Dynamic Infrastructure: The Cloud within the Cloud

Lori MacVittie, 2009-18-02

When folks are asked to define the cloud they invariably, somewhere in the definition, bring up the point that "users shouldn't care" about the actual implementation. When asked to diagram a cloud environment we end up with two clouds: one representing the "big cloud" and one *inside* the cloud, representing the infrastructure we aren't supposed to care about, usually with some pretty graphics representing applications being delivered out of the cloud over the Internet.

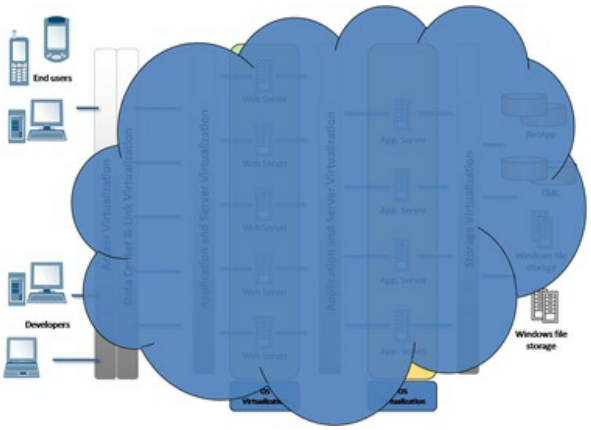

But yet some of us need to care what's obscured; the folks tasked with building out a cloud environment need to know what's hidden in the cloud in order to build out an infrastructure that will support such a dynamic, elastic environment.

It is the obscuring of the infrastructure that makes cloud *seem* so simple. Because we're hiding all the moving parts that need to work in concert to achieve such a fluid environment it appears as if all you need is virtualization and voila! The rest will take care of itself.

But without a dynamic infrastructure supporting all the virtualized applications and, in many cases, infrastructure such an environment is exceedingly difficult to build.

WHAT'S HIDDEN IN THE CLOUD

Inside the "cloud within the cloud" there are a great number of pieces of infrastructure working together. Obviously there are the core networking components: routers, switches, DNS, and DHCP, without which connectivity would be impossible.

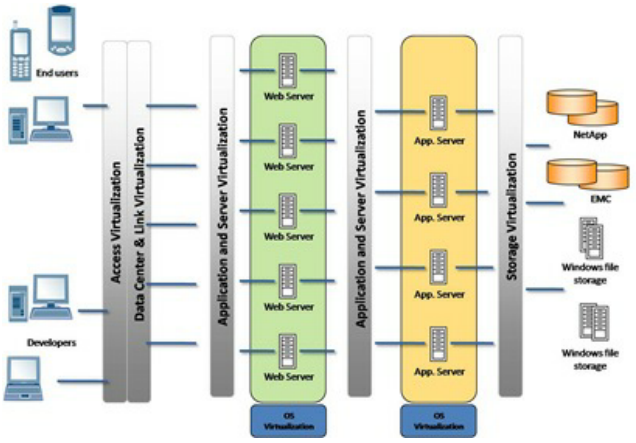

Moving up the stack we find load balancing and application delivery infrastructure; the core application networking components that enable the dynamism promised by virtualized environments to be achieved. Without a layer of infrastructure bridging the gap between the network and the applications, virtualized or not, it is difficult to achieve the kind of elasticity and dynamism necessary for the cloud to "just work" for end users.

It is the application networking layer that is responsible for ensuring availability, proper routing of requests, and applying application level policies such as security and acceleration. This layer must be dynamic, because the actual virtualized layers of web and application servers are themselves dynamic. Application instances may move from IP to IP across hours or days, and it is necessary for the application networking layer to be able to adapt to that change without requiring manual intervention in the form of configuration modification.

Storage virtualization, too, resides in this layer of the infrastructure. Storage virtualization provides enables a dynamic infrastructure by presenting a unified view of storage to the applications and internal infrastructure, ensuring that the application need not be modified in order to access file-based resources. Storage virtualization can further be the means through which cloud control mechanisms manage the myriad virtual images required to support a cloud computing infrastructure.

The role of the application networking layer is to mediate, or broker, between clients and the actual applications to ensure a seamless access experience regardless of where the actual application instance might be running at any given time. It is the application networking layer that provides ***network and server virtualization*** such that the actual implementation of the cloud is hidden from external constituents. Much like storage virtualization, application networking layers present a "virtual" view of the applications and resources requiring external access.

This is why dynamism is such an integral component of a cloud computing infrastructure: the application networking layer must, necessarily, keep tabs on application instances and be able to associate them with the appropriate "virtual" application it presents to external users. Classic load balancing solutions are incapable of such dynamic, near real-time reconfiguration and discovery and almost always require manual intervention.

Dynamic application networking infrastructure is not only capable but excels at this type of autonomous function, integrating with the systems necessary to enable awareness of changes within the application infrastructure and act upon them.

The "cloud within the cloud" need only be visible to implementers; but as we move forward and more organizations attempt to act on a localized cloud computing strategy it becomes necessary to peer inside the cloud and understand how the disparate pieces of technology combine. This visibility is a requirement if organizations are to achieve the goals desired through the implementation of a cloud computing-based architecture: efficiency and scalability.

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