

MULTI-TENANT PUBLIC CLOUD: THE PROS AND CONS OF STAYING SINGLE

Considering how many cloud choices are available, it may seem like the public option is the simplest one to pick. But that may not always support the performance and utilization that your particular workloads require.

Public cloud data centers will account for 61.2% of total spending on cloud infrastructure deployment in 2017, according to market research firm IDC. But while it may seem as though public cloud is the best way to go, it involves fundamental tradeoffs that may not be acceptable or desirable to your business.

Because public cloud is a multi-tenant environment in which you share resources with other companies (the other tenants), not all workloads thrive on it. Certain applications and business functions require dedicated infrastructure to ensure performance, availability, and the right degree of control. For that, you may need to turn to private cloud or more traditional data center options.

HYPERSCALERS SET THE TONE

The public cloud services offered by hyperscale providers such as Amazon Web Services, Microsoft Azure, and Google Cloud Platform have achieved remarkable market share and a reputation for providing services at a scale and price that are hard to beat.

"Hyperscale operators typically have hundreds of thousands of servers in their data center networks, while the largest, such as Amazon and Google, have millions of servers," says [Synergy Research Group](#). Synergy estimates that in 2016 the top 24 operators "accounted for 68% of the cloud infrastructure services market (IaaS, PaaS, private hosted cloud services) and 59% of the SaaS market."

So, it's not surprising with the cost and scale advantages that public cloud adoption is on the rise. Aggregate public cloud and private cloud spending is [projected to exceed](#) traditional data center spending in 2020. Private cloud spending will continue to increase through 2020, although public cloud spending will grow more.

MULTI-TENANT FAULT LINES

Many enterprise cloud users first became familiar with the multi-tenant cloud model through the adoption of Salesforce.com's software-as-a-service (SaaS) application. By sharing infrastructure across customers, Salesforce disrupted the traditional market for enterprise software applications, greatly reducing licensing costs and removing the customer pain point of installing and maintaining applications.

The traditional multi-tenancy SaaS application involved a single instance of an application with dedicated shares sliced up across multiple customers. A more advanced model uses virtualization to run different instances of its application for different customers, providing much of the benefit of a single-tenant application, while still sharing infrastructure across many customers.

But there are downsides to this model, in particular the cascade effect when a major public cloud service suffers an outage. It's not just individual customers of those providers that suffer; it's also the many customers of each SaaS provider using that public cloud platform.

While a multi-tenant service should have inviolable processes to ensure that each customer's data is walled off from the other

tenants, there is the potential for [noisy neighbor](#) syndrome. One customer's demand for resources could impinge on other customers if the infrastructure is not sufficiently elastic to accommodate spikes.

ADEACA, provider of a single-tenant cloud ERP, [points out other disadvantages](#) of multi-tenancy, including:

- You can't customize the code of a multi-tenant application.
- The provider can force through upgrades without your signoff.
- Providers could delay or remove features to satisfy requests of their largest customers.

Another concern for enterprises is the evolving state of [data sovereignty](#). Requirements vary from country to country and are subject to modification, often at the whim of the host country. The ability of public cloud providers to move data around for purposes of efficiency and recovery could put you in jeopardy unless there are strict requirements in the provider contract that govern this issue.

"Public clouds," writes *Network World's* [Brandon Butler](#), "provide access to an almost infinite amount of infrastructure resources without any upfront investment required, and the ability to use cutting-edge technology available first from public cloud vendors. The perceived lack of control and security can be a big turn-off for some customers though."

SOME SHOULD JUST STAY SINGLE!

Single-tenant systems have the benefit of dedicated resources assigned to each customer, eliminating any possibility of neighbor contention and providing greater customer control over customization and upgrades.

Now, many enterprises are going to equate "dedicated" with "extra cost." But that's not necessarily the case, writes [Suit Bhavsar](#), who is responsible for managing and developing Rackspace's server infrastructure to support many industry sectors and workloads.

"When your application runs at a steady state in demand patterns, the total cost of ownership for dedicated hosting is lower than it would be on a public cloud," according to Bhavsar. "Depending on your scaling needs, if your traffic patterns are cyclical or additional compute and storage resources are needed during peak times, bursting into the cloud may be more economical." In addition, he says, dedicated infrastructure is cheaper to run over the long haul due to advances in power management, cooling, and server efficiency.

Datamation's Cynthia Harvey provides an excellent analysis of the [pros and cons](#) of public cloud and private cloud, observing that private vs. public cloud doesn't have to be an either-or choice. "Some organizations find that a hosted private cloud provides a good balance between the relative strengths and weaknesses of private vs. public clouds," she writes. "These environments are managed and run by a third-party vendor, but the physical infrastructure is dedicated to the use of one particular organization."

SOMETHING FOR EVERYONE

Each company's needs are different, but luckily there's a cloud option model for everyone. "Each deployment model—public, private, or hybrid cloud—has its own strengths and weaknesses, making it uniquely suited for particular use cases," says Harvey.

Ultimately, what will matter are performance and the ability to utilize the best cloud option that matches the needs of any workload at any given point in time. For help in finding the right path, visit [Rackspace](#).