

# F5 Friday: Elastic Applications are Enabled by Dynamic Infrastructure



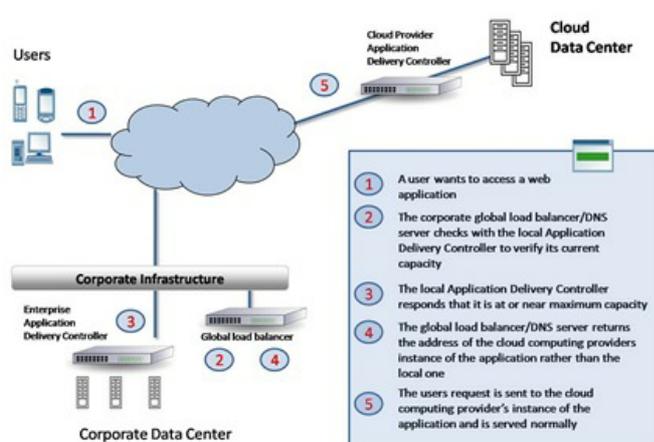
Lori MacVittie, 2010-03-09

You really can't have the one without the other. VMware enables the former, F5 provides the latter.

The use of public [cloud computing](#) as a means to expand compute capacity on-demand, a la during a seasonal or unexpected spike in traffic, is often called cloud bursting and we've been talking about it (at least in the hypothetical sense) for some time now.



When we first started talking about it the big question was, of course, *but how do you get the application in the cloud in the first place?* Everyone kind of glossed over that because there was no real way to do it on-demand.



## OVERCOMING the OBSTACLES BIT by BIT and BYTE by BYTE

The challenges associated with dynamically moving a live, virtually deployed application from one location to another were not trivial but neither were they insurmountable. Early on these challenges have been directly associated with the [difference in networking and issues](#) with the distances over which a virtual image could be successfully transferred. As the industry began to address *those* challenges others came to the fore. It's not enough, after all, to just transfer a virtual machine from one location to another – especially if

you're trying to do so *on-demand*, in response to some event. You want to migrate that application while it's live and in use, and you don't want to disrupt service to do it because no matter what optimizations and acceleration techniques are used to mitigate the transfer time between locations, it's still going to take some time. The whole point of cloud bursting is to remain *available* and if the process to achieve that dynamic growth defeats the purpose, well, it seems like a silly thing to do, doesn't it?

As we've gotten past *that* problem now another one rears its head: the down side. Not the negatives, no, the *other* down side – the [scaling down side of cloud bursting](#). Remember the purpose of performing this technological feat in the first place is dynamic scalability, to enable an *elastic application* that [scales up and down on-demand](#). We want to be able to leverage the public cloud when we need it but not when we don't, to keep really realize the benefits of cloud and its lower cost of compute capacity.

## FORGING AHEAD

F5 has previously proven that a live migration of an application is not only possible, but feasible. This week at [VMworld](#) we took the next step: elastic applications. Yes, we not only proved you can burst an application into the cloud and scale up while live and maintaining availability, but that you can also scale back down when demand decreases. The ability to also include a [BIG-IP LTM Virtual Edition](#) with the cloud-deployed application instance means you can also consistently apply any application delivery policies necessary to maintain security, consistent application access policies, and performance.

The complete solution relies on products from F5 and [VMware](#) to monitor application response times and expand into the cloud when they exceed predetermined thresholds. Once in the cloud, the solution can further expand capacity as needed based on application demand. The solution comprises the use of:

- [VMware vCloud Director](#)  
A manageable, scalable platform for cloud services, along with the necessary APIs to provision capacity on demand.

- [F5 BIG-IP® Local Traffic Manager™ \(LTM\)](#)

One in each data center and/or cloud providing management and monitoring to ensure application availability. Application conditions are reported to the orchestration tool of choice, which then triggers actions (scale up or down) via the VMware vCloud API. Encryption and WAN optimization for SQLFabric communications between the data center and the cloud are also leveraged for security and performance.

- [F5 BIG-IP® Global Traffic Manager™ \(GTM\)](#)

Determines when and how to direct requests to the application instances in different sites or cloud environments based on pre-configured policies that dynamically respond to application load patterns. Global application delivery (load balancing) is critical for enabling cloud bursting when public cloud-deployed applications are not integrated via a virtual private cloud architecture.

- [VMware GemStone SQLFabric](#)

Provides the distributed caching and replication of database objects between sites (cloud and/or data center) necessary to keep application content localized and thereby minimize the performance impact of latency between the application and its data.

I could talk and talk about this solution but if a picture is worth a thousand words then this video ought to be worth at least that much in demonstrating the capabilities of this joint solution. If you're like me and not into video (I know, heresy, right?) then I invite you to take a gander at some more traditional content describing this and other VMware-related solutions:

 [A Hybrid Cloud Architecture for Elastic Applications with F5 and VMware – Overview](#)

 [Hybrid Cloud Application Architecture for Elastic Java-Based Web Applications – Deployment Guide](#)

 [F5 and VMware Solution Guide](#)

If you do like video, however, enjoy this one explaining cloud bursting for elastic applications in a hybrid cloud architecture.

#### Related blogs and articles:



- [Bursting the Cloud](#)
- [vMotion Layer 2 Adjacency Requirements](#)
- [Cloud-bursting and the Database](#)
- [Cloud Balancing, Cloud Bursting, and Intercloud](#)
- [Cloud Balancing, Reverse Cloud Bursting, and Staying PCI-Compliant](#)
- [Virtual Private Cloud \(VPC\) Makes Internal Cloud bursting Reality](#)
- [How Microsoft is bursting into the cloud with BizTalk](#)
- [So You Put an Application in the Cloud. Now what?](#)
- [Migrate a live application across clouds with no downtime? Sure ...](#)

- **Just in Case. Bring Alternate Plans to the Cloud Party**
- **CloudFucius Asks: Will Open Source Open Doors for Cloud Computing?**
- **The Three Reasons Hybrid Clouds Will Dominate**
- **Pursuit of Intercloud is Practical not Premature**

---

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | [f5.com](http://f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)