



White Paper

Keeping Your Head Above the Cloud: Seven Data Center Challenges to Consider Before Going Virtual

One of the pillars of the Virtual Data Center is virtual platform infrastructure, or the virtual machine; however virtual platforms are dependent on many other, oft forgotten components of both the physical and virtual data center.

by Alan Murphy

Technical Marketing Manager, Virtualization



Contents

Virtualization Changes Everything	3
<hr/>	
Virtual Machine Deployment Challenges	4
In Case of Emergency, Break Glass	4
<hr/>	
Catching the Falling Sky	9



Virtualization Changes Everything

Referred to by many names—OS virtualization, server virtualization, kernel virtualization, hypervisors—virtual machine platforms, such as VMware ESX and Microsoft Hyper-V, are typically the first form of virtualization introduced to the data center. These technologies are becoming more inexpensive and less complex every day, making them attractive technologies to be the first virtual “guinea pigs” in the data center. Factors such as consolidation, cost savings, dynamic provisioning, and fluid migration are driving most IT shops to experiment with some form of virtual machine product today, and the advent of large-scale infrastructure systems has moved these back-room experiments and development environments into full, public-facing application infrastructure systems.

Throughout the past two years we’ve seen these products deliver complete micro-data center solutions, almost a “data center in a data center” model (think Russian nesting dolls), beginning the prophetic movement towards changing the way applications compute in the data center. These new infrastructure platforms provide virtual machine (VM) management (OS virtualization), virtual switching and routing (network virtualization), virtual storage in the form of flat-file VM disk images (storage virtualization), and system management with tools like Virtual Center (management virtualization) all in one conveniently packaged product.

Although products like VMware’s ESX and Virtual Center provide multiple tools to start building out the virtual data center (VDC), successfully implementing these technologies typically requires a re-architecture of the existing data center. Virtual machine platforms are possibly the most disruptive data center technology of the past 20 years (since the migration and standardization to Ethernet and IPv4). Adding a virtual machine platform is almost like adding a new floor in the middle of a completed skyscraper, without moving or demolishing any part of the existing structure.

In spite of all of the tremendous benefits they bring to the data center, virtual machines also add complexity, scale, and management challenges. Most of the time, these challenges appear in other parts of the data center, such as on the application delivery and storage networks. OS virtualization can be a powerful tool to improve efficiency, cut costs, and increase agility; however implementing virtualization without considering its impact on surrounding IT resources can lead to catastrophic failures, much like trying to shim in the new skyscraper floor without reinforcing the floors above. While the virtualization of physical machines (often referred to as P2V, physical to virtual) can put a tremendous strain on

“Virtualization is a proven software technology that is rapidly transforming the IT landscape and fundamentally changing the way that people compute.”¹



application and storage networks, it doesn't have to cause disruptions within your critical infrastructure. By planning for the migration to virtual machine platforms and recognizing the challenges that are inherently part of the virtualization process, you can achieve a seamless virtual migration, managed as part of the process rather than a clean-up effort after the skyscraper has collapsed and crumbled to the ground.

Virtual Machine Deployment Challenges

In Case of Emergency, Break Glass

There are typically seven areas of concern, or pain points, that accompany any major virtual platform implementation or migration. Depending on the scope and nature of the project (size, scale, production versus staging, and so on) these critical areas may either appear immediately or over time, potentially months to years later. Often these issues aren't seen during staging and testing and only appear when the virtual machines must take on the same production load as physical machines. These critical pain points represent two cornerstones of the data center: network and storage.

Although this list of pain points can seem a bit daunting and may seem to shine a negative light on virtual platforms, all of these issues can be overcome by understanding, addressing, and planning for them before and during the virtual migration process. To avoid surprises and show-stoppers during a production it is important to educate yourself on what impact these virtual platforms will have on your data center. In addition, inter-departmental communication and collaboration is also critical to solving new virtualization issues, as many of the concerns discussed in the following sections cross typical IT barriers, requiring that members from different teams work together to solve these problems. You will mostly like find that all IT groups—network admins, system admins, application experts—will need to be part of virtual migration planning and deployment.

Lack of Performance and Availability—Resource Starvation

Virtualization moves many I/O tasks tuned for hardware to software via the hypervisor. For example, if an application is optimized for a particular chipset, once the OS is virtualized that chipset becomes a virtual software component and many of your optimization efforts can be lost. The virtualization translation

IT Team Ownership: Network, Application, Server

Benefit of Solving: Unlimited and dynamic virtual machine scale



layer is now responsible for translating that optimized code for the software chip then back to the physical chip or CPU running on the underlying hardware. I/O intensive applications, like cryptographic processing applications for SSL, don't fare well when virtualized because of this translation layer.

Outside of the virtual hardware layer, these performance issues can cause application and storage network resources to be depleted at a much faster and often unanticipated pace. Virtual machine saturation (too many virtual machines running on one physical host) caused by virtualization sprawl can cause unanticipated resource constraints in all parts of the data center. With a physical machine running a network application, that application can have access to the full resources of the network card. Once 10 virtual machines are running together on a single host—sharing a single hardware network card and streaming together through a software switch—that same virtualized application now has fewer networking resources available than it did before. This can lead to overall network performance issues, reduced bandwidth, and increased latency; all issues the application might not be able to deal with. Even smaller issues such as IP address availability can be impacted by virtualization sprawl.

Lack of Application Awareness—OS Virtualization Doesn't Solve Application Virtualization

One of the limitations of hypervisor- and kernel-based virtualization solutions are that they only virtualize the operating system. OS virtualization does not virtualize nor is it even aware of applications that are running on the OS. In addition, those same applications don't realize they are using virtual hardware on top of a hypervisor. By inserting this extra software management layer between the application and the hardware, the applications might encounter performance issues that are beyond their control.

Virtual infrastructure platforms typically include software that can migrate live virtual machine instances from one physical device to another; VMware Distributed Resource Scheduler (DRS) and VMotion are examples of live migration solutions. Like basic OS virtualization, these migration tools are unaware of the application state, and also have no insight into the Application Delivery Network. For example, VMotion may move a virtual machine running a web shopping cart application from one host server to another without taking into consideration the current persistent state of user's carts, how many connections are coming into the shopping cart, or the network load that this particular virtual machine is currently handling. This migration can cause a lapse in availability; connections to the application and application persistence for the user can be lost during this live

IT Team Ownership: Network, Application, Server

Benefit of Solving: Application awareness as part of the virtual machine fluidity and migration



migration, resulting in failed transactions, lost shopping cart data, and frustrated users. And while VMotion attempts to move the live image to a host server that is under less load, VMotion doesn't measure the network load of that host device. It might move a live image to a machine with more available CPU cycles but less available network capacity.

Additional, Unanticipated Costs—The Virtual Solution May Cost More than the Physical Problem

Two of the primary drivers for virtualization are cost reduction and data center consolidation, however implementing a complete virtual machine solution in the data center and migrating physical machines to virtual machines does not come cheap. On the surface this is a real value proposition for virtual platforms: fewer physical servers mean less capital expenses, less hardware, less rack space, less cabling, less cooling, less energy, and so on. The cost savings associated with the physical to virtual migration is not linear, however. Capital expenses such as new hardware platforms to support large numbers of virtual machines, licenses for the virtual software platforms, and—in the case of virtual sprawl—OS licences for the operating systems running on the virtual machines can quickly outscale the purchasing budget. Once virtualization hardware and software is acquired, operational expenses can grow unbounded; headcount can increase or existing staff may require training to administer the new virtual machine platforms. Management of these new tools can be a long-term recurring cost, especially if the virtualization is done in-house. There can be additional growth requirements for the application and storage networks as these virtual machines begin to burden the existing infrastructure. Unexpected and unplanned costs can be a serious problem when implementing or migrating from physical to virtual machines, hindering or even completely halting deployment.

Unused Virtualization Features—Batteries Not Included

New virtual platforms include many advanced networking technologies, such as software switching and support for VLAN segmentation. These features, however, are localized and isolated to the virtual machine platforms and not integrated with the rest of the network infrastructure. An enterprise may have purchased VMware specifically for the DRS and VMotion live migration features, but then realize that their current networking infrastructure can't support these new features, or that they have to re-architect physical portions of their network to take advantage of virtual software switching features. Often, new technologies perform flawlessly in development and staging but they are unable to scale to production levels once deployed. These new platforms may have problems integrating with existing application and storage networks, requiring a re-design of the data center from

IT Team Ownership: Server, Management

Benefit of Solving: Limited exposure for CapEx and OpEx budget overruns

IT Team Ownership: Network, Server

Benefit of Solving: Full functionality of virtual machine platforms, complete flexibility



the ground up in order to implement one virtualization technology.

Storage integration issues tend to arise as soon as virtual machines are moved into production environments. First and foremost, network storage is a requirement for virtual platforms that implement live machine migration; direct attached and local storage will only work for running local virtual machines. While many enterprise storage networks include technologies for data replication that span multiple geographic data centers, virtual machine migration tools are often limited to local storage groups. VMware DRS and VMotion are unable to migrate live machines outside of the local storage domain unless additional storage virtualization tools are implemented between Virtual Center and physical storage.

Overflowing Storage Network—Managing Explosive Growth

Although converting physical machines to virtual machines is an asset for building dynamic data centers (fluidity, disaster recovery, and backups), hard drives become extremely large flat file virtual disk images—a typical Microsoft Vista Ultimate install consumes ~15 gigs of local storage for the OS alone, not counting data files. Consequently, file storage can quickly become unmanageable. Operating system and data files that typically reside on internal storage in physical server environments are often moved to shared network storage during the virtual migration process. This can challenge the extension of existing storage as well as break growth models that may already exist for storing data on the network. In a recent report from Enterprise Strategy Group, 54 percent of the customers surveyed reported some increase in storage capacity specifically tied to OS virtualization and the installation of virtual platforms. Eighteen percent of responders reported an increase of physical storage greater than 20 percent beyond their storage needs for physical servers².

An enterprise may have a model in place for how much growth is expected and how to provision that growth, but might not factor in the exponential size requirements needed for virtual machines, or be able to spin up enough new storage hardware to support a new virtual cluster. Virtual machine platform sprawl can dramatically increase the number and size of unused files as virtual machines are spun up and spun down throughout the virtual data center. It is common architecture to archive and store “gold” virtual images that are cloned as needed into function-specific images and copied to another part of the storage network. This architecture can create “parked” images on shared storage, once gold images are upgraded or abandoned clones, that are very infrequently used.

IT Team Ownership: Storage, Server

Benefit of Solving: Scalable virtual machine storage, reduced OpEx for new storage hardware



Congested Storage Network—Clogged Data Pipes

By design, the portable nature of OS virtualization can dramatically increase data traversing the storage network. For the same reason that virtual machine disk files can overrun physical storage, once these images are made portable it becomes trivial to move these VM images across the network from one host to another or from one storage array to another. It can be a challenge to prevent flooding of the storage network when planning a large VM migration or move. And as virtual sprawl and unmanaged virtual machines begin to appear in the data center, unplanned virtual machine migrations can literally bring the network to a standstill, even on a LAN.

Moving these large virtual disk images outside the LAN can cause extensive delays and flood what are typically much smaller WAN connections. In addition, the type of shared storage used behind the virtual infrastructure impacts and may amplify data migration performance issues. NFS, for example, is typically more susceptible to performance degradation than other storage networking solutions such as Fibre Channel or iSCSI.

Management Complexity—Breaking the Single Pane of Glass

Throughout all areas of the data center, managing virtual machines as part of the complete management solution can be a struggle at best. If you are already using tools to manage the application and operating environments of each physical machine—either by using clients on the systems or by measuring application metrics on the wire (such as latency, response time, and so on)—management of the virtual machines themselves doesn't have to change your management infrastructure. Virtual machines will report the same types of metrics as physical machines. The management challenge with virtual machines appears in two forms:

1. The addition of two new components that need to be managed: the hypervisor and the host system. Neither one of these devices exist in the physical server world are not part of existing data center management solutions, but they do have a major impact on the performance of virtual machines. Managing these devices and insight into these performance differences is critical.
2. Managing your virtual machines, application network, and storage network together: Many virtual machine platforms include built-in management tools, some of them highly sophisticated, such as VMware's Virtual Server. While these tools do provide essential management tasks, such as live migration of virtual guests from one host to another, they don't take into account any external information. For example, VMware DRS can invoke a virtual machine live migration based on the CPU and memory constraints of

IT Team Ownership: Network, Storage, Server

Benefit of Solving: Bandwidth costs, efficient bandwidth usage, data mobility speed

IT Team Ownership: Server, Management

Benefit of Solving: Unified Virtual Data Center management, service virtualization



the host, but it doesn't look at the internal network (virtual switch) nor the external network (physical host port).

With physical servers, there is a line segregating ownership and management responsibilities. The network team owns the network fabric (switches, routers, VLANs, IPs) and the server team owns the servers (hardware, OS, application, uptime). OS virtualization blends these responsibilities onto one platform, blurring the lines of ownership and management. Switches are now software instead of hardware and part of the same core systems that run the server OS and application code. Do system owners inherit the virtual network, or do network owners inherit the virtual systems?

Catching the Falling Sky

Despite the challenges, implementing virtual infrastructure platforms isn't all bad news. It is possible to implement a successful virtual infrastructure migration if the above considerations are addressed before, during, and after the virtual roll-out. Most importantly all parts of the data center must be a factor in the transition plan. It's impossible to be successful if the only action taken is to replace physical machines with virtual clones without considering other physical and virtual components of the existing data center. If the goal is to move to a complete Virtual Data Center, then stepping through the five states of virtualization in the VDC Maturity Model³ requires managing each new virtual component individually and then integrating that new technology into the rest of the data center. This will ultimately lead to a virtualization-ready IT infrastructure—one built to be fluid, dynamic, and provisionable. A data center then can expect and accept hundreds or thousands of new servers, new IP addresses, new routes, new disks, new files, new storage servers—temporary or permanent, virtual or physical.

Proper design, planning, knowledge of the impact of virtualization on the data center, as well as an understanding of how to implement virtual machines in a manner that is complimentary and compatible with the rest of the data center is critical for successfully moving toward a truly Virtual Data Center. Once these issues are addressed they can easily be managed and factored into the migration plan, ultimately resulting in successful virtual machine and platform deployments and a dynamic, provisionable Virtual Data Center.

¹ VMware, <http://www.vmware.com/virtualization/>

² ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007

³ "The VDC Maturity Model – Moving Up The Virtual Data Center Stack", <http://www.f5.com/pdf/white-papers/vdc-maturity-model-wp.pdf>

White Paper

Keeping Your Head Above the Cloud: Seven Data Center Challenges to Consider Before Going Virtual



**F5 Networks, Inc.
Corporate Headquarters**

401 Elliott Avenue West
Seattle, WA 98119
+1-206-272-5555 Phone
(888) 888BIP 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