

Can You Tell One Cloud from Another?

Ed Gubbins

New Paradigm Resources Group, Inc.¹

January 10, 2011

With the possible exception of “net neutrality,” “the cloud” is the most abused technology buzzword of the last few years, applied to almost anything that involves the Internet. One of the more troublesome aspects of this phenomenon is the way in which people often refer to “*the* cloud” as if it’s one thing. It’s not. As cloud computing proliferates and is applied to ever more disparate needs, it will be important for service providers to discern the various segments of the markets and the very different service models that share this term.

Start with what is perhaps the most fundamental and immediate distinction that must be made in any discussion of “the cloud”: private vs. public. So-called “private clouds” – virtualized environments within an enterprise’s own dedicated data infrastructure – are distinct enough from public clouds – those provided by others using infrastructure that is shared among multiple customers -- that’s it’s difficult to have a meaningful conversation about “the cloud” without first identifying which one of the two you mean. It’s akin to talking about trends in the restaurant industry and those in your own home kitchen in the same breath; some of the technology may be the same, but the business models and surrounding concerns are night and day. However, that important distinction is somewhat muddied by the fact that most enterprise cloud adopters are taking a hybrid approach – simultaneously developing public and private clouds that, at any given stage, may or may not be directly integrated (at this stage in the technology’s adoption, usually not).

Here’s another key distinction: Cloud computing (from here on, I’m talking about the public kind) is sometimes a managed service and sometimes not. Rackspace Hosting, for example, one of the rising stars of the cloud computing services business, only began offering a managed version of its offering in December 2010. Though some cloud services are gradually evolving in the direction of managed services, much of their early adoption has been driven by the likes of Amazon and others with little to no hand-holding of customers. In fact, a lot of enterprise adoption thus far has occurred from the bottom up, without enterprise CIOs knowing about it. Employees need computing resources and don’t want to have to beg IT managers for it, only to be denied because budgets are tight or delayed by bureaucratic roadblocks. So they go to Amazon or Google and purchase IT resources directly, over the web, with a credit card. Over time, they start talking to their coworkers about how well it went, and the company starts doing more cloud computing. Then they start using it for more important tasks, which requires more reliability and accountability – and in particular, more security. So they begin to talk about the benefits of a more substantive, formal relationship with a

¹ *New Paradigm Resources Group, Inc. (NPRG) is a Chicago-based research and consulting firm focused on identifying, analyzing and forecasting emerging technologies and industry trends, fostering mission critical decision-making processes for service providers, technology developers and financial institutions. The author is Senior Analyst at NPRG and can be reached at egubbins@nprg.com or (312) 980-4744.*

cloud provider, and they start asking about managed services. As NPRG points out in its [recent report](#) on data center strategies, the distinction between managed and unmanaged services is an important one for service providers entering the cloud services game from the telecom and IT space, since one of the advantages they might hope to exploit is an existing relationship with CIOs and IT managers. For providers of hosting services such as Rackspace, the co-existence of managed and unmanaged offerings is nothing new, as managed and unmanaged hosting services have been around for some time.

As cloud computing covers more ground, the services themselves will differentiate to the point that offerings on one end of the spectrum may bear little resemblance to those on the other end. Some offerings will be bundled with the network transport services of the provider; some will tout network independence – or the choice and redundancy of multiple networks -- as a selling point. Some will involve critical applications with real-time performance demands that require very low latency and thus are sensitive to proximity and physical distance; some, like Amazon's, will be ambivalent to geography. Steve Smith, the CEO of Equinix, staked out *his* company's turf in a recent speech to investors: "We're most interested in proximity-based, high-priority apps; they'll demand the lowest level of latency," he said. "We're less interested in cloud-based apps that are back office[-based] and don't require density or choice of network."

Some cloud computing offerings don't even translate well from one customer size to another. Verizon Business had its own internally developed cloud computing service on the market for several months before it decided to add Terremark's down-market cloud offering to its own portfolio, acknowledging that Verizon's home-grown service was so specifically tailored for large enterprises that it would take too much re-engineering to adapt the service for small and medium businesses.

Over time, differentiators will evolve around industry verticals and other customer segments. For example, cloud services in the health care industry may adapt to meet specific demands for regulatory compliance and security. Services aimed at corporate collaboration may focus more on low latency. Government customers are likely to dictate the specific geography of some cloud offerings, as Verizon illustrated by occupying a big chunk of the DC-area data center owned by Terremark, which gets much of its revenue from the feds' aggressive exploration of cloud computing. And of course, services aimed at residential users will focus on entertainment and media – perhaps including content delivery networking elements as well -- with a particular view toward managing subscriber identities across a range of consumer devices. Meanwhile, momentum is also growing for open-source cloud computing; in January InterNAP, a provider of telecom and IT services, became the first to announce a major storage offering based on OpenStack, the open-source cloud-computing platform developed by Rackspace and NASA.

As the differentiating aspects of these services become more pronounced and better understood, it will be helpful for the industry to develop new terms that delineate their differences. What we call "the cloud" today will, in 10 years, go by many different names.