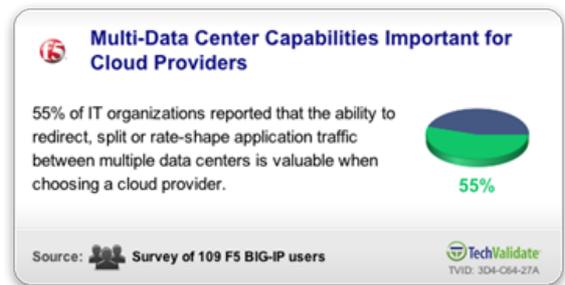


It is the dynamic nature of a “cloud controller”, an [application delivery network](#) (a.k.a. [load balancer](#)), that makes this scenario work. The application delivery controller seamlessly virtualizes *all* application resources – whether internal or in this private cloud – and presents them as a single resource. Those resources can be located anywhere as long as they are accessible via the internal network *somehow*. That somehow, in this case, is a cloud provider’s virtual private cloud solution. The application delivery controller adjusts in real-time (on-demand) to add or remove resources as necessary based on the total capacity available and the total capacity required.

The pool of resources used to provide the capacity to scale an application doesn’t need to be local or physical. It can be virtual, physical, local, or remote. It can be a combination of any of the four. A single pool can comprise a mix of resources and be utilized according to organizational policy and needs comprising variables across the network, the application, and even the business. Using a virtual private cloud via a secure (and one hopes [optimized/accelerated](#)) tunnel to a cloud provider an organization can easily extend its available resources out into the cloud seamlessly, with little to no disruption in the network or application network architecture.

The concept of a VPC can result in the tangible implementation of *internal* cloud bursting, as opposed to *external cloud bursting*. Both are equally valid uses of the cloud and result in the realization of the same goal: the ability to leverage on-demand compute resources from an external cloud provider seamlessly to achieve on-demand scalability of applications in the most cost-effective manner possible.



PERFORMANCE IMPLICATIONS

Worried that remote resources might affect performance? That’s absolutely a possibility given the architecture. So move the resources. That’s one of the benefits of a virtualized, cloud infrastructure, isn’t it? You assign the resources you need – that meet your needs and requirements – in the right place at the right time. Maybe that reporting application, which is going to take some time running anyway, can be moved via the VPC to the cloud and its resources re-allocated to your external facing web application.

This, too, illustrates the importance of a dynamic infrastructure to achieving the gains in efficiency and cost effectiveness touted by cloud and virtualization. The infrastructure needs to be able to support the reallocation of resources on-demand and, if possible, it should be capable of doing so itself based on configured policies that outline when and where and how resources should be allocated to meet both operational and business goals.

It may be the case that if the connectivity via the VPC into the cloud is optimized and accelerated enough there won’t be a performance degradation. That’s a variable that’s wholly dependent on the speed of the underlying connection, the application, and the load on the infrastructure as well as the application in question. The combination of caching and a VPC for some types of resources – scripts, images, and other static content – might turn out to be the best investment you’ve made in years in terms of performance *and* cutting costs. Using cloud-based storage and serving of static content would free up resources on internal application and storage infrastructure and drastically reduce the costs associated with that content. Then add some [caching](#) to your local unified application delivery network and you can improve performance *and* further reduce the cost of serving that content because it will only be accessed periodically, keeping your cloud-related resource consumption to a minimum.

But I digress. The point really is that there are plenty of exciting ways to leverage a virtual private cloud like that offered by Amazon and the fact that its IP addresses and resources are *private* is not in any way indicative of the maturity of or usefulness of the offering. On the contrary, this type of offering is necessary for enterprises to start moving forward and taking advantage of cloud computing.

- F5/Gartner Application Delivery Network Webcast
- The network ain't big enough for the both of us
- The Context Aware Cloud
- The House that Load Balancing Built
- Building an elastic environment requires elastic infrastructure
- Cloud Balancing, Cloud Bursting, and Intercloud
- Cloud Computing Makes Servers Obsolete
- The Virtual Public-Private Cloud Connection

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