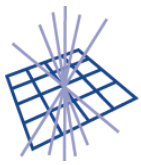

GridPP Project Management Board

GridPP – COMMERCIAL CLOUD COSTS

Document identifier :	**GridPP-PMB-169-GridPP5_CommercialCloud Costs**
Date:	**10/2/14**
Version:	**2**
Document status:	**Final**
Author	**AH,DC**

GridPP 5 Cloud Computing Costs Comparison

In addition to the ongoing development work on running GridPP workloads using a private cloud, it's a reasonable question to ask what the costs would be to run those same workloads using public clouds. The firs thing to mention is that it makes little sense simply to transfer the grid-based working model to the cloud, rather than creating a new model more suited to cloud resources. A second caveat is that the figures presented below are based on extrapolations from those presented by the commercial providers, instead of benchmarks from real experience using their services. For processing power in particular they tend to refer to “virtual CPUs” and “compute units”, which bear a non-straightforward relationship to physical cores. The types of instance (VM) available vary between companies too, again making comparisons more difficult.

Here is a table of the three instance types chosen, all based on European pricing:

Vendor/Type	vCPU	Compute Units	CPU	RAM	Storage	Costs
AWS c3.large	2	7 ECU	E5-2680	3.75 GB	2 x 16GB SSD	$0.171 per hour
Google n1-highcpu-4	4	11 GCEU	2.6GHz Intel Sandy Bridge or newer	3.60 GB	N/A	$0.292 per hour
Azure Large (A3)	4	N/A		7 GB	~400GB	£0.153 per hour

Google offers Persistent Storage for $0.04/GB/month and Microsoft does not list prices for storage capacities above 9 PB/month.

The CPU and Disk information is taken from the table supplied by Pete Gronbech. How much it would be to provide the same resources from the three cloud vendors was calculated. The total HEPSPEC06 achieved across all the GridPP Tier 2 sites is 269,693 and the total storage is 13,049 TB. The figures from the vendors do not have any special discounts applied and in particular AWS 'spot pricing' has not been used. The CHEP 2013 presentation from Tony Wong at BNL showed that runtime was increased by 50% when using spot pricing, owing to the inherent unreliability caused by price fluctuations. The calculations assume the HEPSPEC06 performance of an Intel E5-2660v2, which is generous to some of the vendors and harsh on others (and to reiterate, they are based on extrapolation, not real benchmarking).

Vendor	Annual Compute Cost	Annual Storage Cost	Annual Transfer Cost	Notes
Amazon	£5,476,300	£5,253,648	£31,026	EBS
Google	£4,675,671	£3,820,835	£43,920	Provisioned Storage
Microsoft	£4,016,268	N/A	N/A	

The costs shown are for 100% usage across an entire year (8760 hours).

In addition to these raw costs, over-provisioning would be needed to provide more compute resilience (e.g. using some resources based in the US as well as Europe) and some staffing would be required to manage and administer the clouds. It may also be desirable to provide data duplication, using a mostly read-only service such as Amazon Glacier.

Here is another example using the vendors' respective “high memory” instances. The instance details are:

Vendor	vCPU	Compute Units	CPU	RAM	Storage	Costs
Amazon m2.xlarge	4	13	Intel Xeon family	17.1	1 x 420	$0.460/hour
Google n1-highmem-4	4	11	2.6GHz Sandy Bridge or newer	26	N/A	$0.549/hour
Azure A6	4	N/A		28	N/A	£0.408/hour

The resulting costs are:

Vendor	Annual Compute Cost	Annual Storage Cost	Annual Transfer Cost	Total	Notes
Amazon	£14,731,567	£5,253,648	£31,026	£20,016,241	EBS
Amazon Heavy Reserved	£9,420,869.5	£5,253,648	£31,026	£14,705,543.5	One year of 'heavy usage' reserved instances
Google	£8,790,902.2	£3,820,835	£43,920	£12,655,657.2	Provisioned Storage
Microsoft	£10,710,048	£7,472,977.2	£64,062	£18,247,087.2	Geographically Redundant Storage

Again, these are rather simplistic equivalents of what we have currently at GridPP sites.

Amazon makes the following claims about storage/networking in their documentation about instance types:

> ***2 HI1 instances can deliver more than 120,000 4 KB random read IOPS and between 10,000 and 85,000 4 KB random write IOPS (depending on active logical block addressing span) to applications. The maximum sequential throughput on is approximately 2 GB/s read and 1.1 GB/s write.**
> ***3 HS1 instances can deliver 2.4 GB/s of 2 MB sequential read performance and 2.6 GB/s of sequential write performance.**
> ***4 Instances launched into the same cluster placement group are placed into a non-blocking 10 Gigabit ethernet network**

(taken from https://aws.amazon.com/ec2/instance-types/)

Performance

HS1 instances are $4.9 per hour i.e. roughly 10X the cost listed above, while HI1 instances are $3,1 per hour. Network performance for the c3.large instance type exemplified in the first set of costs is listed as “moderate”. To achieve what's listed as “high”, the c3.2xlarge type would be needed, which is $0.683 per hour, or to ensure all instances are on the same network (as in the note 4 quoted above), the c3.8xlarge type would be needed, which is $2.732 per hour, about 5 times more than the m2.xlarge high-memory type in the second set of costs.

Google's documentation for I/O performance is here: https://developers.google.com/compute/docs/disks and the implication is that to reach the levels mentioned by Andrew would add significantly to the monthly cost.

Joe Masters Emison at Network Computing has run a series of synthetic benchmarks on various public cloud providers:

Azure:
http://www.networkcomputing.com/cloud-computing/iaas-performance-benchmarks-part-5-micro/240164597

Google:
http://www.networkcomputing.com/cloud-computing/iaas-performance-benchmarks-part-4-googl/240164495?pgno=1

Amazon (c3 instances):
http://www.networkcomputing.com/cloud-computing/iaas-performance-benchmarks-part-3-aws-c/240164236

though these are all of single instances rather than a cluster.

Anecdotally, I've read a lot of comments from people saying that any workload involving significant I/O performs badly on AWS, which is really designed for low-usage web hosting tasks, and I imagine the same would apply to the other vendors.

Outages

All the public cloud providers have suffered significant outages. Examples include:

Amazon:

April 20th 2011 - http://aws.amazon.com/message/65648/
June 29th 2012 - https://aws.amazon.com/message/67457/
October 22nd 2012 - https://aws.amazon.com/message/680342/
December 24th 2012 - http://aws.amazon.com/message/680587/

The last one caused problems with Netflix streaming and they referred to those here:
http://techblog.netflix.com/2012/12/a-closer-look-at-christmas-eve-outage.html
and a response (with an element of self-interest/self-promotion) from RightScale:

http://www.rightscale.com/blog/cloud-management-best-practices/aws-outage-lessons-learned-if-netflix-can-suffer-so-can-you

Google:

January 18th 2014 - https://groups.google.com/forum/#!topic/gce-discussion/56LsUSfEQdc

(Google Compute Engine has not been in production for very long)

Microsoft:

February 29th 2012 - http://blogs.msdn.com/b/windowsazure/archive/2012/03/09/summary-of-windows-azure-service-disruption-on-feb-29th-2012.aspx
July 26th 2012 - http://gigaom.com/cloud/microsoft-pins-azure-outage-on-network-miscue/
February 22nd 2013 - http://www.theverge.com/2013/2/22/4019772/xbox-live-and-windows-azure-suffering-from-extended-outages
October 30th 2013 - http://www.pcworld.com/article/2059901/microsofts-windows-azure-cloud-hit-by-worldwide-management-interuption.html

Here is a screengrab of a graph showing two months' worth of Amazon cloud status:

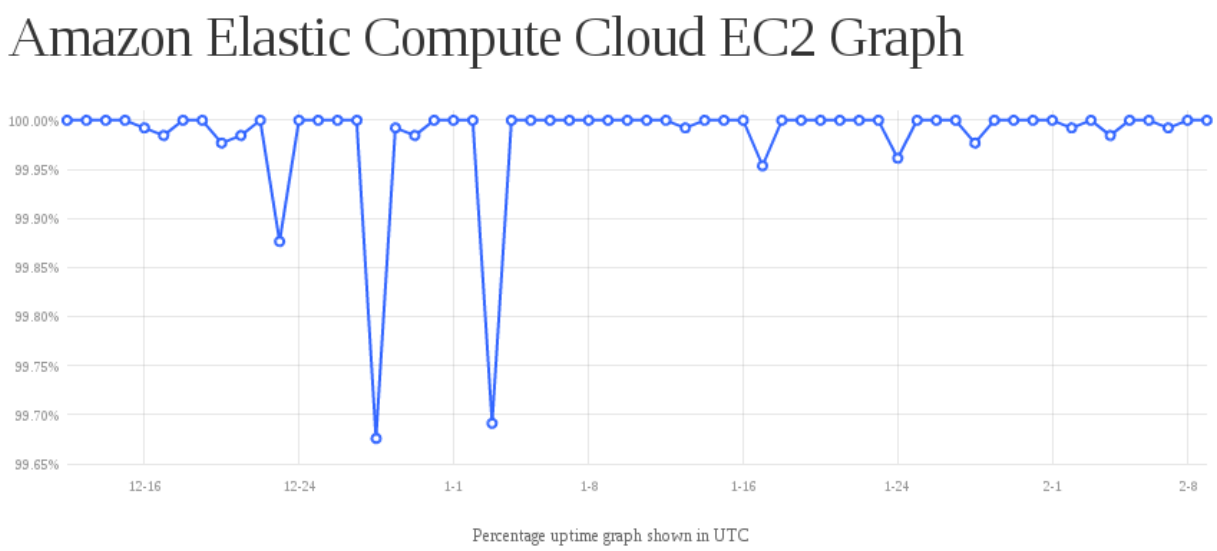

At the very least it would be prudent to have multi-region coverage from a single vendor, and much better to be using multiple vendors in multiple regions, to mitigate these sorts of risks.

Further Information

Amazon Web Services

http://aws.amazon.com/ec2/instance-types/

https://aws.amazon.com/ec2/pricing/

Google Cloud

https://developers.google.com/compute/docs/machine-types

https://developers.google.com/compute/pricing

Microsoft Azure

http://www.windowsazure.com/en-us/services/virtual-machines/

http://www.windowsazure.com/en-us/pricing/details/virtual-machines/