Galvin's All Things Enterprise

The State of the Cloud

PETER BAER GALVIN

Peter Baer Galvin is the CTO for Corporate Technologies, a premier systems integrator and VAR (www.cptech.com). Before that, Peter was the systems manager for Brown University's Computer Science Department. He has written articles and columns for many publications and is co-author of the *Operating Systems Concepts* and *Applied Operating Systems Concepts* textbooks. As a consultant and trainer, Peter teaches tutorials and gives talks on security and system administration worldwide. Peter is also a Lecturer at Boston University and Senior Contributor to *BYTE*. Peter blogs at http://www.galvin.info and twitters as "PeterGalvin." pbg@cptech.com

It is with a tear in my eye that I change the name of this long-running column from Pete's All Things Sun to Galvin's All Things Enterprise. The tear comes from a longing for the good ol' days when Sun was a major contributor to technology innovation. With the purchase of Sun by Oracle, the use of the company name Sun no longer makes sense. And while I'm changing some things, why not change the scope of the column as well? In fact, truth be told, some of the previous columns weren't really Sun-focused, but strayed from the course to cover interesting technology topics that were just at the periphery of Sun. So at the prompting of Rik, the editor of *;login:*, both the name and the scope of this column have changed. Certainly, Oracle will be a topic from time to time, but there is a lot of innovation happening in the IT space, and with the new name, it will be within the purview of this column. But what topic should be first for this expanded column? How about the most hyped (or is that over-hyped?) new area of IT innovation: cloud computing?

Everything is cloudy these days. Between old-school vendors having (or claiming to have) cloud-centric products and services and startups that may or may not have the next great cloud thing, it's impossible in the IT space to avoid hearing, reading, getting marketing about, and generally being bludgeoned by cloud. In fact, the use of the word "cloud" in all things technical has caused some to come down with cloud-itis—one more mention of something "cloud" could cause a serious injury (or cause them to seriously injure someone).

Which of course leads me to write this column about cloud. But wait, didn't I just admit that the world is overly cloudy? Indeed. However, this column is not going to introduce some great new cloud thing that you didn't know you couldn't live without until I told you about it. Rather, this is intended to be a pragmatic cloud sanity check. What are IT managers doing about cloud, and what aren't they doing? What seems to be a cloud-based improvement on the old way of doing things, and what is just pie in the sky? And just what is "cloud" anyway? Read on to see my take on all things cloud.

What Is Cloud Computing?

Cloud computing is many things to many companies' marketing departments, and even analysts have trouble agreeing on what constitutes "real" cloud computing. For example, some definitions include some kind of remote access requirement, while others say that the component must be part of a shared infrastructure to be considered cloud. Gartner has the most sane definition [1]: cloud computing is "a style of computing where scalable and elastic IT-enabled capabilities are provided

'as a service' to external customers using Internet technologies." At a more detailed level, I believe a cloud-based solution has to include these aspects:

Elastic/on-demand/scalable—needs to have the ability to rapidly scale to meet potential reasonable demand
Service-based—facilities provided as services rather than ad hoc or fixed implementations
Shared—used by multiple entities concurrently: for example, multiple internal groups or external companies
Metered/monitored—either charge for use or monitored for use (chargeback or viewback)
Internet technologies–based—enabling potential access from anywhere; location independence; secure remote access; and, in general, the benefits brought by Internet technologies

These requirements allow for a general cloud definition to include both public cloud and private cloud versions. *Public cloud* is the more common and the first salvo in the cloud wars, but *private cloud* is also becoming a major player and needs to be part of any cloud discussion. Public cloud is a service provider making resources (CPU cycles, disk blocks, applications, etc.) available from its datacenter, meeting the above requirements. A private cloud solution is similar but executed within a company's datacenter, with the added qualification that it is under their security control.

What of the other forms of cloud that are making the rounds? *Hybrid cloud* is less a cloud form than a cloud strategy. Hybrid is the use of both public and private cloud computing. A given project, for example, might use both, or separate projects might involve one or the other. The consensus seems to be that, rather than there being a mass migration to the public cloud, for example, this hybrid form of cloud computing will be the mainstream for the significant future.

Another major cloud computing form is the *virtual private cloud*. In essence this form gives IT managers more control over the security and manageability of a public cloud solution by segregating part of the public cloud into a more private cloud-like environment. It is between public and private cloud in terms of costs, security, and manageability.

Those of us who are jaded computing veterans recognize in cloud computing many aspects of other previous-generation solutions. In fact, the Wikipedia entry about cloud computing [2] rightly points out that there are aspects of autonomic computing, client-server model, grid computing, utility computing, peer-to-peer, and service-oriented computing. I would go further and say that many of these ideas have been incorporated in thin-client computing, mainframe, online-service providers, and time-sharing computing. If you doubt this, consider that you could rebuild CompuServe (the old online-service provider) by buying CPU cycles from one public cloud provider and disk space from another.

Why Cloud?

If cloud computing is similar to previous computing models, what makes it different? Note that in the definitions given so far, there is nothing about lower cost. However, that is one of the driving factors. The lower cost comes from efficiencies of scale of the large cloud providers, and also from competition. Competition is the key difference between the old models and the cloud model. Not only are the cloud providers competing with one another for price/performance and features,

but for the first time, they are providing competition to IT management. This is a seismic change. Never before could a business manager use her credit card to pay for infrastructure to host her new technology-enabled business offering, but that is just what is happening in companies worldwide. Note that this is not necessarily a good thing, as IT management is responsible for SLAs, governance, and policies. For example, a company might have a policy that states that "all tier 1 data must be replicated from the production site to the DR site"; such a policy would more than likely be disregarded by the business manager. Who will be responsible when a problem occurs because of the violation of this policy?

How, then, is IT management going to respond to competition? Some are fighting against it, issuing policies that say cloud facilities cannot be used by business components outside of IT. And sometimes this is for good reason, as discussed below. Others are letting nature take its course, enjoying the surprised look on the business manager's face when they get the unexpectedly large credit card bill ("How could $0.12 per hour add up to that?!"). Others are using that competition to their advantage, harnessing cloud solutions to lower their costs and increase the features, functions, capabilities, rapidity of response, and elasticity available to them by using cloud services to meet business needs. They are moving to provide IT as a service to their customers (employees at their business), and those services are frequently composed of hybrid cloud components.

Of course, with any new computing solution there are those who rush to be at the bleeding edge and those who wait for solutions to mature, see if they stand the test of time, and then use them as they fit into their IT strategy. Equally obvious is that not all cloud technologies, solutions, vendors, and products are right for every IT manager. VDI (Virtual Desktop Infrastructure) is the current case in point. Many sites are exploring VDI, but many are finding that it is proving difficult to execute at any cost savings over their current desktop architecture. Some sites are moving to VDI in spite of this, discovering that it solves other thorny issues such as security and easy (and secure) remote access. Others are evaluating new solutions as they appear, waiting for one that fills their set of needs at their needed price.

That leads to another interesting point about cloud computing—it is without a doubt driving a new, high level of innovation in IT. Startups are being funded to provide cloud solutions, and to provide solutions to the new problems that cloud computing is creating. Existing companies are revamping their existing offerings to make them more pertinent in the cloud epoch (or at least having their marketing spin stories about how their offerings are cloud-relevant).

What's Next

So far this column has defined cloud computing and has given reasons why it is important to IT. In the next edition of Galvin's All Things Enterprise, I'll discuss examples of companies making good uses of cloud solutions, I'll give reasons that others should be or are avoiding the cloud, and I'll finish with a detailed list of IT aspects to consider when deciding whether a given service should be public cloud-based, private cloud-based, hybrid strategy-based, or kept as is on current infrastructure.

References

[1] http://www.gartner.com/it/page.jsp?id=1035013.

[2] http://en.wikipedia.org/wiki/Cloud_computing.