THE CLIPPER GROUP
Navigator™

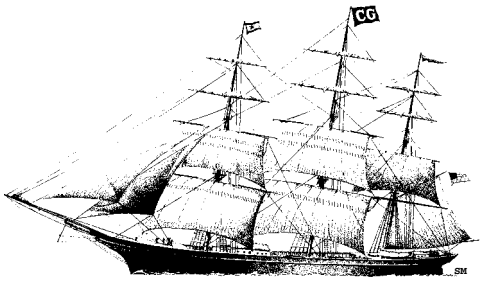

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008054LR October 15, 2008

The Dirty, Little Secret of the Data Center — Data Corruption is Alive on Your Disks *(…and what you can do about it)*

Analyst: David Reine

Management Summary

We all seek a perfect world. The luxury automobile that gets 100 miles per gallon, with gasoline priced at $.99 per gallon, or less. Would you like a delicious chocolate dessert that has only 10 calories and no grams of fat? How would you like a tree in your back yard that grows $20 bills? You need to go shopping – just go out back and pluck a few "leaves" off the tree and enjoy yourself at the mall. Usually, when we fantasize about perfection, we dream of a biblical paradise such as The Garden of Eden, or the "stately pleasure-dome" of Xanadu as described by Samuel Coleridge in *Kubla Khan*. Very rarely do we associate "a perfect world" with an enterprise data center.

If there *is* one word that you would *not* use to describe the enterprise data center, it is *idyllic*. With server sprawl running rampant, the data center is being overrun by a variety of under-utilized, open systems servers, obtained from diverse vendors, and deployed with a single application and, in effect, taking advantage of only 15%-20% of the enterprise investment in compute capability. In addition, each of these servers has access to its own data storage array, either internal to the server platform, or cabled to it in a Directly Attached Storage (DAS) configuration. **Instead of a Garden of Eden, the IT staff must deal with dedicated islands of storage in a turbulent sea of complexity**. In order to simplify the management and maintenance of this diverse processing architecture, the enterprise has authorized the CIO to consolidate both the server and storage environments in an attempt to lower the total cost of ownership (TCO) of the IT infrastructure. The CIO expects this consolidation, and virtualization, of the infrastructure to reduce the number of platforms, and complexity, in the data center, as well as improve its reliability and stability, reduce the administrative overhead required to manage it and reduce the amount of energy required to run and cool the environment. While the CIO may expect these changes to lower the TCO, they are not expected to recreate Eden. However, RAID, Inc., a storage provider for the high-performance computing (HPC) market, as well as many government agencies, has delivered their version of *Xanadu* to help their high performance and government accounts create a more reliable, and less expensive, storage architecture.

Developed in conjunction with NEC, Xanadu was designed to be a modular, open systems storage array, configured with SAS and SATA drives for high performance and scalability, but also engineered for the highest levels of reliability and availability normally found in Fibre Channel (F.C.) arrays, required by the IT staff to ensure data integrity. RAID specified Xanadu to provide the same data protection as F.C. drives, including protection against *silent data corruption*. To learn more about Xanadu, please read on.

IN THIS ISSUE

HPC Data Center Storage Needs

The rampant proliferation of open systems servers and storage platforms within the data center has created an enormous problem for the data center staff. With the typical older x86 server only using 15% to 20% of its computational capability, the IT staff is wasting a significant amount of their IT resources, not to mention the energy required to keep this complex network of servers and storage platforms operational, and talking to each other. By consolidating multiple servers onto a new multi-socket, multi-core platform, and by virtualizing multiple applications onto a single processor, or a single core, the typical data center can make better use of their IT dollars and reduce the carbon footprint of their IT infrastructure. The consolidation of the infrastructure, however, also implies a paradigm shift for the storage environment. Instead of multiple platforms, each with its own internal or DAS storage arrays, the data center must now create shared pools of data, eliminating the isolated islands, consolidating storage onto storage area networks (SANs), with data arrays deployed to service all of the applications accessing them.

What are the implications? First, with all, or a major portion, of the storage on a single array, the reliability of that platform becomes significantly more important. A single failure of a shared storage array can create significant damage to the operation of the enterprise. Any enterprise that is totally dependent on a storage array must be sure that they deploy a system with the highest level of reliability and availability to ensure business continuity for their environment. Fortunately, there are many storage platforms, from a variety of Tier 1 storage providers, that can satisfy the reliability needs of any enterprise: *DMX* from EMC, *DS8000* from IBM, and the *XP* family from HP's *StorageWorks*, to name a few. Unfortunately, these high-end F.C. arrays come with very large price tags, making them much less attractive to the enterprise with a requirement for reliability and scalability, and a limited budget.

Data integrity is a major issue for every enterprise, especially dealing with High Performance Computing (HPC), as well as many government agencies, that consider a single read error "one too many". Performance and scalability are two more factors that are important. With the Tier 1 solutions, you can get the fastest F.C. drives, running at 15K RPM. For many data centers, however, the high capacity, yet lower speeds, of a Tier 2 solution (with SATA drives) meet the enterprise performance requirements, but lack the necessary reliability to meet corporate demands. The newer SAS drives, with less capacity than SATA drives, have more functionality, but not the same as their F.C. brethren. Furthermore, the new storage platforms must meet the same "green" standards as the server network; wasted energy simply is not acceptable.

Data integrity, however, must be the central focus behind the reliability of the IT infrastructure. **If you cannot trust the data input into application processes, you certainly cannot trust the results obtained!** When applications receive data from their storage system that is corrupted, it leads to incorrect results or worse, system crashes. Even though disk arrays include many layers of error checking and correction applications, users still receive erroneous data as a result of misdirected writes, torn writes, data path corruption, and parity pollution. The IT staff has to cope with a set of disk errors that are not caught by the integrated error checking from the drive, itself. These data errors go undetected without read verification after writing or regular searching for anomalies, and therefore go unreported. They are referred to as *silent data failures*[1], even though the cause of the failure might have been a failed write. This silent data corruption is difficult to detect and diagnose and, unfortunately, is fairly common in systems that do not have an *extended data integrity feature* (*EDIF*) to protect against these errors.

So, what is a CIO committed to changing the operational paradigm of the data center to do? Well, one thing the CIO can do is to broaden his or her perspective and look outside what could be deemed the mainstream of storage providers and checkout RAID, Inc., a storage solution provider with a unique perspective on SATA storage reliability.

Who is RAID, Inc.?

If I were to ask, "What is RAID?" many people might rightly answer that it is a *Redundant Array of Independent Disks*. "Who is RAID" is a very different question. One might assume they are an arrogant company trying to claim that they invented RAID technology. A better answer would be that they are a technology provider intent upon *improving upon* RAID technology to provide a truly 100% customized, secure, storage solution. Founded in 1994, RAID has positioned itself to solve the unique storage problems of HPC data centers, along with enterprises, government agencies, and educational institutions. **RAID's goal is to ensure the avail-**

[1] See *Silent data corruption in SATA arrays: A Solution*, by Josh Eddy, NEC, August 2008, and available at http://www.raidinc.com/pdf/Silent%20Data%20Corruption%20Whitepaper.pdf.

ability of trusted critical digital content, in real-time, *all the time*. Their strategy is to provide the industry's only concierge service and support model, to ensure the highest level of personalized service and support to resolve *any* issue. Based in Massachusetts, RAID, Inc. has the industry's first *Storage Network Operations Center*, and the only real-time monitoring and managed service offering, *StorageWatch*, that not only lowers the TCO of the data center, but also increases data availability. With a technical support center operating on a 24x7x365 basis, RAID recognizes the critical nature of each customer's storage assets, assigning a dedicated support team to handle any issue that interferes with the delivery of critical data to the server network. In fact, a personal support engineer can be assigned, in order to build a level of trust and deliver dedicated support and a continuity of knowledge about the deployed platforms. These building blocks have enabled RAID to develop leading-edge solutions and services that assure performance, reliability, and functionality. With a "customer-first" philosophy, RAID has become an industry leader.

Now, RAID has teamed with NEC to deliver an improved storage foundation, a complete storage solution based upon the most reliable storage platform available, *Xanadu*.

The Xanadu Storage Array

Xanadu, based upon the *D3* version of NEC's *D-Series* storage array[2], is a 2U, modular, rack-mount storage solution designed with the reliability and availability features required by the most sophisticated data center, but populated with industry-standard SAS and SATA HDDs, either DAS or on a FC storage area network (SAN). With up to 12 drawers, with 12 HDDs in each, Xanadu can support up to 144 TB of mission- or business-critical data, with 1TB SATA drives. On the other hand, the data center can deploy Xanadu in a single module and as few as three 73GB SAS drives, growing *as required, when required, on demand*, with dynamic provisioning to preserve resources and protect enterprise investment. Xanadu can support 4-12 4Gbps FC host interfaces, with four 3Gbps SAS interfaces to the drive enclosures, for a throughput of 12Gbps. Xanadu employs complete redundancy at the component level with an optional battery backup cache to protect the system in case of power failure for up to 72 hours. With SATA drives installed, Xanadu supports a *RAID TM* triple mirroring architecture, in addition to standard RAID 5 and RAID 6, to protect enterprise data. Moreover, Xanadu provides the staff with an EDIF that calculates a Data Integrity Field (DIF) for each disk sector. This DIF acts like a super-checksum and is checked on the read and write of every sector to identify corrupted data. Thus, EDIF can also avoid the problem of parity pollution that prevents the detection of silent data corruption that other arrays do not. With SAS drives, Xanadu also supports RAID 0, 1, 10, 3, 3DP (Double Parity), and 50. Xanadu also comes with NEC's self-healing *Phoenix* technology to monitor Xanadu continually for any abnormal behavior that could indicate an imminent failure. Xanadu can diagnose and repair most drive failures dynamically. These features provide the data center with the data integrity and scalability they require to manage and grow a viable data store without affecting the availability and performance the enterprise requires.

Xanadu also enables the data center to lower TCO and preserve valuable natural resources by reducing power consumption. With the adoption of MAID[3] technology, Xanadu can power off disks that are not in use, reducing the waste of electrical energy on devices not needed by current applications.

Conclusion

The enterprise IT staff is constantly seeking innovative ways to reduce the TCO of the data center. Increasing the scalability and performance of the SAN, while at the same time simplifying its management, is one way to accomplish that. However, no cost reduction technique will be of value if it adversely affects the reliability and availability of the digital assets of the enterprise.

With Xanadu, RAID Inc., delivers the type of reliability required by the data center for a trusted environment. Xanadu delivers the data integrity and scalability that any enterprise or government agency requires to enable it with dynamic scalability while reducing energy consumption. Xanadu protects the enterprise investment in storage capacity at the same time that it is helping to protect the environment. If your data center needs to provide the highest levels of data integrity with the lowest TCO, it may not be paradise, but Xanadu will certainly ease the burden on your IT staff and improve the quality of data delivery.

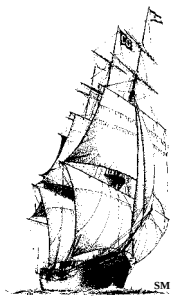

[2] See **The Clipper Group Navigator** dated April 18, 2007, entitled *NEC D-Series Storage Line Spans Midrange Through High End*, which is available at http://www.clipper.com/research/TCG2007060.pdf.

[3] Massive Array of Idle Disks.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group Captain's Log**, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.