THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2013022R Updated November 18, 2013

Simplifying Long-Term Storage of Historical Data — Spectra Logic Introduces BlackPearl

Analyst: David Reine

Management Summary

Philosopher George Santayana once said: *Those who cannot remember the past are condemned to repeat it*. We all learn from our mistakes and the mistakes of others. We also can pick up pointers from the past that can indicate what to do next in the event of specific actions. In chess, for example, we learn that every specific opening move from white has an experience-based countermove for black. Interestingly, some seemingly counterintuitive countermoves have been proven successful based only on a later series of moves. The key here is to learn all of the potential moves and the counters to each. Unfortunately, you would need to have a rather voluminous library of chess matches, indicating the moves of both sides, with the resulting success or failure, not to mention an equally great memory. The same thing is true for life, in general. In fact, we study history to learn where any course of action may lead us, both good and bad. Again, we cannot possibly remember everything, so it is important to have a library of research material available with the means to access it, quickly and affordably.

This is exactly the problem that every enterprise data center is currently experiencing: **What to do with all of the *big data* that is being amassed on storage devices across the enterprise and around the globe?** Historically, most enterprises have saved important information to disk or tape for long-term preservation, but this approach tends to be unsuitable for archiving very large volumes of objects that need to be selectively retrieved (individually or in a set) and often based on the content of what is stored within. It just costs too much to store one or more copies of mostly inactive data on disk and just keep doing this forever, especially if only a very small fraction of the data is accessed in any given year.

Even with innovations in compression and data deduplication, how can enterprises be expected to preserve and protect the polyglot of data forever to assist us in future decision making, and, of course, while staying within a constrained budget? Formerly, it was said that data declines in value over time. However, today, in many industries, the opposite is true. These days, the IT staff often has a mandate to continue to protect and archive what might loosely be described as *historical data*. Storing a single copy is becoming an insurmountable operational and budgetary challenge for many enterprises, let alone storing multiple copies, whether intentionally (for protection or multiple uses) or not. Unfortunately, preserving data on some form of always-spinning media (like disks) or its upstart, more expensive new brother, solid state devices (SSDs), very quickly can become very expensive because of the repeated acquisition costs (generation after generation) plus ongoing power and cooling expenses, floor space, etc.

Most of the historical data being created and archived these days is of a semi-structured or unstructured nature, but as you will see, it can be structured data, as well. What we need is some new thinking on long-term storage that will simplify how historical data is stored and will lower, significantly, the cost of preservation (when compared to disks). That is the subject of today's report.

Combining historical data with LTFS and other standards, the new *BlackPearl Deep Storage Appliance* from Spectra Logic enables enterprises to use inexpensive media to archive all data. To learn more about BlackPearl, please read on.

IN THIS ISSUE

Enterprise Archiving Challenges

The long-term preservation of data is, right now, a serious problem for the IT staff in just about every enterprise data center of every industry, but most notably in cloud services where social media runs amok, media and entertainment, oil and gas exploration, and life sciences. With a high-resolution camera in the pocket of almost everyone, semi-structured[1] and unstructured[2] data are being generated at a precipitous rate. In fact, **long-term storage has become an issue with several challenges for any data center that needs to retain massive amounts of data.** Four questions define the challenge.

- *How much data needs to be preserved?*
- *For how long?*
- *How will it be located at a later (often much later) time, i.e., is it searchable?*
- *How can we afford to acquire, deploy, and maintain such an archive?*

Summing these up, how can the enterprise preserve the voluminous data being generated by the data center without completely overwhelming the IT budget?

The Historical Problem Decades in the Making

The nature of enterprise archiving requirements is defined by the needs of the business; thus making them hard to ignore. However, this potentially-overwhelming set of storage challenges is exacerbated by the technical methods by which we store much of our data today. Here's why.

Most semi-structured and unstructured data is stored in a file system that retains files in a hierarchy of *folders* (also called *directories*), which exist on physical devices (closely analogous to a *file drawer* in a metal file cabinet). The rigid structure and names of the physical devices, folders, and the files are how we navigate to the right containers that hold the data that we want for a specific purpose. There's only one problem with this, you need to be diligent and painstakingly detail-oriented to keep track of what you put *where* (and possibly also *why*) in order to find what it is that you later will seek. Even folks who satisfy both requirements often find themselves asking "where did I put that file?"

The first reason that our traditional methods are problematic is simple – some things can be categorized in multiple ways – for example, by product ID, release date, or project team. In another example, how should you store and later retrieve your digital copy of your completed 2012 tax return? If you are like me, you have multiple storage devices (think "disks" for now). First you might need to decide on which physical device to put it. Then, you need to decide whether you are organizing by year and then subject ("2012" and then "Taxes" and then maybe "Returns") or by subject "Tax Returns" and then by year). You can see how quickly this can become a big chore and mental challenge, especially when you have hundreds or thousands (or millions or billions) of things to store and keep, and *maybe*, occasionally use.[3]

This leads to the second reason for wanting/needing to change the way that you deal with the problem. The number of things being stored is becoming very large; too large, in fact, for any one person to oversee, especially when many items may be stored for decades or longer and more than one person is involved.

Thus, in the name of simplicity and expediency, the very familiar (and now seemingly outdated) file-based storage hierarchy is being replaced for many data collections by one that is *content-based*. The focus no longer is on physical representations and topical hierarchical structures but now is based on what is being stored.

There are two ways to look at this. First, everything being stored, from a photo to a document to an Internet order to a video, is treated as an *object*. Objects can come in all sizes and be used for many purposes; yes, files likely are the most prevalent object today. One thing makes each object unique and that is a *unique identification key*, whether meaningful (like an order number) or arbitrarily created (like a hash key). The object (often a file and a good representation for visualizing this in your head) and the key are presented to a storage system, which dutifully stores it somewhere and promises to give it back when you present the identification key (and also satisfy the usual security parameters). The result is a one-layer-deep storage system with a limited number of storage and access commands, like PUT (store), GET (retrieve), and DELETE. (If

[1] Like office documents, spreadsheets, and PDFs, which have a defined structure to them when they are stored, but vary in internal structure (like section headings or named variables), content, and length.

[2] Like photographs and MRIs, which only may be bit-defined objects.

[3] While this may seem like a gratuitous or sarcastic remark, it does represent a most common reality of archiving. Most archived data is not used regularly and some of what is stored may never be accessed. The problem is that you do not know, in advance, what will be needed in the years or decades ahead.

Exhibit 1— Universal Characteristics of Generic Deep Storage

Deep storage is a place (i.e., a storage device of some sort) where data objects are stored for a long time; some call it "cold storage" Its primary characteristics are listed below.

- **High capacity** – the ability to hold many PBs.
- **High scalability** – the ability to keep on growing (in terms of the PBs being stored).
- **Sufficiently fast** – it needs to be fast enough for the kinds and volumes of data that are envisioned for long-term preservation; this probably is measured in minutes and not is seconds or fractions thereof.
- **Long-term durability** – the ability for the media to last a long time, measured in decades.
- **Long-term viability** – the ability of the storage infrastructure to last for many years (without a fork-lift upgrade).
- **Very low cost** – where the total cost of ownership per PB is very low (in comparison to other viable storage options).
- **Evolving** – where the TCO per petabyte is getting lower because of improving technologies.
- **Energy efficient** – where the watts per TB stored per year is very low (in comparison to other storage technologies); this is about power and cooling.
- **Space efficient** – where the amount of space required per PB is very low (in comparison to other storage technologies).
- **Reliable** – where stored data is inherently safe, stable, and persistent.
- **Diligent** – where the storage infrastructure increases its reliability by problem detection and error correction before a problem arises.
- **Secure** – where data can be encrypted in transit and at rest.
- **Upgradable** – where transition to denser media or new storage technologies can be accomplished automatically (and without pain).
- **Self-describing** – where a set of data can be independently read and its content able to be recategorized.
- **Simple** – where the storage solution and architecture are easy to comprehend and use.
- **Sufficient** – where the solution does what is needed (in a straightforward way) without a lot of unneeded baggage.
- **Integratable** – where the solution fits into the data center without a lot of new requirements for networking (protocols), location planning, or special handling.

Source: The Clipper Group

you are thinking of REST, you are on the right track. More on REST later – in Exhibit 3.)

Most applications and data uses require more than this (a single key to a single object), which means that a piece of middleware needs to keep track of who stored what (and when and for how long, etc.), who is allowed to access or delete it, what is the object being stored (might be a description or specific keywords, like customer number or map location), and what is in the content that might be used to identify (and request) it. This might be the result of categorizing the object by scanning it for keywords or familiar images or might be the result of something slightly more structured (like reading the header or description often embedded in the stored objects).

It is worth noting that – in the description above – little concern has been given to how or on what the objects are being stored; in this case meaning on what kind of storage device. That's because this approach can and should (at least conceptually) be *device neutral*. Whether the data is being stored on solid-state devices (like computer memory or Flash disks), rotating disks, optical storage devices, or (even) magnetic tape, this one-level deep storage architecture should not care, except for reasons of *performance* (speed to store, retrieve, and transfer) and *total cost of ownership (TCO)*.

Performance and TCO are important considerations for determining where (on what kind of device) to store data. One usually is inversely proportional to the other. For example, if you want something fast, that will cost a lot; if you want to store something at a lower cost, it is likely to be slower to store and retrieve. For frequently used data, performance usually trumps TCO. While no one wants to pay for more speed than necessary, causing humans to wait for critical information that is needed "immediately" (i.e., to

determine if the car seen speeding has been stolen) usually means that a more-performant physical storage device will need to be used. Conversely, when data collections are very large (think super-large, either because of the quantity of objects is high or the size of the objects is very large, or both), cost becomes an important factor and time retrieval becomes secondary. This especially is true when the size of the object is very large (for example, a 2-hour high-resolution video), because the time to transfer the object may be measured in minutes or hours and no affordable mass storage device can make that happen in an instant. Also, the cost of storing something is multiplied by how long you plan to store it. Long-term storage (think, decades) forces one to reverse the previous thinking and make TCO/year as the primary determinant (rather than instantaneous retrieval).

This paper is focused on long-term storage (years, decades, or longer) of many (very many) objects, potentially very large objects. Thus, we are very concerned about how to deliver an object-based storage solution with very low TCO per unit of storage capacity per year or decade. This might be expressed as *TCO/TB/year* or *TCO/PB/decade*. However, you might also want to consider *TCO/object/period* of time, especially when your objects have some uniformity of size. There are many characteristics of this kind of long-term storage. See Exhibit 1, at the top of the previous page, for a useful checklist, with definitions.

Keeping the TCO low for storing and retrieving petabytes of historical data held for the long term is a major IT challenge today. This was discussed at length in the most recent Clipper study comparing the archiving costs of both disk and tape over only nine years[4]. That study showed that the infrastructure costs per petabyte were much lower for tape than disk.

Today, many folks in the data center and in the lines of business are facing the challenge of finding the most appropriate IT infrastructure that can receive and catalog data objects (whether structured, semi-structured, or unstructured) from a variety of sources and locations and place it in a single data stream for transmission for long-term storage at the lowest possible cost, perhaps resources that already are in place in your data center, to ensure the ability to protect and retrieve these enterprise assets whenever needed. In this paper, one solution for meeting the many challenges is explored – *BlackPearl Deep Storage Appliance* from Spectra Logic.[5]

Spectra Logic BlackPearl

Everyone knows that when a hurricane is approaching, the sale of portable power generators goes up. A doctor or medical researcher, however, can use historical data to review previous medical studies to draw new or non-study related conclusions from older studies. **Enterprises everywhere need to retain their historical data so that they can leverage the past in order to make better decisions in the future.**

Spectra Logic's *BlackPearl appliance* facilitates just that. **BlackPearl has been designed to preserve and protect object data while reducing the total cost of archiving data that may be accessed rarely but needs to be retained for future use, potentially forever.** It does this while delivering historical data retrieval times that usually are measured in minutes.[6] It enables the deployment of a lower-cost tape-based object storage system, enabling the data center staff to consolidate their structured, semi-structured, and unstructured data onto a single platform.[7] Combined with *LTFS*[8,9], BlackPearl adds portability to long-term object storage. It enables the data center to manage technology migrations, moving data from one format or media to another. This can be critical when an older technology becomes obsolete and requires new media to enable continuous access to the data. See Exhibit 2, on the next page, to learn how BlackPearl satisfies the

[4] See the issue of *The Clipper Group Calculator* dated May 13, 2013, entitled *Revisiting the Search for Long-Term Storage – A TCO Analysis of Tape and Disk*, and available at http://www.clipper.com/research/TCG2013009.pdf.

[5] In order to lessen confusion, Spectra's product will be called "BlackPearl" (without the added "Deep Storage") and "deep storage" (usually without the capitalization) will be used to refer to a generic place where objects are stored for the long term at a very low cost. In other words, all uses of "deep storage" from this point forward will be generic.

[6] If the data being retrieved is voluminous, whether due to the size of the object and/or the number of objects being retrieved, the actual time to move the archived data to the appropriate destination (disk, SSD, etc.) for its use likely will be longer. The point here is that if you need to retrieve data in seconds or less, then a cartridge in a tape library is not the appropriate medium.

[7] BlackPearl treats all objects the same, regardless of what is stored within the object.

[8] Linear Tape File System (LTFS) refers to both the format of data recorded on magnetic tape and the implementation of specific software that uses this data format to provide a file system interface to data stored on magnetic tape.

[9] See **The Clipper Group Navigator** dated June 28, 2010, entitled *Dealing with Cool and Cold Data – and Getting It Just Right*, which is available at http://www.clipper.com/research/TCG2010031.pdf.

Exhibit 2 — How BlackPearl Meets the Universal Requirements for Deep Storage

Building on the checklist shown in Exhibit 1, here is an item-by-item rundown on how BlackPearl satisfies the requirements.

- **High capacity** – With T-Finity, BlackPearl has the ability to hold 3.6 exabytes of data.
- **High scalability** – Based upon your needs, more capacity or performance can easily be added; simply deploy more drives or enable more slots. When you reach the limits of your initial configuration simply add more frames. For example, a single T950 frame holds up to 920 LTO cartridges, scalable to 10,020. A single T-Finity frame holds the same number of LTO cartridges or up to 594 TS1140 cartridges, while a fully configured T-Finity can grow to 50,100 LTO or 38,115 TS1140 cartridges.
- **Sufficiently fast** – With up to 24 LTO drives available in a single T950 frame, with a maximum of 120 in eight frames, or a T-Finity with up to 960 drives in a library complex, BlackPearl can deliver the level of performance required by any enterprise, up to 2.2 PBs/hour (compressed).
- **Long-term durability** – Both LTO and TS1140 media is rated for 30 years of use.
- **Long-term viability** – The scalability and upgradability of Spectra Logic libraries ensures decades of stability.
- **Very low cost** – The roadmap for both LTO and TS11x0 media supports continuing capacity growth. The track record for both shows comparable new cartridge cost regardless of capacity, driving down the cost per TB.
- **Evolving** – Future technology improvements ensures a reduced TCO for BlackPearl users for the coming decades. In addition to the regular pace of generational improvements for both LTO and IBM TS11x0 drives and cartridges, BlackPearl can support other high-density media, as they become available.
- **Energy efficient** – Energy consumption for power and cooling for both LTO and TS11x0 configurations continues to drop for each generation, as performance and capacity increase, especially in comparison to other storage technologies.
- **Space efficient** – As cartridge, and frame capacity increase for each generation, floor space per TB for each library decreases.
- **Reliable** – Both LTO and TS11x0 have extremely high MTBF figures.
- **Diligent** – Spectra Logic's BlueScale technology closely monitors library performance in terms of problem detection and error correction before a problem arises, increasing reliability.
- **Secure** – All Spectra Logic drives come with both encryption and WORM protection to ensure data protection.
- **Upgradable** – When higher capacity media, such as LTO-7, becomes available, migration to the new media can be performed automatically, without operator intervention.
- **Self-describing** – While BlackPearl is not content aware, BlackPearl's clients can (and almost certainly will) categorize what is being stored via DS3. BlackPearl contains an object catalog that includes DS3-client-created metadata information and physical storage locations.
- **Simple** – Both BlackPearl and Spectra Logic library architectures are easy to comprehend and use
- **Sufficient** – BlackPearl does its job automatically (in a straightforward way) without any unnecessary overhead.
- **Integratable** – BlackPearl can drop into just about any enterprise data center without new (need?) for networking requirements, location planning, or special handling. By connecting to the clients via Ethernet, BlackPearl makes networking easy and usually without added cost for the clients.

Source: The Clipper Group

requirements listed in Exhibit 1.

How Does BlackPearl Deliver Deep Storage?

BlackPearl builds upon and enhances Amazon's *Simple Storage Services* (*S3*) protocol[10] to provide for the easy archiving of

[10] S3 provides a simple web services interface that can be used to store and retrieve any amount of data stored on disk from anywhere on the web. S3 has become a *de facto* standard.

Exhibit 3 — DS3 — Under the Hood

Spectra Logic has enhanced Amazon's S3 with REST to enable large amounts of bulk data to be moved to tape for long-term storage, naming it *Deep Simple Storage Service* (*DS3*). The enhancement has added commands to the S3 set that have been specifically designed to optimize the transport of data objects with tape, using bulk commands to PUT and GET data from the client to deep storage, or vice versa. DS3, however, also will work with any media that might eventually replace tape for deep storage objects. Spectra Logic uses Amazon's RESTful S3 interface concept as the basis for DS3, since S3 has become an industry standard for moving data around in the cloud and the Internet. **Spectra Logic is now making tape appear to all as a vast Web depository.**

DS3 provides a simple and modern approach to efficiently moving large bulk files to and from deep storage. It is a data transport and communications interface that enables software clients to direct and manage bulk storage read or write operations with deep storage. Archived data is accessed over an open interface, such as REST interfaces and web protocols. REST is an architectural style that abstracts the architectural elements within a distributed hypermedia system. REST ignores the details of component implementation and protocol syntax in order to focus on the roles of components, the constraints upon their interaction with other components, and their interpretation of significant data elements. REST has emerged as a predominant design mode for Web APIs. REST enables data to be retrieved by whatever format is being used deep into the future. RESTful storage generally is being accepted as the "right" way to move data to the cloud as well as supporting large numbers of objects (and files) internally.

DS3 is available today only via BlackPearl and all Spectra Logic T-Series tape libraries, but Spectra Logic expects that other vendors will adopt DS3 as an open standard.

Source: Spectra Logic

large amounts of bulk, or sequential, data using an enhanced standard RESTful[11] interface, which Spectra calls *DS3* (for *Deep Simple Storage Services*). DS3 enables the BlackPearl server appliance to act as the intermediary that enables standardized or custom applications to direct and manage bulk storage write operations to tape, as well as other forms of deep storage that may be deployed in the future. Spectra Logic hopes that DS3 will become a *de facto* standard through widespread adoption by other vendors and application developers. (See Exhibit 3, above, for more on DS3.)

BlackPearl uses a high-performance client/server architecture, enhanced with solid state disks (SSDs) to cache data from DS3 sources (which Spectra Logic calls "DS3 clients"), as it moves that data from a high-speed online environment to deep storage. The key word in this sentence is "moves". While data protection strategies and backup processes make additional copies of supposedly active data, archiving is about permanent relocation of data that likely will be inactive (probabilistically[12]), usually driven by a desire to lower the TCO of holding data for the long term.

The DS3 client groups data into collections of like objects. These data sets are sent to BlackPearl, to be cataloged and stored into "buckets", each of which is a logical repository for storing and managing like objects, Newly arriving objects are collected in BlackPearl and held until it determines that it should move the like objects (in bulk) to their final storage destination, which is on tape. Each tape cartridge may reside within only one bucket, but a bucket may extend across many tape cartridges. BlackPearl's batching process is designed to optimize backend throughput and maximize media capacity utilization.

Two kinds of buckets are currently supported: *nearline buckets* for objects contained within a Spectra library, and *offline buckets* for objects stored on media that has been ejected from the library. Offline storage is typically used to store data to be retained for disaster recovery or dispersed across a wider geography in case of a localized disaster. At present, BlackPearl supports both LTO-5 and LTO-6 and IBM TS1140 and TS1140 tape media and drives (and will support later generations that also will support LTFS) stored in any of Spectra Logic's libraries from the *T50e* up through the *T-Finity*.[13] At present,

[11] Representational State Transfer. See Exhibit 3, above.

[12] The great unknown is whether or when or even why a historical data set will be retrieved and then be accessed. Probabilistically (as a percentage of all archived data), the vast majority of archived data will not be used in any given year.

[13] IBM drives and media may be used only in mid-range T-Series T380 or larger SL libraries.

BlackPearl only works with Spectra Logic libraries.

BlackPearl contains an object catalog that includes DS3-client-created metadata information and physical storage locations. Using LTFS, it enables the storage of data in an open, self-describing format. This also enables BlackPearl to migrate data stored for the long-term to new media formats, as they become available. It provides the management of the deep storage system, including inventory, object locations, retries, and error handling. BlackPearl uses a TCP/IP interface over either a 10GbE or *InfiniBand* communications line between the DS3 client and BlackPearl server. The appliance then connects to the deep storage environment (a Spectra Logic tape library) over a Fibre Channel or SAS connection.

Data in deep storage via BlackPearl is retained as objects that are self-describing and written in an open file format. Each cartridge is LTFS-enabled and compatible. This means that a tape from one library may be moved to another library and read independently without a problem. It does this by treating each object as a file, thus allowing any file to be read like any other file in a file system (with the usual security restrictions).

BlackPearl and Spectra Logic's libraries enable the enterprise to avoid over-provisioning, with a "pay-as-you-grow" environment. Additional capacity easily can be added, when it is needed. Deep storage also enables the enterprise to avoid significant ongoing operational costs that often result from long-term commitments to cloud storage providers, with never-ending and always-growing monthly charges and with the potential risk of placing data that is out of your control, far away, and possibly very slow to retrieve. A Spectra BlackPearl solution can significantly reduce the cost and restrictions of storing long-term, relatively inactive data, providing long-term data preservation for as little as $.09 per gigabyte, with an ROI compared to the cost of using a public cloud storage alternative in less than one year, according to Spectra Logic using its list pricing.

The client is the interface between the file and application environment and the object-based deep storage. There are generic clients that have been written by Spectra Logic and its business partners also are writing application-specific clients. The key here is that you have many choices, including writing your own application client for using DS3.

BlackPearl can track copies of data stored on nearline devices as well as in offline and/or offsite tape libraries, with end-to-end data integrity verification for long-term data protection and preservation. It leverages and integrates with *Spectra BlueScale*[14] tape library management interface features, which supply functionality such as data encryption onto the tape media, tape drive cleaning, system error detection, and queued ejects to optimize the export of data to offsite storage.

Conclusion

Today, the modern enterprise has no problem amassing more and more data, not necessarily because they have to, but because they can. How to preserve that data in the most economical fashion is another problem. One option, the external cloud, often presents an issue because you are relying on a third party to protect your assets. With the recent failure of Nirvanix, the reputation of cloud storage has taken a big hit. In fact, it has turned a cloud customer's worst fear into reality. How do you retrieve your long-term storage? With Spectra Logic's BlackPearl, the data center can change your archive from a distant black hole to valuable content.

With network costs continuing to grow, large data transactions can consume a lot of time and a lot of money. In fact, today, it costs less to FedEx your data on tape than to transmit it over a WAN. If you are experiencing the kind of growth that can cripple a budget, and an economical long-term object-storage solution is important to you, then you need to investigate BlackPearl from Spectra Logic. It may be the answer to your long-term historical data storage and protection quandary.

[14] See **The Clipper Group Navigator** dated October 30, 2012, entitled *Spectra Implements New Data Protection Management Facility for Encrypted Data*, available at http://www.clipper.com/research/TCG2012025.pdf.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first year, is an independent publishing and 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 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.)***

Regarding Trademarks and Service Marks

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.