COMPANY BRIEF

DataDirect Networks

Date: April, 2008

Author: Steve Duplessie, Sr. Analyst and Terri McClure, Analyst

Abstract: DataDirect Networks has built a reputation developing ultra-high performance storage arrays in niche markets such as high-performance computing and rich media – and suddenly finds itself in the perfect position to take advantage of the explosion of data being generated in the Internet computing era.

Overview

DataDirect Networks has been one of the best kept secrets in storage for over a decade. Much of that is its own doing; DataDirect has historically been happy to stay out of the mainstream "storage" market and instead focus on providing high-value to smaller markets. The company started out in 1988 with a focus on supercomputing, then expanded its focus into broadcast and rich media as it was a natural fit - the large content files, extreme bandwidth, capacity and performance requirements of supercomputing are the same attributes required to store and deliver rich media effectively.

Now the wave of Web 2.0 is upon us – a natural fit for DDN. As Web 2.0 connectivity and technologies become utilized by more and more commercial computing entities', DDN is significantly expanding its market opportunity. Enterprises deploy rich media solutions to better interact with, market to and serve their customers. Today's commercial storage architectures, designed for IO-intensive applications, were not built to support large file based content. Furthermore, most of the existing file-based architectures have been designed to support a different objective – namely being able to store and deliver millions to billions of relatively small, random files to thousands of concurrent users. Very large files such as video and other rich media formats put entirely different demands on a storage system, and as such it is unlikely that general purpose systems of old will be able to adequately meet the demands of the new wave.

DataDirect Networks built its business in the supercomputing niche. DDN is installed on 40 of the top 100 cluster computers in the world, including the BlueGene/L, the fastest supercomputer in the world. The company is also successful in video streaming, broadcast and post production, installed at over 400 television stations and media sites. With more than 200 employees and an annualized revenue run rate of $100M, the company, unlike most of its competitors, is profitable.

The product line offers some of the densest packaging on the market, squeezing 1.2 PB into only two racks. The huge capacity requirements for rich media could mean huge power and cooling bills, so DDN recently introduced MAID functionality, and the systems can spin down inactive drives at a very granular level. Drives can be taken offline and spun down in increments of one, a few, most or all of the drives in the array. DDN has made all this capacity more affordable by leveraging both SATA and SAS drives, offering the ability to mix and match drives in the same enclosure to meet performance and capacity requirements.

Scale-out Challenges

The affordability and high capacity of serial ATA (SATA) disk drives has driven a number of organizations to consider large-scale SATA deployments in recent years. While a number of legacy storage systems have been modified to support SATA drives, early adopters have learned that there are a number of performance and reliability issues to be considered.

Compared to FC drives which are designed for enterprise-class duty cycles and reliability, SATA drives were originally designed for lightly loaded desktops and laptops. As a result, SATA drives are considered far less

reliable when used within large-scale systems—particularly when used for scale-out applications requiring hundreds, if not thousands, of drives. While real world data is still being assimilated, unproven reliability combined with architectures designed for far different I/O profiles lead to issues and problems for traditional transaction-based systems. Specifically, any array leveraging SATA disks in a large-scale deployment will run into issues, which include:

- Silent data corruption
- Intermittent slow-downs
- Capacity and performance limits
- Reliability concerns
- Poor performance scalability
- Excessive power, cooling and space consumption
- Long RAID rebuild times

DataDirect mitigates these SATA challenges in their S2A architecture a number of ways:

Silent corruption occurs when a disk drive sends the wrong data to the host computer, but doesn't flag the data as corrupted. Enterprise-class drives (FC and SAS) have mechanisms internal to the drive to detect these kinds of problems – so they may send bad data, but at least they let you know that they're doing it. That's "Loud" Data Corruption as some call it. However, since SATA drives are engineered to be very low cost, they don't have these built-in error detection mechanisms.

SATAssure™ corrects silent SATA data corruption on-the-fly: Since real-time RAID 6 is implemented within the systems DirectRAID FPGA HW design, every read operation not only reads the data, but both parity drives that store ECC code to ensure accuracy. If a problem is encountered the system will automatically correct the data before forwarding the data to the host. Then the system will re-read the drives to determine if it was a transient error or something needed to be repaired on the disk media. All this happens seamlessly in real-time as the array is self-healing.

Drive Journaling reduces the need for long RAID rebuilds With drive journaling, when the drive goes offline, the system journals all writes that would have gone to that drive in the S2A appliance while attempting to recover the drive. After a successful recovery, only the data in the journal is required to be rebuilt – not then entire disk. This feature works against all the disks in a system – even an entire system outage. Multiple drives may go offline but when corrected, only the journal entries are rebuilt – which can save hours or even days of time.

FastData™ avoids slow-downs associated with intermittently slow SATA drives. This is related to how data corruption issues are managed. In each RAID set, or tier, there are 10 drives – 8 contain the user data and two contain error correcting code. All 10 are read as a set when a user requests data, but only 8are required to fulfill the request. If a read request comes in and two drives of the ten don't respond at all because they're performing internal calibrations, it doesn't matter as the S2A has enough information from the eight drives that did respond to reconstruct the user's data and get it to the host computer. That means in a Petabyte S2A system with 1,000 drives, you could have up to 200 drives responding slowly and still achieve full bandwidth, low latency response times.

The system is designed to operate at 6 GB/s for reads and 6GB/s for writes - simultaneously. The data path through the entire S2A appliance is implemented in FPGA hardware using a real-time state machine design – so they don't just offload parity calculations like most other arrays. Every step of the way from the host ports to the disk and vice versa is handled in hardware. This gives the system not only tremendous speed, but also great Quality-of-Service, or predictability. Third, the system is architected to communicate with drive enclosures in such a way as to parallelize all the operations. Each operation on the array touches a tier, or group of 10 drives. Each drive in the tier has its own channel to the S2A appliance, thus never has channel contention when performing an I/O operation. The system speaks with all 10 drives in the tier simultaneously as a group. The net result is extreme performance that is sustainable under conditions that cripple most other arrays.

Solving these problems enables DDN to gain the economic leverage of SATA without compromising performance, integrity, and reliability – something other vendors will have to also do regardless of the application or I/O type.

New Market Dynamics

The market is changing. The Internet era of computing is upon us and commercial enterprises are going to get dragged into it whether they like it or not. Web 2.0, cloud computing and SOA generated content/data will coexist in commercial enterprises along with both transactional and distributed systems – and will require a mix of price/performance/functionality that differs from both. Most traditional storage players can't (yet) support rich media/high bandwidth, as they are optimized to support transaction processing which requires very different performance characteristics. Just as there became a separate – but additive – type of data to contend with during the arrival of the distributed era (file versus pure block-based transactional data), data generated in the Internet era will also have brand new characteristics. Just as file-optimized storage devices found their way to the mainstream commercial markets to co-exist with traditional core system devices, next-generation arrays capable of addressing the specific attributes and requirements of today's Web 2.0 generated data will also be required.

Don't take this to mean the incumbent players are ignoring the web 2.0 market opportunity and challenges. Within a relatively short time we expect that up to 85% of all data under management in the commercial sector will be created as rich digital content. Just as file data dwarfed transactional data in short order, large-file, rich media data of the Internet computing era will do the same inside organizations. Further, where large orders are still measured in TB's in transactional or distributed environments, they are measured in multiple PB's in Internet computing environments. Large incumbent vendors will not ignore this opportunity, but face the challenge of addressing these new requirements while maintaining their positions in the other markets. This is the window of opportunity for DDN and others to make a name for themselves.

Attacking the Market

DataDirect Networks has made some moves to take advantage of the changing market dynamics. DDN has been content to be a quiet company focused on profitably dominating relatively small market niches. Recognizing the opportunity around Web 2.0 applications and data, DataDirect is investing in bench strength to realize the promise of the market shift. The company has elected to take the gloves off and make some noise in order to educate the market to the fact that DDN has already built products with the attributes required to enable customers to react – and then profit – from their own Web 2.0 initiatives. DDN rightfully acknowledges that just because it may have the best stuff at the right time it won't matter unless they can let the world know about it. While they have the answers, it is imperative that they tell the story or they will quickly be drowned out by those with much deeper marketing pockets.

On the product front, with the introduction of the S2A 6620 DDN broadened its product portfolio to better address the lower end of the rich media market – and an easy to digest entry point for the mass market of commercial computing. The 6620 leverages the same S2A architecture as its bigger brethren, a stable foundation now in its 8th generation, bringing enterprise functionality to the entry level. This system is targeted directly at the traditional media market (CG rendering, media asset management, and both standard definition and high definition editing and archiving) as well as the corporate world. The 6620 is an easy to consume array that non-media companies will find to be a perfect device to store and deliver their own rich media offerings – or those of their customer. DDN knows that if they can seed the market with 6620's their customer will find the product so easy to use and so perfectly in tune with their explosive rich media demands, that demand will explode.

On the distribution front, the company signed an OEM agreement with IBM in Q4 of 2007 that is just starting to kick in. It has an OEM agreement with SGI and Dell is already a reseller. It has been aggressively hiring sales and service personnel to support market expansion.

The Bottom Line

Web 2.0 is all about rich media and big files – where throughput and performance are going to matter. Data is created at the furthest ends of the globe and needs to be accessed in the same way. Write performance and read

performance will matter equally. DataDirect is becoming a high-performance backend component of the Web 2.0 world. It's already added some well known names to the installed base including Xbox Live, Shutterfly, Slide (who creates Facebook applications) and Kodak EasyShare Gallery. These are all multi-PB environments and all have totally unpredictable growth characteristics. Enterprises are going to need to become competent at dealing with this new type of data. Don't expect to see these enterprises throw out their HP, IBM, EMC and NetApp gear and replace it with DataDirect Networks stuff -- mainframe, client-server and rich media are all going to have to live together for a long time to come. But you can certainly expect DDN to start showing up as a legitimate player alongside these guys in the enterprise.

At their simplest, the infrastructural attributes of Web 2.0 are:

- **Infinite Scale** – in real-time, dynamically, with little to no human intervention
- **Self-Management** – infrastructure needs to automatically re-balance and optimize itself – without human intervention
- **Self-Healing** – infrastructure needs to withstand failures and automatically adjust /heal itself
- **Perpetually decreasing commodity costs** – accelerate and leverage declining costs – as soon as they occur

DataDirect is ahead in its ability to scale performance with capacity, and leverages commodity hardware and MAID to drive down costs. If it continues to add resources, enhance products, and execute with sales and marketing, DataDirect Networks could find itself a market leading provider for the Internet computing era in the same way that EMC did in the transactional era and NetApp did in the distributed era. As Web 2.0 applications enter the mainstream of commercial computing and the market kicks into high-gear, DDN intends on staying out front.