Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER
Navigating Information Technology Horizons℠
Published Since 2007 Report #TCG2009021 April 22, 2009

Understanding the Dimensions of Cloud Infrastructure in Order to Harvest the Benefits

Analyst: Anne MacFarland

Management Summary

Recent business constraints on capital expenditures for information technology have increased a business' appetite for alternative technology scenarios. Many of these scenarios, from outsourcing to service bureaus to software-as-a-service (SaaS), have been clumped together as *cloud computing*. The cloud concept came from the use of the cloud in network diagrams to represent *someplace out there*.[1] It is a comfortable concept when it refers to Internet sites. It can be less comfortable when it refers to sensitive business processes.

People see many things in clouds in the sky depending on where they're standing. It's probably appropriate that a variety of approaches and models each claim to represent the cloud-computing paradigm. Some of the confusion is due to the marketing hype used to make each and every offering compelling – but often cloud lacking in fine print about risks and ramifications. Some is due to the different business pressures that are driving contemplation of multi-sourcing – realities like user mobility and variability of demand in addition to the inevitable control of costs. More confusion comes from the diverse perspectives of different *stakeholders* who find the cloud paradigm of IT access attractive**. The ambiguity of cloud computing definitions will persist because what cloud is depends on who you are, what you want to do, what you want the cloud (or several clouds) to do, how you want to pay for it, and what your business requirements and parameters happen to be.**

A more serious source of confusion is that, in a world used to the simple demands of more, better, and faster, other comes as a different kind of decision. It is a change in consumption modes – like a change in eating from fresh foods to frozen dinners or restaurant meals. **Cloud computing is, by its nature, a change from knowing and doing more to knowing and doing less.**

Once you get IT out of local-only mode, new choices are possible. Quality of service may be more attainable (or the penalties for non-performance more clear and financially stated) as a contracted service than as a part of internal operations. Costs may be more predictable, and, in some cases, might be more affordable.

By organizing the confusion into dimensions, it is possible to parse cloud computing in an orderly way – considering the mode change in the light of each dimension. For more details, please read on.

IN THIS ISSUE

[1] See the December 18, 2008, issue of Clipper Notes entitled *As the Cloud Turns - A Cloud Computing and Storage Exposé*, available at http://www.clipper.com/research/TCG2008063.pdf.

Origins of Cloud Computing

As a technology concept, cloud computing marries *componentization* and *declaration*[2], two imperatives made possible by copious band-\width, cheap memory, and plentiful computing cycles. Instead of vast interconnected complexity, you can have pools of components that can be enlisted to support a service. Declarative components are both *discoverable* and *autonomous*, i.e., *self-managing*. Deployed in *pools* or *ensembles*[3], they can speed up application deployment and orchestrate a process' time to completion. With both componentization and declaration, *what you get* can be specified without specifying *how you get it*. When this is combined with virtualization, the result is better asset utilization and more transparent swap-in, swap-out of elements. With adequate administrative procedures and software, IT management costs should be reduced.

As a sourcing option, cloud computing leverages several precursor IT strategies. (See discussion in Exhibit 1, to the right.) When you add in current requirements for user mobility, the evolution of client to browser, the tiering of applications (and their data), the rise of role-based customized presentation, and the move to add components in the forms of widgets and services, you obtain a set of key elements that, potentially, make clouds very useful.

A Paradigm Shift

The paradigm shift from *datacenter as a support* function to *computing as a service* is significant. Cloud computing is the latest in the cadence of paradigm changes that are buffeting the data center towards a service orientation. Sufficiency, in any service environment, is judged by the result (the impact of the whole), and insufficiency by the particular shortcoming. By contrast, in the past, a particular killer application was the bragging point and the aggregation of IT was what we often deemed to be complex, expensive, and insufficient.

With services and components that are declarative rather than dependent on other elements, data center operations can terminate more easily those that are ineffective. This truly is desirable, but it means that all parties must learn to think more atomically and be more specifically declarative about requirements. Cloud computing will satisfy requirements, but it cannot be expected to

Exhibit 1 — A Short History of Disassociated Computing

Service Bureaus, back in the days when all computing was centralized, were a way to outsource certain workloads, like payroll, that were time sensitive, or data processing that was not affordable on the premises. Many of the business processes supported by these service bureau workloads were subject to frequently-tweaked government regulations. Using a specialized bureau that could tweak efficiently across all of its customer accounts, like payroll, was more efficient and effective.

Partitions and symmetric multi-processing (SMP) were the first steps from batch processing towards a more orchestrated approach to complex workloads.

Grid Computing was an emulation of Scott McNealy's "The Network is the Computer" mantra. It managed single servers as a multi-box domain to meet the need for temporary capacity for projects and tests.

Distributed processing in the form of "Scale-out Computing" (SMP being "Scale-Up") caught on during the recent boom market as an attractive way to achieve scale and resiliency with failure-prone components. In the educational analogy in Exhibit 2, think of temporary classrooms and distance learning.

Service Oriented Architecture is another topographical change, initiated in the development realm. Open standards have allowed large complex applications to be decomposed into a collection of composed services. This decomposition allows easier evolution and customization – just as passenger trains can add specialty cars.

Cloud computing builds on the heritages above to meet the many needs to retrench, not just expenditures, but also the ambitions to do everything. This strategy is encouraged and abetted by the pervasiveness of partnering, which makes isolationist thinking (and, perhaps a requirement for unilateral control) an operational liability.

---

[2] A *declarative approach* separates *what is to be done* from the *how to do it*. It requires more pervasive intelligence, which is often called maturity. (Think of when you mother stopped leading with the *how* and just issued a peremptory *what*.) It pervades the world of technology in everything from microprocessor design to industry standards and application development (where it fosters fewer dependencies and better documentation). It underlies SOA and its precursors, which have popped up over the decades whenever the cadence of new parameters (such as regulations) was too rapid or resources too constrained to make change a decorous process.

[3] The latter term can include application-focused configurations of infrastructure aggregates (servers, switches, storage), as well as simpler infrastructure pools. Ensembles may be a more effective way to support partitions of very large databases or massive applications, such as those from SAP or Oracle.

address what we should have been thinking.

From the deployment perspective, cloud environments represent a different build mentality. With a default assumption of user access by browser or portal that may be as lightweight as a cell phone or as rich as a workstation, the separation of concerns is different from the traditional thin-client to thick-client spectrum. Deployment has to support *both or either*, as is necessary. This is a change for those who expect *done* to be a persistent state.

As an operational concept, cloud computing is the latest effort to create a data center operational model that will support all stakeholders that must be satisfied in a reasonable way. Previous service provider models were insufficiently flexible. They did not support either the frequency of change or the many dimensions of demand variability that current business operational models and current operating conditions demand.

The pragmatic trade-offs of sourcing applications and other technology is analogous to the educational systems with which we try to education our use on inadequate budgets and often-declining tax revenue. (See discussion in Exhibit 2, below.) Contemplation of cloud opportunities should be just as pragmatic.

Dimension 1: Business Factors that Affect Sourcing

Variability of Demand

The need for faster (but still safe) implementations of new forms of business support is met most effectively by a leveraging of parallelism in development, testing, and all the other disciplines that keep data and application quality high. In order that *beta* implementations be productive, testing must pervade earlier development stages – but few organizations, now that operations run around the clock, have the capacities or the budgets to support periodically-needed capacity.

The social nature of today's Web, together with fickle customers, and self-service, drive surges in demand for certain applications. The economic downturn complements the surges with crevasses. It is less and less attractive, financially, to size data center infrastructure for a single tenant's guesses at peak demand**.** *Grid Computing* (as described in Exhibit 1) educated the market to a better approach to meet variability of demand. Partitioning (and, now, virtualization) offered data centers the opportunity to size capacities more closely to need.

Exhibit 2 — An Education Analogy

When roads were rudimentary, there were many one-room schoolhouses. As roads improved and proliferated, more options became available. Over time, a rich landscape of educational options has arisen. Depending on the density and prosperity of the population and the local employment opportunities, most towns offer elementary schools, middle schools (some regional), high schools (some specialized), and vocational schools (some very specialized). These include both public (more standardized) and private (some with particular disciplines and philosophies). Most of these new options are defined by other than the geography they serve. (Think of declarative components.) We need this richness, or think we do, to move our children into a mode of financial independency in the form of a first job, or to a well-chosen college that will widen his or her options.

When SCSI cables were only a meter or two in length, computers were, of necessity, collocated in a data center equivalent of a one-room schoolhouse. Every cable was a point of failure, and centralized huge boxes, akin to large city schools, were most easily supported. Now, a combination of connectivity (the Internet) and telephony-style disciplines of resilience over grid-like mesh topography have expanded IT sourcing options to a richness as great as diverse as the U.S. educational landscape. For many, the choices to be made are just as difficult.

A business, like any community, has certain prerequisites – for security (safety) and reliability (both of the service delivered and of the sourcing organization). External sourcing of IT functions a lot like sending a child off to college. You will always worry, but the alternative of keeping children at home is not always feasible, attractive, or effective.

Any sourcing decision is political, cultural and, inevitably, emotional. With more and more organizations going "virtual," sourcing trade-offs are an increasingly important part of business (and community) planning.

What is of primary importance in any sourcing decision is knowledge of the topography of requirements and parameters and recognition of priorities.

User Mobility and Risk Avoidance

With the Internet's increasingly widespread availability, more work is done away from the traditional office. Internet-based applications hosted by an external source suits the demands of these mobile workers for access, as well as the corporate data center's need for risk avoidance. If the sourcing contract can specify qualities of service with penalties for falling short, so much the better.

Economic Intolerance for IT Support Costs

More elements of infrastructure means more frequent change, as elements fail or come to the end of their productive life or lease. One strategy for coping with frequent change is that of a *service bureau* (as described in Exhibit 1). This is familiar to most of us at this time of year as the *TurboTax imperative*. With the tax code constantly changing, it is easier to leave it to the experts to change the *what* we work with, and just enter the required data.

One can, of course, keep an expert in-house, and most organizations have many experts – but this may not be the most productive use of a man-month or woman-hour. Separately sourced, expertise can perform at a scale that makes it meaningfully productive or beneficial for all parties. This approach also is appropriate for patches and upgrades – things that benefit from the attentions of a specialist.

Dimension 2: Stakeholder Perspectives

The first dimension delineated above addressed strategic imperatives for cloud. There also are *operational imperatives*, best described by the stakeholders who bring their needs to the table. Consider what stakeholders are driving the talk of cloud computing at your organization. That profile will narrow the focus on what kinds of cloud computing you want to consider.

Consumers and Leveragers of Cloud

Developers

Developers see temporary use of cloud capacity as a way to meet deadlines. Most are moving to a collaborative mode that speeds the process and improves documentation. It also improves inclusion of new customer requirements, and supports the development of customized services.

Developers are important stakeholders at this stage of cloud computing, because they can make service elements that work in specific situations. If developers do not favor it, a cloud will not thrive. Thus, developers are in the forefront of every vendor's mind, and the focus of many cloud ambitions.

Application Vendors

Application vendors see cloud as an efficient and effective alternative to delivering shrink-wrapped products that can be incompletely deployed and used, or customized to the point where traditional upgrades to new versions becomes difficult. Software-as-a-Service (SaaS) offerings fall in this area. Many vendors offering SaaS consider their infrastructures to be clouds, in that they are irrelevant to consumption.

Cloud offers application vendors an alternative to those unpleasant situations. It provides an opportunity to add well-targeted tutorials and other user-side enhancements, as well as the ability to upgrade applications with less overt tyranny involved. While, in theory, cloud computing cannibalizes application developers' traditional revenue streams, many have found that the new markets created by pay-as-you-go opens to them more than they might lose in perpetual license revenue.

Local IT/ Data Centers

Corporate/institutional IT has the often-unenviable task of supporting business needs, whatever they become. The full array of cloud extents described in the following section of this bulletin offer them opportunities to offload initiatives and imperatives that are ineffective (usually a cost issue) to support internally. This can be everything from sporadic needs for HPC[4]-like capacity to applications that are only sporadically used, yet must be constantly available.

End Users (people or applications)

People want to leverage cloud computing from wherever they are, using whatever device they have available. Applications (think of those on the hand-held device in your delivery-person's hand) have similar requirements. The shipping company wants to be able to upgrade those hand-held devices, as well.

Cloud Providers

Service Providers

For the service providers that do the hosting, the cloud is a new incarnation of their traditional model, replacing annual fees (plus higher fees for capacity on demand) with a purer, usage-based charging model. The latter structure, by its nature, makes sense for occasional users of a service, and expands market opportunities. Multi-tenancy, one of the basic cloud assumptions, has long been a

[4] High Performance Computing.

good strategy to reducing costs. It avoids dedicated hardware provisions that made profitability at a reasonable cost elusive. With virtualized pools of resources, multi-tenancy can be supported more efficiently, and new services can be added with less cost and business risk.

Data Centers

A traditional data center can also be a cloud provider. Many excellent internally-developed applications have become standard in a particular industry, and are both licensed commercially and delivered as a service. Many data centers will find a cloud model an attractive additional revenue stream, particularly if its patterns of use (evenings and weekends, perhaps) can complement local IT demand.

Funders

Funders are found in all consumer, leverager, and provider communities. Consumers and leveragers want the clouds they use to add little overhead. Providers want to make sufficient profit to underwrite the rapid evolution that is needed to keep them attractive. For it all to work, the business model must be profitable, or worthwhile, to all parties. These days, this must be true sooner, rather than later.

Dimension 3: Extents of Cloud Capability

Bare Metal Cloud

A bare metal[5] cloud satisfies the need for periodic capacity common in high performance computing (HPC), test, and development, and by companies who have predictable bursts of unusual activity at certain times of the year. For them, a bare metal cloud offers raw capacity. In this bare metal mode, customers procure and schedule access to capacity, from a service provider. Customers are responsible for installing and managing the application as well as its removal and clean up. This is similar to classic grid operations, for which there are standard and open source grid tools to be used.

Platform-as-a-Service

Platform-as-a-Service is basically an amplification of the bare metal cloud, where an operating system and sometimes other amenities, have been pre-installed.

Basic Application Cloud

Salesforce.com offered the first basic application cloud ten years ago. It was a better way to serve salespeople who were highly mobile, time constrained, and not fond of end-of-the-day paperwork. As more software vendors adopted software-as-a-service delivery mode, application clouds have proliferated. The cloud paradigm offers users freedom from the burden of patches and upgrades. As browser functionality improves, the ability to work offline and resynch with each reconnection to the cloud is also supported, bringing the best of both worlds to many users.

As the target device for applications moves from laptop to mobile phone, the smaller footprint will call for more cloud services.

Storage Cloud

The basic storage cloud being offered by Amazon, and others, offers storage capacity, for the most part with minimal value-adds. The storage components can offer the basic storage replication needed for data protection, and the customer is charged for ingestion and retrieval (though, often, not for data deletion). Where file storage is involved, there is either no file system involvement, or files are organized by the loud purveyor's "global" file system (often federated). Recently, database clouds have come to market, most of them offering archiving services for older, static data.

Cloud Data Security

Data issues come up frequently in cloud discussions, particularly between naysayers. Most of the issues are not unique to cloud computing. In any situation, business data must be protected everywhere – at rest and in motion. The prevalence of encryption and the decreasing performance penalties of using it do a lot to address data safety fears. For particularly sensitive information, the use of masking or tokens[6] can reduce the exposure further.

Data access must be controlled. Data tracking must be ensured. Standards like XBRL[7] and other metadata disciplines can ensure this. Of course, the metadata will also have to be protected from corruption.

---

[5] *Bare metal* refers to hardware unadorned by installed software or operating system.

[6] *Masking* is a one-way process by which sensitive information is replaced by a permanent substitute value of the same data type (so it can be used for testing). In token schemes, the sensitive data is replaced by a token by which the sensitive data can be retrieved when it is needed. Masking addresses the risk to data at rest. Tokens address both data at rest and data in motion.

[7] *XBRL, or eXtensible Business Reporting Language*, is a standard that uses metadata to express the definitions of data (often business financial values). There are several global initiatives whose goal is to reduce definitional inconsistencies within specific industries.

Data Ownership

A side issue, but an important one, is the question of who owns your information, and who can gather information about how you use your application. If you are getting an application as a service, the application vendor has a significant interest in monitoring how you use the application, to guide future development in terms of features and usability. A business must decide whether this is a good think or a source of corporate risk. Usually, it is a little of both.

Most SaaS vendors will offer the option of mirroring or batching company data to the customer site, giving both a locally available record and the added resiliency of information in two locations that is needed in regulated industries and, usually for a smaller range of records, a good business practice for any organization.

Accessorized Cloud

This could be a cloud dedicated to a business process – focused on a set of applications that feed information to each other. These often come as back-end functionality (database or content repository) fronted by a tier of customer-facing access applications. Vendor visions of this space typically include middleware and development platform elements. These additions allow development-in-context and support a finer-grained development approach than that taken by the use of a bare metal cloud.

Full-Featured Cloud

This is an area of speculation – of looking to the future. In health care, a full-featured cloud might involve the ability to aggregate data sources, blogs, and other elements concerning a particular disease and to target selected views of the aggregate at particular audiences – doctors, researchers, nurses, patients, and families. Such a resource would go far beyond the basic definitional information given by *WebMD* or *Wikipedia*. In education, the ability to aggregate applications and data sources to better support targeted learning activities could provide new ways to import expertise on demand into school systems. The level of detail at which metrics can be harvested lets services be evaluated with precision.

Organizations in each and every industry have visions of what a full-featured cloud could do. However, there are many steps and many kinds of standardization that must be accomplished to realize the dreams.

Take-Aways

The service-provisioning model dictates what kinds of policies and control systems should be part of basic cloud operations and which should be specified by service levels in a cloud contract. This varies greatly across the classes of cloud described above. It varies particularly in what kinds of operational management must be provided by whom.

Most clouds are designed to keep both the business model and cloud operations predictably manageable and highly scalable. High security needs can complicate both of these operational modes.

Dimension 4: Cloud Operational Enhancements

The Business Model

At its heart, cloud computing is a *service* business model, not a *support* model. This is a matter of orientation. If you are charged with a support function, you will do the best to provide that support, using sourcing to do more with less. This is the tradition data center philosophy. The service provider is looking to build a business. Vendor loyalty and OS religion are not as important as availability and scalability. Survival depends on a well-chosen focus and optimization. Who wants to pay for a mediocre service?

Architects of an internal cloud will look at the technologies with which they have to work and make the best cloud possible with that. Those looking to build an external cloud can look at the services they want to provide and then determine the best/cheapest/most-efficient infrastructure with which to provide it.

This is like the difference between agriculture and home gardening. In home gardens, you seldom plant what you don't want to eat, and you size the garden to the expected need (of course, you include your neighbors and friends, because a small garden is just puny). Agriculturalists, by contrast, plant according to what sells well, constrained by what can grow in their location.

Cloud Economics

To be sustainable, the value of cloud, after all the faster and better and more nimble is said and done, must also come in cost savings. We are early enough in cloud computing to be in definitional mode (what is in it, why would it be a good idea), but it is not too early to consider costs.

Why should cloud be cheaper than other modes of computing? The pervasive virtualization adds many more layers of abstraction, each of which comes with a license fee. These many more layers must be managed and may be attacked or suborned. Cloud computing's flexibility, together with proper instrumentation and business rules, lets automation be better tailored to speed IT and business processes without sending either to hell in a hand basket – but automation, monitoring and analytics software are not cheap. It is only when cloud computing can be done at a large enough scale to reduce the costs of IT operations that we will look back and wonder why we ever did things the old-fashioned way.

Trade-offs

Intelligence vs. Commodity

Cloud hardware components will be whatever works best in the given situation. The structure and composition of a cloud will depend on the challenge to be met and the leveragability of particular platforms – the special features and the different approaches to parallelism that each, by its inherent microprocessor design, affords. This is on the functional side.

In terms of cloud management, the more intelligent the components are – the more they can automatically perform routine processes – the simpler the management of the whole becomes. Storage arrays invoking replication of their contents is a good and common example. Automated data migration is another. More advanced capabilities in servers, such as workload balancing and prioritization, will be essential to driving down the costs of cloud computing.

What kinds of management are needed depend on whether the cloud is private (i.e., limited in domain) or public (a commercial entity with open enrollment that may be subject to catastrophic growth or retrenchment). The latter obviously is more daunting. Management needs also depend on *where* the cloud is. A private cloud within an organization is a different challenge than a cloud spread redundantly across multiple Internet data centers in multiple geographies. Both have their benefits and limitations.

Homogeneity vs. Heterogeneity

How heterogeneous is a cloud? Do they have to be uniform clone farms? It depends on what you want to do. Anything that provides quality services at a reasonable price is possible; infrastructure becomes transparent.

What is more determinative about the kinds of cloud described in this section is what kinds of management, security, access limitations, governance and other control systems are provided by the cloud. If, as in the bare metal computing option in Dimension 3, they are not, then they must be addressed as a separate consideration or customer provided. That may or may not limit the cost and operational savings that an offload to cloud computing will provide.

Cloud Ramifications

Simplicity

The disassociated mode of cloud computing permits a focus on applications and workloads (what supports the business) and, separately, presentation details (what supports the workers). This separation of concerns allows the IT cloud to focus on knowing the workloads and applications well enough to wring out both costs and excess energy use. It allows the separate area of presentation mode to allow self-service by users to set up custom presentation elements without affecting the application. It requires just as much expertise if not more as classic IT operations. You may have thought that cloud would let you care about less, but it will also enable and force you to care about new things.

Mobility

The *where* of users is never going to settle down. The character of the *where* will be all over the map – from the data center to mid-point aggregations of sensor information to requestor applications on cell phones. All possible situations in a particular scenario must be addressable.

Elasticity

Cloud computing lets an organization consume what it needs when it needs it. It requires no entourage of cook and refrigerators equivalent. The elasticity may make some application affordable, particularly if they are rarely used.

Responsiveness

This is at the heart of the matter. As situational applications, applets, and widgets proliferate, addressing user QoS metrics (the most important metrics) cannot be done only by the traditional bottom-up approach. It must include user-side metrics and forensics. The system as a whole is, in effect, a responsive entity, not a clockworks. "Not my problem" and finger pointing do not serve in an environment where the value comes from the system as a whole. The standard of operations must be better.

Cloud Consequences

Application Coalescence

One consequence of cloud computing may be the coalescence of applications into a less fragmented and friction-filled space, albeit one now eased by software adapters and transformers. The institution of clouds allows more opportunities for situational user of the separation of concerns that keeps computing coherent, manageable, and isolated from dangerous situations. SOA does this in a bottom-up way, through aggregates of components. Cloud computing permits more of a top-down definition of useful collections that can expand, contract, or change as is useful.

Organizational Normalization.

All human activity, including business, runs in cycles. Expansion is followed by consolidation. Ambitious innovation is followed by a more pragmatic move into adjacent markets. In business process (and in data center operations), the desire to optimize is often followed by a need to simplify that challenges the previous optimization.

The greatest value of the cloud approach is not what it does or does not do – but that it can support organizations at any and all phases of those cycles. By doing so, it makes technology, something that in some situations had become part of the problem, once again part of a solution. Think of how cloud computing could change how, when, and where you use technology to assist your business operations.

Conclusion

Satisfaction with cloud operations will most likely be found in applications that have not been highly customized – or in newer applications where the front-side customization is achieved in a browser, aided by mash-ups and widgets. Cloud operations will expand how customization can be done.

Cloud operations are appropriate for well-developed operations, rapid application development, and to meet surges in demand. They are also an appropriate model for applications supporting business operations by mobile workers, and a welcome alternative to arming these workers with laptops that fragment the system-of-record and potentially expose corporate information far and wide.

Looking forward, cloud computing, as both a technology and a business model, allows technology to be used more flexibly and in a wider variety of markets to ease the pains of organizational (and, beyond the bounds of business, personal) existence. Think, then, about how cloud computing can change what you can do.

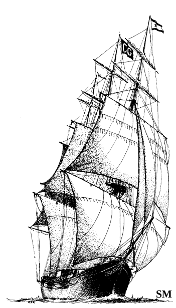

About The Clipper Group, Inc.

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

***Anne MacFarland* is Director of Data Strategies and Information Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" 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.