



Participant Questions and Answers

1. Can we expect SLAs/availability to be the same from dedicated to public cloud environments?

Not in the near term, though it depends on the internal SLA reference point which can obviously vary. The SLA for Amazon EC2 is for 99.95% uptime during the service year, which may be appropriate for some enterprise workloads, but historically not for 'mission-critical' applications. There is still a fair degree of variability in SLA terms across public cloud vendors, and no consistent market expectation has emerged as of yet. Instead of focusing on SLAs, many enterprises are starting to look to architect redundancy in their cloud solutions, so for example if there is an outage at one service provider data center, workloads automatically roll over to a second data center or service provider.

2. If server utilization on dedicated environment is over 80% utilization, does it make sense to virtualize?

Probably not, at least from a workload cost perspective. The 'best-in-class' utilization for a private cloud we see is around 60%, and as mentioned in our presentation 15% is about average. There may be other mitigating factors that might drive one to virtualize in an 80% utilization scenario, but cost would probably not be one of them.

3. How can we achieve transparency with AWS?

The answer depends on the type of transparency to which you are referring. AWS provides fairly clear line-of-sight into pricing, usage and billing. The bigger challenge for enterprises achieving transparency with AWS is around forecasting future consumption and usage.

4. What are the factors that are going to drive down the cost of public cloud services?

The largest factor that will drive down the cost of public IAAS services is vendor competition. Existing cloud service providers are enjoying attractive margins. With new IAAS service providers entering the market on a weekly basis, vendors are likely to drive pricing down to capture and retain customers. In addition, as more enterprises migrate to the cloud, service providers will gain additional scale economies that they will be able to pass through to customers as well.

5. Slide 12 - what is the purpose of showing the suppliers here? Why these suppliers? Is MS a viable solution for IaaS?



Participant Questions and Answers

Microsoft Azure is a Platform-as-a-Service (PAAS) solution that seeks to deliver customers the ability to manage and optimize workloads across differently delivery models (public / private / hybrid), and deliver the economic benefits we describe. While it operates at a different level in the stack than pure IAAS services, we presented it as a potential future option for capturing hybrid bursting economics.

6. It appears your view of private cloud is that it can only exist in a dedicated environment. Is this true or do you also believe it can exist in a multi-tenancy environment, similar to the public cloud.

Private clouds can exist both in on-premise and hosted models, but by definition they are dedicated to a single customer. Some service providers are beginning to offer private cloud services based on 'converged infrastructure' or standardized private cloud 'building blocks' (see VCE) that enable the service provider to offer hosted private cloud services, but with some of the management cost efficiencies seen with multi-tenant public cloud services.

7. Private cloud allows control of data security - how such addressed with public cloud environment? Are they robust enough to meet bank requirements?

It depends on the specific banking requirements to which you refer. In public cloud models, customer data resides in shared, multi-tenant environments, which creates issues and conflicts with some regulatory mandates and industry standards. With private cloud models data remains in a dedicated, single customer environments. It is important to note that public cloud vendors are actively and aggressively working to provide new features and capabilities that help address regulatory compliance issues in the US and abroad.

8. On p. 15, you mention Amazon pricing is going down. Are other providers following the same price trajectory downward? How much are prices dropping? How quickly are they likely to drop going forward.

We believe competition for like services will drive down pricing. Amazon is seeing new entrants come into the market with different and many times lower pricing than Amazon, and as such, either proactively or reactively, adjust pricing to meet market conditions. Tata Communications entered the market last October with a very competitive price structure that included free inbound data transfer.

That being said, continue to watch as each cloud supplier jockeys for position and competitive advantage. For instance, Amazon has introduced dedicated instances and as such, priced that to the market, which is more expensive than their on-demand EC2 instances. Once a supplier enters the market with a different value



Participant Questions and Answers

proposition, could be security, compliance, etc., that pricing will not compete with that of Amazon EC2.

How quickly prices will drop going forward will be determined by competitive pressure and the suppliers' ability to obtain more competitive cost structures.

9. Isn't Amazon more of a "private" supplier versus a corporate IaaS supplier? How is Amazon's offering reducing costs as related to corporate IaaS suppliers?

By definition, (and these definitions can vary based on who you ask) Amazon is a public (multi – tenant) IaaS supplier. Only recently did they introduce dedicated instances, which moves them much closer to providing "private" (single tenant) cloud instances.

Savings are realized because a client only pays for the Amazon instance when they have it turned on, and there is no billing when they turn it off. While per unit cost of owning vs. renting from Amazon may be similar, or even more expensive for Amazon, the fact that you only pay when you are using it drives down the total cost.

Thus, this pricing methodology is exceptionally attractive to those use cases that only need it once in a while. Using the cloud, a customer does not need to physically buy a server, maintain that server, manually provision it for each project, etc.

Application test and development are two of the initial use cases that have beaten a path to the public cloud door. They can spin up a VM in a matter of minutes, use it for the required period, and then simply shut it down when they are finished.

Also, there are many more advantages than simply the pay per use economics. The ability and convenience to spin up a VM with the appropriate OS size provides time utility that customer do not have in-house or with an outsource provider.

10. What are the challenges facing enterprises when contracting for these services? How do they best protect themselves?

It is important for all enterprises to take an inventory of their needs when moving applications to the cloud, recognizing that not all application will have the same requirements. This inventory will lead you down a path to specific cloud providers and drive the price points. Generally speaking, clouds are created for specific use cases. Not all clouds can or will meet the enterprises requirements. Know your sacred tenants, and if they cannot be met, don't move those workloads to the cloud.



Participant Questions and Answers

Once needs are understood, it will be important for the enterprise to do their homework, or hire a consultant firm, to identify those cloud provider(s) that best meet their needs.

If your needs lead you to a path that allows you to use a public or community cloud, take them for a test drive. You can most likely start out in a pay per use methodology. Once you are satisfied, then you can lock-in lower pricing by signing up for a 1 or 3 year (in the case of Amazon) contract.

11. Use of cloud computing frequently involves reliance on relatively immature IT vendors. How should organizations manage the enterprise business risk of cloud computing?

On April 21 of this year, this question was displayed very prominently. That is the day that Amazon's east coast data center all but shut down. However, not everyone who had VM's in that data center was stymied. The key for those that were not negatively affected, such as Netflix, was that they designed for failure.

'Why were some websites impacted while others were not? For Netflix, the short answer is that our systems are designed explicitly for these sorts of failures. When we re-designed for the cloud this Amazon failure was exactly the sort of issue that we wanted to be resilient to. Our architecture avoids using EBS as our main data storage service, and the SimpleDB, S3 and Cassandra services that we do depend upon were not affected by the outage.'

Source: The Netflix "Tech" Blog

In designing for failure, it is important to match the need with the effort. Business critical applications require one effort, where test and development require another.

For more, the article below from zdnet provides an overview of protecting yourself in the cloud.

<http://www.zdnet.com/blog/projectfailures/cio-analysis-examining-amazons-cloud-failure/13152>

12. What is gross savings range of the 45% savings for burst management? What % of demand is bursty?

The 45% (actually I believe we have 40% in our presentation) savings is an estimate of the saving an enterprise can achieve (versus virtual environments) if they are able to burst peak demand to the cloud. In both our analysis, and in the eBay example, these savings are very consistent. The savings come from downsizing the base environment (think hardware and related expenses) because you are now able to



Participant Questions and Answers

burst peak demand into the cloud. In the eBay example, they estimated going from 1900 servers, down to 800 for the same workload, then burst peak demand into the cloud. They estimated monthly savings of \$1.1M due to the reduction of base infrastructure, and a peak bursting cost of only \$36k per month.

The height and duration of the peak that is to burst to the public cloud will significantly alter the savings potential. Generally speaking, if you have an environment that is running at <30% utilization, it may be very possible to achieve savings in the 35-50% range.

13. Slide 18 - is the implication here that customers should expect a contract structure where their pricing decreases over time as the supplier increases their ability to provide their service at a lower cost?

Not necessarily. The implication of this slide is that there are multiple ways by which you can burst a workload, and that over time, it will become more automated and easier. As the technology improves and application migration costs are driven down, this will lead to more bursting and the realization of the cloud economics.

14. Seems this reasoning needs one more layer. If my business process can go to the cloud and my appl will run on a cloud infrastructure, why not convert it to a SaaS solution for additional gains - if they exist?

SaaS solutions are the wild card in this equation. With the possible exception of CRM and Salesforce.com, it is still unclear if other types of applications will enjoy SaaS success. We consistently hear of the desire to move work to a SaaS environment, but see varying degrees of success. Netsuites, an accounting and ERP SaaS application, has enjoyed a great deal of success, however you can argue that it has not seen the same level of success as Salesforce.com. We believe there are friction points that inhibit success, including such items as the level of integration required between other applications, security concerns and application criticality.

15. Based on the cost models / savings presented, is cloud a more suitable application delivery model for green field projects or for more mature organizations?

It would make sense that green field application projects assume the cloud is where they need to focus. The benefits are too great to ignore. So yes, you will see a great deal of cloud application come from the green field projects.

That said, it does not mean one should ignore the more mature applications. The analysis provided by both Everest, and eBay projections would suggest that there is an attractive business case for transforming mature applications. As packaged



Participant Questions and Answers

applications reach their end of life, the newest version will be cloud enabled. Thus, in the next 4-6 years, it is expected that you will experience an explosion of cloud architected applications and options.

16. When using the cloud specifically to cover peak load, how do you move the related data to the cloud in time?

Enterprises are doing this today with peak / event based workloads. To meet the needs of a large event, could be a large sporting event, or ticket sales for a world tour, those companies are taking that event to the cloud. They are manually loading the VM, and related data to the cloud with the appropriate size to manage that event. After the event is over, they take down the VM, and off load the data. It is a very manual solution, but it is working and saving these companies large amounts of money in the process.

Other examples including running your applications in one cloud environment and using the cloud for your DR. In this situation, you may be able to have your VM loaded in your steady state environment, as well as your cloud environment. If your DR solution includes mirroring the data, then it will be continually loaded in the two locations. As your dedicated environments approaches the threshold, you begin to off load workloads to the public cloud, where the data already exists. Companies such as RightScale and Enstratus have and are developing such solutions with some proven use cases.

Being able to dynamically burst your application from one environment to another is the Holy Grail of cloud computing. Everest Group, Ebay, and others are working on understanding those solutions. The technology is very close, if not already there, with some solutions to overcome latency concerns.

17. What are typical workloads that are "early movers" to the cloud? e.g. Analytics?

The workloads that we have seen move to the cloud are generally those that are project based, with defined timeframes and compute needs. Those include:

1. Analytics
2. Development
3. Testing
4. High intensity computing
5. Batch jobs

Other workloads that have jumped directly to a SaaS model include CRM, and to some extent, ERP (Netsuites).



18. Who/what will drive the wide scale adoption of cloud - infrastructure providers like Amazon, low costs or Enterprise apps and system software companies like SAP, Microsoft?

The enterprise/customers' demand for the cloud economics will drive both the system software companies as well as the infrastructure service providers to mature solutions for wide spread adoption. Keep in mind that the lines between the two are becoming quite blurred. Microsoft is clearly entering the infrastructure service provider space with Office 365 and Azure. HW providers are competing against their traditional customers, and service providers are acquiring the technologies to enable cloud infrastructures. In addition to that, the enterprise is working with both to create solutions to enable the economics they so desire, as is the case with eBay betting on two horses, eBay and Rackspace.

19. Further to this last question, how are enterprises managing risk around "limited liability" and "no 3rd party right to audit"?

To truly realize the benefits of cloud computing, an enterprise must be willing to use a standardized, shared environment along the way. That environment is engineered to meet specific needs. Based on the way it is engineered, the provider may not extend favorable liability clauses etc. Most (if not all) public clouds today are you have a right to use, according to the suppliers terms of use, which means you can take it or leave it.

Today, you will find providers that will provide more "outsourcing" like terms for liability and right to audit in the hosted private cloud space. you should expect to see similar terms when they build community clouds, built for a specific purpose in mind.

20. Not sure of results of eBay case study. If 25% of cost is variable, how can they reduce overall costs by 40%?

This is the exact dilemma that all traditional data centers face - How do you make a highly fixed cost environment variable? Remember this fixed cost environment is made of a great deal of excess capacity to accommodate peak workloads. What eBay is working on with Rackspace and Microsoft is the ability to take that peak capacity and "burst" it into the public cloud. Once they crack that code, eBay will be able to physically reduce the size of their compute environment. In the case they presented, eBay wants to go from 1900 servers for the selected workloads, down to 800 servers. This will save them \$1.1M/month, (mostly fixed costs that they no longer need) while their public cloud expense (for that peak load) will go to \$36k per month (which becomes variable costs). eBay must not feel they can do this with all workloads, else they would get significantly greater savings than 40%.



21. Has there been any instance where the enterprise size their DC for the lowest load and then take the peaks using a private cloud. Has that been successful or are there drawbacks in this model?

We have not come across where this has happened yet. However, this is exactly what enterprises and suppliers alike should do, and are working towards. In the eBay example that we provided, they were working with Rackspace and Microsoft (in the words of Neal Sample, betting on two horses) with eight fairly large workloads to enable them to significantly reduce the size of the operating environment of these workloads. In the presentation, they wanted to reduce the compute environment from 1900 servers to 800, a \$1.1M/month savings, while incurring a public cloud bill of only \$36k/month. And that is just how disruptive cloud computing can be.

22. Who is Ebay bursting into? I think I have heard Azure....to the point where ebay wanted an Azure cloud on their premise

In the example that was provided in our webinar, eBay said "they were betting on two horses", working with both Rackspace and Microsoft. It is likely that eBay is also doing a great deal of other work with both of these providers.