



Moving to a Private Cloud? Infrastructure Really Matters!

Analyst: David Reine

Management Summary

The 21st Century has dawned with many new concepts altering the ways we live, including the ways we travel. We are learning that there is more to flying than just the schedule and the ticket price. The first revelation deals with the changing paradigm of service within the shared resource of an airplane. There used to be two classes of service: First Class and Coach. Now each airline has a travel program with assorted perks and more classes than your average university. This multiplicity of choices includes offerings in each of first-class, business, and coach. We used to take the quality of service (QoS) on our flights as a given. If you travel in First Class, you got more legroom and seat width for comfort, a better meal, and free drinks. In coach, you got less legroom, a more mundane meal, and you had to pay for your adult beverage. We did not have to worry about getting to the airport two hours early, gate security, and baggage fees. We now are learning about special services that are available from the airlines (and TSA) to help us to adapt to this new paradigm. We can sign up for an airline credit card and have some baggage fees waived. There are express lines at security to enable first-class passengers and frequent flyers easier access through the maze. We can even pay extra to sit in an exit row or on an aisle. Furthermore, the airlines have upgraded their fleets were more energy-efficient engines (to lower the total cost of operation) and twin-aisle aircraft (to increase the number of passengers per flight and doubling the number of aisle seats), all to increase revenue.

We are also learning that there are multiple options for the things that we do every day, like storing our important digital images. We can do this in a photo album, digitally on our PC or tablet, or on the Internet in an amorphous, shared resource known as *The Cloud*, relieving us of the necessity to buy and manage our own private storage resource. Merriam-Webster defines a *cloud* as “a visible mass of particles of condensed vapor suspended in the atmosphere of a planet.”¹ In the world of IT, however, we have adapted that concept, creating “compute” and “storage” clouds, changing the “*visible mass*” to a “transparent mass of resources”, and changing *floating in the atmosphere* to “moving at lightning speed across the Internet”! A public cloud could be anywhere or even distributed to many unknown locations. You do not need to know nor do you really want to care. The Cloud isn’t about locality. It is about quality of service delivery, cost, and whether the services consumed satisfy our objectives. In fact, many of us may already be using a public cloud, unwittingly, via accounts with Google, Facebook, Apple, etc. For the enterprise, you need to select the right QoS to mitigate the inherent risks or you face the problem of losing data and the ability to execute operationally. Thus, the Cloud provider has to know and care about the reliability and security intrinsic to the environment. Cloud computing has become a model for the enablement of a shared IT resource, consisting of servers, storage, networks, operating systems,

IN THIS ISSUE

➤ Data Center Infrastructure Concerns....	2
➤ The Clouded Future	4
➤ Characteristics of “Good” Cloud Infrastructure.....	5
➤ What Type Of Cloud Do You Need?.....	5
➤ Conclusion	8

¹ See <http://www.merriam-webster.com/dictionary/cloud>

critical applications, and services. These elements of a common pool of IT resources can be deployed rapidly for use by the community with a minimal amount of management effort or interaction by the provider of the cloud service. Cloud computing is becoming an important transition, and a paradigm shift in the delivery of IT services – one that promises large gains in efficiency and flexibility at a time when demands on data centers are growing exponentially. The tools, building blocks, solutions, and best practices for cloud computing are evolving, and challenges to deploying cloud solutions need to be considered.

It is easy to see why *The Cloud* is a near-perfect metaphor for consuming, providing, and managing *virtual resources* (“VRs”). Unfortunately, providing and managing VRs is not without challenges and, sooner or later, virtual resources have to be mapped into real infrastructure – one in which physical assets need to be managed and optimized.

Thus, *The Cloud*² may be a panacea for some, but for others, it is real work that needs to be done (where “the rubber meets the road”, so to speak). If that point of friction is what you do for a living (i.e., you manage and optimize IT infrastructure) then you can’t assume away the issues of real assets, no matter how many times you use that “cloud” word each day. If this rings a familiar bell, then please read on to learn why the underlying cloud assets really matter, especially to enterprises but also to many smaller businesses, as well.

Data Center Infrastructure Concerns

Whether you use a private cloud (owned or leased shared physical IT assets) or public cloud (rented virtual IT assets), The Cloud needs to be able to provide sufficiently high performance and high availability (i.e., be capable of meeting the expectations for Service Level Agreements (SLAs) for the user communities). The IT resources must be scalable, in order to meet daily and seasonal needs, and it must be easy to deploy and manage. Almost always, it also is presumed that The Cloud will:

- Cut IT expense, thus lowering the total cost of ownership (TCO);

² In this paper, “The Cloud” will be a high-level concept for IT resource delivery.

- Accelerate time to value (for applications and data);
- Enable the users to gain near-immediate, on-demand access to IT resources; and
- Enable a new (and better) IT-delivery business model.

However, these goals will be hard to deliver upon without the proper planning for the long term. The Cloud provider needs to keep the “big picture” in mind and you need to ensure that you are happy with that set of services and promises. Remember that you are responsible for attaining goals over the short and long terms, i.e., you are responsible for assessing and managing the risks of your cloud strategy.

The amount of IT infrastructure is rising – not for just the enterprise data center, not for just the mid-sized data center, but for every data center. Many enterprises are building massive infrastructures, running hundreds or even thousands of servers with petabytes of data distributed over multiple tiers of storage and often across vast distances. If there is a failure in the infrastructure, as has happened repeatedly in 2011³, you will rapidly understand why the quality of infrastructure matters.

With the continuing growth of applications (in terms of quantity, complexity, and IT resources needed to run them), with the growing expectations of users (to analyze and do more in less time), and with unstoppable growth in data (to be stored, protected and made available when and where it is needed), the demands for IT infrastructure have never been greater, and will continue to grow. Thus, it is not surprising that the concept of and hopes for The Cloud to “come to the rescue just in the nick of time” are omnipresent.

There are two ways to look at this phenomenon. Either it is an illusion that promises salvation from the depths of disaster (i.e., a “universal panacea” for all of IT’s ills) or it is a hope with some real merit, as long as attention is paid to the details (i.e., somewhere along the chain of IT responsibility).⁴

The Cloud as Panacea

Pardon the sarcastic tone, but it must be nice to be a *Cloud Believer*, whose mantra might be

³ 2011 saw crashes occur to the largest clouds, including Google Docs, Microsoft’s Office 365, and Amazon’s EC2.

⁴ This will be discussed further on the next page in the section entitled *The Cloud as Real Work*.

summarized in this way:

- I don't want to sweat the details.
- Just give me (and my colleagues, partners and customers) the IT resources that we need, when we need them, and in sufficient quantities and qualities to satisfy our needs.
- Let someone else worry about procuring, delivering, managing and optimizing the IT infrastructure (resources) and data where and when they are needed. Just give me access to my applications and data in a properly-secured, "just-in-time" fashion.
- Oh, by the way, do all of this for the same budget as last year, which really means do it for less, because there is more to do this year: do it with more concern for quality and protection, and the need to do it for much more data, as well.

If you are sensing a healthy level of skepticism, you are on the right track. All of this might be possible, if every piece (of infrastructure) easily fell into place, if the glue that holds it together is strong, flexible, and not too costly, and if someone takes responsibility for doing all of the hard work, including optimizing across many objectives (not just cost of delivery, but also qualities of service, protection, and coupled with continual cost improvements). **Someone has to do the real work of delivering, managing, and optimizing the virtual resources.** If you, as a Cloud Believer, want to place that responsibility elsewhere, proceed knowingly and carefully! Otherwise, the only optimization that might result may be the flow of cash from your pockets.

The Cloud as Real Work

If you made it this far, then you have a healthy sense of skepticism about what it takes to make all of The Cloud's hopes and dreams into a functioning and affordable business reality. This is somebody's "real work". Even if you outsource all of the IT resource supply, provisioning, delivery, and management, you need to know that someone is sweating all of the important details. The problems and challenges may change (with time and application of new methods and technologies), but they do not ever go away (i.e., the challenges only evolve and never really disappear). We just handle and manage them better, increasingly through poli-

cy-driven automation.

If your requirements are so simple (or, maybe, just so mundane) that you can't imagine thinking about the underlying issues, then maybe you have no choice but to believe that your solution is in The Cloud and that you can afford what it will cost to buy your service delivery there (i.e., somehow you hope that it becomes more affordable as your volume of business activity (and data) increases). Nonetheless, if IT is not a necessary core competency for your business, make sure that you rely on partners you can trust over the long term (to cover your back when you are looking the other way). For almost every enterprise, large or small, an Internet presence now is essential for success.

If your requirements aren't so simple or mundane or even if they are, but the quantities and volumes of work and data are very large, then you need to "stick your head under the hood", because infrastructure really does matter. Whether you are doing a private cloud, a public cloud, or a hybrid cloud, the underlying infrastructure is important to you, because:

- It may enhance the likelihood of getting done what is needed (to enterprise expectations).
- It may reduce the direct (asset) costs and indirect (administration, networking, energy, and floor space) costs for which you must pay (directly and/or indirectly).
- It may affect your ability to grow when you need to grow, without suffering growing pains.

Back to Reality

If cloud infrastructure really matters, how and why does it matter? When you look under the hood of a car, today you find more than an engine. It is now a system, maybe better called "an integrated system". It is hard to tell where one set of components ends and the next begins. We all know, of course, that not all of the components are interchangeable. We also care very little about each of the components (individually). **What matters is the integrated system. Does it solve our problem? Does it get us where we want to go, in a manner that satisfies expectations?** Rest assured, many designers and engineers have put the solution together to work as a whole. They are sweating the details for you, so that you just can get in and drive away. Conversely, if someone does not cover

all of the details, you are exposed to a potential disaster.

Virtualization and Consolidation

Most of us have been through two sets of tactical IT campaigns – *consolidation* and *virtualization*. First came the consolidations of servers, storage, and networks. We consolidated servers for many reasons, but especially because of the availability of more powerful servers (that needed to be shared to be economical). We consolidated storage and networks to reduce vast excesses in provisioning (i.e., a lot of resources were over-allocated and not being used) and to make them more sharable.

Virtualization aided us in these consolidations. It allowed us to run many applications on a server (by allowing partitioning of the cores, primarily), thus lowering the cost of delivery to applications and users. Storage virtualization allowed us to share disparate resources in meaningful ways, both by tiering and by the homogenization of heterogeneous devices. Network virtualization allowed for load distribution and sharing. All of this virtualization made the physical resources easier to manage.

Doing IT Better

OK, this is your present reality. Now, you want to do it better. Better, in this case, means differently. No, you don't want to abandon consolidation and virtualization but you want to deliver them differently. Maybe a good example of this is the change can be found in the automobile industry and the quickly multiplying set of options in how we buy transportation.

Yes, we have many choices: traditional engines (gas and diesel), hybrid engines, and all-electric propulsion systems. Is one inherently better than the others are? Of course, this is a continuing debate, but each has its strengths and weaknesses. However, none of them really changes the paradigm of driving. However, you could change the ownership experience and get a *ZipCar*, when you need one and only for as long as you need it. For some, this makes economic sense. Alternatively, you could take public transportation (rail, bus, or taxi) or let someone else drive you around (say, in a chauffeured limousine).

The Cloud is a different IT resource delivery system. Private, hybrid, and public clouds are like cars and buses; there are advantages and disadvantages to each. You have to know what advantages you seek and how you will mitigate the corresponding risks.

For IT infrastructure, you seek the best ways to deliver the needed “transportation” (the IT resources that get the app, data, and/or user from here to there, or from start to finish) while satisfying a myriad of requirements. If you think that there is no difference between the underlying integrated delivery systems (in terms of resulting effectiveness and efficiency, while satisfying business (SLA) requirements), then you should leave the driving to someone else ... and take what results with a smile. Otherwise, you need to look for the best infrastructure (in that integrated system under the hood). This means thinking both technologically and economically. Think of it in terms of a possible paradigm shift. If you are going to change your economic model for IT infrastructure delivery, you need to look at the overall delivery system (the functioning whole) and the involved components (the building blocks). If you don't think this is necessary, then go back to Page 2, reread “The Cloud as Panacea”, and call it a day. Otherwise, read onward to learn how and why cloud infrastructure matters.

The Clouded Future

Weather analogies aside, your IT future will be clouded and transparent. Cloud thinking is not new, even though the term seems only to be several years old. It is about sharing pooled resources ... transparently and securely. It is about separating logical resources from physical resources. This is true for both server resources (processing power), storage resources (tiers of storage), and networking resources (the data highways). All of these resources are made transparent (you don't necessarily know where they are physically located or how they are shared) by virtualization.

The hoped-for economic advantage comes from not over-provisioning (over-allocating) resources in advance of their need, without any apparent detriments that come from sharing a pooled set of physical resources. Usually, there also are many savings to be had by less laborious (or even fully automated) provisioning. Always present are issues of control and analyses of benefits and costs. Sharing of resources is good, as long as everyone gets what they expect (including adequate separation among those sharing the resources).

The Cloud can be an ideal way to deliver application services, to improve IT efficiency

and flexibility, and to lower the TCO of that service delivery, but these will be realized only if sufficient attention is paid to the enterprise's requirements and the infrastructure details. Using The Cloud for your IT infrastructure should be transparent, but this is more than a statement about separation of logical views from physical reality. **The Cloud must be transparent operationally, i.e., service delivery requirements and economic expectations must be met.** If it works, but not as operationally expected and/or at a much higher cost per unit of operation, then the situation is bleak for everyone involved. Thus you have to proceed with more than good intentions ... you need to make sure that you know enough to ensure the outcome.

If your business is looking for choice and also to control the cloud capabilities required in its environment, there is an essential set of characteristics that constitute the foundation of The Cloud. These characteristics must all be present, to some degree: *reliability*, *performance*, *availability*, and *security*. By definition, a public cloud must be designed to support a wide variety of user requirements; thus, it usually is more generic in nature than a private cloud, which can be designed to support the specific requirements of a single enterprise, business unit or even a specific application. Some business units will need to control their own cloud resources, in order to receive the services they require and not some universal blueprint that is presented as a one-size fits all cure-all!

Characteristics of "Good" Cloud Infrastructure

Doing the Hard Stuff

To make this happen, The Cloud requires a sound virtualization platform and an equally sound virtualization management capability to control virtual resource sprawl and to provide the functionality that the enterprise requires for its critical applications. This might include a high-availability SLA (for quietly failing over to other equipment, for non-disruptive data mobility along to other storage (near or far), and for a simplified management environment, to ensure service delivery). You will need to have a sound risk-mitigation plan in place before deploying critical applications and/or data to The Cloud, in order to ensure security within a multi-tenancy environment. The Cloud solution may need some sort of identity federation, or single-sign-on policy, in order to simplify access

across the network.

Most public clouds do not meet all of the requirements of the enterprise because they cannot provide all of the business- and mission-critical SLAs. **Therefore, many enterprise data centers will need a private cloud upon which to build or host an application set and higher quality of service infrastructure to support these critical applications.** However, some applications are well suited for the public cloud, especially where the requirements are not mission-critical.

The Need for an Enterprise-Class Cloud

While The Cloud may represent the delivery of IT services, it still requires a solid infrastructure foundation to provide the automation, flexibility, management, provisioning, performance, scalability, virtualization, and lowest cost per use in an on-demand environment. The Cloud infrastructure needs to include virtualized servers, along with shared storage and networking devices. Most enterprises utilize multiple application tiers, from commodity applications such as file and/or print services, to business-critical applications such as email and advanced analytics, to mission-critical business applications such as ERP, manufacturing, database, etc.

A good *Enterprise Cloud* needs to have a solid architectural foundation, in order to be able to mitigate availability issues before you require backup and recovery services. The ability to survive an availability interruption is crucial. When you consider a Cloud, check with the vendor on the reliability, availability, and serviceability (RAS) of the infrastructure. These RAS functions should be built-in and not bolted on, to ensure a good integration among these critical services. The resources must be flexible enough to meet the changing needs of a growing enterprise. The Cloud must be able to address all of these tiers.

These are important services and the Enterprise Cloud must also be able to provide the following.

- **Backup and disaster recovery services** – reducing the impact when a disaster strikes,
- **Compression and data deduplication services** – to reduce the amount of data being transmitted and stored within the architecture and, thus, reducing storage costs, and
- **Data security services** – to protect enterprise and customer information

while at rest and in transit.

The data center needs easy-to-use system management tools to improve usability. Both monitoring and management technology is absolutely essential in mission-critical environments that demand high availability, such as supporting social networking applications, high-performance computing, and more. This is especially true for environments where the user tends to be impatient. *What will happen if you do run out of a critical resource? How easy will it be to add additional resources?* Your Cloud must be flexible enough to adapt to the changing business needs of the enterprise, dynamically.

Managing The Cloud

Do you know which cloud services you require? Do you need a public cloud or a private cloud? Some of the services that you need may depend upon a more robust infrastructure than provided by the typical public cloud. Hardware vendors have extended their management tools and reliability features to include increased flexibility. Let's take a look at a variety of the characteristics available in the cloud.

The Cloud is as much a management or acquisition paradigm as it is an architectural paradigm, at least from the users' perspectives. The Cloud can, and does, host applications to provide on-demand self-service, broad network access, the creation of a pool of shared resources, "limit-less" scalability (within reason), and a method to measure and bill for the utilization of these services. **Most public clouds do not offer useful QoS guarantees, which is why you likely will need a private cloud.** Rarely does one buy a guaranteed response time from a public cloud provider. More typically, you are buying a slice of a server, memory, and storage and network capacity. In most cases, you don't even know what kind of server you are getting. For a useful QoS guarantee, the data center needs to deploy a private cloud. The application set and data that the IT staff wishes to run in the Cloud will determine which of these services is required along with the type of cloud required. Not only are there different levels of service delivery, there also are different qualities of service delivery. These will make a big difference in what is procured and what it costs. If you want to pay the least, that is likely what you will get!

What Type Of Cloud Do You Need?

The cloud can take many forms: *IaaS*, providing *Infrastructure as a Service* to the user

community; *PaaS*, providing a *Platform as a Service*; *SaaS*, providing *Software as a Service* to enable the data center to invoke cloud-based applications, such as security applications to protect their data. The enterprise can deploy a private cloud dedicated to and operated by a single organization, a public (shared) cloud owned by a service provider and available to the general public at large, or a hybrid cloud consisting of elements of both.

- **IaaS** provides the capability to deploy multiple virtual machines (VMs) rapidly and to share applications and data between staff, partners, and customers. It is the vehicle used to demonstrate proof of concept rapidly, in order to get to value faster than the competition.
- **PaaS** enables the IT staff to move application development off-site, with a development environment. It can include multiple tiers of servers (based on CPU type, GHz rating, core count, memory, etc.) and can include a heterogeneous storage capability (FC, SAS, SATA, etc.), enabling multiple tiers of storage (to optimize the utilization of resources and lower the TCO by using diverse storage infrastructure with common access, management, and provisioning, and offering a connection with enterprise data center storage devices).
- **SaaS** is often referred to as an *on-demand software service*. It has a delivery model where both the application and its data are hosted on a centralized system, commonly in the cloud, and accessed by users via a thin client, through a web browser, over the Internet.

The Cloud should be federated (using standards to communicate between multiple platforms), automated (not requiring administrator intervention for mundane tasks), and client-aware (i.e., application aware). Moving toward that more-complete end of the service spectrum will require a focus on three basic characteristics of cloud computing: *efficiency*, *simplification*, and *security* – and on solutions that use standard interfaces and are interoperable with each other. **A successful cloud environment, consisting of servers, storage, networking, etc., depends upon the success of the application(s) and/or data that lives within the Cloud. That success depends largely upon the infrastructure upon which it is deployed.**

Critical Success Factors

The critical success factors to that infrastructure are clearly its *performance*, *reliability*,

security, availability, serviceability, agility, and manageability. Some of these capabilities, such as security and guaranteed QoS, are the justifications for deploying a private cloud. The end-user, whether employee, partner, or customer, does not know from where the resources are coming; they just need to know that they are reliable and secure. In fact, most, if not all, users of The Cloud have no idea what physical resources make up their cloud. To them, speed of access is important; location is not. However, location can be significant to the data center staff responsible for viability and performance.

Data protection is one of the critical success factors that are of prime importance for any enterprise that is utilizing the cloud for data storage. The cloud must invoke a well-defined set of policies to reduce, or hopefully eliminate, the risk anxieties of off-site data storage. This needs to be done in the planning phase for cloud storage, in order to create a firm basis from which compliance can be implemented and monitored by the security-conscious enterprise. The isolation of data and applications from other users is critical in the deployment of a successful cloud. The cloud must be able to customize capabilities according to the needs and priorities of the enterprise. The cloud must enable all organizations to get what they need when they need it. This includes protection from intrusion and from simple data failure.

Key Requirements for an Enterprise Cloud

What are some of the requirements expected in all clouds? What do we really need to consider when deploying an enterprise cloud? These requirements can be categorized into three infrastructure categories: servers, storage, and networking.

In terms of servers, The Cloud's infrastructure needs to be highly sharable and scalable with regard to the number of threads per core and the number of cores per CPU, in order to provide exceptional performance to multiple VMs. It must be resilient and capable of full utilization in terms of server resources in support of the consolidation and virtualization of critical applications, including both processing resources and memory. Scalable memory may be an essential requirement for application performance. The platform must be designed to support on-demand expansion, capable of handling a large number of application iterations. While server acquisition cost is not the most important factor in the cloud acquisition pro-

cess, it is important; as low-cost is essential. Higher qualities of service (like resilience and the ability to fail over to a remote location) will cost more, but this may just be what your enterprise requires.

In terms of storage, The Cloud's infrastructure must account for the self-provisioning of storage resources through automation, its compression, and smart allocation. It must also provide vendor neutrality in order to support a heterogeneous mix of storage tiers, enabling the migration of data from Tier-1 to a more efficient or more performant tier, helping to improve the flow or to lower the TCO. These resources must be virtualized across multiple arrays, vendors, and data centers – then pooled together and accessed from anywhere, compared to non-virtualized enterprise arrays. It should have a strong storage management capabilities that can automatically recognize usage patterns in the virtual disk and migrate the hottest (most used) data between tiers, from Tier-1 to SSD, for example.

Storage usage should be paid for on a *what's-used* basis; end users must be much more aware of the impact of their consumption and service level choices. Finally, storage services must be standardized, either done generically by tier of QoS or to a specific piece of hardware and class of drives.

In terms of networking, The Cloud should provide a converged, high-performance network, designed to improve business agility and resiliency. It must have superior bandwidth to support a myriad number of applications and cloud migration activities. Networks are potential points of failure and the cloud must be optimized for availability and a peak demand requirement.

Key Features of a Private Cloud

Private clouds are being deployed by a variety of enterprises looking to improve security behind their own firewall and to provide the enterprise-class services that they require. They are usually designed with a fit-for-purpose architecture. They are being driven by the continually-expanding business demands on enterprise IT. More and more data centers find themselves facing real limits, either from a lack of power, lack of space, lack of network bandwidth, or lack of server capacity. Expanding the existing infrastructure could perpetuate the significant problems facing the data center staff.

A private cloud can provide a “guaranteed”

SLA to a variable supply of compute resources (up to the maximum available), something that may not be available from the provider of the public enterprise cloud. It should have a cloud management system designed from the ground up specifically for your cloud environment. The data center staff may need to supplement the current systems architecture as it migrates from an existing virtualized environment into a hybrid environment.

The private cloud may need a layer of software on top of the existing architecture to facilitate the differences in existing servers, storage, and networking, and to enable scalability and on-demand access to cloud applications and data. Most likely, it must be able to support a heterogeneous mix of servers and storage arrays, virtualizing storage over a wide range of vendors, including EMC, Hitachi, HP, IBM, NetApp, Oracle, and others. This enables The Cloud to utilize existing enterprise resources and ensures that the solution architects can choose the right tool for the job. Be wary of colleagues and vendors who believe *one size fits all* when it comes to Cloud infrastructure for application delivery. This infrastructure could be driven by *Windows*, *Linux*, or a *UNIX* operating environment, or any hardware platform, from x86 to Mainframe. The storage could be a mix of FC, SAS, SATA, SSD, even LTO-5 or enterprise tape for long-term archiving.

Conclusion

Cloud devices must be location independent, deployed with standardized interfaces for application flexibility. The cloud has to be reliable, sustainable (easy to maintain), and secure in support of a multi-tenant environment. The cloud infrastructure needs to be capable of self-service and be highly automated, including a capability for monitoring all resources and metering system use. It must provide the *right* virtual resources for the *right* application at the *right* time.

Flexibility demands that The Cloud support virtual resource mobility in order to sustain a high-level of application utilization and disaster avoidance. Thin provisioning, automated tiering, and application-integrated snapshot and mirroring are also essential components that are dependent on a sound infrastructure basis. The Cloud must include a centralized management capability, driven from a single pane of glass, to monitor cloud health, capacity utilization, and

performance. Automated provisioning and fail-over, along with charge-back, are features that may be required in an enterprise cloud.

The Cloud provides the capability for the enterprise to grow through transformation. This transformation enables a new delivery model through the use of mobility within the cloud using a solid foundation based upon a sound infrastructure. The most obvious expectations from the implementation of a cloud environment are reduced TCO, improved efficiency, and application and data mobility. Whoever is providing the service must be driving cost out of the equation, while ensuring that the QoS meets the required SLAs. Nonetheless, the lowest-cost provider may not be meeting your QoS requirements. You need to make sure.

Sharing the infrastructure goes one step further than simply getting high utilization from virtualization. *Now* may be the time for your enterprise to rethink its business model. *Now* may be the time for your data center to change the paradigm of how they consume IT resources.

Whenever anything in the infrastructure fails, whether a private cloud or a public cloud, the IT staff immediately looks for someone to blame. If you have used a heterogeneous approach for your private cloud needs to create a competitive environment, you can also be sure of a lot of finger pointing. That is why some in the IT industry prefer to use a single vendor for their IT needs. Some refer to this as “one throat to choke”. However, some might see this as an indicator of failure! Something has gone wrong and you are looking at someone to blame. On the other hand, with a single source you can have “one hand to hold” to guide you through the potential quagmire of a system failure. Whichever approach you prefer, you need to be sure of the functional integrity of the infrastructure. **The quality of infrastructure really does matter, as does attention to all of the requirements and details!**

The Cloud awaits you. Go for it and make it work for you and your enterprise.



About The Clipper Group, Inc.

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

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

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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.