



Choosing the Right Archival Solution — HP Leads the Way

Analyst: Dianne McAdam

Management Summary

Enterprises have always had to need to store information. This information can take many forms. It can be customer records, employee records, financial records, or design specifications for the latest product. In the past, these enterprises have had internal policies that define how long this information would be kept. For example, an aircraft engine manufacturer would keep the design specifications for a particular engine until the last engine of that type was removed from service. Other corporations keep employee records for two years after the employee has left the company, but remove the mailbox and delete all the employee's emails within days after they have left the company.

Today, many corporations are re-examining how long to keep information; good governance practices dictate that information be retained for longer periods of time than were previously viewed as acceptable. In addition to internal policies, government regulations are dictating how long certain information must be retained. Hospitals and clinics must comply with regulations that compel them to retain patient records in some cases for decades and make those available on an ad-hoc basis for discovery, clinical or research requests. Financial corporations, ruled by SEC 17a-4 policies, are now required to keep emails that discuss financial transactions for at least seven years. The Sarbanes-Oxley Act is forcing companies to update their solutions for archiving important records.

Not only are we keeping information for longer periods of time, but also we are in a period of unprecedented data growth. Many enterprises report that data is growing at 50% to 100% per year or more and these growth rates do not show any signs of abating.

Today, many enterprises need to store more and more data for longer periods of time but storing the data is the easy part. Retrieving the data in a timely fashion can be difficult, if not impossible, without the right solution.

The Need for Archiving

A large corporation located in the Northeastern part of the United States was involved in litigation and was required to produce all of the emails sent to and from a client employed by another firm. The company knew that all emails had been backed up for years, so they initially believed that reproducing these emails would not be a problem. The reality was that each email backup tape for the last several years had to be restored to disk. After the email environment was restored from backup tape, all emails in the restored environment had to be scanned to determine if any of the emails were from or to the client in the lawsuit. Then the next backup tape was restored and the search began again. This process continued until the last five years of backup tapes were restored. It took over six months, required the services of an outside consulting company, and the entire process was estimated to have cost over one million dollars. In the end, this company was not sure that all of the emails had been recovered. This costly project could have been accomplished in hours, and days, not months, without requiring outside help, if a proper archiving solution was put in place.

Read on to find on more about the benefits of proper archiving solutions and how HP can help enterprises retrieve many different types of data quickly.

IN THIS ISSUE

➤ The Need for Archiving	1
➤ What <i>Exactly</i> is an Archive?	2
➤ Archiving and ILM.....	2
➤ Advances in Archiving Technology	3
➤ Archiving Solutions Today	4
➤ An Integrated Archival Architecture	4
➤ HP Offers Many Solutions to Solve the Archival Problem	
➤ Choosing the Right Archival Solution ..	5
➤ Conclusion	7

What *Exactly* is an Archive?

Many times, the terms *backup* and *archiving*¹ are used interchangeably, but backup and archiving are two different processes with different end results. Backups have been designed to protect data and be part of a recovery strategy for the first 72 hours after a disaster. If a file or volume is accidentally deleted or corrupted, then a backup is used to restore the contents of that file or volume. Archives, on the other hand, are designed to store data for a specific period of time so that it can be accessed for business purposes. Archives can contain personnel records for employees that are no longer with the company. It can contain financial records that were used to create the quarterly and annual statements. It can contain medical images and records for patients that have been treated months (and years) ago. It can contain emails that contain proposals and contracts. Archives can contain lots of different data. However, archives are not designed to back up and restore data. They contain data that must be managed intelligently and protected for a specified period of time. The data needs to be properly disposed of when the retention period expires. Well-designed archives allow companies to search and retrieve information quickly based on specific criteria.

As a general rule of thumb, backups cannot be easily searched; archives can. The company in the earlier example that required six months to find designated emails discovered that backups are poorly suited to search and retrieve specific data. The need to search and retrieve information is not limited to emails; all mission critical applications contain information that may need to be searched and retrieved.

Application backups back up the same information over and over again. Archives, unlike backups, provide one source of the truth; that is, archives only retain one instance of the email, document, or transaction. For example, all backups must be searched to find all financial transactions initiated by one broker. A consolidated archive need only be searched once. Having one archive to search saves time, money, and resources.

Do You Need an Archive?

There are multiple reasons that may lead a corporation to archive their information.

- **Compliance** - If your company is regulated by legislation, such as HIPAA or Sarbanes-Oxley, then you definitely need an archival solution that can ingest different types of information and allow authorized personnel, such as compliance officers, to search the archives quickly and easily. Some government agencies require that documentation

must be produced within hours or days, not months. Failure to produce the required documentation in the time allowed can result in heavy fines, negative publicity, and lost business.

- **Discovery** - Contract disputes and grievances by disgruntled employees require supporting documentation be produced to settle the claims. Without intelligent archival solutions, it can take a lot of time and money to find the supporting data. After restoring years of backups, the company in our example was never sure that all the documentation was located. If the opposing counsel had documentation that was not available to the corporate attorneys, it could affect the outcome of the lawsuit negatively.
- **Competitive advantage** - Managers and planners can search archives to retrieve information about current and past market sales and use this information to predict future trends. Alternatively, they can review the buying patterns of very large (and very small) customers to target sales campaigns. Archives can turn collected data into information that gives companies a large competitive advantage.
- **Reduce costs** - Archives allow enterprises to migrate emails, files, and transactions to lower cost storage reducing storage and maintenance costs.

An intelligent archival solution can ensure that all necessary information be easily searched and retrieved. Archival solutions are an essential part of good business practices.

Archiving and ILM

Point solutions for managing information help solve very specific problems but often have hidden costs and because they are not part of an overall strategy for information management result in increased costs and overall inefficiencies. Many companies have recognized these challenges and are now looking at *Information Lifecycle Management (ILM)* as a strategy to capitalize on their information and use it as part of the overall business strategy for growth, governance, and profitability. An Information Lifecycle Management strategy recognizes that information can have multiple forms and that it has different levels of business value depending on time, business function and other business policies. ILM captures, manages, retains, and delivers information as defined by the business requirements. The promise of ILM is to reduce the cost of ownership while meeting business and regulatory requirements. Archiving should not be reviewed as a separate process but must be integral to an enterprise's ILM strategy.

ILM recognizes that information has a life cycle with changing business requirements over time. Throughout its life, data may be moved to different storage platforms to align with these changing business requirements.

Archiving as a component of an ILM solution

¹ See **Clipper Notes** issue dated February 1, 2007, entitled "Archiving – Do You Need It?" and available at <http://www.clipper.com/research/TCG2007018.pdf> and The **Clipper Notes** issue dated February 1, 2007, entitled "Archiving- Choosing the Right Architecture" and available at <http://www.clipper.com/research/TCG2007019.pdf>.

provides many benefits, including the following.

- **Reduced total cost of ownership** – Data is migrated from more expensive primary storage to more cost effective storage can defer storage acquisitions. Migrating data to other storage platforms can reduce the duplication of data across the infrastructure and increase storage utilization, while driving down the overall cost of storage.
- **Operational benefits** – Migrating data to different storage platforms increases operational efficiencies by improving access to data. End users experience better response times and operational tasks such as backups and restores occur more quickly.
- **Improved management** – Centralized management reduces administrative overhead and ensures that information receives the proper level of security and data protection.

What Data Should Be Archived?

What data should be retained in archives? That answer depends on the enterprise, its governing regulations, and its internal policies. Archives can contain many different types of data. Data can be classified as one of three types – *unstructured data*, such as files and documents, *semi-structured data*, such as email messages, and *structured data* stored in databases.

Advancements in Archiving Technology

There are many different archiving solutions on the market today. Many of these solutions have similar features. Data is indexed as it is ingested into the archiving system, allowing administrators and managers to search on various metadata, such as *date created*, *date accessed*, or *file type*.

Metadata that is available from many applications is limited and cannot provide the level of searching required by many enterprises today. Business line managers often must decide what information must be added to the indices without knowing what indexes will be relevant many months or years into the future. This creates a false sense of security that the information will be easily accessed when it is needed. Archival solutions provide enterprises with comprehensive search capacities that ensure that information can be retrieved now and in the future.

Email Archiving Solutions

Email is no longer only a communication tool. It is integrated into many business processes today and contains critical information that must be preserved. Email messages are often the subject of regulations and must be archived so they are safely protected, accessible and disposed of according to business and regulatory policies.

Email archiving solutions are important to satisfy requests to produce emails during the discovery process but there are many other benefits as well. For example, in the production environment, the primary

email storage can be reduced in size, making backups and restores run quicker. Email administrators, as well as the users, are not burdened with managing email quotas. Individual mailboxes can continue to grow in size; the email archiving solution will move emails from the production application to the archive based on policies, such as *date received* or *size*.

Email archiving systems are available from several vendors. Some of these systems only support *Microsoft Exchange*, while others also support other email environments, such as *Lotus Domino*. Some solutions “detach” attachments from emails and store them separately. These implementations can save storage space when the same attachment is sent to multiple users. Some systems can index data within attachments; others cannot.

Database Archiving Solutions

Databases, unlike email messages, have a very structured format. Databases grow over time. Take the example of a customer order database for an office supply company. New orders are entered in the database every day. These sales orders age over time and their access requirements diminish. In time, these older orders meet the business requirements of a “closed” transaction. Depending on the type of data, the format can be very complex. For example, 30 or more tables may be involved in capturing data for a single transaction.

The database gets larger and larger over time. It takes longer to back up and it takes longer to restore. Tables and indexes grow larger and performance suffers. Many methods for backing up, purging and restoring information are designed to increase performance, but end up adding complexity and increasing the cost of accessing valuable business information. Software vendors have developed database-archiving products that allow database administrators to weed out inactive transactions and move these transactions (along with its associated indexes and table structures) to lower cost storage.

The main production database with its active order set has now been “trimmed down”. Backups and restores on the production database run faster. Performance improves and stays consistent over time (this makes it easier to meet service levels). The secondary active archive database is updated infrequently and only needs to be backed up when the contents change. Database archiving products improve performance while simplifying administration and management.

File Archiving Solutions

Many archiving solutions focus on managing email messages. However, there are many files that must also be archived for compliance or business reasons. These files can take different forms, such as *Microsoft Office* documents and engineering designs. Nevertheless, these files have the same requirements as email messages. They must be stored for a specific period of time and they must be easily searched.

Moving inactive files to less expensive storage

saves storage costs while improving the backup and recovery times for the files remaining on more expensive disk systems.

Archiving Solutions Today

There are numerous solutions on the market today to archive data. Choosing the right one can be a difficult task. Understanding the architecture of the traditional archiving solution can help you not only choose a solution that will solve today's problem but also continue to support the growth of tomorrow.

There are several components within a traditional archival system.

- *Applications server* (such as an email server);
- *Archiving middleware software* (running on its own server), which receives the data or messages;
- *Database and search engine servers, software, and storage*; and
- *Content Addressable Storage (CAS) software, server, and storage*.

Data flows through each component before it is stored on the CAS storage. The middleware software receives the data from the application server where the process of categorizing the data begins. This middleware platform needs sufficient storage to hold the data temporarily until it is turned over to the database or search engine. The database server receives the data and indexes it. The results of the index operation are stored allowing for future searches. The data is then finally stored on CAS storage.

The applications server, middleware software, database and search engines, and content addressable storage can reside on different platforms, using different operating systems, supporting different protocols. Some solutions require IT to purchase the components from different vendors and then piece together the solution.

This patchwork of different products can create interoperability, security, and performance problems today and in the future. First, extensive testing is required to ensure that all products work well together. Later, maintenance on one system requires extensive testing to be performed again to ensure that the maintenance does not affect the performance or integrity of the entire solution.

Many archival solutions protect data from being modified when it is finally stored. Yet some of these solutions do not prevent data from being modified as it is processed by the various components. For example, data could be modified when temporarily stored by the middleware software. Database tables can be deleted which prevents data from being located.

While security must be an important concern, performance should also be an important criterion to ensure that today's solution can support future growth. Each component can be a potential bottleneck.

What is missing from these point solutions is an integrated architecture that can provide security while

ensuring performance in one *unified architecture*.

An Integrated Archival Architecture

There are integrated archival solutions on the market today that provide a more unified approach. The components are integrated under one "cover" that provides security as soon as the data enters the system. In addition, these solutions incrementally can add processing power and/or storage, allowing the systems to scale to support future growth.

Enterprises that have a large amount of data to be archived or expect the amount of archival data to increase over time should evaluate integrated solutions over "patched-together" point solutions. These solutions will not only scale to meet future performance needs but save money as well. Point solutions may initially cost less, but as additional servers, storage and software licenses are added, the cost for the solution may greatly exceed that of an integrated system and typically experience problems in scaling capacity and performance.

Integrated solutions make management and administration easier. Solutions that can archive all types of data – unstructured, semi-structured, and databases require auditors, for example, to search only one repository and not several. Retention and security policies are managed within one framework. Since these systems may reside in the main data center, they must be replicated to another location to ensure that a disaster in the main data center does not prevent access to the data. One unified platform makes it easy to replicate all archived data to a second location. (See the exhibit on the next page to help you choose the right archival solution.)

HP Offers Many Solutions to Solve the Archival Problem

HP understands that archival data can take many different forms. They also understand that the tremendous data growth experienced by enterprises today requires an archival platform that can provide consistent high levels of performance as the amount of data continues to grow. HP's answer to this archival problem is to offer solutions tailored to enterprise business needs.

HP StorageWorks Reference Information Storage System (RISS)

RISS consists of many *smart storage cells* and that is the secret ingredient that gives RISS its very scalable and flexible architecture. Each smart cell has its own dedicated processor, storage, and indexing intelligence. Smart cells are federated to form a "grid-like" storage fabric. Smart cells are grouped into domains that can service an application. *Need to service more applications?* Just define another domain. *Need more storage?* Just add more smart cells. Adding storage adds more processing power and indexing capabilities at the same time, which is why RISS is able to ingest, index, and search large amounts of data.

Choosing the Right Archival Solution

Choosing the right archival solution can be a difficult task. The following questions may help to narrow down the choices.

Hardware

- *How many servers and what types of servers are required?*
- *Are WORM (Write Once Read Many) devices supported? WORM support is critical for data that must be stored in an unmodified format.*
- *Can additional storage be added nondisruptively?*

Operations and Support

- *What support is available from the vendor? If this is a multi-vendor solution, will one vendor take ownership for problem determination? Who gets the first call when problems occur?*

Supported Environments

- *Which data types are supported? Unstructured? Semi-structured? Structured?*

Architecture

- *How is the solution implemented? Will the solution consist of various point solutions? If so, have the various vendors certified each other's solutions?*

Performance, Scalability, and Availability

- *How many messages/files/transactions can be stored per hour? How many files/messages/transactions can it store? Thousands? A hundred thousand? Millions? Billions?*
- *Indices are created when data is stored to provide quick access. Where are the indexes stored? How much storage is required to store them? If the index is corrupted, can it be rebuilt?*
- *How quickly can the software search through 100,000 messages/files/transactions to find all of the information about "Bob Jones"? When the amount of information increases from 100,000 to ten million or 100 million, is search performance degraded? How much storage can be supported?*
- *Does the solution detect and remove multiple copies of the same file or message? This is usually called "single instance store". Does the solution detect and remove multiple copies of the same attachments? Storing only one instance of the same message or attachment can significantly reduce storage requirements. Are other data reduction techniques, commonly called data deduplication, used?*

Management and Administration

- *How is the system managed? Can it be managed remotely? Can administrators be assigned different levels of access?*

Discovery

- *Can messages/files/transactions that have been archived easily be searched and then quarantined so that they are available for discovery? Is the control of that information only accessible to the compliance officer, legal department, etc.?*

Security

- *How are administrators authenticated into the system?*
- *Can administrators be granted different levels of access to allow or restrict access to data?*
- *How are the indexes protected? Can they be modified/deleted by disgruntled employees?*

Policies

- *How granular are the policies? Can different policies be set for different departments or for different groups of workers?*
- *Can retention periods be extended when necessary? This is important when a discovery or audit process is taking place and there is a danger that data required for the discovery or audit process may expire before the process is completed.*

Replication to a Remote Location

- *Does the solution have policies to route data both locally and to a remote location? If not, can existing replication products be used to replicate the contents of the archive to a secure, remote location? For disaster-recovery purposes, archival storage must be protected at a remote location.*

RISS is not just an email archiving solution, or a database archiving solution, or a file archiving solution. It supports all three. RISS integrates with applications through its API framework. That allows RISS to support not only HP applications but those from third party software vendors as well. Since RISS can support structured, semi-structured, and unstructured data, enterprises have one platform to manage all of their archives. That simplifies operations and management. Moreover, since smart cells can be added or removed from domains as needed, RISS can satisfy constantly changing business and regulation requirements.

Data is secure as soon as it enters RISS. RISS creates a unique 160 bit content ID using the SHA-1 hashing algorithm for each object. SHA-1 is a cryptographic hashing algorithm designed by the National Security Agency (NSA) and published by the National Institute of Standards and Technology (NIST) as a United States Government standard.

Users can be assigned different levels of access for each domain. Some administrators may be granted access to view a domain containing archived database data, for example, but cannot view any email domains. Corporate attorneys may be granted access to search across all domains to find all emails and files that are required during the discovery process.

Each object within RISS is mirrored to another smart cell for data protection. One RISS system can replicate all or selective objects to another RISS system for disaster recovery. In fact, the two RISS systems can operate as an active/active pair. Each RISS can service local requests, while serving as the backup for the other RISS. In addition, RISS can be backed up to other media, such as WORM tape, using traditional backup software.

HP RISS supports a wide range of application partners and is tightly integrated with HP's own data capture and migration applications. The first HP application is *Reference Information Manager (RIM) for Messaging*, which supports Exchange and Domino environments.

RIM for Messaging

RIM for Messaging mines emails from individual mailboxes and sends those messages to RISS. In a compliance environment, it can capture all the journaled emails in the mail system. To enhance e-mail management, RISS can be configured to selectively archiving email from specific mailboxes, and personal folders (PSTs) using rules defined by management based on age, size, sender, recipient, or specific keyword. This operation leaves behind a pointer in the mail server to redirect access to the archived files to RISS. This dramatically reduces the load on the email servers to manage historical emails and maintains the same interface to the users as they interact with email. In fact, with RISS managing older emails, access to older emails can be faster than from an email server.

RIM for Messaging provides enterprises with an email compliance solution, but it also saves storage and

administration costs. Emails no longer must reside on expensive primary storage but can be moved to lower cost RISS storage. Since emails are mined from Exchange and Domino servers, administrators do not have to enforce quotas on mailboxes – removing a constant source of frustration between email administrators and those that cannot keep their inboxes under the pre-set limits. Now that the primary email storage is pared down in size, backups and restores of the primary email environment can be dramatically reduced as well as increasing performance.

Email attachments are stored separately from the message. If an attachment is sent to ten different recipients, only one copy of the attachment is saved. RISS also uses block-level data reduction software to eliminate redundant data streams within and across objects. These two techniques reduce storage requirements and the total cost of ownership while providing a simplified way to search all emails and attachments.

RIM for Databases

Reference Information Manager for Databases is HP's suite of products for managing and archiving databases. RIM for Databases consists of several components. These include *RIM for Databases Relocater* and *RIM for Databases Encapsulated Archive*.

RIM for Databases Relocater manages accelerating data growth by relocating closed transactions and other infrequently accessed data to an easily accessible online archive database. This is generally data that is two-to-seven years old. This significantly reduces the size of the production system and improves application performance, enabling you to do the following.

- **Identify** inactive data based on business rules, activity history and retention policies
- **Relocate** inactive business transactions to an online archive database on less expensive storage
- **Retain application transparency** providing online, real-time user access to a combined view of production and archived data through native application and reporting interfaces

RIM for Databases Relocater can boost application performance by as much as 70 percent. It improves application availability and stability and cuts IT infrastructure costs. What's more, it retains transparent, online real-time access to both the production and archive data.

RIM for Databases Encapsulated Archive stores inactive data (approximately seven to twenty five years old) in an archival store, where it can be moved to increase performance, decrease costs and be managed for compliance.

Supported databases are Oracle and Sybase. Based on the development kit, any homegrown or packaged application can be integrated into RIM for Databases. HP offers pre-packaged integration for the *Oracle e-Business Suite* and *PeopleSoft ERP*.

For example, the financials database contains all customer orders from seven to twenty-five years ago.

RIM for Database Encapsulated Archive will encapsulate those older customer records in an XML format and store these records on RISS. These transactions can now be searched through an XML query.

Since RIM for Database Encapsulated Archive stores records in an XML format, enterprises no longer have to be concerned about keeping older versions of databases (or applications) available to retrieve old records. These archived transactions are now independent of application, database, or operating systems software. Administrators can move records from RIM for Database Relocater to RIM for Database Encapsulated Archive or can move records directly from the primary database to RIM for Database Encapsulated Archive.

RIM for Databases Analyzer is included in the platform. It is an enterprise dashboard that intelligently monitors systems, and forecasts and models data growth. It helps create and enforce data retention policies and can model the impact of data retention policies without impacting production data.

HP StorageWorks File Migration Agent Software

HP StorageWorks File Migration Agent (FMA) software moves inactive files residing in Windows environments to lower cost storage, including the RISS platform. Files moved to RISS are content searchable and the life cycle of these files can be managed through retention policies.

HP Archiving Solutions for Healthcare

HIPAA regulations have many hospitals and medical institutions evaluating their current methods of storing medical images. In many of these institutions, each department has purchased and implemented its own imaging systems. The result is that there are numerous silos or standalone collections of images throughout the hospital. These images cannot be easily shared and each department is responsible for backing up the data locally and providing disaster recovery copies. Centralizing all medical images within one repository can reduce cost and reduce management complexity while providing higher levels of data protection. HP has developed a solution called *HP Medical Archiving Solution (MAS)* to address the problems of managing isolated imaging systems. This solution is designed to work with existing PACS and RIS systems.

HP Medical Archiving Solution (MAS)

HP has designed a grid-based medical system based on HP ProLiant servers and HP StorageWorks disk arrays. These integrated solutions are available in three different capacities – 5 TBs, 10 TBs, and 20 TBs. MAS architecture is modular to create configurations that are highly resilient, scalable, and configurable to meet hospital needs. HP MAS nodes communicate using TCP/IP networking. Medical images are ingested through the Gateway, which presents the grid as a network drive to the workstations and client servers.

These images are indexed by the Control Node, which maintains the metadata and policies for each group of images. Images are sent to the Storage Node where they are stored on the disk arrays. One Storage Node can replicate the images to a second Storage Node in a remote location. The Admin Node provides centralized management and administration and maintains audit information for compliance.

With the centralized archive that MAS provides, doctors in remote clinics can share images with doctors at the main hospital. All data is protected, replicated to remote sites for disaster recovery, and encrypted, if needed. When more capacity is required, additional storage can be added to the grid.

If You Need Help

Centralizing your archival data in one repository, through RISS or MAS, requires planning to ensure that the right data is archived with the right policies. For companies that need assistance, HP offers professional services to help IT and business units plan and successfully implement centralized archival solutions.

Conclusion

Regulations, such as SEC 17a-4 and HIPAA, require financial and health care institutions to retain data for a specified period of time. Enterprises in regulated industries must comply with regulations or face hefty penalties and negative publicity. Enterprises that are not regulated must also gain control over their data that should be archived. Enterprises, both large and small, may be required to produce documentation to fulfill legal discovery.

If you do not have an archival solution in place, then you need to evaluate and implement a solution now. Enterprises must choose archival solutions with a critical eye to the future. Remember that the solution chosen today will be retaining information for many years and must be able to accommodate the continuing growth of data.

HP has developed a suite of products that deliver the performance and scalability required by enterprises with high data growth. Put HP on your short list of vendors to evaluate when investigating archival solutions.



About The Clipper Group, Inc.

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

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

About the Author

Dianne McAdam is Director of Enterprise Information Assurance for the Clipper Group. She brings over three decades of experience as a data center director, educator, technical programmer, systems engineer, and manager for industry-leading vendors. Dianne has held the position of senior analyst at Data Mobility Group and at Illuminata. Before that, she was a technical presentation specialist at EMC's Executive Briefing Center. At Hitachi Data Systems, she served as performance and capacity planning systems engineer and as a systems engineering manager. She also worked at StorageTek as a virtual tape and disk specialist; at Sun Microsystems, as an enterprise storage specialist; and at several large corporations as technical services directors. Dianne earned a Bachelor's and Master's degree in mathematics from Hofstra University in New York.

- ***Reach Dianne McAdam via e-mail at dianne.mcadam@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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

Disclosure

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

Regarding the Information in this Issue

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