

# DevCentral Architecture



Jason Rahm, 2012-30-10

Everyone has surely (don't call me Shirley!) at least been exposed to *THE CLOUD* by now. Whether it's the—I'll go with interesting—"to the cloud!" commercials or down in the nuts and bolts of hypervisors and programmatic interfaces for automation, the buzz has been around for a while. One of F5's own, cloud computing expert and blogger extraordinaire (among many other talents) [Lori MacVittie](#), weighs in consistently on the happenings and positioning in the cloud computing space. F5 has some wicked smart talent with expertise in the cloud and dynamic datacenter spaces, and we make products perfectly positioned for both worlds. With the release of all our product modules on BIG-IP VE last year, it presented the opportunity for the DevCentral team to elevate ourselves from evangelists of our great products to customers as well. And with that opportunity, we drove the DevCentral bull onward to our new virtual datacenters at [Bluelock](#).

## Proof of Concept

We talked to a couple different vendors during the selection period for a cloud provider. We selected Bluelock for a couple major reasons. First, their [influential leadership in the cloud space by way of CTO Pat O'Day](#). Second, their strong partnership with fellow partner VMware, and their use of [VMware's vCloud Director platform](#). This was a good fit for us, as our production BIG-IP VE products are built for the ESX hypervisors (and others in limited configurations, please reference the [supported hypervisors matrix](#)). As part of the selection process, Bluelock set up a temporary virtual datacenter for us to experiment with. Our initial goal was just to get the application working with minimal infrastructure and test the application performance. The biggest concerns going in were related to the database performance in a virtual server as the DevCentral application platform, DotNetNuke, is heavy on queries. The most difficult thing in getting the application up was getting files into the environment. Once we got the files in place and the BIG-IP VE licensed, we were up and running in less than a day. We took captures, analyzed stats, and with literally no tuning, the application was performing within 10% of our production baseline on dedicated server/infrastructure iron. It was an eye-opening success.

## Preparations

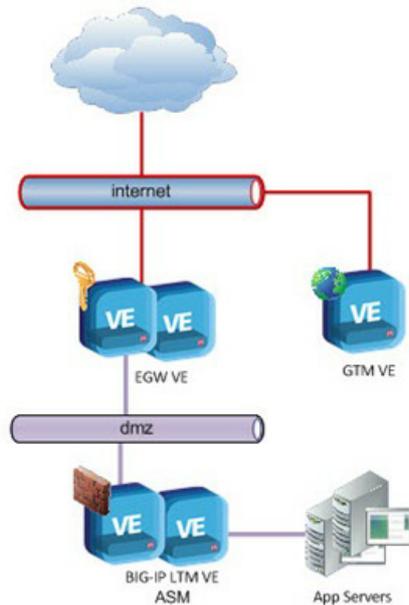
Proof of concept done, contracts negotiated and done, and a few months down the road, we began preparing for the move. In the proof of concept, LTM was the sole product in use. However, in the existing production environment, we had the LTM, ASM, GTM, and Web Accelerator. In the new production environment at Bluelock, we added APM for secure remote access to the environment, and WOM to secure the traffic between our two virtual datacenters. The list of moving parts:

- Change to LTM VE from LTM (w/ ASM module); version upgrade
- Change to GTM VE from GTM; version upgrade
- Change to Edge Gateway VE from Web Accelerator; version upgrade
- Introduce APM (via Edge Gateway VE)
- Introduce WOM (via Edge Gateway VE)
- Application server upgrade
- New monitoring processes
- iRules rewrites and updates to take advantage of new features and changes
- dns duties

There were many many other things we addressed along the way, but these were the big ones. It wasn't just a physical –> virtual change. The end result was a far different animal than we began with.

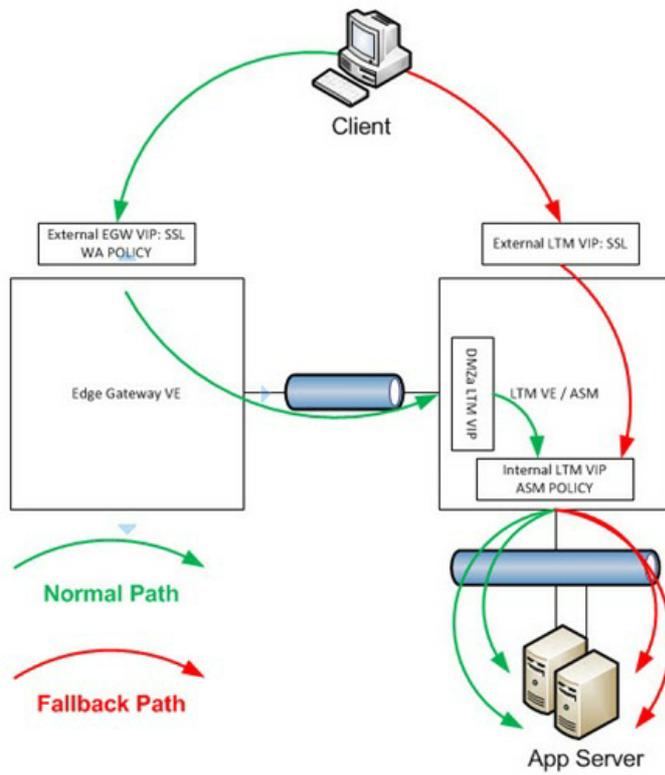
## Networking

In the vCloud Director environment, there are a few different network types: external networks, organizational networks, and vApp networks, and a few sub-types as well. I'll leave it to the reader to study all the differences with the platform. We chose to utilize org networks so we could route between vApps with ~~minimal~~ no additional configuration. We ended up with several networks defined for our organizations, including networks for public access, high availability, config sync, and mirroring, and others for internal routing purposes. The meat of the infrastructure is shown in the diagram below.



## Client Flow

One of the design goals for the new environment was to optimize the flow of traffic through the infrastructure, as well as provide for multiple external paths in the event of network or device failures. Web Accelerator could have been licensed with ASM on the BIG-IP LTM VE, but the existing versions do not yet support CMP on VE, so we opted to keep them apart. SSL is terminated on the external vips, and then the ASM policy is applied to a non-routable vip on the LTM, utilizing iRules to implement the vip targeting vip solution. This was done primarily to support the iRules we run to support our application traffic without requiring major modifications that would be necessary to run on a virtual server with a plugin (ASM,APM,WA,etc) applied.



If you're wondering about the performance hit of the vip targeting vip solution, (you know you are!) in our testing, as long as the front and backside vips are the same type, we found the difference between vip->vip and just a single vip to be negligible, in most cases less than a tenth of a percentage point. YMMV depending on your scenario. You might also be wondering about terminating SSL on BIG-IP VE. There is no magic here. It's simple math. In our environment, the handshakes per second are not a concern for 2k keys on our implementation, but if you need the dedicated compute power for SSL offload, you can still go with a hybrid deployment, using hardware up-front for the heavy lifting and VE for the application intelligence.

## Zero-Hands Expansion

One of the cooler things about living in a virtual datacenter is when it comes time for expansion. We originally deployed our secondary datacenter with less gear and availability, but additional production and development projects warranted more protections for failure. Rather than requiring a lengthy equipment procurement process, it took a few emails to quote for new resources, a few emails for approvals, and the time to plan the project and execute. The total hours required to convert our standalone infrastructure in our secondary vDC to an HA environment and add an application server to boot came in just under thirty, with no visits required to any datacenter with boxes, dollies, cables, screw drivers, and muscles in tow. Super slick.

## Conclusion

F5 BIG-IP VE and the Bluelock Virtual Datacenter is a match made in application paradise. If you have any questions regarding the infrastructure, the deployment process, why the sky is blue, please post below.

---

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)