THE CLIPPER GROUP
Captain's Log™

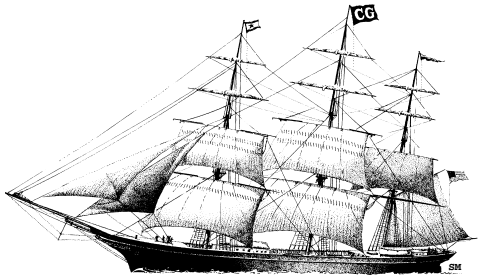

Navigating Information Technology Horizons℠

Published Since 2001 Report #TCG2008063 December 18, 2008

As the Cloud Turns — A Cloud Computing and Storage *Exposé*

Analyst: Mike Kahn

For Better or Worse

Lately, we have seen a lot written and said about *Cloud Computing* and *Cloud Storage*, often under a variety of other names, *like co-location provisioning*, *on-demand provision*, and *Internet provisioning*. Sometimes, the focus is on providing a *buffer* or *cushion* (i.e., the *added capacity*) to what you might provide in-house but, these days, many folks are talking about delivering *infrastructure and/or applications on the Web*. Now, add this *Captain's Log* to the collection. This isn't another tech paper explaining what the cloud is all about and why you have to move, post haste, into *Cloud Mode*. This is an explanation of why ***Cloud Infrastructure* may not be fundamentally different from what you should have been doing, increasingly, over the last handful of years**. Therefore, you might call this bulletin a *Cloud Exposé*.

To get quickly to the bottom of this murky topic, **you need to ask yourself if Cloud Infrastructure is about *where something resides* and *whether you care where something resides*.** Most of the chatter is about *where* something resides, such as *where is my data* or *where does my processing take place*, even if the general answer is *somewhere out there*. For the moment, I'd like to refocus that discussion to *whether you care* where something resides.

Any enterprise employee who has had a PC or notebook understands the two likely answers to the question of where: (a) either "on my computer" or (b) "somewhere else".[1] If you drew a graph showing your confidence in or comfort with each of the locations, it might look like this:

Exhibit 1 – Confidence in Location

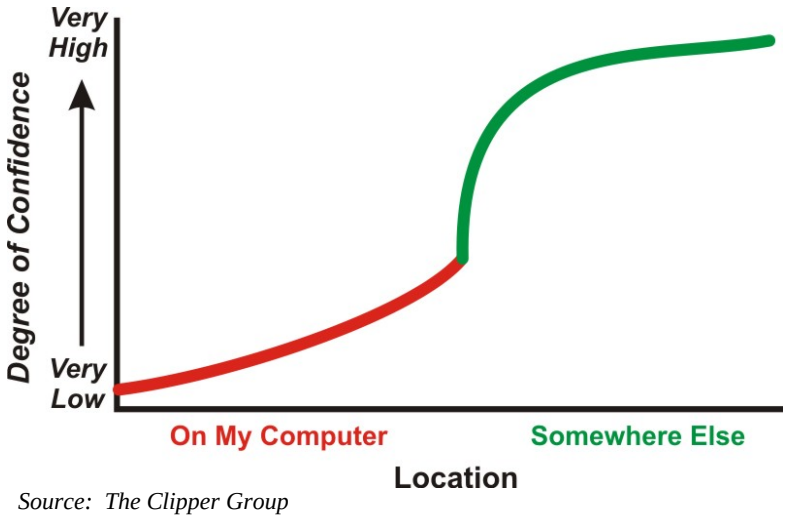

Source: The Clipper Group

[1] There might be a third state of existence called *in transit*, such as on a USB memory stick or as an email attachment, but let's ignore this third state for now.

What this abstraction shows is that you know that you have less confidence in your personal computer as a place to run applications, or to store your data, than you would have if it was running on an enterprise-class server or the data was stored somewhere else, say on an enterprise-class storage device.[2]

While a few of us may do a better (or worse) job of managing, protecting, and optimizing our PC or notebook than the rest of us (which explains the range of confidence of "On My Computer" in Exhibit 1, on the first page), there probably are none of us doing this to *enterprise-class standards*. What this means is that we all know that *enterprise-class* means something other than *what we do on our local computer*.

By agreeing with this, **we believe that *enterprise-class* – assuming that it is managed to higher policies and better qualities of service (QoS) than we do on our own local computers – is superior regardless of where (locationally) it is running.** In a world where all of the good words have been used for many IT meanings, let me sin further by calling this *location transparency*. **Thus, any location, other than on our personal computer or notebook, usually is better than running on my personal machine, as long as we are connected to that other location.**[3]

For five decades, the best of those "other" locations was the *enterprise data center*. Its QoS always has been better than our personal machines or the server (or *mini-computer*, if you are old enough to recall these), or network storage device in a closet at the end of the hall. **In terms of location transparency, *away* always has been better.**

As I write this, a song keeps playing in my head: ***Somewhere Out There***[4], which was the Best Song of the 1988 Grammy Awards.[5] **That's where the better IT resource always has been – *somewhere out there*.**

More correctly, it has been *Somewhere In There* (referring to its residence inside some enterprise data center…*somewhere*). OK, you might now get a sense of where I am going with this all. **We have been relying on the concept of a *better somewhere out there* for a very long time.** Of course, we have been (regularly) distracted by the promises and fantasies of owning it all right in front of where we are sitting, at least until we realized that we now were responsible, to some degree, for the care, feeding, and protection of our applications and data.

However, if one relates Cloud Infrastructure primarily to *Somewhere Out There*, then we've been doing it for decades. So what makes it different now?

- ***First was the notion of location.*** *Somewhere In There* refers, at least on initial consideration, to the local (i.e., within the enterprise) data center (wherever it is located physically).
- ***Second was the notion of ownership***. The closer that we were, physically, to the application processing and data, the more that we felt that we owned it and could use it for our personal benefit, i.e., we could control it.
- ***Third was the notion of possession.*** We felt more enabled, or more powerful, if we could take it with us into our daily battles.

Weaning us from caring about location, ownership, and possession was a first order instantiation of *virtualization*, if you define it as separating the realities of the physical world from its logical representation. It appeared to us (logically) as if it (especially the data) was residing with us, even when it was physically elsewhere; well, almost. As became most clear when we took a notebook computer[6] on an airplane, there were and still are times when we care about physical possession or closeness. Ignoring that for the moment, **IT organizations have been providing an *internal cloud* for years.**

- ***Is the difference related to the vehicle for accessing the processing power and data?*** If I am using a browser to give me a representation of the application and my data that is running and stored somewhere *in the network*, is that really any different than running it all on a laptop? Yes, but that might not be enough to exclude the enterprise data center from the definition of Cloud Infrastructure. You see, at least as an explanation to those of you under 40, that is exactly the mode that we were in with *smart terminals* in the 1980s. The smart, albeit proprietary, terminal managed the local representation

[2] Unless you are trying to keep something private, such as personal data that you do not want to put "out on the network."

[3] There are many variations especially over the decades that could be considered here, but they are unnecessary to comprehend the arguments herein.

[4] Written by James Horner, Barry Mann, and Cynthia Weil.

[5] Somewhere Out There (James Horner song)." *Wikipedia, The Free Encyclopedia*. 1 Sep 2008, 15:39 UTC. 6 Nov 2008 <http://en.wikipedia.org/w/index.php?title=Somewhere_Out_There_(James_Horner_song)&oldid=235609166>

[6] Or a cell phone with a web interface, where both the processing power (i.e., the running applications) and the data are not physically present.

(think, *GUI* or *presentation layer*) to applications and data running elsewhere.

- ***Might the difference be the speed of the networks?*** Quite possibly, but this is not enough to rule out the enterprise data center as the engine behind Cloud Infrastructure. Internal networks are almost always faster than external networks, which many associate with Cloud-based infrastructure. So, network speed, especially rapidly-improving wireless network speeds, may increase the transparency of location and thus the acceptability of working (or living) within the cloud; it isn't enough to explain the apparent distinction between then and now. So, our search for clarity continues.
- ***What about the costs of providing access to servers and storage?*** While it is implied by many that the costs will always be less for externally-procured Cloud provisioning, I am not convinced that this is true on a large scale.
- ***Related to this might be the apparent universality of open server platforms*** (employing "industry standard" processors from Intel and AMD), but I don't buy that argument either. More likely, it has something to do with the software vehicles being deployed (Java, XML, open source applications, etc.) to work with standards-based browsers but, again, there's no difference between the data center's cloud and an external cloud.

Three Factors

OK, I've been leading you, intentionally, on a bit of a wild goose chase. I think I know the answer. For many, it involves three factors: *ownership*, *responsibility*, and *concern*.

- **The *ownership factor* refers to the underlying assets in the cloud.** *Does the enterprise own them or are they part of a leased provisioning or service acquired from a third-party cloud provider?*
- **The *responsibility factor* is about who is making guarantees about the qualities of service?** *Is it an internal organization (i.e., within the enterprise) or an external organization (contracted by the enterprise)?*
- **The *concern factor* is derived from the first two factors.** *Who has to worry about the provisioning and making it all work (within budget or contracted price)?* Many believe that there is less concern if you leave it to someone else.

Thus, the new hot topic of Cloud-based Infrastructure may be based more on the issues surrounding *outsourcing* than anything else. No doubt, many smaller businesses and even mid-sized enterprises may have no good reason (i.e., no economies of scale or proprietary advantage) to provisioning their own Cloud infrastructure from within their own enterprise data center. In this sinking economy, it may be easy to presume that every enterprise should get out of the data center business and leave it to the "professionals."

However, I disagree on four counts.

- ***It's not that simple.*** The key to the Cloud infrastructure is transparency. Your enterprise may benefit from holding its own keys to the Cloud.
- ***It may not be cheaper, over the long run.*** (Isn't it always cheaper to own a car than to lease one?) Maybe not as fashionable, but these hard times require rethinking a lot of our bad habits.) Outsourcing (the Cloud) appears easier and less expensive, at least if you only look out a few years. However, most of us need to consider the long term.
- ***There are many more levers to pull, if you manage your own Cloud(s).*** However, if you are just "another customer" of the Cloud Service, you are stuck, at best, with the terms and conditions clearly delineated in your service contract.
- ***The answer might be "both".*** A federation of internal and external virtualized sources (clouds) might be the best answer. With

Exhibit 2 — SaaS: Another Side of the Same Coin?

You might be thinking ***Cloud Infrastructure isn't really about the underlying infrastructure; it's about delivering an application in a locationally-transparent way.*** *Software as a Service (SaaS)* might be the phrase that immediately comes to mind.

For some, possibly smaller, enterprises, the thought of acquiring all information access and processing through a singular SaaS portal may be very attractive. Most certainly, this is a form of delivery via *the Cloud*. However, the prospect of procuring and delivering a number of applications from a variety of SaaS providers might be less attractive, in terms of the user experiences, the audit trails, the economics of outsourcing, etc.

When considered as the "big picture" rather than its individual components, it seems that – especially for larger enterprises – this still is a question of *outsourcing*, i.e., whether to do so, or not, and why.

virtualized services, you probably won't be able to tell the difference.

Hey, Wait A Minute – What If Your Locally-Managed Infrastructure Is The Problem?

This is an unfortunate truth for many smaller enterprises. The chore of providing high-quality systems infrastructure (think about 1-10 servers) may never justify the cost and/or effort required to meet minimally-acceptable levels of service delivery. Competence usually requires both sufficient experience and enough critical mass, where the latter refers both to *units of infrastructure* and *to the folks who manage the provisioning*. In this case, outsourcing the servers and storage may be the best answer. Then you can focus on the applications (i.e., transparently using the infrastructure).

Conclusion

Cloud thinking is good, as long as you don't let your thinking become cloudy. For larger enterprises, outsourcing it all may not be the best answer, especially in the long run.

Remember, your customers and users already see your applications as being delivered transparently. That makes this primarily an economic and strategic decision, not a technical one. ***Somewhere out there*** **might just be** ***somewhere in here.*** **You shouldn't care! And that is the key point behind location transparency.**

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

Mike Kahn **is Managing Director and a cofounder of The Clipper Group.** Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 16 years at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation - Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics startup; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan, where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial "121" 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***, 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.