

HYPERSCALE VS. ENTERPRISE-CLASS IAAS CLOUD

Which Deployment Option is Right for You?

Cloud computing has taken off over the past few years, and for good reason. By being able to offload mission-critical computer, storage and/or network services to a third-party provider able to operate at a grand scale, organizations gain immediate access to a wider array of IT power and resource at a faster rate than they would by going it alone and expanding on-premises networks. After all, there's a reason more than half of all enterprises with 1,000 plus employees will likely increase cloud spending in 2015.¹

Of the three main cloud deployment models (software-as-a-service, platform-as-a-service and infrastructure-as-a-service), IaaS remains an enterprise fixture, providing organizations with a solid cloud foundation upon which everything else can be built and maintained. It has proven to be immensely popular with businesses in a variety of industries, with research firm Ovum predicting that 80 percent of global enterprises will have IaaS in place by 2016.²

While enterprises are increasingly coming around to the benefits of IaaS, as yet there is no consensus on how IaaS is deployed and consumed. At the moment, organizations have two main options: hyperscale clouds offered by the likes of Amazon and Google that are consumed on an instance basis, and enterprise-based clouds that function on a resource pool-based deployment model.

¹ <http://www.forbes.com/sites/louiscolumnbus/2014/11/26/computerworlds-2015-forecast-predicts-security-cloud-computing-and-analytics-will-lead-it-spending/>

² <http://blogs-images.forbes.com/louiscolumnbus/files/2015/01/Ovum-Cloud-Adoption-Forecast.jpg>

HYPERSCALE VS. ENTERPRISE-CLASS IAAS CLOUD: WHICH DEPLOYMENT OPTION IS RIGHT FOR YOU?

PAGE 2

When deciding between the two, it's critical to think about all key variables that come with one option or another. In particular, how do hyperscale versus enterprise-class clouds compare in regard to availability, cost, flexibility and support? By keeping these four key variables in mind when choosing an IaaS solution, organizations can be sure they get the right cloud for their needs.

Breaking down hyperscale and enterprise-class cloud computing

Before determining how each option breaks down as far as the four key variables are concerned, it's helpful to establish a firm breakdown between what organizations get from an enterprise-class cloud and from a hyperscale cloud provider.

Solutions like Amazon Web Services and Microsoft Azure offer IaaS on an instance basis, which essentially means that compute and storage is purchased in pre-configured bundles and clients get the equivalent of bare-metal resource. This is perhaps the easiest way to obtain IaaS, but the availability design must be built into the application layer. This makes it a low-cost option, but its customers typically have little to no choice about what hardware is used and how data is shared and disseminated. For individuals like developers, hyperscale clouds facilitate easy horizontal scaling, especially for read-only applications, just without high availability guarantees.³

In contrast, enterprise-class clouds offer consumption on a resource pool basis. Under this model, the cloud service provider offers a bundle of resource (typically compute, storage and network) bundled together in a package that is primarily handled and overseen by the provider.⁴ This delivery model enables scale from within a set pool of pre-established resource, and often the virtual machines, or 'instances' can be easily expanded or contracted as needed within the pool. It also provides for dedicated segregation within the IaaS environment,⁵ enabling guaranteed isolation of applications when necessary. Availability design is built into the lowest layer of the stack, and it can support both horizontal and vertical scaling. It is typically not the easiest IaaS model to deploy, but more often than not it is highly available thanks in part to guaranteed service-level agreements.

Comparing IaaS deployment models based on scalability and availability

Hyperscale clouds that deploy on a per-instance basis are built with scalability in mind, enabling its customers to easily purchase more resource on demand when necessary. Since they are not location-dependent solutions, workloads in hyperscale clouds can be made highly available with data distribution and resource deployment across 'availability zones.' For organizations with horizontally scaling applications, hyperscale clouds provide great redundancy options and scale to ensure availability and performance.

Both hyperscale and enterprise-class clouds are built with scalability in mind, just at different rates. While an enterprise-class cloud does support horizontal scaling, it can enable its customer to make small, granular requests to better fine tune their IaaS to meet their needs. It also provides a highly available arrangement - one that is often guaranteed and spelled out in an SLA - for supporting mission-critical data and applications. This all comes at a cost, however, as it means in-house IT staff need to have enough knowledge of the cloud to take full advantage of the granularity of a resource pool cloud. Service providers also typically charge more for access to all of these benefits.

A good use case for an enterprise-class cloud is in banking, as credit unions and banks need a database for all of their transactional information that can be accessed from anywhere. This makes this data warehouse an ideal candidate for cloud hosting, but the highly sensitive nature of this financial information means high availability and granular control are critical here. As such, a bank or credit union would likely turn to an enterprise-class cloud in order to have both a highly scalable and a highly controlled environment for this data.

**HYPERSCALE CLOUDS
FACILITATE**

EASY HORIZONTAL SCALING

³ http://blogs.gartner.com/lydia_leong/2012/12/05/cloud-iaas-slas-can-be-meaningless/

⁴ <http://blogs.technet.com/b/yungchou/archive/2013/08/08/resource-pooling-virtualization-fabric-and-cloud.aspx>

⁵ <http://www.informit.com/articles/article.aspx?p=2093407&seqNum=2>

Expenses associated with both IaaS types

As their name entails, hyperscale clouds are built to be all encompassing and deployed at great scale. Since they target a large and diverse user base, each customer is contributing to only a small portion of the overall revenue pool. This enables hyperscale clouds to charge relatively low fees to get started. But, since these cloud providers price their offerings on an as-needed basis that is mostly catering to small workloads, costs related to scaling or heavy usage can be high.⁶ This deployment model can also easily exacerbate data sprawl-related concerns, creating untold copies of data with little or no oversight every time an expansion is purchased. Hyperscale cloud providers also typically charge their customers to move data.⁷

Upon initial use, an enterprise-class cloud can be more expensive than a per-instance solution. But, using and expanding an enterprise-class cloud comes with no or very few fees related to scaling and data movement, making costs much more of a fixed concern. In addition, as the old saying goes, you get what you pay for. The upfront fee provides more vendor support and granular control, on top of the decreased risk for spiraling fees at scale.

How flexibility differentiates enterprise-class and hyperscale clouds

Per-instance cloud solutions, in order to appeal to as wide of an audience as possible, have a number of core offerings on sale at any given time. This provides flexibility in allowing customers to choose the option that is right for them.

Their flexibility and pricing model is akin to what customers can expect when they go to a clothing store to purchase shirts. Just as shirts come in a small number of set sizes (small, medium, large and extra large), hyperscale clouds has a set number of predetermined offerings for customers to choose from. This can be an ideal starting point for clients with little or no prior understanding of what they might need from an IaaS solution, it is not customizable and can create confusion when internal needs do not align neatly with the set provider options.

Enterprise-class clouds, in comparison, are built with flexibility in mind. Each resource pool can be custom-built and altered for each client on an as-needed basis. This way, expansion or remodeling can happen exactly as is required by a customer without having to build something new and migrate data each time.

AN ENTERPRISE-CLASS
CLOUD COMES WITH NO OR

VERY FEW FEES

⁶ <http://searchcloudcomputing.techtarget.com/feature/Climbing-out-of-an-AWS-EC2-cost-sinkhole>

⁷ <https://aws.amazon.com/blogs/aws/aws-data-transfer-price-reduction/>



Which type of provider has the most knowledgeable staff?

Hyperscale cloud service providers need to have some of the most knowledgeable people in the cloud industry on their side to build out at a pace needed to support all of their customers. The only problem, however, is that these staff members are not typically client facing. Per-instance clouds operate on a “set it and forget it” kind of principle, in which provider involvement in the day-to-day cloud activities of its customers is minimal. This arrangement is ideal for keeping costs low, but not for those that need more support in the transition to and use of the cloud.

For teams that either lack internal expertise or would rather rely heavily on a third-party provider for all of their cloud needs, enterprise-class clouds are an ideal option. Resource pool-based service providers work on a per-client basis and often offer expertise and advice as part of the contract. According to a recent survey of global CIOs conducted by CSC, many IT leaders are increasingly looking to further leverage their relationships with their vendors⁸ - a request that plays into the strengths of resource pool-based cloud computing. Plus, with SLAs at their base, enterprise-class clouds enable firms to take advantage of a highly available solution without having to devote a lot of internal resource toward cloud uptime and maintenance.

Which Deployment Option is Right for You?

The answer to this question is not cut and dry, and it all depends on what an organization needs. While there is consensus on the benefits of IaaS - there's a reason deployed IaaS instances are expected to grow at a compound annual rate of 13 percent between 2013 and 2018⁹ - no one option works well for everyone. For firms in need of a low-cost option that is easy to deploy, a per-instance option may be the best bet. Those prioritizing reliability, customization and support from their cloud solution will likely be far happier with an enterprise-grade cloud. By looking beyond just the consumption model and also inspecting an offering's availability, scalability, cost, flexibility and level of available staff support, organizations can be certain they select the best IaaS deployment model for their unique needs.

About Faction

Faction is an enterprise-class IaaS cloud service provider offering private, public, & hybrid cloud solutions through channel partners. At Faction we supply cloud the way you want it with extreme performance, deep control, and broad customization capabilities. When you join the Faction fold, you take back the keys to your kingdom. Reign as supreme commander in chief of your cloud. No compromises. No exceptions.

Faction is the only cloud that offers patented plug-and-play direct connections (via layer 2) into its cloud resulting in huge time savings (no time spent re-configuring everything)! With Type II SSAE 16 and SOC 1 & 2 compliant cloud nodes in eight geographies across the United States and in Europe (Seattle, Santa Clara, Denver, Chicago, Atlanta, New Jersey, New York, and the United Kingdom), Faction offers both Cisco UCS and Open Compute platforms, is a Platinum-level NetApp Service Provider, and is VMware vCloud® Powered. For more information, visit www.factioninc.com or call (855) 532-4734.

RICK VINCENT

DIRECTOR OF SOLUTIONS ENGINEERING

About the Author

As Director of Solution Engineering for Faction, Richard Vincent is responsible for architecting robust IaaS solutions, for channel partners and enterprises, enabling them to conduct business with compelling differentiation. Prior to joining Faction, Mr. Vincent served as a Storage Solution Executive at Dell, performed as an independent cloud consultant, and held the position of Director, Data Center Computing for Global Technology Resources, Inc. GTRI. Mr. Vincent is a graduate of the University of Georgia, College of Business and is fluent in the Japanese language.

⁸ http://assets1.csc.com/cio_survey_2014_2015/downloads/CSC.Global.CIO.Survey.IG.pdf

⁹ http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.pdf

