THE CLIPPER GROUP
Navigator™

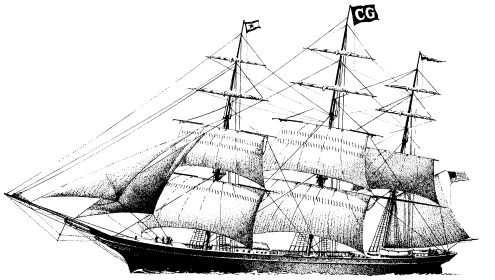

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2003018 April 30, 2003

Managing More Storage with Less Effort — IBM Unveils Its *SAN Volume Controller*

Analysts: Mike Kahn and Mike Fisch

Management Summary

Storage has become the biggest pain for enterprise IT infrastructure managers. In budgetary terms, the ongoing overhead of managing storage arrays far outstrips the outlay for the capital assets. In human terms, non-stop stress dominates the reactionary existence of almost every SAN administrator.

Living with pain requires a tactical response and a strategic approach. The tactical response focuses on coping with an inescapable and unpleasant reality. The strategic approach focus on escaping the previous inescapable. SAN administrators crave pain relief, but not usually to the point of source elimination, because of its career implications. Macho thinking dominates. "I am tough; I can take the pain," they might say. Really, they seek any vehicle for pain relief that keeps them in to loop and preserves their super-hero status. Deaden the pain, and I can do more…" or "enhance my powers, and I can do more in spite of the pain."

Removing the source of pain is a greater challenge, both medically and for storage. (Think about fixing the source of back pain, over taking medicine or acupuncture to block the resulting pain from a herniated disk.) **For storage, tight budgetary constraints at a time of never-ending storage demand require more from storage assets, both installed and to be procured, and managing them by more automated means. This is the inevitable destination that is a journey, not a one step process.** The first round of automation will provide some immediate and important benefits, but do not think of this as your only return on investment. **It can be the gateway leading to the next generation of storage infrastructure and delivered storage services – with automated discovery and provisioning, self-monitoring, -maintaining, and -optimizing to policy guidelines – with guaranteed delivery to selected policy requirements, and the ability to bill users accordingly.**

IBM has an ambitious grand plan for storage automation and management, all keyed to reducing the costs of managing and delivering storage. Not an all-or-nothing proposition, the piece parts will be available in an open architecture. Recently, and with much anticipation, IBM announced its *TotalStorage SAN Volume Controller* (*SAN-VC*, for short), discussed for some time as the "Virtualization Engine" and "Lodestone". When it ships later this quarter, SAN-VC will be a combined offering of software and hardware that resides in your storage area network between your storage arrays and servers. **SAN-VC will coalesce the various attached arrays into a single storage pool and deliver secure blocks of data to server applications, just as before, but with significantly less human involvement.** IBM's initial offering has several limitations, like the ability to connect to only a short list of arrays, but this is common with first releases. Read on to learn more about IBM's SAN Volume Controller.

IN THIS ISSUE

Virtualization Revisited

During the last few years, the IT industry has alternately exalted the phrase "storage virtualization", pooh-poohed it, and handled it with caution and hesitation – sometimes all at the same time.[1] This schizophrenic behavior is not the fault of the concept itself. In fact, storage virtualization is a solid, sound concept that delivers significant value in IT infrastructure. The problem is one of semantics and usage. **Unfortunately, the term virtualization has been overused, misapplied, and frequently misunderstood.** So let's start by setting the record straight.

First, virtualization is about masking complexity through abstraction. This sounds lofty, but it's not, really. Think about making a phone call. You simply dial a number, and the telecommunications network automatically routes the call. You don't have to direct it through the labyrinth of switches, networks, points of presence, and so forth that link your phone and the person you want to reach. The telecommunications network "virtualizes" and masks this complexity, so users only need to deal with logical abstractions, or more simply – phone numbers.

Instead of connecting callers, storage virtualization helps connect servers and storage. It consolidates and abstracts capacity at the block/LUN level[2] and presents to application servers as a single, dynamic pool. In a storage area network (SAN), virtualization can occur on a broad scale because a SAN interconnects all of the servers and storage devices. This breaks the traditional one-to-one relationship between servers and storage and allows capacity to be shared, partitioned, expanded, and reallocated as needed. **As a result, storage management is simpler and less disruptive, utilization is higher, and total cost of ownership is lower.**

Intelligence in the Network

Virtualization is good - and even necessary - in a SAN environment. The next question is where to perform this function. In fact, the same question applies to other storage services, like point-in-time copy and remote mirroring. These are software functions that can conceivably run on any hardware in the data path. Traditionally, they have run on servers or storage arrays, though the scope of these approaches is limited to that particular server or storage array. In environments with multiple servers and storage arrays, and especially different operating systems or storage vendors, it is cumbersome to manage all of the software instances on all of the devices. **If the goal is to simplify management, to do more with less effort, a better approach is to disaggregate the software "intelligence" from both servers and arrays and consolidate it in the SAN.** Here's why:

- **The network provides a common ground through which all data passes.** Software running there has the opportunity to monitor and act upon it.
- **The scope of software features can encompass all servers and storage devices.** Rather than being limited to particular servers or storage arrays, features like virtualization and replication can be applied SAN-wide, to all servers and arrays, whether new or legacy. Software value is maximized, and enterprises are free to use any equipment they want. No more vendor lock-in.
- **Storage management is simplified.** It is easier to manage one super-instance of software in the network than many, possibly different ones on all of the servers and/or storage arrays.
- **TCO is lower.** The bottom line is the bottom line, and the characteristics above work to lower storage TCO, especially as the infrastructure scale. Acquisition costs are likely to be lower, given the procurement flexibility, and operating/management costs are certainly lower.

In short, **moving intelligence into the network is a compelling and more efficient approach for many storage functions.** As more enterprises see the benefits of this next-

[1] For an earlier discussion on this topic, see *Storage Virtualization in 2001: A Space Odyssey* in **The Clipper Group Explorer** dated April 9, 2001 at www.clipper.com/research/TCG2001002.pdf.

[2] Distributed or consolidated file systems can perform a similar virtualization function at the file level. See *Consolidated File Systems – Relieving the Pains of Scale in Data Storage* in **The Clipper Group Explorer** at www.clipper.com/research/TCG2003003.pdf, dated January 30, 2003.

generation storage architecture[3], it will gain significant traction in the marketplace.

IBM's SAN Volume Controller

The *IBM TotalStorage SAN Volume Controller* (SAN-VC) embodies this more efficient approach to storage networking. **SAN-VC is an intelligent, network-resident platform that delivers universal block virtualization, replication services, and centralized storage management.** Any supported server or storage array connected to the SAN can take advantage of these services. SAN-VC consists of hardware and software components.

Hardware

The individual SAN-VC servers or *nodes* connect to host servers and storage over a Fibre Channel SAN. It is an *in-band* architecture, so data passes through a node when traveling between servers and storage over a switched fabric. The nodes are 1U, rack-mountable IBM x-Series Linux servers with:

- 4 FC (2 Gbps) interfaces,
- 4 GB of cache memory,
- Dual Pentium 4, 2.4 GHz processors

In the first release, SAN-VC will support the high-end *IBM TotalStorage Enterprise Storage Server* (a.k.a. "*Shark*") and midrange *IBM FAStT Storage Servers,* connecting to servers running the *Windows, Linux, AIX, Solaris* and *HP-UX* operating systems. Expect other storage arrays later this year. **Customers have the flexibility to mix and match equipment as needed within the support list – without adding complexity.** That is the uniqueness of this technology.

The SAN-VC hardware is designed for data-center class availability. Nodes are configured in clusters of at least two nodes, so there is active-active fail-over, i.e., all data is still accessible if a node fails. It supports multiple fabric paths to servers and storage with load balancing and automatic fail-over. Writes are mirrored in cache across node-pairs. External intelligent UPS units provide backup power, and nodes flush cache to an internal disk if power fails. Data integrity checks (such as CRC, ECC, parity, and chipkill) occur throughout the system. Finally, maintenance and firmware upgrades are non-disruptive.

Performance also caters to the high end. A node pair can reach 140,000 read cache-hit, 40,000 read cache-miss, and 25,000 write inputs/outputs per second (IOPS).[4] IBM's internal benchmarks show negligible differences in response time for random reads performed with versus without SAN-VC in the data path. This should alleviate concerns about the potential performance impact of an in-band solution.

Furthermore, **SAN-VC can expand its aggregate performance and throughput in an almost linear fashion by adding nodes to the SAN.** It can grow with an enterprise's increasing data requirements while still preserving its basic value proposition – universal virtualization and storage services.

Software

More specifically, SAN-VC performs the following storage software services in the network:

- **Storage Virtualization** – It coalesces all storage into a dynamic pool, so administrators can add and reallocate capacity and migrate data transparently among multiple, heterogeneous servers and storage arrays. It lessens administrative tasks, improves utilization, and virtually eliminates disruptions to application servers, even when upgrading or replacing equipment. SAN-VC also supports Managed Data Groups, a vehicle for putting disks into logical groups. As such, it offers some rudimentary capabilities for defining storage tiers with different price/performance characteristics; expect more to follow. Furthermore, virtualization capabilities are granular enough to combine parts of multiple physical LUNs and present it as a single virtual LUN.
- **Remote Copy** – SAN-VC can synchronously mirror data to a remote site to protect against local disasters. The source and target may be different storage models, providing an option for less-costly mirrors.
- **Flash Copy –** This point-in-time copy feature is useful for non-disruptive and LAN-free tape backups, data warehouse loading, and application testing. Periodic

[3] For more details, see *Intelligent Storage Networks – Creating a More Cost-Effective Storage Infrastructure* in **The Clipper Group Explorer** dated February 22, 2002, at www.clipper.com/research/TCG2002006.pdf.

[4] Of course, actual performance varies, depending on configuration and environment.

copies can also be used to quickly roll back to a prior point in time in case of corruption or file restores.

- **Centralized Management** – A central console provides a single point of configuration and control, regardless of how many SAN-VC nodes are deployed in the SAN.

All of these features help enhance the levels of storage service delivered to users and applications.

Benefits

SAN-VC delivers business benefits along two key dimensions – time and money. By simplifying and streamlining storage management tasks, administrators save time and can effectively handle much more storage per person. **It is a stress reliever that helps them to deliver better storage service levels to the business.** In turn, users can benefit from more responsive and available applications, saving them time, as well.

In business, time is money, and all of the redeemed time means a better bottom line. An enterprise does not have to hire more or as many administrators, saving on the largest component of storage TCO (operating costs and management). If users, in general, are more productive and generating more revenue per person, then that is good too.

Furthermore, SAN-VC can cut back on infrastructure costs. Storage utilization in enterprises is typically low (around 40-60% for many applications), and virtualization helps get better use from these resources. The ability to deploy storage tiers with Managed Data Groups can also save money. This permits more precise alignment of storage resources according to the requirements of users and applications.[5] Moreover, SAN-VC could potentially save on server- or array-based software licenses by centralizing intelligence in the network.

More to Come

Though the initial release of SAN-VC supports only IBM storage, later releases will include third-party arrays. IBM will also continue to expand software functionality. One possibility is the dynamic management of *white space*, or allocated but unused LUN capacity, that would further improve resource utilization.

IBM's Storage Vision

SAN-VC is actually part of a larger storage architecture that IBM is unfolding. Later this year, it is expected to include a common file system that's been called *Storage Tank*, to deliver file sharing and centralized management among heterogeneous servers. Under the Tivoli brand, IBM has and is enhancing software for SAN management, backup and restore, storage resource management, and policy-based automation. IBM will also leverage the SMIS/CIM open management standard in a vendor-neutral device management tool. **All of these come together in IBM's grand vision of an *on-demand storage environment* – one that is open, integrated, virtualized, and autonomic (i.e., self-configuring, self-healing, self-protecting, and self-optimizing).** Its goal is to effectively meet the storage service level requirements of an enterprise while minimizing costs.[6]

IBM's storage vision may be sweeping in scope, but it is not rigid or exclusive. An enterprise may consider SAN-VC on its own merits without committing to the whole vision, or it may use another vendor's virtualization product in its open stack of storage software, although IBM will build in synergies that make the whole of storage architecture greater than the sum of its parts, as an incentive.

Conclusion

If you would like to manage more storage with less effort, saving both time and money, consider IBM's new SAN Volume Controller. It represents a next-generation storage solution backed by the credibility of an established vendor. SAN-VC is also a building block toward even greater efficiencies as IBM completes its vision of an on-demand storage environment.

[5] See *Tiered Storage Classes Save Money – Getting The Most Out Of Your Storage Infrastructure* in **The Clipper Group Explorer** dated August 29, 2002, at www.clipper.com/research/TCG2002030.pdf.

[6] See *E Pluribus Unum – The Oneness of IBM Storage* in **The Clipper Group Navigator** dated December 13, 2002, at www.clipper.com/research/TCG2002049.pdf.

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 Authors

Mike Kahn **is Managing Director and a cofounder of The Clipper Group.** Mr. Kahn is a thirty-year veteran of the computer industry. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation — Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull — Director of International Marketing and Support; at Honeywell Information Systems — Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. — a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com*** **or via phone at** ***781-235-0085 Ext. 21. (Please dial "1-25" when you hear the automated attendant.)***

Michael Fisch **is Director of Storage and Networking with The Clipper Group.** He brings over seven years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at Mike.Fisch@clipper.com or at 781-235-0085 Ext. 25. (Please dial "1-25" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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.