

XIV: Storage 2.0

Date: January, 2008

Author: Steve Duplessie, Senior Analyst and Mark Peters, Analyst

Abstract: Today, the data center's ability to meet or exceed the requirements of the business is being challenged like never before. Infrastructure "fluidity" is required at the server, communications and data layers in order to meet today's demands while preparing for an increasingly unknown future. XIV has architected data layer infrastructure solutions specifically designed for the next-generation data center.

Traditional Systems Won't Work Any Longer

The problem with creating a data center service bureau within IT is not people—it's boxes. Stovepipe infrastructure has been the only way to reasonably provide any quality of service level whatsoever to a business or an application. Unfortunately, it is also the least efficient approach from virtually every method of measurement—economic, equipment and management efficiency and, perhaps most importantly, the ability to handle changing requirements.

We have reached an inflection point whereby IT is no longer able to meet the needs of the business unit in any reasonable fashion—which will make it increasingly difficult for the business to use IT as any sort of strategic weapon. The business wants it all—rapid deployment and quality of service guarantees—but won't pay for more people, more data centers or the time necessary to clean the slate. In short, IT is at a crossroads.

Infrastructure is rigid and inflexible by design. "Mission Critical" has been the mantra for infrastructure designers for decades. If the application is worth running, it must run on the most robust systems available. That was true when everything we did was predicated on transactional systems—systems where we could quantify the amount of money lost on downtime or the true cost of even a single transaction.

Times have changed, however. As an information-centric society, we no longer produce transactional data at anywhere near the rate of "digital content." All data begins its life as transactional or dynamic—in a state of flux. At some point, regardless of the data type (database records, objects, files, e-mails, etc.), eventually each of the data "blobs" will become persistent—or fixed. It is this fixed digital content that is currently creating the majority of infrastructure issues today—because of the sheer volume of data created, regulatory compliance, corporate governance and the simple fact that everyone and everything can create digital content almost instantly.

Traditional systems were not designed for unknown requirements. They were designed for just the opposite. Digital content is at direct odds with this philosophy. Core systems built for transactional integrity, speed and reliability are fantastic, but when we ask systems based on these design principals to meet the needs of the new data generation, they inevitably fail.

The New Requirements

Web 2.0 is the catch-all for all things in the next phase of the "everything connected to everything and everyone" Internet. What it means to business and infrastructure is simple to understand: *infrastructure must be able to dynamically deal with scale in any dimension (in real-time, ideally) using inexpensive commodity components with variable attributes (performance, reliability, protection, etc.) to effectively manage itself.*

This is what Infrastructure 2.0 is all about—and it isn't a pipe dream. We have been making strides towards these objectives systematically over the years in each layer of the infrastructure—and clearly the IP networking layer has seen the most advancement of all. With server virtualization technologies taking off, the server layer has seen rapid progress in its ability to create and deploy virtual machines almost instantly—while providing fluidity to those

machines so that an application is unaffected by physical change to a large degree. Applying these constructs to the data layer is a critical next step.

XIV

XIV is an Israeli engineering company led by industry legend Moshe Yanai, who created Symmetrix and is widely credited with architecting one of the most successful products in the history of IT infrastructure. Later on, Yanai set out to build the next great thing in storage.

XIV Nextra

Nextra, XIV's product, is a SAN storage system built around next generation cluster technology to provide dynamic scale, auto-tuning and self-healing to the data layer. The architecture is based upon commodity components to provide lower cost.

The Architecture

Unlike the majority of storage devices which are built upon dual-controller, active/active or active/passive architectures, the Nextra architecture is based on totally independent clustered "nodes"—each designed to perform data storage operations independently and as a member of a "community."

In simple terms, the system is comprised of "Interface Modules," which are low cost servers providing a Fibre Channel or iSCSI connection to the host networks (they present the volumes and data services such as remote replication, migration, etc.), and "Data Modules," which are low cost commodity servers that contain SATA drives and perform the data service execution, as well as caching and snapshot functionality. The systems use commodity GbE switching as a backplane, and can be extended to multiple racks using 10 GbE technology.

The Interface Modules act as data "routers"—keeping track of the distributed data map and pointing any request at the right resource to fulfill it. Theoretically, Interface Modules can scale infinitely as there is no traditional "heartbeat" or chatter between cluster nodes during normal operation.

Data is distributed across every drive within the cluster and total capacity is pooled, so dissimilar drives are able to be fully utilized. The system breaks stored data—regardless of the virtual volume it is a part of—into 1 MB chunks and distributes them across every drive within the system. This eliminates hot-spots and guarantees linear utilization of all disks, without any user intervention.

The system mirrors every write to two distinct disk drives in different data modules. Further availability can be had by replicating (synchronous or asynchronous) to another system—either remote or local. In this way, there are no "spare drives," only spare capacity.

In the event of a failure—of a disk, a cache or an entire node—the system automatically recreates redundancy. Unlike most RAID implementations, this system claims rebuild times of less than 30 minutes for a 1 TB disk—with zero performance degradation.

The company believes it has solved the SATA performance-related issues by the way it distributes data across every spindle, and the way it leverages cache. This architecture can also support future RAID implementations, if customers push for it.

Nextra's snapshot capabilities provide practically an unlimited number of differential snapshots—created instantly—with no effective performance overhead. This is possible because the snapshot functionality is handled within each data module, providing massive CPU and bandwidth power for snapshot redirect-on-write operations.

The Bottom Line

ESG Lab hasn't tested XIV's Nextra and therefore can't verify any of the claims made by the company. However, over 40 customers are currently running it and it has been difficult for us to find flaws—even off-the-record. The theory behind an architecture like this is sound—it solves the issues that the world of Web 2.0 creates. If the stated rebuild times and performance capabilities prove to be true, it is revolutionary.

We all know that having great stuff is no guarantee of success. Israeli technology companies have often proven to have superior engineering abilities, but rarely reach any great level of financial success. The ability to bring technologies to the market in a way that makes them consumable is as serious a challenge as engineering them in the first place. Without global support, for example, will anyone really risk their business on it?

XIV will get the benefit of the doubt because of Moshe—he has earned it. Whether it can use that to leverage success on a global basis remains to be seen.

Finally, this is not the only company developing products for the next generation data center. Most of the successful emerging storage players have a story that is being leveraged in a world requiring scalable, idiot-proof infrastructure. It is only a matter of time before the market dictates that Google and Amazon are not the only companies that should be able to take advantage of commodity components and real-time scale—and that will mean the big players will also have to eventually react to these new demands. Like Google and Amazon, storage infrastructure built to support new digital content requirements does not mean it will necessarily replace storage built for transactional systems. Those companies run their critical business applications on core systems—just like the rest of us. Unlike the rest of us, they realized that the bulk of their information did not require the same treatment, and set out to find new infrastructure to support those new attributes. They couldn't find it—so they had to build it. Other companies are now waking up to these realities and designing for the new wave.

We contend that this new breed of digital content/Infrastructure 2.0 company will find that there is an entirely new accretive opportunity within the market—that this type of product will be additive to users, and not necessarily competitive. Clearly there are negative implications for providers who want users to continue to buy the wrong products for their requirements, but Darwinism has a way of working that out by itself eventually. Having said that, users will (and should) continue to buy infrastructure that meets the needs of a variety of data types and is able to change over time. One size fits all has proven to be false.

What is perhaps most interesting in this endeavor is that Moshe is arguably the inventor and catalyst of the entire monolithic transactional storage array—which created an industry as well as hundreds of billions of dollars of value and investment throughout the 90's—and now he is leading a charge that flies directly in the face of that model in order to address today's new realities.

NOTE: As this brief was set to be published, IBM has announced that it has acquired XIV. In short, it appears to be an intelligent move for IBM as the company did not have these capabilities. Also, the concern regarding XIV's ability to reach a global market is now off of the table. There will be more analysis on this event forthcoming.