

Is Your Glass of Cloud Half-Empty or Half-Full?



Lori MacVittie, 2010-10-06

If we look at cloud in terms of what it does offer instead of what it doesn't, we may discover more useful architectures than were previously thought to exist.



I have a fairly large, extended family. While I was growing up we gathered at our grandparent's home during the holidays for, of course, a meal. Grandma would put extra chairs around the table but because she had five children (and spouses) there really wasn't any room for us grandchildren. So we got to sit ... at the little kid's table. Eventually we weren't "little kids" any more and we all looked forward to the day we could sit at the "big" table with the adults.

Now grandma was a stickler for time, and dinner was served at exactly twelve noon. Not 12:01, not 11:59. 12:00. Exactly. If you weren't there, that was just too bad. So it inevitably it was the case that someone wasn't on time and it was then that, in age-descending pecking order*, some of the "kids" got to sit at the "big" table. Until the Johnny-come-lately adults showed up, at which point we were promptly banished back to the kids table.

This "you can sit at the big-table unless a grown up needs your place" strategy is one that translates well to a hybrid [cloud computing](#) strategy.

THE LITTLE KIDS' TABLE = THE CLOUD

There are myriad surveys out there regarding the inhibitors to cloud adoption. At the top is almost always security and control. CIOs are quick to indicate they do, in fact, have interest in the cloud and its purported operational benefits, but they aren't necessarily willing to risk the security and availability of business-critical applications to get them.

As has been previously mentioned in [Why IT Needs to Take Control of Public Cloud Computing](#), it may be that IT needs to adopt the view that the *data center* is the "big" table at which business-critical applications are deployed and the cloud, as the "little kids' table" is where all non-critical application end up when the "big table" is full. A kind of cloud bursting, if you will, that assumes business-critical applications have priority over local data center compute resources. Non-critical applications may be initially deployed locally but if business-critical applications need additional compute resources then non-critical workloads must give up their local resources and move to the cloud.

This strategy treats cloud as it is today, as [compute on-demand and little more](#). It assumes, moreover, that the application needs very little "care and feeding" in terms of its supporting application and application delivery infrastructure. A little non-specialized [load balancing](#) for scale and a fat pipe might be all this application really needs. That makes it perfect for deployment in an environment that caters to providing cheap "utility" compute resources and little else because it can be migrated – [perhaps even while live](#) – to an off-premise cloud environment without negatively impacting the business.

That would not be true of a business-critical application for which there are strictly defined SLAs or compliance-related polices, many of which are implemented via complex integration with other components, systems, and applications internal to the data center. Migrating a business-critical "application" is a lot more complicated and time-consuming than a non-business critical, non-integrated application. That's because the former requires either (a) migration of all related components and supporting infrastructure or (b) a secure, optimized tunnel to the off-premise cloud computing environment that enables the continued use of integrated application and network components.

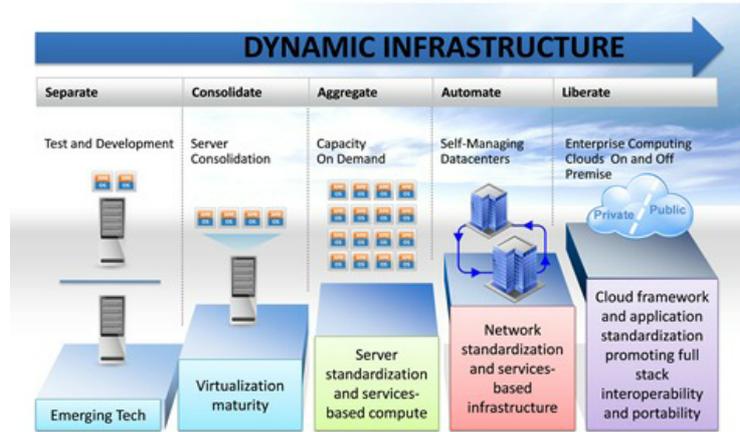
CLOUD (IM)MATURITY DRIVING FACTOR

The immaturity of cloud computing environments with regards to the availability of enterprise-class *infrastructure* services continues to be a root cause of cloud “reluctance.” Without the ability to deploy a critical application in an environment similar to that of the local data center, CIOs are going to be understandably cautious. But for applications that don’t need such a complex network of support infrastructure cloud computing is well-suited for deployment and doing so is certainly – at least from a CAPEX and long-term OPEX point of view – the most appealing option available.

Before cloud can mature, before we reach the “network standardization and services-based infrastructure” [we need the standards](#) upon which standardization will be based. Interestingly enough, that doesn’t necessarily mean *industry* standards. The speed at which various standards organizations are moving today makes it highly probable that organizations that are moving more quickly will develop their own set of standardization that will eventually form the basis for industry standards. Some might argue, in fact, that this is the way it *should* happen, as organizations are the ones that use and exercise the [Infrastructure 2.0 APIs](#) and frameworks currently available across the infrastructure spectrum to integrate and manage infrastructure components in their own data centers.

Without those standards and the resulting infrastructure services, in order for organizations to reap the benefits of cloud computing we should probably stop looking at cloud with a “glass is half-empty” view and take a “glass is half-full” perspective, instead. Don’t look at cloud in terms of what it doesn’t offer, but instead what it *does* offer: inexpensive, easily and rapidly provisioned compute resources. Compute resources that can serve as overflow for non-critical applications when the really important applications need more compute power.

* *As the oldest grandchild I was good with this order of operations, of course.*



F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com