F5 White Paper

VMware DRS: Why You Still Need Assured Application Delivery and Application Delivery Networking

VMware Infrastructure products provide the next generation virtual platform for the new data center, but they don't virtualize the network or application delivery. F5 BIG-IP LTM works with VMware to provide truly virtualized Application Delivery Networking.

by Alan Murphy
Technical Marketing Manager, Virtualization

Contents

Virtual Machine Mobility

Virtual Operating Systems, or virtual machines (VMs) as we've come to know them, started out as isolated containers forced to run on individual physical machines. These physical machines were usually desktop devices, and virtual machines enabled desktop users to run multiple operating systems locally. This was a great technology for trying out different operating systems, testing new deployments, and even adhering to corporate IT policy while running another OS. These technologies were limited in data center adoption, however, due to their explicit requirement of VM location. Once installed and running on a particular physical machine, a VM had to stay on that machine unless it was powered down and manually moved to another physical machine and spun back up. And due to their resource consumption, running many virtual machines on one physical machine in the data center didn't scale as well as expected. Ultimately, this became more expensive to support and deploy than it saved in capital costs.

Fast forward five years and the virtual machine landscape is quite a bit different and much more advanced. Products were developed and released—such as VMware ESX and VirtualCenter—that have moved virtual machines from the desktop to the data center and decoupled the physical requirement of running a virtual machine and virtual disk (VMDK) on singular physical hardware. In addition, it is now the standard to physically separate and remove the computing portion of the virtual machine (CPU, RAM, I/O) from the virtual disk image; the running machine is hosted on a standard server while the virtual disk is stored on some type of shared storage and delivered as needed to the virtual machine via the storage network. Due to these architectural changes it is now possible, and extremely easy, to move running virtual machines between hosts without worrying about the physical disk image or losing state information. Further advances in VM mobility technology enable the same migration to occur with the virtual disk without impacting the running instance of the image. A VMDK file can be moved from one storage disk to another while the virtual machine continues to run uninterrupted on a single host. VMware has led the virtual pack in bringing products to market that facilitate and manage virtual infrastructure migration, however these products focus solely on virtual machines and disks. Like many other virtualization technologies in the data center they remain in silos, ignoring issues such as the network and end-user experience.

Elastic Data Center Limitations—Scaling the Cloud

By creating a near-complete replica of the physical data center on top of virtual platforms, VMware has revolutionized the role of virtualization in the data center and opened the door for what many people are calling the New Data Center. Like all revolutionary changes though, VMware does expose some major limitations in how virtualization is managed with other technologies in the data center.

DRS—Where's the Network?

One of the critical feature functions lacking from VMware products such as DRS and HA is factoring in network-based information as part of resource-scaling decisions. As DRS and HA monitor VM performance across multiple hosts, they are only monitoring what are known as computing resources: CPU and RAM. These two factors are influenced by the guests on each ESX host in the form of running processes, I/O tasks, disk access, and other virtual machine resource requests on the same host. Since networking is not a computing resource, DRS doesn't consider host or VM network resource information when deciding where and when to move a VM with VMotion. Networking constraints are not classified as computing resource shortages.

For example, a VM on one ESX host may consume too many CPU cycles by terminating and managing SSL connections for a secure web application. DRS will correctly detect this CPU resource spike and relocate the SSL-processing VM guest to a new host, one that is using fewer CPU resources for its other running guests. Although other guests on the destination host may be using less CPU and RAM, those VMs could be handling and managing thousands of network connections and saturating the vSwitch, virtual network, and virtual ports on that host. Network saturation and load aren't factors in VMware's resource calculations; DRS will drop a VM running the SSL web app on a host that has no available networking resources to give to that guest. In this example, DRS has moved a CPU and network-heavy VM from one ESX host to another, ultimately taking away resources from the migrated VM rather than freeing up resources.

The key to solving this type of problem is to factor in network information along with CPU and RAM usage. Application servers need both computing and networking resources whether they are physical or virtual machines. The

combination of CPU, RAM, and network resources are critical even though they aren't always interdependent. It's important to discretely monitor each of these components as part of an Application Delivery Network. By integrating networking knowledge and specific data on how applications are running across the network, BIG-IP® Local Traffic Manager™ (LTM) alleviates many of the network challenges virtualization administrators face by implementing DRS and HA alone.

Network Virtualization

BIG-IP LTM virtualizes the network for products such as ESX, DRS, and HA; all application networking data sent to and from virtual machines passes through BIG-IP LTM. Key metrics for network utilization, such as bandwidth utilization, latency, connections, application availability, and so on, are managed by BIG-IP LTM for every pool member and Virtual Server. This information is published via iControl®—the F5 SOAP-based API—and SNMP and can be consumed by VirtualCenter in order to make intelligent decisions based on network information. VirtualCenter can instruct DRS to move an image to a host with the least amount of network load, based on the amount of network traffic data to each guest running on that host, as well as CPU and RAM load. If new images are needed to mediate resource constraints on existing VMs, these images can be dynamically added as BIG-IP LTM pool members once they are spun up by VirtualCenter and are available on the network.

Application Awareness

One of the most important data center components is mostly ignored by OS virtualization and VMware products: applications. Most virtual infrastructure technologies are designed for just that, infrastructure; yet applications running on virtual machines are at the mercy of the newly virtualized infrastructure. The F5 BIG-IP platform focuses on the application and how those apps are running on the network through the data center. When placed in front of application servers running on virtual machines, BIG-IP LTM can manage connections, traffic to individual VM applications, and also load balance as resources become constrained. As DRS and HA move images around the virtual platform infrastructure, BIG-IP LTM proxies those connections as they're coming in, allowing DRS/HA to move the image to a new host without losing application data or connection state information. Even if the machine is unable to handle new connections due to CPU or RAM issues,

BIG-IP LTM will divert connections to another virtual machine as necessary, making a network remediation based on computing scarcity. BIG-IP LTM can also provide application-specific health status to VirtualCenter as part of the overall network status, providing the same level of security, optimization, and availability provided for physical machines.

Resource Offloading

Once machines are moved from physical to virtual environments, any new resource constraints tend to have an immediate impact on applications running on the virtual machines. For example a single ESX host may be hosting 20 virtual machines, each running SSL-enabled web applications that need to all share CPU cycles to terminate those SSL connections. By offloading key network services, such as terminating SSL at the edge or by implementing intelligent caching and compression, BIG-IP LTM can take some of the application processing needs off of the virtual machine, host CPU, and RAM, and provide more processing power for the applications. If more resources are offloaded from the virtual machines and their running applications, there is less need for tools like DRS to manage the individual guests and hosts because there are fewer resource droughts. Processes such as caching and compression can be moved off of the virtual machine applications and on to BIG-IP LTM, allowing the virtual servers to serve new content as needed rather than spinning cycles and flooding the virtual network with redundant and bloated data.

Connection Management

When DRS or HA migrates a machine with VMotion, only Layer 2 information, such as the MAC address and port assignment, is moved with the virtual image from one virtual switch to another; this allows new network connections and application requests to find the relocated guest on the new ESX host and vSwitch. Connections that are application-specific, however—such as SSL persistence, host-based connection management, shopping cart data, or any process that depends on the state of the connection to the application server—could be lost in the VMotion migration. No Layer 3 (IP addresses), Layer 4 (protocol and port), or Layer 7 (application) data is migrated with the VM. BIG-IP LTM proxies connection data across all networking layers for applications running on virtual BIG-IP LTM machines. As new application connections are managed by BIG-IP LTM, it will verify that the application is running and available on a particular VM pool member. During a VMotion migration, the VM being moved will not be available as a pool member; BIG-IP LTM will direct connections to one of the

other available pool members. BIG-IP LTM will also manage stateful connections for the migrated machine and proxy those until the particular image is available again as a pool member. If a user passes shopping cart information to an application running on a VM and that VM is migrated to another ESX host, vSwitch by DRS, and VMotion, BIG-IP LTM will store that information until the host comes back online—with no noticeable impact to the user's session.

Conclusion

Just like physical data centers, Virtual Data Center platform products, such as VMware ESX, DRS, HA, and VMotion, are implemented to run applications. Yet these new platform technologies focus on virtualizing the infrastructure hosting the applications; for the most part, these products aren't concerned with the applications themselves. Just like a standard L2 switch doesn't care what IP addresses are flowing through each port, virtual platform products don't care what applications are running on the virtual machine guests. In order to fully take advantage of these new technologies, application delivery availability still needs to be the ultimate goal and focal point for these new data centers. Utilizing an Application Delivery Controller such as BIG-IP LTM for network and application server virtualization provides the missing component to the Virtual Data Center: virtual network infrastructure. When coupled with VMware's data center platform tools, BIG-IP LTM provides application networking security, optimization, and availability for the virtualized infrastructure. These tools work together to create the next-generation data center and enable a true Virtual Data Center to be fluid, mobile, and have limitless growth.

F5 Networks, Inc.
Corporate Headquarters
401 Elliott Avenue West
Seattle, WA 98119
+1-206-272-5555 Phone
(888) 88BIGIP Toll-free
+1-206-272-5556 Fax
www.f5.com
info@f5.com

F5 Networks
Asia-Pacific
+65-6533-6103 Phone
+65-6533-6106 Fax
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
+44 (0) 1932 582 000 Phone
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

F5 Networks
Japan K.K.
+81-3-5114-3200 Phone
+81-3-5114-3201 Fax
info@f5networks.co.jp

VMware DRS: Why You Still Need Assured Application Delivery and Application Delivery Networking 08/08