Spotlight

Hey Amazon! Mind the GApps

Analyst: William Fellows
Sector: Enterprise Software »» / Storage & Systems »»
Date: 10 Apr 2008

Cloudy forecast

Few, if any, conversations about cloud computing take place without reference to **Amazon's** EC2, S3 and SimpleDB – components of Amazon Web Services (AWS). The cloud category exists because of AWS. So what is cloud computing? The essence of it, says Amazon, is that you don't know where stuff is happening. EC2 and S3 are compute and storage, respectively, in the cloud.

The 451 Group's view of cloud computing is one that allows 'execution venue' (where a transaction takes place) to be part of a cloud-computing service-level agreement. For many organizations – due to regulatory, compliance or taxation – location matters. As far as we are concerned, the cloud, briefly, is IT infrastructure as a service. An expanded view is that the cloud is a combination of abstraction at every layer, the automation of many management chores, and the 'third way' – i.e., computing that is available neither wholly in-house nor entirely outsourced, but located at the optimal point between the two. The idea is that the user doesn't have to know or necessarily care. The experience delivers on expectations in terms of information, functionality and processing quality.

Above all, we see cloud computing as a new way for service providers to offer IT, and a new way for enterprises and other entities to consume it. As such, monitoring, metrics and chargeback – as well as grids, clustering, virtualization and automation – are critical enabling technologies. Utility is a precursor. Public cloud models like AWS are envied, if not emulated, now by enterprises seeking similar benefits for their own IT, and likely being used by skunk works and test bed projects internally. Moreover, shared infrastructures built on top of virtualized environments with utility economic models are a sweet spot of activity for enterprise IT end users in the financial services industry (FSI).

As such, there's an argument to be made that the notion of private infrastructure clouds is as real as public clouds. After all, what do you call something operated within an enterprise or organization that is, in every respect, identical to a cloud – but that runs on an intranet?

Google App Engine

With all the momentum AWS has gathered over the last six of years (300,000 developers and still in beta), we were beginning to wonder if anyone else should bother ... 'when all at once I saw a crowd' and presto, Amazon is joined in the Web-based application hosting business by **Google**, which has launched a beta version of App Engine.

The key to Google App Engine is that it gives developers access to the same building blocks that Google uses for its own applications, though initially they can only do this by packaging their wares in Python code (the same as Google uses). Users' applications will run in Google's software stack using Google's file system, and can use Google's email interfaces and Google Accounts. Specifically, Google AppEngine offers dynamic Web serving; persistent storage via BigTable

(database access) running on its GFS file system (data store services) with queries, sorting and transactions; automatic scaling and load balancing; Google APIs for authenticating users and sending email; and a local Python development environment for Windows, Linux or Mac OS X. It doesn't support Ruby or PHP at this point.

The big reach beyond AWS and other hosting models is that App Engine takes care of configuration management by checking request logs, inspecting the status of the application, pushing new versions of it, monitoring it and integrating all the pieces. Trying to take away the configuration management from developers is a big challenge, but for small businesses, it removes a layer of complexity. Enterprise users report too much complexity, too many suppliers and too few capable management tools, but are unlikely to cede control to Google to this degree. Not surprisingly, the 10,000 free beta App Engine accounts that were up for grabs went in seconds.

Thunderheads: AWS vs. App Engine

Taken altogether, the Amazon AWS components represent a loosely coupled architecture. Developers can choose to tie together EC2, S3, SimpleDB – or not. S3 is widely used as a backup and recovery system. Google's approach is more inflexible (some developers have said 'armtwisty'). BigTable can't be used independently, for instance. App Engine does provide developers with the full effect of leveraging its stack. The developer must be willing to run its entire stack on it; i.e., hand Google a URL along with their code, and trust Google to do everything else. Unlike AWS, it isn't a virtual machine provider and doesn't give the user a raw virtual machine.

Quite simply, Google App Engine demands control, much more so than does AWS. And this, it seems, is the point. Small companies will be able to more easily get into the Web applications business in return for handing over control of their applications. They will depend on Google. Which begs the question: Just how much control are companies willing to cede? App Engine is designed to find out. Through App Engine, Google will know those companies backwards. If it sees something it likes, what will it do? Buy or build. As some folk have already remarked, given that many businesses would like to be bought by Google, this is probably more a feature in this system than a bug!

So AWS and App Engine are cheap and flexible and don't do everything – and they don't care. A flyby view suggests App Engine is going to appeal to folks who want to get applications up quickly – as a rapid application development tool – while AWS has more appeal to traditional IT shops looking for more general-purpose hosting venues. It has more potential on-ramps for these users. The point is their simplicity, flexibility and economics are going to increasingly appeal to datacenter managers looking to take out IT cost. Enterprises such as the **The New York Times** and the **NASDAQ** are AWS users. Enterprise end users in the FSI sector look enviously at models they'd like to either use (but currently can't, due to regulation, organizational issues like ownership, or because of legacy IT requirements) or would like to replicate internally. With Google joining the cloud, the power window for other players just got smaller by a huge factor. They'll need to be highly differentiated to get a start.

Grid computing comes in from the cold

Google's App Engine provides automatic scaling and load balancing but is not, Google insists, 'a grid computing solution.' It argues: 'We don't run arbitrary compute jobs. You don't reserve resources, or machines, or RAM or a number of CPUs, or anything like that.' Amazon has never touted grid credentials as such either, though it has consistently and pragmatically reached out to grid users by addressing the various industry venues where grid folk gather. The basis of its success, it says, is straightforward: 'managing growth.'

So neither Amazon nor Google characterizes their offering as grid, or indeed grid-based, but it's

hard to see them as anything radically different. If the vision of grids is obtaining on-demand computational resources from distributed sources in the same way we get electricity today, then Amazon and Google walk, talk and act like grids. The problem is grids' longtime, seemingly hard-wired (but largely unfair) association with the worlds of science and academia, of high-performance computing (HPC), costly kit and Monte Carlo risk analysis, and computational fluid dynamics. But virtualization and 'gridonomics' have seen users moving beyond these HPC fiefdoms and deploying grid techniques for service-oriented infrastructures. Cloud computing is the realization of those dreams.

Clouds of hosters

Savvis, **Rackspace's Mosso**, **Terremark**, **Joyent**, and smaller players such as **Xcalibre Communications** with its FlexiScale, are using virtualization to support managed hosting customers on shared infrastructure. They leverage virtualization to offer infrastructure services that are flexible enough to be billed according to usage. **Verizon** and **SunGard** offer servers dedicated to a single enterprise customer that may have many virtual machines running on that hardware. **AT&T** has begun to offer shared infrastructure services at the low end, and dedicated virtual machines further up, in addition to its dedicated managed services.

UK provider **Colt Telecom** is building out a dozen new managed services using grids and virtualization on shared or dedicated resources, while **BT's** long-awaited 21C Network is due to be unveiled next month. BT is consolidating its datacenters into a single logical virtual datacenter using **VMware**, and will host customers on it via virtual private LANs.

Should these companies worry about flying into the two thunderheads approaching? Not yet. Becoming a service provider will take time. The hiccoughs that Amazon's EC2 and S3 have suffered are testament to that. Enterprises have (in some cases) regulatory and (in most cases) cultural and organizational issues aplenty to work through. The midmarket enterprise and SMB sector will continue to be the sweet spot initially.

The ecosystem

It will be interesting to see – if Google permits this – what kind of window emerges for third parties to arbitrage App Engine. The Amazon ecosystem, which is positively encouraged by Amazon, includes companies such as **Coghead**, **Relationals** and **Bungee Labs**. Application and software providers including **EnterpriseDB Corp** and Relationals' Longjump use Amazon on-ramps like **Elastra**, **Rightscale** or **CohesiveFT**, which enable software to run in the AWS cloud (and elsewhere).

Microsoft and Yahoo

At this point, **Microsoft's** SQL Server Database Services looks more like a capable Internet relational database than Amazon's SimpleDB or Google's BigTable. The database service for the Internet is due next year. It will synchronize data on multiple devices supporting offline use, and support stored procedures over time. Microsoft is also working on BizTalk Services and has beta versions of Synchronization Framework – both hosted services targeted at developers. The question is, when will we see an overarching hosted application platform: Office in the Cloud anyone? Meanwhile, Microsoft has recently filed a number of patents on cloud computing concepts, which together with its Dryad technology are collectively targeted at ways to parallelize applications without a developer needing to be a concurrency specialist. But patenting cloud computing? Someone's giving bad advice.

Yahoo is a lame duck. If Microsoft bags Yahoo, the question is whether it can commercialize Hadoop and Yahoo's other cloud activities. Yahoo's tie-in with **Tata Group** looks increasingly tough

as a differentiator because of the number of incumbent system vendors already coming at cloud from a grid and supercomputing point of view.

More cloud formations

Other companies that provide an API and SDK to their services include **Salesforce.com's** force.com platform, **OpSource**, **Morph Labs**, Microsoft Office SharePoint Server, **Claranet**, Relationals' Longjump, **Xcerion's** iCloud, **AppJet** and even **Blackboard** and **Facebook**. They have all developed some form of SDK for developers to build APIs on. They depend on members for use and have built their own clouds.

3Tera, which arguably leads another class of cloud companies, provides software for creating virtual private datacenters. 3Tera's AppLogic software has been positioned as technology for enabling companies to get into the software-as-a-service (SaaS) business without having to build any infrastructure themselves. It combines clustered hosting and the utility compute-type of delivery. Now it's also targeted at enterprise end users. It has partnered with managed hosting provider **Layered Technologies**, and more recently storage service provider **Nirvanix**.

Both CohesiveFT, which isn't hard-wired to AWS and enables users to create application and middleware stacks for virtual servers, and Elastra, as it diversifies toward on-premise software, will increasingly bump up against 3Tera.

It's hardly surprising that virtualization vendors **Citrix** and VMware are on board cloud computing. Heck, they're supplying the flying carpet. VMware characterizes the cloud as a service delivery approach that's applicable outside, inside and in-between organizations. Citrix concurs, but expects to become 'big in cloud' this year. It will match VMware's VSSP virtualization licensing scheme for hosting providers later this year. Moreover, it plans to become a provider, supplier to and partner of cloud companies over time.

SaaS isn't in and of itself a cloud, so it remains to be seen to what extent **BEA Systems** (Genesis), Microsoft (Oslo), **Oracle** (Fusion 11g) and **SAP** (NetWeaver 7.1) will reach for cloud computing in their next-generation enterprise SaaS offerings. They're all tilting that way. Storage-in-the-cloud companies are numerous and include **Box.net**, **EMC Mozy**, **Dell** (DataSafe), **Hewlett-Packard** (Upline) and JungleDisk. Google Gears and **Adobe** Air are examples of applications that will run in an offline mode on workstations or laptops.

Private / public cloud vendors

Cisco Systems' cloud vision is centered around Data Center 3.0, a product plan to transform the datacenter into a virtualized environment providing anytime, anywhere access to content on any device. It's got a bunch of new switches supporting this vision that can dynamically provision virtualized resources.

Dell Datacenter Solutions has announced its intent, but has yet to bring its first private infrastructure cloud computing offerings into view. HP's Data Center Transformation portfolio is a suite of products and services aimed at helping companies modernize and consolidate existing datacenters, or shift operations to a cloud-based service hosted by HP. HP's AIaaS (Adaptive Infrastructure as a Service) lets customers host applications in HP datacenters optimized for Microsoft Exchange and SAP. Hosted at two HP datacenters, this EC2-like facility is available only to existing outsource customers at this point. HP's Flexible Computing Service also offers some interesting hosted and managed in-house models, while cloud services are one of the consolidated R&D projects over at **HP Labs**.

IBM appropriated the term cloud early, and its branded products – such as Blue Cloud (BladeCenter plus Tivoli software) – are targeted at enterprise end users seeking to create private

cloud infrastructure. At this point, IBM's cloud activities at this level are more 'marketecture' than a service offering, however. IBM has also got numerous academic and research projects in hand including those with Google, Wuxi in China, the **Georgia Institute of Technology**, **Ohio State University** and its own Dublin Blue Cloud initiative.

Sun Microsystems' $1-per-CPU-hour Sun Grid Compute Utility (aka network.com) has been around for some three years now, and has about 50 apps (mostly open source) prequalified. Since it's been refocused on SaaS, there's been a revival of interest in it. Meanwhile, Sun is also working on Project Caroline, a hosting platform for development and delivery of dynamically scalable Internet-based services – aka cloud computing. **Fujitsu Ltd**, meanwhile, has tapped **Akamai** to help build its platform for cloud computing. Other cloud infrastructure suppliers include **Liquid Computing**.

Is cloud computing disruptive innovation?

Cloud computing certainly has the characteristics of being disruptive: it's a cheaper, more convenient, simpler solution that meets the needs of the low end of the market (AWS and App Engine as opposed to fixed-term, complex, take-it-or-leave-it-priced hosting is just one example). Taken together, cloud technologies support a business model that is already creating a disruptive impact. But can it eventually take over incumbents' mainstream markets? The principle of 'the innovator's dilemma' suggests that even if incumbents themselves tinker with the disruptive technologies, ultimately they can only service their existing business lines, and their mainstream markets will eventually be disrupted. The question is not whether Microsoft or other incumbents can see this happening. It's abundantly apparent. Indeed, many if not all of the incumbents are already tinkering with their own cloud operations. The question is whether they can do anything about it.

The answer – at least as far as the enterprise goes – leads us back to the law of disruption itself: that technology disruption happens exponentially while societal change is incremental (linear). The 'law' suggests that this disparity often leads to physical capabilities that far exceed the capacity that societal norms can handle. Our argument is that in the enterprise, the adoption of cloud computing is as much dependent on the maturity of organizational and cultural (including legislative) processes as the technology per se. The typical enterprise IT end user right now does not expect utility model billing across shared infrastructure. Enterprises will take time to accommodate new IT practices. Will this give the incumbents the breathing space they need to change before they are overtaken?

Clearly, this is only the beginning. Look out for The 451 Group's forthcoming position paper on cloud computing, where we'll examine cloud visions, how they are connected and what cloud will mean for users, vendors and investors alike.

Search Criteria

This report falls under the following categories. Click on a link below to find similar documents.

Company: Amazon.com

Other Companies: 3Tera, Adobe Systems, Akamai Technologies, AppJet, AT&T, BEA Systems, Blackboard, Box.net, BT Group, Bungee Labs, Cisco Systems, Citrix Systems, Claranet, Coghead, Cohesive Flexible Technologies, Colt Telecom, Dell Inc., Elastra, EMC Corp, EnterpriseDB Corp, Facebook, Fujitsu Ltd, Georgia Institute of Technology, Google, Hewlett-Packard, HP Labs, HP Labs, IBM, Joyent, Layered Technologies , Liquid Computing, Microsoft Corporation, Morph Labs, Mosso, MozyPro, Nasdaq Stock Market Inc, Nirvanix, Ohio State University, OpSource, Oracle, Rackspace Managed Hosting, Relationals, RightScale, Salesforce.com, SAP, Savvis Communications, Sun Microsystems, SunGard, Tata Group, Terremark Worldwide Inc., The New York Times Company,

Verizon Communications, VMware, Wuxi Tai Lake Industry Investment and Development Company, XCalibre Communications, Xcerion, Yahoo

Analyst: William Fellows

Category: Enterprise Software, Storage & Systems

Sector:
Infrastructure management
Hosted services / Managed hosting
Enterprise networking / WANs

Featured links

451 Market Insight Service

Elastra's on-ramp to cloud is rolling
(28 May 2008)

Merrill's grid, utility... cloud?
Grid computing has underpinned Merrill Lynch's enterprise computing strategy. It's created a utility offering on it, which looks like a pretty cloudy infrastructure on our Cloud Index. (19 May 2008)

It's raining databases
The emergence of cloud computing and Web-based applications has prompted a sudden increase in the number of database options facing enterprises. How can you tell a data query service from an online distributed database? (14 May 2008)

Arjuna targets private cloud systems, with EC2 to follow
The transactional systems specialist locks in cloud computing as its starting position. (9 May 2008)

Etelos reverses onto the public stage as it nurses major platform ambitions
The SaaS apps and development tools player has taken an unusual route to going public via a reverse merger. As more daylight shines on its still-fledgling operations, has Etelos got what it takes in technology and strategy to be around for the long haul? (2 May 2008)

ElephantDrive's online backup is more than a front end to Amazon
The startup has gained a good deal of traction through its early integration with Amazon's S3, but, as it aims to secure its first external funding round, it emphasizes that it's more than a front end to the popular storage service. (21 Apr 2008)

Microsoft tiptoes toward online database services
The company's SQL Server Data Services online database offers hints at a new direction for the software giant, but it is not giving up on on-premise software just yet. (16 Apr 2008)

Google gets itself an enterprise Salesforce(.com) and tight Apps, CRM integration
Once again, the two vendors are deepening their relationship, with the emphasis this time around on hooks into each other's software. How long before those ties result in Google really sweeping Salesforce.com off its feet in a permanent embrace? (14 Apr 2008)

Google powers up App Engine, an application hosting platform for Web developers
Web search giant Google is opening up its internal infrastructure as a platform for developers to build Web applications. Will Google transform the platform as a service market with this offering? (11 Apr 2008)

3Tera presses on, sees SOA as new opportunity

The company has overcome the passing of its cofounder and president/CEO Vlad Miloushev to register impressive early growth and now has its sights set on SOA as a new source of opportunity for its grid computing operating system. (4 Apr 2008)

Verari bolsters its blade server business with renewed push into storage

A new business unit has been formed at Verari to focus the surprising growth in storage revenue it's seen over the past few years. By approaching storage from the server angle, it's avoiding direct competition from EMC and NetApp. (3 Apr 2008)

Cloudsmith wants to be the assembly line in the cloud for software publishers

Writing software is easy, but packaging it may not always be. Cloudsmith wants to provide an on-demand infrastructure to overcome this hurdle in distributing software. (2 Apr 2008)

The Open Grid Forum: Time called in the Last Chance Saloon

The goals and objectives of the OGF have lacked clarity as priorities within the grid community shifted. But has the 'Tiger Team' come up with the right answers? (1 Apr 2008)

Elastra delivers database clustering for Amazon's cloud

The new startup has officially launched its Cluster Server product, promising to bring the potential benefits of cloud computing to the enterprise. (31 Mar 2008)

EMC buys Aussie ITSM software vendor Infra

(10 Mar 2008)

Grid, utility, cloud: Enterprise infrastructure is the initial focus of IBM Blue Cloud

Using Blue Cloud could help manage large-scale datacenters and get better utility – the IBM way. (16 Jan 2008)

Head in the cloud: CohesiveFT now does component assembly for virtualized platforms

From Financial Technology virtual appliances to Flexible Technology component assembly, Cohesive's Elastic Servers target application virtualization on the datacenter. (25 Oct 2007)

---

451 TechDealmaker

A Rackspace redo: Hoster eyes IPO (again)

If at first you don't succeed... build a $360m business that's profitable every year and attracts tens of thousands of customers... then try again to hit the public market. And do it through a Dutch auction. (1 May 2008)