THE CLIPPER GROUP
Explorer™

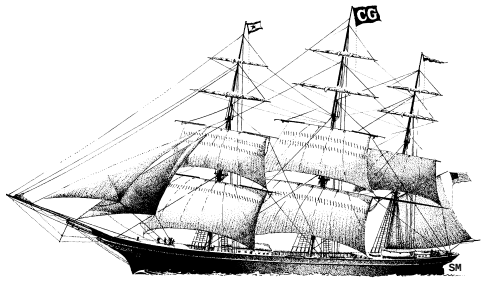

Navigating Information Technology Horizons℠

Published Since 1996 Report #TCG2003054 October 16, 2003

Intelligent Storage Networking — Poised for Broad Adoption

Analyst: Michael Fisch

Management Summary

Why would you use an automobile instead of a horse? A word processor instead of a typewriter? A calculator instead of a slide rule? The answers to these questions seem obvious today, but there was a time when the answers were not obvious – even controversial. The newfangled, unproven devices presented a challenge to the way people had traditionally done things. **They required people to change, and there was legitimate debate about whether it was worth it.** But after a period of trial and uncertainty, the superior technologies slipped into the mainstream and eventually became the norm.

Enterprise storage is in a similar period of uncertainty about the issue of intelligent storage networking. **While networked storage (i.e., SAN and NAS) has become the norm, especially in larger IT environments, the next evolutionary step calls for moving some storage intelligence from servers and storage arrays into the *network*.** This new approach is gaining momentum, though still in the early stages. To be sure, just because it is new does not mean it will carry the day – there must be clear and compelling reasons to adopt.

A close look reveals that such reasons do exist. Though the defining characteristics of intelligent storage networking tend to be technical, the main reason for adopting it is not. Storage is essentially about meeting the service level requirements of the business at the lowest possible cost. Enterprises need information to run like automobiles need gas, and this information must be appropriately stored, protected, and made accessible. Not all data is of equal value, and not all users and applications require the same level of performance, availability, and recoverability. The job of IT providers is to make sure the storage infrastructure delivers the right service levels to the right constituents at the lowest possible cost. This is a tricky balancing act – better service costs more, but insufficient service levels can slow or stop the applications and business processes they support. Both extremes are costly. **The beauty of intelligent storage networking is that it can deliver all required service levels more easily and cost-effectively.** It reconfigures the entire storage-value equation, which is why there is so much excitement about it. It is a technical architecture whose impact is ultimately economic and bottom-line oriented.

The *how* and *why* of intelligent storage networking are worth examining, especially now as the ferment of broad adoption is likely close. Read on for a closer look.

IN THIS ISSUE

- State of Storage Networking 2
- Curing What Ails You 2
- On the Road to the Future 4
- Risks, Costs, and Implications 6
- Conclusion 7

State of Storage Networking

Networked storage has grown quickly over the last five years. Beginning as an alternative to direct-attached storage (DAS), it promised to break open the exclusive server-storage relationship and pave the way for broad storage consolidation on a network. Both SAN[1] and NAS[2] lowered storage total cost of storage ownership (TCO) and gave enterprises leverage to keep up with the increasing data requirements of the Information Age. These are now mainstream technologies, and networked storage is the predominant enterprise storage architecture in the market. However, its value has only begun to be tapped. There is room for more innovation and increased efficiency, and intelligent storage networking presents itself as a logical next step.

While today's storage networks can quickly, reliably, and securely move data from point A to point B, this new approach calls for going a step further by placing features like *virtualization*, *replication*, *data migration*, *provisioning*, and *data management* in the network's fabric. If a data center has these software capabilities today, they usually run on host servers or in storage arrays, and each location has its pro's and con's. **Intelligent storage networking would disaggregate these features from the end points and place them in the network – the common ground that joins servers and storage.** This would mean that the software runs, instead, on an intelligent switch, appliance, or platform in the data path.[3] These valuable features could then apply to a much broader set of servers and storage systems, even from multiple vendors.

There are well over a dozen startups offering intelligent storage networking hardware and/or software. The market leaders have also jumped in the water, and most now offer such products or will in the near future. Vendors are clearly moving in this direction. *Should enterprise IT departments follow suit?*

The answer depends on whether there is a compelling reason to adopt. With your enterprise's particular requirements in mind, you must weigh the benefits of intelligent storage networking versus the risks and costs of adoption, and then see which way the scale tips. To help with the analysis, several factors are considered:

- **Does intelligent storage networking address the pains that IT departments feel today?** The technology should remedy shortcomings in real-world enterprise storage environments.
- **Is intelligent storage networking consistent with the overall future direction of computing?** This is a forward-looking and conceptual question, but it matters because storage is evolving quickly. The fast pace of information growth and technology innovation guarantees that the infrastructure on your data center floor will be quite different in three or five years. So, it is reasonable to ask whether this new technology positions your enterprise for a smooth transition going forward, or takes it on a technological tangent that may make it harder to change with the times.
- **What are the risks, costs, and implications of making the change?** The answer to this question will depend on your IT situation as well as the particular product(s) and vendor(s) in consideration. Risks and costs are always involved in change, but there should be a sense of proportion that clearly tips the scale in favor of the benefits.

Before proceeding, keep in mind that intelligent storage networking is an architecture, not a product. Though this report refers to specific features, the purposes and capabilities of individual products will vary. This is intended to be a broad overview.

Curing What Ails You

An important, pragmatic test of this technology is whether it solves real-world problems. That is, does intelligent storage networking cure what ails you today? Like the old *Whack-A-Mole* game in the arcades, storage costs and complexities pop up in many places, such as low resource utilization, time-consuming management, insufficient data protection, and so on. The test is how well it can whack them back into their hole and out of sight.

Poor Resource Utilization

Poor utilization of storage resources is a widespread problem in enterprise data centers, and it's a waste of money. Unless an IT department has implemented specific tools and processes to boost utilization, the majority of storage capacity is typically unused. Lacking appropriate tools, administrators are often unaware of the extent of this problem. The result is a lot of unnecessary storage purchases and their associated environmental and operating costs.

Administrators can have much greater

[1] Storage area networks (SANs) provide access to block-level storage typically over a Fibre Channel network.

[2] Network-attached storage (NAS) provides file system access typically over an IP network.

[3] Most intelligent storage networking devices operate at the block level, which implies a SAN, so this report focuses primarily on that aspect. There are such platforms for NAS devices, however.

visibility and control over storage capacity if it is virtualized. Some may think the virtualization term is overused, but it best describes the important ability to pool resources through logical abstraction. This allows storage to be managed as a single, large entity, rather than numerous, disconnected, discrete ones. It masks the physical complexities of underlying resources and makes them more flexible and easier to manage. Think of dialing a phone number, which is a logical representation of someone connected to a telecommunications network. You don't have to know how to route the call through the labyrinth of switches, hops, and points of presence to reach the intended person – *that's virtualization*![4]

The storage virtualization function can reside anywhere in the SAN data path, though some places are arguably better than others. A storage array can virtualize and partition its own capacity into LUNs[5] and volumes of various sizes and protection schemes (e.g., RAID 1, 5, 10), but the scope is limited to that particular array and the servers connected to it. A host server with a volume manager can virtualize LUNs connected to it, even potentially from multiple storage arrays. However, the scope of virtualization is limited to that server and the LUNs that have been zoned and masked[6] to it. Practically speaking, it only sees a small portion of total storage. Virtualization based on storage arrays or servers is like scanning the horizon from the foot of a mountain – you can see, but the view is limited. However, **virtualization in the network can encompass all servers and arrays.** Like the guru sitting on top of a mountain, it sees all.

As a result, it can aggregate distributed storage systems into a single, more manageable, and better-utilized entity. In contrast, non-virtualized DAS or SAN-attached storage becomes fragmented into multiple islands of under-utilized capacity. It cannot be shared, and each island must carry its own overhead capacity. However, the single view of network-based virtualization helps liberate isolated capacity through sharing and minimize wasted overhead. It also eliminates the associated operating and environmental (i.e., power, cooling, and floor space) costs of wasted, unused capacity. Another useful feature is real-time visibility into capacity usage and server/application mappings. This information is useful for more precise storage management. If virtualization support includes legacy equipment, enterprises can also get more mileage out of older storage systems[7] by connecting them to the virtualized pool instead of relegating them to early retirement. **Through virtualization in the network, enterprises do much more with what they have.**

Time-Consuming Management

Ongoing management costs, especially skilled labor, is by far the largest component of TCO over the life of storage. This includes activities from maintenance to adding or reallocating capacity to coordinating backups and restores. In the face of rapid and continuous data growth, anything an enterprise can do to improve the productivity of their storage administrators is good – not only for the well-being of employees, but also for the bottom line.

Moving intelligence into the network maximizes the scope and applicability of software features, especially those that require real-time processing in the data path. Instead of deploying and managing multiple, possibly different software suites for multiple server operating systems and/or storage models, why not have just one super-instance in the network that applies to everything? It does not matter whether you are talking about virtualization, replication, data migration, mirroring, provisioning, backup and restore, or anything else, **a centralized function with a single point of management that encompasses diverse enterprise infrastructure will save management time and costs.** Less is more, in this case.

Application Downtime

Application downtime is always undesirable, and the degree of impact depends on the nature and scale of the supported business process. If the corporate e-mail server for a small business stalls for an afternoon, it can annoy workers and negatively impact productivity. If the transaction processing application for a major stockbroker fails for even a short period, the losses in terms of revenue and customer ill will would be tremendous. So whether it means an inconvenience or a catastrophic business loss, downtime should be avoided.

Placing intelligence in the network can help mitigate storage downtime. There are several causes of application downtime: equipment fail-

---

[4] See *Storage Virtualization in 2001: A Space Odyssey* in **The Clipper Group Explorer** dated April 9, 2001, at http://www.clipper.com/research/TCG2001002.pdf.

[5] Logical units of storage, a singularly-defined collection or space.

[6] Zoning and LUN masking are security features for partitioning a SAN and securely mapping a LUN to a particular server, respectively. They prevent the dangerous situation of one server accidentally overwriting another's data, and they lessen storage traffic.

[7] Assuming support costs do not become prohibitive.

ure, software bugs, data corruption/loss, operator error, and disruptive management tasks like firmware upgrades, capacity expansion, and backup. These causes can occur at any point in the data path, from the application server to the disk drive. However, by inserting a layer of abstraction between servers and storage, administrative tasks like provisioning, volume expansion, backup, and data migration can be non-disruptive to the application. This avoids *planned downtime*. Furthermore, by facilitating data replication through point-in-time copies and mirrors, it can speed recovery from data corruption or equipment failure (i.e., *unplanned downtime*). The additional flexibility of network-based intelligence accelerates the repair and recovery process, in general, such as by quickly allocating a volume to a replacement server after a failure. While these availability features can reside in servers or storage arrays, the network may be a more flexible and economical alternative, especially in larger IT environments.

Inadequate Data Protection

The issue of data protection, or ensuring data is there when you need it, is worth exploring in more detail. Data protection is a function of the *three Rs of resilience*[8]:

- **Redundancy** – Multiple copies to prevent data loss,
- **Remoteness** – Remote copies to survive local disasters, and
- **Recoverability** – Fast recovery from an interruption.

By making it easier and less expensive to perform all three Rs, **intelligent storage networking allows an enterprise to do more (and perhaps what it should have done all along) to ensure business continuity and guard against major data loss.** For instance, less-costly storage systems can be used as targets for point-in-time copies and synchronous/asynchronous mirrors, instead of being locked-in to using the same model or vendor, as in most array-based solutions. The smart use of periodic copies, real-time mirrors, and tape backup can further the three R's of resilience. Again, a centralized approach also removes barriers of complexity for establishing and maintaining these capabilities across heterogeneous servers and storage systems.

Data Lifecycle Management

Data (or information) lifecycle management is on track to become a major storage trend in itself. The idea behind it is to take a cradle-to-grave view of information, from creation to archival to deletion, and move data to the optimal price/performance storage tier throughout its lifecycle. It takes advantage of the fact that the value of data changes based on a number of factors including time, application importance, and special events that may impact the value of data. Its objective is to more precisely optimize storage service levels and minimize costs.

Data lifecycle management presumes the existence of tiered storage classes, or multiple storage tiers with different price/performance characteristics.[9] It also presumes the ability to copy and move data non-disruptively among the tiers. Intelligent storage networking can facilitate both by (1) making it easier to deploy and manage multiple tiers of storage and (2) facilitating data replication and migration across the tiers. **The functional centralization and virtualization it delivers can greatly simplify data lifecycle management.**

Feeling Better

Placing intelligence in the network is not a cure-all, but it does address a number of the critical problems that enterprises face. It plays a role in improving the economics of storage by lowering acquisition, management, and environmental cost. It also qualitatively improves storage service levels in areas like availability and recoverability. **So, intelligent storage networking passes the most important test of solving real-life shortcomings in today's storage systems.**

On the Road to the Future

Another important test considers whether this technology lies on the road to the future. Even among rival IT vendors, there is remarkable agreement about what the future of computing generally looks like. This shining city on a hill is sometimes called *utility computing*; others refer to it as adaptive infrastructure, on-demand computing, etc.[10]; but the concept is fundamentally the same: IT infrastructure will be like a public water or electricity utility. **Customers "turn it on" and receive computing services in an on-demand,**

[8] See *Business Continuity Goes Better With SANs - The 3 Rs of Resilience* in **The Clipper Group Explorer** dated January 25, 2002, at http://www.clipper.com/research/TCG2002003.pdf.

[9] See *Tiered Storage Classes Save Money – Getting The Most Out Of Your Storage Infrastructure* in **The Clipper Group Explorer** dated August 29, 2002, at http://www.clipper.com/research/tcg2002030.pdf.

[10] See *In Search of - Utility Computing* in **The Clipper Group Explorer** dated June 2, 2003, at http://www.clipper.com/research/tcg2002024.pdf

pay-per-unit[11] manner. They just specify the resource (computing, storage, networking) and the quality of service (performance, reliability, and cost). In turn, the utility automatically and transparently delivers and accounts for the needed resources. It flows like water from a faucet!

This may sound futuristic, but the industry is converging on the utility concept faster than many realize. It represents the best approach for delivering:

- The right resource to the right constituent at the right time,
- Efficient resource utilization,
- Simple management,
- Adaptability to meet changing requirements,
- Accountability for resource consumption and delivery, and
- Low total cost of ownership.

Simply put, **utility computing meets business requirements much more effectively than the hard-to-manage point systems and silos of computing that dot the landscape of enterprises today.** This emerging paradigm is generally associated with the attributes below. The question is whether intelligent storage networking can help deliver these utility attributes more fully, and therefore act as a stepping-stone to the future.

Virtualized

Virtualization is an important utility characteristic because it allows computing resources to be pooled and more easily managed. As previously discussed, placing this function in the network turns the storage infrastructure into a single, more manageable entity, which helps enable the utility.

Dynamic and adaptable

By all measures, we live in turbulent times, and businesses must be adaptable to survive. It follows that the IT infrastructure that supports business processes must also be adaptable.[12] Scaling, partitioning, and allocating resources should be automatic, fast, and non-disruptive. Reconfiguring systems should be like automatically shifting gears in a car while traveling down the road, not like taking the vehicle to a repair shop to rebuild the engine. The infrastructure should be flexible enough to meet multiple, unique service levels based on the needs of individual applications and users. It should not assume that "one size fits all." **Infrastructure should be fluid, not fixed.**

The broad scope of network-resident software does deliver more flexibility. To establish a remote mirror, schedule point-in-time copies, provision a volume, etc., the administrator does not have to do it differently for different servers or storage systems. There is a unified management approach, so adapting and reconfiguring the infrastructure is simpler. It also speeds up tasks like adding more servers or storage or deploying a new application, which has a positive impact on a business's ability to scale or reconfigure itself. **Network-based intelligence makes the whole storage infrastructure more dynamic and contributes to the adaptability of IT and, therefore, to the entire business.**

Heterogeneous and Open

"Heterogeneous and open" means multiple vendors and products can participate and interoperate in an IT infrastructure. It is the opposite of proprietary systems and their result – vendor lock-in and higher costs. Just as free markets encourage competition and benefit consumers, open IT encourages competition and innovation among vendors that benefit enterprise customers. It delivers "power to the people," if you will, and gives greater choice and flexibility in components, systems, and vendors.

By making storage functions independent of storage vendor/model and server operating system, intelligent storage networking adds openness to a SAN.[13] For instance, functions like replication that reside in a storage array are generally only compatible with storage arrays of the same make and model. These features running on host servers would also be limited to particular operating system(s). Both options can constrain procurement flexibility, but if these functions reside in the network, one is free to mix and match servers and storage. **Network-resident software provides greater openness and heterogeneity than the alternatives**, making the storage infrastructure more utility-like.

Unified and centrally managed

This is absolutely essential to the utility concept, though possibly the most difficult to accomplish. A SAN infrastructure is an interconnected web of devices that must work

[11] See *The Accounting Pendulum Swings at Storage (Or, Why the Taxman Cometh?)* in **The Clipper Group Explorer** dated October 31, 2002, at http://www.clipper.com/research/TCG2002043.pdf.

[12] See *Times They Are a-Changin' – Adaptability Is Supreme* in **The Clipper Group Explorer** dated June 6, 2003, at http://www.clipper.com/research/TCG2003028.pdf.

[13] The openness of a particular solution depends on the variety of components with which it can interoperate.

individually and with each other to deliver the final output – storage as a service. Like an orchestra conductor, administrators have the time-consuming job of managing and coordinating all of them. This becomes increasingly difficult as the infrastructure grows and the number and diversity of devices increase. Heterogeneity and scale actually make matters worse, *unless* they are accompanied by unification and centralized management. This is the glue that ties together the computing utility. Historically, IT has been centrally managed (e.g., mainframe) or heterogeneous (e.g., open, distributed systems), but not both simultaneously. *That's the challenge!*

A single super-instance of software in the network and single view into storage definitely contributes to unification and centralized management. The alternative, as mentioned above, is to manage many different software suites running on various servers and storage arrays. The single view of storage and its usage also contributes to accountability for users' consumption of resources and facilitates measures like charge-back. So, intelligent storage networking helps deliver the critical utility characteristic of unification and centralization.

Automated

Automation is the frosting on the cake of utility computing. It builds on the ability to monitor measure IT activities by taking action automatically based on rules and policies, thus minimizing the need for human intervention. Its usefulness is directly related to the features and functions that it can automate – that is, it magnifies the features in place. If an IT infrastructure substantially offers the preceding characteristics, then automation can push it over the line to full utility computing. Network-resident functionality can be included in the automation.

A Milestone on the Road

For decades, information technology has continuously become more powerful, efficient, manageable, and cost-effective. A computer that began as a room full of vacuum tubes now fits on a sliver of silicon in a wristwatch. The nature of the technology, the intensity of industry competition, and the dynamism of free markets give us every reason to believe this trend will continue. If you trace this progressive path to its logical conclusion, you will find utility computing at the end of the line. It is simply a more efficient and cost-effective way to do things. (This is not to say that the utility won't keep improving once we arrive. It will – nothing is static!)

If you hold up intelligent storage networking against the future backdrop of utility computing, you will find that it fits – that it makes sense as a milestone on the road toward this better future. It improves on current architectures and satisfies important utility characteristics. This is important because **intelligent storage networking represents a significant architectural change, and no enterprise wants to invest in something like that might eventually lead it on a divergent path, away from the mainstream of computing evolution.** On the contrary, it appears to lie right in the middle of the road.

Risks, Costs, and Implications

Intelligent storage networking stands out in the light of both present needs and future trends, but it is new and involves change. What are the risks, costs, and implications of moving to this architecture?

First, it is worth clarifying that this is not about buying a technology or architecture for its own sake – for all the nice bells and whistles. **If you clear away all but the essentials, the overriding objective of a storage infrastructure is to meet the service level requirements of the business while minimizing TCO.** On one hand, nobody wants to spend more than they have to on storage. Executives certainly have other designs for scarce cash resources. On the other, storage is not an area one can cheat or neglect because information flow is mission-critical to a modern enterprise. **It is a matter of finding the right balance between service levels and cost.**

Therefore, enterprises will probably not consider intelligent storage networking for its own sake, but as a means to improve storage service levels and/or lower costs. They will have specific storage needs and objectives in mind and may consider it as part of a potential solution. So, what are the implications of such a move?

Risks

Intelligent storage networking is on the early side of the adoption lifecycle, though not as early as you may think. Some products have been on the market for two or three years, with a good number of customer installations. It is also not the exclusive domain of exciting-but-uncertain start-ups, either. Most major storage vendors offer such a product or will in the near future. **So the technology is not coming out of left field, and the categorical risks of going down an obscure, tangential path are relatively small.** The adoption risk is more directly related to specific vendors and products, not the general architecture.

A second risk factor derives from the broad scope of the solution. Ideally, it encompasses the

entire storage infrastructure in order to deliver a powerful form of unification and centralization. The downside is that if something goes awry, the adverse effects would also be broad. Therefore, the platform ought to be highly available, as in redundant components, clustering, and non-disruptive upgrades. A gradual approach to deployment would also be reasonable. Run it on a subsection of the IT environment for a while, such as a SAN island, and then incrementally expand to full coverage as comfort levels rise – baby steps to intelligent storage networking!

Costs

The cost of a solution is incremental to a SAN deployment, and it depends on the particular product and requirement. If a SAN is already in place, then you could add a solution to it, with some reconfiguration, as long as the solution supports the existing operating systems, switches, and storage systems. (Check with the vendor.) If you are starting with direct-attached storage, then the cost to get to intelligent storage networking will include a SAN and the intelligent platforms.

Implications

Intelligent storage networking addresses problems associated with scale. Small IT environments, such as a single storage array plus one or few servers, would not experience the kind of shortcomings that it could address. **Thus, the most likely candidates for this technology are medium and large enterprise environments.**

So, which storage functions are most appropriate to place in the network? There are no hard and fast rules here. Since the solution resides in the data path, functions that require real-time I/O processing are promising candidates. The most compelling are virtualization/volume management, replication (i.e. point-in-time copy and mirroring), and data migration. Additionally, write caching, data lifecycle management, provisioning, backup and restore, and other storage and data management functions are good possibilities. RAID processing is likely to remain in the storage arrays, since it is so closely tied to the disk arrays.

Within this general category, there are also several different product architectures from which to choose, each with their respective advantages:

- **In-band, software-only:** This solution resides in the SAN (between host servers and storage) and runs on standard *Linux* or *Windows* servers. It can offer a broad, integrated suite of features at a competitive price.
- **In-band, purpose-built platform, integrated software:** This SAN-resident solution also has a broad, integrated feature set but runs on a specially-constructed hardware platform (including switches and directors) that delivers very high performance and availability – it is a high-end solution.
- **In-band, purpose-built platform, open software:** Similar to the previous solution, it has an open API that allows third parties to write software to run on it. The feature set depends on which software is available and what the customer wants.
- **Out-of-band, integrated software:** This solution uses host-based agents and a management/metadata server that resides outside of the data path. It offers an integrated feature set, as above, but it is technically a host-based solution.

Finally**, a good time to consider an intelligent storage networking solution is during a storage consolidation project, which may also include a SAN.** While re-architecting the infrastructure and buying new equipment, it is a convenient opportunity to consider it. You may be able to enhance the anticipated benefits of storage consolidation (simpler management, higher utilization, greater flexibility) with an intelligent solution.

Conclusion

This up-and-coming trend is poised for broad adoption. Intelligent storage networking solves real problems today by taking cost and complexity out of networked storage. It is also future-proof, by all appearances, and its relevance will increase, not diminish, as the macro trend of utility computing unfolds. On balance, the benefits of improved storage economics and service levels outweigh the costs and risks of embracing intelligent storage networking. Therefore, it is a question of when, not if, the market takes off.

If this technology is appropriate for your enterprise needs, you may want to adopt sooner rather than later. There are competitive advantages for early movers. Alternatively, you may prefer to play it conservatively, wait, and move with the pack. It is a business decision that you must make based on your choice of a particular risk/ reward profile. With a better understanding what these risks and rewards are, you are now better equipped to decide. **In any case, keep your eyes on intelligent storage networking - because it is coming!**

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

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

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

Regarding Trademarks and Service Marks

Disclosure

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

Regarding the Information in this Issue

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