

Cloud Price Index October 2014 - \$2.56 is the market price

Analyst: Owen Rogers

17 Oct, 2014

In the first edition of the Cloud Price Index, the average hourly price for the 'typical' Web application is \$2.56, with the hyperscalers (AWS, Azure and Google Compute Engine) slightly cheaper at \$2.36. However, in our simpler Virtualization Price Index, the roles are reversed: the VPI is \$0.73 across all providers, whereas the hyperscalers are slightly more expensive at \$0.78.

The 451 Take

The conclusions from this edition of the Cloud Pricing Index are very general: discounts can make substantial savings, larger providers have a greater range of services, and the big names can be cheaper than the hyperscale providers. The CPI and VPI benchmarks of \$2.56 and \$0.73, respectively, provide a basis for evaluating how to price and how to consume cloud services. But the interesting analysis comes into play once a new round of price-cutting takes place. Will the headlines match the typical expenditure of a typical application? We expect further price cuts to take place in November, when a new CPI will be published.

To examine the 'real world' cost of cloud over time, 451 Research has created a Cloud Price Index. Like a consumer price index, our Cloud Price Index is made up of a basket of goods, but in our case it is a specification of the services required to operate a typical Web server application. Changes in the CPI over time will reflect how changes and differences in pricing in the cloud industry are reflected in real-world use. As well as a CPI, another index - the Virtualization Price Index - will be used to assess changes in price of a simpler application:

- The Cloud Price Index: A multi-service, three-tier cloud application consisting of Linux Virtual

Machines, object storage, block storage, relational database, NoSQL database, load balancing, access control lists and snapshot backup in a resilient architecture.


- The Virtualization Price Index: A single-service three-tier Web application consisting of differently sized virtual machines in a nonresilient architecture.

Those who provided quotations can submit standard pricing and cost-optimized pricing based on reserved instances, volume discounts, etc. We analyze the index in two manners: first, considering just the hyperscale providers of AWS, Microsoft Azure and Google Compute Engine, and second, the cheapest providers that have submitted quotations.

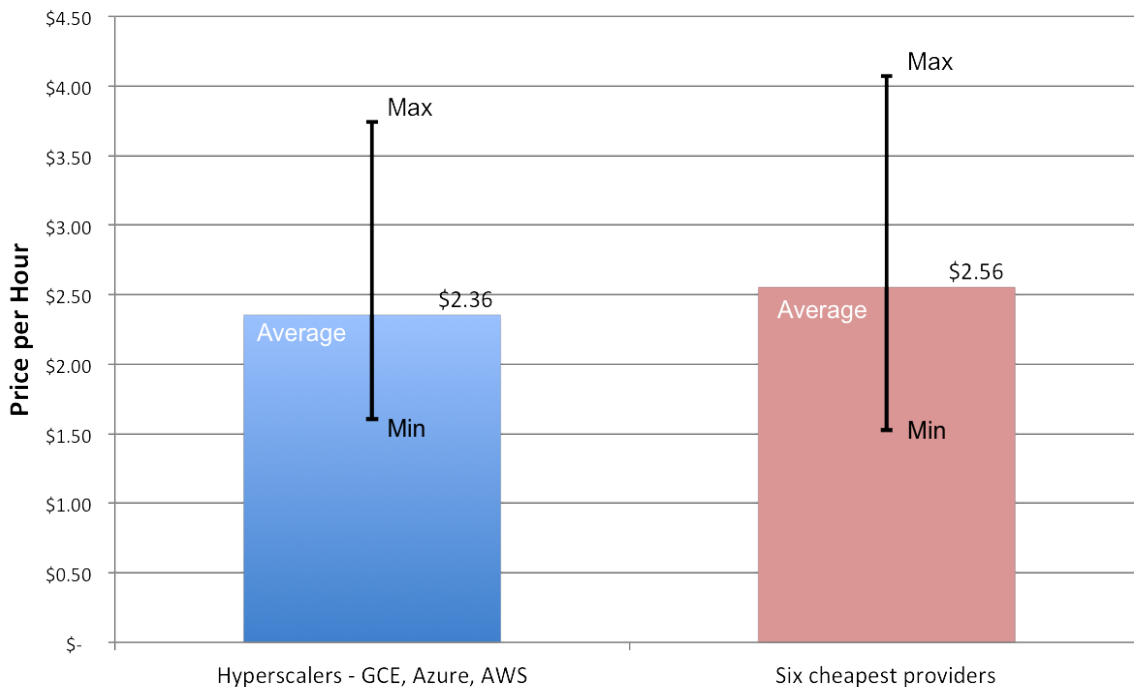

For a full explanation of the methodology and results, please read this companion document.

The Cloud and Virtualization Price Indexes as of October 3, 2014


Cloud Price Index (CPI)



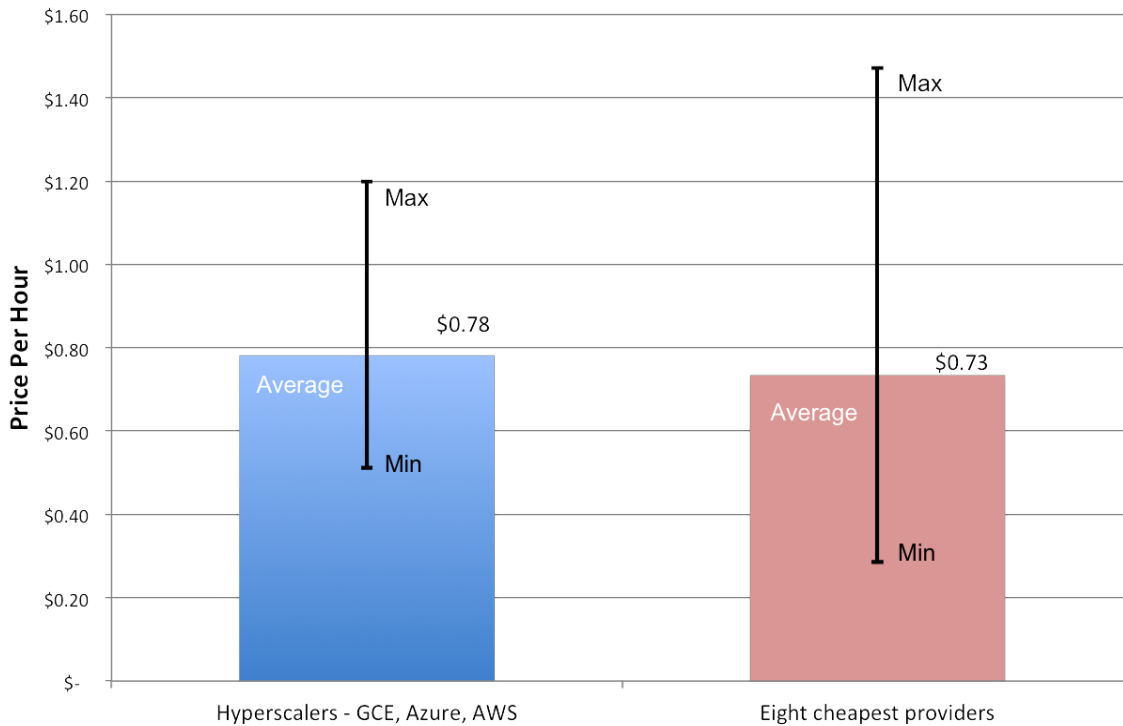

Average Saving for Commitment



Virtualization Price Index (VPI)



Average Saving for Commitment



Contributing companies

For the CPI, Amazon Web Services provided a public-facing quotation of \$1,170 using standard pricing and \$602 using optimized pricing, a normalized saving of 49%. Google responded with a public-facing quotation of \$2,731.02 using standard pricing and \$2,282 using optimized pricing, a normalized saving of around 16%. Microsoft Azure did not provide a quotation, so an online estimate based on quoted rate card pricing was used. Verizon, Swisscom and Windstream all provided detailed quotations for the CPI. Dimension Data, LeaseWeb, Colt and 2nd Watch contributed data not included in the index.

For the VPI, Amazon Web Services, Google Compute Engine and Microsoft were not approached for quotations because the simplicity of the specification made its pricing trivial. UpCloud responded with a public-facing monthly cost of \$208 for its Chicago datacenter. Gandi responded with a standard price of \$315 and an optimized price of \$222. Lunacloud, Internap and Peak 10 all responded with detailed quotations.

Analysis

The average cost of running our multi-service cloud application is \$2.56 per hour, or just \$1,865 per month, which includes bandwidth, storage, databases, compute, support and load balancing in a nongeographical resilient configuration. At this hourly price for an application that potentially could deliver in excess of 100,000 page views per month, it's easy to see how cloud is a compelling proposition. The hyperscale providers are generally able to offer this application at a better price point of \$2.36 (\$1,722 per month), and have a small min/max range compared with all cloud providers included. However, it's worth noting that the cheapest provider was not a hyperscale provider – a sign, perhaps, that enterprises should investigate all options and not just the big names.

Those hyperscalers that provided discounting details were found to be on average 32% cheaper than on-demand, with an average saving of 26% achieved across all providers. AWS offered a large discount of 49% in our bursting scenario, in relative terms. This backs up one of the main findings of the original Cloud Codex: on-demand pricing enables low-risk experimentation, but alternative pricing schemes are the way to save money in the long term. The most expensive hyperscale provider and overall provider was more than twice the price of the cheapest (\$3.64/\$1.60 and \$1.52/\$4.07, respectively).

The average cost of running our simple cloud application is just \$0.73 per hour, or \$536 per month, for all virtual machines and bandwidth. Again, such a small hourly price shows why cloud services are so attractive. Interestingly, the average price from the hyperscale providers was higher at \$0.78 per hour, or \$570 per month, suggesting that for simple applications, the smaller names may be cheaper. Hyperscale providers were able, on average, to provide end users with savings of 29%, typically through easily purchasable instruments, such as reserved instances, sustained-use discounts and subscriptions. Across the board, an average saving of 23% can be achieved through commitment and tough negotiation. The range of the price of quotations received for the VPI was significant: The most expensive was more than double the price of the cheapest (\$1.20 compared with \$0.51).

Differences in performance, support, geographies and portfolios; a small sample size; and a single use case mean the CPI and VPI aren't a measurement of value or relative competitiveness of individual providers, and certainly aren't a perfect comparison. We are also trusting providers to submit honest, accurate and realistic quotes. But it does, at least, draw a line in the sand to help us understand changes in cloud pricing.

General findings

Through the process of data collection, the experience has given us some more quantitative findings that we will cover in a future report:

- Few companies could provide the specification requirements of the CPI. The CDN pricing component was removed from final quotations because it could not be adequately addressed by all providers. The hyperscale providers are unrivalled when it comes to breadth of cloud services. Windstream, Google and AWS all included CDN capability in their quotations.
- Many companies responded that they didn't have enough time to provide quotations. This suggests that the industry is far away from simple, easily understood pricing. Many participants in the CPI and VPI provide calculators to determine pricing.
- Many companies responded that they couldn't provide this information. This suggests that open and transparent pricing is also yet to be fully realized. Most, but not all, participants in the CPI and VPI provide public pricing.

Reproduced by permission of The 451 Group; © 2015. This report was originally published within 451 Research's Market Insight Service. For additional information on 451 Research or to apply for trial access, go to: www.451research.com