Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE
FOR THE INFORMATION TECHNOLOGY BUYER

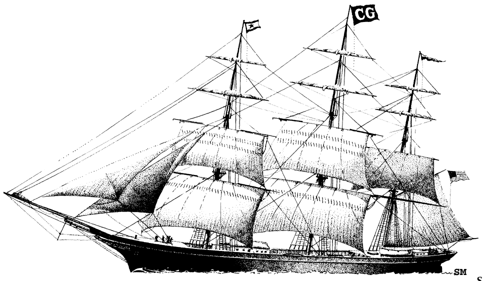

Navigating Information Technology Horizons℠

Published Since 2007 Report #TCG2011032 September 30, 2011

Disasters Remind Us Why We Need a High-Availability Data Center — *How Prepared are You?*

Analyst: David Reine

Management Summary

When we wake up in the morning, there are a set number of events, or processes, which we expect to occur. First, we shower (preferably with *hot water*) and then we get dressed, eat breakfast, and leave for work. When we get into our car, we have certain expectations. First, when we push the ignition button (or, turn the key), we expect the car to start – not sometimes, not usually, but *every time*! When we get to the end of the driveway, we step on the brake; we expect the car to stop: not sometimes, not usually, but *every time*! ***We expect these functions to work every time!*** If the car doesn't start, we are inconvenienced, but it is not a disaster. A failure to stop when the brake is engaged carries a much higher risk than a failure to start. If the car does not stop, *that* could become a disaster! In fact, most of these risks could be mitigated if we bring the car in for service on a regular schedule. I suspect that the cost of preventative maintenance is less than the cost of time/opportunities lost when a failure does occur. The question that you need to ask is: ***Can I afford not to improve the availability of my transportation? But, also, shouldn't I have a "Plan B" for getting to work, if anything (like an accident) takes my vehicle out of service?***

On a business basis, what would happen if there was a multi-day power failure affecting your data center, or worse, a natural disaster, say, a hurricane, affecting a wider geographic area? Is your enterprise equipped to provide an alternative solution quickly, especially for an enterprise data center with mission- and business-critical applications deployed on sophisticated servers with any number of hardware reliability, availability, and serviceability (RAS) features? Have you even identified if you have any single-points-of-failure (SPOF), or multiple points, that could cause your data center to lose the capability to respond to business and user needs? Many enterprise data centers have deployed more capable (i.e., more sophisticated and expensive) servers because they offer higher-availability in case of a component failure. Are you looking at everything that can cause a failure? Are you using sophisticated enough servers and are you using them properly? ***Can you afford not to?***

Usually, the more that you spend, the greater the protection and the higher the availability (often measured in minutes of unplanned down time per year). Many enterprises have installed a high-availability (HA) solution for server fail over, in the event of a power outage. Many of the server failover options, however, have a shared storage component, potentially a mission-critical SPOF. ***Are you prepared for the loss of your data center? Exactly how well prepared are you?***

Although high-availability and disaster-recovery options can be expensive, they may be far less costly than suffering lost business and/or damaging

IN THIS ISSUE

customer relations when a major outage occurs. Only you can determine how much pain your enterprise can tolerate. Operationally, this is measured in time duration (i.e., days, hours, or minutes without access to critical business processes and data). However, this usually can be translated into financial costs (like losses of revenue and impact on profit, costs of litigation and amelioration, etc.). ***Do you fully understand what a prolonged systems and/or data outage would cost your enterprise?***

Ultimately, this will come to a weighed discussion of how much your enterprise can afford to pay to reduce the potential losses that might result from any outage, let alone a prolonged one. The costs of prevention (or sizable reduction of the scope and duration probabilities) of a prolonged outage must be weighed against the costs that would be avoided (by not being down for a prolonged time). This is no different from considering a variety of insurance options to decide what makes the most sense for your personal situation.

Let's be clear here. While the probabilities of a disaster may be small, just like the very low probability that you won't survive your commute home, the odds of a power outage are significantly higher. Some are related to disasters, to be sure, while others are just major interruptions.

However small that probability is for a major interruption, the questions basically are very similar for your enterprise and family:

- ***How well prepared is your enterprise (or family) to continue without its critical systems and data (or you)?***
- ***How long will it take the enterprise (or your family) to recover?***
- ***Will there be serious consequences as a result?***

Once you understand the consequences (and costs, thereof), you are prepared for the other half of the analysis: ***What could and should you do to lessen the damages that the consequences will bring and how much does it make sense to spend to ameliorate the risks?***

None of this is easy, either personally – say, in terms of buying life insurance (i.e., there are so many options and personal considerations) – or in beefing up your enterprise's IT infrastructure (i.e., so many technical options and business considerations). This issue of Clipper Notes will help you define the risk factors that might cause you to deploy an integrated, high-availability solution for your data center, one that enables you to mitigate against those risks and protect the data that is so integral to a stable environment. ***You will have to answer that most pertinent question: Can your enterprise afford not to protect its most critical applications, systems, and data from the effects of a disaster?*** To learn more, please read on.

Understanding the Need for Risk Mitigation

Many enterprises have just gone through the tedious task of virtualizing much of their data center infrastructure, consolidating multiple mission- and business-critical applications across one or more multi-processor, multi-core servers.[1] These servers might access a common (shared) database, placing even greater emphasis on maintaining the viability of that infrastructure. Before, if one server went down, the data center might lose one application. A mad scramble likely would ensue as the data center staff, administrators, and managers scurried about to either restore or replace the failed platform. Today, with multiple applications co-habiting the same physical server, the data center might lose a large number or even all of their critical applications with the failure of a single server or storage array, creating a potentially disastrous event, and possibly putting the staff into a recovery mode that they may be ill prepared to handle, and sending everyone else on an extended coffee break. ***Can your enterprise afford to take this risk?***

The Risks of Accidents

A robust data center infrastructure, including both the server and storage environments, is an essential part of the foundation for critical, virtualized environments, including cloud computing. However, in order for that infrastructure to be protected adequately from a localized situation or disaster, it needs to span multiple data centers. An accidental/unexpected loss of power can put an entire data center out of commission, not to mention an entire metropolitan region. (For example, how will your skilled IT staff reach your data center(s), if roads are impassible or unsafe to travel?) Occasionally, accidents will

[1] The principal difference between *mission-critical* and *business-critical* is the tolerable time it takes to reach a recovered state (i.e., being recovered sufficiently to operate somewhat normally). Mission-critical applications, business processes, and data clearly have a lower tolerance for being down than do business-critical ones. Mission-critical likely has a tolerance measured in seconds and minutes while business-critical likely is measured in hours or longer.

happen, causing the unintentional loss of critical parts of the IT environment (or the entire environment), which is why they are called "accidents". ***Are you prepared for an accident?***

The Risks of Disasters

Then there are man-made and natural disasters, both physical and virtual. Some are caused by disgruntled employees and hackers out to destroy whatever they can, because they can! Others are caused by human errors. Of course, many are caused by external happenings that interrupt the power and/or networks upon which all IT activity depends or the building where the data center resides. Natural disasters are caused by no one other than Mother Nature. These, too, can and often do affect entire data centers and necessary utilities.

We seem to be having an epidemic of these of late: earthquakes in Indonesia (2004) and Haiti (2010) causing massive destruction of property and killing hundreds of thousands; the earthquake/tsunami in Japan which caused loss of life, massive damage, and a nuclear disaster, similar to the accident at Chernobyl in 1986 (but Mother Nature was not responsible for that one). In fact, on this September 12th, there was an explosion at a nuclear waste treatment center in southern France that caused the shutdown of two nuclear power plants, which shows us, once again, the potential interruptions that enterprises must be prepared to handle.

The U.S. has not been exempt, with hurricanes Katrina (2005) devastating New Orleans and a major portion of the southeast coast, and Irene and Lee (2011) flooding many and wreaking havoc along the East Coast. There has been any number of tornados and floods tearing apart the Midwest and South, along with assorted earthquakes along the West Coast. Recently, there were multi-day brownouts in parts of Michigan. Also recently (on September 8th), a major power outage knocked out electricity to more than 6 million people in California, Arizona, and northern Mexico, taking two nuclear reactors offline, leaving people sweltering in the late-summer heat and disrupting flights at the San Diego airport.

Put together, the unexpected, prolonged outages at affected data centers potentially represent the loss of billions of dollars in business revenue due to the lack of a viable high availability (HA) options and/or disaster recovery (DR) planning and preparedness or, more simply put, **the inadequacy of the existing DR plans for enterprise-critical IT infrastructure.**

This is not meant to trivialize the loss of human life and widespread devastation but to highlight the economic risks at play for any enterprise that is not prepared to react to IT outages resulting from these emergencies. In most, if not all, of these disasters, there was little that others could do to prevent the losses that occurred. **However, there is much that an enterprise can do to prevent massive loss of revenue due to an IT outage from any cause**!

Senior executives need to understand both the business concerns and the IT concerns for their critical IT infrastructure and need to identify the risks involved in maintaining both the business and IT infrastructures.

- ***Do they understand what impact that an outage of an SPOF can have on the enterprise?***
- ***Do they know where their SPOFs are?***
- ***Is there a policy for geographic dispersion?***
- ***Is there a policy for compliance testing?***
- ***What will be the cost to your business if it is off the Internet for one week, one day, or even one hour?***
- ***How widespread is the economic impact?***
- ***What is the impact on brand and image?***
- ***How many customers are at risk to lose, if you are offline?***
- ***How long can your enterprise function, if it can't access its critical data?***
- ***Are you risking the very existence of the enterprise because customers can't trust the enterprise's reliability?***
- ***Is there something (anything) that the data center staff, administrators, and managers can do to mitigate these risks?***
- ***If there is, does it make economic sense to do just that?***
- ***Will it cost more to implement an HA or DR solution today, or will it cost more to attempt to repair the damage tomorrow (if and when it happens)?***

Let's take a look at what should concern every CEO and CIO in every enterprise around the world: ***integrating effective high-availabil-***

ity and disaster-recovery processes into existing IT infrastructure with the lowest possible incremental Total Cost of Ownership (TCO).

Line of Business Concerns

First, let's take a look at the concerns affecting your lines-of-business. From an application standpoint, the data center must maintain the highest levels of availability possible to mission- and business-critical applications and data. Any outage *will* affect the daily (ongoing) operation of the lines of business and affect the enterprise on the bottom line. However, outages will occur, even on the most highly reliable server platforms. Recovery from a single system failure depends upon how rapidly the automated and/or manual processes take to restore normal operations – in order to minimize the effect of that outage on a critical application. **The goal (i.e., the amount of time expected for recovery) is called the *recovery time objective (RTO).***

There is an unfortunate rule of thumb, here: ***the shorter the recovery time, the more costly the solution.*** Any effective high-availability or disaster-recovery solution will need to keep the outage to a reasonable (cost-effective) minimum, usually measured in terms of minutes or even seconds. If you have to measure a mission-critical application outage in terms of hours, or even days, you might want to expect a significant loss of revenue and, perhaps, a permanent loss of customers. In fact, the cost of a single outage easily might exceed the total cost of the disaster recovery solution, resulting in a potential ROI period of *zero*, if the data center experiences an outage that day! ***Can you afford not to protect your critical IT infrastructure, even if the probabilities are low?***

The business staff also has to guard against the loss of business transactions, and the potential loss of customer and other enterprise-critical data. A good data protection plan needs to anticipate a system crash or a loss of power and protect the storage with the frequent replication of critical business data at frequent intervals, in order to minimize the effort to recreate databases and files. *How frequently?* As frequently as necessary to keep information flowing! **The *recovery point objective (RPO)*** measures the amount of data that must be recovered in the event of an unplanned outage, in units of time. This is established by the policies put in place to restore data to a meaningful point in time (with no more than a specified amount of lost data). The data center needs a *Business Continuance and Disaster Recovery (BC/DR) Plan*, with RPO serving as a key factor for business-critical applications**.** ***Can you afford not to be able to achieve an effective RPO?***

The lines of business also need to worry about the ease and completeness of the integration of the enterprise applications (like CRM and ERP) with the HA and DR solutions. This integration must be tailored carefully to your environment. Anything less could result in significant delays in the implementation and rollout of your applications after a recovery. Finally, the integrated solution must be affordable on a total cost of ownership basis.[2] ***Can you afford not to have an integrated solution?***

Data Center Concerns

Rather than concern for servers on an application basis, the data center staff, administrators, and managers have to be concerned with the availability of the entire data center from a platform perspective: *they must keep all of the critical servers and storage arrays operational, r-egardless of the application set or how they have been virtualized and where they physically exist.* This was hard enough when you could tie physical resources (servers and storage, in particular) to the critical applications running on them. The effects of virtualization are positive in many ways. By simplifying BC/DR planning and by virtualizing storage configurations, cross-site failover procedures may get simpler. There is also the potential for a negative effect, as it can complicate other aspects of a HA/DR solution. Most definitely, you do have more to consider but you probably have available more tools and vehicles to solve your HA/DR challenges.

Even with all of the redundancy available in the mission-critical servers, the data center staff still needs to be prepared for a broader failure: an entire network or, even, the entire data center.

- ***What happens when power goes out?***
- ***Is there a UPS[3] deployed?***
- ***Is it necessary to support everything in data center during short-term outages?***
- ***Are all of your critical processes and active data still available when a primary site has no power or worse?***

[2] You may need to look at the *acquisition costs* (including hardware and software) and operational costs (including implementation costs, IT staff costs, and bandwidth costs.).

[3] Uninterruptible power supply (or system).

- ***Is there an alternate data center nearby?*** [4]

All data center staff, administrators, and managers need to be prepared to transfer processing capability to another local data center for immediate recovery, if one is still operational. ***Can your enterprise afford not to be prepared?***

In fact, there are significant costs that must be considered in having sufficient bandwidth to keep metropolitan and remote data centers in sync with the primary data center. This includes not only the (sometimes large) bandwidth needed to back up these sites, but also sufficient bandwidth to fail over in case of an emergency. ***Can you afford not to?***

What happens if the outage is due to flood, earthquake, or regional power failure (think of the recent effects of Hurricanes Irene and Lee)?

- ***Could any single site or campus/metro/regional fail over system survive an outage of that geographical scope?***
- ***How would your data center be affected?***
- ***Do you need a remote data center, 50, 100, or even 1000 miles (80 to 1600 kilometers) away from the primary data center, in order to keep the enterprise operational in the event of a widespread disaster? You also may need to consider different options based on local geographic, governmental and population factors, which may vary by country or region.***

This remote location could possibly also serve as an archive location. ***Are you prepared for such a widespread disaster? Can you afford not to be?***

Perhaps your enterprise needs both a local and a remote, secondary (or tertiary) site. An additional nearby (local) site may be necessary to provide a high-availability, rapid response for near-instantaneous automated fail over, while the remote site is required to provide disaster recovery relief. A three-site configuration is a more expensive alternative, and quite clearly not for everyone, especially from a TCO viewpoint. **However, if a prolonged outage could put your enterprise out of business,** ***can you afford not to have at least three sites?***

Data center managers especially must be aware of the skill set necessary to transfer total systems operation, both servers and storage to another location, even more so when the alternative site may be across the country or farther. You cannot just assume that your existing IT staff resources will have the skills and tools necessary to deploy a high-availability and disaster-recovery solution. Your DR solution needs to include the capability to install, setup, manage, and conduct regularly scheduled tests of the recovery process. For some industries, these processes may need to be certified, i.e., the data center needs to prove and document that the processes in place will work, when needed. ***Can you afford not to be prepared?***

Let's assume for the moment that your enterprise has recognized the potential danger that a magnitude 10 disaster (think very big!) can have on the survival of the data center, and the very enterprise itself, if a metropolitan area is widely affected. Even two local data centers may be compromised. Perhaps, the data center staff already has reached out to the vendors in the computer industry to find the very best products available from a wide variety of vendors. In fact, each member of your staff probably has been trained in some aspect of the executing within several HA and DR scenarios. ***Is there any certainty (or, even better, a guarantee) that all of these pieces will play together harmoniously, like a symphony orchestra, or will it sound more like many first-round tryouts for American Idol?***

You must ensure the viability of the solution as a whole and test it on a regular basis to ensure that that viability remains solid. ***Can you afford not to?***

Requirements of an HA/DR Solution

Exactly what components of the data center environment need to be addressed to ensure that the infrastructure remains available? First, you need to have platforms with a cluster-aware operating system that has the capability to monitor system heartbeats[5] and to take the necessary actions to affect recovery, based upon predetermined policies for automated switchover. The more the fail over process is automated, the

[4] This usually is defined to be within *synchronous communication distances*, practically up to about 100 kilometers, about 60 miles. If an additional data center is within this distance, then data can be mirrored with what is at the primary data center.

[5] Between itself and its fail-over "mates", whether within the single data center, or in nearby or remote additional data center locations.

better. This does not mean that the decision to fail over has to be automated but that once a decision has been made to fail over, the processes need to be executed according to a preplanned and prioritized procedure (usually, some form of a script).

The HA/DR solution must have the flexibility to meet current and future requirements; this includes scalability for the servers, storage, network, and applications. In addition, the network must have sufficient bandwidth to support today's and tomorrow's throughput needs, while minimizing the cost of that bandwidth.

The involved operating systems also need to support virtualization and its many fail-over options. A centralized, cluster-aware topology management scheme, with cluster-aware security and automated testing, is a big advantage. Obviously, a complete set of infrastructure components must be enabled for high availability and disaster recovery (i.e., not just servers).

The HA/DR solution should be pre-integrated with the hardware infrastructure to ensure ease-of-use for deployment and daily operation. If it is feasible and practical, the entire solution should come from a single supplier, in order to minimize what the IT staff has to integrate into a coherent, testable whole; the more components from different suppliers, the harder this is. That is like getting parts of a jigsaw puzzle from different sources and trying to get them to resemble a picture of your (unique) design and imagination. Integration includes the ability to work with any of the popular applications that proliferate the enterprise data center, such as those from Oracle and SAP, and also to support workload balancing, in order to maximize utilization of available system resources (across multiple sites).

The HA/DR solution needs to support effective synchronized data replication for local, or metropolitan, recovery and possibly should in-

Exhibit 1 — Critical Success Factors Often Overlooked When Considering Disaster Recovery Solutions

- ► **How easily does it integrate with the existing IT infrastructure?**
  - ***Is it readily customizable?***
    - Does it support virtualized servers and storage?
    - Is it aware of critical applications (like SAP, email, homegrown, etc.)?
    - Does it support synchronous and asynchronous mirroring, as required?
  - ***Does it do event monitoring of servers, storage, and networks?***
    - Does it handle actions based on policies?
    - Does is handle automated and manual fail over processes?
    - Does it allow operational monitoring of the status of system components, with drill down options?
  - ***Is it sufficiently flexible to meet current and future needs?***
    - Does it scale as hosts, storage, networks, and applications increase in number and complexity?
    - Is there a single point of control for multiple hosts, storage, networks, and applications?
    - Does it support multiple options for backend storage?
- ► **Are the most important risks being mitigated?**
  - ***Are recovery time and data loss being minimized to a point that is deemed reasonable and acceptable?***
  - ***Have implementation and vendor risks been mitigated or at least minimized?***
- ► **Have the all of the relevant costs of ownership and operation been considered adequately?**
  - ***Have broad issues and costs for the needed bandwidth been considered?***
  - ***Have all of the relevant infrastructure costs been considered?***
  - ***Have all of the relevant operational costs been considered?***

Source: The Clipper Group, Inc.

clude effective asynchronous data replication for remote access at long distances (when the metropolitan systems and storage have been affected). Effective implies replication processes that satisfy all RTO requirements of the enterprise. The HA/DR solution needs to support a heterogeneous mix of storage platforms from many storage providers, including the leading high-end array vendors, such as EMC, Hitachi, and IBM.

The enterprise needs to consider all of the costs associated with the HA/DR solution, not just the acquisition costs. The IT staff needs to include the expenses associated with bandwidth, operations, and implementation, as well. ***How well does your HA/DR solution help to contain these potentially large ancillary costs?***

Conclusion

Although high availability and disaster recovery options can get expensive (on a TCO basis), they do not match the costs to the enterprise associated with a total system failure. This especially is true if you also need to implement a high-availability solution that includes your SAN. The profitability of your enterprise, and, moreover, your job, may depend on keeping these systems available, at least *almost* all of the time. Cobbling together a piece-part solution from multiple vendors may not provide a satisfactory response in terms of addressing everything that is critical to keeping your enterprise viable to all of your enterprise's users, customers and partners. **Finding an integrated solution, eliminating all finger pointing, may be the best alternative, once you have identified the issues.**

Can your enterprise continue to offer its services if your data center is destroyed by an earthquake, flood or tornado? How long can you stay in business without access to needed customer data?

- ***Is it enough*** to have a single data center with a fail over system within it?
- ***Is it enough*** to have a primary and secondary data center within the local area to provide processing power in the event of a fire or other localized event?
- ***Is it enough*** merely to have a remote data center in the event of a regional disaster?
- ***Have you done enough*** for your business to survive? ***How much protection is enough?***

The answers clearly depend upon the risks that your enterprise is willing to accept when compared against the costs of mitigating the related risk.

The data center staff, administrators, and managers must determine what the minimal definition (requirements) of "survival" is for their environment (i.e., the critical applications, business processes, and data that they support). They need to determine what operational parameters are required for survival. This involves the IT infrastructure, the capability of one or more suppliers to replace or rebuild that infrastructure, and the demands of senior level management. The staff needs to keep in mind any legal requirements to keep the enterprise compliant with industry standards and governmental regulations.

The goal of this bulletin was to raise the visibility of the key requirements for high availability and disaster recovery. You need to review the requirements that have been identified and match them against the solution that you have already deployed or plan to deploy. To assist you, see the often-overlooked list of critical success factors presented in Exhibit 1, on the previous page. These questions should help you assess the completeness of your HA and DR considerations.

When it comes to ensuring the stability of your critical applications, data and business processes, and compliance with governmental and industry regulations, how much insurance is enough? If you don't know, go ask your boss and his or her boss, etc., until you know the answer. Otherwise, you are doing contingency planning in the dark, and you won't get any credit for the protections that you are providing (but you will get the blame, when the HA/DR processes come up short).

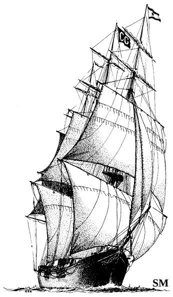

About The Clipper Group, Inc.

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

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

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

Disclosures

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

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

Regarding the Information in this Issue

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