

ON OUR
RADARIQT
IN·Q·TEL

DEMYSTIFYING CLOUD COMPUTING

By Greg Shipley

Cloud Computing has gone mainstream, and by now most people have heard the pitch: the IT landscape has forever been changed by improvements in automation, economies of scale, developments in “big data” analytics platforms, and advances in virtualization technology. Pundits would have us believe that Cloud Computing will revolutionize the way computing is done — simply embrace the new paradigm and IT nirvana awaits. Few would dispute that serious change is afoot, but is the road ahead really *that* simple?

We think not. While there have been notable changes throughout the industry, the drivers behind these changes have been in play for some time. The Cloud phenomena has been more evolutionary than revolutionary. There is also reason to believe that this transformation is still very much in its adolescence, with substantial changes still ahead. The wise strategists will identify when to use these evolving technologies, and perhaps just as important, when not to. With the rapid growth of data sources and data holdings, the pressure to do more with less, and a dizzying set of claims regarding what is possible, it has never been more important to separate the fact from the fiction.

The Only Constant is Change

You needn't be an IT professional to notice the seismic changes occurring within the IT industry. Compaq, DEC, SGI, and SUN Microsystems are all names of the past. 2010 saw Apple surpass Microsoft as the highest valued technology company in the world.

Google now employs over 30,000 people, Facebook claims over 800 million users, and Amazon has grown from an online bookseller to a data service provider for thousands of organizations across the globe. Smart phones, now designed to be tethered to remote computing farms, possess features that meet and often exceed functionality found on personal computers, and will soon eclipse PCs in total units sold. Next generation petabyte-scale data processing technologies such as Hadoop originated not from industry heavyweights like EMC, IBM, and Oracle, but from small teams of developers working to solve big challenges in an open manner.

At the heart of each of these changes lies pieces of the Cloud Computing paradigm:

Elasticity — Amazon's retail business demanded vast computing resources for seasonal consumer “bursts,” but those systems went idle during off-peak periods. Embracing a services-centric deployment model and figuring out how to scale up and scale down resources



With the rapid growth of data sources and data holdings, the pressure to do more with less, and a dizzying set of claims regarding what is possible, it has never been more important to separate the fact from the fiction.

led to the creation of Amazon Web Services. Dozens of on-demand data products emerged as a result of those efforts and they now support thousands of organizations.

Distributed workloads and broad network access — Apple’s voice recognition technology, Siri, combines local device input but limited processing power with more advanced processing and data services “in the Cloud.”

Embracing open source models — Amazon, Facebook, Google, LinkedIn, Twitter, and dozens of other companies have open source technologies at the heart of their computing infrastructure, and their businesses could not run without it.

Successful organizations have shown the ability to embrace and inspire change. Other success factors include adoption of open source initiatives, use of “off-brand” hardware, rigid standardization, fixed workloads, shifts in development processes, new labor metrics, and perhaps the most important, a willingness to challenge the status quo. These forces might be considered “behind the scenes” elements, but without them this new IT order could not have materialized.

Forces at Play

Many Cloud heavyweights have focused on the technology areas where the most value can be delivered, and have driven the cost of everything else down by accelerating commoditization wherever possible. This is good news for many IT purchasers, but bad news for technology companies sitting on or near the commoditization curve.

One vehicle for executing this plan is open-sourcing software components that aren’t considered core intellectual property, and spreading out the workload required to maintain these projects. There is no shortage of open source poster children: the

creation of Hadoop at Yahoo, Facebook and Twitter’s role in getting Cassandra off the ground, Amazon’s use of Linux and the Xen hypervisor, and the large assortment of organizations behind OpenStack, to name a few. But these efforts have had unintended consequences. For example, the openness of the Linux operating system, its support for non-x86 based CPU architecture, and its subsequent rise in the embedded systems space paved the way for it to become the Android operating system. We’ve also seen this “develop and release openly” approach slowly begin to move beyond software.

Several years ago, Google surprised IT professionals around the globe when it began releasing select information regarding its data centers and server designs. At the time, conventional thinking prescribed data centers with raised floors, servers from brand-name suppliers, redundant power supplies, redundant disk arrays, and redundant network connections. Google decided to forego the traditional data center build-out, leveraging industrial shipping containers, stripping down systems to only the raw essentials, lowering per-server unit costs by avoiding buying from the big brands, skipping RAID arrays, eliminating redundant power supplies on system boards, and even removing the need for chassis enclosures. It was different.

In 2011, Facebook’s IT professionals not only embraced elements of Google’s model, they took it a step further by designing and open-sourcing the CAD diagrams for their hardware system specifications, actively encouraging others to contribute to the effort. The concept of open-sourcing in Silicon Valley is no longer limited to just software; hardware is now officially on the table.

But there is a dark side to some of these dynamics. Just as the rise of open source technology can be swift, the descent can be equally as stunning. What

happens if the primary contributors for an open source initiative all work for the same organization? What happens if that organization switches directions, re-prioritizes, or abandons the project? What happens to the cost of a software or component if the project becomes orphaned and the organization now using it has to foot the bill for further development? The health of an open source project is directly related to the resiliency of the team behind it.

The Challenge of Assumptions

Cloud Computing technologies and providers alike are rarely concerned with legacy systems or data, and most are barely over a decade old. Data is often “clean,” as it was carefully curated upon ingestion, and providers can frequently control what data they will accept. Workloads are often static. Applications are designed to run on a limited number of fixed components, many of which were built within the last decade. In general, many of these organizations operate relatively unburdened by legacy systems and data sets.

By contrast, large corporations and government agencies don’t have the luxury of living “legacy-free.” Data holdings are anything *but* clean, and adversaries aren’t going to comply with requests to standardize. Perhaps most notable, while successful analytical processes are critical to many “Web 2.0” businesses, the level of mission criticality varies a bit from traditional IT definitions. Missing trends related to spikes in Cheetos® purchases or failing to stream the latest Harry Potter movie differs from the ramifications of botching financial transactions, disrupting the operation of critical control systems, or failing to connect the dots between bad actors.

That’s not to say that Cloud Computing technology or providers are unreliable, and as time goes on our worlds will continue to move closer together. But assumptions must be understood, risks weighed, staffing and support factored into the cost equations, and the threat of orphaning measured. Some Cloud technologies won’t exist within acceptable risk parameters, and some amount of adaptation might be required.

The Way Forward

Critical questions still remain. How do we effectively manage risk in these new environments? What is the optimal business model for supporting open source software, and which companies will get it right? Can you realize true cost savings without changing the labor part of the equation? Are Cloud Computing and Virtual Desktop Infrastructure (VDI) simply the second coming of the mainframe? And will we ever agree on terminology?

As we look ahead, we believe there are at least four things organizations should consider when weighing their Cloud options: *1) Factor in the assumptions and new realities.* The use and dependence on open source technology, changes in labor force dynamics, and assumptions about workloads and use cases remain critical to the assessment process. *2) Consider the labor impact of a technology area.* Salaries and skill sets are different with many of these technologies, and there are far fewer qualified professionals in emerging technology areas. *3) Identify and assess the security controls that are the most critical.* Security controls are important in any computing environment, but focusing on quality vs. quantity is essential. *4) Question the key drivers.* Cost is often touted as the primary driver for Cloud initiatives, and there is no doubt that, in some situations, reductions are achievable. But not in all. The IT road is littered with technology transitions that delivered improvements but didn’t actually reduce costs. Understanding the Total Cost of Ownership (TCO) remains essential.

In the coming year, In-Q-Tel’s efforts in the areas of Cloud Computing and Next Generation Infrastructure will remain focused on identifying essential Cloud “enabling” technologies, studying successes and failures, cooperating with technology startups to develop optimal business models, and working to identify what is accurate, and what is not, on this new frontier. There is an opportunity to leverage some of these new technologies and models, but we must remain smart about not just the *what*, but the *when*, the *where*, and the *how*. **Q**

Greg Shipley joined IQT in 2010 as a Technology Vice President within the Information and Communication Technologies Practice, where he is responsible for IQT’s Cloud and Next Generation Infrastructure investments. Greg also helps guide IQT’s investments in information security areas. Prior to joining IQT, Greg was the founder and Chief Technology Officer for Neohapsis, an industry leader in information security and IT risk management. Greg also ran the Chicago test lab for Network Computing magazine, was a contributing editor for Information Week magazine, and spent over a decade testing and reviewing technology on behalf of Fortune 500 companies.