

Amazon Web Services

Already the premier online shopping destination for everything from toilet paper to televisions, many would argue that Amazon could comfortably continue its operations strictly selling tangible goods and never think twice about entering other markets. Jeff Bezos, however, has more ambitious intentions. Bezos, the Princeton computer science alum that founded Amazon in 1994 intending to build “Earth’s largest bookstore” has recently gone on to build one of Earth’s most robust ‘cloud-based’ on demand computing platforms: Amazon Web Services (AWS). Launched in 2006, AWS has allowed individuals and enterprises alike to leverage Amazon’s existing network architecture and software expertise to quickly and easily host their own applications and services on Amazon’s hardware. Whereas traditionally businesses would have to accurately forecast system usage and in turn purchase, configure, and install a set number of servers, AWS allows users to share core system resources with thousands of other users on demand, as needed.

As Amazon was the pioneer in inaugurating online commerce, it too has become the pioneer in introducing a paradigm-shift from traditional, self managed IT services to a dynamically sized “cloud computing” model¹. As this idea becomes increasingly adopted by existing IT companies and startups begin to try to replicate Amazon’s success, AWS must in many ways differentiate itself from competitors in features and actively build a reputation in the IT industry as being a serious, long-term contender.

Amazon’s History

Amazon, founded in 1994 originally as an online bookstore in founder Jeff Bezos’ Seattle garage, has matured to a \$79 billion dollar company² revolutionizing the way customers research and buy goods online. In 1999, Amazon introduced its sales platform as a service to other retailers wishing to outsource the laborious and overwhelming notion

¹ Robert Hof, “Jeff Bezo’s Risky Bet,” *Business Week*, November 13, 2006
<http://www.businessweek.com/magazine/content/06_46/b4009001.htm>

² In market capitalization, as calculated on December 2, 2010

of building their own web interface and replicating Amazon's industry leading fulfillment system (**Exhibit 1**). Amazon has gone so far as to even allow users to sell their used items on the Amazon website – alongside Amazon's new items. By charging commissions and flat fees for use of these services, Amazon became not just a competitive eCommerce retailer but a go-to eCommerce marketplace all together.

Amazon has since moved on to enter industries ranging from digital media distribution (Amazon Digital Downloads), web search (A9) to even human resource markets (Amazon Mechanical Turk). Essentially, whenever Amazon believes it can improve on the status quo, new industries are entered without hesitation.

Amazon Web Services

Amazon Web Services was born of the idea that network infrastructure is becoming an increasingly commoditized service and should behave as other utilities do, be they water or electricity. This means that such services should be available at any capacity on a moments notice and without necessarily forecasting demand. Amazon meets this expectation in both of its key AWS products. Amazon's Elastic Cloud Computing (EC2) platform allows applications to run on an instantly scalable number of processors on demand, while Amazon's Simple Storage System (S3) allows access to a practically infinite allocation of disk space on demand. Where as projects must normally anticipate a set amount of demand and invest in infrastructure capacity sufficient to accommodate the maximum feasible demand, even though it may only be utilized 2% of the time for instance, the Amazon EC2 platform allows applications to use as much processing power as they need at any given time, scaling up and down parallel to demand. Similarly, while conventional projects may dictate that a maximum capacity of storage be available at any given time, regardless of utilization, S3 allows applications to scale storage needs exactly in parallel with demand. Users never pay for more than their share of occupation. This is all possible because of the implementation of a seemingly infinite, shared resource pool.

These products grew out of Amazon's existing demand and expertise in information technology (IT) and are therefore immediately suited for extremely large

scale applications, while inherently ideal for smaller projects as well. While most IT hosting companies charge up front set up and usage fees together, and maintaining one's own servers requires tremendous up front capital expenditure (capex) and ongoing support costs, Amazon began AWS by charging directly in proportion to usage (Amazon EC2 charges anywhere from \$0.10 to \$0.80 per processor hour³ while S3 charges up to \$0.14 per GB per month of storage, with bandwidth costs of \$0.10 to \$0.15 per GB of bandwidth downloaded or uploaded⁴). This inexpensive, pay-as-you-go price scheme eliminates the risk associated with investing in technologies never before tested, encouraging system administrators and curious programmers to play with the service at extremely low costs. Even before its 'official' launch, public betas were promoted to programmers and system administrators to play with the services at no cost, encouraging users to build applications into AWS and later lock-in users who couldn't easily rebuild their projects elsewhere.

The 'Cloud-Computing' Market

In sum, the 'cloud computing' market, as defined by Gartner research, will be worth an estimated \$68.3 billion in 2010 alone, up 16% from last year's \$58.6 billion. Gartner Research further estimates that an additional \$112 billion will be invested in the next four years by companies moving their existing operations into the cloud. While these figures envelop many different types of services (say, Google Docs and the Salesforce.com Customer Relationship Management software), it is a safe trajectory from which we can gauge the direction of the Platform-as-a-Service (PaaS) industry: growing at a steady and dramatic pace outwards.

Since Amazon first entered this market in 2006, many refined IT companies have quickly followed suit. Some immediately differentiated themselves as high-level niche services while others developed similar, low-level network clouds. High-level services imply that the user is developing only application software, while server software (such

³ Amazon EC2 Pricing <<http://aws.amazon.com/ec2/pricing/>> as accessed on December 3, 2010

⁴ Amazon S3 Pricing <<http://aws.amazon.com/s3/pricing/>> as accessed on December 3, 2010.

as which model of Linux to operate) is considered low-level⁵. Salesforce.com entered the market with a high-level platform catered to internal business operations called *force.com*. Based on a pre-built software ‘stack’, *force.com* allows novice programmers to quickly assemble applications for internal business processes or buy and customize existing applications available in the *force.com* marketplace. Similar to AWS, all applications are run on Salesforce.com’s externally managed network infrastructure and immediately scale to demand.

While the Amazon EC2 platform allows offline computation using a seemingly infinite pool of processors available on demand, Google’s solution instead simply offers a ‘host’ in which users can upload software code and have them run like regular websites, but on Google’s servers. This host, Google AppEngine, is only available to operate 2 programming languages and automatically dictates how much power is required as the website operates and bill accordingly. After passing a threshold of free usage, users are charged \$0.10 per CPU hour, \$0.10 per GB of bandwidth incoming and \$0.12 per outgoing GB. All together, storage is charged at a rate of \$0.15 per GB⁶. Unlike AWS, customization opportunities are very minimal, making it only suited for customer-facing web-applications.

Microsoft, whose Windows Server software held a controlling 73.9% market share in fourth quarter of 2009⁷, clearly has a lot at stake if they were to not enter to PaaS market. Offering a suite of services similar to AWS, Microsoft introduced the Windows Azure Cloud Services Platform in February 2010 in absolute direct competition of AWS. Compute instances are charged at a rate of \$0.12 per hour, with storage costing a fixed rate of \$0.15 per GB and bandwidth costing \$0.10 in and \$0.15 out per GB⁸. Unlike AWS, Microsoft does not offer structured volume discounts for scaled use, but instead offers quantity discounts for monthly to yearly subscriptions. Because this has to be paid up

⁵ Armbrust, Michael, “A View of Cloud Computing,” Communications of the ACM, April 2010, <<http://cacm.acm.org/magazines/2010/4/81493-a-view-of-cloud-computing/fulltext> >

⁶ Google AppEngine Blog, < <http://googleappengine.blogspot.com/2009/02/new-grow-your-app-beyond-free-quotas.html> > as accessed on December 3, 2010

⁷ Mary-Jo Foley, Behind the IDC data, *ZDNet*, < <http://www.zdnet.com/blog/microsoft/behind-the-idc-data-windows-still-no-1-in-server-operating-systems/5408> > as accessed on December 3, 2010

⁸ Pricing, Windows Azure, <<http://www.microsoft.com/windowsazure/compute/default.aspx>>

front, Microsoft seems to be relying heavily on their existing brand goodwill in the server industry, hoping users will demonstrate immediate trust in their service.

Performance and Differentiation

In first introducing AWS, the product was launched as a free, unlimited public beta to encourage experimentation and adoption⁹. Because no immediately comparable products existed until the introduction of Windows Azure, Amazon enjoyed excited adoption by startups and established companies large and small. Startups like Animoto, an online photo to video-slideshow building website, built themselves entirely on AWS and were able to elastically manipulate anywhere from dozens to thousands of computing instances on the fly, as demand shrunk and grew uncontrollably. As the graph in **Exhibit 2** depicts, with computing instances on the Y axis and time on the X axis, deploying an additional roughly 3300 processing cores on 2 days notice (from 4/16 to 4/18) would be virtually impossible unless they were purchased and installed far beforehand. Similarly, AWS recently hosted live video feeds of the birth of a baby elephant at a zoo in Belgium. As **Exhibit 3** illustrates, AWS was able to produce processing cores on demand as activity at the zoo fluctuated throughout a 2 day time period¹⁰.

Not all AWS operations have to be customer facing, however. Many applications for EC2 have been in research or operational capacities. Pharmaceutical giant Eli Lilly and Company, for instance, recently began using EC2 to “quickly and cost-effectively process tremendous amounts of research data.”¹¹ More recently, Netflix has moved away from managing its own datacenter to putting its own applications on EC2. As Netflix Cloud Architect Adrian Cockcroft explained, “most everything that scales with customers” will be moved to AWC, including “encoding movies for streaming.”

Likewise, plenty other examples have surfaced where in AWS has been able to dramatically cut overhead, shift expenditures from capital to operating expenses (which

⁹ Inside Amazon’s EC2 < <http://www.zdnet.com/blog/btl/inside-amazons-ec2/3541>>

¹⁰ EC2 and Wowza Media Support, *Amazon Web Services Blog*, <<http://aws.typepad.com/aws/2009/05/ec2-and-wowza-media-support-belgiums-largest-live-streaming-event.html>>

¹¹ News Release: AWS launches Amazon EC2 for Windows, < <http://phx.corporate-ir.net/phoenix.zhtml?c=176060&p=NewsArticle&id=1216597>>

would maintain strategically higher Return on Assets figures and better match cash expenses to cash revenues), and allow margin for unpredictable traffic.

Amazon has managed to differentiate itself in this industry by not only being the first to market, but also by offering incredible incentives to join (with an introductory ‘free tier’ pricing strategy, allowing easy and costless registration and experimentation, potentially leading to lock-in), a product set that easily outcompetes traditional web hosting and self-hosting schemes, and the ability to operate on a low ‘system level,’ allowing EC2 to accommodate an incredible range of technologies and operations.

Amazon’s Web Services Future

Although market share data does not yet exist for the vaguely defined ‘cloud-computing’ industry, it is clear that when it comes to elastic computing Amazon holds a sizeable lead based on the caliber of companies currently using the system. In last year’s annual shareholders meeting, Amazon displayed an impressive list of companies ranging from ESPN to the University of Oxford who currently employ AWS (**Exhibit 4**). By maintaining this loyal customer base and continuing to offer attractive ‘beginner’ AWS accounts for free, Amazon will insure that its grasp on the utility computing market will not be relaxed.

While Amazon began as a simple eCommerce website, it should continue delivering value in any market it sees it could make a material difference in. As it has literally defined the Platform-as-a-Service industry, Amazon should continue investing and innovating in AWS while also exploring other internal business operations that could be similarly abstracted and provided at a cost to other customers.


Exhibit 1:

[Privacy + Security](#) | [Cookies + Targeted Advertising](#) | [Terms + Conditions](#)
[California Privacy Rights](#) | [About This Site](#)

© 2010 Target.com. All rights reserved.
The Bullseye Design and Bullseye Dog are trademarks of Target Brands, Inc.
Powered by Amazon.com

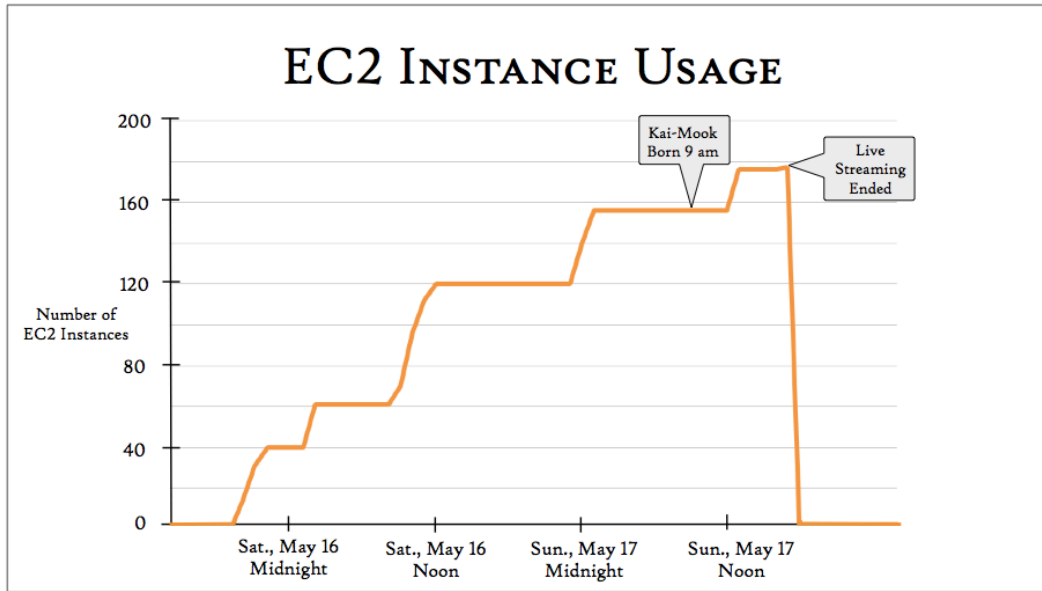
Source: Target.com

Exhibit 2: Animoto Instance Usage



Source: Amazon Web Services Blog, <http://aws.typepad.com/aws/2008/04/animoto---scali.html>

Exhibit 3:



Source: Amazon.com presentation, "Annual Meeting of Shareholders," 2009

Exhibit 4:



Source: Amazon.com presentation, "Annual Meeting of Shareholders," 2009