white paper

The Reality Gap

How SAN Hype Ignores Real Customer Needs

by
Stuart Berman
Chief Technology Officer

Vixel Corporation
January 2002

VIXEL

real world solutions for storage networking

www.vixel.com | 1·888·AT·VIXEL

Introduction

It's often hard to distinguish fact from fiction with SANs.

With vendors over-promoting what's possible at the expense of what's truly practical, hype has become a bit of a habit.

We often don't know—or even notice—the difference between the products of today and the promises of tomorrow. Will iSCSI provide a silver bullet for running SANs over IP? Is InfiniBand really all its cracked up to be? And what about complete Enterprise storage consolidation over Fibre Channel's director switches?

These are just some of the topics that get vendors pumping speeds and feeds, and revving their hype meters to the red line.

It's all very interesting, and sometimes compelling. But it's confusing and often counter productive.

What's missing is a clear view of what *customers* really need from a SAN. In the customer view of the world, the issues are simple, direct and real: What can I do with SANs that I could not do before? What is the value in terms of my business? What does it cost? How does it work? And where are the hidden liabilities that could hinder scalability, security and management?

It's a view that's not nearly as fun as churning out hype, but it has much more validity and value from the customer perspective.

In most of the examples that follow, the barometer pressure exerted by hype sharply exceeds the gauge of reality. But in a few poignant instances, the promise may actually *undersell* the true customer potential. To keep things fair, we'll take a look at a mix of acronyms, architectures, standards and boxes. And in each instance, we'll ask the same questions: Is it real? Is it necessary? And is it cost-effective for customers?

A Quick Look Back

Before we review the technologies that may or may not lead to the future, let's take a look at the evolution of the SAN market. Identifying past progress and present conditions helps to paint a clearer perspective on future challenges.

In the Beginning...

We have just finished the first stage of SAN evolution – the Homogenous Era. Here, small SAN islands were populated with pre-tested and integrated components from single vendors.

The real challenge was bringing SANs to life, learning how they operate and helping them grow.

For the most part, these SANs included multiple servers attached to one or more RAID storage subsystems deployed in a redundant manner, i.e., two separate interconnect SANs between the servers and the storage subsystems.

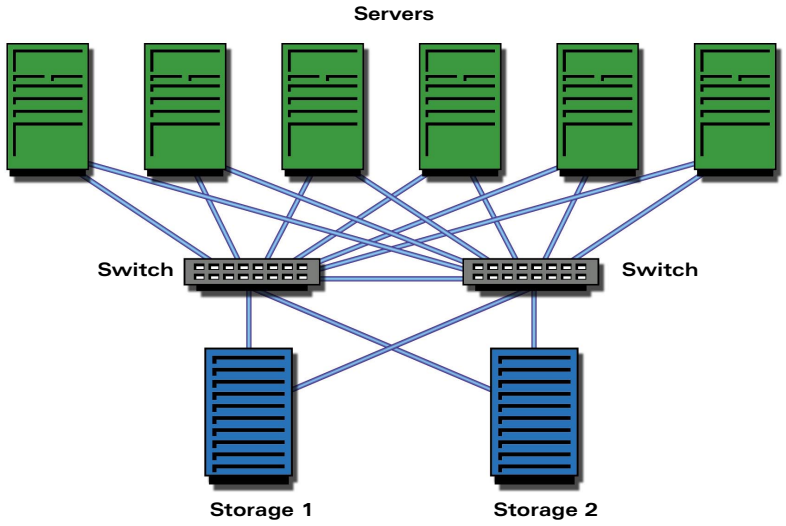

During this era, a law of SAN economics started to emerge: In cost-effective deployments, the cost of the SAN infrastructure is about 5 to 15 percent the total system cost. Unfortunately, not all deployments were cost-effective. Many customers overpaid for SAN infrastructures, since solutions were pre-packaged from a single vendor and were often over-engineered and over-priced.

No SAN Is An Island...

In the second stage of evolution, another law of SAN deployment will appear. This is the law of leverage, where customers will want to leverage their SAN based storage assets as much as possible. Here SANs will embrace IP networks and include storage virtualization.

Until now, Fibre Channel has worked well in high performance storage applications, such as database-to-storage access. It is natural that in the future, lower performance applications residing on IP Ethernet networks – such as mail servers, network directory servers, etc. – can benefit from storage consolidation and central management.

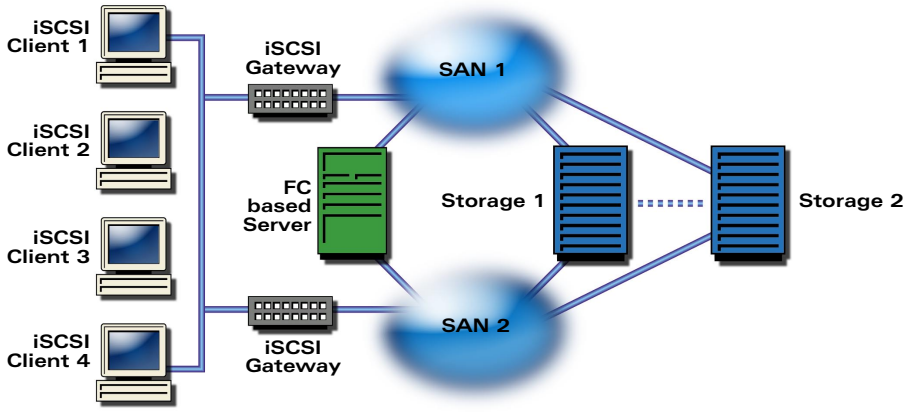

Storage virtualization (the broad term given to logically managing storage as a single entity by abstracting the physical connection of the storage to the servers) will one day provide the ability to purchase anyone's storage and manage it the same way.

There are many types of storage virtualization – host-based, network-based, symmetric and asymmetric – with pros and cons to each. The important concept here is simply to make core storage management operations (such as storage snap shots, backup, replication and mirroring) easier to manage as one entity, rather than spreading them across several different devices from several different manufacturers.

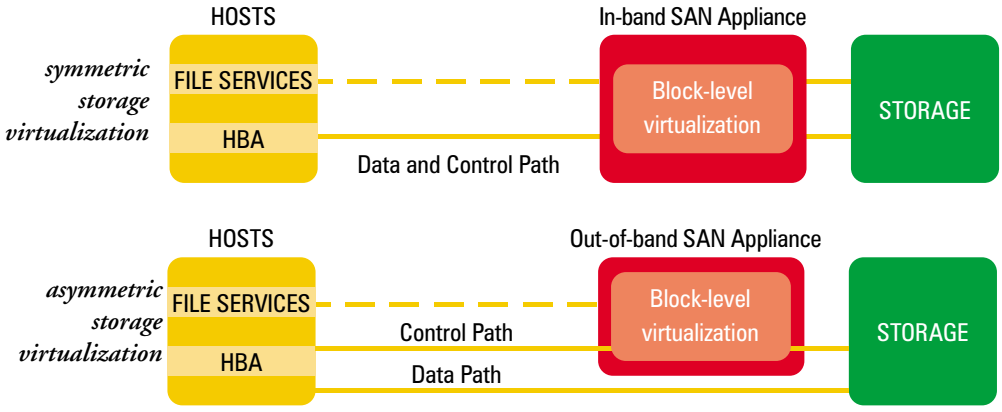

Naturally, there are scores of issues to resolve before large scale SANs with thousands of devices can be deployed effectively and securely. Throughout this evolution, industry experts and knowledgeable customers agree that by starting small and deploying SANs with just a few to a few hundred devices, companies can make optimum use of the infrastructure and gain essential practical experience.

Below are topics of conversation in the SAN industry. To help illustrate the gap between industry hype and customer needs, you'll find meters rating hype vs. reality for each topic. In a couple of instances, you'll find an additional "future reality meter" that indicates the technology's future potential for solving real customer needs.

IP Storage—Too Many Choices

There's an old saying: "When you don't know where you are going, every road gets you there."

Such is the case with hybrid IP and Fibre Channel-based SANs. Not only do all roads promise to get users to a world of happy coexistence, the roads even sound alike—FCIP, mFCP, iFCP, FC-BB, iSCSI. The opportunity has attracted a lot of attention, with more than 200 vendors working on iSCSI solutions. One of the reasons for such a high level of interest among vendors is that IP storage provides entree into the high growth storage networking arena without requiring core Fibre Channel based technology.

To make sense of this alphabet stew, let's identify what are the most common problems that need to be solved. From the user's perspective there are two. First, how do you interconnect the disparate SANs that have already been deployed? Second, how to you consolidate storage from servers that are on an IP network?

The simple answer to the first problem — interconnecting SANs that are already deployed – is Fibre Channel. But the answer gets more complicated when the SANs are on different campuses and another network is involved. If the SAN islands are being connected over the Internet (which has a high error rate and low performance) the answer is FCIP with its lower performance but TCP-based retry capability. If the SANs are being connected over lease lines, then it's FC-BB, which can scale to very high speeds but does not utilize the TCP protocol, with its inherent rate limits.

The solution to the second problem — consolidating storage from IP-based network servers into storage subsystems — is iSCSI. By promising to bring block storage access capability to servers on a low-cost IP Ethernet network, iSCSI helps leverage the investment in expensive SAN based storage subsystems. And that's good news in these tough economic times.

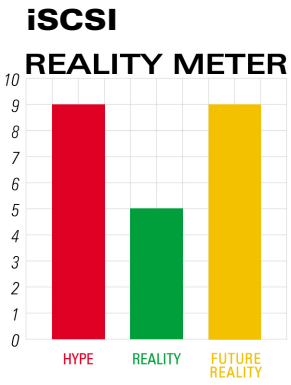

Meter Reading:

iSCSI is a technology, providing an on-ramp for IP Ethernet-based servers to access a Fibre Channel SAN. Although it makes for dramatic copy, until the day that efficient NICs or 10 gigabit per second technology is introduced, iSCSI will really only enable the growth of Fibre Channel based SANs. In addition, since iSCSI is a new technology, it will still have to go through the definition, interoperability, deployment and management phases. That being said, with over 200 companies working on iSCSI solutions plus the leveraging of incumbent IP networks, iSCSI should make great strides in the makeup of storage area networks.

mFCP, iFCP

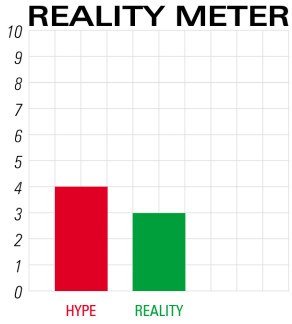

Meter Reading:

Both mFCP and iFCP are good ideas with a critical showstopper. The idea behind mFCP and iFCP is to connect Fibre Channel devices using the ubiquitous IP network. The problem is that the Fibre Channel network is now twice the speed of the Gigabit Ethernet market, so the user still has to purchase more expensive Fibre Channel host bus adapters and the standards do not define connection to Fibre Channel switch ports.

FCIP and FC-BB

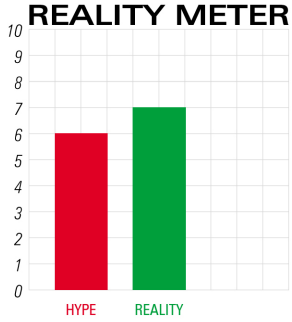

Meter Reading:

FCIP is an IETF standard that bridges two SANs over an IP network using the TCP/IP protocol. This protocol has the advantage of implementing error correction and detection (i.e., retries if the IP network has high error rates). This is an ideal way to bridge SANs over a lower performance and higher error rate IP network. FC-BB is the Fibre Channel Back Bone standard and describes ways to bridge SANs over many types of networks. One of the standards in FC-BB describes an IP encapsulation method that does not require a retry method (i.e., it relies on the upper level SCSI error methods). This IP encapsulation can be accomplished in hardware and can scale to multi-gigabit rates. The only requirement is that the network must be high speed with low error rates, such as fiber optic networks typically implemented in leased dark fiber and MAN environments.

Fibre Channel Arbitrated Loop Switches—The Nanotechnology of SANs

With all the talk about SANs providing business continuance advantages there is one element that is always overlooked, the economics of SANs. Business continuance has as much to do about availability as it does to do with a profitable business model. So, what's hindering even wider deployments? For the benefits of SAN infrastructures to truly penetrate the small departments segment of the market, they must become lower cost and easier to deploy and manage. It's difficult enough finding IT personnel with SAN knowledge at the departmental level. It's harder still to cost-justify their employment. Let's repeat that little fact: switched SAN interconnect is too expensive today for a majority of small environments that could benefit from their deployment. Yet SANs within departments have many benefits, since these environments still contain a lot of SCSI-based storage and need solutions that are of similar complexity to deploy and manage. The answer to this potentially huge market is a Fibre Channel Arbitrated Loop Switch or simply "loop switch."

Loop switches are similar to Fibre Channel switches, with two important distinctions: they are of much lower cost, and they do not need the complexity and interoperability of fabric services. Instead, they depend on the attached devices for functions like discovery. Loop switches are similar to fabric switches in that they contain non-blocking switch cores, support zoning and performance monitoring on each port, etc. But since they don't need complicated fabric services, they avoid the management overhead and complexity required for deployment and maintenance of SANs that comes with these services. In fact, loop switch devices are contained on a single ASIC whose cost is so low they are being used in the backplanes of storage subsystems.

Single ASIC loop switches are a disruptive technology to Fibre Channel switches — where cost and/or manageability have become barriers. These environments include low cost NT and Linux server deployments, video streaming environments, and environments where the many to few rule applies (i.e., many servers to few storage subsystems — where many is a dozen or fewer servers). But the best part of loop switch technology is they are interoperable with all Fibre Channel devices.

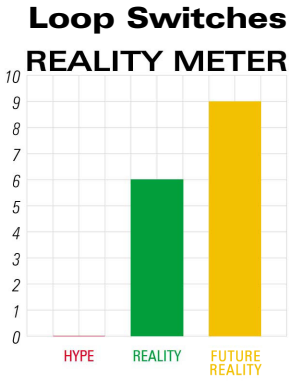

Meter Reading:

While switches are installed in a great number of SANs, a majority of those SANs contain fewer than six servers connected to one or two RAID storage subsystems. These are ideal environments for the robust, high performance, low maintenance and easy to deploy loop switches. Plus, these devices are being embedded in a large number of storage subsystems. While today's reality meter has yet to rev up, in the future loop switches will take the entry level SAN infrastructure by storm.

SAN Quality of Service —Waiting and Wondering

In the world of LANs, Quality of Service (QoS) is generating real benefits by taking a best-effort IP delivery network and providing the ability to offer differentiated services and enhanced security. The result is more revenue and control of the IP infrastructure. But how can QoS be used in SANs?

First-generation SAN build-outs focused primarily on separating storage from servers and enabling new and faster applications. The SAN infrastructure treated all devices similarly and provided no access control, no guaranteed service, and little-to-no security. These functions had to be implemented in either the server or the storage subsystem.

To scale and to share current SANs, QoS parameters such as access control and, at the very least, measured service need to be deployed. SAN policies are needed to provide the translation from business agreements (for example, Service Level Agreements) to the instructions SANs will use to implement services, i.e., QoS parameters.

Without QoS, multi-user SAN business models are not very attractive, since expensive storage assets cannot be adequately shared. The QoS solutions introduced in the future must enable higher levels of security in the SAN as well as monitoring and controlling SAN resources. A concrete example of required enhanced security includes the ability to enforce — in hardware — server or user access to certain storage assets or LUNs. The only current switch-based SAN access control is zoning, which works on a device level not on a more needed LUN level. Another example of where QoS is required is the ability to monitor traffic between SAN endpoints which gives a more useful perspective than link-level monitoring when measuring Service Level Agreements. To enforce SAN Service Level Agreements a further step of limiting maximum bandwidth between two SAN devices must be assured. No vendor has addressed this capability.

The important question to ask is where is the QoS function going to reside: in a switch, the server, the storage subsystem, the host OS or a separate appliance?

SAN Quality of Service

REALITY METER

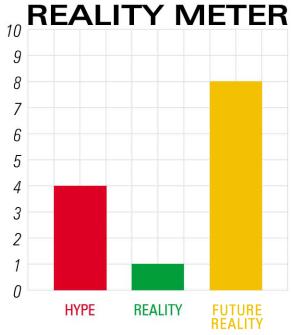

Meter Reading:

QoS: You gotta have it. Without QoS, scaling SANs won't be practical or affordable. While the hype around QoS is relatively small, so is the reality of current solutions. There's plenty of room to improve security, access control and guaranteed levels of service. Over the next year, both the hype and the reality of QoS will greatly increase.

To InfiniBand...And Then Beyond

Okay, believe the hype.

InfiniBand is going to play a major role in scaling current SANs and lead to the next big evolution of SANs. How's that for reality?

InfiniBand (IB) describes an architecture that's based on a switched-fabric concept and is aimed at next generation I/O systems and data centers. The key value proposition behind IB is that it's a fast and scalable architecture used for interconnecting servers and shared I/O systems. By moving I/O out of the server box and onto a switched network, IB technology results in what has been termed "deconstruction of servers," which is anticipated to be a foundational design principle of future data centers. IB is ideally suited to replace all the 1U rack-and-stack servers in the data center compute tiers. IB will increase the "rack compute density" and optimize the power, cooling and other redundant elements currently in the 1U solution.

InfiniBand is supported by more than 200 companies, including the major server vendors. This is huge. Forget about the IB press articles trying to validate the technology from a speeds-and-feeds perspective. IB is about power and control. What's at stake is the next generation data center and control of the infrastructure. The struggle is primarily among server vendors, storage vendors, and network vendors, and secondarily SAN infrastructure vendors. IB will become a very important structural element to connect servers to storage, since there are many more servers than storage subsystems in any scaled SAN and IB is a server-side technology optimized for server interconnection.

With IB technology, it is not hard to see 10,000 servers attached to a dozen or more large storage subsystems primarily through an IB core with edge or leaf FC gateway switches. IB only gets interesting at 4x or 10 Gb/s link speeds and this will be deployed in a mass scale in 2004.

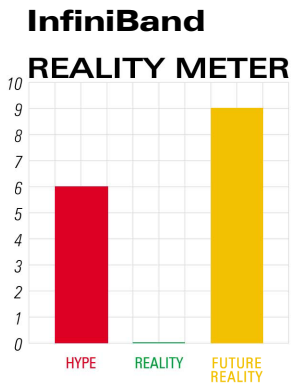

Meter Reading:

Every now and then, hype reflects reality. In this instance, the deal is real. Offering tremendous economies of scale and massive industry support from titans like Compaq, Dell, HP, IBM, Intel, Microsoft and Sun, IB becomes a real candidate for the backbone transport of choice. The real discussion (see below) is what happens to big, expensive and proprietary Fibre Channel directors when low cost, scalable and interoperable InfiniBand architectures appear.

Fibre Channel Directors— The Crusade to Build Battlestar Galactica

It's bound to happen. Vendors have success selling small to medium sized Fibre Channel switches and small port count directors and pretty soon they want to build the Battlestar Galactica director and create the black hole of SAN interconnection — at 60 percent margins, by the way. Sure, you can buy 128x128 by 2 Gb/s crossbar switch ASICs, put them in a hot plug chassis and create 128 port switches (by the way at least one leading vendor is not doing this standard practice, and thus building an inferior blocking architecture!). But what customers want is scalability above the 128 ports, more like thousands of ports for the cost they will shell out for a director based infrastructure.

We really need to step back and address the motivations for implementing such a huge SAN in the first place. Users want storage consolidation, central management, the ability to use common tools to manage storage, and the promise of scalability in terms of both server connections to storage and storage capacity. And don't forget they have to do this all without increasing IT headcount.

But, at a certain point, economies of scale break down for all technologies. Management gets complex, performance decreases, equipment has physical limits, and other technologies begin to emerge with improvements in cost, interoperability, maturity and practical usage. Look at how the computer industry moved on from its sole reliance on mainframes to clustered small server solutions.

So, let's look at the practical considerations of total Enterprise storage consolidation over the mother of all Fibre Channel SANs.

Many issues arise when looking at a Fibre Channel director strategy for Enterprise storage consolidation, but the one big problem is scale and maturity of standards. Fibre Channel directors do not scale in terms of ports, nor does the software required to manage even homogeneous large environments (hundreds to thousands of servers). Maturity of standards is a very big impediment for large scale SAN deployments. For example there is only one routing algorithm defined for Fibre Channel and over five for IP networks. The five are required because when you scale networks and build infrastructure hierarchy, a single algorithm is no longer optimum. The same can be said about device discovery, management and security. All of these protocols and methods are mature in the IP world, developed by a myriad of different companies. In the Fibre Channel arena there are very few companies creating the standards, none of which have experience in scaling networks to the size required for true Enterprise consolidation. Bottom line, the industry can't wait for the current pace of these standards to mature when there are more viable IP based solutions working today.

Many legacy Fibre Channel vendors have announced that their switches will also support multiprotocols such as IP and IB, but in reality will only do so as an add-on. By being an add-on, the hybrid Fibre Channel/IB/IP switch will have a higher cost, be less scalable and have interoperability issues when compared to a pure IP or IB fabrics with leaf interconnects to legacy Fibre Channel networks. For multiprotocol backbone class switches, my money will be on the incumbent networking vendors.

Fibre Channel Directors

REALITY METER

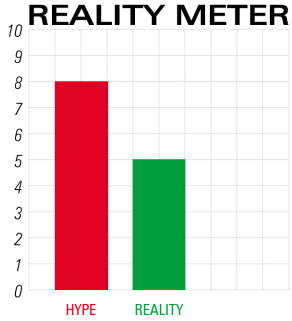

Meter Reading:

"Scale remains the Holy Grail," according to Gartner Group. But the crusade to build the Battlestar Galactica of Fibre Channel directors doesn't bring us any closer to finding it. So, what can users do to consolidate storage for a large number of servers? The simple answer is to look beyond Fibre Channel directors. Instead, implement several small SANs in a cookie-cutter fashion, looking at some of the other technologies discussed above. Utilize both Arbitrated Loop Switches and iSCSI to provide economical access to Fibre Channel storage and experiment with storage virtualization. Initially connect these SANs together with either small redundant switches or a small Fibre Channel director. Over time, these deployments can be migrated and consolidated to an Enterprise level System Area Network composed of a multiprotocol backbone fabric that will economically leverage existing IP or new InfiniBand infrastructures and leverage current IT expertise.

Conclusion: Don't Hide Your Head in the SAN

Now's a time for clarity in the SAN market — not hype.

As economic uncertainty continues to roil the technology industry – trimming IT budgets, slowing the adoption of new products and spreading a gospel of fiscal conservatism – the SAN market cannot afford to veil potential pitfalls behind promotional promises.

Besides the very real threat of stalling the entire SAN market, false promises can translate into very real costs for customers. Too much hype can dupe customers into replacing perfectly acceptable solutions with untried, unproved and over-priced technology.

Don't believe the hype. Now is an expensive time to be driving a forklift through your company. Rip-and-replace is fine for commodity parts, but invasive and excessive for a SAN network.

Instead, be smart. Focus on what you need. And ask three simple questions: Is it real? Is it necessary? And, is it cost effective?

Here are a few final suggestions…

- **Think Big, But Start Small**

  Don't try to get too big too fast. Deploy small SAN islands, get familiar with the technology, and gain practical insights.

- **Embrace Diversity**

  Avoid getting locked into single vendor solutions especially with technologies that are advancing with break neck speed in areas where a lot of innovation is yet to be done. Multi-vendor deployment enhances competition, exerts price pressure and forces innovation. It also ensures interoperability – which is something that will benefit all customers.

- **If It Sounds Too Good To Be True, It Probably Is**

  Be skeptical. Trust your instincts, but rely on your experience, and the experience of others.

- **Remember, No SAN Is An Island**

  While early deployments of SANs were typically small and pre-packaged from one vendor, the future of large-scale, cost-effective, IP-based, multi-vendor interoperable SANs is just around the corner. Be cautious, but stay optimistic. Deployed properly, even the largest-scale SANs will prove real, necessary and cost-effective for customer needs.