# Data Center Virtualization Q&A

with Erik Giesa, F5 Networks VP, Product Management & Marketing

**Q What's driving the need for data center virtualization?**

A We know that if business continuity is a key objective of an organization, it means that operations are up and running 24x7. Best practices suggest using geographic redundancy to establish multiple data centers or sites located in different geographic regions, each with replicated applications and data. Do you need to replicate everything? No, not necessarily, just those things that are deemed mission critical. Some organizations will feel that the bulk of their applications and data are mission critical, whereas others will have a smaller subset.

**Q How do you implement geographic redundancy?**

A You can implement geographic redundancy in a number of ways. You could choose to deploy multiple sites and use something like Veritas' Smart Location or EMC's Replication Storage to duplicate applications and data, which is a significant investment.

Today, most IT folks still build redundant sites as a backup and manually manage data replication and failover to the secondary site when needed. So they have their site sitting there inert as an insurance policy, but also as a non-performing asset. By virtualizing data center resources at both sites, you can turn non-performing assets (with the exception of a disaster) into an ongoing available asset that will function in a distributed scenario to achieve maximum reliability and performance regardless of location. For example, in an active-active data center configuration, you could do data replication, upgrades, and maintenance on a more-frequent basis, increasing your overall uptime and time-to-market for services.

There are other benefits to virtualization when you look at the data center itself. Let's say you need maximum availability and high performance for your applications and data. You can deploy one, very reliable midrange server with RAID and redundant power supplies that cost you half a million dollars. However, you're still going to have a single point of failure because it's a single system. You could also try to achieve those business objectives by throwing very expensive hardware at it, trusting that all the components will keep running.

A better practice is to virtualize your server and application resources, which is much more cost effective and a better overall architecture. Instead of deploying that very expensive mid-range system, virtualize multiple, low-cost, high-performance servers with applications and data, so when one server goes down, you're not impacted. This gives you the opportunity to achieve high availability and performance without breaking the bank.

**Q What do you have to consider to virtualize your data center?**

A A key part of the answer to this question is the architecture itself. It starts with the application. Can this application be deployed in a manner that it can be virtualized? Does it support clustering or are there tools that help it support clustering so that each application instance recognizes state? If that's the case, that application is a great candidate for virtualization within the broader context of the application delivery network framework.

Can the underlying applications be replicated in real time between redundant sites so that they can resolve requests at any site at any time, ensuring that the data is current? If you can't replicate the data in real time, there might still be an opportunity to virtualize redundant sites if the data being served doesn't require up-to-the-minute freshness. There are a lot of scenarios where that does makes sense. What day-old data is acceptable? Ultimately, you have to look at the underlying application infrastructure to determine what you can virtualize. The same is true for virtualizing connectivity and links.

You also have to consider the amount of data and performance during the replication process. In this case, the primary challenge is not about the bandwidth or link capacity. The challenge is how much of that data can be concurrently transferred or put into the pipe while eliminating protocol communication overhead. We've seen customers with OC-3 connectivity between data centers, but their replication process only uses a fraction of that pipe. They have a lot of data to transfer and it just trickles into the pipe, so replication literally takes days to complete – it's just not efficient. Fortunately, there are solutions out there that use symmetrical WAN acceleration to mitigate this situation. So replication processes that took days to finish now get completed in hours. That's a better model and a better use of the underlying infrastructure, which includes available bandwidth.

**Q What are the benefits of data center virtualization?**

A From an architecture standpoint, there are so many benefits to virtualizing your resources that deliver applications. The savings are profound such as better use of infrastructure, 99.999% availability, and simplified management – it really boils down to better operational efficiently.

With virtualization, there's efficiency in the underlying hardware requirements. In essence, you need less hardware or less expensive hardware to do the same work. You can get 5 times the performance for a third of the cost when you compare a mid-range system to a server farm that cost you about $3,000 a piece. If I can put 10 of those low-cost servers in a virtualized resource pool, I've got 5 to 10 times the power of the most powerful mid-range system at a third of the cost. By virtualizing my servers, I not only realize a tremendous cost savings, but I have a much better architecture for availability and ongoing maintenance. If I need to bring one server down it doesn't impact the others, and I can gracefully add in and take out systems to support my underlying architecture.

For things like ongoing maintenance and management, there are significant efficiencies to be realized. For redundant active-active data centers that are managed by an intelligent DNS system, I can very easily bring down one data center for maintenance without affecting the other data centers or impacting users.

The benefits of virtualization really runs the gambit – ongoing maintenance and management, reduction of hardware acquisition costs, and better architecture for availability, security, and performance. This is why virtualization is really becoming the standard for how you design your IT resources for the future.

**Q What stage do you think most businesses are at in their quest to reach data center virtualization?**

A For complete data center virtualization and multi-data center virtualization, I would say most enterprises are at about 10-15%. Of course, you'll have pockets where people are virtualizing their application instances using things like VMWare or clustering. It's not across the board. But then we see virtualization on the combined server and application level using products like F5's BIG-IP Local Traffic Manager. Virtualization is pretty much a given in the Web world. At the application server and database level, it's still nascent. There's still a lot of room to build that proven model at the Web tier down in the application and data tiers. However, most of our large enterprise customers already have some form of SAN virtualization in production. You also see virtualization principles at work as far as connectivity, topology, and access. So virtualization really isn't a new concept.

What is new is thinking about all the points in the WAN and LAN infrastructure where you can realize virtualization benefits regardless of where you started out. Consider this – you want your worldwide employees to securely access your network and applications at any time from any device and from any location. Sometimes sites go down for maintenance, connectivity reasons, and disasters. If you provide worldwide access, but it's only available 95% of the time and it's under performing 98% of the time, you're not really achieving your goal of worldwide access 24x7. This is where virtualization integrated with access technologies like SSL VPN comes into play. Virtualization of distributed access devices that route users to the best possible site, which hosts your SSL VPN access control, gives them access to applications and network resources without any interruption of service. Routing users to the best available site is completely transparent and doesn't require updating client software or reconfiguring clients, which we know is fraught with problems. Again, virtualization is a better model. So think about virtualization from a holistic architectural approach to fully realize its benefits.

**Q How does F5 products help organizations virtualize their data center to leverage the benefits of virtualization?**

A When it comes to virtualization, F5 believes in the notion of abstracting the layer up from the applications and leveraging the network to perform services on the applications' behalf. From this layer, F5 products can deliver end-to-end application security, performance, and availability in the most operationally efficient manner. F5 enables organizations to realize the benefits of virtualization by giving them the underlying architectural building blocks to create an infrastructure that supports virtualization.

### Availability

You can implement a secondary data site and use the BIG-IP Global Traffic Manager to monitor application health to automatically failover to another data center, maintain service interdependencies in a SOA application, and automatically route users to the best site. Even if a user's closest, highest-performing site is unavailable, the user is transparently routed to another site for the best possible experience.

Say you have a need for multi-homing, but you have to provide reliable connectivity to your applications for both public (customers) and private (employees) use. The BIG-IP Link Controller monitors your ISP links to automatically route inbound and outbound traffic to the best-performing link based on link health, performance, and cost, business policy, user location, or topology. To ensure the quality of service for any user regardless of their endpoint, integrated Rate Shaping reserves bandwidth for priority applications and compression accelerates the performance of the session over those links.

Although the BIG-IP Local Traffic Manager uses load balancing to keep your applications available across your networked devices, the most common scenario is distributing application traffic across an array of web servers that host standard web traffic, including e-commerce traffic. The BIG-IP load balances HTTP connections to the Internet or the company's intranet web applications as well as to HTTPS. Using F5's TMOS full proxy architecture, you can apply a virtualization model across any type of application whether it's a SIP, RTSP, TCP, XML or UDP-based application. If it's IP traffic, we can see into the application stream and apply security, availability, and acceleration services.

### Security

Instead of spending so many cycles, resources, and effort on securing each and every potential vulnerability, why not virtualize your IT resources and centralize their management by defining security policies based on who is requesting access (client type, endpoint security, integrity of client, SSL credentials), the type of device requesting access (operating system, firewall, AV), the type of encryption (3DES, AES), and what transactions are allowed (what is authorized by that user's role), and what is allowed for the application (inputs, characters, links, cookies, etc.) This saves a tremendous amount of time, gives you a better security posture, and improves your auditing capabilities because it centralizes these functions. You can do this using the BIG-IP Local Traffic Manager with the Application Security Manager module and FirePass to not only protect your applications and network resources, but also provide secure remote access to your resources no matter where your employees are located.

Virtualization also places an additional layer in front of hackers. F5's TMOS full application proxy sits between hackers and any BIG-IP device to cloak the IP addresses of your resources. That's a huge benefit because hackers are not talking directly to your resources.

### Performance

Although you have no control over public users' endpoints, there are things that the BIG-IP can do asymmetrically to benefit client sessions using technologies like TCP/IP optimization, compression, caching, and offloading your servers to do SSL processing and client authentication.

You can optimize cross-WAN intra-application communication occurring on the backend and accelerate and optimize those connections by deploying F5's WANJet end to end. In this way, you can use WANJet to optimize data replication across data sites, bringing data consistency much closer to real-time in almost all situations. To significantly reduce the time and bandwidth required to synchronize multiple sites, WANJet works with underlying technologies such as SAN replication, database log shipping, transactional replication, file copies, etc.

To be productive, mobile workers depend on accessing their critical business applications via the Web. F5's WebAccelerator uses an asymmetric deployment to optimize Web applications for users where you can't or don't want to deploy an endpoint device. WebAccelerator's employs unique technologies such as Express Loader and Express Connect to make inefficient protocols like HTTP more efficient, which dramatically improves performance without making changes on the client side or in the application itself. Again, it's a better model for achieving your performance objectives.

### Conclusion

When you consider virtualizing your IT resources, you have to consider all the critical junctures of your network topology. What is your current environment? Do you have multiple data centers, do you currently multi-home or provision multiple ISP links from different providers? Do you have applications that you want or could virtualize? Where are your users coming from, the branch office, overseas, or remotely from the road? And are those users private employees, public users, contractors, suppliers, and customers? What devices are they coming from to access your applications? Do you want to accelerate and secure application sessions across this wide variety of usage scenarios? And, of course, what are your business goals, objectives, and SLAs?

Enabling you to virtualize your resources is the greatest value of F5's product strategy. Regardless of where you're starting out or your constraints, F5 has the solutions that function as architectural building blocks to virtualize your resources. With F5, you can leverage the benefits of virtualization and keep your applications secure, fast, and available in the most operationally efficient manner.

### About Erik Geisa

Erik is the Vice President of Product Management and Marketing for F5 Networks, where he is responsible for driving product management and marketing strategy for all products families, including BIG-IP, FirePass, TrafficShield, WANJet, WebAccelerator, and iControl. While at F5, Erik has successfully driven the product interoperability and integration strategy with key vendors like Oracle, Microsoft, Siebel Systems, and BEA, leveraging F5's unique iControl API as an open standard to integrate F5's products with their enterprise applications. In addition, Erik and his team have driven and launched some of the most innovative and profitable products for F5 including the award-winning BIG-IP platform with the unique TMOS architecture. Prior to F5, Erik led product management and marketing teams at Holistix, Inc., a web systems management company, and WRQ, Inc., where he launched and ran the fastest growing product line in WRQ's 15 year history.

## About F5

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast, and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 9,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to **www.f5.com**.

**F5 Networks, Inc.**
**Corporate Headquarters**

401 Elliott Avenue West
Seattle, WA 98119
(206) 272-5555 Voice
(888) 88BIGIP Toll-free
(206) 272-5556 Fax
www.f5.com
info@f5.com

**F5 Networks**
**Asia-Pacific**

+65-6533-6103 Voice
+65-6533-6106 Fax
info.asia@f5.com

**F5 Networks Ltd.**
**Europe/Middle-East/Africa**

+44 (0) 1932 582 000 Voice
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

**F5 Networks**
**Japan K.K.**

+81-3-5447-3350 Voice
+81-3-5447-3351 Fax
info@f5networks.co.jp