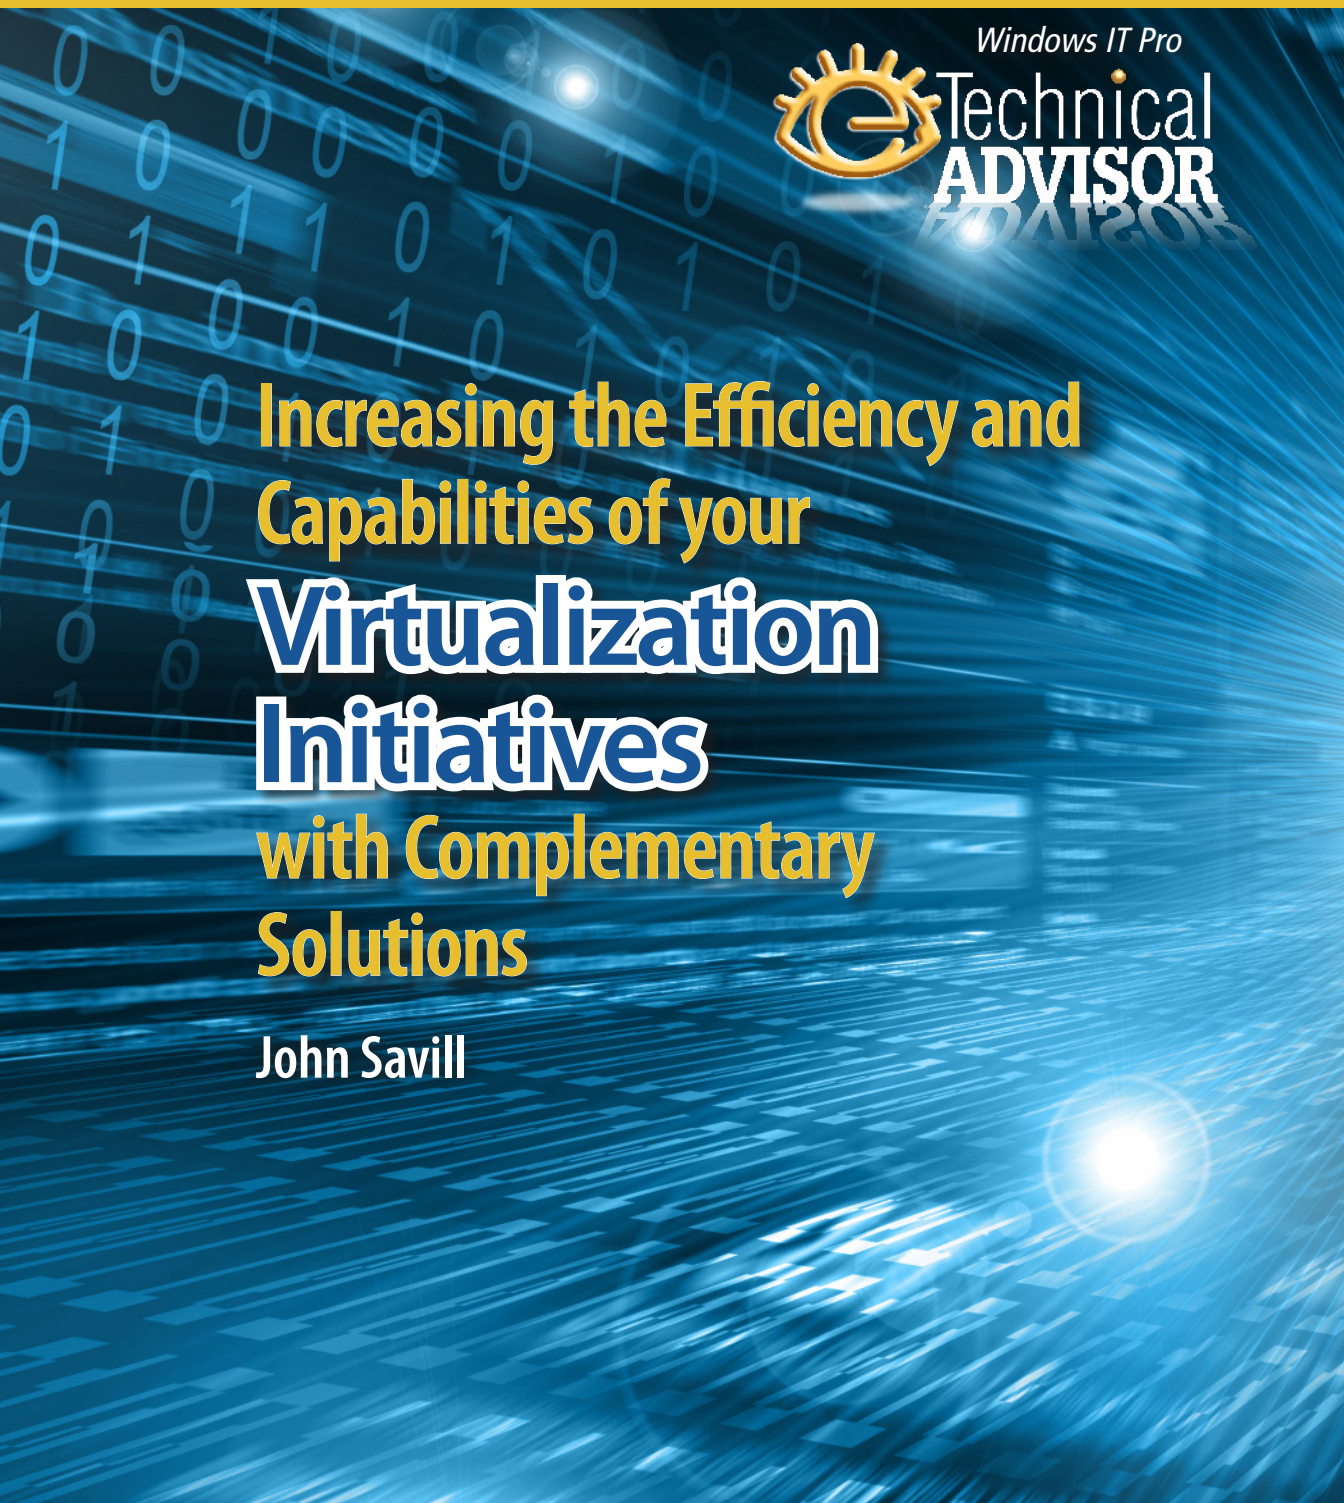

Increasing the Efficiency and Capabilities of your Virtualization Initiatives with Complementary Solutions

John Savill

sponsored by ca technologies ARCserve®

CONTENTS

Increasing the Efficiency and Capabilities of your Virtualization Initiatives with Complementary Solutions

By John Savill

Virtualization has truly ushered in a new age of datacenter architecture. Organizations are able to reduce the amount of hardware required, increase the utilization of resources, and cut down on costs associated with facilities, power, infrastructure, and IT operations—all while increasing availability and resiliency for systems and applications.

Hypervisors on the market today offer technologies to maximize the density of virtual machines in the hosts, enable high availability, simplify management, and increase business capabilities. In this Technical Advisor, I'll explore the path organizations can take to achieve virtualization nirvana, which includes three main stages:

1. Adopting server virtualization for new systems and migrating existing physical systems
2. Managing, protecting, and optimizing the virtual environment
3. Enabling better service availability and resiliency through virtualization of non-virtual systems

Moving to server virtualization

Adopting server virtualization initially starts very carefully as the technology is evaluated and tested to ensure it meets the needs of the IT organization and business. Adoption then proceeds fairly quickly as IT teams migrate the easy-to-migrate physical boxes (the low hanging fruit), and non-critical systems. Typically, however, after the initial fast adoption of virtualization, migrations stall. Often, IT teams are stalled by a lack of detailed information about systems and an uncertainty about the best way to proceed. And when dealing with production systems, IT teams are hesitant to perform any action that might impact production availability and their long-term employment.

Understanding your existing environment is critical to a successful virtualization migration and to keeping the virtualization momentum going. Understanding not only the servers you have but also their operating system version; the applications that run on them; configured resources such as CPU, memory, disk, and network connectivity; and most importantly the actual resource utilization of servers can help you best understand the right ways to consolidate and identify systems that would coexist well.

Microsoft provides a free discovery tool for your environment that will provide details about the system configuration, the utilization statistics gathered over a period of time, and suggested consolidation configurations. The tool, the Microsoft Assessment and Planning Tookit (MAP), is a free download that requires only a single desktop or server operating system to host the MAP server component and database. MAP runs in an agentless mode for all the servers in the environment and gathers information using a passed set of credentials that have administrator privileges on the target boxes. Using WMI, MAP gathers information and performance metrics at a defined time interval for a configured duration.

Other tools are available to gather system information and utilization. Here are key factors to consider:

- Run utilization statistic collection for a sufficient amount of time to get a true sample of system usage in normal and busy times. If systems have peek loads at week or month end, it's important to gather usage data during those times or the consolidation guidance will be based on incomplete data which will likely cause issues once the systems are virtualized in production.
- Don't set performance gathering too frequently because any monitoring can impact performance on the monitored system. Typically, this would be a very low load, but you should also consider the additional load on the network if you gather very frequent performance data from every system in your datacenter.
- Where possible, automatic gathering of system dependencies is very useful—or you should manually create these relationships between systems because this will be important when you virtualize. You'll need to decide

whether certain systems should be co-located on the same virtualization host for best performance or kept on different hosts to maximize availability. For example, if you have three domain controllers you would want to make sure all three are not running on the same virtualization host. You should configure anti-affinity rules when you virtualize the domain controllers to make sure the domain controllers never run on the same host.

Once you understand the servers you have in your environment, their utilization and relationships between them, you can architect how to place your virtual machines, how to divide your virtualization hosts into highly available clusters, and which groups of VMs to place on which cluster groups. When determining the virtualization architecture a major focus used to be on which VMs to place on which hosts. That consideration is still a part of the process today, but current virtualization solutions include technologies to automatically balance virtual machines on all the hosts in a cluster for best performance. This means that even if you put VM A on Host 1 it's entirely possible that the automated distribution technologies of the virtualization solution will move VM B to host 2, if that would better balance resource utilization. So while it's important to understand the load of VMs for total cluster sizing, it's more important to configure rules for virtual machine affinity (keeping VMs together) and anti-affinity (keeping VMs apart). When VMs are automatically balanced you maintain your relationships to ensure best performance and availability.

You might have some systems that use a lot of resource, such as all the processing and memory of a physical system, and you would not consolidate them with other systems. Should you leave those very large systems as physical instances? The answer is, "It depends." If your virtualization platform can present all the machine's resources such as all the processing cores as virtual CPUs and all the memory to a single VM (except for the amount of resource used for the hypervisor and management partition), then virtualizing the system can still be a good idea. Although you don't gain consolidation you do gain abstraction of the operating system and applications from the underlying hardware, which gives you portability for the system. This portability in itself is a huge benefit of virtualization and removes dependencies on specific types of hardware.

In terms of actually performing your physical to virtual migration, you'll see three main types of systems:

1. Systems that are easy to migrate, such as file servers, domain controllers, and network services. Essentially, these are systems that are fairly self-contained and that you can easily move to a VM or create a new instance of the system in a VM and turn off the physical instance (which is common with domain controllers). It's important to note that if you migrate your testing systems to virtualization and not the production systems, the testing you perform won't be representative of production. You should keep your test and production systems on similar platforms where possible. Moving testing systems to virtualization to ensure functionality in a virtual environment makes sense and should be part of your migration to virtualization.
2. Systems that are more complex, such as Exchange, SQL Server, and other multi-tiered applications that are more business critical but for which guidance exists on virtualization best practices.
3. Custom applications that no longer have support capabilities within the organization and third-party applications where virtualization guidance does not exist.

Organizations will work their way down the list outlined above to migrate from physical to virtual and it's very important as you start hitting the second and particularly the third groups to do research to check for virtualization guidance and also support statements for systems to find if virtualization is even supported. For systems that might not support virtualization, you must decide whether migration to virtual will still be performed and it's important to understand the potential support complexities associated with running in an unsupported configuration and the likelihood a vendor will ask for the problem to be recreated in a supported (i.e. physical) configuration. This is one area where complementary solutions to virtualization can bring real value. We are about to explore going physical to virtual (P2V), which is necessary to move to virtualization. But being able to go virtual to physical (V2P) is a great capability to have in these support situations and many other scenarios. With V2P you can take a virtual machine and migrate it back to physical hardware, getting the application back into a vendor-supported state.

Following the discussion above, let's say you've identified which servers to virtualize, documented the relationships between systems, and understood the best practices and support statements. Now how do you take the operating system, applications, and data on a physical box and make it a virtual machine on one of your virtualization hosts. As mentioned earlier, one way is to simply reinstall the operating system into a VM, reinstall the application, and then migrate the data and application settings. But this approach is a huge amount of work, it can lead to the application being unavailable for a significant amount of time, and it's prone to having misconfigurations creep in that cause problems either straight away or further down the running of the environment. The preferred approach is to take the complete system environment—which includes the operating system, applications, and data—and moving it directly into a virtual machine: a physical to virtual, or P2V migration.

You can perform a P2V migration in a number of ways, but this approach must include:

1. Creation of a virtual machine matching either the configuration of the physical host, a modified set of virtual hardware based on user overrides, or based on performance data gathered from actual requirements because many physical boxes are over provisioned for what is actually needed.
2. Capturing the physical disk drive's content and migration of all content to virtual hard disks in the virtual environment. This capture may be performed while the source (the physical host) operating system is running through technologies such as Volume Shadow Copy Service (VSS)—which ensures the data backed up is complete and in a usable state—or it might require the operating system on the source to be shut down and booted from an alternate capture environment such as Windows PE. If the operating system is running on the source, shutting down the application might be required to avoid any loss of state between when the source is captured and when it is started on the virtual target
3. Once the disks from the source are captured, you can update the operating system environment in the virtual environment to remove physical hardware settings and drivers so that it's ready to run in a virtual environment.
4. If the physical source operating system is running it is shut down and the VM started, completing the P2V migration.

In the process above the application would be unavailable while the P2V migration is performed and assumes that you wish to switch over to the virtual machine as soon as you perform the P2V. An alternate approach is to create the virtual machine using P2V, perform a sanity check that the VM seems healthy, then manually perform the switch. Another option is to create the VM using P2V and then keep the VM synchronized with the physical machine, either ongoing using replication or periodically using delta replication.

This ability to perform the P2V is a key feature to look for from complementary solutions to virtualization platforms that typically don't natively have this capability. Typically, P2V solutions offer very different functionalities, so look at how P2V is performed, at options for bulk migrations and synchronization (you don't want to migrate thousands of operating systems one at a time), and the reporting available to track migration status.

Once your environment is virtualized, look for technologies that monitor the usage of the virtual environment, that help balance the placement of the VMs, and that help show trending to VM usage to help you plan for future growth and avoid running out of resources. You also can leverage technologies to help monitor VM usage to chargeback virtual environment usage to the business units that requested the VMs.

Keeping your virtual environment protected

When you have 1,000 physical boxes, if one box fails you have a single failure of a single operating system instance that you have to fix. The situation is not good but it affects one system. When you use virtualization we might have 50 different operating systems running on a single physical host, which means if a host has a problem you lose 50 operating systems, the applications running on those operating systems, and access to data those applications may be serving.

Most virtualization platforms have high availability capabilities centered on the hosts that should give your protection from the scenario described above. For example, you would create a cluster of hosts that have a shared set of resources. VMs in the

cluster can run on any host that is a cluster member. If a host in the cluster fails, the VMs running on that failed host are automatically moved and started on other hosts in the cluster.

Typically, these cluster scenarios rely on a shared set of storage that holds the VMs that normally reside in a single datacenter. But it's becoming increasingly common that organizations want site resiliency and the ability to keep business running if the primary datacenter is lost.

To enable high availability between sites, you need a technology to replicate the data stores that contain your VMs to storage at the alternate datacenter. On top end, Storage Area Networks (SANs) provide the ability to perform this type of replication. But most organizations will require a software-based solution that will replicate the data from one set of storage to another set over the network. With the storage replicated to a second location, additional hosts at the second location can run the VMs from the replicated storage, if the primary datacenter is unavailable.

With the emergence of more cloud services, another option is available that involves using public cloud service providers to host the replicas of your data instead of using a second datacenter. Using a public cloud provider allows an organization to have offsite protection for their environment without an actual second site; plus, some solutions will allow you to take the replicated storage and enable the VMs to run on a virtualization platform hosted in the cloud (e.g., an Infrastructure as a Service (IaaS) provider such as Amazon EC2). This means your datacenter is replicated to the public cloud and in the event you lose the site the VMs start in the cloud automatically and clients are redirected to the cloud-based version of the VM. This type of cloud replication and hosting has a cost associated so organizations may choose to only replicate mission-critical VMs rather than the entire environment—which again stresses the need to understand the dependences between VMs ensure not only that critical VMs are replicated but also any VMs that may be needed for those critical VMs to function.

Also don't forget normal backups of your virtual machines because replication is great for high availability and disaster recovery in site failure scenarios. However, a logical data corruption gets replicated quickly as does a deletion of data. Backups provide the ability to recover from data corruption and loss, as well as enable organizations to keep point-in-time views of the data for flexible restore operations and long-term archival for corporate and regulatory reasons. When performing backups consider if just backing up the entire VM at the virtualization host level gives the right level of desired granularity for restore operations. For example, as if you would ever want to restore the entire VM or maybe files from the VM, or if you might need application-aware restorations such as restoring a specific table from a SQL database, which will likely require the backup to be taken from within the VM.

Interestingly, you also can use public cloud services for offsite storage of your backups by replicating your backup data to the public cloud—and for long-term storage of your backup data instead of having to maintain large amounts of onsite disk storage.

When you implement backups and, where needed, data replication to either an alternate datacenter or a public cloud provider you benefit from increased resiliency from failure, increasing the availability of your IT resources while ensuring protection from data loss and corruption.

Using virtualization to increase availability for the entire IT infrastructure

Once you've discovered, analyzed, and migrated your operating systems to well-architected clusters of virtualization hosts—and those hosts have been protected with replication and backups to local resources and synchronized with a public cloud provider or a second data center—you can feel good about your virtual efforts. You've made a truly an amazing accomplishment: You've streamlined the IT infrastructure that is virtualized for your organization, saved it money thanks to reductions in many areas, and increased its resilience to failure. If every system in your datacenter was virtualized then your journey would be at an end. But the reality is that organizations will still have many systems that won't be virtualized, either because of supportability, time, or scalability.

Even for the non-virtualized systems you can still increase their availability using your virtual infrastructure. You can use the same technologies that let you perform the P2V migration of your systems to replicate your physical systems to a virtual environment that you can use in disaster scenarios in either your secondary datacenter or the public cloud. Although you might not have

virtualized these systems for day-to-day running, a contingency virtual instance is preferable to no system at all in a disaster. Many organizations leverage third-party complementary solutions to enable this type of physical to virtual replication capability because it enables a disaster recovery capability without having to maintain a set of identical physical hardware that would be required for a physical-to-physical replication solution.

And don't forget about normal backups of these physical systems. Consider leveraging a single solution to protect your virtual environments and your physical environments. Protection for your client operating system data might also be possible with the single solution, which in all cases can then be replicated to a public cloud or second location for complete resiliency to a site failure.

Final thoughts

In this Technical Advisor, I've looked at key phases in your journey from physical to virtual and it's critical to not underestimate the value of good discovery of your systems and relationships before embarking on a virtualization effort. Using the right tool can turn a very painful discovery and migration effort into a far simpler and, more importantly, more successful endeavor. Once you are virtualized, put the right technologies in place to protect your virtualization investment. Fortunately, solutions are available that can provide great levels of protection and even site resiliency without organizations actually needing a second data center, thanks to integration with public cloud services.

John Savill is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He also is the author of the popular *FAQ for Windows*, and a senior contributing editor to *Windows IT Pro*. John is currently writing his fifth book, Microsoft Virtualization Secrets which will be out in 2012.