THE CLIPPER GROUP
Navigator™

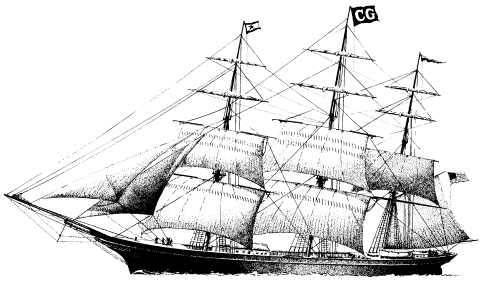

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2010007 March 4, 2010

The Mainframe and Its Storage — The Search for Optimized Infrastructure

Analysts: Mike Kahn, Jim Baker and David Reine

Management Summary

Traveling on a group tour to a set of destinations can be fun and can allow you to see the sights and have the experiences that you seek. Of course, it usually is a compromise of multiple dimensions. You have traded the possibility of complete satisfaction of your requirements (the perfect itinerary of destinations, the amount of time spent at each, accommodations that are tailored to your desires, etc.) for the convenience, simplicity and, typically, lower costs associated with group travel. For many travelers, group travel is a willing and desired compromise. The traveler may not be seeking the optimized travel experience but an easily achievable and seemingly more affordable one.

Optimization is not a word that glibly slides off the tongue of a travel agent, if you even have one in this era of Internet everything. However, it is a frequently used word when it comes to delivering the infrastructure that runs an enterprise's mission-critical applications. **In the data center, there are many focal points for optimization.**

1. ***Infrastructure optimization* – Is the data center getting the best possible returns on investment from its hardware, software, etc., and the people who manage it?**
2. ***Application optimization* – Are the applications being delivered in the manner that allows them to work well and to the priorities established by enterprise policies?**
3. ***Customer and user experience optimization* – Are their needs being met at the levels of expectation set, often described in terms of quality of service?**

This list could be longer but is sufficient to draw a conclusion. **Trying to optimize simultaneously in three dimensions is difficult. For many, this results in scaling back expectations by making compromises**, just like those folks who "settle" for group travel. You may get to where you want to go but it might not be the optimized experience that you might really seek, unless your goal is to make it as easy as possible.

As we all know, some folks are more focused on optimization than most others are. If there are Mainframes in your data centers, those folks who focus on optimization are not far away. **These days, you don't have a Mainframe because *you have to*, it is because *you want to*. You want that *optimization masterpiece* at the heart of your enterprise, pumping information and transactions to all involved – reliably, securely, and flexibly.**

All of this takes more than a supersized server. It requires an integrated solution that encompasses storage and systems software. Read on to see why this is true and what is required.

IN THIS ISSUE

- In Search of Optimization 2
- How the Enterprise Benefits from Storage Optimized for Its Mainframe 3
- IBM's DS8000 Leads the Way 4
- More Synergy via IBM's ProtecTIER Deduplication Gateway for System z 6
- Conclusion 7

In Search of Optimization

So, if the just-presented three dimensions of optimization (infrastructure, applications, and customer/user experience) describe the domain of the problem, what are the characteristics that we seek in a solution? Three come to mind.[1]

1. *Infrastructure coherence*
2. *Optimized execution (performance)*
3. *Interlocking architectures*

These will be the framework for discussing the Mainframe and its storage.

1. Infrastructure Coherence

We live in New England, where centuries ago, would-be farmers set out to convert virgin land (undeveloped forests and pastures) into farmland. The evidence is everywhere in the mostly straight and mostly erect stone walls that dot the landscape. Created when the farmer cleared the land for growing crops, most of the stone walls ended up at property boundaries or the side of a road, so that the growing area could be maximized. Yes, it meant that every stone moved had to be painstakingly relocated to the edge of the field (often, the property boundary) but the effort was worth it, especially to the many generations to come. Although much of the old farmland has been converted to other, supposedly more productive uses, today's lesson is about the longevity of these walls. The main reason for their longevity is pretty simple – *the pieces were put together with forethought and care*. There was structural coherence in those walls. Not only were the pieces fit together to maximize the volume consumed[2], each stone provided support for the ones around it and stayed in place without the use of mortar. Let's call this quality *infrastructure coherence*.

Infrastructure coherence is very important to the data center and its enterprise, as well. Over the last 35 years, the pendulum has swung all the way out and back on the value of infrastructure coherence. When the IT world was very proprietary, almost all parts of a computer system came from a single vendor, because there was no plug compatibility (whether for connecting storage, terminals, or printers). These tightly-knit "systems" exhibited strong infrastructure coherence, as the pieces fit together better than any possible alternatives. That was good for the enterprise, except possibly for the lack of procurement advantages that resulted.

Minicomputers, PCs, networks, increasing levels of standardization for peripherals (including storage) all pushed the pendulum to the opposite extreme. Costs per unit of computing and storage did come down but the administrative burden increased significantly as enterprises now had to integrate the more-open components into their customized solution. "Open" and "small" was felt to be good and proprietary offerings, especially large ones, were seen as evil to be avoided. Unfortunately, the coherency of the IT infrastructure declined to the lowest common denominators. Today, big is not seen as bad, at least to the degree of decades ago. With the "rediscovery" of server virtualization for commodity platforms, partitioned large(r) servers are seen as a way to optimize IT budgets. This brings us back to the present round of discussion on the importance of infrastructure coherency.

Let's make this as simple and clear as possible. **Greater infrastructure coherency is better than lesser infrastructure coherency, unless the costs of achieving it are way out of line.**

2. Optimized Execution/Performance

Today, we are all "speed junkies"; nothing seems to be happening fast enough to keep pace with the work that has to be done. We want "to do more" with fewer resources, especially time. *Tempus fugit! (Time flies!) Hurry up!*[3]

Waiting for something to happen so that you can proceed with the next step is at the top of the list of unproductive behaviors to be avoided. If a high-cost asset (like a big server) has to wait regularly for data to reach it (say, from a storage device), that delay can add up to a lot of wasted resource (and extra spending). While the Main-

[1] Yes, answers tend to come in threes. While some may see philosophical or religious reasons, we think that most of us have trouble thinking about more than three interdependent variables at the same time. Of course, there are more than three variables that we might discuss, but we won't, because it will get messy and very quickly be hard to comprehend. There is always another bulletin to write!

[2] As you could only make a mortarless, straight wall a couple of feet high before you had to widen the base and angle the walls to keep it from collapsing.

[3] The chorus from an old Broadway tune "Hurry Up, Hurry Up" from *The Pajama Game* said it so well:

Hurry up, hurry up, hurry up, hurry up
Can't waste time, can't waste time, can't waste time, can't waste time
When you're racing with the clock
When you're racing with the clock
And the second hand doesn't understand
That your back may break and your fingers ache
And your constitution isn't made of rock
It's a losing race when you're racing with the
Racing racing racing with the clock

Lyrics by Richard Alder and Jerry Ross.

frame is very good at allocating and using resources, all resources (think "storage devices") are not the same and the net amount of work done is affected by these differences.

3. Interlocking Architectures

IT product vendors find themselves divided on how to develop and deploy their products into the marketplace.

- *Do they develop a part of the solution (say, a standalone storage system) to run with any server (or set of servers)?*
- *Do they desire to optimize the efficiency of those servers, especially where their products are not dominant?*
- *Do they do everything to make the resulting integrated system operate optimally?*
- *Or, do they limit their objectives and focus on making the best general-purpose or multi-purpose product?*

These are important things to know when you are buying parts of an integrated system, especially one where you are responsible for its optimization.

We all know that the Mainframe is different in many ways from other kinds of servers. It manages its resources far more granularly and with greater levels of logical and physical protection and separation. It uses protocols that are used by few other systems (like FICON for communicating with peripherals, primarily storage devices). Thus, Mainframe storage inherently is different architecturally and in the degrees of those differences. The more that a Mainframe storage device "interlocks architecturally" with the Mainframe, the more opportunities there are to optimize what is being done. *Are tightly interlocking architectures a bad thing (because they might be seen as proprietary) or a good thing (because of the potential for optimization)?* This question is both technical and philosophical

How the Enterprise Benefits from Storage Optimized for Its Mainframe

Storage is akin to a server's dancing partner. Working together, they can be as a singularity. Separately, they are just two unattached players on the floor of the data center.

Well-architected storage can coax more productivity out of the Mainframe, and vice versa. That's what creates interesting possibilities with the potential for a synergistic payback.

It seems easy, make that easier, to keep the two at arms' length, by not utilizing the features that constitute well-conceived, interlocking architectures, but it also usually means giving up some valuable capabilities.[4] Each can take the often heard attitude of "not my job" when it comes to making the other partner more effective. Unfortunately, this leaves folks in the data center with a bag full of fasteners and some big problems to solve. And there can be a lot of finger pointing among the vendors when something goes wrong. More than any other server-storage pair, the Mainframe and its storage really benefit from working together. Here's why.

Mainframes and big storage systems each are complex entities, with more inside and more going on than with other high-end products. It is a story that extends beyond the large-scaling that separates them from the competition. Their operating software is more complex and each is capable of delivering higher levels of quality and service, including better reliability and security. These are the reasons that enterprises pay more for them. More is expected and they deliver more. So, you might say that each is optimized to excel in its own domain. That is good and to be valued heavily.

But why stop here? Each has the potential to make the other more productive. In baseball, the pitcher and catcher can seem to be somewhat autonomous, with each doing their job as well as each can. However, the greatest effect is when they work together as a team, using the same strategy and approach for retiring the opposing batters.

The Mainframe can make storage more productive by understanding how its resources are used[5] and optimized, by sensing when it is overloaded and attempting to work around it, and, most importantly, by understanding how its software works and optimizes. The same is true in reverse. If the storage system understands the Mainframe's workload management and the priorities of the load ahead, especially for movement of large amounts of data, it can modify its behavior as well. It's like two kids on different ends of a seesaw. By watching the cues from the other and the pattern of movement of the seesaw, each can improve their experience and the experience of their partner.

That's what this is all about, even for IT infrastructure – *delivering more by working to-*

[4] Ultimately, the value of the synergies makes the effort to use the special features worthwhile or makes operations easier in other ways. For example, using dynamic volume expansion is easier than quiescing work, taking a volume offline, expanding it, and then bringing it back online. In this case, it probably is better to let the system do the work.)

[5] Especially, cache. Unless the Mainframe understands how the array caches, and vice versa, each is forced to operate blindly.

gether. Let's look at the *IBM System Storage DS8000* to see what synergies are in play and leveraged with the Mainframe, IBM's *System z*.

IBM's DS8000 Leads the Way

IBM's DS8000[6] is the storage subsystem designed not only for the most demanding general-purpose enterprise storage requirements, but also with special attention for the System z, especially the *z/OS* environment. **It is intentional that many features from the DS8000 can be exploited by System z and vice versa, to deliver optimized execution to both.** Because these are IBM products, IBM has planned, architected, engineered, designed, built, tested, supported, and documented them together. Even before the engineering stage started, IBM architects have identified important points of synergy that could be exploited for enhanced performance and long-term competitive product differentiation. This is a prime example of infrastructure coherence.

Because IBM controls the I/O channel and the interfaces used on System z and on the DS8000, it can exploit embedded capabilities not yet shared with other vendors, such as those "reserved for future use".[7] Using this system of interlocking architectures gives users advanced technology now rather than the possibility that it might be available in the future. Moreover, it reassures them in the process, because it is coming from the original source.

Constant Drumbeat of Features and Functions Important to the Enterprise Class

The DS8000 has all the functionality features expected in an enterprise-class storage system today. In fact, it defines the category more than any other vendor does, although its competitors do offer some features that are not on the DS8000. The completeness of the IBM storage portfolio is the envy of its competitors. Customers can find virtually any functionality they need - *now*. Let's look at several relevant examples from the DS-8000's feature list, categorized for easy discussion.

Performance Features

- ***High Performance FICON (or zHPF)*** uses advanced protocols to reduce channel utilization significantly between the host and its storage subsystem. More data gets through the pipe faster, thereby using the resource more efficiently and improving overall application performance. Unused capacity in the channel represents "headroom" or growth potential for additional users needing to be served or for new I/O-heavy applications.
- ***High Performance FICON – Multitrack*** (introduced in 2009) reduces channel utilization for channel programs accessing more than one Model 3390 track of data. (The first release of zHPF was limited to operations no larger than a single track.) More data can be moved per channel access thereby increasing throughput and efficiency. Information transfers are denser, meaning less time and processing spent on overhead.

Availability Features

- ***Basic HyperSwap*** allows a single System z server to swap I/Os from one local disk system to another local disk system while applications remain online, in case a disk system fails. It can also protect the user while repair operations are completed – unbeknownst to the user community - since performance is not degraded while in the swapped position. Basic HyperSwap is standard in z/OS. Although it is not as comprehensive as IBM's *GDPS HyperSwap* capability, it thoroughly can satisfy the needs of many enterprises. Note that *Geographically Dispersed Parallel Sysplex* (GDPS) satisfies the difficult multi-location networking, timing, reliability, and data sharing requirements of the world's largest and most complex multi-site installations.

Growth Features

- ***Dynamic Volume Expansion for 3390s*** allows volume sizes to be increased while data remains online. Users can avoid an application outage while the volume's capacity is increased. The importance of staying online all the time cannot be underestimated in today's Internet economy.

---

[6] For some background on the DS8000, see **The Clipper Group Navigator** dated August 30, 2006, entitled *The IBM DS8000 Series of Enterprise Storage is Fast, Scalable, and Now Tiered*, available at http://www.clipper.com/research/TCG2006078.pdf.

[7] Like any vendor, IBM is free to develop new features, functions, and interfaces into its products. IBM does not have to disclose the technology and innovation that makes these possible, although at times it has chosen to do so (as with contributions to the open source community). Similarly, nothing forces IBM to license the detailed specifications for features, functions, or interfaces with its competitors, although at times it has done so (presumably for a negotiated fee). To be clear, it is the *specifications* that are being shared, not the underlying code or technology. These limitations on what it shares or licenses and when and how it chooses to do so means that IBM's product competition is always behind IBM in the compatible features that are being delivered (because it takes a long time to get from a licensed specification to a released compatible product) and because the competitors do not choose to provide all of the features that IBM offers.

- *Larger 3390 volumes (also known as Extended Address Volumes or EAVs)* support storing more data on the same or fewer volumes. This feature also helps avoid reaching the ceiling for User Control Block (UCBs), an architectural constraint for z/OS. An incremental advantage of Extended Address Volumes is easier management of data.
- *Space-efficient volume copies* reduce the capacity needed to store point-in-time copies of volumes in the same disk system. Only changed data needs to be captured rather than full copies of all data in the volume.

Configuration Flexibility Features

- ***Use of Solid-State Disks*** as another disk media type in the storage subsystem.[8] These storage devices are extraordinarily fast because, as integrated circuits, there are no moving parts. Yet they can be organized as if they were typical data files on traditional "round, brown rotating" storage. With immediate I/O availability and lightning speed, they are often used for applications (or portions of applications) needing state of the art performance. Putting a database index onto SSDs ahead of a very deep data warehouse is typical. Another immediate usage model is Information *Life Cycle Management or ILM*, wherein most recent data is posted on SSDs while older information is migrated to lesser-performing storage tiers and ultimately retired to tape.

  Once deployed, SSDs require simple yet sophisticated management. IBM has addressed this need by a feature called *DFSMS recognition of SSDs*. (DFSMS, a portion of the z/OS operating system that focuses on storage, helps customers identify high-performance solid-state disks for controlling data set allocation, and for reporting purposes.)

Security Features

- ***Self-Encrypting Disks*** - This feature protects data against unauthorized access if individual disks are removed, or when an entire disk system is redeployed or retired. Because System z is deployed in so many mission-critical environments, this security enhancement is appreciated by many users. Making the encryption of data into an embedded function of the drive itself means that it is automatic and avoids the necessity of purchasing add-on hardware or software plug-ins. The result is ease of management and maintenance on the part of busy z/OS administrators.

Protection Features

For customers with multiple locations, IBM provides protection features that help maintain an additional real-time copy of production data in one or two other locations. Data is protected because local disasters, such as flood or fire, do not necessarily compromise both locations simultaneously. The surviving location can continue to operate until it is reasonable to "fail back" to the primary location (after repairs have been made, for example). And regional disasters, such as wild fires or hurricanes, can be mitigated by making distances between data centers even farther. **IBM products for two-site protection are *Metro Mirror*, *Global Mirror*, and *z/OS Global Mirror). Three-site solutions are built using these products.** Enhancements to these technologies make them ever more functional.

- ***Remote Pair FlashCopy*** provides benefits for Metro Mirror's remote mirroring. When a FlashCopy command is issued to the local disk system to make a copy of a volume or a data set, the Remote Pair FlashCopy capability will automatically send the command to the remote disk system. This has two significant benefits: it reduces link utilization and improves data synchronization when a FlashCopy target is also a Metro Mirror source by avoiding sending each local disk system track as it is copied. In addition, it facilitates the recovery process if operations are moved to the remote disk system.
- ***z/OS Global Mirror–extended distance FICON*** - *zGM* is a long-distance remote copy feature optimized for System z and used by many customers. It is supported by DS8000 and also by other high-end disk systems. *Extended distance FICON* improves FICON protocols so that costly channel extenders from third parties are not needed to span long distances.
- **z/OS Global Mirror–*multiple readers*** – zGM is enhanced by *multiple readers*, which improve performance by parallelizing the processing of writes to be transmitted to the remote disk system. Also, availability is improved (via increased remote site data currency), as are performance and ease-of-management.
- **z/OS Metro/Global Mirror Incremental Resync after HyperSwap** – *Incremental Resync* enhances the DS8000's existing three-system,

[8] For a tutorial on solid-state disks, see the February 10, 2009, issue of Clipper Notes entitled A *New Tier of Storage Appears - Faster, Solid-State Drives State Their Case,* available at http://www.clipper.com/research/TCG2009006.pdf.

multi-target business continuity solution, which uses the DS8000's Metro Mirror and z/OS Global Mirror (formerly known as *XRC*). The GDPS HyperSwap process lets a Metro Mirror secondary (remote) disk system quickly substitute for a Metro Mirror primary (local) disk system that has experienced an outage. Following a HyperSwap, Incremental Resync only needs to send changed data from the secondary Metro Mirror system to the z/OS Global Mirror target at the third site (instead of having to send and then resynchronize the entire volume (or volumes) of data from the secondary Metro Mirror system to the third site, which could take hours or even days). Now, with this Incremental Resync capability, resynchronization can be accomplished in minutes. Not only can Incremental Resync greatly reduce the resynchronization time and bandwidth requirements after a HyperSwap, it can lead to more effective data consistency between long-distance sites very quickly, which limits the risk exposure, should another outage occur (e.g., in the case of rolling power outages).

Follow the Leader

Clearly, IBM is driving the Mainframe storage agenda. However, there are several big storage vendors in the market, each with good product pedigrees as well. Hypothetically (and rhetorically) speaking, *why would the user not opt for them and their products rather than IBM?* The short answer to this important question is that you can get these features from IBM *today* and they form a more complete Mainframe storage solution set *today*. When buying from a competitor, you may have to wait for these features. **However, the demands on the enterprise are here *right now*. So why wait? All of the features described above are shipping right now from IBM.**

Why are competitive offerings so often delayed? It is the nature of being a follower instead of the leader. Typical following behaviors from competitors require the competitor to always come from behind in their attempt to catch up with the leader. In addition, competitors must perform compatibility engineering that creates inherent speed-to-market disadvantages. The typical competitor has many extra steps to deliver an equivalent capability, often taking a long time after IBM has brought the capability to market.

A competitor's first step is deciding what to build and how it fits into the investment profile for competing product developments. There is no small element of risk in this decision because they will be coding to specifications of external interfaces, not internal ones. Negotiating (with IBM) for published specifications can be time consuming (and, presumably, costly). Although IBM occasionally sells its written specifications to competitive vendors, it does not sell code. Specs take considerable time to publish. Most of the time, such specifications become available for licensing about 90 days after initial announcement. While there may be continuing arrangement(s) in place for IBM to offer the specs for new features to its competitor(s), the competitor's decision to proceed may also take time, followed by even more time to assess the specification, assemble the engineers, design a compatible solution, and write, test, and deliver the code. The time advantage is a potent reason to consider buying from IBM.

Competitors often do not have the laboratory facilities to test their own interpretation of the IBM feature. Or, if they do, they may not have enough capacity in the lab to test for scalability. Enterprise class users expect enterprise scale testing from their vendors.

Now there are always reasons why a competitor may have a preferred solution and when that procurement decision may be perfectly reasonable. For example, a competitor may have a unique (for now) answer to a problem that IBM has not yet addressed. Such examples are usually short-lived, however, as IBM turns its attention and engineering to deliver the most complete solution.

Often, IT executives prefer to have a two-vendor shop, "just to keep the playing field honest and keep pricing flexibility". Although there has been some merit to this argument, it is less meaningful today than in decades past because some kind of data mobility between competitive storage systems is offered by each of the competitors. Thus, the staging of one competitor against another (on features, pricing, etc.) lends itself to be more effective when looking to replace an entire generation of storage (rather than when just procuring the next increment of capacity).

More Synergy via IBM's *ProtecTIER* Deduplication Gateway for System z

We have seen, above, how IBM has brought together the mission-critical applications available on System z with the functionality inherent in the DS8700. One major item that has been missing from this synergy has been *data deduplication*, the data reduction technology that eliminates redundant data and helps to manage data growth. With

the acquisition of Diligent Technologies in 2008, IBM added to their open systems portfolio one of the leading innovators in data deduplication technology. With products like the *TS7650G ProtecTIER Deduplication Gateway & Appliances*[9], IBM has established itself as a leader in data reduction in the open systems arena. Now, with the introduction of the *TS7680 ProtecTIER Deduplication Gateway for System z*, IBM enables Mainframe data centers around the world to avail themselves of the same technology.

IBM TS7680 ProtecTIER Deduplication Gateway for System z

Just about every large enterprise is experiencing unparalleled growth in terms of the storage of electronic data, often reaching 100% per year over the last few years. While this may have slowed down a little, due to the stagnating economy, the fundamental trends of making more copies (both snapshots for backup and clones for other uses, plus increased sharing of data and email attachments) continue to put pressure on budgets, floor space, and energy

Managing this growth is causing a significant burden on the IT budget due to the TCO of the storage infrastructure. This TCO includes the ac - quisition and management cost of the storage systems, as well as the costs associated with backup administration, shrinking backup windows, demanding service levels, disappearing floor space and the energy required to run the infrastructure and cool the environment. Moreover, that assumes that there is a sufficient amount of energy available to keep all of your servers humming, disks spinning, and the lights on.

With data deduplication, the data center can store up to 25 times more data on disk, significantly reducing backup and recovery times, improving the reliability of data center backup operations, reducing the TCO of the IT infrastructure (especially the amount and cost of energy), and enabling the enterprise to increase their data retention time on disk. The TS7680 ProtecTIER Deduplication Gateway is a high-performance, highly-scalable inline data deduplication solution that extends this capability from the open systems world into the System z Mainframe data center, providing a FICON interface for the host to gain access to a heterogeneous Fibre Channel storage world. IBM maximizes the options for the enterprise to attach deduplicated storage. IBM has not only qualified the *DS8000, XIV, and DS5000* storage systems, but they have also certified selected open system third party options, as well. **IBM's broad collection of storage solutions provides that Mainframe data center with the opportunity to unify their information infrastructure.**

In order to accomplish this, IBM had to integrate z/OS functionality into the ProtecTIER Gateway. This enables the Mainframe data center to take advantage of ProtecTIER's revolutionary and patented *HyperFactor* data deduplication technology. This data deduplication solution provides enterprise-class performance, scalability, and proven enterprise-level data integrity to meet the disk-based data protection needs of the Mainframe data center while enabling significant infrastructure cost reductions. This serves as a further example of infrastructure coherence and optimized execution through interlocking architectures.

The TS7680 appears to the host as an automated tape library with *3592* tape Model J1A device in 3590 emulation mode. No host application or tape management changes are required. It provides up to 1PB of native disk capacity using 256 virtual drives with up to one million virtual volumes. It can achieve a maximum throughput level of up to 500MB/sec with dual controllers to provide enterprise-class availability, enabling the data center to store up to 25 times more data. This is an end-to-end solution, as System z and the TS7680 ProtecTIER have been architected, designed, and tested to work well together.

Conclusion

Innovation waits for no one. Do not forget that IBM is always staging for their next big enhancement and it has more stakes in this game as the provider of Mainframe servers, storage, and software than those who only supply the storage for the Mainframe. In addition, the importance of an obscure feature may be well known to IBM (internally) but not necessarily yet visible to or understood by competing vendors.

Before you make your next Mainframe storage procurement, take a close look at IBM's storage offerings. If you are looking to optimize your infrastructure, you'll be glad you did.

[9] See **The Clipper Group Navigator** dated February 25, 2008, entitled *Reversing the Requirement for Storage Growth – IBM Consolidates and Simplifies Tier-2 Storage*, available at http://www.clipper.com/research/TCG2009008.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

David Reine **is a Senior Contributing Analyst for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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.)***

Jim Baker **is a Senior Contributing Analyst for The Clipper Group.** Mr. Baker specializes in storage hardware and software, the relationship between technology and its application in business, and emerging technological trends affecting users and vendors in the Information Technology community. He recently joined The Clipper Group after three decades in hardware and software product marketing and product management at companies as large as EMC and Honeywell Information Systems and as small as IPL Systems and Datapoint Corporation. He most recently served as Research Manager for Storage Software at IDC in Framingham, MA. Mr. Baker earned a Bachelor of Science in Marketing degree from the University of Illinois at Urbana-Champaign and an MBA, also from UIUC.

- ***Reach Jim Baker via e-mail at jim.baker@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" when you hear the automated attendant.)***

Mike Kahn **is Managing Director and a cofounder of The Clipper Group.** Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 17 years at Clipper. 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 startup; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, at 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. 121. (Please dial "121" 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

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.