys have been resorting to scale-out because, really, who's going to build an Intel server that can hold 20 terabytes of RAM? And even if it could, how do you protect it?

So we've come up with an algorithm here that effectively uses

about 1/50th the RAM and can deliver the same 4K block size in-line deduplication. This is one of the reasons we can build such high scale systems in such a small size. The Gemini X takes advantage of that technology, and so will the Gemini F, as part of running the HALO 2014 OS.

**DCIG: What environments are you seeing that are pushing high IOPS?**

**Thomas:** It's definitely geared more toward folks that just need huge amounts of capacity in a single domain. A lot of our customers, like our biggest one, which has 100 Gemini systems – they have no interest in actually presenting that as a single logical name space. They really do want 100 different name spaces, because of the way they're doing their scale-out. But they're a very sophisticated cloud provider, they can do very specific fancy things. For general purpose enterprise that doesn't have that level of sophistication – they're used to having their 500 terabyte hard drives or whatever – they need something that can present as one big box, and that's where this guy plays.

For example, a Fibre Channel port on a good day can do 100,000 IOPS. You're not going to get a million IOPS into a single server unless you're prepared to stick ten Fibre Channel cards in that server – which is going to be a challenge – and then run everything perfectly parallelized and all this other stuff. So our thought process in supporting four million IOPS is that we're going to need to support dozens or maybe hundreds of physical machines. And at 100,000 IOPS a port, that actually works out to about four million IOPS because you can have up to forty host ports on the Gemini X.

It's not so much that there's any one application that can come close to that, but you need to maintain a reasonable sort of IOPS-per-terabyte/IOPS-per-port kind of ratio. That's what our rationale is on achieving four million, because if you look at the Gemini F, by itself it's doing north of a million

in a read-write balance scenario. So on an IOPS-per-terabyte basis, the Gemini F is actually better, because when you do a cluster scale-out like this, you're going to have at least a little bit of latency from the cluster grid. We've kept latency very low because of the Flash Director design.

*In part two of this interview series, we discuss micro-latency and how the Gemini X performs against the competition.*

---

## **The Ways in Which Flash Memory Appliances will Impact How Organizations Buy and Manage Storage**

As DCIG makes its final preparations for the release of its inaugural Purpose-Built Flash Memory Appliance Buyer's Guide, we have had a number of conversations internally about what the criteria for product inclusion and exclusion in this Buyer's Guide will be. As we do so, our conversation almost always turns to ways in which these purpose-built flash memory appliances will impact organizations and their decision making and buying habits.

Probably the worst kept secret in the storage industry is that flash memory in whatever form it takes is almost always faster, quieter, consumes less power and produces less heat than comparable HDD-based storage systems. It is these benefits that every flash memory provider on the planet banks on to displace today's existing storage arrays.

Yet what is not so intuitive is how these purpose-build flash

memory appliances change the storage conversation. While they impact application performance, they change the dynamics of how organizations will buy, manage and even depreciate storage. Consider the following:

- **Flash as cache or storage.** One of the ways that traditional storage arrays compete against this emerging class of flash memory appliances is by deploying flash as a tier of cache or “Tier 0” storage inside their storage arrays. To get data on this flash memory tier they use an automated storage tiering feature that many include with their arrays. They argue (rightfully I might add) that 5% or less of the data stored on the storage array is active and will actually benefit from the performance that flash memory offers.

*The shortcoming of this argument is, “How does the storage array determine which of the data on the array is active?” Yes, they have their algorithms and formulas but do they work equally well for all applications they host. The answer to that is nebulous at best.*

*This highlights an important advantage that purpose-built flash memory appliances have over traditional storage arrays: they keep flash simple. In other words, if data is stored on a flash memory appliance, it is stored on flash. This eliminates the guesswork as to if an application will perform better and makes it easier to understand (albeit potentially more costly) the benefits of flash.*

- **Solid state disks (SSDs) or flash memory.** In flash’s initial iteration it was “easy” to deploy. All you had to do was replace a hard disk drive (HDD) with an SSD in a storage system and, voila, you had flash. However the costs, overhead and risk of this approach have surfaced. SSDs cost more than HDDs, organizations still have to get data on the SSDs once they are deployed and the

storage array has to be programmed to recover the data to another SSD should the primary SSD fail.

*This is not to say storage manufacturers are not taking these concerns about deploying SSDs in their storage arrays into account. But the decision to deploy SSDs in storage arrays is no longer as “easy” as it once was as the costs, risks and overhead of SSDs have become better known.*

*Purpose-built flash memory arrays have their own set of challenges but the primary advantage they again offer is their simplicity. Aside from the assurance that all data is stored on flash, purpose-built flash memory appliances are specifically designed to manage flash and account for its idiosyncrasies (garbage collection, wear leveling, etc.) whereas their traditional storage array counterparts are not.*

- ***The future roles of storage engineers and architects.*** Almost every enterprise organization as well as a storage reseller of any size relies upon storage engineers and architects to support and/or manage their storage infrastructures. But as purpose-built flash memory appliances are deployed, how does that change their role? It is unlikely it makes them obsolete but it should seriously reduce the amount of time they spend troubleshooting application performance issues.

*As such, what do they do now? Develop strategies that optimize where data is placed across different tiers of storage? Focus in on data protection and disaster recovery? The answer will clearly vary by organization but troubleshooting and managing application performance issues should fall off of or to the bottom of their list of items to manage.*

- ***The future of storage networks.*** FC and Ethernet storage

networks have become almost inextricably linked with enterprise storage deployments. But with the performance that purpose-built flash memory appliances offer, FC and Ethernet switches start to become a performance bottleneck. So what happens next? My guess is that this will give organizations impetus to deploy 10 Gb Ethernet or 16 Gb FC but who knows? 40 Gb Infiniband is already available and some early flash memory providers like [Nimbus Data Systems](#) already tell me that they ship systems equipped with Infiniband that is on par with their Ethernet and FC shipments.

- **Changing how one depreciates storage arrays.** One of the recurring themes I hear regarding why many are reluctant to deploy flash is its high cost. However what many fail to consider is how improvements in flash arrays extend their viability of these systems beyond the normal three (3) year storage system life span to five(5) or even ten (10) years with minimal or no decrease in performance during that period of time(we are talking microseconds in improvement in the years to come.)

*That being the case why do organizations still want to depreciate these systems over a three (3) year period? It is probably time for organizations to re-evaluate the depreciation period assigned to these arrays and move to at least a four (4) if not a five (5) year depreciation cycle.*

Flash memory does a lot more than improve application performance. It forces an almost top-to-bottom change in how organizations need to think about their storage infrastructures from how they should deploy them to how they should depreciate and manage them. However the real trick for each organization will be to arrive at a flash memory strategy that is right for them which may be the most difficult feat of them all.

---

# Why FCoE and iSCSI Trump Infiniband in Today's SSD Deployments; WhipTail CTO Interview Series Part V

In this final installment of our blog series on WhipTail Technologies, a Solid State Drive ([SSD](#)) array provider with some impressive features and capabilities, I am continuing my discussion with WhipTail Technologies Chief Technology Officer, James Candelaria. [Last time](#), we looked at how WhipTail implements software RAID on its devices. Today, we will be discussing the different transport protocols supported by the WhipTail array and why the FCoE and iSCSI protocols trump Infiniband in today's SSD deployments.

**Ben:** *We have talked a bit about the wide array of hardware you both use and support. Can you tell me a little bit about the storage protocols supported by WhipTail?*

**James:** [WhipTail](#) is fully multi-protocol. We support fibre channel ([FC](#)), [iSCSI](#), [CIFS](#) and [NFS](#), as well as the SCSI [RDMA](#) protocol. SSDs are all about latency. It is all about the microsecond response time. The minute you get onto a fiber channel or Ethernet fabric, you are talking about 120 microseconds. You are talking about doing a Remote Directory Memory Access ([RDMA](#)) over IP and you are talking about maybe three to ten micros – massive latency advantage there. So we are fully multi-protocol serviceable.

**Ben:** *And that is because of where you are in the stack naturally, right?*

**James:** Exactly.

**Ben:** I did a DCIG Buyer's Guide looking at [Enterprise Scale-out Storage](#). I was surprised how little Infiniband there is out there in the market.

**James:** So am I. It is a great transport. But you know what, I understand. **People just do not want another Layer 2 typology in their environment.** It costs money to maintain it. Even though the Infiniband switches and Host Channel Adapters today are dirt cheap, there is another cable I have to run. I have to find somebody who understands what a subnet manager is. It is not without its intangible costs.

**Ben:** I think that is why a lot of places are going with Fiber Channel Over Ethernet and iSCSI – just because they understand Ethernet.

**James:** Yes, Ethernet is a lowest common denominator in the data center. To be honest with you, when you can make iSCSI go end-to-end in the 120 – 150 microseconds realm, do you really need to go much faster? You have to ask yourself that question. I am a speed freak at heart. I race cars on the weekend, so I like to push the envelope.

But in the data center, it is not about that. It is about meeting your business objectives. I honestly think that Whiptail combined with a high speed Layer 2 fabric you already own, like 10 gigabit Ethernet, meets 90 percent of all business objectives. So I try to be less emphatic about my Layer 2 choices. We all have our preferences.

**Ben:** Do you see WhipTail as being part of the VDI revolution right now?

**James:** VDI's been a huge success for us because some of our corporate DNA. If you look at some of our senior executives, myself, our CMO Brian Feller, Andy Flesch, our VP in sales, all of us come from a virtualization background. So we get it.

And we are effectively able to leverage our knowledge and contacts in that space to solve a glaring hole in the VDI deployment model. Which was “we got all the brokers, we could fix all the display problems, but no one ever told us we need to account for storage performance!”

One day somebody woke up and said “we actually have to do this.” We were standing right there with the solution and it has been very successful for us. That is for a couple reasons. Obviously first – right place, right time. But the other guiding factor is that there is a new type of buyer that has been empowered to buy storage.

***The VDI buyer has zero allegiance to the incumbents.*** He has never bought storage before. That buyer has been told “you have to make VDI work, you have to get 5,000 seats up tomorrow.” He has never bought a [NetApp](#). He has never bought a [Symmetrix](#). He has never bought a Hitachi [USP V](#).

He is looking for something to solve his problem. He is not looking for a storage platform. So that has been one of the reasons we have been really successful in that market.

But as the product progresses and we mature as a company, other markets are obviously naturals for us to get in to – transaction processing, transaction. We have been very, very successful there, too.

We have a lot of customers that are Oracle shops that found us just too attractive – it was not really their first line of thinking, but somebody approached them and they said “*Hey, how about an extremely low latency piece of storage to speed up that report?*”

They gave it a shot and they took report run time down from 24 hours to one hour. All of a sudden their eyes lit up.

I had one customer in particular bought one for a specific use case, and then bought four more because he just decided

everything in my data center needs to be this fast -just started buying them like they were candy.

So we have been successful in marketing towards the application owners with problems – tactical solution right now for that report that does not run very well, that PDI performance crunch, that one online transaction database thing.

And just again, I cannot really talk about futures, but watch the space – we are going to go from a tactical company to a strategic company in the near future. So there is some pretty cool stuff coming.

**Ben:** *Another thing I find interesting about WhipTail is that, by extending the life of these flash RAMs, you are becoming a force multiplier.*

**James:** That is actually a phrase we have used around the office quite a bit is “force multiplier.” Alone, SSDs and NAND can only do so much. But when you couple them with an error mitigation strategy, they can become so much more to the data center, more than just a point solution that you think is going to be disposable at the end of the year because of wear, which will stop you from deploying them widely.

Once you get those constraints under control, you can start thinking about using them much more widely across every workload that you have that needs to go faster. Let’s face it, IT, the governing factor on business intelligence has always been performance.

**Ben:** *I think what is unique here is that you are catching something that is not just improving performance; it is drastically improving the lifespan of the product that you are using.*

**James:** Exactly. WhipTail is one of those rare things where you get additional performance and you actually get additional

endurance simultaneously, which is fairly rare. Intel is able to do it with processors and what not. But on everything else they will usually go together. It is like the old adage: &nbsp; p; “Good, fast, cheap – pick two.”

**Ben:** *Thank you, James. This has been a lot of good information.*

**James:** *Thanks, Ben. I appreciate your time.*

In [Part I](#) of this interview series, James explained the SSD garbage collection problem and how WhipTail handles it.

In [Part II](#) of the series, James discussed how WhipTail is optimizing SSD performance while minimizing the deficiencies of MLC flash.

In [Part III](#), Candelaria and I discussed how WhipTail deals with manufacturer variations in SSD drives.

In part [IV](#) in this series, Candelaria explains how and why WhipTail uses software RAID in its SSD appliance.

---

## **SSDs : Challenges, Opportunities and “Gotchas”**

As I mentioned in a number of the [blogs](#) that I posted last week while attending Storage Networking World ([SNW](#)), there was a heavy emphasis on solid state drives (SSDs) at the conference. However, during the many presentations that I attended and conversations that I had about this technology, SSD vendors revealed some key “gotchas” about SSDs. They also shared how SSDs stand to impact the hard disk drive (HDD) market as well as the market for memory as well. So here, in

no particular order, are some of the new challenges and opportunities that SSDs create as well as what to watch out for.

**No Management Software.** This was probably the most glaring deficiency in many of the new SSD products that I saw at SNW. Every new SSD provider was focused on the IOPS and throughput that its appliance could provide but there was little or no talk about what options users had to manage the data once they had stored on SSD. A couple of the SSD providers I spoke with are having discussions with storage virtualization providers such as [FalconStor Software](#) and [RELDATA](#) for this exact reason.

**Potential for data loss.** Store data to disk or tape and put it on the shelf or even just turn off the storage system on which it is stored and one can be reasonably sure that the data will be there when you turn it back on. SSD does not come with any such assurance.

One SSD provider told me that on a brand new SSD, one can be reasonably certain that data written to it will be there for 10 years if it is powered off. However an SSD that is 80% “worn” (i.e. – has had 80% of its allocation of writes consumed) may only be able to preserve the data for about a year.

This may be of particular concern in those situations when an SSD is physically placed in a server and the server is turned off for a period of time. However I also got the sense that this SSD feature is in a state of flux as SSD providers are working hard to overcome this deficiency and the potential for data loss may vary widely by SSD provider and what steps they have taken to correct it.

**Predictive failures.** This same “wearing” of SSDs that leads to data loss is not an entirely bad thing. Unlike traditional HDDs where you are essentially rolling the dice as to when the HDD will fail, this is less so the case with SSDs.

While SSDs still have the potential to fail at anytime, since they have no moving parts, the potential for unexpected failure is far less. Rather SSD providers are finding that by monitoring the number of writes to SSDs, they can predict when an SSD is coming to the end of its life and advise users to proactively replace drives.

***Saturation of existing network interconnects.*** One SSD provider, [WhipTail](#) Technologies, spoke about how its SSD appliance could easily saturate its two 4 Gb/Sec FC connections and even suspected the new 8 Gb/sec FC standard could be a bottleneck in performance intensive environment. (WhipTail was not yet using 8 Gb FC because it has not found the 8 Gb/FC drivers sufficiently mature.) This is leading it to look more seriously at introducing an Infiniband interface into its appliance.

***Be suspicious of Iometer results.*** Most SSD vendors had Iometer prominently displayed in their booths at SNW showing results of 100,000+ IOPS. Where users need to exercise caution is to not be unduly swayed by these results as Iometer can be configured a number of different ways. One can assume that to achieve the results displayed in these booths, Iometer was configured in the most optimal way possible, probably with 100% reads and zero writes.

***Replace memory on servers.*** Much attention has been given to SSDs replacing HDDs in the near future but it is also conceivable SSD could replace memory as well or at least reduce the amount of memory that systems need. While SSD is not as fast as DRAM, it may be a “good enough” replacement for memory on many application servers that organizations can justify making the switch and eliminating HDDs and memory altogether on them.

---

# Gartner Analyst Sees VMware Security as Immature; Wachovia Decouples Processing and I/O; Flashbacks to the Fall IBTA

As I write this blog entry, I am currently on a flight to New York City to attend the last day of the fall 2008 [Storage Decisions](#) conference. While I intend to post a blog entry about my experiences at SD this Friday, the flight is giving me some time to go back to last week and share some additional thoughts and insights I gained while attending the InfiniBand Trade Association ([IBTA](#)) Tech Forum in Las Vegas on Monday, Sept 15. While infiniband was obviously covered as part of this forum, it was done so in the larger context of what virtualizing the corporate infrastructure means and how that will contribute to how companies construct and manage their data centers in the future.

First, I want to highlight some tidbits in the keynote that Gartner's John Enck gave at the opening of the IBTA Tech Forum. Initially he challenged the assumption that [VMware](#) is a secure platform that is immune to attacks and viruses. He mentioned that security is still a point of immaturity in VMware and that VMware issued a [fix](#) for known vulnerabilities in its product just a week prior to VMworld. While I believe that VMware is more secure than Windows was in its early days, more vulnerabilities are likely to be discovered as it is more widely deployed.

Another interesting tidbit that Enck shared had to do with some of the new management functions that server virtualization makes possible. Enck specifically mentioned that because the images of each individual VMware virtual machine (VM) is stored in what is called a [VMDK](#) (Virtual Machine Disk Format) file, it opens the door for companies to assign metadata to these files including service oriented metadata. In so doing, companies can begin to manage and assign service priorities to VMDK files in the same way they manage other files. It is not yet clear how this will evolve over time but it makes sense that priority levels will become part of the VMDK file's metadata. In this way, VMs will be automatically assigned to the appropriate tier of server and storage hardware simply by looking at the setting at the metadata.

Personally I found Enck's presentation one of the best I have heard on the current challenges of server virtualization as well as its future promise. His comments were on-target based upon what I know while providing an honest and realistic assessment of the infiniband market as it relates to virtualization in general.

Another presentation that I also wanted to comment on came from Citrix's CTO, Simon Crosby, who is responsible for XenServer's oversight and direction. Those of you who think that VMware has already won the server virtualization battle may be less convinced after hearing Crosby talk. While I am not going to dispute that VMware has market momentum and is gaining corporate mindshare, Crosby said that [XenServer](#) has the world's largest high performance computing (HPC) cluster deployed at the Department of Defense (10,000 nodes deployed worldwide) and Google also has a very large Xen implementation going on.

One of the more interesting thoughts that Crosby shared during his presentation had to do with what occurs when the operating system becomes aware it is virtualized. Operating systems are

now oblivious to the fact that the relationship between it and the underlying hardware no longer exists as before. Using operating systems that are “virtualization aware” creates a paravirtualization environment which opens the door for a profound change in how companies manage their infrastructure. However Crosby did not get into the details as to what “profound” changes he saw occurring.

Another point that Crosby shared had to do with hypervisor performance and how the different virtualization hypervisors scale as the number of cores in the physical machine increase. In his opinion, performance on all hypervisors is about equal where there are only 2 cores in the physical machine; when 4 cores are present you start to see performance on one hypervisor start to fall and when 8 cores are present, he said that Xen starts to separate itself from the crowd. He attributes this to two things. First, Xen is fully 64-bit aligned and said that Xen helps to drive the 64-bit roadmap. Second, its core code base is under 100,000 lines which make for faster processing such that the performance overhead on the physical machine from supporting Xen hypervisor can be under 1%.

Crosby concluded his presentation with a in-depth discussion of how important I/O is going to become in the coming years. While I want to dedicate a separate blog entry to this after I have done some more research on this topic, he stressed that it is important that companies begin to put the right network interface cards (NICs) in their virtualized servers. There are now essentially 4 classifications or levels of NIC cards (0, 1, 2 & 3) and it is NIC cards that support levels 2 & 3 that companies should give priority. He said the level 2 is where the big jump was made as the hypervisor can now begin to directly access buffers in the NIC card and reserve specific buffers for specific VMs. In so doing, this can provide a tremendous boost in performance on specific VMs.

Some final thoughts I wanted to share come from Jacob Hall,

the Chief Architect for [Wachovia](#) Corporate and Investment Banking. First, he is creating an infrastructure where the I/O is decoupled from the chassis as he sees embedding I/O in the chassis as decreasing the life time of the chassis. By decoupling it, he can scale and upgrade processors and I/O independently of one another and take advantage of advances in each without impacting the other. Second, he thinks the entire industry needs to accelerate its adoption of advances in I/O technology. Advances in other areas of technology happen in over 1 – 4 years while advancements in I/O take 7-8 years to find their way into the datacenter. In his opinion, this is too long. By increasing the adoption rate of new I/O interconnects, it not only increases performance but reduces energy costs.

---

## **Cisco Bows Out of Infiniband; Appliances Join the Virtual World; Notes from Day 2 at VMworld**

Day 2 at [VMworld](#) has come and gone and probably my biggest regret was that I had to miss this morning's keynote by VMware's new CEO, Paul Maritz. In reading through some other blogs this evening about the event and assuming [Storagezilla](#) called it right, it was a doozy essentially declaring open war on other operating systems.

In any case, my day was focused on catching up with a number of vendors to get some of the latest behind the scenes scoop in the storage world. In fact, as one walks into the exhibitor hall in VMworld, it is hard not to mistake this conference for a storage conference. HUGE booths from EMC, NetApp, Dell, NEC, Symantec and Emulex greet attendees as they enter the exhibitor hall. While I never did a formal count of how many storage vendors were exhibiting, it had to be fully half of those exhibiting were in some way connected to storage. Even [TechTarget](#) had its own booth tucked away at VMworld which I found slightly amusing considering it does its

own set of conferences.

I started the day by meeting with Jon Toor from [Xsigo Systems](#) to catch up on what was going in its world. For those of you unfamiliar with Xsigo Systems, it offers an InfiniBand I/O Director that virtualizes server I/O. Its [VP780 x2 I/O Director](#) connects directly to servers via Infiniband but provides FC and Ethernet ports so users can keep their existing Ethernet and FC infrastructures in place.

Since Xsigo Systems had only released its first product in Q4 of 2007, I was curious to hear how far it had progressed not just in terms of market acceptance but if it was actually in production anywhere. Not only is it in production, an airline reservation system with stringent SLAs is already using it in its production environment. This is remarkable for a product that is less than a year out of GA to already have such a major win. While Toor could not disclose who it was, he did say the client migrated to this new InfiniBand infrastructure in well under six months.

This bodes extremely well for Xsigo Systems specifically and the introduction of InfiniBand into the data center in general because competing 10 Gb Ethernet is still a good year away (that is a best case scenario). Further, part of the reason Xsigo Systems won this account was because Cisco told the prospect that Cisco had halted its own internal project work on InfiniBand. This left Xsigo Systems as the only competitive solution that could provide an answer for this customer's problems.

Another interesting tidbit of news I picked up was from a noontime conversation I had with [Xiotech](#). Apparently they are winning some deals against [Dell EqualLogic PS Series](#) systems because of poor performance issues on these systems and replacing them with FC. Now I am not sure what the exact details are around these wins (old EqualLogic systems, customers mixing iSCSI traffic with normal network traffic, improperly configured systems, virtualizing multiple arrays) or their accuracy because everything I have heard about Dell EqualLogic in the past has been fairly positive. But needless to say this news surprised me as EqualLogic systems hold up pretty well under most normal conditions.

A final interesting note that I picked up from a presentation given by Jacob Hall from [Wachovia](#) on Monday at the [IBTA](#) Tech Forum and then again heard from a briefing I received from STORServer is what appears to be the start of trend of introducing server virtualization on appliances. In the former case, Wachovia wanted to virtualize the thousands of Wachovia desktops but could not justify virtualizing all of them and putting them on the data center floor.

So what he did instead was virtualize one in every four. Assuming I understood him correctly, he put one high powered PC in place of every four, connected monitors and keyboards to it and used VMware's [VDI](#) to present four virtual desktops to the end-users. He did this using an appliance from a Canadian company (whose name he did not disclose). In so doing, he kept the costs out of his data center, reduced the amount of hardware and power his company needed and still met his end-users expectations.

[STORServer](#) announced something similar this week in its Sept 15 [press release](#) regarding its V-Series Data Protection Appliance. Normally STORserver just provides on Data Protection Appliance a self-contained, turnkey

version of Tivoli Storage Manager (TSM) but if customers desired other STORserver products (archiving, CDP, etc), they would need to purchase other appliances with those features on it. But by now bundling VMware on its appliances, it can create different virtual machines on its appliance that hosts different applications.

While the possibilities of this new virtualization appliances are intriguing, companies need to keep in mind that this a new frontier for VMware and the licensing for these appliances has not yet caught up to the reality. STORserver's VP of Business Development, Bob Antoniazzi, says that VMware normally only sells its virtualization license to the end-user, not the provider. These virtualization appliances change the model and VMware has not yet come out with a means to address it. However Antoniazzi didn't seem to care. "We're offering these appliances with VMware starting this week so this should help to put some pressure on VMware to figure the licensing issue out," says Antoniazzi.

---

# Day 1 at VMworld Begins with a Look at the Future Pain of Server Virtualization: Attending the IBTA Tech Forum

It's day one at [VMworld](#) in Las Vegas and while the day for me began in Omaha NE at 4:30 am CST before landing in Las Vegas around 7:30 am PST, I did not join the throngs basking in the VMworld love fest. Instead I spent the day educating myself more about the topic of Infiniband by attending the InfiniBand Trade Association's ([IBTA](#)) annual tech forum that was held at [Harrahs](#) (Harrahs is adjacent to [The Venetian](#) where VMworld is being held). The reason that I elected to first attend the IBTA Tech Forum and not VMworld is simple. Everyone already knows that server virtualization is the BIG thing. What everyone doesn't know or understand is why Infiniband is making a case to become the next big thing in another form of virtualization: Virtualizing server I/O.

If one looks strictly at the numbers of people attending the IBTA Tech Forum (55 - 60) versus the throng of 6,000+ descending upon VMworld (my best guess - VMware would not disclose the actual number), you would think I had missed the boat. But judging by the number of grey hairs in the room (the number 55 - 60 seemed to coincide with the average age of those in attendance at the IBTA) versus the vast numbers attending VMworld that were in their 30's, I doubt it. Many in attendance at the IBTA were the same individuals virtualizing servers and infrastructure before it was in vogue to do so. Now these same individuals are pointing out the new

problems that server virtualization creates and building a case as to why InfiniBand used as a server I/O interconnect is the logical technology to solve them.

I choose the term “server I/O interconnect” carefully because InfiniBand advocates made a critical mistake in the past when they initially marketed Infiniband and they admit it, at least privately. The initial noise around Infiniband was that it would become the data center interconnect and eventually replace Ethernet in the data center. As anyone in the data center knows, that’s ludicrous. So not only did it show the Infiniband folks were drinking a little too much of their own Kool-aid, it hurt the credibility of InfiniBand in an industry that has a long memory and is slow to forgive.

This time around it does not want to get into an “InfiniBand versus Ethernet” debate. It will lose that debate every time and it knows it. By positioning InfiniBand as complementary to Ethernet and avoiding the “rip and replace Ethernet” concern, it better positions itself to make an entrance into data centers.

The good news for InfiniBand is that demand for high speed, high throughput, low latency interconnects is on the verge of exploding. The bad news is no one knows about it (other than the 60 people who showed up for the tech forum. The big challenges that InfiniBand faces in addressing this are:

- Positioning itself against the forthcoming Data Center Ethernet (10G) as complementary to it and not a replacement for it.
- Overcoming Cisco’s marketing of Ethernet at the executive level
- Providing mature Infiniband management tools
- Educating people on how InfiniBand differs from Ethernet
- Agreeing among themselves on a common message as to why it is so compelling and how companies should first look to deploy it

So how close is InfiniBand to having a realistic play in enterprise data centers? Best guesses put it anywhere from 1 - 4 years away. The reason for the delay is that the wide scale corporate adoption of server virtualization is just getting under way as only about 7% of corporations have adopted it according to Gartner’s VP John Enck. It is not until a large number of companies implement server virtualization and start to experience some of the pain associated with server virtualization creates that companies will begin to understand the performance issues associated with server virtualization. More specifically, I/O on virtualized servers starts to surface as a performance bottleneck that InfiniBand is better suited than Ethernet to solve.

While this is not the right time to get into all of the technical details of why that is the case, the overhead of managing I/O impacts virtual servers such that hardware utilization of CPU and memory can top out at around 40% using Ethernet due to the latency involved waiting for responses from network and storage resources. InfiniBand’s higher throughput (40Gb/sec) and lower latency eliminates this I/O bottleneck plus it becomes more cost effective and uses less power than Ethernet.

VMworld is undoubtedly where the buzz is. Goofy VMware signs are everywhere; EMC is giving T-shirts and \$100

gas cards away, and some vendor was giving away green stuffed animals that looked like Kermit the Frog (for what reason, I have no idea). Yet the real action today was behind the scenes as the Infiniband industry refines its message and shores up its product lines to get ready to go into the enterprise because everyone knows the next generation of server virtualization problems in the form of I/O congestion is coming. The big question is will it be Ethernet or InfiniBand to which companies turn to solve this problem? Right now the IBTA wishes they knew the answer to that question.

---

## Exploration into Infiniband

I believe a new way on thinking should be applied to the deployment on Infiniband technology in the storage landscape. Most of you probably think of Infiniband as predominately a backend transport for storage, and/or the interconnection mechanism for high compute clusters (HPC). Or, “Oh yeah, I heard something about that 5-6 years ago, isn’t that only used in super-computing or giant research labs?” Let’s first break down some of the basics of Infiniband:

- It is a channel-based, switch fabric, interconnect architecture for compute and I/O
- It is extremely high speed and low latency fabric (10Gb/s and 20Gb/s now with 40Gb/s coming)
- It leverages the ability to consolidate storage, clusters, networking and management over the same fabric

Some of the real grass roots benefits in the deployment of Infiniband are:

- Higher performance than any other I/O transport available today
- Massive consolidation of cabling infrastructures to your servers
  - As one Infiniband HCA (Host Channel Adapter) can service compute, clustering, storage, and networking
  - (Think about that for a second, especially when you consider the amount of cabling required for a VMware server this can reduce the amount of heat generated in the backend of a server-rack, as well as provide a well-groomed, clean and easy to manage cable plant.)

- Infiniband is inherently much more cost effective from a power consumption point of view as well.

There are a number of very attractive use cases for Infiniband which can apply to many aspects of your current storage infrastructure. The augmentation of your existing storage network with the use of an Infiniband gateway to provide a transport for I/O from your servers into the Fibre Channel and/or iSCSI fabrics. This allows you to maintain your protection in those FC and iSCSI investments. Based on the performance characteristics of Infiniband, you will not see the bottlenecks as have been seen with many other traditional gateway technologies (e.g. FC-to-SCSI Bridges).

When you sit down to start planning your next-gen storage landscape for your organization, I encourage you to do some research, talk to your manufacturers, and see how their roadmaps and product sets align with Infiniband. I think you will be presently surprised how many of them are already using it and plan to extend its reach.

I know many of you are probably saying to yourself, “Fibre Channel and Ethernet are way too embedded into my environment to even begin to think about adding an Infiniband layer”. Just weigh some of the basic facts I have mentioned in this blog like low latency, speed, traffic aggregation for FC and Ethernet, and cost, and keep your mind open.

There are a number of useful links that provide a great amount of detail around Infiniband:

<https://www.infinibandta.org/home>

<https://www.oreillynet.com/pub/a/network/2002/02/04/windows.html>

In future segments I will cover these use-cases for Infiniband in greater detail so you will have all the ammunition you need to move forward.

---

# Is FCoE a Diabolical Plot?; Musings on SNW's Day 2 FCoE Announcements

I initially intended to share in this blog posting what I learned from my briefings on Day 3 of Storage Networking World (SNW). However I've had some more time to digest the news surrounding the [Fibre Channel over Ethernet \(FCoE\)](#) announcements at SNW on Tuesday and the more I think about it, the more this whole FCoE strikes me as a huge setup to lock users into [Fibre Channel](#) (FC) that is being carefully orchestrated by the FC industry. Though this was hinted at about a year ago in an [article](#) that appeared on Computerworld's website, the roadmap and agenda of how vendors like [Brocade](#), [Emulex](#) and [Qlogic](#) and, to a lesser extent, [Cisco](#) and [Intel](#), intend to do so over the next 10 years is more clear.

My understanding is that 8 Gb/s FC represents the end of the upgrade cycle for the current generation of FC technology. Whether enterprises are running 1, 2, 4 or 8 Gb/s FC, the underlying optics are essentially the same allowing for interoperability between new HBAs and existing FC cables and directors. Most importantly, the FC infrastructure did not need to dramatically change from generation to generation as upgrades occurred or new products were released.

However those days are over. Again, as I understand it, the next FC upgrade cycle in data centers beyond 8 Gb/s, whether it is to 10 or 16 Gb/s FC, is going to require a rip-and-replace of the current data center FC infrastructure. With that looming, FC vendors knew they needed to cooperate and collaborate to keep FC viable regardless of which FC technology path users choose. Otherwise when users start to take a long, hard look at the pros and cons of FC versus [InfiniBand](#) during the next data center refresh cycle, 40 Gb/s

InfiniBand stands an above average chance of replacing FC.

So to avert this, my guess is that the FC vendors concocted a plan: Use FCoE to connect all enterprise servers, get a few analysts on board to endorse the idea and then convince end-users to take their eyes off the longer term ramifications of using FCoE. By getting enterprise users to bite on FCoE and spend the next few years connecting their remaining 85% of their servers to existing FC SANs, users are locked into FC for the next 10 years until the next disruptive technology comes along.

Now with the remaining 85% of the servers in the data center running FCoE, the most logical upgrade path for users for the original 15% of servers and storage is FC. Then regardless if the next FC upgrade is 10 Gb/s or 16 Gb/s FC, when the inevitable rip-and-replace comes in 2 – 4 years, FC lives on and InfiniBand remains a niche market.

Tuesday's announcement had less to do with what's best for the end users and everything to do with [preserving Brocade's, Emulex's and Qlogic's core FC business](#). To do so, they needed Intel and Cisco to come on board, support it and promote it. If this FCoE initiative fails and users actually start to compare the benefits of InfiniBand to FC and realize that they can get 10x the benefits at the same cost as FC. FC and InfiniBand could swap places. Then FC could become the new niche market and InfiniBand may begin to dominate in the data center.

Look for my notes and thoughts on my Day 3 SNW briefings and meetings on Monday.

---

# Cisco, Emulex, Qlogic, etc Announce 10 Gbps FCoE While Mellanox says “We raise 4x” and Announces 40 Gbps Infiniband; SNW Day 2 Recap

Day 2 of Storage Networking World at SNW is now in the books. Right now I’m frantically trying to keep up with the blogging updates between briefings, interviews and just generally shooting the breeze while roaming the hallways of the Rosen Creek Shingle hotel where the conference is being held. Here are the highlights from my rounds at day 2 at SNW.

[Xiotech](#) made the first “earthshaking” announcements of the day at 7:00 am which mostly had those I spoke to shaking their heads trying to figure out what the announcement meant. The announcement centered on their new patented [Intelligent Storage Element \(ISE\)](#) technology that they acquired from Seagate last November that will, according to Xiotech, “virtually eliminate the need for service, scale from one terabyte to one petabyte and dramatically boost performance”.

Exactly how ISEs do this on [Xiotech’s Emprise](#) storage systems is unclear to me at this point. Their press release is long on the business benefits but vague on how it works. An analyst report put out by [ESG labs](#) may shed some further light on how this technology works and what its value is to end users though I still need to read through it to determine to what extent they address it.

Next up was a briefing with [Data Direct Networks’](#) Josh Goldstein. The most significant point of this conversation was the uptick Data Direct Networks has seen in the last two

quarters in the number of its storage systems shipping with [Infiniband](#). While Josh could not provide me with the exact numbers, he said that he had not closely examined the numbers for a couple of quarters and when he looked at them a couple of weeks ago, the numbers surprised even him.

A briefing with [COPAN System's](#) CEO Mark Ward followed. It was good to hear COPAN Systems has finally gotten off of its "We have MAID" soap box and moved on to selling product, instead of telling everyone about the wonders of MAID. While it shared some new evidence that using MAID would extend the life of hard drives from 600,000 hours of MTBF (mean time between failure) to 4 million hours, it more importantly has increased its employee count from 60 to 180 and is growing sales correspondingly.

Early in the afternoon [Cisco](#), [Emulex](#), [Qlogic](#) and everyone else's brother was supposed to be the 2nd "earthshaking" announcement of the day. Together they announced a new category of [Converged Networks Adapter \(CNA\)](#) cards that enable both storage and network traffic to be combined into a single unified fabric. To do this, they are using the [Fibre Channel over Ethernet \(FCoE\)](#) protocol that converges Ethernet-based LAN and [Fibre Channel](#)-based SAN connectivity to a 10 Gb/s Enhanced Ethernet network.

Usually I am big fan of the "Ethernet for Everything" principal but not in this particular case. Despite their combined marketing muscle, their predominant presence in data centers and the ubiquitous nature of Ethernet, they are going to miss the mark short term.

I say this because [Mellanox Technologies](#) (who also announced a [FCoE CNA w/Host Offload on Monday](#)) had on April 1, 2008, announced a 40 Gb/s Infiniband card. Though this announcement definitely lacked the fanfare surrounding the FCoE announcement, 40 Gb/s beats 10 Gb/s according to my numbers. Equally importantly, there are emerging switch

technologies from companies like [Xsigo Systems](#) that capitalize on the virtualization I/O benefits that Infiniband provides, Infiniband's roadmap is much more robust than Ethernet and its price point is the same or less than these new FCoE. While Mellanox's VP of Product Marketing, Thad Omura, told me that FCoE really won't gain much traction until 2009, I see that as a best case scenario.

My final briefing was with [Bocada's Nancy Hurley](#) who provided me some insight into how

Bocada is adapting to the recent changes in market conditions for the [data protection and recovery management \(DPRM\)](#) market. Bocada Enterprise is finding itself engaged in new battles with [EMC](#) and [Symantec](#) at the enterprise level with backup reporting. Hurley says Bocada is often winning in head-to-head bake-offs but when pricing is introduced, EMC and Symantec are including or packaging their data protection management software for "free" as part of larger backup software or storage deals.

To address this, Bocada is starting to switch gears and position [Bocada Enterprise as a Unified Management Platform](#). It is capitalizing on the agentless architecture and analytics already found in Bocada Enterprise and adding new features to specifically manage [Microsoft Data Protection Manager \(DPM\)](#). To introduce this functionality, Bocada has signed a 6,000 seat agreement with Microsoft where it will test drive this new functionality.

Look for my day 3 SNW thought and comments by mid-day tomorrow.

Please email Jerome Wendt directly at [blog.master@dcig.com](mailto:blog.master@dcig.com) with comments and corrections.

---

# Fibre Channel over Ethernet (Part 2 of 3)

Last month I did some research and evaluation of Fibre Channel over Ethernet (FCoE). In my [Part 1 of 3](#) I shared some elements that can encourage the use of FCoE in the data center.

During my research I spent about an hour on the phone with Mike Krause, Fellow Engineer at HP. Mike and I talked about few things related to FCoE, shared storage and network fabrics. I asked Mike about creating a shared fabric using InfiniBand, because InfiniBand requires a single card type and InfiniBand switches that take frames to their respective networks. Mike countered saying that InfiniBand is a great option, but it introduces a new, third architecture to the data center. I realized immediately that it was orthogonal to existing data and storage networks. Mike further commented that he liked InfiniBand as an option, but the original intent was to replace PCI as the primary peripheral interconnect within servers etc. Thus, it made sense that InfiniBand wasn't creating a shared fabric, just moving the cards to another location. Mike finished our discussion about InfiniBand by saying "Ethernet is the logical consolidation of [the] fabric."

We started getting into the details of FCoE from Mike and HP's perspective. HP has a gold *mind* with Mike on board as a Fellow Engineer. Mike's salient points about shared fabric and blade systems with respect to cable consolidation was fascinating. He made it very clear that cable consolidation should not be a factor in choosing FCoE as the devices emerge through 2009. Cable consolidation is best left to [Blade Systems](#), a rack-n-stack consolidation effort.

Mike echo's industry thoughts that the world is Ethernet

already and everything is connected by Ethernet. Therefore, Ethernet is a good answer to shared fabric and FCoE is the right answer to transition Fibre Channel to the Ethernet shared fabric. Moreover, Mike made it clear that the data center is not run as a single sub-net, therefore it is not practical to assume entire data and storage administrative teams will consolidate in the next few years. In fact, Mike, along with [Claudio and Bill from Cisco](#), believe consolidating data and storage networks, as well as administrative functions, will be done incrementally. Servers would still have four cards, just four 10gig Ethernet cards. Two cards for storage network and two cards for data networks (need to have redundancy). Mike conveyed that using one card for both storage and data networks will happen by necessity, such as during a disaster recovery or other dire circumstances as a contingency plan!

It was really great talking to Mike and I would encourage anyone interested in learning more about FCoE, InfiniBand and HP try to talk with Mike at any of HP's events or [Storage Networking World](#). We'd be happy to pass on your comments as well!

---

## **Fibre Channel over Ethernet (Part 1 of 3)**

This week I've spent much of my time talking to Cisco, HP, EMC and [Fibre Channel Industry Association](#) Chairman Skip Jones. Monday afternoon I spent time talking to Claudio DeSanti, Technical Leader, Advanced Architecture and Research and Bill Lulofs, Product Manager, Data Center Business Unit, both at Cisco Systems.

When I received the assignment to review the FCoE specification and compare it to iFCP, FCIP and iSCSI (block protocols over data networks) I was thinking it might be boring, I was very wrong. After just a few short minutes with Claudio and Bill I knew I was talking to a pair of very intelligent and thoughtful business technologists. I had a good idea of what I wanted to talk about. I had already read these articles and all the comments:

In this blog, [The uncomfortable marriage of Fibre Channel and iSCSI](#), Julian Satran's comments were most helpful prior to my interviews with the aforementioned companies.

This Question and Answer article where [Brocade, EqualLogic tussle over Fibre Channel and iSCSI supremacy](#) helped me realize that FCoE offers two types of green. Doug's comments made it clear there is green to be saved by by maintaining FC assets, employees and software. John's comments made it clear there is green to be had with envy by the iSCSI proponents.

Bill and Claudio helped me understand there is room for FCoE in the data center, citing these items as the tops:

1. Data center focus for low-latency, low-cost per bit and no gateway required storage network
2. Reduction of power by consolidating I/O cards to 10gig Ethernet
3. Preservation of fibre channel physical assets, administrative skills and management software
4. Native Fibre Channel over Ethernet; nothing changes just the leveraged buying of I/O cards and switches
5. iSCSI requires overhead of transport and network layer, when not needed
6. Consolidation of I/O using evolutionary approach; building off of iSCSI

Be sure to read the upcoming Trends article in Storage Magazine's November edition where Jerome and I set the

baseline for FCoE. In the final installment of this blog series I'll talk more about the six items above, but I'm hoping to see this technology in action at the FCIA booth at [Storage Networking World](#).

In part 2 of 3 I'm writing about InfiniBand, more on Cisco and my interview with Mark Krause, Fellow Engineer at HP.